A Forcing Money Pump for Naive Agents[⋆]

Johan E. Gustafsson[†]

ABSTRACT. Suppose you prefer prospect A to prospect B, B to prospect C, and C to A. The standard argument that such cyclic preferences are irrational is the money-pump argument, which shows that, if you have those preferences, you are open to exploitation. The standard version of that argument relies on an exploitation scheme that is a forcing money pump for myopic agents (who are rationally required to go along with each step of the scheme) but only a permitting money pump for naive agents (who are merely rationally permitted to go along with each step of the scheme). So it may seem that there is no forcing money pump for naive agents. In this paper, however, I put forward a forcing money pump for naive agents.

Is it irrational to have cyclic preferences? For instance, suppose there are prospects A, B, and C such that you prefer A to B, B to C, and C to A. Are these preferences irrational? There is a standard argument that such cyclic preferences are irrational. This argument, the money-pump argument, goes like this.

Suppose you start off possessing A (that is, if you were to turn down all trades, you would end up with A). Now, an exploiter offers a trade from A to C. Since you prefer C to A, you are rationally required to accept this trade. Next, the exploiter offers a trade from C to B. Since you prefer B to C, you are rationally required to accept this trade too. Finally, the exploiter offers a trade from B to A^-, where A^- is A after making a small payment.

If you prefer X to Y, then you prefer X to Y with some margin so that your preference is robust to some sourings of X and some sweetenings of Y. A *souring* of a prospect X is a prospect just like X except that it is certainly inferior in one dimension that the agent cares about. And a *sweetening* of a prospect X is a prospect just like X except that it is certainly superior in one dimension that the agent cares about. (We can let the dimension be money, but this isn't needed for the argument.) So

⋆ Forthcoming in *Philosophical Studies*.

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

the exploiter's offer of A given a small payment can be thought of as a souring of A. That souring offered by the exploiter, A^-, is a sufficiently small souring of A which you still prefer to B (and each of A and C is preferred to A^-). Accordingly, since you prefer A^- to B, you are rationally required to accept this last trade. But now you have been money pumped — you ended up paying for what you could have kept for free.[1]

Since you were rationally required to accept each trade, this is a forcing money pump. A *forcing* money pump is a money pump where the agent is rationally required, at each step, to go along with the exploitation.[2]

But the argument assumed *myopic choice* — that is, it assumed that you are rationally required to accept a trade if it is preferred to your current position, not taking into account future trades.[3]

A rival approach is naive choice. According to *naive choice*, rational agents consider, at each choice node, the prospects of all available plans and assess which of these prospects are choice-worthy in a choice between all of them and then choose in accordance with a plan to end up with one of the choice-worthy prospects.[4] Notably, each plan is considered without taking into consideration whether the agents would later depart from it.

To see how naive choice works and how it handles the standard money pump, we diagram the decision problem as follows:

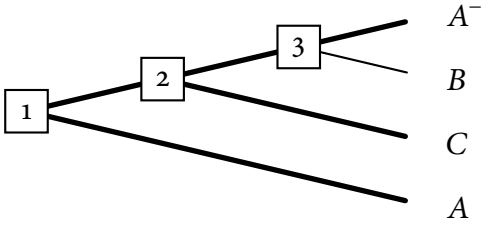

Here, the squares represent the choice nodes at which the trades are of-

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

[2] Gustafsson 2022, p. 27.

[3] Following Dow 1984, p. 96 and McClennen 1990, pp. 11–12.

[4] Pollak 1968, pp. 202–3 and Hammond 1976, p. 162. The analogous form of consequentialism that prescribes following the plan that leads to an outcome that is no worse than the outcome of any other available plan is known as *possibilism*. (See Feldman 1975, p. 260 and Jackson and Pargetter 1986, p. 233.)

fered. And you accept a trade at a choice node by going up and you turn it down by going down.

At node 1, naive choice requires that you consider which among the prospects of all available plans are choice-worthy. That is, you need to consider which of A, A^-, B, and C are choice-worthy in a choice between them. Yet, since each prospect is less preferred than one of the others, no option is rationally permitted according to the Maximization Rule:[5]

> *The Maximization Rule* It is rationally permitted to choose a prospect X if and only if there is no available prospect Y such that Y is preferred to X.

So, to guarantee that there will be at least one rationally permitted option, we can rely on the following rule:[6]

> *The Uncovered-Choice Rule* It is rationally permitted to choose a prospect X if and only if there is no available prospect Y such that Y is preferred to X and, for all available prospects Z, it holds that Y is preferred to Z if X is preferred to Z.

