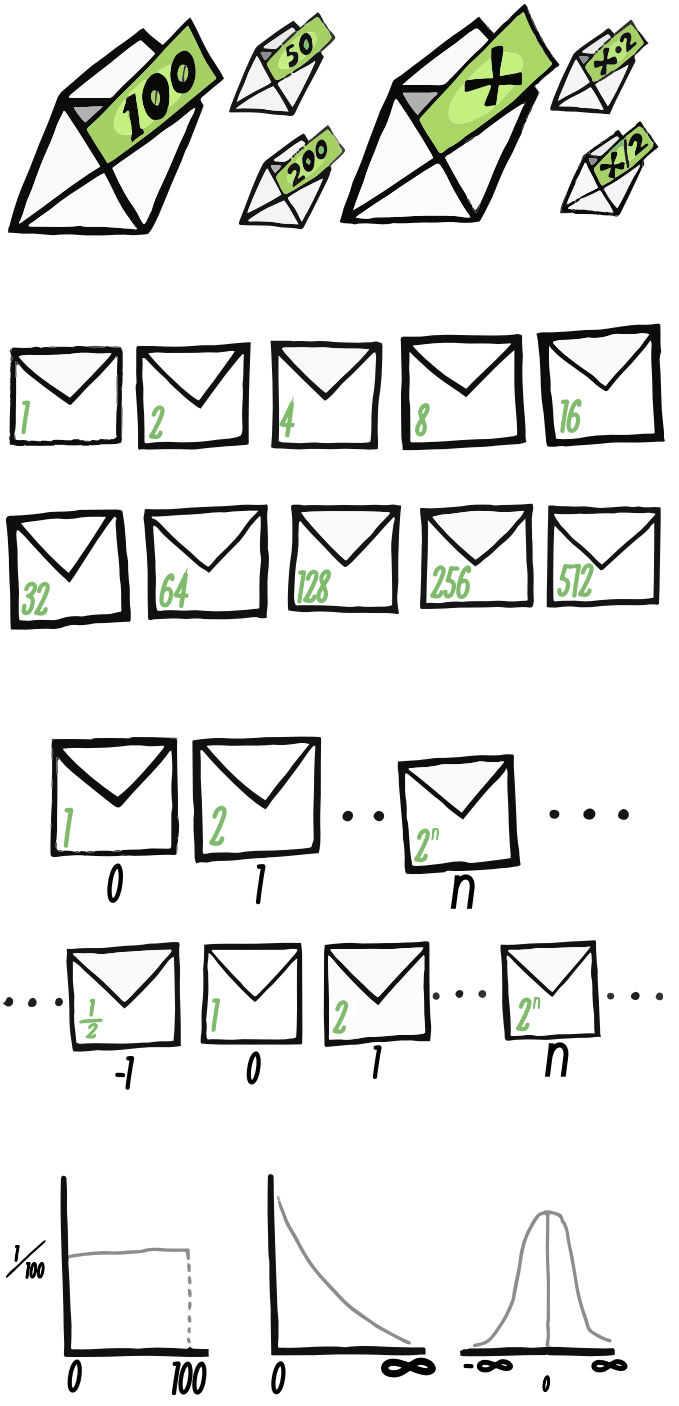
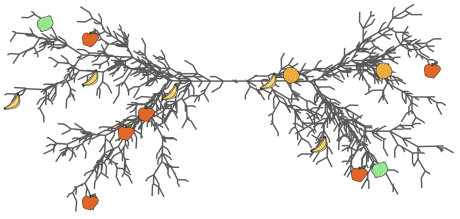
Envelope Trade-Up

