Money pumps for agents who are ambiguity-averse or ambiguity-seeking

Adam Elga and Johan E. Gustafsson

January 9, 2026*

Abstract

We argue that under general conditions, agents who exhibit Ellsberg-style ambiguity-aversion (or ambiguity-seekingness) are subject to money pumps: they will give up, for no gain, what they could have kept for free. The argument differs from existing results in three ways. First, it gains generality by depending on only minimal assumptions about the target agent. Second, unlike standard approaches that appeal to backward induction, the argument does not assume that rational agents expect that they will choose rationally after having made an irrational choice. Third, the argument is particularly simple.

Keywords: Ambiguity aversion, Ellsberg paradox, backward induction, sequential exploitability, money pump.

1 *Introduction*

Not all evidence is created equal.

Sometimes evidence is *unambiguous*. For example: you are uncertain whether a particular radioactive sample will decay in the next hour, and are fully informed about the chance that it will. Or: you are uncertain whether a frequently-performed medical intervention will be successful, and have decisive and plentiful statistics about its effectiveness in an appropriate class of similar circumstances. Or: you are uncertain about the outcome of

*Forthcoming in *Journal of Philosophy*. Thanks to Nabil Al-Najjar, Charlotte Elga, David Elga, Zachary Goodsell, Jacob Nebel, Agustín Rayo, Gideon Rosen, Katie Steele, an audience at the University of Texas at Austin, and two anonymous *Journal of Philosophy* referees. For research environment support, A. Elga thanks the Carpinteria Public Library.

a coin toss, and symmetries in the toss setup clearly justify having equal confidence in heads as tails.

Other times evidence is *ambiguous*. For example: you are uncertain how an esoteric physical process will turn out, but have no idea about the chances that govern it. Or: you are uncertain whether a never-before-attempted medical intervention will be successful, but have no relevant statistics. Or: you are uncertain whether there is extraterrestrial life, but it is unclear how much relative weight to assign to various aspects of your evidence.[1]

Some of us are *ambiguity-sensitive*: we care about the difference between ambiguous and unambiguous evidence. For example, we may be *ambiguity-averse*: all else equal, we would prefer to avoid having our lives and livelihoods depend on matters about which our evidence is ambiguous.

Imagine having to choose a medical procedure. There is an "unambiguous" procedure as well as two "ambiguous" ones.[2] Each procedure is certain to either produce immediate death or a complete cure. The unambiguous procedure has been studied in depth: its chance of success for you is exactly 50%, and as a result you are 50% confident that it would succeed. The ambiguous procedures are novel: no one has any idea how likely each is to succeed. But you are informed of this: exactly one of the novel procedures would succeed for you.

Which procedure do you prefer? If you prefer the unambiguous procedure to each of the ambiguous ones, you are ambiguity-averse (and thereby ambiguity-sensitive).

A less dramatic example concerns two urns full of red and black balls.[3] You are informed that a ball will be randomly drawn from each urn, and

[1] On the contrast between unambiguous and ambiguous evidence, see for example Ellsberg 1961, Joyce 2005, Kaplan 1996, Keynes 1921, Levi 1974, 1980, Walley 1991. (As noted in Machina and Siniscalchi 2014, p. 732, it is shown in LeRoy and Singell 1987 that the above distinction does not quite correspond to the distinction Knight (1921) famously drew between "risk" and "uncertainty", since Knight claimed that agents have what amounts to precise subjective probabilities even in cases that he counted as involving "uncertainty".)

[2] Compare to Safra and Segal 2022, p. 1.

[3] Ellsberg 1961, pp. 650, 653. As noted in Machina and Siniscalchi 2014, p. 731n6, this example was independently discovered by Fellner (1961, p. 673).

that the proportion of red balls in the "unambiguous" urn equals one half. As a result, you are 50% confident that a red ball will be drawn from that urn. You are told nothing about the proportion of red balls in the other — "ambiguous" — urn. Now you must choose between the following: (1) a ticket worth $100 if a red ball is drawn from the unambiguous urn, (2) a ticket worth $100 if a red ball is drawn from the ambiguous urn, and (3) a ticket worth $100 if a black ball is drawn from the ambiguous urn.

Which ticket do you prefer? If you prefer the unambiguous ticket to each of the ambiguous ones, you are ambiguity-averse (and thereby ambiguity-sensitive).

2

Is being ambiguity-sensitive rationally permissible?

We will argue that under general conditions, ambiguity-sensitive agents are subject to money pumps: they will give up, for no gain, what they could have kept for free. So, if it is irrational to be vulnerable to a money pump, then it is irrational to be ambiguity-sensitive.

Our argument builds upon a rich literature arguing that under various conditions, ambiguity-sensitivity makes one vulnerable to money pumps (see appendix A). Since different parts of the literature target different theories of ambiguity-sensitivity, it is worth mapping out the terrain in order to situate the present contribution.

According to some theories, ambiguity-sensitivity is associated with *incomplete* of preferences: cases where an agent doesn't prefer either of two options to the other, but also does not regard them as equally preferable. Several important generalized models of decision-making fall in this category, including Levi 1980 and Bewley 1986. Gustafsson (2022, forthcoming) argues that incomplete preferences make one vulnerable to money pumps.

According to other theories, agents can be ambiguity-sensitive while having complete preferences. Many important generalized models of decision-making fall in this category, including Choquet expected utility (Schmeidler 1989), maxmin expected utility (Gilboa and Schmeidler

1989), α-maxmin expected utility (Ghirardato and Marinacci 2001), "epistemic reliability-weighted" expected utility theory (Gärdenfors and Sahlin 1988), models employing second-order probabilities (Klibanoff et al. 2005, Nau 2006, Segal 1987, 1990), and Source Theory (Baillon et al. 2025).

