

Franz Dietrich and [Christian List](#)

The two-envelope paradox: an axiomatic approach

Article (Accepted version)
(Refereed)

Original citation:

Dietrich, Franz and List, Christian (2005) The two-envelope paradox: an axiomatic approach. [Mind](#), 114 (454). pp. 239-248.

DOI: [10.1093/mind/fzi239](https://doi.org/10.1093/mind/fzi239)

© 2005 [Oxford University Press](#)

This version available at: <http://eprints.lse.ac.uk/6050/>

Available in LSE Research Online: June 2008

LSE has developed LSE Research Online so that users may access research output of the School. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LSE Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain. You may freely distribute the URL (<http://eprints.lse.ac.uk>) of the LSE Research Online website.

This document is the author's final manuscript accepted version of the journal article, incorporating any revisions agreed during the peer review process. Some differences between this version and the published version may remain. You are advised to consult the publisher's version if you wish to cite from it.

The Two-Envelope Paradox: An Axiomatic Approach

Franz Dietrich and Christian List¹

May 2004

There has been much recent discussion on the two-envelope paradox. Clark and Shackel (2000) have proposed a solution to the paradox, which has been refuted by Meacham and Weisberg (2003). Surprisingly, however, the literature still contains no axiomatic justification for the claim that one should be indifferent between the two envelopes before opening one of them. According to Meacham and Weisberg, "decision theory does not rank swapping against sticking [before opening any envelope]" (p. 686).

To fill this gap in the literature, we present a simple axiomatic justification for indifference, avoiding any expectation reasoning, which is often considered problematic in infinite cases. Although the two-envelope paradox assumes an expectation-maximizing agent, we show that analogous paradoxes arise for agents using different decision principles such as maximin and maximax, and that our justification for indifference before opening applies here too.

1 The two-envelope paradox

Two indistinguishable envelopes each contain a dollar amount in $\{1, 2, 4, 8, \dots\}$, represented by the random variables X and Y .² One envelope contains twice the amount of the other, but you do not know which one is which. The envelope containing X is presently in your possession. Specifically, we here assume that the joint distribution of X and Y is as follows. For all $k \in \{0, 1, 2, 3, \dots\}$,

$$P(X = 2^k \& Y = 2^{k+1}) = P(X = 2^{k+1} \& Y = 2^k) = \frac{2^{k-1}}{3^{k+1}} \text{ (Broome 1995).}$$

You are now offered a choice between keeping your envelope and receiving the value of X or switching to the other envelope and receiving the value of Y . Which alternative – “keep” or “switch” – is more attractive? The following principle is intuitive:

Indifference before opening. You are indifferent between X and Y , not knowing the value of either.

¹Addresses for correspondence: F. Dietrich, Center for Junior Research Fellows, University of Konstanz, 78457 Konstanz, Germany, franz.dietrich@uni-konstanz.de; C. List, Department of Government, London School of Economics, London WC2A 2AE, U.K., c.list@lse.ac.uk.

²We here assume that you have a preference relation \succsim over the set of discrete real-valued random variables called *lotteries* (such as X and Y). The corresponding strict preference relation \succ and indifference relation \sim are defined in the usual ways, i.e. for all lotteries L_1, L_2 , $L_1 \succ L_2$ if and only if $L_1 \succsim L_2$ and not $L_2 \succsim L_1$; $L_1 \sim L_2$ if and only if $L_1 \succsim L_2$ and $L_2 \succsim L_1$. To represent preferences after opening your envelope, we also assume that, for each information E (a non-empty set of possible worlds), you have a preference relation conditional on E , denoted \succsim_E . The corresponding strict preference relation and indifference relation are defined as above.

However, it can be shown that, conditional on any amount in your envelope, the expected value of the other envelope is higher than the given amount in yours, more precisely:

$$E(Y|X = x) = \begin{cases} \frac{11}{10}x & \text{for } x > 1 \\ 2 & \text{for } x = 1 \end{cases}, \text{ which is greater than } x \text{ for all } x.$$

So, if you open your envelope and look inside it, then, to maximize the expected amount, you should switch to the other envelope:

Switch after opening. You strictly prefer Y to X , knowing the value of X .

