White Paper

Enhancing the Efficiency of Pharmaceutical R&D

Optimizing success rates.

ALEX ROLAND, PhD, Associate Principal, R&D Transformation Consulting, IQVIA
WILLIAM FOX, Consultant, R&D Transformation Consulting, IQVIA
ANN BAKER, Vice President, R&D Transformation Consulting, IQVIA

Table of contents

Introduction

During much of the last two decades, the volume of new product approvals failed to reflect the magnitude of industry R&D investment and pharmaceutical company productivity continues to be a topic of concern among industry analysts today. Furthermore, recent pricing and regulatory developments — such as the Inflation Reduction Act in the US[1] and anticipated changes to exclusivity periods in the EU[2] — are expected to reduce prospective returns on R&D investment for many investigational products,[3] presenting significant additional risk to overall R&D productivity.

Broadly speaking, R&D productivity can be thought of in terms of two components: R&D efficiency and R&D effectiveness.[4] The R&D efficiency component describes the cost of translating R&D inputs — such as ideas, effort and investment — into R&D outputs, or more specifically the approval of new pharmaceutical products. The R&D effectiveness component on the other hand, describes the quality or commercial value delivered by each output (Figure 1). In an environment in which external pressures will increasingly limit R&D effectiveness, taking action to enhance R&D efficiency becomes imperative in maintaining or improving overall productivity.

Whilst there are several underlying efficiency levers that can be targeted in order to improve efficiency, modelling shows that optimizing success rates — the likelihood of successfully transitioning investigational assets through each phase — usually provides the largest opportunity.

In this paper, we discuss the importance of R&D success rates and highlight some of the ways in which companies can act to optimize success rate and attrition profiles so as to maximize the efficiency and productivity of their R&D functions.

Modelling R&D efficiency

As previously described,[5] we calculate R&D efficiency in cost per new product approval terms, broadly following the methodological approach described by Paul et al, 2010.[4] Across a cohort of 14 major pharmaceutical companies, we use company-specific between-phase success rate estimates to calculate the number of pipeline assets required at each phase of R&D in order to achieve a single asset approval, taking into account pipeline attrition. We then apply estimated 'out-of-pocket' between-phase costs to each of these prior assets. To account for the time value of money,[6] expenditure is capitalized to the point of approval based on company-specific between-phase cycle times. The R&D efficiency metric therefore represents the average capitalized, attrition loaded cost of achieving a single new product approval for each company.

Figure 1: Elements of R&D productivity

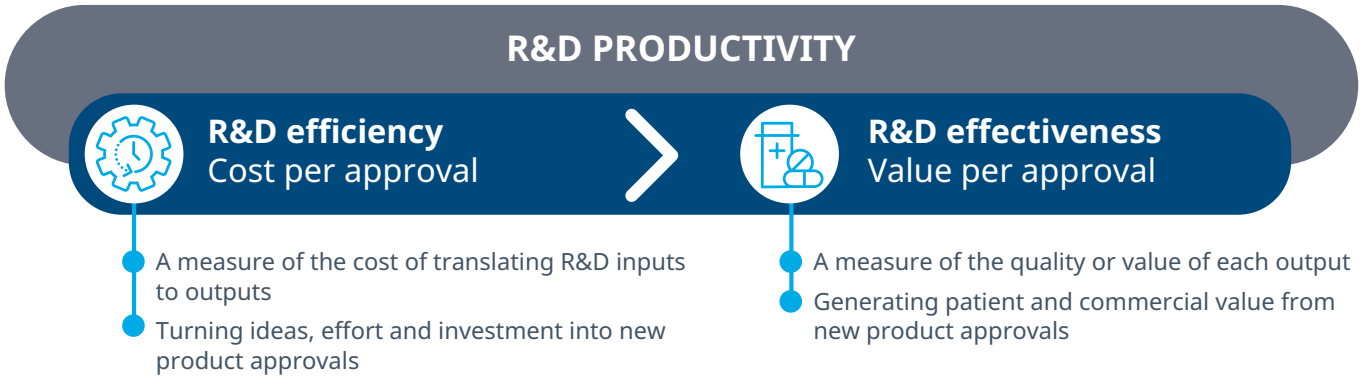

Success rates are usually the most important driver of R&D efficiency

Most companies have a significant opportunity to improve cost-per-approval

Using our previously-described R&D productivity model,[5] we have found that improvements in between-phase success rates — particularly in mid-to-late-stage R&D — typically have a greater impact on overall cost per new product approval than improvements in any other efficiency lever. The paramount importance of success rates in R&D efficiency is also evident in published reports based on syndicated benchmarking data,[4] albeit at an industry aggregate level.

Our analyses indicate that almost all major companies could achieve improvements to productivity by optimising success rates in R&D, although to varying degrees. Moreover, these opportunities are demonstrably tractable. Several robust precedents of dramatic success rate improvement have been disclosed by major pharmaceutical companies.[7,8] Nonetheless, optimizing success rates can be complex in practice and there are several important pitfalls to be avoided in doing so.