Seidenfeld (1988, pp. 281–283) proves an important result targeting ambiguity-sensitive agents whose preferences are complete, meet certain structural conditions (involving lotteries with externally-determined probabilities), and satisfy a stochastic dominance condition.[4] The result entails that if such agents are ambiguity-averse and employ a strong form of "backwards induction" reasoning, they will be subject to money pumps.[5]

Al-Najjar and Weinstein (2009) also target ambiguity-sensitive agents whose preferences are complete. One of their examples involves an ambiguity-averse agent who chooses "naively": the agent doesn't take into account that they may in the future make a choice that looks inadvisable by their present lights. In the example, the agent ends up making a *dominated* sequence of choices — a sequence that yields a worse outcome, come what may, than a competing available sequence.[6]

Epstein and Le Breton (1993) target agents whose preferences satisfy a modified version of Savage's (1954, pp. 18–40) axioms. Their results entail that such agents sometimes violate "dynamic consistency", the principle that one's initial preferences over what to do in a particular choice situation don't change if one ends up in that situation (Hammond 1976, p. 166, Epstein and Le Breton 1993, p. 10).

Given these and similar results, the prospects for a theory of ambiguity-sensitivity that avoids vulnerability money pumps look dim. A bold theorist might declare victory, concluding that ambiguity-sensitivity in general entails vulnerability to money pumps. But a cautious theorist might worry.

What about agents who are ambiguity-sensitive, but who do not sat-

[4] Machina 1989 proves several related results.

[5] On backward induction reasoning see section 5, especially note 30.

[6] What about agents who choose in a "sophisticated" way (see note 28)? Al-Najjar and Weinstein, pp. 263–265 claim that requiring that agents exercise sophisticated choice "escapes the paradox of choosing dominated acts". So they do not take their examples to show that sophisticated ambiguity-averse agents are subject to money pumps.

isfy the conditions required by the above results? For example, what if their preferences are complete, and so they are immune to the setups in Gustafsson 2022, forthcoming? What if they do not employ the strong form of backward induction needed for the result in Seidenfeld 1988, but also do not choose in the naive way targeted by Al-Najjar and Weinstein 2009? What if their preferences do not satisfy the version of the Savage axioms assumed by Epstein and Le Breton 1993?[7]

It would be good to have an argument that these agents, too, are vulnerable to money pumps. It would be good for the argument to make only minimal assumptions about the target agent, beyond the assumption that they are ambiguity-sensitive. And it would be good for the argument to be simple. Such an argument would also address the worry that some undreamt-of future theory will allow ambiguity-sensitive agents to avoid vulnerability to money pumps.

Our aim is to give such an argument.[8] Our first step is to characterize ambiguity-sensitivity.

3 *Characterizing ambiguity sensitivity*

In the presence of ambiguity-aversion, it cannot be assumed that an agent's uncertainty can be represented by an ordinary probability function. (Per-

[7] The axioms assumed by Epstein and Le Breton 1993 include Savage's (1954, p. 31) "P4". P4 is a stakes-independence axiom sometimes glossed as "independence of beliefs from tastes" (Abdellaoui and Wakker 2020, p. 4) or "likelihood payoff independence" (Ghirardato 2002, p. 87). For considerations in favor of allowing beliefs to be stakes-sensitive (and hence against imposing P4 as a rational requirement), see Karni 1996, Gilboa 2009, 126–127, Armendt 2010. Models that allow for violations of P4 include Karni and Schmeidler 1993 and Wakker and Zank 1999.

[8] A helpful way to situate our project is in relation to Eva and Stern 2023, which provides weak, qualitative conditions that entail a version of belief dilation (Seidenfeld and Wasserman 1993). Analogously, we provide weak conditions that entail vulnerability to money pumps. In both projects, the central proof depends on a setup in which an agent learns a biconditional of the form: [unambiguous proposition] iff [ambiguous proposition]. And in both projects, the use of weak conditions avoids getting into the details of how exactly imprecise beliefs are to be represented. As a result, each project yields lessons not just for existing theories but also future theories that meet the relevant weak conditions. Thanks here to an anonymous referee.

haps such uncertainty should be represented using interval-valued probabilities, by sets of probability functions, or by other constructions.)[9] So rather than stating assumptions in terms of probabilities (and, thereby, getting into details about how ambiguous probabilities should be represented), we instead state our assumptions in terms of preferences between acts.

Let '$x \succsim y$' denote that x is at least as preferred as y. Then we define preference and indifference in the standard way:

$x \succ y =_{\text{df}} x \succsim y$ and it is not the case that $y \succsim x$.

$x \sim y =_{\text{df}} x \succsim y$ and $y \succsim x$.

Consider an agent whose preference relation $\succsim$ is transitive and complete over a set of "Savage acts" — gambles understood as functions from states of the world to outcomes.[10]

Suppose that the agent's evidence about proposition H is less ambiguous than their evidence about proposition U. For example, H might be the proposition that a particular coin lands heads, while U might be a hard-to-assess proposition about the existence of extraterrestrial life.[11] Or H might be the proposition that a red ball is drawn from an urn containing exactly 1/3 red balls, while U might be the proposition that a red ball is drawn from an urn containing an ambiguously-determined fraction of approximately 1/3 red balls (Ellsberg 1961, pp. 653–654; 2001, pp. 42–45).

Suppose that the agent is ambiguity-averse in the sense that they prefer a bet in favor of H to a bet in favor of U, while also preferring a bet against H to a bet against U. More precisely, suppose that there are outcomes a and b such that $a \succ b$ and:

[9] See for example Baillon et al. 2025, Gilboa and Schmeidler 1989, Gärdenfors and Sahlin 1988, Joyce 2005, Levi 1980, Walley 1991. For concise surveys of formal models of ambiguity aversion, see Etner et al. 2012, Gilboa and Marinacci 2013, Machina and Siniscalchi 2014, pp. 750–777. For a user-friendly introduction to such models, see Gilboa 2009, pp. 145–169. For a survey of proposals for updating beliefs that exhibit ambiguity sensitivity, see Machina and Siniscalchi 2014, pp. 793–796.

[10] Savage 1954, p. 14. For money-pump arguments that preferences should be transitive, see Gustafsson 2022, pp. 24–44 and Gustafsson forthcoming.

[11] We will understand propositions as sets of states of the world. In related literature, sometimes sets of states are referred to as "events".