Top (Basic Setup): The left scenario shows concrete values: one envelope contains $100, with potential alternatives of $50 or $200. The right scenario generalizes this with variables: envelope X could correspond to either $x/2$ or $2x$ in the alternative envelope. This applies regardless of specific amounts.
Second (Bounded Finite): Eight envelopes containing $1, 2, 4, 8, 16, 32, 64, 512$ represent a bounded geometric sequence. In this finite setup, edge effects matter: the smallest value (1) can only be paired with (2), while the largest (512) can only be paired with (256). These boundary constraints resolve the paradox.
Third (Semi-Infinite): The sequence $1, 2, 4, \ldots, 2^n, \ldots$ extends infinitely in one direction. This introduces asymmetry: there's a smallest possible value but no largest. The lower bound affects the expectation calculation.
Fourth (Bi-Infinite): The sequence $\ldots, 1/2, 1, 2, 4, \ldots, 2^n, \ldots$ extends infinitely in both directions. This symmetric case is where the classical paradox is the most clear, as there are no boundary effects to break the apparent symmetry.
Bottom (Probability Distributions): Three different priors over envelope values: (1) Uniform distribution over $\{1, 2, \ldots, 100\}$ assigns equal $1/100$ probability to each value. (2) Exponentially decaying distribution over positive integers gives decreasing probability as values increase. (3) Gaussian distribution over all real numbers.



Envelope Trade-Up

The Envelope Paradox presents two envelopes where one contains twice the money of the other. After selecting one envelope, seemingly valid probabilistic reasoning suggests an expected gain by switching (averaging x/2 with 2x), regardless of which envelope was initially chosen. This symmetric conclusion creates a logical inconsistency since perpetual switching cannot be optimal. The paradox arises from improper application of expected value calculations to scenarios with unbounded distributions or when conditional probabilities are not properly accounted for. Resolving the paradox requires distinguishing between known values and variables, recognizing when probability distributions are ill-defined, and understanding the limitations of calculations with potentially infinite quantities.



Two Envelopes Paradox ◦ Switching Expectation ◦ Incompatible Baselines ◦ Bounded Distribution Solution ◦ Improper Prior Problem ◦ Bertrand's Chord Paradox ◦ Random Without Context ◦ Boundary Cancellation ◦ Symmetric Model Zero Gain ◦ Kraitchik 1930 ◦ Simulation Verification

"How wonderful that we have met with a paradox; now we have some hope of making progress."

— Niels Bohr, c. 1958

Envelope Trade-Up

The envelope paradox traces its origins to Belgian mathematician Maurice Kraitchik, who proposed a related puzzle in his 1930 book *Les Mathématiques des Jeux*. Kraitchik's version involved two men comparing the value of their neckties, with the winner giving his necktie to the loser as consolation. He also discussed a variant where the men compared the contents of their purses, assuming each contained between one and some large number of pennies with equal probability.

The puzzle gained mathematical attention when John Edensor Littlewood mentioned it in his 1953 book on elementary mathematics, crediting the idea to physicist Erwin Schrödinger. Littlewood's formulation involved cards with numbers written on both sides — a player sees one side of a random card and must decide whether to flip it over. His version made explicit the role of improper prior distributions in generating the paradox.

Martin Gardner brought the puzzle to widespread attention in his 1982 book *Aha! Gotcha*, presenting it as a wallet game between two equally wealthy individuals. Gardner's formulation captured the essential difficulty: each person could construct identical arguments for why the game favored them, yet by symmetry, the game had to be fair. Remarkably, Gardner confessed that while he could analyze the problem correctly, he struggled to pinpoint exactly what was wrong with the switching argument.

The modern envelope paradox resurfaced as probability theory matured into a rigorous mathematical discipline. The twentieth century developments of measure theory, decision theory, and Bayesian analysis provided new tools for understanding why such problems arose and how they might be resolved.

You are presented with two identical envelopes — you are told that one envelope contains twice the money of the other. No other information is given. No other information is given. You select one envelope at random and open it, revealing $100. At this point, you are given the opportunity to switch.

