Economics Bulletin

Volume 0, Issue 0

Comparison of the costs of drug development under different methods of calculation

Olga Rozanova
Gaidar Institute for Economic Policy; HSE University (ICEF)

Abstract

This paper compares two methods for estimating drug development costs: the conventional cost per expected approval and the probabilistic cost to achieve at least one approval with a given confidence level. The paper shows that, for typical success rates and high confidence levels, the probabilistic approach requires 2–4 times more trials—implying that standard average-cost figures can underestimate reliable investment by a similar factor. Extending the model to multiple phases, the analysis demonstrates that while the gap between the two objectives narrows over time, stochastic transitions often leave portfolios underfunded for subsequent phases. Thus, expected-value estimates carry a high risk of failing to meet confidence-based portfolio goals, with direct implications for research funding and policy design.

I thank the Editor and the anonymous Referee for their valuable comments, which have helped improve this paper. I also thank peers in the field for discussions that inspired the comparative focus of this work. The usual disclaimers apply.
Citation: Olga Rozanova, (2026) "Comparison of the costs of drug development under different methods of calculation", *Economics Bulletin*, Volume 0, Issue 0, pages 15-24
Contact: Olga Rozanova - olga.v.rozanova@gmail.com.
Submitted: November 19, 2024. **Published:** April 27, 2026.

1 Introduction

The process of developing a therapeutic drug is usually long and expensive. This process is risky since not all clinical trials result in a successful outcome. Conventional estimates of drug development costs account for expenditures on unsuccessful clinical trials (i.e., "failures").[1] Typically, the cost of developing one new therapeutic entity is calculated as the spending on the number of trials that, on average, yield exactly one approval. This expected-value approach is used in foundational works such as DiMasi et al. (2003) and Adams and Brantner (2006), as well as in many other studies.[2]

An alternative methodology estimates costs by determining the expenditures needed to guarantee at least one success with a specified confidence level (e.g., 90%). For example, Veugelers and Zachmann (2020) calculate the number of trials required to obtain at least one or at least three COVID-19 vaccines with 99% probability under optimistic and pessimistic scenarios. Related work includes Gouglas et al. (2018), who employ a stochastic optimization model to estimate the minimal costs of successfully advancing at least one vaccine per disease through Phase 2a trials. Glickman et al. (2006) estimate the probability of developing at least one antituberculosis drug from the project portfolio existing at the time and simulate how this probability changes if the number of trials doubles. More recently, Athey et al. (2020) derive the optimal number of COVID-19 vaccine trials and, as an intermediate step, illustrate how the probability of at least one success depends on the number of candidates. Rey et al. (2022) likewise model the probability of at least one success as a function of the number of candidates. In this framework, the cost of developing at least one new entity with a target probability can be obtained by multiplying the required number of candidates by the average cost of one clinical trial.

Although both approaches — expenditures per expected success and expenditures per success with a given confidence — appear in the literature, their results have not, to the best of my knowledge, been systematically compared.[3] Consequently, it is unclear how cost estimates for developing a new entity differ depending on the methodological choice. This paper aims to fill this gap.

Comparing these estimates clarifies the true economic commitment required for drug development. Such analysis can inform regulators about the scale of investment needed to achieve a successful outcome with a target confidence level. In turn, this knowledge supports policy aligned with the United Nations Sustainable Development Goal 3 (Good Health and Well-Being), specifically target 3.b, which calls for supporting research and development of medicines for diseases affecting developing countries and ensuring access to affordable treatments.[4]

[1] Although most studies in this literature account for unsuccessful trials, exceptions exist. For example, Light et al. (2009) estimate the costs of developing two specific vaccines—RotaTeq (Merck) and Rotarix (GlaxoSmithKline)—using public information and interviews with senior informants. Their approach tallies the number of subjects in each clinical phase for these successful vaccines, but does not include costs from trials for rotavirus vaccines that ultimately failed.

[2] Other works that estimate development costs based on the expenditures required for one success in expectation include Hansen (1979), DiMasi et al. (1991, 2004, 2007), Adams and Brantner (2010), and Vernon et al. (2010).

