A Universal Money Pump for the Myopic, Naive, and Minimally Sophisticated

Johan E. Gustafsson*

Draft: February 10, 2024 at 4:35 p.m.

ABSTRACT. The money-pump argument aims to show that cyclic preferences are irrational. The argument can be based on a number of different money-pump schemes that vary in what needs to be assumed about the agent. The Standard Money Pump works for myopic and naive agents, but not for sophisticated agents who use backward induction. The Upfront Money Pump works for sophisticated agents, but not for myopic or naive agents. In this paper, I present a new money pump, the Universal Money Pump, that works for myopic, naive agents, and sophisticated agents. Moreover, the Universal Money Pump (just like the Upfront Money Pump) also works for minimally sophisticated agents who need not assume that they will choose rationally at nodes that can only be reached by irrational choices.

Suppose you prefer A to B, B to C, and C to A. Your preferences are cyclic. It may seem obvious that such preferences must be irrational. But there are examples of seemingly rational preferences of this form.

For a first example, suppose that A, B, and C are bottles of wine and that A is pricier than B and B is pricier than C. You can't taste the difference between A and B nor between B and C, but you do prefer C to A for its superior taste. And, when you can't taste any difference, you prefer the cheaper bottle. So you prefer A to B and B to C (and, as mentioned, C to A).[1]

For another example, suppose that (as a champion of democracy) you prefer what a majority of the people prefer. About a third of the people prefer A to B and B to C, about a third prefer B to C and C to A, and about a third prefer C to A and A to B. So, following the majority, you prefer A to B, B to C, and C to A.[2]

* I would be grateful for any thoughts or comments on this paper, which can be sent to me at *johan.eric.gustafsson@gmail.com*.

[1] Armstrong (1939, p. 457n1) and Dummett (1984, p. 34) present similar examples.

[2] Condorcet 1785, p. lxi; 1976, p. 54, Dodgson 1876, pp. 8–12, and Black 1948, pp. 32–3.

For an auditory example, suppose that A, B, and C are the first, second, and third thirds, respectively, of a seemingly continuously ascending Shepard tone loop — which, in fact, ends up at the same pitch it started.[3] And suppose you prefer a segment to another segment if it appears to have lower pitch. Then you prefer A to B, B to C, and C to A.

A final example concerns outcomes with variable populations. Let A be an outcome where people have very good lives, and let C be an outcome just like A except that there are some additional people with significantly worse lives but still good lives. Since C is just like A but with an addition of good lives, you prefer C to A. Now, let B be an outcome with the same people as C except that everyone has equally good lives and the average quality of life is higher than in C, though lower than in A. Since B is more equal and has a higher average quality of life, you prefer B to C. But, since the quality of life is higher in A than in B, you prefer A to B.[4] Hence you prefer A to B, B to C, and C to A.

Letting '$X \succ Y$' denote that X is (strictly) preferred to Y, we can represent your preferences as follows:

(1) $A \succ B \succ C \succ A$.

Your preferences violate the following form of acyclicity:[5]

> *Three-Step Acyclicity* If $X \succ Y \succ Z$, then it is not the case that $Z \succ X$.

Given the above examples of seemingly rational cyclic preferences, we may wonder whether Three-Step Acyclicity really is a requirement of rationality. The standard argument that it is so is the money-pump argument. A *money-pump argument* for an alleged requirement of rationality (such as Three-Step Acyclicity) is an argument that otherwise rational agents who violate the requirement would in some possible situation end up paying for something they could have kept for free even though they knew in advance what decision problem they were facing.[6]

The standard version runs as follows: Suppose that, initially, you can walk away with A — that is, you end up with A if you turn down all trades.

[3] Shepard 1964, pp. 2347–9.

[4] This is the mere-addition paradox. See McMahan 1981, pp. 122–3 and Parfit 1982, pp. 158–60.

[5] Samuelson 1947, p. 151 and Sen 1977, p. 62.

[6] Gustafsson 2022, pp. 1–2.

First, you are offered a trade from A to C. Since you prefer C to A, you accept this trade. Then you are offered a trade from C to B. Since you prefer B to C, you also accept this trade. Finally, you are offered a trade from B to A for a small payment. Since you prefer A to B with some margin, there should be some small amount of money that you would be willing to pay to get A instead of B. Likewise, since you prefer C to A, you should prefer C without any payment to A with a payment. Accordingly, there is a soured version A^- of A such that

(2) $A \succ A^-$, and $C \succ A^- \succ B$.