(1) $\boldsymbol{H} \succ \boldsymbol{U}$,

(2) $\overline{\boldsymbol{H}} \succ \overline{\boldsymbol{U}}$,

where for any proposition E, the boldface symbol $\boldsymbol{E}$ denotes a bet in favor of E — the act that yields a (the preferred outcome) if E is true and yields b otherwise.[12] (We use bars above propositions to denote negation, so the symbol $\overline{\boldsymbol{E}}$ denotes the act that yields a if E is false and yields b otherwise.) Preferences satisfying (1) and (2) reflect ambiguity aversion because they constitute an inclination to favor some bets regarding the truth of H over corresponding bets regarding the truth of U.[13] If instead H were *more* ambiguous than U, then such preferences would reflect ambiguity-seekingness.[14]

We characterize ambiguity sensitivity with the condition that (1) and (2) are satisfied for some outcomes $a \succ b$. An advantage of this condition is that it allows our argument to assume relatively little about an agent's preferences. For instance, the argument is compatible with the view that (one's evidence for) every non-trivial proposition is somewhat ambiguous. It is compatible with the existence of agents who are ambiguity-averse with respect to some pairs of propositions but ambiguity-seeking with respect to other pairs. It does not impose the structural conditions (relating to lotteries with externally-determined probabilities) that are imposed in Seidenfeld

[12]This pattern of preferences is exhibited in the medical procedure and urn examples in section 1.

[13]Conditions (1) and (2) are in effect identified in Ellsberg 1961, p. 655, and are similar to Condition (7) from Baillon et al. 2025, p. 6, which notes that "similar conditions have been used in many models in the literature".

[14]It might be suggested that (1) be replaced by the slightly weaker condition

(1*) $\boldsymbol{H} \succsim \boldsymbol{U}$.

Behind this suggestion is the thought that a situation in which $\boldsymbol{H} \sim \boldsymbol{U}$ and $\overline{\boldsymbol{H}} \succ \overline{\boldsymbol{U}}$ also involves ambiguity sensitivity. We have some sympathy with this suggestion and have two comments. First, the above situation violates a plausible symmetry constraint: that, if an agent is ambiguity-averse with respect to H and U (and outcomes a and b), they are also ambiguity-averse with respect to $\overline{H}$ and $\overline{U}$ (and outcomes a and b). Second, at the cost of slightly complicating the argument from section 4 and appealing to an additional dominance assumption, all of our conclusions go through if we replace (1) by (1*). See appendix B.

1988, pp. 281–283. It does not assume a strong form of backward induction reasoning. It does not even require the agent to satisfy the modification of the Savage axioms assumed in Epstein and Le Breton (1993).[15,16]

Since the condition is so inclusive, the argument to follow shows that a correspondingly wide group of agents are subject to money pumps.

So, if ambiguity sensitivity is rationally permissible, then so are preferences satisfying (1) and (2).[17] We will now argue that, given plausible background assumptions, any agent with such preferences is vulnerable to a money pump.

4 *Ambiguity-sensitivity leads to preference reversals*

We proceed in two steps. In this section, we show that, if one is ambiguity-sensitive, then one is subject to an anticipated preference reversal. In the next section, we will show that, if one is subject to an anticipated preference reversal, then one is vulnerable to a money pump.[18]

By an anticipated preference reversal (for short: "preference reversal") we mean a situation in which one realizes that one's preference between acts X and Y would reverse upon finding out that proposition E is true,

[15]For example, Epstein and Le Breton (1993) includes Savage's "P4"axiom (see footnote 7), and it is possible to satisfy both (1) and (2) while violating P4.

[16]The present argument also does not assume that the target agent's preference relation is biseparable in the sense of Ghirardato and Marinacci 2001, p. 865. Several important generalized models of decision-making (under very weak background conditions) entail biseparability, such as Choquet expected utility (Schmeidler 1989), maxmin expected utility (Gilboa and Schmeidler 1989), and α-maxmin expected utility (Ghirardato and Marinacci 2001) (as noted by Maccheroni et al. 2022, p. 690n2). So the present argument does not assume that the target agent's preferences are representable in any of those theories, either.

[17]The exact characterization of ambiguity sensitivity is subtle and disputed (Schmeidler 1989, pp. 582–583, Epstein 1999, pp. 581–586, Ghirardato and Marinacci 2002, Ghirardato et al. 2004, Baillon et al. 2021, Baillon et al. 2025) and we do not claim that (1) and (2) together constitute a necessary and sufficient condition for ambiguity sensitivity. Instead, we claim that, if any interesting type of ambiguity aversion or ambiguity seekingness is rationally permissible, then so is satisfying (1) and (2) for some $a \succ b$.

[18]The general idea that preference reversals can lead to money pumps goes back at least as far as Lichtenstein and Slovic 1971, p. 53n3. See also Hammond 1976.

and would also reverse upon finding out that E is false:

$X \succ Y$, but $Y \succsim_E X$ and $Y \succsim_{\overline{E}} X$,

where, for any proposition E, the subscripted symbol $\succsim_E$ denotes the preferences the relevant agent would have upon learning E.[19] With the help of some reasonable assumptions, we will show that ambiguity sensitivity forces the existence of a preference reversal.

To introduce the needed assumptions, return to our ambiguity-sensitive agent. Assume that the agent's preferences remain transitive and complete even when they gain new evidence, that the agent's preferences over outcomes do not change over time, and that the agent is fully informed about all of the above and about what choice situation they are in.[20] Some terminology: Say that two acts *agree on proposition E* exactly when the acts produce the same outcomes at every state compatible with E.[21] Say that a proposition E is *null* if any two acts that agree on $\overline{E}$ are equally preferred. Agents thus treat null propositions as negligible when comparing acts.[22]

Assume that for any proposition E that is not null, once the agent has learned E their preferences between acts is entirely determined by the outcomes of those acts at states compatible with E. This assumption, *Learning Elimination*, in effect says that states ruled out by the agent's evidence are eliminated from consideration (when it comes to comparing acts). Learning Elimination is an undemanding and plausible constraint on how the agent takes evidence into account.

