Newcomb, Frustrated

Rhys Borchert
University of Arizona

Jack Spencer
Massachusetts Institute of Technology

Abstract

This paper develops a hybridization of Newcomb's Problem and the Frustrater (Spencer and Wells 2019), underscoring how difficult it is to reconcile the rationality of taking both boxes in Newcomb's Problem and the rationality of taking the envelope in the Frustrater.

1.

Some decision theorists have a pair of intuitions. They think that two-boxing in Newcomb's Problem is rational, and they think that taking the envelope is rational in the following:

> <u>Frustrater</u>
>
> There are two boxes, A and B, and an envelope. The agent has three options — they can take any of the three. The envelope contains $40. The boxes together contain $100. The distribution is frustrating: a reliable predictor made a prediction about what the agent would choose, placing $100 in A/B if they predicted that the agent would take B/A, and placing $50 in each box if they predicted that the agent would take the envelope. The agent knows all of this.[1]

[1] This decision problem from Spencer and Wells (2019) is a variant of Ahmed's (2014) Dicing with Death.

The two intuitions are hard to reconcile. The two leading approaches to decision theory are Evidential Decision Theory (EDT) and Causal Decision Theory (CDT), and neither vindicates both. Suppose that we have a set of possible worlds, $W = \{w_1, \ldots, w_n\}$, a set of dependency hypotheses, $K = \{k_1, \ldots, k_m\}$, a set of options, $A = \{a_1, \ldots, a_k\}$, and suppose that C and u are the agent's credence and utility functions, respectively. EDT prescribes maximizing V, where $V(a) = \Sigma_W C(w|a)u(w) = \Sigma_K C(k|a)V(ak)$. CDT prescribes maximizing U, where $U(a) = \Sigma_K C(k)V(ak)$. But two-boxing in Newcomb's Problem does not maximize V, and taking the envelope in the Frustrater does not maximize U.[2]

The question arises whether the intuitions can be reconciled, and this paper gives new reasons for thinking not. It formulates a new decision problem, a hybridization of Newcomb's Problem and the Frustrater, which, in two ways, underscores how hard reconciling the intuitions is.

First, the hybridized decision problem challenges a weakening of the common ground between EDT and CDT. Two-boxing in Newcomb maximizes U; taking the envelope in the Frustrater maximizes V. Someone who wants to reconcile the intuitions thus might hope, conservatively, that the following, though not a complete decision theory, is anyway true as far as it goes:

Disjunction. An option is rational only if it maximizes either V or U.

Or, even more conservatively, might hope that the following is true as far as it goes:

[2] In the Frustrater, $U(a_E) = 40$, and $U(a_A) + U(a_B) = 100$.

Weighted Average. An option is rational only if it maximizes some weighted average of V and U (i.e. $\alpha V + (1 - \alpha)U$, for some α, $0 \leq \alpha \leq 1$).[3]

But Weighted Average fails, we will argue, if both intuitions are correct.

Second, the hybridized decision problem challenges prominent attempted reconciliations. Gallow (2020) and Spencer (2021b), partly in hopes of vindicating both intuitions, have formulated newfangled approaches to decision theory. Both decision theories vindicate both intuitions, and neither entails Weighted Average. But if the hybridized decision problem is a counterexample to Weighted Average, then it is also a counterexample to Gallow and Spencer.

2.

We build up to the hybridized decision problem in steps. The first step instantiates Newcomb's Problem:

<u>Newcomb Envelope</u>

There is an envelope that contains at least \$40. The agent has two options — they can take the envelope, or pay \$20 for the envelope. A reliable predictor, pre-rewarding paying, made a prediction about what the agent would choose, placing an additional \$40 inside if they predicted that the agent would pay for the envelope, and doing

[3] For sympathetic discussions of Weighted Average, see e.g. Nozick (1993) and MacAskill (2016).

nothing if they predicted that the agent would take the envelope. The agent knows all of this.

The second step modifies the Frustrater by doubling the sum of money distributed between the boxes.

