

Expected Value

[Read this note online](#)

Expected value is the probability-weighted average of every possible outcome of a decision. If you repeated the decision a thousand times, expected value tells you what you'd average per attempt. It's the simplest mathematical tool for thinking under uncertainty, and it beats pure intuition for most everyday choices.

The expected value formula is short, the calculation is fast, and most people who learn it start spotting bad bets they used to take and good bets they used to dismiss. The catch: expected value treats a 50% chance of losing your house the same as a 50% chance of losing a dollar. Survivability matters separately, and any serious decision framework has to layer that consideration on top of raw expected value.

Expected value needs outcomes defined on a consistent basis. A cost paid in every case must be included whether the favorable event happens or not.



Think of $E[X]$ as the balance point of the outcomes when weighted by their probabilities.

EXAMPLE 1: FAIR SIX-SIDED DIE

Let X be the outcome when a fair six-sided die is rolled.

Outcome x_i	Probability p_i	Product $p_i x_i$
1	1/6	1/6
2	1/6	2/6
3	1/6	3/6
4	1/6	4/6
5	1/6	5/6
6	1/6	6/6
Total	1	21/6

$$E[X] = 21/6 = 3.5$$

→ The long-run average outcome is 3.5.

EXPECTED VALUE

The expected value (mean) of a discrete random variable is the long-run average outcome.


$$E[X] = \sum_{i=1}^n p_i x_i$$

x_i = possible outcome
 p_i = probability of x_i
 n = number of outcomes



For a fair coin with +1 for H and -1 for T, the expected value is 0. (Break-even game)

VISUALIZING WEIGHTED OUTCOMES ($p_i x_i$)



$E[X]$ is the balance point (average) of these weighted contributions.

EXAMPLE 2: FAIR COIN GAME

Let $X = +1$ for Heads (H), -1 for Tails (T).

Outcome x_i	Probability p_i	Product $p_i x_i$
+1 (H)	1/2	$(+1)(1/2) = 1/2$
-1 (T)	1/2	$(-1)(1/2) = -1/2$
Total	1	0

$$E[X] = 0$$

→ On average, you break even.

USEFUL PROPERTIES

- $E[aX + b] = aE[X] + b$
- $E[c] = c$ (constant)
- $E[X_1 + X_2] = E[X_1] + E[X_2]$ (linearity)

Expected value formula and weighted outcomes bar chart modern textbook illustration

The Formula

For outcomes $x_1, x_2, \dots, x_n$ with probabilities $p_1, p_2, \dots, p_n$:

The probabilities must sum to 1. The outcomes can be money, time, units sold, customers retained, lives saved, anything you can put a number on.

For a continuous random variable with probability density function $f(x)$:

The integral version is the same idea expressed for the continuous case. Most everyday decisions live in the discrete world; the integral version becomes essential in finance, physics, and statistics where outcomes vary smoothly.

$$E[X] = \sum_{i=1}^n p_i \cdot x_i = p_1 x_1 + p_2 x_2 + \dots + p_n x_n$$

$$E[X] = \int_{-\infty}^{\infty} x \cdot f(x) dx$$

Worked Example: A Project Bid

You spend 5,000 dollars to prepare a bid. Winning the contract has a 40% probability and brings 30,000 dollars in profit before the bid-preparation cost. Lose the bid (60% chance) and you lose the 5,000 dollars in prep cost.

Positive expected value. If you face many similar bids over a year, the expected profit is 7,000 dollars per attempt under these assumptions. The math says bid.

The deeper insight: expected value also tells you the maximum you should rationally spend on a bid prep. If preparing this bid took 14,000 dollars instead of 5,000, the EV would flip negative ($0.4 \times 16,000 + 0.6 \times (-14,000) = 6,400 - 8,400 = -2,000$) and the math says walk away. Expected value isn't just a yes/no signal; it's a calibration for how much resource is worth committing.

$$E[\text{bid}] = 0.4 \times (30,000 - 5,000) + 0.6 \times (-5,000)$$

$$E[\text{bid}] = 10,000 - 3,000 = 7,000$$

Worked Example: an Insurance Premium

Insurance pricing is expected value made visible. A 500,000 dollar policy covers a 0.2% annual risk:

The actuarially fair premium is 1,000 dollars per year. Insurers charge more (say 1,500 dollars) to

fund overhead, profit, and the cost of capital they hold against catastrophic claim years. You pay because you're transferring catastrophic risk, not because the bet is mathematically favorable.