The idea here is to avoid ruling out all options. So we don't want to rule out an option as irrational just because it is less preferred than some other option. An option X that is less preferred than an option Y can still have an advantage over Y if X is preferred to some option that Y is not preferred to. The idea is to rule out the options that are less preferred than some other option without having this kind of advantage over that option. These ruled-out options are dominated in terms of preference comparisons by the more preferred option. Every option that is not dominated in this way but is still less preferred than some other option has the advantage over that option by being preferred to some further option while the other option is not. So we count these options as rationally permitted along with options that are not less preferred than any option.

Given your preferences, the Uncovered-Choice Rule entails that A, B, and C are all choice-worthy. But A^- is not choice-worthy, since A is preferred to A^- and, in addition, A is preferred to each option that A^- is preferred to (that is, B).

So, at node 1, it's rationally permitted to go up and also rationally permitted to go down. At node 2, the prospects of the available plans are A^-,

[5] Uzawa 1956, p. 37.

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

B, and *C*, which are all choice-worthy according to the Uncovered-Choice Rule. So, at node 2, it's likewise rationally permitted to go up and rationally permitted to go down. Finally, at node 3, you are rationally required to go up, since you prefer A^- to *B*. (The rationally permitted choices according to naive choice, combined with the Uncovered-Choice Rule, are marked with thicker lines in the diagram.)

Accordingly, it's rationally permitted to go along with each step of the exploitation scheme. But it's not rationally required to do so — you are also rationally permitted to turn down the offers at nodes 1 and 2. So, rather than a forcing money pump, this is just a permitting money pump. While the standard money-pump argument is forcing given myopic choice, it is not forcing given naive choice.

A problem with this permitting money pump is that it is not a reliable exploitation scheme. While you may end up with A^-, you may also end up with *A* or *C*. And, if you end up with *C*, the exploiter has merely traded you *C* for *A*, which needn't be to their advantage (and you prefer it to your initial possession).[7] A *permitting* money pump is a money pump where, at each step, the agent is rationally permitted to go along with the exploitation.[8]

There could be permitting money pumps that are still risk free. To state what a risk-free money pump is, let a *non-prohibiting* money pump be a money pump where, at each step, the agent is not rationally prohibited from going along with the exploitation.[9] (A non-prohibiting money pump can fail to be a permitting money pump in case, at some step, it's neither rationally permitted nor rationally prohibited to go along with the exploitation — that is, rationality would be silent on whether to take that step.) A *risk-free* money pump is a non-prohibiting money pump where the agent, without making any irrational choice at any step, can only end up with either the prospect of what they would get if they turned down all

[7] See Ahmed 2017, p. 1011 and Gustafsson 2022, pp. 16–17. This problem also affects Cantwell's (2002, pp. 328–9; 2003, pp. 388–9) three-way money pump in which naive agents are rationally required to end up paying for something they could have had for free. But they are not rationally required to end up paying for a specific alternative. So the three-way money pump does not work as a reliable exploitation scheme. The agents, starting off possessing *A*, may end up trading *A* for C^-, which they prefer to *A*. And then the exploiter has not pumped the agents for money, but rather just sold them *C* in exchange for *A*.

[8] Gustafsson 2022, p. 27.

[9] Gustafsson 2022, p. 27.

offers or a souring of that prospect.[10] A risk-free permitting money pump would be a plausible exploitation scheme even if it isn't forcing, since the exploiter either suffers no loss or makes a gain. Given naive choice, however, the standard money pump is not risk free.[11]

The standard money-pump argument does not work at all given sophisticated choice. *Sophisticated choice* requires rational agents to use foresight and take into account what they predict will be chosen at future choice nodes when they make their current choices. Such agents predict that they would go up at node 3. So the choice at node 2 is then effectively between A^- (going up) and C (going down). Since they (plausibly) prefer C not only to A but also to A^-, they would go down at node 2. And then the exploitation is blocked. Nevertheless, there are forcing money pumps that do work against sophisticated agents.[12]

Hence it may seem that there are forcing money pumps for myopic and sophisticated agents, but not for naive agents.[13] But, as we shall see, there is a forcing money pump for naive agents.

Let A^-, B^-, and C^- be sourings of A, B, and C respectively. And let A^+, B^+, and C^+ be sweetenings of A, B, and C respectively. These sourings and sweetenings are soured and sweetened by a sufficiently small amount so that (i) each one of A, A^-, and A^+ is preferred to each one of B, B^-, and B^+, (ii) each one of B, B^-, and B^+ is preferred to each one of C, C^-, and C^+, and (iii) each one of C, C^-, and C^+ is preferred to each one of A, A^-, and A^+. And any prospect is preferred to any souring of that prospect

[10] This is much the same as McGee's (1999, p. 257n2) notion of a gentle Dutch book.