So we let the third offer be an offer to trade from B to A^-. Since you prefer A^- to B, you accept this final offer and end up with A^- (that is, you pay for A) even though you could have walked away with A (that is, you could have keep A for free).[7]

We can diagram this set-up with a decision tree:[8]

The Standard Money Pump

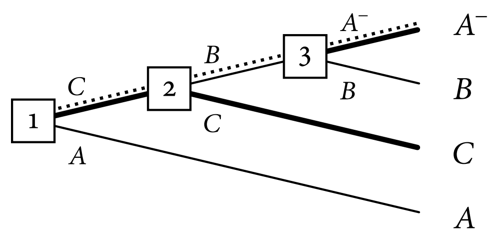

$A \succ B \succ C \succ A \succ A^-$, and $C \succ A^- \succ B$.

Here, the squares represent the choice nodes where you are offered the trades. Accepting a trade corresponds to going up at a node, and turning a trade down corresponds to going down.[9] The outcome on the upper right of each square is what you get if you accept the trade. The outcome on the lower right of each square is what you get if you turn the trade down.

The standard version of the money-pump argument, as presented earlier, assumes that the agent is myopic. *Myopic* agents make their choices under the assumption that they will walk way from all future trades. If

[7] Davidson et al. 1955, p. 146, Edwards et al. 1965, p. 273, and Pratt et al. 1965, ch. 2, p. 10.

[8] Rabinowicz 1995, p. 393.

[9] Following Rabinowicz 2008a, p. 152.

we treat the walk-away option (that is, what you get if you turn down all future trades) as what you currently possess, then being myopic can be thought of as choosing between possessions without taking future choices into account.[10] (The myopic choices are marked by a dotted lines.)

Nevertheless, the Standard Money Pump also works for naive agents. *Naive* agents (i) consider the outcomes of all available plans and assess which of these outcomes are choice-worthy in a choice between all of them and (ii) choose in accordance with a plan to end up with one of these choice-worthy outcomes, without taking into consideration whether they would later depart from that plan.[11] To be a naive agent in the Standard Money Pump, you need to make a choice between all potential outcomes — that is, between A, A^-, B, and C. But, given the preferences in (1) and (2), we can't rely on maximization — that is, we can't rely on the following rule:[12]

> *The Maximization Rule* It is rationally permitted to choose an outcome X if and only if there is no feasible outcome Y such that $Y \succ X$.

The trouble is that, for each potential outcome, there is another that's preferred to it. So, rather than maximization, we adopt the following rule:[13]

> *The Uncovered-Choice Rule* It is rationally permitted to choose an outcome X if and only if there is no feasible outcome Y such that $Y \succ X$ and, for all feasible outcomes Z, it holds that $Y \succ Z$ if $X \succ Z$.

Choosing between all potential outcomes at node 1 with the Uncovered-Choice Rule, you only deem A, B, and C as choice-worthy. So it's rationally permitted both to accept and to turn down the trade at node 1. So you may rationally accept the trade at node 1. Then, at node 2, the potential outcomes are just A^-, B, and C. Choosing between them with the Uncovered-Choice Rule, you deem all of them as choice-worthy. Hence it's rationally permitted both to accept and to turn down the trade at node 2. So you may rationally accept the trade at node 2. At node 3,

[10] See Dow 1984, p. 96 and McClennen 1990, pp. 11–12.

[11] Pollak 1968, pp. 202–3 and Hammond 1976, p. 162.

[12] Uzawa 1956, p. 37.

[13] Schwartz 1990, p. 21. See also Miller 1980, pp. 72–4.

the potential outcomes are A^- and B. Choosing between them with the Uncovered-Choice Rule, you only deem A^- as choice-worthy. So you accept the trade at node 3. Hence, with rationally permitted choices, you end up with A^- even though you could have walked away with A.

The Standard Money Pump for naive agents is an example of a permitting, non-forcing money pump. A money pump is *forcing* if and only if the agent is rationally required, at each step, to go along with the exploitation. A money pump is *permitting* if and only if, at each step, the agent is rationally permitted to go along with the exploitation. A money pump is *non-prohibiting* if and only if, at each step, the agent is not rationally prohibited from going along with the exploitation.[14] Finally, a money pump is *non-forcing* if and only if it is non-prohibiting and, at some step, the agent is not rationally required to go along with the exploitation.[15]

While the money pumps we'll deploy against other kinds of agents will be forcing, we can't do better than permitting money pumps when we deal with naive agents. The trouble, roughly, is that — as long as all of A, A^-, B, and C are potential outcomes — the plans leading to A (the walk-away outcome) will be rationally permitted. And, as long as A is a potential outcome, it can't be rationally required to follow a plan that potentially leads to A^- and not potentially to A. So, when A is a potential outcome, the only way it can be rationally required to follow a plan that potentially leads to A^- and not potentially to A is when, in addition, not all of A, B, and C are potential outcomes. But, if A, B, and C aren't all potential outcomes, we can't exploit the fact that the agent has cyclic preferences over these outcomes.

