

APRIL 2015

McKinsey Quarterly

Brightening the black box of R&D

Eric Hannon, Sander Smits, and Florian Weig

An all-in-one, one-for-all formula to determine R&D's productivity can help companies see how well the function is performing.

The question of R&D's productivity has long resembled a Gordian knot. Look nearly anywhere else in today's corporations, and there's far less difficulty measuring productivity and performance. In manufacturing and logistics, you can get a sense of things just by looking around the production floor, the inventory room, or the loading dock. Even the performance of the advertising budget—once famously opaque—is now, thanks to digital technology, much easier to see.

But the R&D department provides fewer clues. There's no flow of tangible goods through the process, for one thing, but rather a stream of ideas and concepts that resist the efforts of efficiency experts and innovation gurus alike. In the face of this difficulty, most companies fall back on a few well-worn approaches: R&D as a percentage of revenue, the ratio of new products to sales, or the time it takes for new products to reach the market. None of these really gives a good idea of how well the R&D function is performing, either overall or by team—nor is it clear why (or when) any given project might suddenly prove a failure though it had earlier shown every promise of success.

We have endeavored to address this long-standing puzzle. We may not have answered it definitively, but we have developed a formula we believe will be useful to any company that wants to establish and maintain a comprehensive and transparent overview of the R&D

organization's many platforms, hundreds of projects, and thousands of engineers, technicians, program managers, and lab workers. Just as Alexander the Great is said to have undone the Gordian knot by the simple expedient of slicing through it with his sword—rather than trying to unravel it by hand, as others had attempted to do—our formula makes relatively quick, simple work of a knotty problem.

This formula takes a novel approach to measuring R&D outcomes: multiplying a project's total gross contribution by its rate of maturation and then dividing the result by the project's R&D cost. Since proposing this idea, we have worked with several companies to test it and introduced it to a diverse group of approximately 20 chief technology officers (CTOs) and other senior executives in a roundtable setting. So far, the formula demonstrates several virtues. First, it's a single metric rather than a collection of them. Second, it aims to measure what R&D contributes within the sphere of what R&D can actually influence. Finally, by measuring productivity both at the project level and across the entire R&D organization (the latter through simple aggregation), it endeavors to speak to the whole company, from the boardroom all the way to the cubicle. Refinements to the approach may be necessary, but for now at least, the formula seems to represent an advance in measuring R&D's productivity and performance.

The case for a new approach

Before describing the formula in greater detail, let's examine what *doesn't* work in today's approaches to measuring R&D's productivity, and why that matters.

Today's flaws . . .

The most common approach takes the ratio of R&D's costs to revenue. This method divides revenue from products developed in the past by what's currently being spent on products for the future. That might be useful in a stable or stagnant company whose prospective revenues are expected to grow very steadily or to remain flat. But for any other company, this assumption is artificially pessimistic for investing in future growth and falsely optimistic when the product pipeline is weakening. Indeed, repeated studies

have shown no definite correlation between this R&D ratio and any measure of a company's success.¹

Not that anything better has been proposed in the past—and not for lack of trying. One academic paper² found no single, top-level metric and therefore recommended that companies instead use a suite of metrics at different levels of the organization.

. . . and why they matter

Maybe at one time, R&D's productivity mattered less. But today, myriad competitive forces drive down R&D budgets, and nearly every company we know—even those investing heavily in growth—continues to ask the R&D organization to achieve more with the same or fewer resources. (One CTO admits that his method is “to keep turning the budget dial down until the screaming gets too loud”; that's when he knows he's hit the right level.)

Meanwhile, as product variations, functional requirements, and customization needs (to say nothing of regulatory demands) proliferate, the complexity and cost of R&D continue to rise. Small wonder friction arises between R&D managers, struggling to articulate the scope of the challenges they face, and other executives, who are frustrated with the rising cost of product development. In some industries, such as semiconductors, where Moore's law is pushing the limits of physics, this friction is acutely apparent.

