

PROBABILITY

Expected Value

The probability-weighted average, why casinos always win, and calculating it for real decisions, with a 10-question practice set and full solutions.

SUBJECT	LEVEL	INCLUDES
Mathematics	High school and early college	Practice set, answer key, revision sheet

Read this note online, with updates: gauravtiwari.org/study-notes/expected-value

INSIDE

- 1 The Formula
- 2 Worked Example: A Project Bid
- 3 Worked Example: An Insurance Premium
- 4 Worked Example: A Casino Roulette Bet
- 5 Where Expected Value Quietly Fails
- 6 Expected Utility: The Standard Fix
- 7 Variance: The Other Half of the Picture
- 8 Quick Mental Checklist
- 9 Where Expected Value Shows Up Beyond Money
- 10 Expected Value in Marketing and Advertising
- 11 Expected Value in Venture Capital
- 12 Expected Value vs Median Outcome
- 13 Expected Value in Insurance Reserves
- 14 Expected Value in Game Theory
- 15 Expected Value of Perfect Information
- 16 Common Pitfalls When Using Expected Value
- 17 Practice Questions
- 18 Answer Key and Revision Sheet

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.



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 $p_1, p_2, \dots, p_n$:

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

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 $f(x)$:

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

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.

Worked Example: A Project Bid

You spend 5,000 dollars to prepare a bid. Win the contract (40% chance) and you net 30,000 dollars in profit. Lose the bid (60% chance) and you lose the 5,000 dollars in prep cost.

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

Positive expected value. If you face many similar bids over a year, you average 9,000 dollars per attempt. 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.

$$E[\text{bid}] = 12,000 - 3,000 = 9,000$$

Worked Example: An Insurance Premium

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

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

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.

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: $18/38 \approx 0.474$.

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

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.

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.

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:

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

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.

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.

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.

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:

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

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.

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× the fund, but most individual investments return zero. The EV math depends on a small number of huge winners (10x, 100x, 1000x) carrying the entire fund.

If you invest \$1M each in 30 startups, expected outcomes might look like: 18 fail (return 0), 7 return 2x, 3 return 5x, 1 returns 50x, 1 returns 200x. 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 $\approx 280\text{M}$ 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.

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.

Q1. 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.

Q2. 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.

Q3. 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.

Q4. 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.

Q5. 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).

Q6. 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.

Q7. 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.

Q8. 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.

Q9. 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.

Q10. 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.

Answer Key

QUICK ANSWERS

Q1 $E[X] = \sum x_i P(x_i)$ **Q2** 3.5 **Q3** $E = 50$; net -50 per ticket **Q4** $E \approx -0.53$ per bet **Q5** expected profit = 100/policy

Q6 unfavorable; $E \approx -3.33$ **Q7** B has higher E (66,000 vs 60,000), but A has zero risk **Q8** player $E = 4.5$; operator expects +0.50/play **Q9** expected cost = 60,000 **Q10** $E = \infty$, yet real willingness-to-pay stays small: EV alone can mislead

Revision Sheet

Definition

$E[X] = \sum x_i P(x_i)$: the probability-weighted average outcome.

Insurance logic

Insurer needs positive EV on average; buyer accepts negative EV to reduce catastrophic variance.

Not a prediction

$E[X]$ need not be a possible outcome (die: 3.5) and rarely equals any single trial's result.

EV is not the whole story

Ignores risk/variance; a lower-EV certain outcome can rationally beat a higher-EV gamble.

Games and gambling

House edge = negative player EV, built in deliberately (roulette green slots, lottery odds).

Law of large numbers

EV becomes a reliable planning number only when averaged over MANY repeated trials.

About These Notes

This note is part of a free study-notes series on gauravtiwari.org covering mathematics, physics, chemistry, and biology: the concepts that keep deciding grades in high school, board, and entrance exams.

2008

Writing and teaching online since

2,000+

Articles published on gauravtiwari.org

125+

Free study notes in this series

M.Sc.

Mathematics, with physics and chemistry training

Gaurav Tiwari is a developer, educator, and founder of Gatilab. The study-notes series exists because most exam prep material is either a full textbook or a thin definition page, and revision needs something in between: one concept, complete, with the traps named and the practice attached.

GET THE FULL SERIES

Every note in the series is free to read online and free to download as a PDF like this one. Browse the complete collection at gauravtiwari.org/free-study-notes-exam-prep and pick the topics your next exam actually covers.

USE IT FREELY

This PDF is free for personal study, classrooms, and study groups. Share the file or the link with anyone it can help. The one thing not allowed is reselling it or republishing it as your own work.



Gaurav Tiwari

gauravtiwari.org