So far, we have only dealt with myopic and naive agents. Some agents, however, are neither myopic nor naive. A problem with the Standard Money Pump is that cyclic preferrers can avoid being money pumped if they are sophisticated, rather than myopic or naive. *Sophisticated* agents, assuming that they choose rationally at all future choice nodes, make their

[14] The distinction between permitting and non-prohibiting money pumps relies on the distinction between an option being rationally permitted and it merely not being rationally prohibited. Suppose that you have a preferential gap between two options (that is, neither option is at least as preferred as the other). It then seems that rationality could be silent — so that neither option is prohibited and neither options is permitted. Contrast this case with the case where you are indifferent between two options. Then, it seems that neither option is prohibited but each option is permitted. Since these cases differ in permissibility but not in non-prohibition, permissibility must differ from non-prohibition. See Rabinowicz 2008b, p. 26 and Gustafsson 2020, pp. 121-2.

[15] Gustafsson and Espinoza 2010, pp. 761–2 and Gustafsson 2022, p. 27.

choices using backward induction.[16] And, to use *backward induction* is to predict what one would choose at later choice nodes and to take those predictions into account when one is choosing at earlier nodes.[17]

To see that sophisticated agents with the preferences in (1) and (2) avoid the Standard Money Pump, suppose that you're a sophisticated agent with those preferences. At node 3, you would accept the trade from B to A^-, since you prefer A^- to B. Taking this prediction into account, the choice at node 2 is effectively between A^- (accepting the trade) and C (turning it down). Since you prefer C to A^-, you would turn down the trade at node 2. Taking this prediction into account, the choice at node 1 is effectively between C (accepting the trade) and A (turning it down). Since you prefer C to A, you would accept the trade at node 1. So you end up with C and avoid being money pumped. (The choices that are recommend by backward induction — assuming rational choices at future nodes — are marked with thicker lines in the decision tree.)

But there are money pumps that work for sophisticated agents. Consider the following decision problem:[18]

The Upfront Money Pump

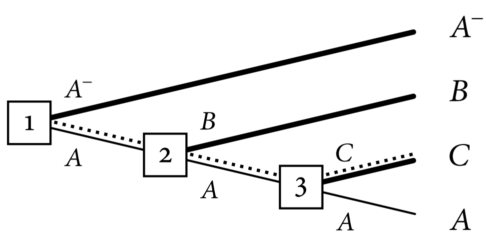

$A \succ B \succ C \succ A \succ A^-$, and $C \succ A^- \succ B$.

To see that sophisticated agents with the preferences in (1) and (2) can't avoid exploitation in the Upfront Money Pump, suppose again that you're a sophisticated agent with those preferences. At node 3, you would accept the trade from A to C, since you prefer C to A. Taking this prediction into account at node 2, the choice at that node is effectively between B

[16] Pollak 1968, p. 203, and Hammond 1976, p. 162.

[17] von Neumann and Morgenstern 1944, pp. 116–17, Strotz 1955–1956, p. 173, and Raiffa and Schlaifer 1961, pp. 7–8.

[18] Gustafsson and Rabinowicz 2020, p. 583. For a similar construction in voting theory, see Moulin 1983, pp. 96–7. Rabinowicz (2000, p. 141) presents another money pump for sophisticated agents, but — unlike the Upfront Money Pump — it does not work for minimally sophisticated agents (defined later).

(accepting the trade) and C (turning it down). Since you prefer B to C, you would accept the trade at node 2. Taking this prediction into account at node 1, the choice at that node is effectively between A^- (accepting the trade) and B (turning it down). Since you prefer A^- to B, you accept the trade at node 1. Hence you end up with A^- even though you could have walked away with A.

The standard form of backward induction on which sophisticated agents rely can be challenged, however. It is based on the dubious assumption that the agent would choose rationally even at future choice nodes that can only be reached by irrational choices.[19] There is, however, a weaker form of sophistication which doesn't rely on this assumption. *Minimally sophisticated* agents choose using backward induction based on what they predict they will choose in the future assuming that, at nodes that can be reached without making any irrational choices, they retain (i) their rationality and (ii) their trust in their rationality at nodes that can be reached without making any irrational choices.[20] Hence minimally sophisticated agents predict that they will conform to the requirements of rationality (leaving open exactly what those requirements are) at future choice nodes in case they haven't, by then, violated those standards. Unlike standard backward induction, minimal sophistication allows (but does not require) that the agents predict that they will choose irrationally at some nodes that can only be reached by violating the requirements of rationality.

In some decision problems, minimally sophisticated agents need not make the same choices as sophisticated agents with the same preferences. Consider, once more, the Standard Money Pump. And suppose (consistently with minimal sophistication) that you predict that you would, irrationally, accept the trade at node 2 even though you predict that you would (rationally) accept the trade at node 3. We mark this predicted irrational choice with a thick dashed line in the decision tree:

[19] Binmore 1987, pp. 196–200, Bicchieri 1988, pp. 334–5, and Pettit and Sugden 1989, pp. 171–4.