Formulated in terms of preferences, Learning Elimination says that, if (i) E is not null, (ii) acts X and X' agree on E and, (iii) acts Y and Y' agree on E, then $X \succsim_E Y$ if and only if $X' \succsim_E Y'$.[23]

[19] Compare to the negation of "Diachronic Coherence" from Hild 1998, p. 237.

[20] Here we identify an outcome with the "constant act" that has that outcome at every state.

[21] Al-Najjar and Weinstein 2009, p. 261.

[22] Savage 1954, p. 24. In typical probabilistic representations, null propositions get probability zero.

[23] Compare to "Null Complements" in Hanany and Klibanoff 2006, p. 285 and "the requirement of fact-based updating" in Al-Najjar and Weinstein 2009, p. 261.

To clarify and motivate Learning Elimination, consider a lottery in which a single number from 1 to 500 is drawn. Imagine comparing two tickets. Each ticket yields a payoff that depends on which number is drawn. Suppose that you learn (with certainty): the number drawn was either 5 or 18. Once you have learned this, then in order to compare the two tickets you only need to consider what payoffs they have for numbers 5 and 18. More generally, once some numbers have been eliminated (ruled out by your evidence), the payoff of the tickets at those eliminated numbers shouldn't affect your assessment of the tickets. And this judgment in no way depends on a prior rejection of ambiguity-sensitivity: it remains just as compelling regardless of any ambiguity in your evidence about the number drawn.

Learning Elimination generalizes the above reasoning, where states of nature play the role of lottery numbers and acts play the role of lottery tickets.[24]

Notice that, if E is not null and events X and Y agree on E, then completeness and Learning Elimination entail that $X \succsim_E Y$ (since completeness entails that $X \succsim_E X$).

Now for the proof. From before, we have the ambiguity-sensitivity assumptions:

(1) $\boldsymbol{H} \succ \boldsymbol{U}$,

(2) $\overline{\boldsymbol{H}} \succ \overline{\boldsymbol{U}}$.

[24] Learning Elimination should be sharply distinguished from Savage's (1954, pp. 21–23) "P2" (sometimes called "the sure-thing principle", or "separability" (Baccelli and Hartmann 2023)). While Learning Elimination concerns what one prefers *after* you learn a proposition E, Savage's P2 concerns one's initial preferences. P2 says that when two acts agree on a proposition H, then one's preference between those acts is independent of their (common) values on H (Ellsberg 1961, p. 648). Learning Elimination does not entail P2, and its intuitive motivation does not depend on the truth of P2. Indeed, many theories that reject or weaken P2 are compatible with imposing Learning Elimination as a rational requirement (Buchak 2013, Ghirardato and Marinacci 2001, Gilboa and Schmeidler 1989, Schmeidler 1989). It is important that the present argument does not assume P2, for such an assumption would be dialectically ineffective. That is because many theorists sympathetic with ambiguity-sensitivity consider canonical cases of ambiguity-sensitivity (such as Ellsberg's 1961, pp. 653–654 two-urn example) to be counterexamples to P2.

Let proposition S be that H and U have the same truth value.[25] Note that S is not null. For, if S were null, we would have $\boldsymbol{H} \sim \overline{\boldsymbol{U}}$ (since $\boldsymbol{H}$ and $\overline{\boldsymbol{U}}$ agree on $\overline{S}$) and also $\boldsymbol{U} \sim \overline{\boldsymbol{H}}$ (since $\boldsymbol{U}$ and $\overline{\boldsymbol{H}}$ agree on $\overline{S}$). It would then follow from (1) and transitivity that $\overline{\boldsymbol{U}} \succsim \overline{\boldsymbol{H}}$, which contradicts (2).

Also, (2) entails that $\overline{S}$ is not null (since $\overline{\boldsymbol{H}}$ and $\overline{\boldsymbol{U}}$ agree on S).

By completeness either $\overline{\boldsymbol{H}} \succ_{\overline{S}} \overline{\boldsymbol{U}}$ or $\overline{\boldsymbol{U}} \succsim_{\overline{S}} \overline{\boldsymbol{H}}$. We will show that either case leads to a preference reversal. First case:

(3) $\overline{\boldsymbol{H}} \succ_{\overline{S}} \overline{\boldsymbol{U}}$.

We have:

(4) $\boldsymbol{U} \succ_{\overline{S}} \boldsymbol{H}$. [By (3) and Learning Elimination, since $\boldsymbol{U}$ and $\overline{\boldsymbol{H}}$ agree on $\overline{S}$, and $\boldsymbol{H}$ and $\overline{\boldsymbol{U}}$ agree on $\overline{S}$]

(5) $\boldsymbol{U} \succsim_{S} \boldsymbol{H}$. [By completeness and Learning Elimination, since $\boldsymbol{U}$ and $\boldsymbol{H}$ agree on S]

Note that (1), (4), and (5) constitute a preference reversal. Second case:

(6) $\overline{\boldsymbol{U}} \succsim_{\overline{S}} \overline{\boldsymbol{H}}$.

We have:

(7) $\overline{\boldsymbol{U}} \succsim_{S} \overline{\boldsymbol{H}}$. [By completeness and Learning Elimination, since $\overline{\boldsymbol{H}}$ and $\overline{\boldsymbol{U}}$ agree on S]

Note that (2), (6), and (7) constitute a preference reversal. Since a preference reversal arises in both cases, we are done.

5 Preference reversals lead to money pumps

We have just shown that our ambiguity-sensitive agent is subject to a preference reversal. It remains to show that, if one is subject to a preference

[25] In other words, S equals [H if and only if U]. The proposition S plays a role in our proof similar in spirit to the role that "$p \leftrightarrow$ heads" plays in White 2009, p. 176, and that "$H \equiv B$" plays in Eva and Stern 2023, p. 618.

reversal, then one is vulnerable to a money pump. To show this, we will need additional assumptions about the existence of "soured" versions of acts and about how the agent's preferences constrain their choices in sequential decision problems.