[3] Partial comparisons exist. For instance, Veugelers and Zachmann (2020) calculate the number of trials needed to guarantee at least three COVID-19 vaccine approvals with 99% confidence and note that this number yields four successes in expectation.

[4] *Sustainable Development Goal 3b: "Support the research and development of vaccines and medicines for the communicable and non-communicable diseases that primarily affect developing countries, provide access to affordable essential medicines and vaccines" (UN General Assembly, 2015).* A clearer under-

To calculate the costs of developing at least one successful drug, I adopt a framework similar to those of Fernandez et al. (2012), Snyder et al. (2023), and Veugelers and Zachmann (2020), all of whom model success probabilities under the assumption that clinical trials are statistically independent.[5]

My approach generalizes this framework by treating success rates as continuous parameters rather than fixed numbers, and by systematically comparing the probabilistic guarantee approach with the expected-value approach across a range of success rates. I derive analytic expressions for: (1) the number of trials needed to achieve one success in expectation, and (2) the number of candidates required to ensure at least one success with a target confidence level. The first expression depends solely on the success rate; the second depends on both the success rate and the desired confidence level. I then plot the ratio of these two quantities across a range of success rates and confidence levels. This visualization directly illustrates how the number of trials required for a probabilistic guarantee diverges from the number needed for an expected success. Because development costs are proportional to the number of trials, the same figure also reveals the corresponding discrepancies in estimated expenditures under the two methodological approaches.

Beyond comparing the overall number of programs required to achieve one success in expectation versus at least one approval with a target confidence, I analyze how the gap between these two objectives evolves across successive development phases — particularly as programs advance toward approval. I further examine cases where the actual number of programs transitioning to the next phase diverges from the expected value due to the inherent stochasticity of the development process. I then compare the expected value and confidence level targets from the perspective of compounding portfolio risk throughout the development process.

I also calculate the probability of achieving at least one success when the number of trials is calibrated for one expected success ($\tilde{N}$). For aggregate success rates (p) between 6.9% and 21.6%, this probability does not exceed 70%.[6] To guarantee at least one success with 90% confidence, the required number of trials ($\hat{N}$) is approximately double $\tilde{N}$.[7] Consequently, for typical success rates, expenditures required for a reliable (90% confidence) outcome are significantly higher than those based on the expected-value approach. This analysis clarifies the resource commitment for drug development and, therefore, helps identify the size of financial incentives regulators should provide to achieve a successfully developed therapeutic drug.

The remainder of the paper proceeds as follows. Section 2 derives the expressions for the required number of trials under both the expected value and confidence-level frameworks. Building on this foundation, Section 3 compares these expressions and analyzes the evolution of the gap between the two objectives through the development process. Section 4 concludes.

---

standing of development costs directly contributes to strategic planning for this target.

[5]For instance, Fernandez et al. (2012) compute "the probability of at least one success among 150 independent programs." Snyder et al. (2023) likewise assume independence in their probability formula for at least one success. Veugelers and Zachmann (2020) also calculate probabilities of at least one (or three) successes under independence. While some studies incorporate correlated trials (e.g., Wu et al. 2012; Dai et al. 2025), the present analysis maintains the independence assumption.

[6]According to different estimates of the aggregated success rates of the clinical trials (either for all indications or for the lead indications) they are in the range from 6.9% to 21.6% (see Wong et al. 2019; Thomas 2016; Hay et al. 2014). It is worth noting, however, that disaggregated estimates of the success rates of the clinical trials (e.g. probability of success for some particular therapeutic groups) can be well outside this range (see Wong et al. 2019).

[7]Again, these calculations are provided for the success rates in the range from 6.9% to 21.6%.

2 Number of clinical trials

2.1 Number of clinical trials needed to get one success on average

Assume that the probability of getting a successful outcome out of each clinical trial is constant and equal to p. That is, p is the probability that the entity successfully advances from Phase 1 to approval. If N clinical trials are conducted, then on average (i.e., in expectation) there are $p \cdot N$ successes as a result of these clinical trials. The expected number of successes is equal to one, if

$$\tilde{N} = \frac{1}{p}. \tag{1}$$

2.2 Number of clinical trials needed to get one success with a given confidence level

Assume there are N clinical trials. Let us find the probability that as a result of these clinical trials there will be at least one successful outcome.