At the source of the frustration is the difficulty of generating lasting R&D-productivity improvements at many companies. One reason is the lack of repetitive tasks, at least compared with other parts of the organization. Another is the more frequent reshuffling of R&D project teams.

¹ Studies include Christoph H. Loch and U.A. Staffan Tappar, “Implementing a strategy-driven performance measurement system for an applied research group,” *Journal of Product Innovation Management*, 2002, Volume 19, Number 3, pp. 185–98, wiley.com; Martin Falk, “Quantile estimates of the impact of R&D intensity on firm performance,” *Small Business Economics*, 2012, Volume 39, Number 1, pp. 19–37, springer.com; and Raphael Braga da Silva et al, “Innovative intensity and its impact on the performance of firms in Brazil,” *Research in International Business and Finance*, 2015, Volume 34, pp. 1–16, elsevier.com.

² See Christoph H. Loch and U.A. Staffan Tappar, “Implementing a strategy-driven performance measurement system for an applied research group,” *Journal of Product Innovation Management*, 2002, Volume 19, Number 3, pp. 185–98, wiley.com.

Moreover, R&D managers usually can't identify troubled projects until they're well into an escalating spate of costly late changes and firefighting. Many of the technical shortfalls of products become clear only just before they are introduced into the market. As a result, it's often hard to determine, in the fire drill that accompanies the last weeks and months of a troubled project, exactly what all the engineering hours were spent on and who spent them.

A new formula

When you dig more deeply into the R&D conundrum, you quickly encounter the problem of measuring what the R&D organization actually accomplishes—the outputs, so to speak. Any formula for productivity by definition divides outputs by inputs. The input variable, in this case, is straightforward: the cost of an R&D project. That's the one used by most existing measures of R&D's productivity and the one we too decided to use.

To capture the outputs—a stickier task—we settled on using, first, the gross contribution of a project and, second, a complementary measure: the rate of maturity, or a project's progress toward meeting its full technical and commercial requirements. We chose these measures for their overall explanatory power and the visibility they provide into certain aspects of the R&D process. They come together in the formula shown in the exhibit.

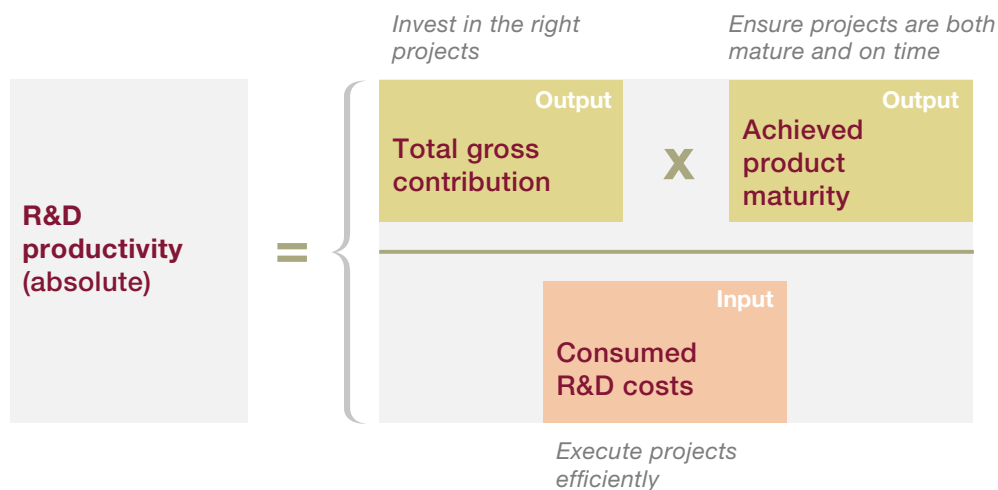
Total gross contribution

We chose total gross contribution as one part of the formula's numerator because it represents, over time, a product's economic value to customers, while keeping fixed costs out of the equation. That allows us to home in on what R&D can directly influence. Also, by looking at the total gross contribution of projects over time, companies can highlight information that helps to evaluate the projects they have in process and to continue or cancel them. That nicely ties the metric to one kind of behavior it's meant to influence.