[11] One might wonder why we care about the perspective of the exploiter rather than that of the exploitee. The idea is that a rational agent should not give an exploiter an opportunity to get money from them (without taking any risks or knowing something the agent doesn't know). The requirement that the exploitation scheme has to be risk free for the exploiter makes it harder to construct a suitable money pump. Without this requirement, the three-way money pump works against all non-resolute agents. But the three-way money pump is not risk free. As mentioned in note 7, the exploiter risks merely trading one option for another, which need not be to their advantage or to the exploitee's disadvantage.

[12] Rabinowicz 2000, p. 141 and Gustafsson and Rabinowicz 2020, p. 583.

[13] There is also a forcing money pump against self-regulating agents. *Self-regulating choice* rationally forbids, if it can be avoided, choosing options that can be followed by a sequence of rationally permitted choices that has a prospect that you would not have chosen in a direct choice between the prospects of all available plans. (Ahmed 2017, p. 1001. See also Cubitt and Sugden 2001, p. 143.) The forcing money pump against self-regulating agents relies on a unidimensional version of stochastic dominance. See Gustafsson 2022, pp. 18–19.

(and to any souring of that souring and so on).

Now, consider the following decision problem:

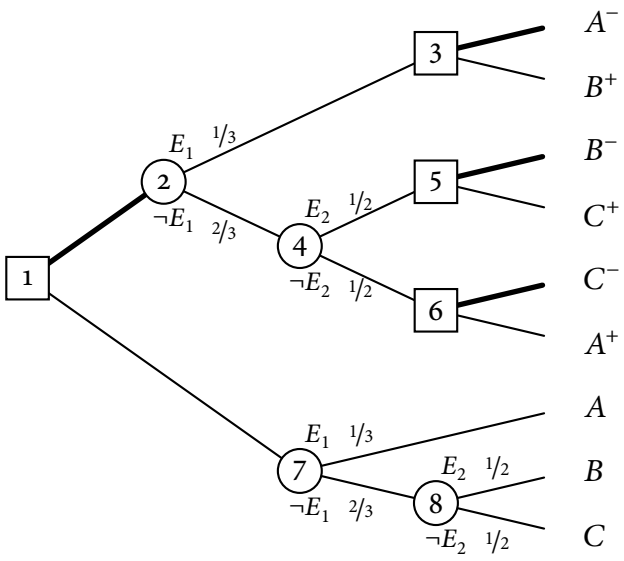

Here, the circles represent chance nodes. At nodes 2 and 7, chance goes up if and only if chance event E_1 occurs. And, at nodes 4 and 8, chance goes up if and only if chance event E_2 occurs. The probabilities of these events are 1/3 for E_1 and 1/2 for E_2, and the events are independent of each other.

Let '$\langle X, Y, Z\rangle$' denote the gamble of getting X if E_1 occurs, getting Y if E_2 occurs but E_1 does not occur, and getting Z if neither E_1 nor E_2 occurs. You start off possessing $\langle A, B, C\rangle$. Then, at node 1, you are offered a trade from $\langle A, B, C\rangle$ to $\langle B^+, C^+, A^+\rangle$. At node 3, you are offered a trade from B^+ to A^-. At node 5, you are offered a trade from C^+ to B^-. And, at node 6, you are offered a trade from A^+ to C^-.

The argument needs to assume two requirements of rationality governing gambles. The idea behind the first is that, if you prefer X to Y, then that preference should be robust to sweetenings of X. It would be weird if a sweetening of X made it less preferred than Y. Plausibly, that sweetening should be more preferred than X. So we adopt the following requirement of rationality:[14]

[14] Note that a sweetening is *certainly* superior in a dimension the agent cares about. So Sweetening Dominance is much weaker than the following more general principle (von Neumann and Morgenstern 1944, pp. 26–7):

> *Sweetening Dominance* If gamble X^+ is a sweetening of gamble X and X is preferred to gamble Y, then X^+ is preferred to X and to Y.

Here, a sweetening of a gamble is defined analogously to a sweetening of a prospect: a *sweetening* of a gamble X is a gamble just like X except that it is certainly superior in one dimension that the agent cares about.

By Sweetening Dominance, we find that

(1) $\langle A^+, B^+, C^+\rangle$ is preferred to $\langle A, B, C\rangle$ and to any gamble that $\langle A, B, C\rangle$ is preferred to.

