

Newcomb dilemma in development management

Newcomb dilemmas show a discrepancy in our rational reasoning, as made clear by comparing Evidential Decision Theory with Causal Decision Theory. In this paper, I look at three versions of the dilemma: the original, highly technical and abstract one plus two more mundane cases. I also account for the general schema of the dilemma possibly appearing in macroeconomic situations. Ahmed (2014) aims to provide a solution for macroeconomic cases that opens room for forming a development management Newcomb dilemma – an imaginary case of electric motor competition between Toyota and Tesla. I argue that Ahmed’s solution may solve the macroeconomic Newcomb dilemma, but it cannot be applied to the development management dilemma. If I am right, similar Newcomb situations could be cropping up regularly in development management, leading to seemingly insoluble strategic decisions having to be made. This may create an inevitable pitfall for development management.

Keywords: *decision theory, development management, Newcomb dilemma, macroeconomics, causal decision theory, evidential decision theory*

Author Information

Zsolt Ziegler, Budapest University of Technology and Economics and Eötvös Loránd University
<https://orcid.org/0000-0001-9222-2265>

How to cite this article:

Ziegler, Zsolt. “Newcomb dilemma in development management.”
Információs Társadalom XX, no. 4 (2020): 7–18.
<https://dx.doi.org/10.22503/inftars.XX.2020.4.1>

*All materials
published in this journal are licenced
as CC-by-nc-nd 4.0*

1. Introduction

In this paper, I argue that in development management, developers may easily face with a — Newcomb — dilemma (Nozick 1969), known in decision theory, bringing about insoluble strategic decisions. In this situation, developers are not able to make an ideal decisions in principle even though their cognitive resources are not bounded on any scale. I argue that the so-called Newcomb Dilemma might regularly occur in development management.

First, I describe the original Newcomb dilemma and show a genuine discrepancy in our rational thought, according to contradictory decision theories. Then I give an account of the more mundane case of a Newcomb dilemma found in macroeconomics. It turns out that the dilemma is not that abstract and can occur in many walks of life. I consider a possible solution to the Newcomb dilemma provided by Ahmed (2014). However, although it may solve the dilemma found in macroeconomics, I argue that Ahmed’s solution cannot be applied to developmental management Newcomb dilemmas. Finally, I provide a somewhat fictional case of Toyota and Tesla in which a Newcomb situation renders Toyota unable to decide whether to develop electric cars or not.

2. The Newcomb dilemma

An ideally rational agent¹ is supposed to choose between taking (and gaining the contents of) (i) an opaque box that is now in front of her or (ii) that same opaque box *and* a transparent box holding \$1000. Yesterday, a machine that has an excellent track record—let’s say 99% right—of predicting agents’ decisions predicted about one’s decision. If the machine made a prediction about the agent that the agent would take only the opaque box (‘one-boxing’), the machine put \$1 000 000 in the opaque box yesterday. The machine did not place anything in the opaque box if it saw that one would take both (‘two-boxing’). The matrix summarises the possibilities of the agent found below:

| | The machine predicts one-boxing | The machine predicts two-boxing |
|------------|---------------------------------|---------------------------------|
| One-boxing | \$1 000 000 | \$0 |
| Two-boxing | \$1 001 000 | \$1000 |

¹ The dilemma is that even an ideally rational agent cannot make the ideal decision owing to the structure of the Newcomb situation.

We have two theories determining two differing decisions in this situation. According to Evidential Decision Theory (EDT), ‘the rational act is whichever available one is the best evidence of what you want to happen’ (Ahmed 2018, 8). So, if the agent acts in accordance with EDT, the agent believes that their act is evidentially relevant to the state that they desire.

□ : The most reasonable decision is to choose the one-boxing strategy.

To see the argument according to EDT, for the sake of simplicity, let us go back to our 99% accurate predicting machine. To gain the maximum payoff, the agent might reason thus: If it is true that the machine is 99% right and 1% wrong, then taking one box has the expected utility 990 000. In this case, the agent is thinking that if they take one box that the machine yesterday predicted with 99% accuracy and the utility value is \$1 000 000, then $0.99 * 1\ 000\ 000 = 990\ 000$. However, if the agent takes two boxes while the machine has predicted one-boxing, the agent must also suppose that there is a 1% chance of machine error. According to this latter scenario, where the expected utility is \$1 001 000, the resulting expected utility is 10 100 ($0.01 * 1\ 001\ 000$). Since the expected utility of the one-boxing strategy is 990 000 but that of the two-boxing strategy is 10 100, the agent must choose one-boxing over two according to EDT. The fact that the agent knows the machine’s predictive power to be 99% provides the best evidence for them to make up their mind.

