

Notebook - Work Rules!: Insights from Inside Google That Will Transform How You Live and Lead

kindle

Bock, Laszlo

Page 2 | Highlight

even the best-designed business plans fell apart when people didn't believe in them.

Page 3 | Highlight

ranking employees based on performance, choreographing job changes for top talent every twelve to eighteen months, and building a global training center in Crotonville, New York.

Page 3 | Highlight

a 20-70-10 performance ranking

Page 6 | Highlight

always surging forward to create an environment of purpose, freedom, and creativity.

Page 8 | Highlight

Why do we spend so much time on recruiting?

Page 12 | Highlight

Here is a sample of the decisions managers at Google cannot make unilaterally: Whom to hire Whom to fire How someone's performance is rated How much of a salary increase, bonus, or stock grant to give someone Who is selected to win an award for great management Whom to promote When code is of sufficient quality to be incorporated into our software code base The final design of a product and when to launch it Each of these decisions is instead made either by a group of peers, a committee, or a dedicated, independent team.

Page 13 | Highlight

If you're solving for what is most fair across the entire organization, which in turn helps employees have greater trust in the company and makes rewards more meaningful, managers must give up this power and allow outcomes to be calibrated across groups.

Page 13 | Highlight

"Managers serve the team,"

Page 13 | Highlight

the default leadership style at Google is one where a manager focuses not on punishments or rewards but on clearing roadblocks and inspiring her team.

Page 14 | Highlight

Performance improved only when companies implemented programs to empower employees (for example, by taking decision-making authority away from managers and giving it to individuals or teams), provided learning opportunities that were outside what people needed to do their jobs, increased their reliance on teamwork (by giving teams more autonomy and allowing them to self-organize), or a combination of these.

Page 16 | Highlight

Indeed, it's when the economy is at its worst that treating people well matters most.

Page 16 | Highlight

All it takes is a belief that people are fundamentally good—and enough courage to treat your people like owners instead of machines. Machines do their jobs; owners do whatever is needed to make their companies and teams successful.

Page 17 | Highlight

Every great tale starts with an origin story.

Page 21 | Highlight

always insisted that hiring decisions be made by groups rather than a single manager.

Page 21 | Highlight

simply to share what they were working on turned into the hundreds of Tech Talks we host each month.

Page 22 | Highlight

We believe it is easy to be penny wise and pound foolish with respect to benefits that can save employees considerable time and improve their health and productivity.

Page 32 | Highlight

three defining aspects of our culture: mission, transparency, and voice. A mission that matters

Page 34 | Highlight

This kind of mission gives individuals' work meaning, because it is a moral rather than a business goal. The most powerful movements in history have had moral motivations, whether they were quests for independence or equal rights.

Page 39 | Highlight

Having workers meet the people they are helping is the greatest motivator, even if they only meet for a few minutes. It imbues one's work with a significance that transcends careerism or money.

Page 40 | Highlight

We all want our work to matter. Nothing is a more powerful motivator than to know that you are making a difference in the world.

Page 40 | Highlight

people see their work as just a job ("a necessity that's not a major positive in their lives"), a career

(something to “win” or “advance”), or a calling (“a source of enjoyment and fulfillment where you’re doing socially useful work”).

Page 41 | Highlight

If you believe people are good, you must be unafraid to share information with them. Transparency is the second cornerstone of our culture. “Default to open” is a phrase sometimes heard in the open-source technology community. Chris DiBona, leader of Google’s open-source efforts, defines it like this: “Assume that all information can be shared with the team, instead of assuming that no information can be shared. Restricting information should be a conscious effort, and you’d better have a good reason for doing so. In open source, it’s countercultural to hide information.”

Page 42 | Highlight

We share everything, and trust Googlers to keep the information confidential.

Page 43 | Highlight

The benefit of so much openness is that everyone in the company knows what’s going on. This may sound trivial, but it’s not.

Page 45 | Highlight

One of the serendipitous benefits of transparency is that simply by sharing data, performance improves.

Page 45 | Highlight

“My most important principle is that getting at the truth ... is essential for getting better. We get at truth through radical transparency and putting aside our ego barriers in order to explore our mistakes and personal weaknesses so that we can improve.”

Page 45 | Highlight

A more subtle objective is to reduce politicking. It’s hard to go behind someone’s back when they

| can later listen to your meetings.

Page 46 | Highlight

| The way we solve the “backstabbing” problem, for example, is that if you write a nasty email about someone, you shouldn’t be surprised if they are added to the email thread. I remember the first time I complained about somebody in an email and my manager promptly copied that person, which forced us to quickly resolve the issue.

Page 46 | Highlight

| Fundamentally, if you’re an organization that says “Our people are our greatest asset” (as most do), and you mean it, you must default to open. Otherwise, you’re lying to your people and to yourself. You’re saying people matter but treating them like they don’t. Openness demonstrates to your employees that you believe they are trustworthy and have good judgment.

Page 46 | Highlight

| All of us want control over our destinies Voice is the third cornerstone of Google’s culture. Voice means giving employees a real say in how the company is run. Either you believe people are good and you welcome their input, or you don’t.