The idea behind the second requirement is that it shouldn't matter in what state of nature an outcome occurs. (Note that A, B, C and their sweetenings need not be final outcomes; an outcome, here, may be a risky prospect, rather than a final, non-risky outcome.) If two gambles are stochastically equivalent (in the sense that, for each possible final outcome, they give you the same chance of that outcome), then they should be preferentially fungible. That is, if one of them is preferred to another gamble, then the other should be so too. Hence we adopt the following requirement of rationality:[15]

> *Transitivity of Strict Preference* If prospect X is preferred to prospect Y and Y is preferred to prospect Z, then X is preferred to Z.

In Sweetening Dominance, X^+ is just like X but certainly superior in one dimension the agent cares about, so — given that X is preferred to Y — there is no reason for X^+ to do worse than X in comparison with Y. And note that, unlike the Weak Principle of Eventwise Dominance (see note 15), Sweetening Dominance allows that neither X nor X^+ is preferred to Y even if both X and X^+ eventwise dominate Y.

[15] It may be objected that Stochastic Fungibility conflicts with

> *The Weak Principle of Eventwise Dominance* If there is a set of events such that (i) the set is a partition of states of nature and (ii) given each event E in the set, the outcome of gamble X given E is preferred to the outcome of gamble Y given E, then X is preferred to Y.

See Arrow 1965, p. 17. The Weak Principle of Eventwise Dominance entails that $\langle A, B, C\rangle$ is preferred to $\langle B, C, A\rangle$. In this manner, as noted in Fishburn 1988, pp. 188–9 (see also Baumann 2022, pp. 15–16 and Williamson 2024, p. 368n20), Stochastic Fungibility and the Weak Principle of Eventwise Dominance jointly rule out cyclic preferences. So, given cyclic preferences, we can't have both principles. Of course, if we accept both principles as requirements of rationality, we have a direct argument that rational preferences are acyclic. (The main benefit of the money-pump argument over this alternative is precisely that we don't need to rely on the Weak Principle of Eventwise Dominance.) A

> *Stochastic Fungibility* If, for each possible final outcome, the probability of that outcome given gamble X is the same as the probability of that outcome given gamble Y, then, for any gamble Z, it holds that X is preferred to Z if and only if Y is preferred to Z.

By Stochastic Fungibility, we find that

(2) $\langle B^+, C^+, A^+\rangle$ is preferred to the same gambles as $\langle A^+, B^+, C^+\rangle$.

And, from (1) and (2), we find that

(3) $\langle B^+, C^+, A^+\rangle$ is preferred to $\langle A, B, C\rangle$ and to any gamble that $\langle A, B, C\rangle$ is preferred to.

Note that $\langle B^+, C^+, A^+\rangle$ is the prospect of an available plan at node 1 (going up at node 1 and down at nodes 3, 5, and 6) and $\langle A, B, C\rangle$ is the prospect of going down at node 1. Given (3), the Uncovered-Choice Rule entails that it's not rationally permitted to go down at node 1. The Uncovered-Choice Rule entails that some option must be rationally permitted.[16] So it entails

notable theory that violates Stochastic Fungibility is skew-symmetric-additive utility theory (see Loomes and Sugden 1982, p. 809 and Fishburn 1988, pp. 197–8; 1989), which adds up pairwise strengths of preference in each state of nature weighted by probability. One argument in favour of the stochastic principle over the eventwise one is the strangeness of caring about in what states of nature the outcomes happen (Bader 2018, p. 500). Hare (2010, pp. 239–41), however, presents a case that throws doubt on stochastic reasoning when some outcomes X and Y are related by preferential gaps — that is, neither X nor Y is at least as preferred as the other. Since we haven't ruled out preferential gaps, this is a worry. But note that none of the outcomes of $\langle B^+, C^+, A^+\rangle$ and $\langle A^+, B^+, C^+\rangle$ (the only gambles to which we apply Stochastic Fungibility) are related by a preferential gap. One notable argument for the eventwise approach that applies in the cyclic case is the Argument from Full Information — roughly: if you knew which event obtains, then you would prefer the prospect that eventwise dominates another; so you should defer to your better-informed self (Hare 2010, p. 242) or recognize the primacy of outcomes (Schoenfield 2014, p. 268) and so prefer the eventwise-dominating prospect. For a response to this argument, see Francis and Gustafsson 2026, pp. 301–7. Francis and Gustafsson's (2026, pp. 307–10) positive argument for stochastic reasoning, however, is not available here, since it depends on transitivity.