The answer to improving productivity is not however to maximize end-to-end success rates across the entire R&D lifecycle from the beginning of discovery-stage work through to approval, for a number of key reasons which we discuss in the following sections.

Attrition is a necessary evil

Pipeline investment decision-making and portfolio prioritization are of course essential elements in pharmaceutical R&D governance. Eliminating the least attractive assets and projects from the portfolio allows limited R&D resources to be focused on the best prospects with the greatest opportunity to benefit patients and deliver commercial returns. A degree of pipeline attrition is therefore a requirement for productive pharmaceutical R&D; however, as discussed further below, how attrition is distributed between different R&D stages has a major impact on R&D efficiency and productivity.

Success rates and probabilities of success

Between-phase success rate is an industry-standard metric describing the proportion of R&D projects or assets that successfully progress from one R&D stage to the next, as a percentage of the total number of progressions plus terminations. Pharmaceutical companies sometimes also cite a secondary, 'success rate to approval' metric. This derived metric is used to examine the cumulative success rate spanning multiple R&D stages and ending at product approval or launch. It is calculated as the product of each component between-phase success rate over a defined time period. Success rates vary by modality, indication, development stage and other characteristics and are also influenced by a wide range of process, governance and strategic factors that drive considerable differences between companies.

Success rates reflect actual, achieved project or asset outcomes and should not be confused with 'probabilities of success' — including PTS (Probability of Technical Success) and PTRS (Probability of Technical and Regulatory Success). The latter represent prospective probability metrics, describing predicted likelihoods of successful progression. They are typically applied in the process of risk-adjusting anticipated future costs and returns for individual R&D project valuations. Nonetheless, success rates (actuals) can be used to inform probabilities of success (predictions) and, conversely, the probability of success predictions across a portfolio of projects will influence future success rates, assuming that those predictions are realistic (Figure 2).