For any act X, let a *souring* X^- of X be an act that is just like X except that it is certainly inferior in a dimension the agent cares about.[26] Assume the *Souring Principle*: that an act is always preferred to any souring of it. Assume *Souring Continuity* — that, for all acts X and Y such that $X \succ Y$, there exists a souring X^- of X such that $X \succ X^- \succ Y$.[27] These highly plausible souring assumptions say in effect: any given act can be made slightly less preferred by combining it with a small penalty.

Suppose that, in any decision problem, the agent chooses rationally at all nodes (choice situations) that can be reached without making any irrational choices. Finally, suppose that the agent realizes all of the above, and exercises "sophisticated" choice: when the agent can predict what will be chosen at future nodes that can be rationally reached, the agent takes those predictions into account in making their present choice. (A variant of the argument applies to agents who exercise either "myopic" or "naive" choice. See appendix C.)[28,29] In other words, sophisticated agents rely on

[26] Gustafsson 2022, p. 5.

[27] Gustafsson 2022, p. 5. Notice that Souring Continuity rules out some preferences according to which ambiguity differences only act as "tie-breakers" when other considerations are *exactly* balanced. We do not count agents with such preferences as ambiguity sensitive. That is because the rational behavior of such agents matches that of some agents who are not ambiguity-sensitive at all (for example: certain classical expected-utility maximizers). In contrast, theorists from Ellsberg on have envisioned ambiguity aversion as motivating agents to, for example, pay some small positive amount in order to avoid ambiguity. Hence ambiguity averse agents sometimes exhibit — for example, in Ellsberg's urn cases — behavior that does not match that of any expected utility maximizer. Thanks here to an anonymous referee.

[28] On sophisticated choice, see Pollak 1968, p. 203, Hammond 1976, p. 162. Choice methods most commonly contrasted with sophisticated choice include "myopic choice" (Dow 1984, p. 96), "naive choice" (Pollak 1968, pp. 202–203, Hammond 1976, p. 162), and "resolute choice" (McClennen 1985, pp. 102–103, McClennen 1990, pp. 12–13, Gauthier 1997, Buchak 2013, pp. 176–183).

[29] For appeals to resolute choice as a way of blocking diachronic exploitation arguments, see Buchak 2013, Gauthier 1997, McClennen 1985. For critical discussion and replies, see

(a restricted form of) backward induction in their decision making.[30]

Given all of the above, suppose that an agent is subject to a preference reversal:

(8) $X \succ Y$, but $Y \succsim_E X$ and $Y \succsim_{\bar{E}} X$.

To show that the agent is subject to a money pump, start by using (8) and Souring Continuity to obtain

(9) $X \succ X^- \succ X^{--} \succ Y$,

where X^{--} is a souring of X^-, which, in turn, is a souring of X. By (8), (9), the Souring Principle, and transitivity:

(10) $Y \succ_E X^-$ and $Y \succ_{\bar{E}} X^-$.

Now, consider the following decision problem faced by an agent who starts out with X^-:

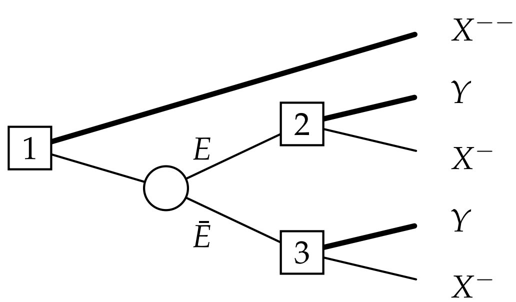

Rabinowicz 1995, Gustafsson 2022, §7.

[30] Selten 1978, pp. 136–138 has challenged backward induction due its recommendation of immediate defection in the centipede game. But see Sobel 1993, pp. 130–131 for a response. Note also that the present argument (unlike standard backward induction arguments) does not assume that the agent chooses rationally at every node in the decision problem. It only assumes that the agent chooses rationally at every node *that can be reached without irrational choices*. This weaker assumption suffices because the decision problem in this section is *BI-terminating*: it is such that standard backward induction reasoning prescribes options that are not followed by any further choices. See Rabinowicz 1998, pp. 100–101. By depending on this weaker assumption, we avoid Binmore's (1987, pp. 196–200) complaint that backward-induction reasoning requires that the agent assumes that they will choose rationally even at nodes that could only be reached by making irrational choices, which is dubious.

Here, the boxes represent choice nodes and the circle represents a chance node. At the chance node, chance goes up if and only if E is true.[31] In this decision problem, the agent can choose to pay a small penalty (give up X^- in exchange for X^{--}) in order to avoid first learning whether E is true and then, afterwards, choosing between X^- and Y.

Suppose, for contradiction, that in this problem *the agent goes down at node 1*. Then since the agent chooses rationally at node 1:

> Nodes 2 and 3 can each be reached without any irrational choices. So the agent would choose rationally at each of nodes 2 and 3.
>
> So, by (10), the agent, upon reaching either node 2 or node 3, goes up and ends up with Y.
>
> Taking this into account at node 1, the agent sees that: if they go down at node 1, they will end up with Y.
>
> Since $X^{--} \succ Y$, the agent goes up at node 1. Contradiction.

We have derived a contradiction from the assumption that the agent goes down at node 1. So the agent goes up at node 1. But to go up at node 1 is to choose X^{--} when they could have kept X^- for free (by turning down all offers — that is, by always going down). So the agent is subject to a money pump.

6 *Conclusion*

We have argued that, under extremely general conditions, agents who are ambiguity sensitive are subject to money pumps. Given that vulnerability to money pumps is a mark of irrationality, it follows that it is irrational to

[31] Despite the terminology, we do not assume that "chance nodes" are associated with chance processes or events with well-defined probabilities. Instead we simply mean nodes at which the agent learns a proposition without making a choice. Hammond (1988, p. 31) calls such nodes "natural nodes".

be ambiguity-sensitive.[32] So we have a response to Ellsberg, who expresses sympathy for an ambiguity-averse agent and claims:

> It will not be true that [ambiguity-averse] behavior . . . exhibits intransitivities, or amounts to "throwing away utility" (as would be true, for example, if it led him occasionally to choose strategies that were strongly "dominated" by others).[33]

We can now reply: under very general conditions, ambiguity-sensitive agents *do* throw away utility.