But since this reasoning applies for any amount you find, it appears that you should switch even without opening your envelope. This seems absurd, given that the two envelopes are indistinguishable. The two-envelope paradox consists in the apparent tension between indifference before opening and switch after opening. (As noted below, this tension is no logical inconsistency.)

2 Meacham and Weisberg's reply to Clark and Shackel

The expected gain from switching, $E(Y - X)$, is mathematically undefined, because the value of the infinite sum of all probability-weighted values of $Y - X$ depends on the order of summation.³ While Meacham and Weisberg deduce from this that standard decision theory does not rank switching against keeping, Clark and Shackel try to rescue the expectation-based approach by considering "nested expectations" such as $E(E(Y|X) - X)$ or $E(E(Y - X|X + Y))$, each of which corresponds to a different order of summation of the mentioned infinite sum. Clark and Shackel claim to have found the "right" order of summation – that corresponding to $E(E(Y - X|X + Y))$. This "right" order of summation results in a zero infinite sum, which Clark and Shackel take as a justification for indifference before opening. One may question, first, whether "nested expectations" such as $E(E(Y - X|X + Y))$ are relevant for choice: they are *not* expected gains *but* expected-expected gains. Second, assuming that nested expectations *are* relevant, it appears problematic that, among the different possible nested expectations, Clark and Shackel choose $E(E(Y - X|X + Y))$ by appeal to the symmetry of both envelopes; yet, as noted by Meacham and Weisberg, the symmetry of both envelopes was to be shown, not assumed.

3 Can indifference before opening be justified by expectation reasoning?

There are three expectation-based ways in which indifference before opening might be justified. The attempt to show that the expected gain from switching, $E(Y - X)$, is

³For a random variable Z with discrete (possibly infinite) range $\mathcal{Z} \subset \mathbf{R}$ (such as $Z = Y - X$), the expectation $E(Z)$ is defined as the (possibly infinite) sum $E(Z) := \sum_{z \in \mathcal{Z}} zP(Z = z)$ *provided* that this sum is independent of the order of summation, and $E(Z)$ is undefined otherwise. This includes the possibility that $E(Z) = \infty$ or $E(Z) = -\infty$, in which case $E(Z)$ is sometimes called a *quasi*-expectation. For instance, $E(Z)$ is always defined if Z takes only positive or only negative values.

zero fails because $E(Y - X)$ is defined not even in the quasi-sense of footnote 3. Clark and Shackel's attempt to consider nested expectations of the forms $E(E(Y|X) - X)$ or $E(E(Y - X|X + Y))$ fails for the reasons mentioned above. The only potential expectation-based justification for indifference before opening is given by comparing $E(X)$ with $E(Y)$, for, unlike $E(Y - X)$, these two expectations *are* defined (in the quasi-sense of footnote 3): they are both infinite and thus equal. So, as $E(X) = E(Y)$, an agent applying expectation maximization even in infinite cases should be indifferent between X and Y . But is it reasonable to apply expectation maximization in infinite cases? Consider the following example. The random variables X and $X + 1$ have infinite and thus equal expectations, yet the second yields an extra dollar *in every state of the world*. While some might bite the bullet and remain indifferent between X and $X + 1$ following expectation maximization, others would strictly prefer $X + 1$ to X , a preference that is reasonably defensible as explained below. An expectation-based justification of indifference before opening, however, would require abandoning this preference.

4 An Axiomatic Approach

We now give an axiomatic justification for indifference before opening without appeal to any (contestable) reasoning involving infinite expectations. Our justification is more general than an expected-based one, since it applies to agents with any preference relation satisfying some mild conditions, including the preference relations of expectation maximizers, maximinimizers and maximaximizers. We also prove that our justification for indifference before opening is consistent with switch after opening.