Figure 2: Success rates and probabilities of success

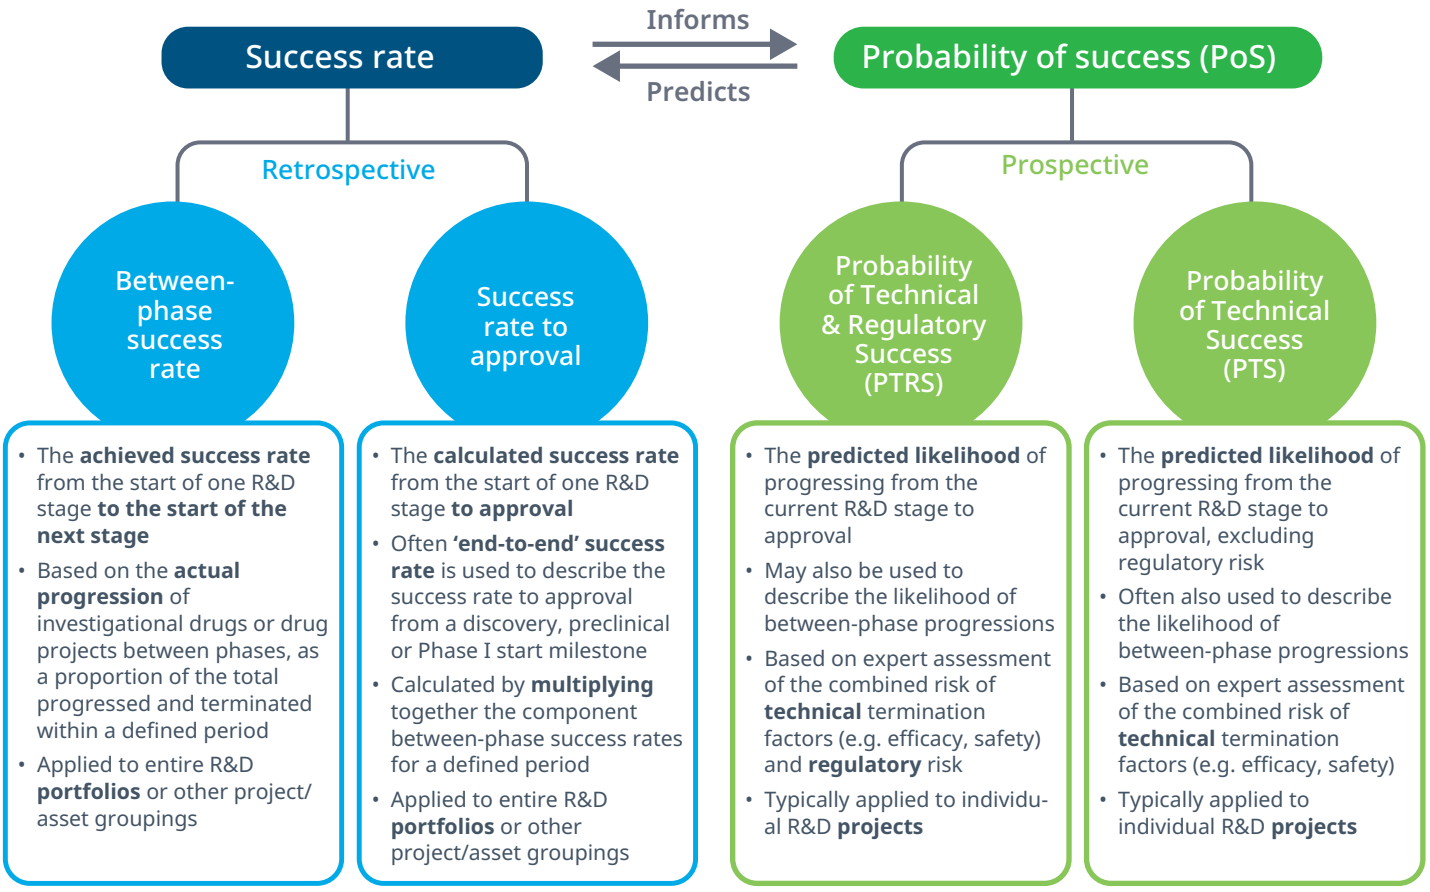

All attrition is not equal

The impact of increasing or declining success rates varies by R&D phase. We have illustrated this using our productivity model[5] for a hypothetical 'typical' major pharmaceutical company, in which all efficiency levers perform at cohort-median levels. The impacts of increasing or decreasing success rates in individual clinical phases, to best quartile or worst quartile levels for the cohort (75th or 25th percentiles, respectively), are shown in Figure 3.

Changes to clinical-stage success rates typically have the greatest cost-per-approval impact in Phase II and the smallest impact in Phase I. In the example, improving Phase II success rate from cohort-median level to best-quartile delivers a saving of $383m in the cost of each new product approval. But if the Phase II success rate were to deteriorate to worst-quartile, each new product approval would require an additional $821m to achieve.

This analysis also indicates that an otherwise-typical company with poor Phase II success rates could save over $1b per approval if it were able to transition from bottom- to top-quartile performance on this metric (based on the difference between 25th and 75th percentile performance; Figure 3). However, in taking steps to achieve this, care would need to be taken to identify and manage the potential for counter-productive impacts in other phases or on other levers, as discussed further in the next section of this paper.

Figure 3: Typical R&D efficiency impact of median clinical-stage success rates moving to top- or bottom-quartile levels (75th or 25th percentile)

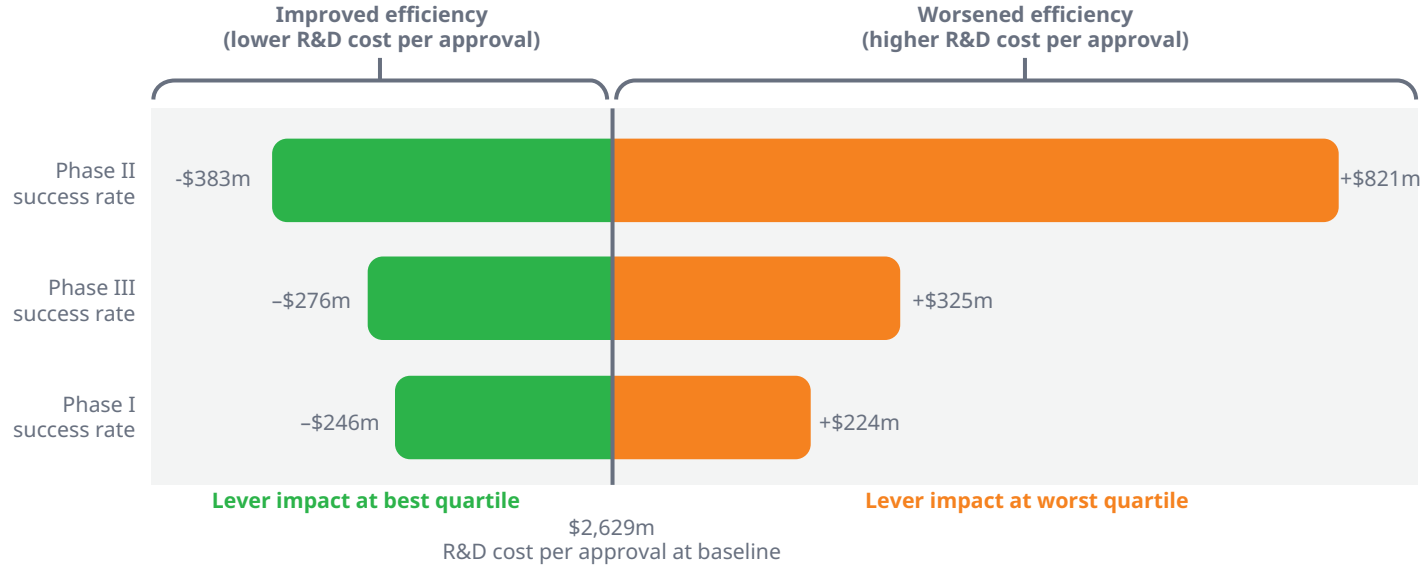

Source: IQVIA analysis.