[32]Of course the thought that money pumps are a mark of irrationality would need to be supplemented by an appropriate analogue of the "converse Dutch book theorem" showing that suitable preferences (that are not ambiguity sensitive, for example) are *not* subject to money pumps. And the many objections to inferring irrationality from exploitability would have to be addressed. On both points see Hájek 2005, 2008.

[33]Ellsberg 1961, p. 663.

Appendices

A Prior work

It is impossible to concisely do justice to the rich literature exploring how ambiguity sensitivity can lead to violations of dynamic consistency conditions. The comments below are necessarily selective and synoptic.

Working in a framework where the objects of choice are "horse lotteries" (functions from states to lotteries with externally specified probabilities) that obey a reduction axiom, Seidenfeld (1988, pp. 281–283) shows that agents who satisfy stochastic dominance and an Archimedean condition but who (due to ambiguity aversion) violate an independence condition (that preference between two given lotteries remain unchanged if each is mixed with a common third lottery) are subject to episodes of "sequential incoherence" in which they make a sequence of choices yielding a lottery that they regarded (at the outset) as inferior to an outcome they could have guaranteed by way of a competing sequence of choices.

Machina 1989, pp. 1636–1638, working in a similar framework, shows that violations of "replacement separability, mixture separability, or the independence axiom" can lead to the choice of a dominated plan (a plan that yields a worse outcome, come what may, than another available plan).

Hammond (1988) proves a subjective expected-utility representation theorem whose main assumption is "consequentialism" — roughly, that a sequence of choices is acceptable in a dynamic decision problem exactly if the plan corresponding to that sequence is acceptable in the "flattened" (normal-form) counterpart of that decision problem (Steele 2010, p. 470). Given Hammond's other background assumptions, it follows that ambiguity-averse agents (whose preferences do not admit of subjective expected-utility representations) must violate consequentialism.

Al-Najjar and Weinstein (2009), working in a framework where the objects of choice are Savage acts, shows that ambiguity-averse agents who choose "naively" (without taking into account that their future choices may be governed by preferences differing from their present ones) sometimes choose dominated plans, but claims that requiring that agents exercise "sophisticated" choice (see note 28) "escapes the paradox of choosing dominated acts" (Al-Najjar and Weinstein 2009, pp. 263–265).

Epstein and Le Breton (1993) work with a modified version of the Savage axioms (Savage 1954, pp. 18–40) from which the sure thing principle (P2) has been removed. Applying a result from Machina and Schmeidler 1992, p. 766, Epstein and Le Breton show that the modified axioms, together with dynamic-consistency assumptions, entail the existence of a representation via a uniquely determined probability function. It follows that, given the modified Savage axioms, preferences not representable by way of orthodox probability functions (including preferences that reflect ambiguity aversion) must violate dynamic consistency.

White (2009, pp. 175–181) considers a case of probability dilation (Seidenfeld and Wasserman 1993) in a framework where belief states are represented by sets of probability function and argues that various candidate updating and choice rules have unattractive features, including anticipated belief changes that violate of a version of the Reflection Principle (van Fraassen 1984, p. 244).

B *Proof using a weakened ambiguity-sensitivity condition*

The goal is to give a version of the proof from section 4 (that, if one is ambiguity-sensitive, one is subject to a preference reversal) that replaces the ambiguity sensitivity conditions

(1) $\boldsymbol{H} \succ \boldsymbol{U}$,

(2) $\overline{\boldsymbol{H}} \succ \overline{\boldsymbol{U}}$,

with the slightly more inclusive conditions

(1*) $\boldsymbol{H} \succsim \boldsymbol{U}$,

(2) $\overline{\boldsymbol{H}} \succ \overline{\boldsymbol{U}}$.

Notation: for acts X, Y, and proposition E, let $(X)_E(Y)$ be the act that matches X on E and matches Y on $\overline{E}$. The modified proof depends on *Strong Dominance*, which says that making an act less preferred at every state in a non-null proposition makes the act less preferred overall.[34] Formally, Strong Dominance says that $X \succ (X^-)_E(X)$ whenever proposition E is not null and act X^- is a souring of X.

Given this assumption, only two changes to the proof are necessary: replace all references to (1) with references to (1*); and replace the treatment of the "first case" starting at (5) with the following reasoning:

Use (4), Souring Continuity, and the Souring Principle to obtain

(11) $\boldsymbol{U} \succ_{\overline{S}} \boldsymbol{U}^- \succ_{\overline{S}} \boldsymbol{H}$,

where $\boldsymbol{U}^-$ is a souring of $\boldsymbol{U}$. Let act $V = (\boldsymbol{U})_S(\boldsymbol{U}^-)$. We have

(12) $\boldsymbol{U} \succ V$. [By Strong Dominance, since $V = (\boldsymbol{U})_S(\boldsymbol{U}^-)$]

(13) $\boldsymbol{H} \succ V$. [By (1*), (12), and transitivity]

(14) $V \succsim_S \boldsymbol{H}$. [By completeness and Learning Elimination, since V and $\boldsymbol{H}$ agree on S]

(15) $V \succsim_{\overline{S}} \boldsymbol{U}^-$. [By Learning Elimination, since V and $\boldsymbol{U}^-$ agree on $\overline{S}$]

(16) $V \succ_{\overline{S}} \boldsymbol{H}$. [By (11), (15), and transitivity]

Note that (13), (14), and (16) constitute a preference reversal.

C *Money pump for naive or myopic choosers*

Section 5 shows that any sophisticated chooser who is subject to an anticipated preference reversal is vulnerable to a money pump. Here we modify that proof to show that the same goes for naive choosers and myopic choosers. *Naive* choosers perform an action consistent with a plan (specification of what to do at each node) that gets them a result that they

[34] This assumption is termed "strong monotonicity" by Baillon et al. 2025, p. 3.

now most prefer, ignoring whether they will later stick to that plan (Pollak 1968, pp. 202–203, Hammond 1976, p. 162, Gustafsson 2022, p. 6). *Myopic* choosers consider each potential trade as if it were their last. (Dow 1984, p. 96, Gustafsson 2022, p. 6).