How do we know what the gross contribution is? Looking back in time, it's easy enough to determine. Thus, when a company calculates a project's rate of maturation (a step we'll describe in a

Exhibit

A simple formula provides companies with a single measure to assess the productivity of the R&D function.



Source: McKinsey analysis

moment), it can determine a completed R&D project's productivity retrospectively.

However, when executives consider a project that's in process or has yet to be started, they don't know whether it will capture its potential gross contribution and must instead rely on a credible and reasonably accurate estimate. The more accurate the forecast, the better the formula will work as a leading indicator. You could even say, from a skeptical point of view, that the formula is only as good as the estimates that go into it—which is true, as far as it goes. But even for companies that tend to be overly optimistic or pessimistic in their business cases, faulty estimates will provide at least a basis for "go/no-go" decisions about different projects. In addition, even a flawed estimate can be used to see, earlier in the evaluation process, whether a project's productivity is dropping relative to the forecast. This is often a reliable indicator that a project won't return its predicted gross contribution.

That said, the formula we propose will work best for companies with incremental R&D processes and less well in start-ups with more uncertain R&D spending.

Achieved product maturity

While a project's gross contribution may be necessary to measure R&D's output, it's not sufficient, because it isn't earned all at once but rather over time. The likelihood that a project will attain the projected gross contribution depends, in part, on the maturity of the product at the time of its market introduction—how close it is to verifying and validating its technical and commercial requirements. (Of course, other factors also influence whether a given product or service captures its full potential, including how well it was marketed and how well the company timed its introduction.) Our experience shows that the closer to full maturity a product is when introduced, the better the chance that it will fulfill its expected gross contribution.

That's not only because the product-maturity rate largely determines time to market but also because late changes to a developing product typically cost more to fix than earlier ones. Such late changes might, for example, require a company to rework expensive tooling or to redesign interface components or features. Higher costs mean a lower gross contribution.

The implication is that companies must be able to assess, in real time, how close their R&D projects are to full maturity. Few companies may in fact have this capability, but a rough-and-ready version of such a system *can* be built fairly quickly, often in two to three weeks. To do so, a company simply looks at critical dimensions (such as cost, functionality, and quality) during each of the quality gates a project passes through in its development. These provide a fair proxy in a rudimentary system if they are reported in consistent fashion throughout a company.

But if we are going to find the precise productivity formula we're seeking, we need a more sophisticated and systematic method—for example, one that checks on a project's progress toward meeting its performance requirements within a narrowing allowable deviation corridor over its lifespan. This method uses technical and commercial metrics specific to each product instead of the more generic metrics used in the rough-and-ready version. It lets companies drill down to the maturity of single components within a project and to zoom out and gauge the maturity of an entire product and service pipeline. (See sidebar, "Absolute or relative productivity?")

Absolute or relative productivity?

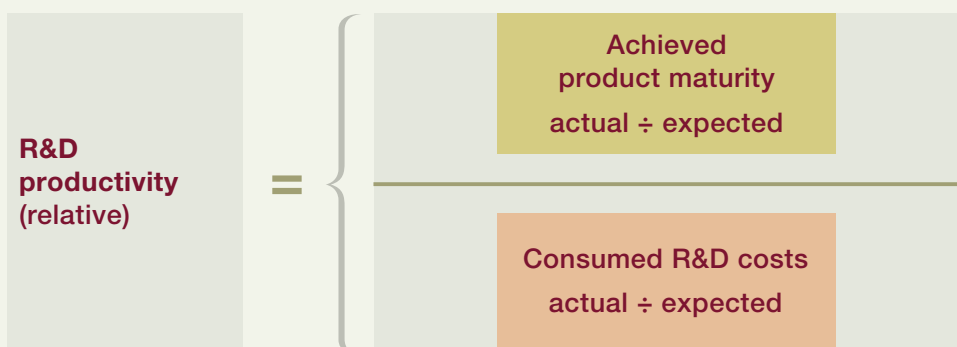
The formula as presented in this article calculates the *absolute* productivity of a company's R&D organization. This means that it calculates a single number for the productivity of all the projects a company undertakes. However, in practice, we have found that it is helpful for companies to calculate a *relative* measure of the formula's maturity and cost factors, as well. This compares progress against a forecast for a single project or set of projects, evaluating execution performance independently of project size and expected returns.