[20] Rabinowicz 1998, pp. 108–9 and Gustafsson and Rabinowicz 2020, p. 583.

The Standard Money Pump
(with a predicted irrational choice)

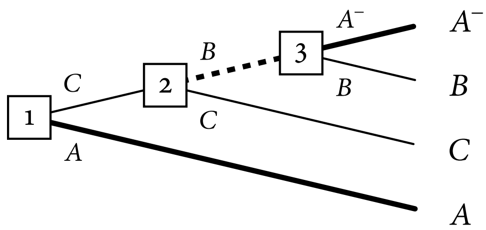

$A \succ B \succ C \succ A \succ A^-$, and $C \succ A^- \succ B$.

Given the prediction that you would accept the trades at nodes 2 and 3, the choice at node 1 is effectively between A^- (accepting the trade) and A (turning it down). Since you prefer A to A^-, you turn down the trade at node 1. Given your prediction, it would be irrational to accept the trade at node 1. Hence the predicted irrational choice at node 2 would follow an initial irrational choice at node 1. So a minimally sophisticated agent with the preferences in (1) and (2) can consistently turn down the initial trade in the Standard Money Pump.

But, in some decision problems, minimally sophisticated agents must make the same choices as sophisticated agents. One such decision problem is the Upfront Money Pump. We will show this by a proof by contradiction.

Given that you're a minimally sophisticated agent, we note that you choose using backward induction based on what you predict you will choose in the future assuming that, at nodes that can be reached without making any irrational choices, you retain (i) your rationality and (ii) your trust in your rationality at nodes that can be reached without making any irrational choices.

Firstly, we assume (for proof by contradiction) that node 3 in the Upfront Money Pump can be reached without making any irrational choices. Then, at all choice nodes, you retain your rationality and your trust in your rationality at these nodes. Accordingly, you would accept the trade from A to C at node 3, since you prefer C to A. But, if so, the choice at node 2 is effectively a choice between B (accepting the trade) and C (turning it down). Since you prefer B to C, the choice to turn down the trade at node 2 was irrational, which contradicts our assumption that node 3 can be reached without making any irrational choices.

Secondly, we assume (for proof by contradiction) that node 2 can be reached without making any irrational choices. Then, at nodes 1 and 2, you retain your rationality and your trust in your rationality at these nodes. Since we have already shown that node 3 cannot be reached without making irrational choices, it follows that it's irrational to turn down the trade at node 2. So you would accept the trade at node 2. But, if so, the choice at node 1 is effectively between A^- (accepting the trade) and B (turning it down). Since you prefer A^- to B, the choice to turn down the trade at node 1 was irrational, which contradicts our assumption that node 2 can be reached without making any irrational choices.

Hence it's irrational to turn down the trade at node 1. So you accept the initial trade and end up with A^-, even though you could have walked away with A.[21]

So neither sophisticated nor minimally sophisticated agents avoid exploitation in the Upfront Money Pump. How do myopic and naive agents fare?

If you are a myopic agent, you turn down the trade from A to A^- at node 1, since you prefer A to A^-. Then, at node 2, you turn down the trade from A to B, since you prefer A to B. Finally, at node 3, you accept the trade from A to C, since you prefer C to A. Hence you end up with C and avoid exploitation.

If you're a naive agent following the Uncovered-Choice Rule at node 1, you only deem A, B, and C as choice-worthy among the potential outcomes A, A^-, B, and C. So you turn down the trade at node 1. Then, at node 2, you deem all the potential outcomes A, B, and C as choice-worthy. So it's rationally permitted both to accept the trade and to turn it down. If you accept the trade at node 2, you end up with B and avoid exploitation. If you turn it down, you will accept the trade at node 3, since you only deem C as choice-worthy among the potential outcomes A and C. And then you end up with C and, likewise, avoid exploitation.

Hence neither myopic nor naive agents are open to exploitation in the Upfront Money Pump. So the money pumps we have looked at don't work for all myopic, naive, and minimally sophisticated agents with cyclic preferences.

Yet there is a universal money-pump that works for myopic, naive, and minimally sophisticated agents. Consider the following decision

[21] Gustafsson and Rabinowicz 2020, p. 585, which adapts the argument from Broome and Rabinowicz 1999, pp. 240–2. See also Rabinowicz 1998, pp. 108–9 and Aumann 1998, p. 103.

problem:[22]

