Expected Value Thinking

Idea In Short

Judge uncertain proposals by multiplying each possible outcome by its probability, summing the results, and subtracting the cost. Require every investment case above a set threshold to show that figure along with its break-even probability, which is the chance of success at which the proposal stops paying. Then test the inputs, because the answer is only as good as the odds and payoffs behind it. Where a loss could threaten the firm, cap the exposure regardless of the average. Used this way, expected value turns debate about whether a bet feels attractive into a discussion of specific numbers that colleagues can challenge, adjust, and revisit as evidence arrives.

Every executive who approves a budget is placing a bet. The product might succeed or fail, the acquisition might deliver synergies or destroy value, and the pilot might scale or stall. Expected value is the oldest tool for comparing such bets on a common footing. It is simple enough to compute on a single page, and it exposes assumptions that otherwise hide inside enthusiasm or caution. This article explains where the idea came from, shows a worked calculation, and covers the conditions under which the average stops being a good guide.

The problem that started it

In 1654, the French gambler Chevalier de Méré put a question to Blaise Pascal about how to divide the stakes in a game that ended before either player won.1 Pascal wrote to Pierre de Fermat, and their letters produced a solution that still anchors probability theory. Both saw that a fair split depends not on the score at the moment of interruption but on the futures that remained open, each weighted by its likelihood.2 Today we call that weighted average an expected value. The shift in thinking matters for management as much as for mathematics. A project's history, including what the team has already spent, tells you little about what it is worth. Its possible futures tell you far more.

The calculation

An expected value is a potential event's magnitude multiplied by its probability, summed across every outcome.3 It describes the long-run average of a decision repeated many times.4 The tool has been used for centuries in gambling, insurance, and investment, and the arithmetic has not changed.

A short illustration shows how it works. Suppose a company can spend $2 million on a product extension. There is a 30 percent chance that it earns $10 million, and a 70 percent chance that it earns $0.5 million from salvaged components and licensing. The expected payoff is 0.3 times $10 million plus 0.7 times $0.5 million, which is $3.35 million. Subtract the $2 million cost and the expected value is $1.35 million. The figure does not predict what will happen. It says that, across many projects with these odds and payoffs, the company would average a gain of that size.

Finding the break-even probability

The more useful number is the probability at which the case stops paying. In the example, set the expected payoff equal to the $2 million cost and solve for the chance of the large outcome. The answer is about 16 percent. If the team believes the real chance is above that, the extension pays on average, and if it believes the chance is below, it does not. This reframes the debate. People no longer argue about whether the project will work. They ask whether the odds of success clear a specific bar, and which evidence would move the estimate across it.

Where the inputs come from

An expected value is only as sound as its probabilities and payoffs. Decision theorists distinguish objective probabilities, which come from known frequencies such as a lottery with one chance in 10,000, from subjective ones, which express a decision maker's degree of belief.5 Business decisions mostly fall in the second group. A team can improve these inputs in a few ways. It can use base rates from comparable projects, ask several people for independent estimates, and record a range instead of a point. It can also list payoffs for each branch separately, including the cost of failure, because teams tend to describe the success case in detail and the failure case in a sentence.

Costs deserve the same care. Timing matters, since a payoff three years away is worth less than one next quarter, and correlated risks matter when several branches fail together. A

model that treats each uncertainty as independent can look comfortable while hiding a single weak point that affects everything.

Testing the sensitivity of the answer

Once the first calculation is complete, a disciplined team varies each input across a plausible range and observes how the result responds. Most models turn out to depend heavily on one or two assumptions, typically the probability of the large payoff and the timing of the first revenue, while the remaining inputs barely change the conclusion. Identifying those dominant assumptions redirects research effort toward the evidence that could genuinely alter the decision, rather than toward refining figures that were never influential. The exercise also reveals the conditions under which a proposal reverses from attractive to unattractive, and those conditions can be written into the approval itself as explicit review triggers. A sponsor who understands that a delay of two quarters or a ten-point drop in the success probability would erase the expected gain will monitor those variables more carefully than one who has seen only a single headline figure.

Repeated bets and the law of large numbers

Expected value earns its keep when a decision repeats. A venture investor who funds forty companies, a lender who approves thousands of loans, and a retailer who tests dozens of promotions each year all see results converge toward the average. In that setting, picking positive expected value bets is a sound policy, even though individual bets fail often. The same logic supports stage gates in a pipeline, where each gate cuts losses early and keeps the portfolio average positive.

One-shot decisions behave differently. A firm that makes a single large acquisition does not live in the long run. It lives with one outcome. The average still organizes the analysis, but leaders should also look at the spread of possible results and ask whether they could survive the worst one.

When the average misleads

Two problems recur. The first is ruin. A bet with a high expected value can still carry a small chance of a loss that ends the enterprise, and no number of future wins repairs that. Leaders should cap exposures so that no single outcome threatens survival, whatever the

average says. The second problem is that money does not measure everything people care about. Decision theorists address this with expected utility, which replaces dollar amounts with the value a decision maker attaches to each result and so allows for risk aversion.

Reputation, regulatory standing, and employee trust all resist dollar conversion. A proposal that offers a modest average gain but risks a safety incident should not pass on the arithmetic alone.

Value in the decisions, not only the payoffs

Drug development offers an instructive case. Risk-adjusted net present value, which discounts future cash flows by the probability of each stage succeeding, is widely used in pharmaceutical portfolio management and venture capital. A recent paper argues that the method has a blind spot: it conflates wrong decisions to advance weak candidates with right decisions to advance strong ones, so it cannot value the tests that tell them apart.6 The lesson extends well beyond pharmaceuticals. In any staged project, part of the value lies in the information each stage produces and the option to stop. An expected value model that ignores those options undervalues cheap early tests and overvalues commitments made before the evidence arrives.

A routine for expected-value reviews

A one-page template works for most proposals. It forces the same questions each time and makes proposals comparable. The steps are:

1. List the possible outcomes, including at least one failure case and one partial success
2. Assign a probability to each and record the source or the reasoning
3. State the payoff or cost of each outcome, with its timing
4. Compute the expected value and the break-even probability
5. Identify the worst outcome and confirm the organization can absorb it

Reviewers then spend their time on the numbers that move the result. Often one probability and one payoff drive most of the answer, and those become the focus of further research. Over several quarters, compare predicted and realized outcomes to see where the team's estimates run high or low.

Limits of the method

Expected value is a way to organize judgment, not a replacement for it. It cannot repair a poorly framed question, and it cannot supply probabilities the team does not have. False precision is a real hazard, since a result such as $1.35 million invites more trust than the inputs deserve. Show ranges, run the calculation at the optimistic and pessimistic ends of each input, and watch for cases where the sign of the answer flips. When it flips, the proposal depends on beliefs that deserve more research before capital is committed.

- 1Problem of points
- 2How a Renaissance gambling dispute spawned probability theory
- 3Expected value (and beyond)
- 4Expected values (mathematical concept)
- 5The history and methodology of expected utility
- 6Tripartite models for estimating the value of drug candidates and decision tools

Summary

Expected value gives leaders a common yardstick for choices whose results cannot be known in advance. It works best when decisions repeat, the inputs are honest, and no single loss can end the enterprise. It works poorly when a proposal carries a threat of ruin, when probabilities are guesses dressed as data, or when the value lies in the decisions the project allows rather than in its final payoff. Keep the arithmetic simple, publish the assumptions, and treat the break-even probability as the main number to debate. Teams that do so waste less time on enthusiasm and spend more on the facts that move the result.