- Since p is the probability of success of each clinical trial, then $(1-p)^N$ reflects the probability that none of the clinical trials has a successful outcome.
- Therefore $[1-(1-p)^N]$ is the probability that at least one clinical trial has a successful outcome.

Let $\hat{N}$ denote the number of clinical trials that guarantee **at least** one success with a given confidence level α. $\hat{N}$ is found from the condition:[8]

$$1-(1-p)^{\hat{N}} = \alpha. \tag{2}$$

Thus:

$$\hat{N} = \frac{\ln(1-\alpha)}{\ln(1-p)}. \tag{3}$$

2.2.1 Confidence level of getting at least one success if $N = \tilde{N}$

Assume the number of clinical trials is equal to the one that guarantees exactly one successful outcome on average (i.e.,in expectation). That is, $\tilde{N} = \frac{1}{p}$. Let us determine the probability with which this number of clinical trials results in at least one success.

As discussed earlier, for any number of clinical trials N, the probability of at least one success is equal to $[1-(1-p)^N]$. So, if $N = \tilde{N} = \frac{1}{p}$, the probability of at least one success

[8] While the actual number of clinical trials conducted is an integer, the theoretical quantities $\frac{1}{p}$ and $\frac{\ln(1-\alpha)}{\ln(1-p)}$ are continuous. To maintain analytical tractability and simplify the exposition, these exact expressions rather than their integer-ceiling counterparts are used. This practice of working with continuous values is well-established in methodological literature — for example, in economic analyses where the number of firms is treated as continuous (Mukherjee 2005). A direct analogue in pharmaceuticals is found in the standard cost-of-development literature. DiMasi et al. (2003), in their widely cited study, estimate the cost per approved drug (C) by dividing the cost per investigational program (c) by the probability of success (ρ): $C = \frac{c}{\rho}$. Here, $\frac{1}{\rho}$ represents the average number of programs required to yield one approval. For instance, they calculate capitalized costs of US $ 100.4 million per investigational drug and, with a success rate $\rho = 0.215$, derive a cost per approved drug of approximately US $ 467 million (DiMasi et al. 2003, Section 5.4). Crucially, they treat the number of programs needed, $\frac{1}{\rho}$, as a continuous variable rather than an integer, precisely mirroring the approach that is adopted in the current model.

is equal to $[1-(1-p)^{\frac{1}{p}}]$. According to Wong et al. (2019), the probability of success of each clinical trial p ranges from 6.9% to 21.6%. This implies that in order to get one success on average the number of clinical trials should be in the range from $\frac{1}{0.216} \approx 4.63$ to $\frac{1}{0.069} \approx 14.69$. Given these values for the probability of success (p) the chances of getting at least one success when the number of clinical trials is equal to $\tilde{N} = \frac{1}{p}$, are in the interval from $(1-(1-0.069)^{\frac{1}{0.069}}) \approx 0.645 = 64.5\%$ to $(1-(1-0.216)^{\frac{1}{0.216}}) \approx 0.676 = 67.6\%$.

Thus, to ensure that the probability of at least one success exceeds 70%, more than $\tilde{N}$ clinical trials are required.

3 Comparison of the number of clinical trials under two approaches

Let us now compare the number of clinical trials that are needed in order to get one success on average (i.e. $\tilde{N}$) with the number of clinical trials that ensures at least one success with some target confidence level ($\hat{N}$).

To compare $\tilde{N}$ and $\hat{N}$ let us consider the ratio $\frac{\hat{N}}{\tilde{N}}$. Based on the expressions for $\tilde{N}$ and $\hat{N}$ in the previous section the ratio of interest is:

$$\frac{\hat{N}}{\tilde{N}} = \frac{p \cdot \ln(1-\alpha)}{\ln(1-p)}. \tag{4}$$

It is worth noting that the ratio $\frac{\hat{N}}{\tilde{N}}$ is also informative about the comparison of the spending on the clinical trials under the two approaches (i.e., when the number of clinical trials ensures exactly one success in expectation or when it ensures **at least** one success with some given confidence level). To see this point, assume that the cost of one clinical trial from Phase 1 to approval is equal to I on average, then the spending on $\tilde{N}$ and $\hat{N}$ clinical trials are equal to $I \cdot \tilde{N}$ and $I \cdot \hat{N}$, respectively. So, the ratio of the expenditures on clinical trials under the two approaches is equal to the ratio of the number of clinical trials (i.e. $\frac{I \cdot \hat{N}}{I \cdot \tilde{N}} = \frac{\hat{N}}{\tilde{N}}$).