Doubled Frustrater

There are two boxes, A and B, and an envelope. The agent has three options — they can take any of the three. The envelope contains $40. The boxes together contain $200. The distribution is frustrating: a reliable predictor made a prediction about what the agent would choose, placing $200 in A/B if they predicted that the agent would take B/A, and placing $100 in each box if they predicted that the agent would take the envelope. The agent knows all of this.

The third step merges Newcomb Envelope and Doubled Frustrater, producing the hybridization.

Newcomb Frustrater

There are two boxes, A and B, and an envelope. The agent has four options — they can take A, B, or the envelope, or they can pay $20 for the envelope. Two reliable predictors made predictions. The first made a prediction about whether the agent would end up with A, B, or the envelope, placing $200 in A/B if they predicted that the agent would take B/A, and placing $100 in each box if they predicted that the agent would take the envelope or pay for the envelope. The second, pre-rewarding paying, made a prediction about whether the agent would pay, placing an additional $40 inside the envelope and inside each box if they predicted that the agent would pay for the envelope, and doing nothing if they predicted that the agent would take a box or take the envelope. The agent knows all of this.

As we verify below, we have a counterexample to Weighted Average and the decision theories developed by Gallow and Spencer if taking the envelope is rational (i.e. among the rational options) in Newcomb Frustrater.[4]

3.

Start with Weighted Average. Let k_X be, for any option a_X, the dependency hypothesis that holds if the conjunction of the predictors' predictions entails a_X, and suppose, for simplicity, that the predictors are perfectly reliable. Then:

$V(a_A) = V(a_B) = 0 < V(a_E) = 40 < V(a_\$) = 60$;

$U(a_A) = C(k_A)(0) + C(k_B)(200) + C(k_E)(100) + C(k_\$)(140)$;

$U(a_B) = C(k_A)(200) + C(k_B)(0) + C(k_E)(100) + C(k_\$)(140)$;

$U(a_E) = C(k_A)(40) + C(k_B)(40) + C(k_E)(40) + C(k_\$)(80)$; and

$U(a_\$) = C(k_A)(20) + C(k_B)(20) + C(k_E)(20) + C(k_\$)(60) = U(a_E) - 20$.

Taking the envelope in Newcomb Frustrater therefore does not maximize any weighted average of V and U. For if $\alpha > 0.5$, then $\alpha V(a_E) + (1 - \alpha)U(a_E) < \alpha V(a_\$) + (1 - \alpha)U(a_\$)$. And since $U(a_A) +$

[4] Something akin to Newcomb Frustrater can be arrived at by merging the Frustrater and XOR Blackmail, a different alleged counterexample to EDT; cf. Levinstein and Soares (2020).

$U(a_B) = C(k_A \vee k_B \vee k_E)(200) + 2C(k_\$)(140)$, either $U(a_A) \geq C(k_A \vee k_B \vee k_E)(100) + C(k_\$)(140) = U(a_E) + 60$, or $U(a_B) \geq C(k_A \vee k_B \vee k_E)(100) + C(k_\$)(140) = U(a_E) + 60$. So if $\alpha \leq 0.5$, either $\alpha V(a_E) + (1 - \alpha)U(a_E) < \alpha V(a_A) + (1 - \alpha)U(a_A)$, or $\alpha V(a_E) + (1 - \alpha)U(a_E) < \alpha V(a_B) + (1 - \alpha)U(a_B)$.

Next, turn to Gallow. Gallow defends a news management theory, which begins with a special pairwise comparison, $I(a_i, a_\mathrm{j}) = \Sigma_K C(k|a_i)(V(a_i\&k) - V(a_j\&k))$: the news value of a_i compared to a_j, from the perspective of the plan to take a_i.[5] It then defines a special sort of comparative news value, $N(a_i, a_j) = I(a_i, a_j) - I(a_j, a_i)$, which it uses to rank options. In Newcomb Frustrater, if we again suppose, for simplicity, that the predictors are perfectly reliable, then:

$N(a_E,a_\$) = I(a_E,a_\$) – I(a_\$,a_E) = 20 + 20 = 40$;

$N(a_B,a_\$) = N(a_A,a_\$) = I(a_A,a_\$) – I(a_\$,a_A) = –20 + 80 = 60$;

$N(a_B,a_E) = N(a_A,a_E) = I(a_A,a_E) – I(a_E,a_A) = –40 + 60 = 20$; and

$N(a_A,a_B) = 0 – 0 = 0$.