The unequal impact of attrition at different R&D stages also has an important consequence for end-to-end success rate (the cumulative success rate from an early R&D stage to approval). Two otherwise-identical companies that display the same end-to-end success rate can have very different costs per approval, depending on how their attrition is distributed between different R&D stages.

Counter-intuitively, it is even possible for an increase in *end-to-end* success rate to result in *reduced* R&D efficiency (increased cost per approval), if attrition is unfavourably distributed between phases. We have illustrated this effect for a hypothetical 'typical' company — 'Company X' — using two alternative scenarios (Figure 4). In both scenarios, end-to-end success rate increases from 11% to 12%. Scenario 1 saves around $300m per approval, whereas Scenario 2 *increases* cost-per-approval by around $100m, due to differences in the distribution of underlying attrition (and therefore success rates) between clinical phases.

This demonstrates that an effective 'quick-kill' approach to R&D — in which processes and technologies are employed to facilitate confident elimination of unsuitable programs at the earliest possible stage of R&D — is of course highly desirable. Indeed, taking actions to increasingly weight portfolio attrition towards early R&D stages can be more impactful than taking actions to improve the overall *end-to-end* success rate in a less focused manner.

Between-phase success rates are interdependent

As discussed above, a degree of attrition is essential in order to enrich the portfolio with the most promising prospective medicines and should ideally be achieved as early as possible in order to maximize R&D efficiency. But in seeking to optimize success rate and attrition profiles, it is important to be aware of interdependencies that exist between different R&D stages.

Consider a company with strong governance processes and advanced predictive technologies in place to enable rigorous early assessment of biological targets and candidate drug quality. Such a company might be expected to have a highly front-loaded attrition profile, with low discovery research success rates and high clinical-stage success rates. In this company, the development-stage portfolio is enriched with high quality targets and promising drug candidates, resulting in high R&D efficiency and productivity.

Figure 4: The impact of end-to-end success rate improvement depends upon how success rates are distributed between phases

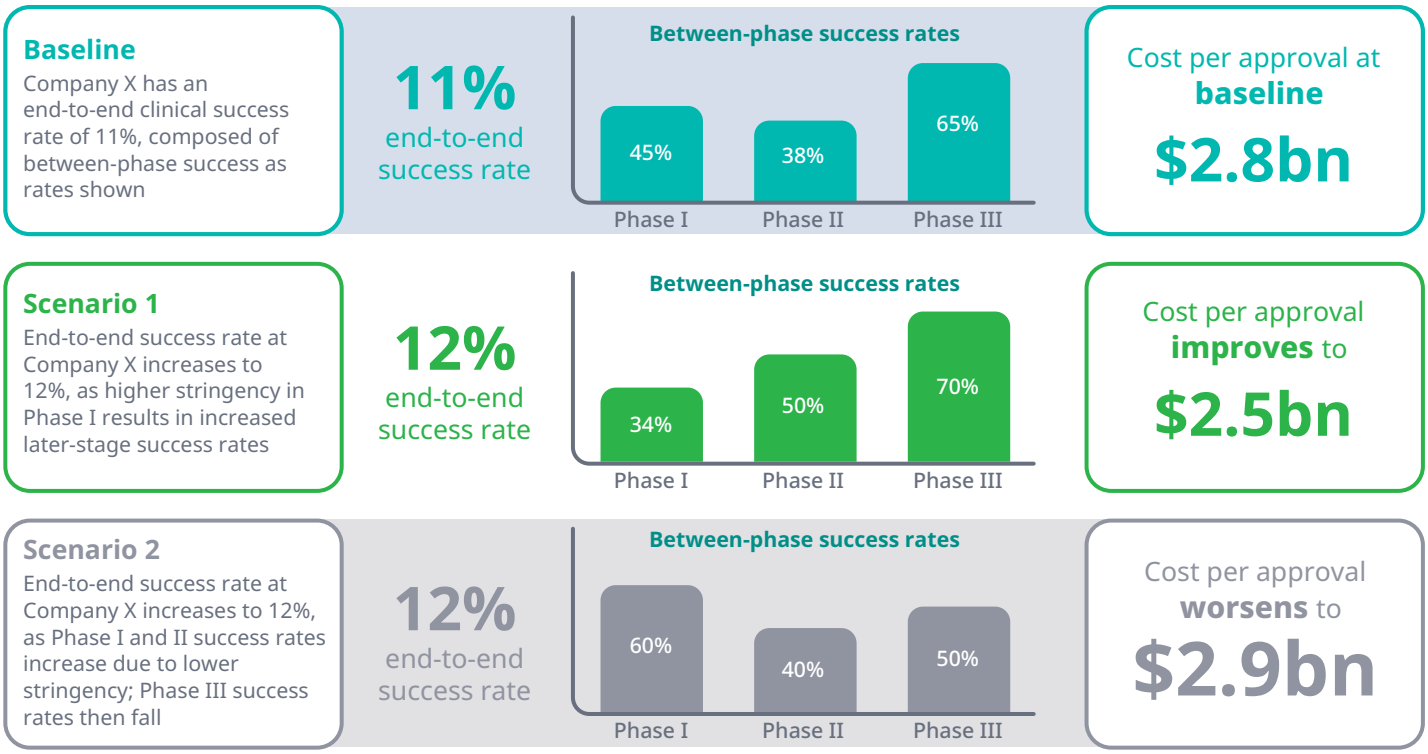

Source: IQVIA analysis.