This is also why people buy insurance against rare, large losses (house fires, life insurance, liability) but not against frequent, small losses (lost umbrellas, broken phones). The EV math runs against the buyer in both cases, but the **variance** reduction is worth the premium when the loss could ruin you. Variance reduction is the product you're actually paying for.

$$E[\text{payout}] = 0.002 \times 500,000 = 1,000$$

Worked Example: A Casino Roulette Bet

American roulette has 38 pockets: 18 red, 18 black, 2 green (0 and 00). A bet on red pays 1:1. Probability of winning: $18/38 \approx 0.474$.

Every dollar bet on red has an expected value of -5.3 cents. Over many spins, the casino's edge accumulates with statistical certainty. Casino games exist precisely because the house engineered a negative-EV game and the gambler's variance lets occasional winners exist long enough to feed the long-term losers.

This is also why advantage gambling, card counting in blackjack, poker against weaker opponents, sports betting against soft lines, works only when the player creates a positive EV through skill or information. Expected value is the line between game-of-chance entertainment and a job.

$$E[\text{bet}] = \frac{18}{38} \times 1 + \frac{20}{38} \times (-1) = -\frac{2}{38} \approx -5.3\%$$

Where Expected Value Quietly Fails

Two games:

Both have $E = 100,000$. Most people prefer Game A. You can't pay rent with expected value. Expected value treats outcomes as interchangeable as long as the average matches. It doesn't see ruin.

This is why **risk assessment** pairs expected value with variance, the Kelly Criterion, and utility functions. Expected value alone is the floor of decision math, not the ceiling.

The St. Petersburg Paradox makes the failure stark. Imagine a coin-flip game where you win 2^n dollars if the first heads occurs on flip n . The expected value is infinite ($\sum 2^n \cdot 2^{-n} = \sum 1 = \infty$), yet nobody would pay even 100 dollars to play. Expected value alone doesn't capture what humans actually value. Daniel Bernoulli proposed the resolution in 1738: people optimize expected utility, not expected value, and utility curves flatten as wealth grows.

Game A: guaranteed 100,000 dollars. **Game B:** 10% chance of 1,000,000 dollars, 90% chance of 0

dollars.

Expected Utility: the Standard Fix

Expected utility theory replaces raw outcomes x_i with utility values $u(x_i)$, then computes $E[u(X)] = \sum p_i u(x_i)$. A common choice for u is the natural logarithm:

The logarithmic utility function captures diminishing marginal value. The first 100,000 dollars matters more than the second 100,000 dollars, which matters more than the third. A risk-averse decision maker maximizes $E[\ln(W)]$ instead of $E[W]$.

For Game A vs Game B above, with starting wealth 100,000 dollars: $u_A = \ln(200,000) \approx 12.21$ vs $u_B = 0.1 \times \ln(1,100,000) + 0.9 \times \ln(100,000) \approx 11.75$. Game A wins on expected utility, matching the intuitive preference.

Expected utility is the formal foundation for risk-adjusted decision making in finance, insurance, and economics. The Kelly Criterion for optimal bet sizing falls out of maximizing logarithmic expected utility.

$$u(W) = \ln(W)$$

Variance: the Other Half of the Picture

Two decisions can share the same expected value and have wildly different risk profiles. The dispersion of outcomes, measured by variance and standard deviation, is what separates them.

For two investments with $E = 10,000$, one with $\sigma = 2,000$ and another with $\sigma = 10,000$, the second carries five times the risk for the same expected return. Sharpe ratio formalizes this idea: return per unit of risk. See the companion notes on [variance and standard deviation](#) and the [normal distribution](#) for the mechanics.