Consider the reasoning that leads to paradox. The envelope you opened contains $100, which must be either the smaller amount or the larger. If $100 is the smaller amount, the other envelope contains $200. If $100 is the larger amount, the other envelope contains $50. Since these two cases seem equally likely, the expected value from switching appears to be:

$$\text{Expected gain} = 0.5 \cdot (+100) + 0.5 \cdot (-50) = +25.$$

This suggests a $25 gain from switching. The same calculation applies regardless of the amount observed, whether $10, $100, or $1000. You should apparently switch every time. But this logic also tells you to switch back again after switching, producing an endless preference loop. The contradiction is that each envelope appears preferable to the other.

Why does this reasoning feel compelling? The calculation follows expected value logic. The probabilities seem reasonable. Without additional information, why shouldn't the observed amount be equally likely to be the smaller or larger? The arithmetic is correct.

©2026 David H. Silver, CC BY-NC 4.0

Yet the conclusion violates intuitions about symmetric problems. The puzzle demands a definitive answer while generating contradictory recommendations.

The error lies in how the observed amount x is interpreted across different terms of the expectation. In one term, x represents the smaller amount; in the other, it represents the larger amount. This reference creates incompatible baselines for comparison.

To make the model coherent, let x denote the smaller of the two amounts. Then the envelopes contain x and $2x$, and each is equally likely to be selected. If you hold x, switching yields $+x$; if you hold $2x$, switching yields $-x$. The outcomes cancel:

$$\text{Average change} = 0.5 \cdot (+x) + 0.5 \cdot (-x) = 0.$$

No advantage arises. The paradox dissolves because the original argument uses expectation without a consistent model.

This becomes clearer in a bounded setup. Suppose the smaller amount is chosen uniformly from $\{2^0, 2^1, \dots, 2^{999}\}$. For most observed amounts A, switching seems to yield a gain of $+0.25A$ since A could be either the smaller or larger envelope value. However, this ignores boundary cases: if $A = 2^0$, switching cannot halve it; if $A = 2^{999}$, switching cannot double it. These rare but extreme boundary effects precisely cancel the average gain across interior values, returning the total expectation to zero.

In the limit as the model becomes unbounded, for example, when x is drawn from $\{\dots, 2^{-2}, 2^{-1}, 2^0, 2^1, \dots\}$, the problem reappears. For any observed amount, switching appears to yield a gain of $+0.25A$. But this assumes all values are equally probable, which is not possible over an infinite set.

A uniform distribution over an infinite number of values cannot exist (Kolmogorov, 1933) — there is no way to assign equal, nonzero probability to infinitely many outcomes and still have the total probability equal to one. Any attempt results in an *improper prior*: a function that resembles a distribution but cannot be normalized.

To reason coherently in such a context, one must use a *probability measure*: a rule that assigns consistent, additive weights to sets of values and sums to one. A measure is *proper* if it satisfies this condition. If it diverges or is undefined, expectations may not exist. Even if each outcome is finite, the global average may be infinite or ill-defined. In that case, expectation ceases to be a meaningful decision tool.

One setup does yield a switching advantage. Fix a value a and flip a coin: if heads, prepare envelopes with a and $2a$; if tails, use a and $a/2$. Hand the envelope containing a to the player. Switching yields either $+a$ or $-0.5a$ with equal probability, giving an expected gain of $+0.25a$. Here, switching is optimal because the model is asymmetric and expectation is applied with explicit conditioning.

The switching advantage depends entirely on the unknown prior distribution over envelope pairs. If smaller amounts are more probable, observing \$100 suggests you likely hold the larger envelope. If larger amounts are more probable, the reverse holds. Without knowing this distribution, rational choice becomes impossible.

Bertrand's paradox (1889) poses the question: what is the probability that a random chord of a circle is longer than the side of an inscribed equilateral triangle? The answer depends

on what "random" means. Select two random points on the circumference and connect them: the probability is $1/3$. Select a random radius and place the chord perpendicular to it at a random distance from the center: the probability is $1/2$. Select a random interior point as the chord's midpoint: the probability is $1/4$. Each method seems reasonable, yet they yield different results. "Random chord" is undefined without specifying the selection procedure.