This distinction between absolute and relative productivity measures is shown in the exhibit. In moving from absolute to relative value, managers must define expected maturity and development cost curves for each project, so that its actual performance can be measured against them. That in turn requires a transparent, credible process for establishing the expected ("should") values. As companies apply the formula to their performance-management systems, they will want to be mindful of which measure best fits the results and behavior they hope to promote at each layer of the organization.

A relative measurement places all projects on a level playing field. The advantage is that their performance can be considered equally, regardless of a given project's size or scope. In this case, all projects are treated in the same way, and bigger ones, with a higher possible contribution, do not receive greater weight. This approach enables a fair comparison of execution ability across projects. However, a relative system does not automatically provide a built-in driver for increased productivity over time. Since an ideal project will always achieve a productivity value of 1, this system must be tuned to provide incentives for improved performance.

For most companies, a combination of the relative and absolute versions has proved to be best. At the highest levels of an organization, the absolute measurement gives a clear overall view of R&D productivity. At the operational level, the relative version tends to give a better view of execution ability.

The formula for measuring relative R&D productivity compares progress against a forecast.



Of course, there's a broader reason, beyond time to market, why the rate of maturity is an important measure of the R&D function's output: designing and maturing the products that the strategy and marketing functions conceive is the primary reason R&D exists.

Integrating the elements

These three elements—total gross contribution, rate of maturity, and cost of R&D—come together in a formula that attempts to quantify R&D's overall performance and to shed light on separate aspects of productivity. This, in turn, facilitates more confident managerial interventions to improve them.

By weighting projects according to their expected gross contribution, for instance, we keep our focus on efforts critical to a company's success, while also articulating the value R&D generates over a defined time period. By tracking the race to a mature product, we make sure R&D gets credit for its value contribution only if it delivers such a product. Projects that reach maturity in timely fashion are acknowledged for having justified the full business case for them. Project teams that launch immature products, which are less likely to capture their full expected gross contribution, get penalized.

The formula's usefulness, then, lies in the way it drives the right behavior. By more heavily weighting projects forecast to make a higher gross contribution, our approach helps focus management's attention on the ongoing projects most critical to a company's future success. Furthermore, the formula encourages a faster time to market, since products that reach maturity more quickly will show a higher level of productivity. Finally, the formula encourages the efficient execution of projects because those that consume less investment will also have a higher productivity value.

The formula in action

Measuring productivity, valuable though that may be, is just a starting point—it won't change R&D's efficiency on its own. The formula must be integrated into existing management processes or lead to the creation of new ones. One company used the approach to perform a one-time analysis looking at all of its R&D projects for

the previous five years. The idea was to establish a baseline R&D-productivity measure that would serve as a yardstick for future efforts. To see how productivity is changing, the company now runs each of its current projects and each of its project teams through the formula two and four times a year, respectively. It will take a few years before the company can trace the results all the way to specific products and their marketplace performance. But already, we can see its benefits when confronting some perennial challenges: setting the direction of R&D, improving the performance of teams, making decisions, and driving change.

Setting direction

A key benefit of this productivity formula is its ability to address, through a single metric, all levels of the organization—from individual engineering teams to the full R&D pipeline. As such, it provides a backbone for an integrated performance-management system that unifies an entire company's R&D efforts. This unity comes with significant flexibility: companies can select separate parts of the formula to gain insights into the different elements of the R&D function and thus to influence both the particulars and the whole.