The box-taking options, a_A and a_B, are the Condorcet winners, tying each other and defeating every other option, so Gallow predicts that a_A and a_B are the only rational options.

Finally, turn to Spencer. Spencer's pluralist theory of rational choice starts with a set of partitions, $\{K^\mathbf{0}, \ldots , K^\top\}$, linearly ordered by granularity. The finest, $K^0 = \{k_1^0, \ldots, k_n^0\}$, is the set of dependency hypotheses, each a proposition that conjoins a hypothesis about the laws of nature with a hypothesis about the history of the world up to the moment of choice. Coarser partitions are arrived at by successively removing slices of history, thus producing shorter initial segments, and then finally removing the laws of nature themselves. The coarsest partition, $K^\top$, is the trivial

[5] Barnett (2022), who defends a theory similar to Gallow's, calls this quantity, 'graded ratifiability'.

partition which has ⊤ as its only element. The set of partitions induces a set of quantities, $\{U^{\mathbf{0}}, \ldots, U^{\top}\}$, where $U^{j}(a) = \Sigma_{K^{\wedge}j}C(k^{j})V(a\&k^{j})$. The set of quantities inherits the linear order of the set partitions. If a is a member of the agent's set of options, then a is said to *stably* maximize U^{j} just if, for any $a_i \in A$, $\Sigma_{K^{\wedge}j}C(k^{j})V(a\&k^{j}) \geq \Sigma_{K^{\wedge}j}C(k^{j})V(a_i\&k^{j})$ and $\Sigma_{K^{\wedge}j}C(k^{j}|a)V(a\&k^{j}) \geq \Sigma_{K^{\wedge}j}C(k^{j}|a)V(a_i\&k^{j})$. An element of $\{U^{\mathbf{0}}, \ldots, U^{\top}\}$ is said to be *stably maximized* just if some option stably maximizes it, and, according to Spencer, an option is rational just if it stably maximizes the finest element of $\{U^{\mathbf{0}}, \ldots, U^{\top}\}$ which is stably maximized.

In Newcomb Frustrater, the finest elements of $\{U^{\mathbf{0}}, \ldots, U^{\top}\}$ are maximized by the box-taking options, a_A and a_B. But none of the elements of $\{U^{\mathbf{0}}, \ldots, U^{\top}\}$ which are maximized by a box-taking option are stably maximized, and all of the elements of $\{U^{\mathbf{0}}, \ldots, U^{\top}\}$ which are not maximized by a box-taking option are maximized (and indeed stably maximized) by paying \$20 for the envelope, $a_{\$}$. Spencer thus predicts that $a_{\$}$ is the only rational option.

4.

The conditional claim — that taking the envelope is rational in Newcomb Frustrater, if two-boxing is rational in Newcomb's Problem and taking the envelope is rational in the Frustrater — enjoys plausibility, we think, even before any argument in its favor has been broached. But we can motivate it, and deepen our engagement with Gallow and Spencer, by considering the following three claims which together entail it:

(1) If two-boxing is rational in Newcomb's Problem, then taking the envelope is rational in Newcomb Envelope.

(2) If taking the envelope is rational in the Frustrater, then taking the envelope is rational in the Doubled Frustrater.

(3) If taking the envelope is rational both in Newcomb Envelope and in the Doubled Frustrater, then taking the envelope is rational in Newcomb Frustrater.