Quick Mental Checklist

This six-step pass takes two minutes and catches most obviously bad decisions. For high-stakes calls, layer in variance, downside protection, and expected utility.

List every outcome you can think of, including the bad ones. Estimate a probability for each. Make probabilities sum to 1. Multiply each outcome by its probability. Sum the results. Positive E means the bet is favorable on average. Before acting, ask: can I survive the worst single outcome? If not, expected value isn't enough. If outcomes vary widely, also estimate variance or worst-case downside before committing.

Where Expected Value Shows Up Beyond Money

Expected value's power is that it scales. The same one-line formula works for a roulette bet, a venture capital portfolio, a vaccine policy, or a recommendation algorithm. The scaling fails only when you forget to ask whether the average outcome is enough, and when one bad draw can end the game.

Marketing. Expected revenue per customer = conversion probability $\times$ order value. Drives bid caps in paid search and decisions about whether a campaign is worth running. **Hiring.** Expected value of a candidate = probability of strong performance $\times$ upside – probability of poor performance $\times$ cost of replacement. **Engineering reliability.** Expected downtime = sum of (failure probability $\times$ time to recover) across components. **Healthcare.** Expected quality-adjusted life years (QALYs) for a treatment combine survival probability with quality-of-life adjustments. **Game design.** Loot tables, encounter rewards, and matchmaking algorithms are built on expected value calculations to balance progression and challenge. **Public policy.** Cost-benefit analysis is just expected value applied to civic decisions, with all the strengths and limitations that implies.

Expected Value in Marketing and Advertising

Every paid-ads platform, Google Ads, Meta Ads, LinkedIn Ads, bids on the expected value of a click or impression. The auction sets the price you pay; your expected revenue per click sets your maximum bid. Bid above EV and you lose money on every click. Bid too far below and you lose impressions to competitors.

The formula:

For a 3% conversion rate, \$80 average order, and 30% margin, the max bid is $0.03 \times 80 \times 0.3 = 0.72$ dollars. Anything above that and the campaign loses money on average. Smart marketers compute EV per channel, per audience, and per creative, then optimize accordingly.

$$\text{Max bid per click} = \text{conversion rate} \times \text{average order value} \times \text{margin}$$

Expected Value in Venture Capital

VC portfolios are built on expected value with extreme variance. A typical seed-stage portfolio has an expected return of 2 to 3 $\times$ the fund, but most individual investments return zero. The EV math depends on a small number of huge winners (10 $\times$, 100 $\times$, 1000 $\times$) carrying the entire fund.

If you invest \$1M each in 30 startups, expected outcomes might look like: 18 fail (return 0), 7 return 2 $\times$, 3 return 5 $\times$, 1 returns 50 $\times$, 1 returns 200 $\times$. EV per investment $\approx (0.6 \times 0 + 0.23 \times 2 + 0.1 \times 5 + 0.033 \times 50 + 0.033 \times 200) \approx 9.3$ million. Total fund EV ≈ 280 M on \$30M deployed. The math works only

because of the extreme tail; ignore it and you'd never make seed-stage bets.

Expected Value vs Median Outcome

For symmetric distributions, expected value (mean) and **median** agree. For skewed distributions, they can be wildly different. Lottery tickets have positive expected value if the jackpot is large enough, but the median outcome is "lose your dollar." Startup investments have positive expected value, but the median outcome is "lose your money."

For one-shot decisions, the median often matters more than the mean. For repeated decisions, the mean is what you converge to. Knowing whether you're playing a one-shot or a repeated game determines which statistic to anchor on.

Expected Value in Insurance Reserves

Insurance companies set aside reserves to cover expected future claims. The reserve is the present value of expected payouts: $R = \sum_t \frac{E[\text{payout}_t]}{(1+r)^t}$ where r is the discount rate.

For a life insurance policy, the expected payout depends on mortality tables (probability of death at each age) multiplied by the policy face value, summed over the policyholder's remaining expected lifetime. Actuarial science is largely the discipline of computing these expected values reliably across portfolios of millions of policies.