[16] For proof, let's say that option X *covers* option Y if and only if X is preferred to Y and to any available option that Y is preferred to. Note that the covering relation is transitive. So, given a finite number of available options, there will be at least some available option that is uncovered. And all such options are rationally permitted according to the Uncovered-Choice Rule. See Miller 1980, p. 75.

that it's rationally permitted to go up at node 1. Accordingly, given naive choice, you are rationally required to go up at node 1.[17]

Suppose that you, following that prescription, do go up at node 1. Then, if E_1 occurs, you are rationally required to go up at node 3, since you prefer A^- to B^+. If E_1 does not occur and E_2 occurs, you are rationally required to go up at node 5, since you prefer B^- to C^+. And, if neither E_1 nor E_2 occurs, you are rationally required to go up at node 6, since you prefer C^- to A^+.

Hence you are rationally required to go up at each choice node, which results in the gamble $\langle A^-, B^-, C^- \rangle$, whereas, if you had gone down at node 1, you would have ended up with $\langle A, B, C \rangle$. Note that $\langle A^-, B^-, C^- \rangle$ is just like $\langle A, B, C \rangle$ except that it is inferior in every state of nature in some dimension you care about. So the prospect of going up at each choice node is a souring of the prospect of going down at node 1. So, even though you could have walked away from all offers and kept $\langle A, B, C \rangle$ for free, you end up paying for it. And, at each step, you were rationally required to go along with the exploitation scheme. Consequently, we have a forcing money-pump argument against naive agents with cyclic preferences.

I wish to thank Wlodek Rabinowicz and two anonymous referees for valuable comments.

[17] This follows on either the duality definition of required to ϕ as not permitted not to ϕ (von Wright 1951a, p. 4; 1951b, p. 37) or the definition of required to ϕ as permitted to ϕ and not permitted not to ϕ (Gustafsson 2020, p. 123).

References

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

Arrow, Kenneth J. (1965) *Aspects of the Theory of Risk-Bearing*, Helsinki: Yrjö Jahnssonin Säätiö.

Bader, Ralf M. (2018) 'Stochastic Dominance and Opaque Sweetening', *Australasian Journal of Philosophy* 96 (3): 498–507.

Baumann, Peter (2022) 'Rational Intransitive Preferences', *Politics, Philosophy & Economics* 21 (1): 3–28.

Cantwell, John (2002) 'The Pragmatic Stance: Whither Dutch Books and Money Pumps?', *Croatian Journal of Philosophy* 2 (6): 319–336.

— (2003) 'On the Foundations of Pragmatic Arguments', *The Journal of Philosophy* 100 (8): 383–402.

Cubitt, Robin P. and Robert Sugden (2001) 'On Money Pumps', *Games and Economic Behavior* 37 (1): 121–160.

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.

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

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.

Feldman, Fred (1975) 'World Utilitarianism', in Keith Lehrer, ed., *Analysis and Metaphysics*, pp. 255–271, Dordrecht: Reidel.

Fishburn, Peter C. (1988) *Nonlinear Preference and Utility Theory*, Baltimore: The Johns Hopkins University Press.

— (1989) 'Non-Transitive Measurable Utility for Decision under Uncertainty', *Journal of Mathematical Economics* 18 (2): 187–207.

Francis, Tomi and Johan E. Gustafsson (2026) 'Stochastic Dominance for Incomplete Preferences', *Analysis* 86 (2): 299–311.

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 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.
Hare, Caspar (2010) 'Take the Sugar', *Analysis* 70 (2): 237–247.
Jackson, Frank and Robert Pargetter (1986) 'Oughts, Options, and Actualism', *The Philosophical Review* 95 (2): 233–255.
Loomes, Graham and Robert Sugden (1982) 'Regret Theory: An Alternative Theory of Rational Choice under Uncertainty', *The Economic Journal* 92 (368): 805–824.
McClennen, Edward F. (1990) *Rationality and Dynamic Choice: Foundational Explorations*, Cambridge: Cambridge University Press.
McGee, Vann (1999) 'An Airtight Dutch Book', *Analysis* 59 (4): 257–265.
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.
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 (2000) 'Money Pump with Foresight', in Michael J. Almeida, ed., *Imperceptible Harms and Benefits*, pp. 123–154, Dordrecht: Kluwer.
Schoenfield, Miriam (2014) 'Decision Making in the Face of Parity', *Philosophical Perspectives* 28 (1): 263–277.
Schwartz, Thomas (1990) 'Cyclic Tournaments and Cooperative Majority Voting: A Solution', *Social Choice and Welfare* 7 (1): 19–29.
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.
von Wright, Georg Henrik (1951a) 'Deontic Logic', *Mind* 60 (237): 1–15.
— (1951b) *An Essay in Modal Logic*, Amsterdam: North-Holland.
Williamson, Timothy Luke (2024) 'A Risky Challenge for Intransitive Preferences', *Noûs* 58 (2): 360–385.