The Universal Money Pump

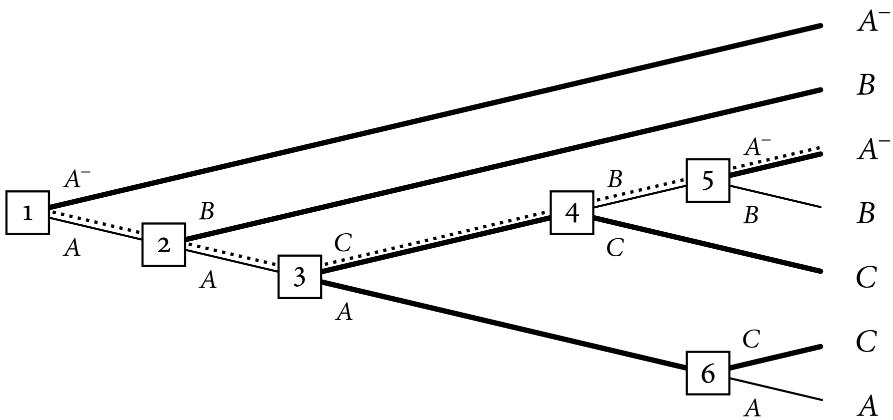

$A \succ B \succ C \succ A \succ A^{-}$, and $C \succ A^{-} \succ B$.

To see that myopic agents with the preferences in (1) and (2) get money pumped in the Universal Money Pump, suppose that you are a myopic agent with those preferences. At node 1, you turn down the trade from A to A^{-}, since you prefer A to A^{-}. At node 2, you turn down the trade from A to B, since you prefer A to B. At node 3, you accept the trade from A to C, since you prefer C to A. At node 4, you accept the trade from C to B, since you prefer B to C. Finally, at node 5, you accept the trade from B to A^{-}, since you prefer A^{-} to B. Hence you end up with A^{-} even though you could have walked away with A.

To see that naive agents with the preferences in (1) and (2) may be money pumped in the Universal Money Pump, suppose that you are a naive agent with those preferences and that you follow the Uncovered-Choice Rule. At node 1, you only deem A, B, and C as choice-worthy among the potential outcomes A, A^{-}, B, and C. So you turn down the trade at node 1. At node 2, the potential outcomes are the same, so it's rationally permitted to turn down the trade (and permitted to accept it). If you turn the trade down, you reach node 3, where the potential outcomes are still the same. So it's rationally permitted to accept the trade (and permitted to turn it down). If you accept, you reach node 4, where the potential outcomes are A^{-}, B, and C. You deem all of them as choice-worthy, so it's rationally permitted to accept the trade at node 4 (and permitted

[22] If you wonder why the trade at node 6 is needed, consult the appendix.

to turn it down). If you accept, you reach node 5, where the potential outcomes are A^- and B. Since you only deem A^- as choice-worthy, you accept the trade. Hence, by rationally permitted choices, you end up with A^- even though you could have walked away with A.

To see that sophisticated agents with the preferences in (1) and (2) get money pumped in the Universal Money Pump, suppose that you are a sophisticated agent with those preferences. At node 5, you would accept the trade from B to A^-, since you prefer A^- to B. Taking this into account, the choice at node 4 is effectively between A^- (accepting the trade) and C (turning it down). Since you prefer C to A^-, you would turn down the trade at node 4. Note, next, that you would accept the trade from A to C at node 6, since you prefer C to A. Taking these predictions into account, your choice at node 3 is effectively a choice between C (accepting the trade) and C (turning it down). So, if you were to reach node 3, you would end up with C. Taking this into account, the choice at node 2 is effectively between B (accepting the trade) and C (turning it down). Since you prefer B to C, you would accept the trade at node 2. Finally, taking this prediction into account, your choice at node 1 is effectively between A^- (accepting the trade) and B (turning it down). Since you prefer A^- to B, you accept the initial trade and end up with A^- even though you could have walked away with A.

To see that minimally sophisticated agents with the preferences in (1) and (2) get money pumped in the Universal Money Pump, suppose that you are a minimally sophisticated agent with those preferences. As before, we note that, being a minimally sophisticated agent, you choose using backward induction based on what you predict you will choose in the future assuming that, at nodes that can be reached without making any irrational choices, you retain (i) your rationality and (ii) your trust in your rationality at nodes that can be reached without making any irrational choices.

Firstly, we assume (for proof by contradiction) that node 5 can be reached without making any irrational choices. Then, at each choice node, you retain your rationality and your trust in your rationality at these nodes. Hence you would accept the trade from B to A^- at node 5, since you prefer A^- to B. But then the choice to accept the trade at node 4 was irrational, since it is effectively a choice of A^- over C, even though you prefer C to A^-. So node 5 can only be reached by making some irrational choices, which contradicts our assumption.

Secondly, we assume (for proof by contradiction) that nodes 4 and 6 can be reached without making any irrational choices. Then — at nodes 1, 2, 3, 4, and 6 — you retain your rationality and your trust in your rationality at these nodes. And, since we have already shown that node 5 cannot be reached without making irrational choices, we find that it must be irrational to accept the trade at node 4. Hence you would turn down the trade at node 4. And, since you prefer *C* to *A*, you would accept the trade at node 6. Then the choice at node 3 is effectively a choice where you end up with *C* regardless of whether you accept or turn down the trade. But then the choice to turn down the trade at node 2 was irrational, since it is effectively a choice of *C* over *B*, even though you prefer *B* to *C*. So nodes 4 and 6 can only be reached by making some irrational choices, which contradicts our assumption.