According to EDT, the expected utility table² of the standard Newcomb dilemma is as follows:

| | The machine predicts one-boxing | The machine predicts two-boxing |
|------------|--|--|
| One-boxing | Exp. Ut.: 990 000 ($0.99 * 1\ 000\ 000$) | Exp. Ut.: 0 ($0.01 * 0$) |
| Two-boxing | Exp. Ut.: 10 100 ($0.01 * 1\ 001\ 000$) | Exp Ut.: 990 ($0.99 * 1\ 000$) |

Causal Decision Theory (CDT), however, suggests that ‘the rational act is whichever available one is most likely to cause what you want to happen’ (Ahmed 2018, 8). So, if another person behaves according to CDT, the agent holds that that person’s actions must have a causal influence on the state that the agent wants.

□□ : The most reasonable decision is to choose the two-boxing strategy.

² Note that the case where the machine predicts two-boxing with 99% prediction accuracy and the agent takes both boxes yields an expected utility of 990. No ideally rational agent would therefore rely on this option as it provides the least expected utility.

CDT determines two boxing according to the following reasoning: whatever the agent is about to choose, the machine has already placed (or hasn't) \$1 000 000 into the opaque box. If you like, *the die is cast*. The prediction of the machine has nothing to do with the decision the agent is about to make. Consequently, the agent faces only two options. First, if the agent follows the one boxing strategy, then she either gets \$1 000 000 or nothing. Second, if the agent acts upon the two boxing strategy, she may gain \$1 001 000 or \$1000. Since the agent's actual choice does not influence the content of the opaque box now, the only reasonable decision is to take both boxes.

The following chart summarises the agent's options concerning whether the machine did or did not place a million dollars into the opaque box. (Note that dashed-line boxes represent the transparent boxes, and black boxes illustrate the agent's actual choice, while grey boxes show what the agent did not pick).

| | The machine placed one million dollars yesterday | The machine did not place one million dollars yesterday |
|------------|--|---|
| One-boxing | | |
| Two-boxing | | |

Importantly, CDT makes use of the principle of causal independence: correlation does not imply causal dependence. To see this with an example: the forecasts of meteorologists today do not cause the weather tomorrow. Meteorologists make predictions based on independent facts that will cause tomorrow's weather. The correlation is established by a common cause agent, namely certain atmospheric conditions. The same is true for the Newcomb dilemma: the machine's prediction does not cause the agent's decision at all. Similarly to weather, the agent's choice is based on a causally independent (earlier) state of the world.

CDT and EDT do not agree over cases where an agent's acts are evidence for states that they do not causally promote, and this is precisely the situation in Newcomb's problem. One-boxing is evidence that you will get \$1 000 000 because it is evidence of the state in which you were predicted to one-box; EDT, therefore, recommends one-boxing. Two-boxing brings it about that you are \$1000 richer than you would otherwise have been; CDT, therefore, recommends two-boxing. Note, though, that EDT and CDT do share some similarities: both theories of rationality aim to maximise expected utility since ideal

agents want to gain the maximum benefit by choosing one (\$1 000 000) or two (\$1 001 000) boxes.

It is worth noting that, according to Skalse, accepting either EDT or CDT sets certain epistemic conditions determining what the agent is supposed to do in the Newcomb Dilemma. “This means that they would be in different epistemic states when they make their decisions, and hence not be facing the “same” decision problem.” (Skalse 2021, 4)

Before we proceed further, we need to look at the general structure of Newcomb dilemmas. There are always two roles in the schema: an expector and a decision-maker. As the following table shows, the expected utility must always be (ii) > (i) > (iv) > (iii) according to the two-by-two options of the decision-maker’s choice and the expector’s prediction.³

| | Expector predicts non-X | Expector predicts X |
|--------------------------|-------------------------|---------------------|
| Decision-maker non-X-ing | i | iii |
| Decision-maker X-ing | ii | iv |

Note also that the probabilities of the expected utilities are $p\theta(i) > p\lambda(ii)$ and $p\theta(iv) > p\lambda(iii)$ and also $p\theta + p\lambda = p1$. Importantly, the expector’s expectation does not depend causally upon what the decision-maker is about to choose because it is always the case that the expector predicts in advance of the decision-maker.