Similarly, should you treat envelope pairs $(50, 100)$ and $(100, 200)$ as equally likely? Or should you treat the amounts \$50, \$100, \$200 as equally likely? Each choice determines a different prior over the smaller value x. The first approach makes pairs equally likely; the second makes values equally likely, implying smaller values are more probable than larger ones when considering pairs. Without specifying the generation mechanism, "random envelope" has no unique meaning, and the "principle of indifference" — treating outcomes as equally likely (Keynes, 1921) in the absence of information — produces contradictions.

The Bayesian approach requires a prior distribution over envelope pairs. Without knowing the generation mechanism, no prior can be justified. Even given a distribution, improper priors (such as uniform over all positive reals) produce undefined probabilities, and heavy-tailed distributions yield infinite expected values. In symmetric setups, observing one amount provides zero information about which envelope contains more. Over repeated trials with any proper distribution, switching gains nothing on average.

Renormalizing Pascal

Pascal's Wager collapses under the same pathology (Pascal, 1670). The canonical argument claims belief in God offers infinite expected utility: either finite loss (if wrong) or infinite gain (if right). The arithmetic seems decisive. But critics note the state space is unbounded. For any deity G granting infinite utility for act A, one can posit an Anti-God penalizing A with infinite disutility. The expected value becomes:

$$EU = \sum_{i=1}^{\infty}[P(G_i) \cdot \infty] + \sum_{j=1}^{\infty}[P(G_{-j}) \cdot (-\infty)].$$

Undefined. The same divergence that breaks the envelope calculation breaks Pascal.

Physical field theories face identical infinities at high energies. The solution is *renormalization*: impose a cutoff scale (Dyson, 1949) Λ, integrate out fluctuations above this threshold as irrelevant to the macroscopic theory. Virtual particles appear transiently but lack the coupling strength to affect observables.

Apply this to theology. Introduce a *social-credal cutoff* ϵ. A theological hypothesis enters the expected utility calculation only if its social magnitude — measured by historical persistence, institutional architecture, mass adherence — exceeds the threshold. Below ϵ, the probability is set to zero.

Ad hoc philosophical fictions function as virtual particles: transient fluctuations in logical space lacking ontological mass to couple with macro-sociological reality. The

infinite tail truncates. The divergent series collapses into a finite choice between established traditions and the null hypothesis.

This appears to commit the *ad populum* fallacy. It does, but only if logical operators possess Platonic independence. They do not. Axioms such as the Law of Excluded Middle or $1 + 1 = 2$, in this context, are not absolute universals. They are emergent properties of shared human neurobiology, intersubjective agreements sustained by consensus among rational agents.

The validity of *logic* is operational, not objective. It rests on the *vox populi* of the species. If 99% of logicians were rewired to reject the Law of Excluded Middle, it would cease to function as truth within our epistemic system. When we call an argument irrational, we point to inconsistencies with shared basic axioms, not disagreements about logic itself. Rationality is the practice of consistency within a common framework.

We operate several layers of abstraction above this foundation. When debating probability theory or physical laws, we invoke principles two or three levels removed from basic logic, mistaking these derived arguments for appeals to cosmic truth. The system functions only because the base layer — the axioms of logic itself — enjoys universal consensus. If basic logic ever fragmented to 60/40 agreement, rational discourse would collapse entirely. We could not even agree on what "inconsistent" means. Uncomfortable as it feels, one must recognize that at the bottom of all our reasoning — why we believe anything from a news piece to a scientific theory, or even logic itself — lies a social standard. There is no easy way to define truth independent of our shared neurobiology.