Page 47 | Highlight

| high-quality decisions and organizational effectiveness. Research on voice has shown positive effects of employees speaking up on decision quality, team performance, and organizational performance.”

Page 49 | Highlight

| Culture matters most when it is tested

Page 49 | Highlight

| mission, transparency, voice—

Page 51 | Highlight

If you give people freedom, they will amaze you

Page 51 | Highlight

Culture isn't static.

Page 53 | Highlight

companies have been successful with both low-freedom and high-freedom models.

Page 53 | Highlight

do you have the courage to treat people the way your beliefs suggest?

Page 53 | Highlight

if you give people freedom, they will surprise, delight, and amaze you.

Page 54 | Highlight

Why hiring is the single most important people activity in any organization

Page 56 | Highlight

Buying companies and then shutting down their products is a recent Silicon Valley phenomenon, awkwardly known as acquihiring. The ostensible purpose is to obtain people who have demonstrated their capabilities by building great products and who otherwise would not join you as employees.

Page 57 | Highlight

two-thirds of mergers and acquisitions fail to create value when the products and businesses are kept alive,

Page 57 | Highlight

labor market for employees isn't as transparent as the market for baseball players.

Page 58 | Highlight

most people simply aren't very good at interviewing.

Page 58 | Highlight

There's ample data showing that most assessment occurs in the first three to five minutes of an interview (or even more quickly),

Page 58 | Highlight

with the remaining time being spent confirming that bias; that interviewers are subconsciously biased toward people like themselves; and that most interview techniques are worthless.

Page 59 | Highlight

Designing effective training is hard. Really hard.

Page 59 | Highlight

90 percent of training doesn't cause a sustained improvement in performance or change in behavior because it's neither well designed nor well delivered.

Page 60 | Highlight

Put bluntly, which of the following situations would you rather be in? A. We hire 90th percentile performers, who start doing great work right away. B. We hire average performers, and through our training programs hope eventually to turn them into 90th percentile performers.

Page 60 | Highlight

Companies continue to invest substantially more in training than in hiring, according to the Corporate Executive Board.⁷⁴

Page 60 | Highlight

“I’m in great shape—I spent \$500 on my gym membership this month?” The presence of a huge training budget is not evidence that you’re investing in your people. It’s evidence that you failed to hire the right people to begin with.

Page 61 | Highlight

the majority of our time and money spent on people is invested in attracting, assessing, and cultivating new hires.

Page 62 | Highlight

you do have to make two big changes to how you think about hiring. The first change is to hire more slowly.

Page 64 | Highlight

Instead of naming a price up front, Salar dreamed up running an auction for every word or phrase a user might search for.

Page 64 | Highlight

our advertisers bid for the position they want in the list of ads, which can cost from less than a penny to more than \$10 per word.

Page 64 | Highlight

My simple rule of thumb—and the second big change to make in how you hire—is: “Only hire people who are better than you.”

Page 65 | Highlight

you also need managers to give up power when it comes to hiring.

Page 65 | Highlight

But even the best-intentioned managers compromise their standards as searches drag on.

Page 65 | Highlight

letting managers make hiring decisions gives them too much power over the people on their teams.

Page 66 | Highlight

One of the delightful side effects of this rigor is that the best people don't always look like what you'd expect.

Page 66 | Highlight

So we started seeking out candidates who had shown resilience and an ability to overcome hardship.

Page 66 | Highlight

The pedigree of your college education matters far less than what you have accomplished.

Page 67 | Highlight

"The broader failing of McKinsey and its acolytes at Enron is their assumption that an organization's intelligence is simply a function of the intelligence of its employees. They believe in stars, because they don't believe in systems."

Page 67 | Highlight

looking for a wide range of attributes, among the most important of which are humility and conscientiousness.

Page 67 | Highlight

"It's an interesting phenomenon: interviewing people better than you and saying 'No.'"

Page 67 | Highlight

It's about finding the very best people who will be successful in the context of your organization, and who will make everyone around them more successful.

Page 67 | Highlight

Refocusing your resources on hiring better will have a higher return than almost any training program you can develop.

Page 69 | Highlight

The evolution of Google's "self-replicating talent machine"

Page 70 | Highlight

the instinct to hew to an objective standard, ideally enforced by having a single, final, central reviewer who is charged with upholding that standard. Today we split that responsibility across two teams of senior leaders, one for product-management and engineering roles and another for sales, finance, and all other roles. And we have one final reviewer of every—yes, every—candidate: our CEO, Larry Page.

Page 71 | Highlight

These are almost always a compromise of quality. The result is that you go from hiring stellar people as a small company or team to hiring average people as a big company. The early days: hiring astounding people at a snail's pace

Page 72 | Highlight

There were many hiring committees, and each would be composed of people who were familiar with the job being filled but didn't have a direct stake in it. For example, a hiring committee for online sales roles would be made up of salespeople, but would not include the hiring manager or anyone who would directly work with the candidate. This was to ensure objectivity.