Underestimating expected payouts leads to insurer insolvency. Overestimating leads to uncompetitive premiums. Regulators enforce minimum reserve calculations precisely because the expected-value math is what stands between policyholders and the company's promise to pay.

Expected Value in Game Theory

In games of strategy, expected value is the standard way to evaluate mixed strategies, strategies that randomize between actions. The Nash equilibrium concept relies on each player's strategy maximizing their own expected payoff given the others' strategies.

Rock-paper-scissors has a unique Nash equilibrium where each player picks each option with probability 1/3. Any deviation creates a positive-EV exploitation opportunity for the opponent. Game theorists use this kind of EV reasoning to model auctions, negotiations, voting, and competitive market behavior.

Reinforcement learning extends the same idea to sequential decision making. The Bellman equation is essentially an expected-value recursion: the value of a state equals the immediate reward plus the discounted expected value of the next state.

Expected Value of Perfect Information

Decision theorists compute the *expected value of perfect information* (EVPI), how much a decision maker should pay to eliminate all uncertainty before deciding. EVPI is the difference between the expected value of the best decision under uncertainty and the expected value of always making the best decision in each scenario.

If perfect knowledge would only marginally improve outcomes, paying for information isn't worth it. If EVPI is large, investing in research, surveys, or testing pays off. This is the formal basis for whether to commission a market study, run a pilot, or buy data, answer "is the EVPI greater than the cost?"

Common Pitfalls When Using Expected Value

Overconfidence in probability estimates. EV is only as good as the probabilities you plug in. Wide confidence intervals on probabilities translate to wider confidence on EV. **Ignoring opportunity cost.** A bet with positive EV is still a bad bet if a better positive-EV alternative is available. **Not accounting for transaction costs.** Trading commissions, taxes, and time costs all eat into expected value. The headline EV is rarely the realized EV. **Treating one-shot decisions as repeated.** EV converges over many trials. For a single decision with a heavy downside, the expected value can be misleading even when positive. **Mixing time horizons.** A "positive EV per day" decision might still bankrupt you if the variance is daily and your bankroll is monthly.

The related [OpenStax mathematics textbooks](#) provide additional explanations and exercises for the mathematical definitions used here.

Keeping Costs Consistent

A project has a 40% chance of earning Rs. 30,000 before a Rs. 5,000 preparation cost. Winning leaves Rs. 25,000; losing leaves -Rs. 5,000. Expected net profit is $0.4 \times \text{Rs. } 25,000 + 0.6 \times (-\text{Rs. } 5,000) = \text{Rs. } 7,000$.

The same expression is $0.4 \times \text{Rs. } 30,000 - \text{Rs. } 5,000$. This shorter form works because the preparation cost is certain. If that cost rises to Rs. 14,000, expected net profit becomes -Rs. 2,000. Positive expected value describes an average under the assumptions; it does not guarantee a gain on the next attempt or account for every risk preference.

The points to check in this example are:

- Outcome values should all be gross or all be net on the same basis.
- Probabilities over the complete outcome set sum to 1.

- A long-run-average interpretation needs relevant assumptions about repeated trials.

Check Your Understanding

Try the questions before reading the answers.

Question 1. What preparation cost breaks even here?

Answer. Rs. 12,000.

Question 2. Can an outcome have negative value in an expected-value calculation?

Answer. Yes. Losses must be included with their signs.

Frequently Asked: How EV Connects to Probability and Variance

Expected value is the first moment of a probability distribution. Variance is the second moment around the mean. Together they describe the location and spread of any random variable. Higher moments (skewness, kurtosis) describe asymmetry and tail-heaviness, but EV and variance are the workhorses for everyday decision making.

The relationship between EV and probability is one-way: probability fully determines EV (via integration or summation), but EV alone doesn't determine probability. Two distributions with the same mean can have wildly different shapes. That's why pairing EV with variance, or with the full distribution, gives a much more useful picture than EV alone.

Practice Questions

Work each question before reading its solution. The set runs from direct recall and substitution to the applied questions that exams actually use to separate grades. The downloadable PDF includes these questions and their solutions.

Question 1. Define expected value and give its formula for a discrete random variable.