Thirdly, we assume (for proof by contradiction) that node 4 can be reached without making any irrational choices. Then — at nodes 1, 2, 3, and 4 — you retain your rationality and your trust in your rationality at these nodes. And, since we have shown that node 5 can't be reached without making irrational choices, it must be irrational to accept the trade at node 4. So you would turn down the trade at node 4. Since we have shown that nodes 4 and 6 cannot both be reached without making any irrational choices, we find that 6 can't be reached without making any irrational choices. Hence it's irrational to turn down the trade at node 3, and so you would accept that trade. But then the choice to turn down the trade at node 2 was irrational, since it is effectively a choice of *C* over *B*, even though you prefer *B* to *C*. So node 4 can only be reached by making some irrational choices, which contradicts our assumption.

Fourthly, we assume (for proof by contradiction) that node 3 can be reached without making any irrational choices. Then — at nodes 1, 2, and 3 — you retain your rationality and your trust in your rationality at these nodes. And, since we've shown that node 4 can't be reached without making irrational choices, we find that it must be irrational to accept the trade at node 3. So you would turn down the trade at node 3. Hence node 6 can be reached without making any irrational choices. So, at node 6, you retain your rationality. And so you would accept the trade from *A* to *C* at node 6, since you prefer *C* to *A*. But then the choice to turn down the trade at node 2 was irrational, since it is effectively a choice of *C* over *B*, even though you prefer *B* to *C*. So node 3 can only be reached by making some irrational choices, which contradicts our assumption.

Finally, we assume (for proof by contradiction) that node 2 can be reached without making any irrational choices. Then, at nodes 1 and 2, you retain your rationality and your trust in your rationality at these nodes. And, since we've shown that node 3 can't be reached without making irrational choices, it must be irrational to turn down the trade at node 2. Hence you would accept the trade at node 2. But then the choice to turn down the trade at node 1 was irrational, since it is effectively a choice of B over A^-, even though you prefer A^- to B. So node 2 can only be reached by making some irrational choices, which contradicts our assumption.

Hence node 2 can't be reached without making an irrational choice. So, at node 1, it is irrational to turn down the trade. So you go up at node 1 and end up with A^-, even though you could have walked away with A.

Hence we have a money pump that works for agents who violate Three-Step Acyclicity regardless of whether they are myopic, naive, or minimally sophisticated.

An advantage of a universal money pump that works regardless of whether the agent is myopic, naive, or minimally sophisticated is that it requires less knowledge on the part of the exploiter. In order for the Standard Money Pump to work, the exploiter needs to know, assuming that the agent is rational, that either myopic choice or naive choice is rationally required. And, in order for the Upfront Money Pump to work, the exploiter needs to know — again, assuming that the agent is rational — that minimal sophistication is rationally required. For the Universal Money Pump to work, the exploiter merely needs to know, assuming that the agent is rational, that it is rationally required that the agent is myopic, naive, or minimally sophisticated. This is a notably weaker assumption; it may hold even if it is not rational rationally required to be myopic, nor rationally required to be naive, nor rational required to be minimally sophisticated.[23]

[23] Is there room for a still more universal money pump? The Universal Money Pump doesn't work against self-regulating agents. A *self-regulating* agent avoids, if it can be avoided, choosing options that may be followed by a rationally permitted sequence of choices that has a prospect that the agent would not have chosen in a direct choice between the prospects of all available plans. (See Ahmed 2017, p. 1001.) As Ahmed (2017, pp. 1002–4) shows, self-regulating agents are not vulnerable to money pumps — assuming that there are no chance events involved. While there are money pumps that work against self-regulating agents, they involve chance events. And, due to these chance events, those money pumps depend on some additional requirements of rationality for choice under risk. (See Gustafsson 2022, pp. 18–19.) Hence a more universal money pump that also works against self-regulating agents would need a set-up with chance events, which would then need to rely on additional requirements of rationality

A notable limitation of the Universal Money Pump, however, is that, unlike the Standard Money Pump and the Upfront Money Pump for myopic and sophisticated agents, there does not appear to be any obvious way to extend it to cover preference cycles over more than three outcomes — while retaining its universality. So it can't be used to defend Acyclicity in general for myopic, naive, and minimally sophisticated agents; it only works for Three-Step Acyclicity. But the standard money-pump argument for Transitivity (that is, the principle that, if X is at least as preferred as Y and Y is at least as preferred as Z, then X is at least as preferred as Z) relies on the money-pump argument for Three-Step Acyclicity, rather than the one for Acyclicity.[24]