The first claim, (1), is uncontroversial. Gallow's decision theory falsifies (2). In the Frustrater, $N(a_A,a_E) = I(a_A,a_E) - I(a_E,a_A) = -40 + 10 = -30$: the agent is certain that a given box contains \$40 less than the envelope, conditional on taking that box, and certain that the envelope contains \$10 less than either box does, conditional on taking the Envelope. In the Doubled Frustrater, $N(a_E,a_A) = I(a_E,a_A) - I(a_A,a_E) = -40 + 60 = 20$: the agent, again certain that a given box contains \$40 less than the envelope, conditional on taking that box, is certain that the envelope contains \$60 less than either box does, conditional on taking the envelope. Gallow's decision theory thus recommends the envelope in the Frustrater, but recommends the boxes in Doubled Frustrater.

The falsification of (2) tells against Gallow. If taking the envelope is rational in the Frustrater, it is so, it seems, because the boxes are frustrated, and the boxes in the Doubled Frustrater are no less frustrated. Doubling the sum of money distributed between the boxes does not weaken the pro-envelope intuitions.

In fact, using Gallow's decision-theoretic technology helps brings out how odd the recommendation to take a box in the Doubled Frustrater is. Consider:

Roomed Doubled Frustrater

The agent has two options — they can take the envelope containing \$40 (a_1), or face the Doubled Frustrater (a_2).[6]

The agent has no uncertainty about what the envelope contains and knows that they can take the envelope if they face the Doubled Frustrater, so facing the Doubled Frustrater should not be bad news compared to taking the envelope. But $I(a_2,a_1)$ — the news value of facing the Doubled Frustrater compared to taking the envelope, from the perspective of the plan to take the envelope — is negative if the agent gives positive credence to taking a box upon facing the Doubled Frustrater. (And if some plausible assumptions hold, $I(a_2, a_1) < I(a_1, a_2)$.)[7]

Spencer's decision theory verifies (2), but falsifies (3). Like EDT, it recommends taking the envelope in the Doubled Frustrater and paying \$20 for the envelope in Newcomb Frustrater. The falsification of (3) tells against Spencer.

Spencer's decision theory recommends taking the envelope in some variants of Newcomb Frustrater, including the following:

---

[6] This argument draws on Spencer (2021a). Also see Rothfus (2022).

[7] Suppose that $C(k_A \vee k_B|a_2) > 0$. $I(a_2, a_1) = \Sigma_K C(k|a_2)(V(a_2k) - V(a_1k)) = \Sigma_K C(k|a_2)(V(a_2k) - 40)$. If the agent takes the predictor to be perfectly reliable, conditional on the conjunction of any dependency and a_2, as well they might, then $C(a_A|a_2k_A) = C(a_B|a_2k_B) = C(a_E|a_2k_E) = 1$, in which case: $V(a_2k_A) = C(a_A|a_2k_A)V(a_Aa_2k_A) + C(a_B|a_2k_A)V(a_Ba_2k_A) + C(a_E|a_2k_A)V(a_Ea_2k_A) = (1)(0) + (0)(200) + (0)(40) = 0$; $V(a_2k_B) = C(a_A|a_2k_B)V(a_Aa_2k_B) + C(a_B|a_2k_B)V(a_Ba_2k_B) + C(a_E|a_2k_B)V(a_Ea_2k_B) = (0)(200) + (1)(0) + (0)(40) = 0$; $V(a_2k_E) = C(a_A|a_2k_E)V(a_Aa_2k_E) + C(a_B|a_2k_E)V(a_Ba_2k_E) + C(a_E|a_2k_E)V(a_Ea_2k_E) = (0)(100) + (0)(100) +(1)(40) = 40$; $I(a_2,a_1) = C(k_A \vee k_B|a_2)(-40) < 0 < C(k_A \vee k_B|a_2)(40) = \Sigma_K C(k|a_2)(40 - V(a_2k)) = I(a_1, a_2)$. But the agent might have different conditional credences. If the agent knows for certain that they will take a box upon facing the Doubled Frustrater, for example, then $C(a_E|a_2k_E) = 0$. If $C(a_E|a_2k_E) = 0$, then, plausibly, $C(k_E|a_2) = 0$, in which case, still, $I(a_2, a_1) = \Sigma_K C(k|a_2)(V(a_2k) - V(a_1k)) = C(k_A \vee k_B|a_2)(-40) < 0$. But if $C(a_E|a_2k_E) = 0$, then $V(a_2k_E) = 100$, in which case: $I(a_1,a_2) = C(k_A|a_1)(40 - 0) + C(k_B|a_1)(40 - 0) + C(k_E|a_1)(40 - 100) = -60 < I(a_2, a_1)$.