Solution. $E[X] = \sum x_i P(x_i)$: each outcome weighted by its probability, then summed. It is not the most likely outcome nor a guaranteed result; it is the long-run average if the random process repeated indefinitely.

Question 2. A fair 6-sided die is rolled. Find its expected value.

Solution. $E[X] = \frac{1}{6}(1 + 2 + 3 + 4 + 5 + 6) = \frac{21}{6} = 3.5$. Notice 3.5 is not even a possible outcome, underlining that expected value is a long-run average, not a prediction for any single roll.

Question 3. A lottery ticket costs 100. It pays 100,000 with probability $1/2000$, and nothing otherwise. Find the expected value of buying one, and interpret the sign.

Solution. $E = 100,000 \times \frac{1}{2000} + 0 \times \frac{1999}{2000} = 50$. Net expected value of playing: $50 - 100 = -50$. Negative expected value means that, on average, every ticket loses money, exactly how lotteries fund their prizes and profit.

Question 4. A roulette wheel has 38 slots (18 red, 18 black, 2 green). A 10 bet on red pays 10 if red comes up, loses if not. Find the expected value.

Solution. $E = 10 \times \frac{18}{38} + (-10) \times \frac{20}{38} = \frac{180-200}{38} \approx -0.53$. The 2 green slots are the house's edge: without them the game would be perfectly fair ($E = 0$). That -0.53 per 10 bet, about 5.3%, is the house edge baked into every spin.

Question 5. An insurer sells a policy for 500/year. Claims occur with probability 0.02 and cost the insurer 20,000 when they do. Find the insurer's expected profit per policy.

Solution. Expected payout: $0.02 \times 20,000 = 400$. Expected profit: $500 - 400 = 100$ per policy. Insurance is a business built entirely on positive expected value from the insurer's side, funded by policyholders individually accepting negative expected value in exchange for reduced variance (protection against catastrophic loss).

Question 6. A game pays 3x your bet if you roll a 6 on a fair die, and you lose your bet otherwise. Is this game fair, favorable, or unfavorable to the player? Compute E for a 10 bet.

Solution. $E = 30 \times \frac{1}{6} + (-10) \times \frac{5}{6} = 5 - 8.33 = -3.33$. Unfavorable: a fair payout for a 1-in-6 event would need 5x, not 3x, the bet. "Fair" games have $E = 0$; this one systematically favors the house.

Question 7. Two job offers: Job A pays a guaranteed 60,000. Job B pays 50,000 base plus a bonus of 40,000 with probability 0.4, else no bonus. Compare expected values and explain why a risk-averse person might still choose Job A.

Solution. Job B's expected value: $50,000 + 0.4 \times 40,000 = 66,000$, higher than A's guaranteed 60,000. But expected value ignores variance: a risk-averse person with real bills to pay might prefer certainty over a higher average that comes with a 60% chance of only 50,000. Expected value is not the whole decision, only 1 input to it.

Question 8. A dartboard-style game charges 5 to play. You win 50 with probability 0.05, 10 with probability 0.2, and nothing otherwise. Find the expected value and the expected profit for the operator per play.

Solution. Player's expected payout: $50(0.05) + 10(0.2) + 0(0.75) = 2.5 + 2 = 4.5$. Operator collects 5 and expects to pay out 4.5, netting an expected 0.50 profit per play. Every carnival game, arcade claw machine, and casino slot is tuned by exactly this arithmetic to guarantee long-run house profit.

Question 9. A factory produces items with a 3% defect rate. Each defect costs 200 to fix. If 10,000 items are made, find the expected total repair cost, and explain why this number is useful for budgeting despite no single item costing exactly that.

Solution. Expected defects: $10,000 \times 0.03 = 300$. Expected cost: $300 \times 200 = 60,000$. No individual production run will hit exactly 60,000 in repairs, but averaged across many production runs, actual costs cluster tightly around this figure (by the law of large numbers), making it a reliable number for setting aside a repair budget.

Question 10. Explain the St. Petersburg paradox: a game pays 2^n if the first heads appears on flip n . Show the expected value is infinite, and explain why almost no one would pay more than a small amount to play.