If rationality itself is consensus-dependent, religious truth within a rational wager must undergo identical phase transitions. A deity with zero social magnitude is a private delusion. It lacks the intersubjective validation required to exist as a probability vector in a decision matrix.

The "Inverted Gods" objection does not refute Pascal's logic. It demonstrates that his system requires boundary conditions. The ϵ cutoff aligns the Wager with the effective field theory of human consensus. Pascal's emphasis on *praxis* becomes pivotal: the "true" religions are those that operationalized belief into action sufficient to cross ϵ.

In practice, the Wager loses meaning for a simpler reason: humans are not probabilistic decision makers (Kahneman & Tversky, 1979). We cross streets and eat apples despite non-zero probabilities of losing our infinitely-valued lives to accidents or allergies (maybe subconsciously we find terms such that they cancel each other out). We institute cutoffs on rare events and extreme losses not through Bayesian calculation but through neural circuits that force decisions with limited data. The theoretical requirement for a cutoff matches the empirical fact that we already use one. We cannot meaningfully operate on probabilities when assumptions and expected rewards are unbounded. The Wager fails not because the mathematics is wrong, but because the mathematics is not the only guidance to action. For problems at the edge of human experience, intuition and experience are paradoxically more valuable and coherent than analytical slaloms (such as this very piece of commentary!).

Exercises

1. Wine/Water Paradox

A mixture contains wine and water in a ratio $x = \frac{\text{wine}}{\text{water}}$, with $x \in [\frac{1}{3}, 3]$. Assuming x is uniformly distributed, compute the probability that the mixture has $x \leq 2$. Now let $y = \frac{1}{x}$, representing the water-to-wine ratio, and assume y is uniformly distributed instead. What is the probability that $y \geq \frac{1}{2}$? Why are the two results inconsistent?
Hint: Use

$$P(x \leq x_t) = \frac{x_t - x_{\min}}{x_{\max} - x_{\min}}, \qquad P(y \geq y_t) = \frac{x_{\max}(1 - x_{\min} y_t)}{x_{\max} - x_{\min}}.$$

Evaluate both at $x_t = 2$, $y_t = \frac{1}{2}$. What contradiction arises?

2. Monty Hall Problem

Three closed doors hide a car and two goats. You choose Door 1. The host, who knows what is behind each door, opens Door 3, revealing a goat, and offers you the option to switch to Door 2. Should you switch?
Hint: The host never opens your door and never reveals the car. Compute:

$$P(\text{Car behind D1} \mid \text{Host opens D3}) \quad \text{vs} \quad P(\text{Car behind D2} \mid \text{Host opens D3}).$$

Use Bayes's theorem. The host's choice depends on your initial pick and the hidden car's placement.

3. Tuesday Boy Problem

A family has two children. You are told that one of them is a boy born on a Tuesday. What is the probability that both are boys?
Hint: Let $\varepsilon = \frac{1}{7}$. Then:

$$\mathrm{P}(\mathrm{BB} \mid \mathrm{B_T}) = \frac{(1 - (1-\varepsilon)^2) \times \frac{1}{4}}{0 + \frac{1}{4}\varepsilon + \frac{1}{4}\varepsilon + \frac{1}{4}(\varepsilon + \varepsilon - \varepsilon^2)} = \frac{1 - (1-\varepsilon)^2}{4\varepsilon - \varepsilon^2} = \frac{(2-\varepsilon)}{(4-\varepsilon)}$$

4. Two-Black-Sides Card Problem

You randomly select a card from a box containing: one black-black card, one white-white card, and one black-white card. You see a black face. What is the probability the other side is black?

5. Sleeping Beauty Problem

Sleeping Beauty is put to sleep. A fair coin is tossed. If heads, she is woken on Monday only. If tails, she is woken on Monday and Tuesday (but forgets Monday). Upon waking, what is her credence that the coin landed on heads?
Hint: Model the wake-up events: {H-M, T-M, T-T}. Normalize using the condition: "I am awake now."