Our argument for indifference before opening is that X and Y are two lotteries that have *identical distributions*, i.e. that take each value with the same probability. We argue that the following principle is uncontestable:

Indifference principle. If two lotteries L_1 and L_2 have identical distributions, i.e. $P(L_1 = x) = P(L_2 = x)$ for all x , then you are indifferent between L_1 and L_2 .⁴

Since in standard decision theories the payoff (outcome) is *all* you care about,⁵ a strict preference between two lotteries should only arise if at least one payoff occurs with a different probability under the two lotteries. More specifically, the indifference principle can be reduced to an even more fundamental principle:

Independence of non-payoff characteristics. Your preference relation between two lotteries L_1 and L_2 depends only on the distributions of L_1 and L_2 ; formally, for any two pairs of lotteries L_1, L_1^* and L_2, L_2^* , where L_1 and L_1^* have identical distributions and L_2 and L_2^* have identical distributions, you weakly prefer L_1 to L_2 if and only if you weakly prefer L_1^* to L_2^* .⁶

⁴In all our axioms, if there is additional information, distributions (and hence expectations, minimal values and maximal values) and also preferences are conditional on that information. If general lotteries and not just discrete ones are admitted, then L_1 and L_2 have identical distribution if $P(L_1 \in [a, b]) = P(L_2 \in [a, b])$ for all intervals $[a, b]$.

⁵It is generally taken to be a defining feature of the *payoff* or *outcome* of an act that it includes all relevant consequences of the act (e.g. Joyce 1999).

⁶See footnote 4.

Proposition. If your preference relation is reflexive and transitive, then the indifference principle is equivalent to independence of non-payoff characteristics.

Proof. Let \succsim be reflexive and transitive. First, assume the indifference principle. To show independence of non-payoff characteristics, consider two pairs of lotteries L_1, L_1^* and L_2, L_2^* such that L_1 and L_1^* have identical distributions and L_2 and L_2^* have identical distributions. Then, by the indifference principle, $L_1 \sim L_1^*$ and $L_2 \sim L_2^*$. So, by transitivity, $L_1 \succsim L_2$ if and only if $L_1^* \succsim L_2^*$. Second, assume independence of non-payoff characteristics. To show the indifference principle, consider two lotteries L_1 and L_2 with identical distributions. By independence of non-payoff characteristics, the preference between L_1 and L_2 is the same as that between L_1 and L_1 . By reflexivity, $L_1 \sim L_1$. Hence we have $L_1 \sim L_2$. ■

If we accept the indifference principle as a fundamental decision-theoretic principle, we can use it to justify indifference before opening.

Proposition. The indifference principle implies indifference before opening.

Proof. To apply the indifference principle, we have to show that, for all $k \in \{0, 1, 2, \dots\}$, $P(X = 2^k)$ and $P(Y = 2^k)$ are identical. The latter holds for $k = 0$ by

$$P(X = 1) = P(X = 1 \& Y = 2) = \frac{2^{k-1}}{3^{k+1}}, \quad P(Y = 1) = P(X = 2 \& Y = 1) = \frac{2^{k-1}}{3^{k+1}},$$

and for $k \in \{1, 2, \dots\}$ by

$$\begin{aligned} P(X = 2^k) &= P([X = 2^k \& Y = 2^{k+1}] \vee [X = 2^k \& Y = 2^{k-1}]) \\ &= P(X = 2^k \& Y = 2^{k+1}) + P(X = 2^k \& Y = 2^{k-1}) = \frac{2^{k-1}}{3^{k+1}} + \frac{2^{k-2}}{3^k}, \\ P(Y = 2^k) &= P([X = 2^{k+1} \& Y = 2^k] \vee [X = 2^{k-1} \& Y = 2^k]) \\ &= P(X = 2^{k+1} \& Y = 2^k) + P(X = 2^{k-1} \& Y = 2^k) = \frac{2^{k-1}}{3^{k+1}} + \frac{2^{k-2}}{3^k}. \quad \blacksquare \end{aligned}$$

Switch after opening follows immediately from the following standard principle.

Finite expectation maximization. For any two lotteries L_1 and L_2 with finite expected values, if the expected value of L_1 is greater than or equal to that of L_2 , then you weakly prefer L_1 to L_2 .⁷

While the indifference principle is consistent with any risk attitude, only risk-neutral preference relations over lotteries satisfy finite expectation maximization. Below we show that switch after opening (and hence a two-envelope paradox) can also be derived without assuming risk-neutral preferences, for instance for preferences based on maximin or maximax.