We appeal to one small additional assumption: *Sweetening Continuity* says that whenever $X \succ Y$, there exists Y^+ such that $X \succ Y^+ \succ Y$, where Y is a souring of Y^+.

The modified proof continues from (8), which we rewrite here:

(8) $X \succ Y$, but $Y \succsim_E X$ and $Y \succsim_{\bar{E}} X$.

From (8) we have by successive applications of Souring Continuity and Sweetening Continuity,

(17) $X \succ X^- \succ Y^+ \succ Y$,

where X^- is a souring of X, and Y is a souring of Y^+. Also from (8) we have by the Souring Principle and transitivity,

(18) $Y \succ_E X^-$ and $Y \succ_{\bar{E}} X^-$.

Now, consider the following decision problem faced by an agent who starts out with Y^+:

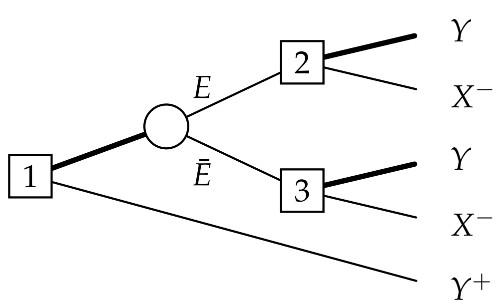

In this decision problem, the agent may at node 1 choose to go up and give up Y^+ in exchange for X^-. At nodes 2 and 3, after finding out whether E is true, the agent may choose to go up and give up X^- in exchange for Y.

If the agent is *naive*, they will go up at node 1. That is because the plan "go up at node 1 and then go down at node 2 or 3" (if implemented!) gets them X^-, which they prefer to the result of implementing any other plan (by (17)).

Alternatively, if the agent is *myopic*, they will also go up at node 1, giving up Y^+ in exchange for X^-. That is because they prefer X^- to Y^+ (by (17)).

Upon reaching either node 2 or node 3, either sort of agent will go up, since they will then prefer Y to X^- (by (18)). So either sort of agent will end up with Y, when they could have kept Y^+ for free.

References

Mohammed Abdellaoui and Peter P. Wakker. Savage for dummies and experts. *Journal of Economic Theory*, 186:104991, 2020.

Nabil I. Al-Najjar and Jonathan Weinstein. The ambiguity aversion literature: A critical assessment. *Economics and Philosophy*, 25(3):249–284, 2009.

Brad Armendt. Stakes and beliefs. *Philosophical Studies*, 147(1):71–87, 2010.

Jean Baccelli and Lorenz Hartmann. The sure-thing principle. *Journal of Mathematical Economics*, 109(102915), 2023.

Aurélien Baillon, Han Bleichrodt, Chen Li, and Peter P. Wakker. Belief hedges: Measuring ambiguity for all events and all models. *Journal of Economic Theory*, 198:105353, 2021.

Aurélien Baillon, Han Bleichrodt, Chen Li, and Peter P. Wakker. Source theory: A tractable and positive ambiguity theory. *Management Science*, February 2025.

Truman Bewley. Knightian decision theory: Part 1. Cowles Foundation Discussion Papers 807, Yale University, New Haven, CT, November 1986.

Ken Binmore. Modeling rational players: Part i. *Economics and Philosophy*, 3 (2):179–214, 1987.

Lara Buchak. *Risk and Rationality*. Oxford University Press, Oxford, 2013.

Gregory K. Dow. Myopia, amnesia, and consistent intertemporal choice. *Mathematical Social Sciences*, 8(2):95–109, 1984.

Daniel Ellsberg. Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics*, 75(4):643–669, 1961.

Daniel Ellsberg. *Risk, Ambiguity and Decision*. Garland, New York, 2001.

Larry G. Epstein. A definition of uncertainty aversion. *The Review of Economic Studies*, 66(3):579–608, 1999.

Larry G. Epstein and Michel Le Breton. Dynamically consistent beliefs must be Bayesian. *Journal of Economic Theory*, 61(1):1–22, 1993.

Johanna Etner, Meglena Jeleva, and Jean-Marc Tallon. Decision theory under ambiguity. *Journal of Economic Surveys*, 26(2):234–270, 2012.

Benjamin Eva and Reuben Stern. Comparative opinion loss. *Philosophy and Phenomenological Research*, 107(3):613–637, 2023.

William Fellner. Distortion of subjective probabilities as a reaction to uncertainty. *The Quarterly Journal of Economics*, 75(4):670–689, 1961.

David Gauthier. Resolute choice and rational deliberation: A critique and a defense. *Noûs*, 31(1):1–25, 1997.

Paolo Ghirardato. Revisiting Savage in a conditional world. *Economic Theory*, 20(1):83–92, 2002.

Paolo Ghirardato and Massimo Marinacci. Risk, ambiguity, and the separation of utility and beliefs. *Mathematics of Operations Research*, 26(4): 864–890, 2001.

Paolo Ghirardato and Massimo Marinacci. Ambiguity made precise: A comparative foundation. *Journal of Economic Theory*, 102(2):251–289, 2002.

Paolo Ghirardato, Fabio Maccheroni, and Massimo Marinacci. Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory*, 118(2): 133–173, 2004.

Itzhak Gilboa. *Theory of Decision under Uncertainty*. Econometric Society Monographs. Cambridge University Press, Cambridge, 2009.

Itzhak Gilboa and Massimo Marinacci. Ambiguity and the Bayesian paradigm. In Daron Acemoglu, Eddie Dekel, and Manuel Arellano, editors, *Advances in Economics and Econometrics: Tenth World Congress: Volume 1: Economic Theory*, volume 1 of *Econometric Society Monographs*, pages 179–242. Cambridge University Press, Cambridge, 2013.

Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2), 1989.

Johan E. Gustafsson. *Money-Pump Arguments*. Cambridge University Press, 2022.

Johan E. Gustafsson. A behavioural money-pump argument for completeness. *Theory and Decision*, forthcoming.