Page 75 | Highlight

our analyses revealed that academic performance didn't predict job performance beyond the first

Page 75 | Highlight Continued

two or three years after college,

Page 76 | Highlight

Bad performers and political people have a toxic effect on an entire team and require substantial management time to coach or exit.

Page 76 | Highlight

“The reality is, there are some employees you should get rid of, but the goal of recruitment should be to have no such employees!”

Page 77 | Highlight

One of the main reasons we focus so much on growing the company is to have enough great jobs for our people.

Page 79 | Highlight

intrinsic motivators, which come from inside yourself. Examples of these include a desire to give back to your family or community, slaking your curiosity, or the sense of accomplishment or pride that comes from completing a difficult task.

Page 80 | Highlight

Oops—our employees don’t know everyone in the world Our overreliance on referrals had simply started to exhaust Googlers’ networks.

Page 81 | Highlight

Breaking down a huge question (“Do you know anyone we should hire?”) into lots of small, manageable ones (“Do you know anyone who would be a good salesperson in New York?”) garners us more, higher-quality referrals.

Page 84 | Highlight

there's more variance in quality within search firms than across search firms, so selecting the individual search consultants you work with is more crucial than selecting the company.

Page 87 | Highlight

Why our instincts keep us from being good interviewers, and what you can do to hire better You never get a second chance to make a first impression" was the tagline for a Head & Shoulders shampoo ad campaign in the 1980s. It unfortunately encapsulates how most interviews work.

Page 87 | Highlight

interviewers make initial assessments and spend the rest of the interview working to confirm those assessments.

Page 87 | Highlight

These small moments of observation that are then used to make bigger decisions are called "thin slices."

Page 87 | Highlight

judgments made in the first ten seconds of an interview could predict the outcome of the interview.

Page 88 | Highlight

immediate impressions based on a handshake and brief introduction predicted the outcome of a structured employment interview.

Page 88 | Highlight

predictions from the first ten seconds are useless.

Page 88 | Highlight

an interview is spent trying to confirm what we think of someone, rather than truly assessing them. Psychologists call this confirmation bias, “the tendency to search for, interpret, or prioritize information in a way that confirms one’s beliefs or hypotheses.”

Page 89 | Highlight

most interviews are a waste of time because 99.4 percent of the time is spent trying to confirm whatever impression the interviewer formed in the first ten seconds.

Page 89 | Highlight

Equally worthless are the case interviews and brainteasers used by many firms.

Page 89 | Highlight

At worst, they rely on some trivial bit of information or insight that is withheld from the candidate, and serve primarily to make the interviewer feel clever and self-satisfied.

Page 89 | Highlight

there is no way to distinguish between someone who is innately bright and someone who has just practiced this skill.

Page 90 | Highlight

typical, unstructured job interviews were pretty bad at predicting how someone would perform once hired.

Page 91 | Highlight

The best predictor of how someone will perform in a job is a work sample test (29 percent). This entails giving candidates a sample piece of work, similar to that which they would do in the job, and assessing their performance at it.

Page 91 | Highlight

The second-best predictors of performance are tests of general cognitive ability (26 percent).

Page 92 | Highlight

most standardized tests of this type discriminate against non-white, non-male test takers

Page 93 | Highlight

Tied with tests of general cognitive ability are structured interviews (26 percent), where candidates are asked a consistent set of questions with clear criteria to assess the quality of responses.

Page 93 | Highlight

There are two kinds of structured interviews: behavioral and situational. Behavioral interviews ask candidates to describe prior achievements

Page 93 | Highlight

Situational interviews present a job-related hypothetical situation

Page 94 | Highlight

combinations of assessment techniques are better than any single technique.

Page 94 | Highlight

The goal of our interview process is to predict how candidates will perform once they join the team. We achieve that goal by doing what the science says: combining behavioral and situational structured interviews with assessments of cognitive ability, conscientiousness, and leadership. xxvi To help interviewers, we've developed an internal tool called qDroid, where an interviewer picks the job they are screening for, checks the attributes they want to test, and is emailed an interview guide with questions designed to predict performance for that job.

Page 95 | Highlight

Tell me about a time your behavior had a positive impact on your team.

Page 95 | Highlight

Tell me about a time when you effectively managed your team to achieve a goal.

Page 96 | Highlight

Tell me about a time you had difficulty working with someone

Page 96 | Highlight

For each component, the interviewer has to indicate how the candidate did, and each performance level is clearly defined. The interviewer then has to write exactly how the candidate demonstrated their general cognitive ability, so later reviewers can make their own assessment.

Page 97 | Highlight

A concise hiring rubric addresses all these issues because it distills messy, vague, and complicated work situations down to measurable, comparable results.

Page 97 | Highlight

the US Department of Veterans Affairs has a site with almost a hundred sample questions at www.va.gov/pbi/questions.asp. Use them. You'll do better at hiring immediately. Remember too that you don't just want to assess the candidate. You want them to fall in love with you.