Variant Newcomb Frustrater

> There are two boxes, A and B, and an envelope. The agent has four options — they can take A, B, or the envelope, or they can pay $20 for the envelope. Two reliable predictors made a prediction. The first placed $80 in the envelope if they predicted that the agent would pay for the box, and placed $40 in the envelope if they predicted that the agent would take a box or take the envelope without paying. The second placed $200 in A/B if they predicted that the agent would take B/A, placed $100 in each box if they predicted that the agent would take the envelope without paying, and placed $140 in each box if they predicted that the agent would pay for the envelope. The agent knows all of this.

The *U*-values and *V*-values of options are the same in Variant Newcomb Frustrater and Newcomb Frustrater, so CDT and EDT treats them equivalently. But Spencer does not. In Variant Newcomb Frustrater, as in Newcomb Frustrater, the finest elements of $\{U^{\mathbf{0}}, \ldots, U^{\top}\}$ are maximized by the box-taking options, and none of the elements of $\{U^{\mathbf{0}}, \ldots, U^{\top}\}$ which are maximized by a box-taking option are stably maximized. But in Variant Newcomb Frustrater, the finest stably maximized element of $\{U^{\mathbf{0}}, \ldots, U^{\top}\}$, a quantity that corresponds to a partition that specifies the first but not the second of the two predictions, thus specifying the contents of the envelope but not the content of the boxes, is maximized (and indeed stably maximized) by taking the envelope and not paying.

In Newcomb Frustrater, however, every partition which specifies anything about the content of the envelope specifies enough about the contents of the boxes to ensure that no option stably maximizes the corresponding quantity, so the only elements of $\{U^{\mathbf{0}}, \ldots, U^{\top}\}$ which are stably maximized are maximized by the V-maximizing option, $a_{\$}$.

5.

Reconciling the two intuitions — formulating a decision theory that predicts both that two-boxing is rational in Newcomb's Problem and that taking the envelope is rational in the Frustrater — is harder than has been appreciated, if taking the envelope is rational in Newcomb Frustrater. And there is reason to think that taking the envelope is rational in Newcomb Frustrater, if both intuitions are correct.[8]

References

Ahmed, A. 2014. Dicing with death. *Analysis*, 74(4): 587–592.

Barnett, D.J. 2022. Graded ratifiability. *Journal of Philosophy* 119(2): 57–88.

Gallow, J.D. 2020. The causal decision theorist's guide to managing the news. *Journal of Philosophy 117*(3): 117–149.

Rothfus, G. J. 2022. A plan-based causal decision theory. *Analysis 82*(2): 264–272.

Levinstein, B.A. and N. Soares. 2020. Cheating death in Damascus. *Journal of Philosophy* 117(5): 237–66

MacAskill, W. 2016. Smokers, psychos, and decision-theoretic uncertainty. *Journal of Philosophy 113*(9): 425–445.

Nozick, R. 1993. *The Nature of Rationality*. Princeton: Princeton University PRess.

[8] Thanks to an extraordinarily helpful anonymous reviewer, whose feedback led to a much improved paper.

Spencer, J. and I. Wells. 2019. Why take both boxes? *Philosophy and Phenomenological Research 99*(1): 27-48.

Spencer, J. 2021a. An argument against causal decision theory. *Analysis 81*(1): 52–61.

Spencer, J. 2021b. Rational monism and rational pluralism. *Philosophical Studies 178*: 1769–1800.