Figure 1 plots this ratio $\frac{\hat{N}}{\tilde{N}}$ as a function of p for various levels of α.

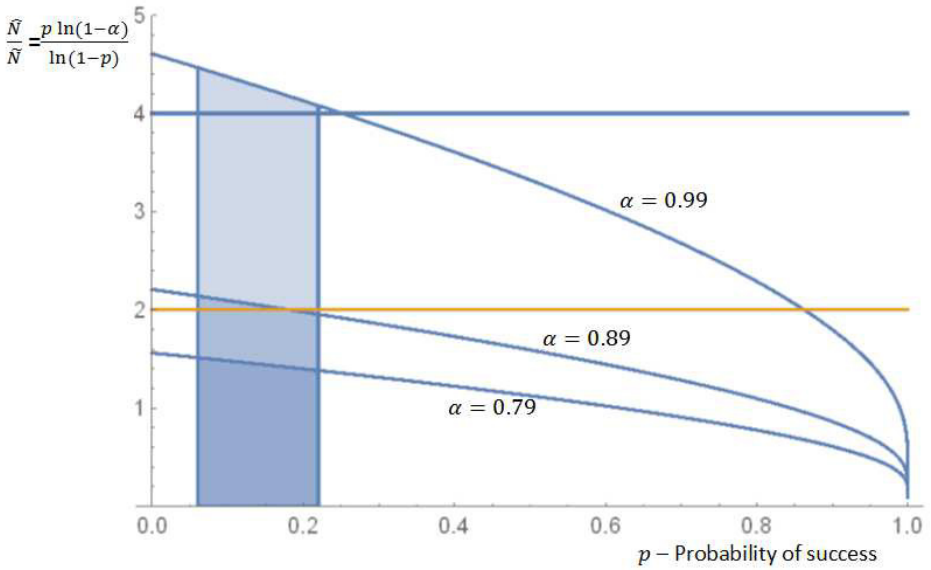

Figure 1: Ratio of the numbers of clinical trials

From Figure 1 it can be seen that for $p \in (0.069, 0.216)$ and $\alpha = 0.89$ (respectively, $\alpha = 0.99$) $\hat{N}$ is about two times (respectively, four times) higher than $\tilde{N}$. It means that spending on the clinical trials that ensure **at least** one success with probability 0.89 (respectively, 0.99) is about two times (respectively, four times) higher than the expenditures on the clinical trials that guarantee one success in expectation.

The analysis demonstrates that a conventional approach to measuring the costs of drug development (i.e., clinical spending needed to get exactly one success in expectation) reports a value that may turn out to be significantly lower than the expenditures needed to guarantee at least one success with a target confidence level. For the success probability range reported in the literature (6.9% to 21.6%) and for confidence levels of 90% or higher, this implies that public support for drug development must be substantially larger if the regulator's goal is to achieve at least one success with a given confidence level, rather than one success in expectation.

3.1 Dynamics of gap across clinical phases

This section analyzes the dynamics of the gap between the "expected value" and "confidence-level" objectives, focusing on two key drivers: how the gap evolves across development phases and how portfolio deviations affect it.

Evolution of the gap across phases

Assume each clinical phase is started with a portfolio size calculated for one expected success (i.e. one approval). The question is: whether the discrepancy between the two strategic goals (i.e. one "success" on average and at least one "success" with a target confidence level) becomes more or less pronounced throughout the drug development lifecycle?