Page 97 | Highlight

Interviews are awkward because you're having an intimate conversation with someone you just met, and the candidate is in a very vulnerable position.

Page 99 | Highlight

there were four distinct attributes that predicted whether someone would be successful at Google:

General Cognitive Ability.

Page 99 | Highlight

understanding how candidates have solved hard problems in real life and how they learn, not checking GPAs and SATs.

Leadership.

Page 99 | Highlight

Google looks for a particular type of leadership, called “emergent leadership.” This is a form of leadership that ignores formal designations—at Google there is rarely a formal leader of any effort.

Page 99 | Highlight

We have a strong bias against leaders who champion themselves: people who use “I” far more than “we” and focus exclusively on what they accomplished, rather than how.

“Googleyness.” We want people who will thrive at Google. This isn’t a neatly defined box, but includes attributes like enjoying fun (who doesn’t?), a certain dose of intellectual humility (it’s hard to learn if you can’t admit that you might be wrong), a strong measure of conscientiousness (we want owners, not employees), comfort with ambiguity (we don’t know how our business will evolve, and navigating Google internally requires dealing with a lot of ambiguity), and evidence that you’ve taken some courageous or interesting paths in your life.

Role-Related Knowledge. By far the least important attribute we screen for is whether someone actually knows anything about the job they are taking on.

Page 101 | Highlight

Once we identified these attributes, we began requiring all interview feedback to comment specifically on each one. Not every interviewer had to assess every attribute, but at least two independent interviewers had to assess each attribute. In addition, we required that the written feedback include the attribute being assessed, the question asked, the candidate’s answer, and the interviewer’s assessment of that answer.

Page 102 | Highlight

Constantly check that your hiring process actually works

Page 102 | Highlight

And just as our products can always get better, so can our hiring machine.

Page 105 | Highlight

Never compromise on quality

Page 106 | Highlight

First, assessment is led by Google's dedicated recruiters rather than line managers.

Page 107 | Highlight

After the resume is screened and selected, the second part of our process is a remote interview.

Page 107 | Highlight

Having professionals do the initial remote evaluation also means that it's possible to do a robust, reliable screening for the most important hiring attributes up front. Often a candidate's problem-solving and learning ability is assessed at this stage. We do this early so that later interviewers can focus on other attributes, like leadership and comfort with ambiguity.

Page 108 | Highlight

In every interview I've ever had with another company, I've met my potential boss and several peers. But rarely have I met anyone who would be working for me. Google turns this approach upside down.

Page 108 | Highlight

but more important is meeting one or two of the people who will work for you.

Page 108 | Highlight

The third key difference in our approach, therefore, is to have a subordinate interview a prospective hire.

Page 108 | Highlight

cronyism, where managers hire their old buddies for their new teams. We find that the best candidates leave subordinates feeling inspired or excited to learn from them. Fourth, we add a “cross-functional interviewer,” someone with little or no connection at all to the group for which the candidate is interviewing. For example, we might ask someone from the legal or the Ads team

Page 108 | Highlight

A Googler from a different function is unlikely to have any interest in a particular job being filled but has a strong interest in keeping the quality of hiring high.

Page 108 | Highlight

Fifth, we compile feedback about candidates in a radically unusual way. We’ve discussed how interview feedback must cover our hiring attributes, and the use of backdoor references.

Page 110 | Highlight

So we continue reporting the individual interview feedback scores but emphasize the average score.⁹³ This has the virtue of eliminating the ability of any single person to blackball any candidate, as well as limiting anyone’s ability to politick for a candidate. Sixth, we rely on disinterested reviewers. In addition to using structured interviews and the hiring attributes, we deliberately include at least three layers of review for each candidate.

Page 111 | Highlight

If the hiring committee rejects the candidate, the process stops there. If they are supportive of a candidate, their feedback is added to the hiring packet and sent to the Senior Leader Reviews.

Page 113 | Highlight

Putting it all together: how to hire the best If you wondered if this takes a lot of Googler time, it

does.

Page 113 | Highlight

Until we hit about twenty thousand employees, most people in the company spent four to ten hours per week on hiring,

Page 113 | Highlight

In 2013, with roughly forty thousand people, the average Googler spent one and a half hours per week on hiring,

Page 114 | Highlight

But by far the best recruiting technique is having a core of remarkable people.

Page 114 | Highlight

So how do you create your own self-replicating staffing machine? Set a high bar for quality. Before you start recruiting, decide what attributes you want and define as a group what great looks like. A good rule of thumb is to hire only people who are better than you. Do not compromise. Ever.

Find your own candidates. LinkedIn, Google+, alumni databases, and professional associations make it easy.

Assess candidates objectively. Include subordinates and peers in the interviews, make sure interviewers write good notes, and have an unbiased group of people make the actual hiring decision. Periodically return to those notes and compare them to how the new employee is doing, to refine your assessment capability.

Give candidates a reason to join. Make clear why the work you are doing matters, and let the candidate experience the astounding people they will get to work with.