Note that finite expectation maximization does not imply the indifference principle. If two lotteries have identical distributions and infinite expected values, such

⁷See footnote 4.

as the two envelopes before opening, then finite expectation maximization is silent on how to rank them, and we require the indifference principle to deduce indifference.

But if the indifference principle and finite expectation maximization together imply indifference before opening and switch after opening, two apparently inconsistent conditions, does this not suggest that we must reject either the indifference principle or finite expectation maximization? We do not think so. The apparent inconsistency between indifference before opening and switch after opening relies on the tacit use of an additional principle, which we reject (see also Chalmers 2002).

Event-wise dominance principle. Let \mathcal{P} be a partition of the set of all possible states of the world into non-empty events. For any two lotteries L_1 and L_2 , if you strictly prefer L_1 to L_2 conditional on observing event E for every E in \mathcal{P} , then you strictly prefer L_1 to L_2 unconditionally.

Switch after opening together with the event-wise dominance principle contradicts indifference before opening. In the two-envelope case, the set of all possible states of the world can be represented as the set of all pairs of values x and y of the lotteries X and Y . Partition this set into the event that $X = 1$ (containing the state $\langle 1, 2 \rangle$), the event that $X = 2$ (containing the states $\langle 2, 1 \rangle$ and $\langle 2, 4 \rangle$), the event that $X = 4$ (containing the states $\langle 4, 2 \rangle$ and $\langle 4, 8 \rangle$), and so on. Each event in the partition corresponds to precisely one observed value of X . Now switch after opening implies that, for every such event, you strictly prefer lottery Y to lottery X conditional on observing that event. But this is exactly the antecedent condition of the event-wise dominance principle. So the principle applies, and you should strictly prefer Y to X unconditionally too. But this contradicts indifference before opening. Furthermore, under finite expectation maximization, the event-wise dominance principle is even internally inconsistent.⁸

On the other hand, without the event-wise dominance principle there is no logical contradiction between switch after opening and indifference before opening.

Is the event-wise dominance principle compelling? Surely, if lottery L_1 were preferable to lottery L_2 *conditional on every possible state of the world*, then L_1 would also seem preferable to L_2 unconditionally. This may be the reason for the intuitive appeal of a dominance principle. But, crucially, saying that L_1 is preferable to L_2 conditional on every event in some partition, as in event-wise dominance, is *not* the same as saying that L_1 is preferable to L_2 conditional on every possible state of the world. Some events in the partition may contain more than one state of the world. For instance, the event $X = 2$ contains the states $\langle 2, 1 \rangle$ and $\langle 2, 4 \rangle$. A partition into events is therefore typically less fine-grained than a partition into states of the world. One may prefer L_1 to L_2 conditional on some event, and yet disprefer L_1 to L_2 conditional on some state of the world within that event. This is precisely what happens in the two-envelope example. You would prefer the other envelope conditional on the event

⁸ Assume finite expectation maximisation and the event-wise dominance principle. Consider first the partition into the events $\{X = x\}$, $x = 1, 2, 4, \dots$. As $E(Y|X = x) > x$ for all x , by finite expectation maximisation you prefer Y to X conditional on each event $\{X = x\}$, $x = 1, 2, 4, \dots$. Hence, by the event-wise dominance principle, you prefer Y over X unconditionally, i.e. $Y \succ X$ and not $X \succ Y$. By exactly the same argument for the partition $\{Y = y\}$, $y = 1, 2, 4, \dots$, you prefer X over Y unconditionally, i.e. $X \succ Y$ and not $Y \succ X$, a contradiction. Similar arguments can be given for maximin or maximax preferences (see below).

that you observe \$2 in yours, but you would disprefer it conditional on one of the two states of the world within that event: the state in which the other envelope contains \$1. As the event-wise dominance principle does not require fine-grained partitions – which are essential to dominance reasoning – the principle does not satisfactorily capture the essence of dominance reasoning.

It is now clear that we reject the event-wise dominance principle. But while we reject this coarse-grained dominance principle, the following fine-grained one may be acceptable:

State-wise dominance principle. For any two lotteries L_1 and L_2 , if you strictly prefer L_1 to L_2 conditional on every possible state of the world, then you strictly prefer L_1 to L_2 unconditionally.