As Figure 1 shows, the ratio $\frac{\hat{N}}{\tilde{N}}$ decreases as the probability of success, p, increases. This implies that the gap between the number of trials needed for "at least one success with a given confidence level" and "one success in expectation" narrows in later phases of drug development.[9] Consequently, the strategic error incurred by targeting the simpler "expected value" goal, rather than the confidence-based goal, diminishes as programs approach approval.

For illustration, consider typical phase transition probabilities from the literature ($p_{12} = 0.64$, $p_{23} = 0.32$, $p_{3A} = 0.49$). As shown in the Appendix (Section A, Table I and Figure A1), the ratio $\hat{N}/\tilde{N}$ declines from approximately 2.1 at Phase 1 to 1.6 at Phase 3 for $\alpha = 0.89$, confirming that the gap narrows as programs advance.

Impact of portfolio deviations

Let us now examine how the gap between the two strategic goals—ensuring at least one approval with confidence α versus achieving one approval in expectation—evolves when the actual number of programs deviates from the expected-value target: Because

[9] In Figure 1 for most values of the success probability p, the ratio $\frac{\hat{N}}{\tilde{N}}$ exceeds 1. However, for certain combinations of the confidence level α and p, the ratio falls below 1. In this latter scenario, while the ratio continues to decrease with increasing p, this now corresponds to a **divergence** between $\hat{N}$ and $\tilde{N}$, not a convergence. Because the ratio is greater than 1 for most parameter values — and specifically for the success probabilities reported in many studies at common confidence levels (e.g., 0.89 and above)— the focus of the current paper is on the case where $\frac{\hat{N}}{\tilde{N}} > 1$. This simplifies interpretation: a **decrease** in the ratio signifies convergence of the estimates, while an **increase** signifies divergence.

transition phase-to-phase probabilities are less than 1, the number of programs entering Phase j (M_j) is a random variable. Even if a firm starts Phase 1 with $\tilde{N}_1$ programs, it may enter Phase 2 with fewer or more than $\tilde{N}_2$ programs by chance. So, the question of interest is how the ratio $R_j = \hat{N}_j/(\tilde{N}_j + y)$ behaves not only for planned values ($y = 0$) but also for random realizations $y \neq 0$.

Let Phase j have a success probability of $\rho = p_{jA}$.[10] The number of programs expected to yield exactly one approval at this Phase is $\tilde{N}_j = 1/\rho$, while the number needed to guarantee at least one approval with probability α is $\hat{N}_j = \frac{\ln(1-\alpha)}{\ln(1-\rho)}$. In practice, stochastic outcomes cause the actual portfolio size at phase j to differ from $\tilde{N}_j$. It is therefore expressed as $\tilde{N}_j + y$, where the integer offset y can be negative (deficit), zero (on target), or positive (surplus).

The ratio

$$R_j = \frac{\hat{N}_j}{\tilde{N}_j + y} \tag{5}$$

measures how the confidence-based requirement compares to the actual portfolio. I analyze how R_j changes as ρ increases (i.e., as programs advance toward approval). For $y = 0$, R_j decreases with ρ (see Figure 1 and Appendix Section A), meaning the goals converge. The general case $y \neq 0$ is examined via the derivative $\partial R_j/\partial\rho$.

Substituting the expressions for $\hat{N}_j$ and $\tilde{N}_j$ gives

$$R_j = \frac{\ln(1-\alpha)}{\ln(1-\rho)\left(\frac{1}{\rho} + y\right)}. \tag{6}$$

Differentiating and simplifying yields

$$\frac{\partial R_j}{\partial \rho} = \frac{\ln(1-\alpha)}{\left[\ln(1-\rho)(\frac{1}{\rho} + y)\right]^2 \rho} \left(\frac{1-\rho y}{1-\rho} + \frac{1}{\rho}\ln(1-\rho)\right). \tag{7}$$

Since $\ln(1-\alpha) < 0$ for $0 < \alpha < 1$ and the denominator outside the parentheses is positive, the sign of $\partial R_j/\partial\rho$ is determined by

$$S(y,\rho) \equiv -\left(\frac{1+\rho y}{1-\rho} + \frac{1}{\rho}\ln(1-\rho)\right). \tag{8}$$

Sign analysis:

1. For $y \geq 0$ (actual portfolio meets or exceeds the expected-value target), $\partial R_j/\partial\rho < 0$. *Justification:* At $y = 0$, $S(0,\rho) < 0$ (see Figure 1). Because $\partial S/\partial y = -\rho/(1-\rho) < 0$, $S(y,\rho)$ decreases with y and remains negative for all $y \geq 0$. Hence, R_j falls as ρ increases—the confidence-based requirement becomes easier to satisfy relative to the actual portfolio.