Page 115 | Highlight

Managers hate the idea that they can't hire their own people. Interviewers can't stand being told

that they have to follow a certain format for the interview or for their feedback. People will disagree with data if it runs counter to their intuition and argue that the quality bar doesn't need to be so high for every job. Do not give in to the pressure. Fight for quality.

Page 115 | Highlight

If you're committed to transforming your team or your organization, hiring better is the single best way to do it.

Page 115 | Highlight

There's one other beneficial effect of hiring this way: In most organizations, you join and then have to prove yourself. At Google, there's such faith in the quality of the hiring process that people join and on their first day are trusted and full members of their teams.

Page 118 | Highlight

Take power from your managers and trust your people to run things

Page 118 | Highlight

Just as important, do you trust your manager? Does she sponsor and fight for you and help you get work done? If you're thinking about taking another job, can you talk to her about it?

Page 118 | Highlight

in fact do many good things. It turns out that we are not skeptical about managers per se. Rather, we are profoundly suspicious of power, and the way managers historically have abused it. A traditional manager controls your pay, your promotions, your workload, your coming and going, whether you have a job or not, and these days even reaches into your evenings and weekends.

Page 119 | Highlight

"Power corrupts; absolute power corrupts absolutely."

Page 120 | Highlight

Managers aren't bad people. But each of us is susceptible to the conveniences and small thrills of power.

Page 120 | Highlight

Watch closely next time. As attendees file in, they leave the head seat vacant. It illustrates the subtle and insidious nature of how we create hierarchy. Without instruction, discussion, or even conscious thought, we make room for our "superiors."

Page 122 | Highlight

Now, budgets may seem different. The whole point of a budget is that you're supposed to stay within it. But at Google you should always, always make room for a truly exceptional person, even if it puts you over budget. And yet many of us have such a built-in respect for following norms that it feels revolutionary to suggest that.

Page 123 | Highlight

Managers have a tendency to amass and exert power. Employees have a tendency to follow orders. What's mind-blowing is that many of us play both roles, manager and employee, at the same time. We each have experienced the frustration of a controlling manager, and we have each experienced the frustration of managing people who just won't listen.

Page 124 | Highlight

If you believe people are fundamentally good, and if your organization is able to hire well, there is nothing to fear from giving your people freedom.

Page 124 | Highlight

The first step to mass empowerment is making it safe for people to speak up.

Page 124 | Highlight

Eliminate status symbols

Page 124 | Highlight

To mitigate our innate human tendency to seek hierarchy, we try to remove the signifiers of power and status.

Page 124 | Highlight

really only four meaningful, visible levels at Google: individual contributor, manager, director, and vice president.

Page 127 | Highlight

Make decisions based on data, not based on managers' opinions

Page 128 | Highlight

Relying on data—indeed, expecting every conversation to be rooted in data—upends the traditional role of managers. It transforms them from being providers of intuition to facilitators in a search for truth, with the most useful facts being brought to bear on each decision.

Page 128 | Highlight

“Don’t politick. Use data.”

Page 128 | Highlight

We use data—evidence—to guard against rumor, bias, and plain old wrongheadedness.

Page 129 | Highlight

Inspired by the show, we try to test myths within the company and debunk them wherever we can.

Page 130 | Highlight

For example, in 2010 our annual employee survey revealed that many engineering Googlers felt that Google didn’t take firm enough action on poor performers. What was really happening was that on

one team of ten people, nine were all looking at the same struggling performer and concluding that no one was doing anything to help them improve or fire them. They didn't see the other five teams of like size where no one was struggling, nor was it often known that managers and People Operations folks were working with the individual behind the scenes. This is classic sample bias, where someone is drawing conclusions based on the small, flawed sample that they happen to see.

Page 135 | Highlight

"If you're not careful, you may learn something before you're done." Find ways for people to shape their work and the company

Page 136 | Highlight

In some ways, the idea of 20 percent time is more important than the reality of it.

Page 141 | Highlight

focus on engagement,¹⁰⁶ which as Prasad Setty explains, "is a nebulous concept that HR people like but doesn't really tell you much.

Page 141 | Highlight

Googlegeist instead focuses on the most important outcome variables we have: innovation (maintaining an environment that values and encourages both relentlessly improving existing products and taking enormous, visionary bets), execution (launching high-quality products quickly), and retention (keeping the people we want to keep).

Page 146 | Highlight

it is a robust, data-driven discussion that brings the best ideas to light, so that when a decision is made, it leaves the dissenters with enough context to understand and respect the rationale for the decision, even if they disagree with the outcome.

Page 146 | Highlight

But hierarchy in decision-making is important. It's the only way to break ties and is ultimately one of the primary responsibilities of management.

Page 149 | Highlight

Pick an area where your people are frustrated, and let them fix it.

Page 150 | Highlight

Improve performance by focusing on personal growth instead of ratings and rewards

Page 151 | Note

The main problem with performance reviews → systems for actually managing people!

Page 151 | Highlight

The major problem with performance management systems today is that they have become substitutes for the vital act of actually managing people.