Resolving the Envelope Paradox

The paradox emerges from treating the observed amount Y as equally likely to be the smaller or larger of two amounts related by a $1:2$ ratio. The naive expectation assumes: $\mathbb{E}[\text{switch}] = 0.5 \cdot 2Y + 0.5 \cdot Y/2 = 5Y/4$, suggesting a gain from switching. But this treats Y as both X and $2X$ in different terms — a misapplication of conditional expectation.

Correct Conditioning

Let X be drawn from prior $f(x)$. The envelope pair $(X, 2X)$ is constructed, and one is presented at random. Let Y denote the observed amount:

$$Y = X \Rightarrow \text{other envelope has } 2Y,$$
$$Y = 2X \Rightarrow \text{other envelope has } Y/2.$$

The conditional probabilities are:

$$\mathbb{P}(Y = y \mid \text{observed smaller}) \propto f(y),$$
$$\mathbb{P}(Y = y \mid \text{observed larger}) \propto \tfrac{1}{2} f(y/2).$$

The expected value of switching given Y is:

$$\mathbb{E}[\text{sw} \mid Y] = \frac{2Y \cdot f(Y) + \frac{1}{2}Y \cdot f(Y/2)}{f(Y) + f(Y/2)}.$$

This depends on the shape of f. When f decays rapidly, $f(Y/2) \gg f(Y)$ for large Y, implying Y is likely the larger value and switching is unfavorable.

Example: Pareto Prior

For $X \sim \text{Pareto}(\alpha)$ with $f(x) = \alpha x^{-(\alpha+1)}$ on $[1, \infty)$:

$$f(Y/2) = \alpha \cdot 2^{\alpha+1} \cdot Y^{-(\alpha+1)} = 2^{\alpha+1} f(Y).$$

Thus:

$$\mathbb{E}[\text{switch} \mid Y] = \frac{2Y + \frac{1}{2}Y \cdot 2^{\alpha+1}}{1 + 2^{\alpha+1}} = Y \cdot \frac{2 + 2^{\alpha}}{1 + 2^{\alpha+1}}.$$

This may exceed or fall below Y depending on α.

Improper Priors

For log-uniform $f(x) \propto 1/x$ on $[1, \infty)$ (improper since $\int_1^\infty \frac{1}{x} dx = \infty$):

$$f(Y) = \frac{1}{Y}, \quad f(Y/2) = \frac{2}{Y},$$
$$\mathbb{E}[\text{switch} \mid Y] = \frac{2Y \cdot \frac{1}{Y} + \frac{1}{2}Y \cdot \frac{2}{Y}}{\frac{1}{Y} + \frac{2}{Y}} = Y.$$

This suggests switching is neutral, but the result is meaningless: the improper prior makes the marginal distribution undefined.

Finite Uniform Model

Let $x \in \{2^0.2^1, \ldots, 2^{N-1}\}$ be uniform. For observed amount $A = 2^m$:

- Interior ($1 \le m \le N-2$): $\mathbb{E}[\Delta \mid A] = \frac{1}{2}(2^m) + \frac{1}{2}(-2^{m-1}) = 2^{m-2}$.
- Boundaries: $\mathbb{E}[\Delta \mid A = 2^0] = +1$, $\mathbb{E}[\Delta \mid A = 2^{N-1}] = -2^{N-2}$.

Global expectation:

$$\mathbb{E}[\Delta] = \frac{1}{N}\left(\sum_{m=1}^{N-2} 2^{m-2} + 1 - 2^{N-2}\right) = \frac{1}{N}\left(-2^{N-3} + \tfrac{1}{2}\right) < 0.$$

As $N \to \infty$, this approaches zero from below, confirming no long-run switching advantage.

References:

Nalebuff, B. (1989). The Other Person's Envelope Is Always Greener. *J. Econ. Persp.*, **3**(1), 171–181.