2. For $y \leq -1$ (actual portfolio falls at least one program short of the expected-value target), $\partial R_j/\partial\rho > 0$. *Justification:* At $y = -1$,

$$S(-1,\rho) = -\left(1 + \tfrac{1}{\rho}\ln(1-\rho)\right) > 0,$$

because $\frac{1}{\rho}\ln(1-\rho) < -1$ for $\rho \in (0,1)$ (see Figure 2). Since $S(y,\rho)$ decreases in y, it stays positive for all $y \leq -1$. Thus, R_j rises with ρ—the shortfall relative to the confidence target widens even as success probability improves.

[10] p_{jA} denotes the probability that the program successfully passes from Phase j till approval.

These results reveal a critical asymmetry:

- When the portfolio is at or above the expected-value size ($y \geq 0$), progression through later stages (higher ρ) naturally aligns the two strategic objectives.
- When the portfolio is deficient ($y \leq -1$), however, further progression *amplifies* the misalignment: the required number of programs for confidence α grows relative to the actual count, making the expected-value objective an increasingly poor proxy for the confidence-based goal. A phase-specific illustration of this widening gap is provided in Appendix C (Figure C1).

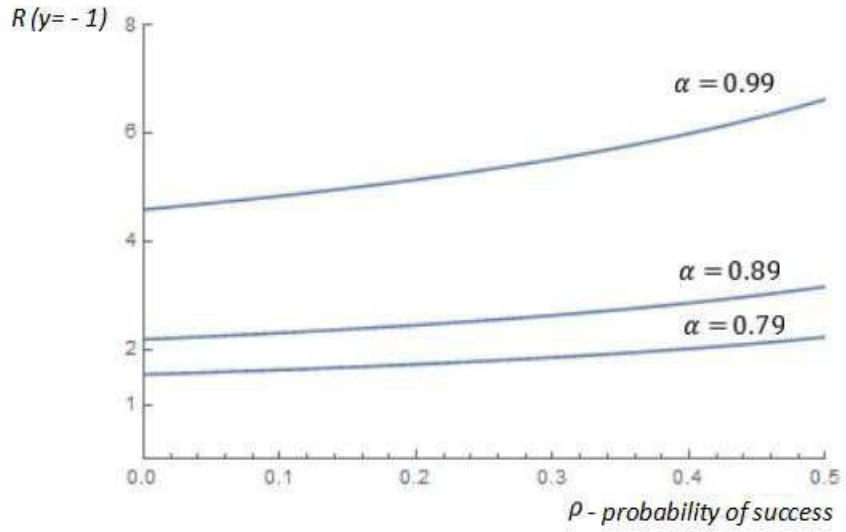

Figure 2: Ratio of the numbers of clinical trials for $y = -1$

Note: For $y = -1$, the ratio is $R = \frac{\ln(1-\alpha)}{\left(\frac{1}{\rho}-1\right)\ln(1-\rho)}$. In this figure, ρ ranges over $[0, 0.5]$. The case $\rho > 0.5$ is not relevant here because $\frac{1}{\rho} - 1 < 1$, implying fewer than one program would, on average, enter the next phase. Since at least one program is required for development to continue, the analysis is restricted to $\rho \leq 0.5$.

Integer constraint robustness

The main analysis treats the number of trials as a continuous variable for analytical tractability. Appendix Section B examines how imposing integer constraints affects the results. While the constrained ratio $\hat{N}/\tilde{N}$ becomes non-monotonic, its overall decreasing trend—and therefore the qualitative conclusion that the gap narrows with higher success probabilities—remains unchanged. Similarly, the increasing pattern of $\hat{N}/(\tilde{N}-1)$ persists under integer constraints, confirming that the continuous approximation does not drive the core findings.