Appendix: A tempting simplification

You may wonder why we need the trade at node 6 in the Universal Money Pump. That is, you may wonder why we don't use the following, simpler decision problem:

The Semi-Universal Money Pump

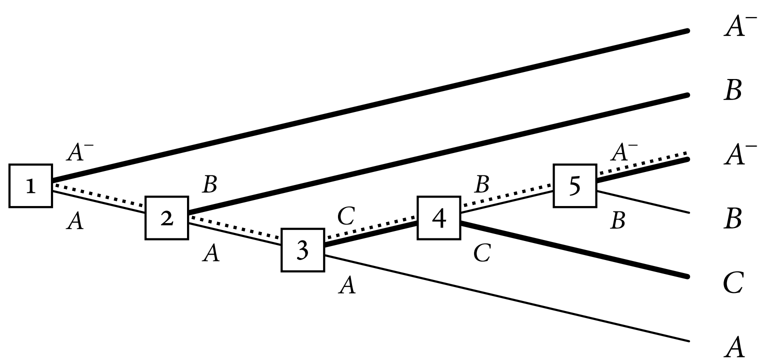

$A \succ B \succ C \succ A \succ A^-$, and $C \succ A^- \succ B$.

This money pump works equally well against myopic, naive, and sophisticated agents. But it doesn't work for minimally sophisticated agents.[25]

---

for choice under risk. But the need for additional requirements conflicts with the motivation for universal money pumps, which is to get by with weaker assumptions than the Standard and the Upfront Money Pump (which don't need any assumptions about choice under risk).

[24] See Gustafsson 2022, pp. 40–42.

[25] Likewise, Rabinowicz's (2000, p. 141) money pump with repeated offers works for

To see this, suppose that you're a minimally sophisticated agent who predicts that you would accept the trade at node 5 and also that you would irrationally accept the trade at node 4 (highlighted by the thick dashed line in the following decision tree):

The Semi-Universal Money Pump
(blocked by a predicted irrational choice)

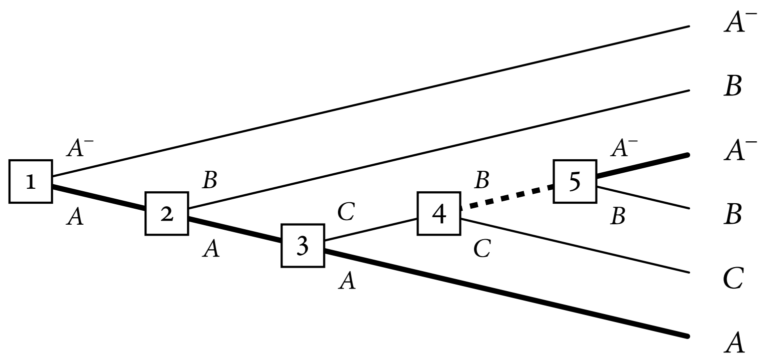

$A \succ B \succ C \succ A \succ A^-$, and $C \succ A^- \succ B$.

Given the predicted choices at nodes 4 and 5, it's irrational to accept the trade at node 3, since it would effectively be to choose A^- over A. So the predicted irrational choice at node 4 would follow an irrational choice at node 3. Hence these predictions are consisted with your being a minimally sophisticated agent. So minimally sophisticated agents with the preferences in (1) and (2) can consistently turn down all trades in the Semi-Universal Money Pump. But they can't block the Universal Money Pump in this manner.

myopic, naive, and sophisticated agents but not for minimally sophisticated agents. See Gustafsson and Rabinowicz 2020, p. 582 and Gustafsson 2022, pp. 10–11.

References

Ahmed, Arif (2017) 'Exploiting Cyclic Preference', *Mind* 126 (504): 975–1022.

Armstrong, W. E. (1939) 'The Determinateness of the Utility Function', *The Economic Journal*, pp. 453–467.

Aumann, Robert J. (1998) 'On the Centipede Game', *Games and Economic Behavior* 23 (1): 97–105.

Bicchieri, Christina (1988) 'Backward Induction without Common Knowledge', *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association* 1988 (2): 329–343.

Binmore, Ken (1987) 'Modeling Rational Players: Part I', *Economics and Philosophy* 3 (2): 179–214.

Black, Duncan (1948) 'On the Rationale of Group Decision-Making', *Journal of Political Economy* 56 (1): 23–24.

Broome, John and Wlodek Rabinowicz (1999) 'Backwards Induction in the Centipede Game', *Analysis* 59 (4): 237–242.

Condorcet, Marquis de (1785) *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*, Paris: Imprimerie royale.

— (1976) 'Essay on the Application of Mathematics to the Theory of Decision-Making', in Keith Michael Baker, ed., *Selected Writings*, pp. 33–70, Indianapolis: Bobbs-Merrill.