Page 152 | Highlight

Performance management as practiced by most organizations has become a rule-based, bureaucratic process, existing as an end in itself rather than actually shaping performance. Employees hate it. Managers hate it. Even HR departments hate it.

Page 152 | Highlight

the degree of gaming that went into the system.

Page 152 | Highlight

The response in vogue today is to surrender. Adobe, Expedia, Juniper Networks (a computer hardware manufacturer), Kelly Services (a temporary worker agency), and Microsoft have all eliminated performance ratings.

Page 153 | Highlight

“We came to a fairly quick decision that we would abolish the performance review, which meant we would no longer have a one-time-of-the-year formal written review,” says Morris. “What’s more, we would abolish performance rankings and levels in order to move away from people feeling like they were labeled.”

Page 154 | Highlight

It’s far, far easier and more common to stick with vague pleasantries.

Page 155 | Highlight

It’s important to have both a quality and an efficiency measure, because otherwise engineers could just solve for one at the expense of the other.

Page 155 | Highlight

everyone’s OKRs are visible to everyone else in the company on our internal website, right next to their phone number and office location. It’s important that there’s a way to find out what other people and teams are doing, and motivating to see how you fit into the broader picture of what Google is trying to achieve.

Page 156 | Highlight

Having goals improves performance.¹¹³ Spending hours cascading goals up and down the company, however, does not.

Page 156 | Highlight

The science on rating systems is inconclusive.¹¹⁵ There’s no strong evidence to suggest that having three or five or ten or fifty rating points makes a difference.

Page 157 | Highlight

Our managers were spending thousands of hours every three months assigning ratings that were ludicrously precise but that weren't an accurate basis for determining pay.

Page 157 | Highlight

measuring performance four times per year.

Page 157 | Highlight

year either assigning ratings, calibrating ratings (I'll explain what that means in a few pages—it's important), or communicating ratings.

Page 158 | Highlight

Consensus was impossible. In the absence of clear evidence, everyone became an expert and there were constituencies arguing for every possible variation. People had strong opinions about questions like whether five or six performance categories are best.

Page 159 | Highlight

People took performance management seriously.

Page 159 | Highlight

Experimentation was vital. In the absence of external evidence, we had to develop our own, working with the leaders of each part of Google to help them test their ideas.

Page 159 | Highlight

Every team we approached was frustrated with the current system, and every team was resistant to doing something new.

Page 160 | Highlight

we stopped doing quarterly ratings, in favor of every six months.

Page 160 | Highlight

a 5-point rating scale: needs improvement, consistently meets expectations, exceeds expectations, strongly exceeds expectations, and superb.

Page 160 | Highlight

Primum non nocere.

Page 160 | Highlight

achieve the same levels of satisfaction, fairness, and efficiency

Page 161 | Highlight

The forty-one points created only an illusion of precision.

Page 162 | Highlight

By simply having more rating categories to choose from, Group B unconsciously, inadvertently, and incorrectly decided that they have almost no star performers.

Page 163 | Highlight

we found that having five categories was superior to having more in at least these two ways.

Page 164 | Highlight

Ensuring fairness

Page 164 | Highlight

the soul of performance assessment is calibration.

Page 164 | Highlight

Google's rating system was (and is) distinctive in that it isn't just the direct manager making the

decision. A manager assigns a draft rating to an employee—say, “exceeds expectations”—based on nailing OKRs but tempered by other activities, like the volume of interviews completed, or extenuating circumstances such as a shift in the economy that might have affected ad revenues. xlv Before this draft rating becomes final, groups of managers sit down together and review all of their employees’ draft ratings together in a process we call calibration.

Page 164 | Highlight

A group of five to ten managers meet and project on a wall their fifty to a thousand employees, discuss individuals, and agree on a fair rating. This allows us to remove the pressure managers may feel from employees to inflate ratings.

Page 165 | Highlight

Calibration diminishes bias by forcing managers to justify their decisions to one another. It also increases perceptions of fairness among employees.

Page 165 | Highlight

Even if you’re a small company, you’ll have better results, and happier employees, if assessments are based on a group discussion rather than the whims of a single manager.

Page 165 | Highlight

recency bias is when you overweight a recent experience because it’s fresh in your memory.

Page 165 | Highlight

We’d start each calibration meeting by revisiting these errors.

Page 166 | Highlight

Calibrators also look at the distributions of ratings across different teams, not to force a single distribution but to understand why some teams might have different distributions.

Page 167 | Highlight

I think it's about fairness. Ratings are tools, simplifying devices to help managers make decisions about pay and promotion. As an employee, I want to be treated fairly. I don't mind someone being paid more than me if they are contributing more.

Page 167 | Highlight

It also means that if someone does exceptional work, they'll be seen not just by their manager, but by lots of managers in the calibration meeting, who together create and promulgate a consistent standard across the company. Ratings also make it easier for people to move across the company. As a manager, I can trust that someone who "strongly exceeds expectations" does great work, whether her last job was working on Chrome or in Sales.