3.2 Practical implications for method selection

The results derived in Section 2 and the comparative analyses in Section 3 provide objective criteria for choosing between the expected-value and probabilistic guarantee approaches in real-world financing and planning. The probabilistic method requires approximately two to four times more trials—and correspondingly higher expenditure—than the expected-value method to achieve high confidence (Figure 1).

Two dynamic effects shape how the two targets relate across the development life cycle:

- Under portfolio shortfalls ($y \leq -1$), the gap between the objectives widens as programs advance, making the expected-value target an increasingly poor proxy for confidence-based planning (Section 3.1).
- Under normal progression, as success probabilities rise in later phases, the ratio $\hat{N}/\tilde{N}$ decreases (see Appendix Section A, Figure A1), meaning the required confidence buffer shrinks and the expected-value estimate becomes a closer approximation to the confidence-level requirement.

Even when Phase 1 portfolios are calibrated exactly to each target, the expected-value approach carries a greater-than-50% risk of downstream deficit (see Appendix Section D, Table II), whereas the probabilistic target is more robust to phase-to-phase attrition.

Consequently, the probabilistic approach seems to be better suited for risk-averse, mission-critical contexts where a minimum probability of success must be ensured (e.g. public funding for high-priority diseases or pandemic vaccine development). In contrast, the expected-value approach remains useful for historical benchmarking, long-term portfolio budgeting, and diversified R&D planning, where the focus is on average returns across many projects over time rather than guaranteeing a specific outcome.

These considerations can inform staged funding strategies: the confidence-based buffer is most crucial in early, high-risk stages, while expected-value estimates become more adequate for late-phase budgeting as the inherent risk diminishes.

4 Conclusion

This paper has systematically compared two fundamental methodologies for estimating the resource requirements of drug development: the conventional approach, which calculates the cost per expected approval, and the probabilistic approach, which determines the cost to secure at least one approval with a target confidence level.

The analytical relationship between these two metrics ($\tilde{N}$ and $\hat{N}$) is derived. It is demonstrated that the ratio $\hat{N}/\tilde{N}$ can be substantial, often ranging from 2 to 4 for typical success probabilities and high confidence levels. This means that widely cited average cost figures can underestimate the necessary investment for a reliable outcome by a factor of two or more.

Furthermore, the analysis is extended to a multi-phase context, revealing two critical dynamic insights. First, the gap between the two objectives narrows as programs advance, meaning the strategic error of using the average-cost target is most severe in the risky early stages of development. Second, due to the inherent randomness of clinical transitions, even a portfolio perfectly sized for one expected approval faces a significant chance of entering the next phase with a deficit. Crucially, when such a deficit occurs, the misalignment between the actual number of programs and a confidence-level target widens with progression into later phases.

Several limitations should be acknowledged. The analysis assumes statistical independence across clinical trials and simultaneous trial launch, whereas in practice outcomes may be correlated and trials initiated at different times. Second, the analysis assumes that all trials in the portfolio have the same probability of success. In reality, even within the same therapeutic area, success rates vary across drug types (e.g., small molecules, specialty drugs, GLP-1 agonists). The analysis also treats the probability of success as exogenous, whereas it may be affected by investments in trial design and execution.

Readers should therefore interpret the results as illustrative under the assumption of homogeneous success probabilities. The relevant conclusion for a real portfolio depends on the actual mix of drug types — a dimension not captured here. Future research could extend this analysis by relaxing the assumptions of homogeneous success probabilities and exogenous trial outcomes, as well as by incorporating correlated outcomes across trials.

References

[1] Adams, C. P., & Brantner, V. V. (2006). Estimating the cost of new drug development: is it really $802 million?. *Health affairs*, 25(2), 420-428.

[2] Adams, C. P., & Brantner, V. V. (2010). Spending on new drug development 1. *Health economics*, 19(2), 130-141.

[3] Athey S., Baker A., Castillo J.C., Glennerster R., Kremer M., Lee J., Snyder C., Tabarrok A. and Tan B.J. (2020). Incentive design to accelerate a COVID-19 vaccine. Available at https://drive.google.com/file/d/1JPNiV7LGnSUUDcXY0sK3VdgpJnR2XH2t/view. Accessed 6 Jan. 2026.