Accepting state-wise dominance has the following advantage. The inconsistency between indifference before opening and switch after opening is avoided, and yet some intuitively plausible instances of dominance reasoning can be accommodated, such as the preference of the lottery $X + 1$ over the lottery X , where you are *guaranteed* an additional dollar in the first lottery in every possible state of the world.

5 A two-envelope paradox with maximin or maximax

Now assume one of the following two decision principles instead of finite expectation maximization:

Maximin. For any two lotteries L_1 and L_2 , if the minimal value of L_1 is greater than or equal to that of L_2 , then you weakly prefer L_1 to L_2 .⁹

Maximax. For any two lotteries L_1 and L_2 , if the maximal value of L_1 is greater than or equal to that of L_2 , then you weakly prefer L_1 to L_2 .¹⁰

As mentioned above, indifference before opening can still be justified by the indifference principle. After opening, a maximaximizer will always switch, so that the two-envelope paradox remains for maximax preferences. A maximinimizer, however, will keep after opening *unless* the observed amount is $X = 1$ (in which case Y must be 2).

To obtain a two-envelope paradox also for maximinimizers, we drop the assumption that the amounts X and Y are powers of 2, and assume instead that the pair $\langle X, Y \rangle$ has as its range the set of all pairs $\langle x, y \rangle$ of rational numbers x and y such that $0 < x < 1$ and $0 < y < 1$.¹¹ Assume further X and Y have the

⁹See footnote 4. Also, one might want to restrict the quantification to only lotteries L_1 and L_2 that are bounded below, i.e. take values larger than some constant k . Otherwise, maximin requires indifference between L and $L + 1$ whenever the lottery L is unbounded below, violating state-wise dominance.

¹⁰See footnote 4. Also, for reasons analogous to those in footnote 9, one might want to restrict the quantification to only lotteries L_1 and L_2 that are bounded *above*. If there is additional information, maximal values are conditional on that information.

¹¹The assumption of rational numbers may be dropped if X and Y are allowed to take uncountably many possible values.

same (marginal) distribution. Then both maximinimizers and maximaximizers run into a two-envelope paradox: they are indifferent before opening (by applying the indifference principle), but after opening a maximaximizer always switches and a maximinimizer always keeps.

This result is interesting in several respects. First, we have identified a two-envelope paradox that involves no expectation reasoning whatsoever, neither before nor after opening. Second, an often criticized feature of the standard two-envelope paradox has been removed, namely the unboundedness of the possible payoffs (or the related unboundedness of the agent's utility function): both amounts range between 0 and 1 and of course have finite expectations. Third, we have removed the tight link between the two amounts X and Y by allowing for a wide range of joint distributions of X and Y , for instance for independence of X and Y (for an expectation-based two-envelope paradox without a dependence, see Chalmers 2002).

6 Conclusion

Standard expectation-based decision theory is unable to rank keeping and switching before opening any envelope. It has to be complemented by an additional principle, the indifference principle, which allows to deal with certain cases of infinite expectations. The indifference principle seems uncontestable, as it can be formally justified from the even more fundamental principle of independence of non-payoff characteristics, which lies at the heart of any outcome-oriented decision theory. The apparent tension between indifference before opening and switch after opening relies on the tacit use of the event-wise dominance principle; we have rejected this principle, while suggesting that one may plausibly keep the weaker state-wise dominance principle. As illustrated by our new two-envelope paradoxes for maximin and maximax preferences, expectation reasoning, unbounded payoffs or a dependence between the two amounts are not essential for the occurrence of the paradox.

7 References

- Broome, J. 1995: "The Two-Envelope Paradox". *Analysis* 55, pp. 6-11.
- Chalmers, D. 2002: "The St. Petersburg Two-Envelope Paradox". *Analysis* 62, pp. 155-57.
- Clark, M. and Shackel, N. 2000: "The two-envelope paradox". *Mind* 109(435), pp. 415-442.
- Joyce, J. M. 1999: *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press.
- Meacham, C. J. G. and Weisberg, J. 2003: "Clark and Shackel on the Two-Envelope Paradox". *Mind* 112(448), pp. 685-689.