Page 167 | Highlight

A fair process for ratings gets you only so far. As a manager, you want to tell people not only how they did, but also how to do better in the future. The question is: What is the most effective way to deliver those two messages? The answer: Do it in two distinct conversations.

Page 169 | Highlight

demonstrated the power of incentives, as well as the debilitating effect of removing the incentives.

Page 169 | Highlight

They went on to demonstrate that intrinsic motivation drives not just higher performance, but also better personal outcomes in terms of greater vitality, self-esteem, and well-being.

Page 169 | Highlight

And Sam would keep arguing until I was so worn down I gave up and assigned him a higher rating. I'm ashamed to admit that, but know that I'm not alone.

Page 170 | Highlight

As long as ratings are directly linked to pay and career opportunities, every employee has this incentive to exploit the system.

Page 170 | Highlight

one study conducted by Maura Belliveau of Long Island University, 120 184 managers were asked to allocate salary increases across a group of employees. The increases aligned nicely with performance ratings. Then they were told that the company's financial situation meant that funds were limited, but were given the same amount of funds to allocate. This time, men received 71 percent of the increase funds, compared to 29 percent for the women, even though the men and women had the same distribution of ratings. The managers—of both genders—had given more to the men because they assumed women would be mollified by the explanation of the company's performance, but that the men would not. They put more money toward the men to avoid what they feared would be a tough conversation.

Page 170 | Highlight

As Prasad Setty explains, "Traditional performance management systems make a big mistake. They combine two things that should be completely separate: performance evaluation and people development. Evaluation is necessary to distribute finite resources, like salary increases or bonus dollars. Development is just as necessary so people grow and improve." 121 If you want people to grow, don't have those two conversations at the same time. Make development a constant back-and-forth between you and your team members, rather than a year-end surprise. The wisdom of crowds ... it's not just for recruiting anymore!

Page 171 | Highlight

When it's time to conduct annual reviews, Googlers and their managers select a list of peer reviewers that includes not just peers, but also people junior to them.

Page 171 | Highlight

In 2013, we also experimented with making our peer feedback templates more specific. Prior to that, we'd had the same format for many years: List three to five things the person does well; list three to five things they can do better. Now we asked for one single thing the person should do more of, and one thing they could do differently to have more impact.

Page 172 | Highlight

To make sure employees' conversations with their managers were more useful, we developed a one-page handout for them to use during their performance conversation. Again, the goal was to make the conversation more specific and tangible. We distributed these handouts to employees

just to be on the safe side; we hoped the managers would cover the right topics, but it didn't hurt to have the employees ready to guide the discussion too.

Page 173 | Highlight

was surprised, maybe even a little embarrassed (I'm supposed to know this stuff!), at how a few small changes could have such large effects across the board. Making the templates more specific reduced the time spent writing reviews by 27 percent, and for the first time, 75 percent of peers felt that writing the reviews was helpful, up 26 points (on a scale of 100) from the prior year.

Page 174 | Highlight

There's just one other twist. Googlers working in engineering or product management can nominate themselves for promotion.^{xlvi} Interestingly enough, we found that women are less likely to nominate themselves for promotion, but that when they do, they are promoted at slightly higher rates than men. This seems to be related to the dynamic that is seen in classrooms: In general, boys raise their hands and try to answer any question. Girls tend to wait to be certain, even though they are right as often as boys, if not more often.

Page 175 | Highlight

We also found that with a small nudge (an email from Alan Eustace to all technical Googlers describing this finding), women then nominate themselves in the same proportion as men.

Page 175 | Highlight

I wanted to update everyone on our efforts to encourage women to self-nominate for promotion. This is an important issue, and something I feel passionately about. Any Googler who is ready for promotion should feel encouraged to self-nominate and managers play an important role in ensuring that they feel empowered to do so. ... We know that small biases—about ourselves and others—add up over time and overcoming them takes conscious effort. ... To monitor this, we also reviewed the last three cycles of promotion data to identify any persistent gaps. ... I will continue to share this data in an effort to be transparent and open about the issue and so we can keep up this positive momentum. Of course, not everyone gets promoted, regardless of gender. If you are not promoted, the committee provides feedback on what to do to improve your chances next time.

Page 176 | Highlight

The only companies I've seen that spend as much time on performance management and promotions as we do are colleges and firms run like partnerships. At both, promotion eventually leads to you becoming something like a permanent part of the family, either as a tenured professor or a partner. Great care is taken in making a long-term commitment to you.

Page 176 | Highlight

First, set goals correctly. Make them public. Make them ambitious. Second, gather peer feedback. There is a range of online tools, not the least of which is Google Sheets, that allows you to create surveys and compile the results. (Type "Google Spreadsheets survey form" into your browser.) People don't like being labeled, unless they are labeled as extraordinary. But they love useful information that helps them do their jobs better.

Page 176 | Highlight

Third, for evaluation, adopt some kind of calibration process. We prefer meetings where managers sit together and review people as a group.