Solution. $E = \sum_{n=1}^{\infty} 2^n \times \left(\frac{1}{2}\right)^n = \sum_{n=1}^{\infty} 1 = \infty$: the formula genuinely diverges. Yet virtually nobody would pay even 20 to play, because the probability of any large payout shrinks so fast (needing 20+ consecutive tails first) that the "average" is dominated by astronomically unlikely events no real decision-maker weighs that way. The paradox reveals that raw expected value can behave very differently from rational human risk preference, which is part of why utility theory was developed.

FAQs

Is expected value the same as the average?

It's a probability-weighted average, not a simple average. If outcomes have different probabilities, weight each outcome by how likely it is before averaging. A simple average treats every outcome as equally likely, which is rarely true in real decisions.

When should I not use expected value?

When one outcome can ruin you. Expected value assumes you can absorb any single loss and keep playing. For one-shot decisions with catastrophic downsides, use expected utility, the Kelly Criterion, or explicitly model survival probability before deciding.

How do I estimate probabilities I don't have data for?

Start with reference classes, find similar past situations and use their base rates. Use wide ranges (10 to 40%, not 25%) until evidence narrows them. Track your predictions over time and check calibration. Bad probability estimates still beat ignoring probability entirely.

Does positive expected value guarantee profit?

No. It guarantees long-run average profit if you can repeat the decision indefinitely with independent outcomes. Single trials can lose money even with positive EV. The whole point of variance and risk management is to handle the gap between expected and realized outcomes.

What's the difference between expected value and expected utility?

Expected value treats money linearly. Expected utility applies a curve (often logarithmic) that captures the fact that the first 100,000 dollars matters more than the second 100,000 dollars. Expected utility models risk aversion; expected value doesn't.

How does expected value relate to fair odds?

Fair odds are the odds at which expected value equals zero. If a bet has 25% win probability and pays 3:1, the EV is zero, those are fair odds. Bookmakers and casinos build their margins by offering odds slightly worse than fair, creating a small negative EV per bet that compounds across all customers.

Can expected value be negative?

Yes, and a negative EV usually signals you shouldn't take the bet, unless you're paying for variance reduction (insurance) or entertainment (lottery, casino). Identify what you're buying when you accept negative EV; if it isn't insurance, variance reduction, or entertainment, walk away.

How does expected value apply to sports betting?

By comparing your estimated win probability against the bookmaker's implied probability. If your model says a team wins 55% of the time and the line implies 50%, the EV is positive and you should bet. Most casual bettors lose money because they don't model probabilities and the bookmaker's vig adds about 4 to 5% negative EV to every bet.

What's the connection between expected value and the law of large numbers?

The law of large numbers says sample averages converge to the expected value as sample size grows. Expected value is the long-run average; the law of large numbers is the guarantee that long-run averages actually approach it. Together they justify why EV is a meaningful planning tool for repeated decisions.

Why do casinos always win in the long run?

Because every casino game is engineered with negative expected value for the player. Even a 2% house edge guarantees the casino wins as the number of bets grows. Players see variance, wins and losses bouncing around, but the expected value drift wins over enough trials. Casinos are essentially operating large-sample statistical machines.

How is expected value used in poker?

Every betting decision is an EV calculation. Calling a bet has positive EV when (probability of winning $\times$ pot size) exceeds the cost of the call. Pro players think in pot odds and implied odds, both EV-based concepts. Poker is essentially a long-form EV optimization game with imperfect information.

What's the difference between EV and ROI?

EV is an absolute expected outcome (e.g., \$9,000 profit per bid). ROI is a ratio (e.g., 180% return on \$5,000 invested). Both are useful: EV for raw decision making, ROI for comparing opportunities of different sizes. They give the same go/no-go answer but different rankings when capital is constrained.

What is the expected value of perfect information?

EVPI is how much you'd theoretically pay to eliminate all uncertainty before deciding. It equals the expected value of acting optimally with full knowledge minus the expected value of acting optimally under uncertainty. If EVPI exceeds the cost of gathering information, it's worth getting the data.