[4] Dai, Y., Lo, A. W., Singh, M., Xu, Q., & Zhang, R. (2025). Estimating Correlations Between Clinical Trial Outcomes Using Generalised Estimating Equations. *Oxford Bulletin of Economics and Statistics*.

[5] DiMasi, J. A., Grabowski, H. G., & Vernon, J. (2004). R& D costs and returns by therapeutic category. *Drug Information Journal*, 38(3), 211-223.

[6] DiMasi, J. A., & Grabowski, H. G. (2007). The cost of biopharmaceutical R& D: is biotech different? *Managerial and decision Economics*, 28,469-479.

[7] DiMasi, J. A., Hansen, R. W., Grabowski, H. G., & Lasagna, L. (1991). Cost of innovation in the pharmaceutical industry. *Journal of health economics*, 10(2), 107-142.

[8] DiMasi, J. A., Hansen, R. W., & Grabowski, H. G. (2003). The price of innovation: new estimates of drug development costs. *Journal of health economics*, 22(2), 151-185.

[9] Fernandez, J. M., Stein, R. M., & Lo, A. W. (2012). Commercializing biomedical research through securitization techniques. *Nature biotechnology*, 30(10), 964-975.

[10] Glickman, S. W., Rasiel, E. B., Hamilton, C. D., Kubataev, A., & Schulman, K. A. (2006). A portfolio model of drug development for tuberculosis. *Science*, 311(5765), 1246-1247.

[11] Gouglas, D., Le, T. T., Henderson, K., Kaloudis, A., Danielsen, T., Hammersland, N. C., Robinson, J., Heaton, P., & Røttingen, J. A. (2018). Estimating the cost of vaccine development against epidemic infectious diseases: a cost minimisation study. *The Lancet Global Health*, 6(12), e1386-e1396.

[12] Hansen, R. W. (1979). The pharmaceutical development process: estimates of development costs and times and the effects of proposed regulatory changes. Center for the Study of Drug Development, Department of Pharmacology and Toxicology, University of Rochester Medical Center.

[13] Hay, M., Thomas, D. W., Craighead, J. L., Economides, C., & Rosenthal, J. (2014). Clinical development success rates for investigational drugs. *Nature biotechnology*, 32(1), 40-51.

[14] Light, D. W., Andrus, J. K., & Warburton, R. N. (2009). Estimated research and development costs of rotavirus vaccines. *Vaccine*, 27(47), 6627-6633.

[15] Mukherjee, A. (2005). Price and quantity competition under free entry. *Research in Economics*, 59(4), 335-344.

[16] Rey, A., Rozanova, O., & Zhuk, S. (2022). Estimating correlations between vaccine clinical trial outcomes. *Journal of Applied Statistics*, 49(13), 3392-3413.

[17] Snyder, C. M., Hoyt, K., & Gouglas, D. (2023). An optimal mechanism to fund the development of vaccines against emerging epidemics. *Journal of Health Economics*, 91, 102795.

[18] Thomas, D. W. (2016). Clinical development success rates 2006–2015. *BIO Industry Anal.*, 1, 16.

[19] United Nations General Assembly. (2015). Transforming our world: the 2030 Agenda for Sustainable Development (A/RES/70/1). https://sdgs.un.org/2030agenda

[20] Vernon, J. A., Golec, J. H., & Dimasi, J. A. (2010). Drug development costs when financial risk is measured using the Fama-French three-factor model. *Health economics*, 19(8), 1002-1005.

[21] Veugelers, R., & Zachmann, G. (2020). Racing against COVID-19: a vaccines strategy for Europe (No. 2020/07). Bruegel Policy Contribution.

[22] Wong, C. H., Siah, K. W., & Lo, A. W. (2019). Estimation of clinical trial success rates and related parameters. *Biostatistics*, 20(2), 273-286.

[23] Wu, S., Crespi, C. M., & Wong, W. K. (2012). Comparison of methods for estimating the intraclass correlation coefficient for binary responses in cancer prevention cluster randomized trials. *Contemporary clinical trials*, 33(5), 869-880.