Page 177 | Highlight

Fourth, split reward conversations from development conversations. Combining the two kills learning. This holds true at companies of any size. All the other pieces of performance management: the number of performance categories, whether the categories are numbers or words, how often to give ratings, whether to do it online or on paper ... it doesn't matter.

Page 177 | Highlight

Focus instead on what does matter: a fair calibration of performance against goals, and earnest coaching on how to improve.

Page 178 | Highlight

The biggest opportunities lie in your absolute worst and best employees

Page 180 | Highlight

The Gaussian distribution is popular with researchers and businesspeople because it describes the distribution of many things: height, weight, extroversion and introversion, the width of tree trunks, the size of snowflakes, the speed of cars on a highway, the rate at which defective parts are made, the amount of customer service calls that come in, and so on. Even better, anything that follows a Gaussian distribution will have an average and a standard deviation, and you can use those to predict the future. The standard deviation describes how likely a certain amount of variation (or deviation) is to happen.

Page 180 | Highlight

The virtue of the Gaussian distribution is also its weakness. It's so easy to use, and superficially it seems to describe so many different phenomena that it's applied in cases where it doesn't describe the underlying reality.

Page 181 | Highlight

The name "power law" is used because if you wrote an equation describing the shape of the curve, you'd need to use an exponent to describe it, where one number is raised to the power of another number

Page 182 | Highlight

human performance in organizations follows a power law distribution for most jobs. Herman Aguinis and Ernest O'Boyle of Indiana University and the University of Iowa explain that "instead of a massive group of average performers dominating ... through sheer numbers, a small group of elite performers [dominate] through massive performance." 130 Most organizations undervalue and underreward their best people, without even knowing they are doing it.

Page 182 | Highlight

What most organizations miss is that people in the bottom tail represent the biggest opportunity to improve performance in your company, and the top tail will teach you exactly how to realize that opportunity.

Page 183 | Highlight

At Google, we regularly identify the bottom-performing 5 percent or so of our employees. These

individuals form the bottom tail of our performance distribution. Note that this happens outside our formal performance management process. We're not looking to fire people: We're finding the people who need help.

Page 183 | Highlight

we don't force a distribution of ratings, because different teams perform at different levels. It would be madness to force the manager of a team of all superstars to rank someone as failing.

Page 184 | Highlight

"You are in the bottom 5 percent of performers across all of Google. I know that doesn't feel good. The reason I'm telling you this is that I want to help you grow and get better."

Page 184 | Highlight

"compassionate pragmatism."

Page 185 | Highlight

The "yank" myth of differentiation says the bottom 10% are summarily fired. In reality, that's rare. More typically, when a person has been in the bottom 10% for a sustained period of time, the manager starts a conversation about moving on. Occasionally, of course, an underperformer doesn't want to go. But confronted with the cold reality of how the organization views them, most people leave of their own accord and very often end up at companies where their skills are a better fit and they are more appreciated.

Page 186 | Highlight

Every current and former Microsoft employee I interviewed—every one—cited stack ranking as the most destructive process inside of Microsoft, something that drove out untold numbers of employees. ... "If you were on a team of 10 people, you walked in the first day knowing that, no matter how good everyone was, two people were going to get a great review, seven were going to get mediocre reviews, and one was going to get a terrible review," says a former software developer. "It leads to employees focusing on competing with each other rather than competing with other companies."

if you believe people are fundamentally good and worthy of trust, you must be honest and transparent with them.

Page 187 | Highlight

Because top performers live in a virtuous cycle of great output, great feedback, more great output, and more great feedback. They get so much love on a daily basis that the extra programs you might offer can't actually make them much happier. More important is to learn from your best performers.^{xlvii} Every company has the seeds of its future success in its best people, yet most fail to study them closely.

Page 193 | Highlight

Career decisions were made fairly. Performance was fairly assessed and promotions were well deserved. Their personal career objectives could be met, and their manager was a helpful advocate and counselor. Work happened efficiently. Decisions were made quickly, resources were allocated well, and diverse perspectives were considered. Team members treated each other nonhierarchically and with respect, relied on data rather than politics to make decisions, and were transparent about their work and beliefs. They were appropriately involved in decision-making and empowered to get things done. They had the freedom to manage the balance between work and their personal lives.

Page 193 | Highlight

manager quality was the single best predictor of whether employees would stay or leave, supporting the adage that people don't quit companies, they quit bad managers.

Page 196 | Highlight

we found that technical expertise was actually the least important of the eight behaviors across great managers. Make no mistake, it is essential. An engineering manager who can't code is not going to be able to lead a team at Google. But of the behaviors that differentiated the very best, technical input made the smallest difference to teams.

Page 196 | Highlight

Gawande then argued that the field of medicine was entering the same stage, where complexity

outstripped human capacity and checklists would save lives.

Page 197 | Highlight

But if we could reduce good management to a checklist, we wouldn't need to invest millions of dollars in training, or try to convince people why one style of leadership is better than another. We wouldn't have to change who they were. We could just change how they behave.

Page 197 | Highlight

a system of reinforcing signals to improve the quality of management at Google. The most visible, a semiannual Upward Feedback Survey,