Considering this company's early R&D in isolation — without considering the broader R&D context — could conceivably result in the research organization being perceived as underperforming. Objectives might then be set to increase discovery success rates and the volume of projects progressing into development stages. If this is achieved at the expense of progression decision rigor, superficially positive outcomes seen in research — as early-stage success rates increase — will result in declining quality of substrate entering the development organization. This in turn will have important negative downstream consequences for development and for R&D as a whole, as late-stage success rates eventually deteriorate and overall efficiency declines.

So in this scenario, reduced rigor in early R&D increases early success rates but ultimately worsens development success rates and overall R&D efficiency. But this may not be the end of the story. As development-stage success rates worsen, it may become necessary to ramp-up business development efforts in order to plug emerging gaps in the pipeline. But we might also envisage a worst-case (but not unprecedented) scenario in which an imperative to quickly repopulate the now-sparse late-stage pipeline increases pressure to advance assets from Phase I to Phase II and from Phase II to Phase III. This may exacerbate the original problem by delaying attrition of low-quality programs to later development stages. Importantly, this misguided strategy may again manifest pseudo-positive outcomes in the short-term, as Phase I and Phase II success rates increase (in line with reduced progression thresholds) long before the downstream decline in Phase III success rates becomes visible.

In seeking to optimize R&D success rates, it is therefore critical to consider the R&D continuum in its entirety, including potential downstream negative effects. Siloed approaches within early-stage R&D functions carry a particularly high risk of counter-productive long-term outcomes for R&D as a whole.

Acting to optimize success rates

One size does not fit all

Although broad conclusions can be drawn on how best to leverage success rates in increasing R&D efficiency at industry-level, in our experience the greatest optimization opportunities can be very different from one company to the next. As such, we do not advocate for a single 'industry standard' approach. Instead, we believe it is essential that efforts to improve R&D efficiency are customized for each company — and are informed by retrospective analysis of terminations, governance decision-making, KPIs, R&D functional objectives and incentives, and other factors. Importantly, data-driven productivity and benchmarking analyses should serve as a starting point, highlighting the critical focus areas for investigation and facilitating generation of key questions and hypotheses to explore.

In the following sections we highlight some generic strategies to optimize success rates (Figure 5); this is by no means an exhaustive list and the relevance of each will vary for different companies.

Enhancing the quality of early R&D outputs

The importance of the decisions that are taken during discovery research is often underestimated when considering the overall health of in-house R&D. Discovery-stage decisions such as target or candidate selection will collectively have critical downstream impacts on late-stage portfolio prospects and overall efficiency. Even the best R&D strategies will not be enough to save a project if the biological target selected during discovery research proves to be unsuitable for the indications of interest.

Although unanticipated project or asset liabilities will continue to present considerable challenges to R&D programs for the foreseeable future, important steps have been — and continue to be — taken to improve researchers' abilities to predict future program risks. Through earlier identification of safety, efficacy, pharmacokinetic/pharmacodynamic and other potential liabilities, leading companies have been able to increasingly 'front-load' attrition into early stages of R&D in order to minimize its impact on overall R&D efficiency, in cost per approval terms.