CTOs can convincingly quantify for their boards any increase, over the preceding year, in the productivity of the entire R&D organization by annually measuring its productivity. By looking only at the numerator, executives can report R&D's overall value contribution. By multiplying the product portfolio's expected gross contribution by the respective increase in maturity achieved over the measured time period, they can determine the total value R&D generates.

And that's not all. By taking the formula's left-hand elements—the total gross contribution and R&D costs of individual projects—executives can develop a metric to help prioritize the overall product-development pipeline and thereby make better portfolio and resource-allocation decisions. (Are critical and valuable projects being deprived? Has organizational momentum allowed bloated projects to consume too many resources?) And by looking at the formula's right-hand elements—the rate of maturation divided by the cost of R&D—executives can better assess the efficiency of working teams. Such transparency is a powerful tool for improving their performance.

Improving teams

In any R&D organization, some teams perform at an extremely high level and others struggle. This range of performance can be difficult to identify, at least objectively. Naturally, individual managers often have an instinct for high-performing teams but lack a means to quantify that performance or to make comparisons.

Publishing a ranking of productivity by using the right-hand elements of the formula—the rate of maturation over the corresponding R&D cost—makes a team’s performance immediately apparent. Obviously, that insight does not, in and of itself, drive improvement. But by enabling investigations into what specific kinds of behavior truly make teams excel, the formula provides an important first step.

Companies can therefore avoid the broad, one-size-fits-all improvement approaches that rightly make executives leery. Particularly in large organizations, it’s almost impossible to improve all the engineering teams at once. The starting points and improvement needs of different projects and teams are simply too diverse. Companies are better off focusing their limited resources on teams with the most potential for improvement. By applying the methodology described here, a company should avoid employees’ “not invented here” hostility toward the practices of external organizations. The practices identified through the formula, after all, are internal to the company that carries out the analysis, and lower-ranking teams can simply walk across the hall, so to speak, to see and learn from their higher-performing peers.

We have seen R&D teams that apply internal practices commit themselves voluntarily to improving their performance (in the most important indicators) by more than 20 percent, on average. One company, for example, significantly increased its ability to hit its technical objectives by implementing a systematic process for the engineering release of a highly complex industrial component.

Making objective decisions

This productivity-based method improves the management of R&D in a third way, as well: by providing an objective and numerical basis for making decisions and setting targets. It bypasses gut-feeling decisions and the sort of arbitrary budget and performance-

improvement targets so often divorced from the reality of R&D challenges. The formula allows executives to better understand the demands they're placing on the function in the context of its historical productivity performance, creating a more reliable budget for the product-development portfolio. When executives know the productivity of individual R&D teams, they can calculate the likely cost of a project, even down to the contribution of individual functional areas.

Driving change

The transparency this system of performance measurement provides is an invaluable companion to any large-scale R&D-productivity initiative. Compared with initiatives in other parts of a company—for example, programs to reduce the cost of materials, where any gain is very tangibly demonstrable in the piece price—improvements in R&D are often ephemeral. In our experience, many large-scale transformations identify millions of dollars in R&D-efficiency benefits only to leave the function's budget unchanged.

Our method allows managers to measure a change program's impact objectively. And even if the R&D budget does stay the same, faster or better development should be reflected in overall productivity. By quantifying the impact of any change program, moreover, executives will be better able to communicate its success in a credible and convincing way.



R&D is one of the few areas that often remain opaque to executives in today's corporations. Quantifying what it actually accomplishes has resisted the efforts of executives and academics alike. By clarifying the outputs, the simple formula proposed here endeavors to generate a single measure companies can use to determine and agree on the R&D function's productivity—the better to assist decision making and to improve performance. ○

Eric Hannon is an associate principal in McKinsey's Frankfurt office, **Sander Smits** is a principal in the Amsterdam office, and **Florian Weig** is a director in the Munich office.