Davidson, Donald, J. C. C. McKinsey, and Patrick Suppes (1955) 'Outlines of a Formal Theory of Value, I', *Philosophy of Science* 22 (2): 140–160.

Dodgson, Charles Lutwidge (1876) *A Method of Taking Votes on More than Two Issues*, Oxford: Clarendon Press.

Dow, Gregory K. (1984) 'Myopia, Amnesia, and Consistent Intertemporal Choice', *Mathematical Social Sciences* 8 (2): 95–109.

Dummett, Michael (1984) *Voting Procedures*, Oxford: Clarendon Press.

Edwards, Ward, Harold Lindman, and Lawrence D. Phillips (1965) 'Emerging Technologies for Making Decisions', in Theodore M. Newcomb, ed., *New Directions in Psychology II*, pp. 259–325, New York: Holt, Rinehart and Winston.

Gustafsson, Johan E. (2020) 'Permissibility Is the Only Feasible Deontic Primitive', *Philosophical Perspectives* 34 (1): 117–133.

— (2022) *Money-Pump Arguments*, Cambridge: Cambridge University Press.

Gustafsson, Johan E. and Nicolas Espinoza (2010) 'Conflicting Reasons

in the Small-Improvement Argument', *The Philosophical Quarterly* 60 (241): 754–763.
Gustafsson, Johan E. and Wlodek Rabinowicz (2020) 'A Simpler, More Compelling Money Pump with Foresight', *The Journal of Philosophy* 117 (10): 578–589.
Hammond, Peter J. (1976) 'Changing Tastes and Coherent Dynamic Choice', *The Review of Economic Studies* 43 (1): 159–173.
McClennen, Edward F. (1990) *Rationality and Dynamic Choice: Foundational Explorations*, Cambridge: Cambridge University Press.
McMahan, Jeff (1981) 'Problems of Population Theory', *Ethics* 92 (1): 96–127.
Miller, Nicholas R. (1980) 'A New Solution Set for Tournaments and Majority Voting: Further Graph-Theoretical Approaches to the Theory of Voting', *American Journal of Political Science* 24 (1): 68–96.
Moulin, H. (1983) *The Strategy of Social Choice*, Amsterdam: North-Holland.
Parfit, Derek (1982) 'Future Generations: Further Problems', *Philosophy & Public Affairs* 11 (2): 113–172.
Pettit, Philip and Robert Sugden (1989) 'The Backward Induction Paradox', *The Journal of Philosophy* 86 (4): 169–182.
Pollak, R. A. (1968) 'Consistent Planning', *The Review of Economic Studies* 35 (2): 201–208.
Pratt, John W., Howard Raiffa, and Robert Schlaifer (1965) *Introduction to Statistical Decision Theory*, New York: McGraw-Hill, preliminary edn.
Rabinowicz, Wlodek (1995) 'To Have One's Cake and Eat It, Too: Sequential Choice and Expected-Utility Violations', *The Journal of Philosophy* 92 (11): 586–620.
— (1998) 'Grappling with the Centipede: Defence of Backward Induction for BI-Terminating Games', *Economics and Philosophy* 14 (1): 95–126.
— (2000) 'Money Pump with Foresight', in Michael J. Almeida, ed., *Imperceptible Harms and Benefits*, pp. 123–154, Dordrecht: Kluwer.
— (2008a) 'Pragmatic Arguments for Rationality Constraints', in Maria Carla Galavotti, Roberto Scazzieri, and Patrick Suppes, eds., *Reasoning, Rationality, and Probability*, pp. 139–163, Stanford, CA: CSLI.
— (2008b) 'Value Relations', *Theoria* 74 (1): 18–49.
Raiffa, Howard and Robert Schlaifer (1961) *Applied Statistical Decision Theory*, Boston: Harvard Business School.
Samuelson, Paul Anthony (1947) *Foundations of Economic Analysis*, Cam-

bridge, MA: Harvard University Press.
Schwartz, Thomas (1990) 'Cyclic Tournaments and Cooperative Majority Voting: A Solution', *Social Choice and Welfare* 7 (1): 19–29.
Sen, Amartya K. (1977) 'Social Choice Theory: A Re-Examination', *Econometrica* 45 (1): 53–89.
Shepard, Roger N. (1964) 'Circularity in Judgements of Relative Pitch', *The Journal of the Acoustical Society of America* 36 (12): 2346–2353.
Strotz, R. H. (1955–1956) 'Myopia and Inconsistency in Dynamic Utility Maximization', *The Review of Economic Studies* 23 (3): 165–180.
Uzawa, Hirofumi (1956) 'Note on Preference and Axioms of Choice', *Annals of the Institute of Statistical Mathematics* 8 (1): 35–40.
von Neumann, John and Oskar Morgenstern (1944) *Theory of Games and Economic Behavior*, Princeton: Princeton University Press.