Figure 5: Key opportunities to optimize R&D success rates

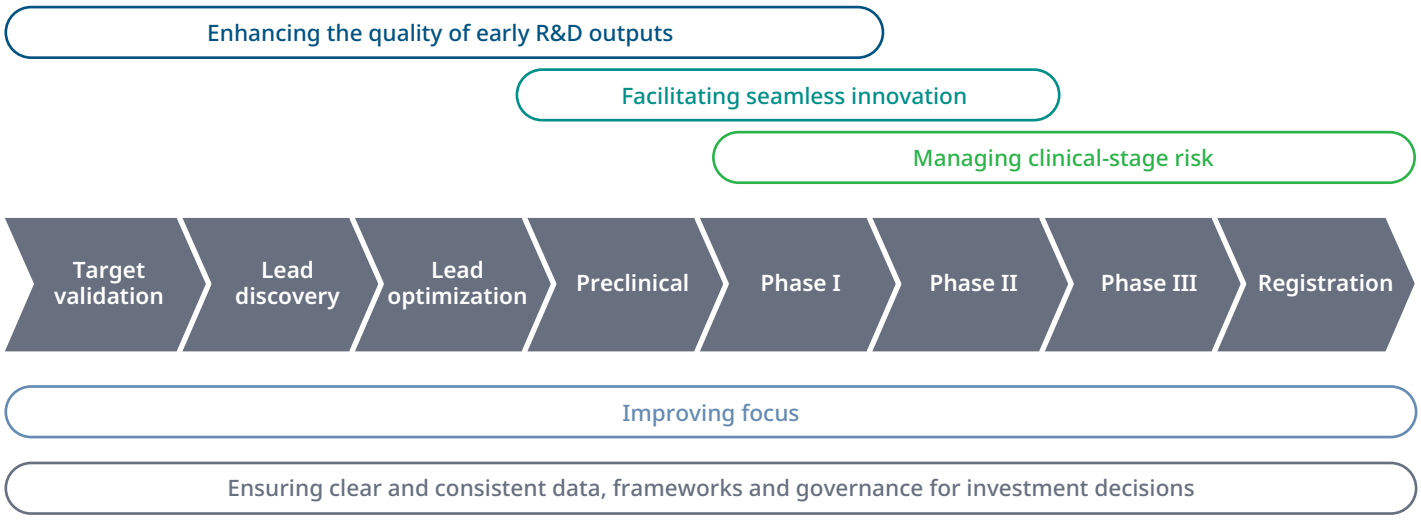

Many pharmaceutical majors have noted the importance of genetic evidence in early R&D decision-making, since focusing on genetically-validated drug targets has the potential to double overall success rates.[9] Whilst some of these targets may not historically have been very amenable to traditional drug discovery, novel emerging modalities are expected to enable prosecution of previously-intractable opportunities, thereby improving associated success rates. Efforts to improve the predictive validity of disease models used in drug discovery may also prove fruitful in further front-loading attrition into discovery research.[10]

Several major pharmaceutical companies have formally implemented approaches to identify signals of likely clinical efficacy in or prior to early clinical studies. Such 'proof-of-mechanism' approaches (distinct from traditional clinical proof-of-concept) have proven to be very helpful in predicting future success.[7,11]

Companies have also improved their ability to predict toxicological liabilities. Historically, investigational monoclonal antibody therapeutics displayed considerably higher early clinical success rates than their small molecule counterparts; however, the success rates of these two modalities have converged over time as the impact of off-target liabilities, associated with less-specific small molecules, has diminished.

Since the huge potential of artificial intelligence in pharmaceutical R&D has been widely discussed elsewhere, we will not go into detail here. Nonetheless, most if not all pharmaceutical majors we have analysed have utilized AI approaches at scale to improve the efficiency of drug discovery and development. Early R&D has been a particularly amenable area of focus and AI has contributed to all of the areas discussed above, informing decision-making from target-selection, through lead generation and beyond.

"We cut back the pipeline significantly and really focused on the quality of our research and development, not on the quantity or volume."

— Mene Pangalos, AstraZeneca EVP BioPharmaceuticals R&D[12]

Facilitating seamless innovation

Given the phase-to-phase interdependencies in portfolio attrition, we believe that it is important to understand and manage the impact of attrition holistically across the entire R&D process. In most major pharmaceutical companies, discovery and preclinical research activities sit within a distinct R&D function, often with Phase I and sometimes Phase II activities included. The R&D baton is then handed to a separate development organization for late-stage R&D (at Phase II, Proof of Concept or Phase III). While there are good arguments in favour of these independent structures, they do present a risk to overall R&D efficiency and — in particular — to a company's ability to manage end-to-end attrition effectively.

There is no shortage of criticism in the peer reviewed literature of the problems caused by a historical emphasis on 'shots-on-goal' R&D, or the idea that maximizing the volume of R&D programs progressing was an effective way to ensure productivity.[13,14] Even so, volume-based delivery objectives remain commonplace in early R&D functions — and individuals are incentivized accordingly. In our view, it is essential to link such volume objectives to quality thresholds in order to minimize the risk of simply delaying likely project failures into later stages of R&D which may be organizationally separate and where the efficiency impact of attrition is much greater. Good decisions — including good termination decisions — are more important than shots-on-goal.

In large organizations, quantitative analyses conducted to support project decision-making in a portfolio context (e.g. PTRS assessments and project valuations) are often heavily focused on near-term commercial opportunities, with early-stage R&D receiving only light-touch support.

While there are understandable practical reasons for this (e.g. paucity of evidence; resourcing constraints; and investor and C-suite focus on near-term assets), this can lead to sudden changes in perspective as projects transition into later-stage development and are subjected to more-rigorous analytics.

Commercial functions have an important role to play. Although it can be difficult to provide a meaningful forecast for an early-stage project that may be a decade away from market (and likely to fail on the way), low commercial expectations are a significant source of mid-to-late-stage attrition. Pragmatic but regular engagement spanning early- and late-stage R&D can help to minimize the delayed attrition that might otherwise result from sudden shocks as improved forecasting becomes available.