Peter Gärdenfors and Nils-Eric Sahlin. Unreliable probabilities, risk taking and decision making. In Peter Gärdenfors and Nils-Eric Sahlin, editors, *Decision, Probability and Utility: Selected Readings*, pages 313–334. Cambridge University Press, 1988.

Peter J. Hammond. Changing tastes and coherent dynamic choice. *The Review of Economic Studies*, 43(1):159–173, 1976.

Peter J. Hammond. Consequentialist foundations for expected utility. *Theory and Decision*, 25(1):25–78, 1988.

Eran Hanany and Peter Klibanoff. Updating preferences with multiple priors. *SSRN Electronic Journal*, 2006.

Matthias Hild. The coherence argument against conditionalization. *Synthese*, 115:229–258, 1998.

Alan Hájek. Scotching Dutch books? *Philosophical Perspectives*, 19(1):139–151, 2005.

Alan Hájek. Dutch book arguments. In Paul Anand, Prasanta Pattanaik, and Clemens Puppe, editors, *The Oxford Handbook of Rational and Social Choice*, pages 173–195. Oxford University Press, 2008.

James M. Joyce. How probabilities reflect evidence. *Philosophical Perspectives*, 19(1):153–178, 2005.

Mark Kaplan. *Decision Theory as Philosophy*. Cambridge University Press, Cambridge, 1996.

Edi Karni. Probabilities and beliefs. *Journal of Risk and Uncertainty*, 13(3): 249–262, 1996.

Edi Karni and David Schmeidler. On the uniqueness of subjective probabilities. *Economic Theory*, 3(2):267–277, 1993.

John Maynard Keynes. *A Treatise on Probability*. Macmillan and Co., London, 1921.

Peter Klibanoff, Massimo Marinacci, and Sujoy Mukerji. A smooth model of decision making under ambiguity. *Econometrica*, 73(6):1849–1892, 2005.

Frank H Knight. *Risk, Uncertainty, and Profit*. Houghton Mifflin, Boston, first edition, 1921.

Stephen F. LeRoy and Larry D. Singell. Knight on risk and uncertainty. *Journal of Political Economy*, 95(2):394–406, 1987.

Isaac Levi. On indeterminate probabilities. *Journal of Philosophy*, 71(13): 391–418, 1974.

Isaac Levi. *The Enterprise of Knowledge: An Essay on Knowledge, Credal Probability, and Chance*. MIT Press, Cambridge, MA, 1980.

Sarah Lichtenstein and Paul Slovic. Reversals of preference between bids and choices in gambling decisions. *Journal of Experimental Psychology*, 89 (1):46–55, 1971.

Fabio Maccheroni, Massimo Marinacci, and Jingni Yang. On the cardinal utility equivalence of biseparable preferences. *Theory and Decision*, 92 (3–4):689–701, 2022.

Mark J. Machina. Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature*, 27(4):1622–1668, 1989.

Mark J. Machina and David Schmeidler. A more robust definition of subjective probability. *Econometrica*, 60(4):745, 1992.

Mark J. Machina and Marciano Siniscalchi. Ambiguity and ambiguity aversion. In *Handbook of the Economics of Risk and Uncertainty*, volume 1, pages 729–807. Elsevier, 2014.

Edward F. McClennen. Prisoner's dilemma and resolute choice. In Richmond Campbell and Lanning Sowden, editors, *Paradoxes of Rationality and Cooperation: Prisoner's Dilemma and Newcomb's Problem*, pages 94–104. University of British Columbia Press, Vancouver, 1985.

Edward F. McClennen. *Rationality and Dynamic Choice: Foundational Explorations*. Cambridge University Press, Cambridge, 1990.

Robert F. Nau. Uncertainty aversion with second-order utilities and probabilities. *Management Science*, 52(1):136–145, 2006.

R. A. Pollak. Consistent planning. *The Review of Economic Studies*, 35(2): 201–208, 1968.

Wlodek Rabinowicz. To have one's cake and eat it, too: Sequential choice and expected-utility violations. *The Journal of Philosophy*, 92(11):586–620, 1995.

Wlodek Rabinowicz. Grappling with the centipede: Defence of backward induction for bi-terminating games. *Economics and Philosophy*, 14(1): 95–126, 1998.

Zvi Safra and Uzi Segal. A lot of ambiguity. *Journal of Economic Theory*, 200: 105393, 2022.

Leonard J. Savage. *The Foundations of Statistics*. Wiley, New York, 1954.

David Schmeidler. Subjective probability and expected utility without additivity. *Econometrica*, 57(3):571–587, 1989.

Uzi Segal. The Ellsberg paradox and risk aversion: An anticipated utility approach. *International Economic Review*, 28(1):175, 1987.

Uzi Segal. Two-stage lotteries without the reduction axiom. *Econometrica*, 58(2):349, 1990.

Teddy Seidenfeld. Decision theory without "independence" or without "ordering". *Economics and Philosophy*, 4(2):267, 1988.

Teddy Seidenfeld and Larry Wasserman. Dilation for sets of probabilities. *The Annals of Statistics*, 21(3):1139–1154, 1993.

Reinhard Selten. The chain store paradox. *Theory and Decision*, 9(2):127–159, 1978.

Jordan Howard Sobel. Backward-induction arguments: A paradox regained. *Philosophy of Science*, 60(1):114–133, 1993.

Katie Siobhan Steele. What are the minimal requirements of rational choice? arguments from the sequential-decision setting. *Theory and Decision*, 68 (4):463–487, 2010.

Bas C. van Fraassen. Belief and the will. *Journal of Philosophy*, 81(5):235–256, 1984.

Peter P. Wakker and Horst Zank. State dependent expected utility for Savage's state space. *Mathematics of Operations Research*, 24(1):8–34, 1999.

Peter Walley. *Statistical Reasoning with Imprecise Probabilities*. Chapman & Hall, 1991.

Roger White. Evidential symmetry and mushy credence. In T. Szabo Gendler and J. Hawthorne, editors, *Oxford Studies in Epistemology*, pages 161–186. Oxford University Press, 2009.