Highly efficient companies often have structures or processes in place to bridge or eliminate the gap between early-stage and late-stage R&D functions, avoiding siloed R&D objectives, minimizing short-termism and promoting a pan-portfolio view. Portfolio decision-making committees, where they exist separately for early- and late-stage functions, are connected — often through overlapping senior level membership. Sufficiently empowered R&D portfolio management or strategic functions are well-placed to play a key role in facilitating seamless innovation, providing an independent overarching view of the R&D pipeline and contributing organizational context to early- and late-stage decision-making committees and individual project teams.

Managing clinical-stage risk

As outlined above, the ideal R&D engine will front-load as much portfolio attrition as possible into early R&D, thereby delivering higher quality (and higher probability of success) substrate into clinical development stages. Nonetheless, it of course remains important to continue to identify, understand and discharge risk (or discontinue development) as early as possible as assets progress through clinical stages.

While not the optimally efficient solution, it is certainly very possible for companies with low Phase I success rates to remain high-performers overall in cost-per-approval terms — as long as their late-stage portfolios are selectively enriched with high probability of success assets as a consequence. But even the best-performing companies still lose around half of their R&D programs to Phase II attrition — and Phase II success rates for most companies are considerably worse than this.

Efficacy and safety reasons account for the majority of program terminations in mid-to-late-stage clinical development, but clinical trial characteristics can also have substantial impact. For example: difficulties in identifying, enrolling or retaining patients, perhaps due to unrealistic inclusion/exclusion criteria or high complexity and patient burden; insufficient investment resulting in a suboptimal trial design, ambiguous outcomes and the unanticipated need for additional intermediate-stage trials. Although none of these examples will come as a surprise to clinical development organizations, they remain active and very real risks which companies must continue to mitigate with suitable operational practices.

" One of the things that we saw from prior portfolio iterations is if you don't have good Phase II data that can inform things like dose and patient population for Phase III, your chance of success in Phase III is pretty low."

— Joshua Smiley, Lilly SVP and CFO[15]

" In the 3 years leading up to 2017, Pfizer's Phase II success rate, which is defined by successful transition into Phase III was 17%, which was well below the industry median and put us in the bottom quartile. I'm proud to say that, today, we have tripled our Phase II success rates on a 3-year rolling average, going from 17% in 2017, to 47% in 2019, to currently 53% in 2020 year-to-date, and we're now among industry leaders in this metric."

— Mikael Dolsten, Pfizer CSO[16]

Programs that are progressed to Phase III should have had sufficient risk discharged in prior phases to be considerably more likely to achieve regulatory approval than to fail; a high failure rate of programs in Phase III or after regulatory submission will dramatically impact a company's overall cost-per-approval — as well as commercial expectations and investor sentiment for its near-term pipeline.

Improving focus

When highlighting productivity improvements to the investment community, pharmaceutical CEOs and R&D leaders have often cited an improved pipeline focus as an important factor in transforming their innovation engines. In practice, this usually reflects a reduction in the number of therapeutic areas being pursued and — as such — can result in substantial organizational change. Company context is of course critical in interpreting portfolio refocusing, but typically a tighter focus offers benefits such as: reduced investment in low productivity areas; more dedicated expertise in priority areas; and concentration of investment on fewer projects.

Investment analysts have heavily criticized some companies in the past for having too little investment and too little expertise spread too thinly across too many projects, meaning that individual projects were not always sufficiently funded to enable fast, high-quality decision-making. Consequently, some R&D programs could languish in-phase or be dependent on sub-optimal evidence for progression decisions.

Abandoning or scaling back a low-productivity therapeutic area can be catalysed by internal underperformance factors, but is more typically driven by low expectations for probability of success or commercial returns for the therapeutic area overall. This is best exemplified by neuroscience — which many major companies abandoned or 'virtualized' over the last 10–15 years due to inherently risky R&D — and infectious diseases, which has generally been a challenging area for commercial reasons. Importantly, both neuroscience and infectious diseases illustrate the impact that emerging science, technology and evolving societal needs can (and should) have on pharmaceutical investment, since both areas are now seeing a resurgence. Several companies have recently re-entered or newly prioritized work in these spaces, demonstrating the importance of external as well as internal context in portfolio rationalization.

Investment analysts have heavily criticized some companies in the past for having too little investment and too little expertise spread too thinly across too many projects.

Ensuring clear and consistent data, frameworks and governance for investment decisions

While a detailed discussion of R&D portfolio management and governance is beyond the scope of this article, it is important to emphasize the importance of certain key aspects of this in considering success rates, attrition distribution and ultimately R&D efficiency.

Decision-making committees have a responsibility not just to ensure investment is directed towards progression of the most promising programs, but also to ensure that the discontinuation of ultimately-doomed projects and assets is achieved at the earliest possible stage. As discussed elsewhere in this article, the cost-benefit of 'front-loading' terminations has a substantial portfolio-level efficiency impact which goes beyond simple consideration of the opportunity cost released upon termination of an individual project.

Robust, structured, transparent and timely decision-making is critical for R&D efficiency — whether those decisions lead to progression of R&D projects or to terminations. Clear, consistent and pre-defined expectations for data, evidence and analyses — including contextual portfolio analyses — must be well-understood by all contributors. Decision-making frameworks and key decision-facilitating analyses should be mandatory, well-known and well understood by project teams and must be applied consistently and fairly across projects. Portfolio analysis, portfolio management, decision science and biostatistics representatives should be seen as important contributors and facilitators of robust progression/termination decision-making, providing key inputs to governance committees, project managers, project leads and their teams and minimizing the potential for bias.

Although there is no one-size-fits-all approach, we have found that large and achievable cost efficiency opportunities exist for the majority of companies. Cost-per-output modelling can be highly valuable in enabling identification of key success rate and other efficiency levers on a company-specific basis.

Conclusions

Between-phase success rates are a critical factor in determining R&D efficiency, having a major influence on the overall R&D cost of each new product approval. Although the value and tractability of success rate improvements varies from company to company, dramatic improvements have been demonstrated by several organizations as a result of systematic transformation efforts.

Actions that increase success rates must however be approached thoughtfully; interdependencies, knock-on effects and potential unintended consequences must be carefully considered — including the risk of counterproductive postponement of attrition to later stages of development. Failure to do so can also result in exacerbation of underlying problems, for example by increasing pressure to reduce project progression hurdles in response to higher-than-expected attrition. Interdependencies beyond attrition should also be considered, such as impacts on cycle times or on per-project progression costs.

Efficiency-transforming activities may have consequences for the distribution of R&D resources. A company that successfully redistributes its attrition through effective predictive approaches and front-loading of terminations into early R&D will evolve a different portfolio shape upon reaching steady-state; more early-stage programs and fewer late-stage programs will eventually be required for each new approval. Mid-to-long-term implications for R&D resourcing must therefore be considered, with discovery-stage activities needing to scale-up to ensure that front-loaded attrition does not result in too few projects to sustain the late-stage pipeline.

Although there is no one-size-fits-all approach, we have found that large and achievable cost efficiency opportunities exist for the majority of companies. Cost-per-output modelling can be highly valuable in enabling identification of key success rate and other efficiency levers on a company-specific basis. The resultant findings may then be utilized to prioritize and guide further in-depth diagnostic analyses, informing customized, company-specific recommendations and ultimately unlocking large cost benefits for R&D organizations.

References

1. Inflation Reduction Act of 2022. Congress.gov (2022) https://www.congress.gov/bill/117th-congress/house-bill/5376/text
2. Reform of the EU pharmaceutical legislation. European Commission (2023) https://health.ec.europa.eu/medicinal-products/pharmaceutical-strategy-europe/reform-eu-pharmaceutical-legislation_en
3. Sullivan SD. Medicare drug price negotiation in the United States: implications and unanswered questions. Value Health 26, 394-399 (2023)
4. Paul SM, et al. How to improve R&D productivity: the pharmaceutical industry's grand challenge. Nature Reviews Drug Discovery 9, 203-214 (2010)
5. A new look at R&D productivity: How small changes to key levers can dramatically impact cost per output. IQVIA white paper (February 2023) https://www.iqvia.com/library/white-papers/a-new-look-at-rd-productivity-how-small-changes-key-levers-can-dramatically-impact-cost-per-output
6. Chit A, et al. The Opportunity Cost of Capital. Inquiry 52, 1-5 (2015)
7. Morgan P, et al. Impact of a five-dimensional framework on R&D productivity at AstraZeneca. Nature Reviews Drug Discovery 17, 167-181 (2018)
8. Fernando K, et al. Achieving end-to-end success in the clinic: Pfizer's learnings on R&D productivity. Drug Discovery Today 27, 697-704 (2022)
9. Nelson MR, et al. The support of human genetic evidence for approved drug indications. Nature Genetics 47, 856–860 (2015)
10. Scannell JW, et al. Predictive validity in drug discovery: what it is, why it matters and how to improve it. Nature Reviews Drug Discovery 21, 915-931 (2022)
11. Morgan P, et al. Can the flow of medicines be improved? Fundamental pharmacokinetic and pharmacological principles toward improving Phase II survival. Drug Discovery Today 17, 419-424 (2012)
12. Brackley P. How AstraZeneca achieved a fivefold increase in its new medicines success rate. Cambridge Independent (January 2019)
13. Cook D, et al. Lessons learned from the fate of AstraZeneca's drug pipeline: a five-dimensional framework. Nature Reviews Drug Discovery 13, 419-431 (2014)

14. Ringel M, et al. Does size matter in R&D productivity? If not, what does? Nature Reviews Drug Discovery 12, 901-902 (2013)

15. Lilly. Morgan Stanley 18th Annual Global Healthcare Conference (September 2020)

16. Pfizer. Investor Day (September 2020)

CONTACT US
Ann Baker
Vice President, R&D Transformation Consulting
ann.baker@iqvia.com

Alex Roland
Associate Principal, R&D Transformation Consulting
alexander.roland@iqvia.com

William Fox
Consultant, R&D Transformation Consulting
william.fox2@iqvia.com

iqvia.com