3. A Newcomb dilemma in macroeconomics

Although the Newcomb dilemma might seem quite abstract, there have been many real-life Newcomb situations. Broome (1990) presented a version of a Newcomb situation that seems to apply to macroeconomics. Let us, then, suppose that the Open Market Committee of the Federal Reserve in the United States is trying to decide whether to expand the money supply or not. The standard theory of macroeconomics teaches that increasing the money supply fosters employment, plus, as a result of the increased amount of money on the market, banks do not have to reserve a huge sum of funds against deposits but can instead provide retail and business credits. The committee is facing a dilemma owing to the strong probabilistic interdependence between the money supply and the public’s expectations of the money supply. From monetary records, it is known that the public can predict pretty accurately – say with 70%

³ Note that since David Lewis there has been vivid discussion as to whether Newcomb’s problem is really two versions of the Prisoner’s Dilemma (see Lewis 1979, 1981; Bermúdez 2013, 2015; Walker 2014, 2015; Weber 2016; Binmore 2021).

precision – what the committee chooses to do. Broome (1990, 488) describes this as follows:

If the government expands the money supply, the people will probably have predicted that, so the result will be inflation. If it does not expand it, they will probably have predicted that too, so the result will be no change. (*status quo* — the author added) The Bolker-Jeffrey theory [i.e., EDT], then, will assign a higher expected utility to not expanding. It suggests that this is the right thing to do. Dominance reasoning, however, shows that the right thing is to expand. That, at any rate, is the conclusion of most authors who have considered this ‘time-inconsistency problem’. The government’s dilemma has exactly the form of the ‘Newcomb Problem’, which first led to the interest in causal decision theory.

Note further that if the public gets surprised because the committee does not increase the money supply, then the result will be a recession. But if the committee manages to surprise the public by increasing the money supply when the public thought it would remain constant, then increased employment will be the most likely outcome.

Let us summarise this in a table. Note that the central bank’s subjective expected values are added to the possible outcomes.

| | Public expects no expansion | Public expects expansion |
|--------------|---|--|
| No expansion | Status quo (9) Exp. Ut.: 6.3 (0.7*9) | Recession (0) Exp. Ut.: 0 (0.3*0) |
| Expansion | Increased employment (10) Exp. Ut.: 3 (0.3 * 10) | Inflation (1) Exp. Ut.: 0.7 (0.7 * 1) |

As before, we consider what EDT suggests in this situation, which is not to expand the money supply. To see the supporting argument, let us suppose that the public can predict the bank’s monetary strategic moves with 70% precision, meaning that they mispredict 30% of the time. Similarly to the original Newcomb dilemma, given the above-presented subjective utility values (‘0’, ‘1’, ‘9’, ‘10’), the committee needs to reason thus: If the *status quo* scenario happens, then the utility value is 9 resulting in the expected utility 6.3 given that the public’s predictive ability is 70% (0.7 * 9). If, however, the central bank decides to expand the money supply while the public predicted the opposite, then the committee needs to assume that the public will err with regard to its strategic moves, which has a 30% chance. According to this scenario, the utility value is 10, resulting in the expected utility of 3 (0.3 * 10). The eviden-

tial principle suggests that *a rational agent does what constitutes their best evidence that they will realise their aims*; therefore, the central bank needs not to expand the money supply.

On the other hand, the central bank can reason based on CDT, concluding that expanding the money supply is the correct decision. The die is cast, the bank may presume, and market participants have already made up their minds as to whether to borrow money to start and expand a business. Accordingly, any decision the committee is about to make will not influence in any respect the public's strategic moves. Therefore, similarly to the meteorologist's forecast and the weather today, the prediction of the market participants has nothing to do with what decision the central bank should make. The public is aware of that, the committee has to take into account two options. If the central bank decides against expanding the money supply, then either *status quo* (9) or recession (0) will happen. If, however, the committee chooses to expand the money supply, then the US economy will either enjoy *increased employment* (10) or face inflation (1). According to CDT, the central bank needs to choose the dominant decision by expanding the money supply.

4. Ahmed's reply to the bank's Newcomb situation

Ahmed (2014) argues that the central bank's Newcomb situation can be solved, and that CDT leads to the right decision, namely to expand. First, to see the argument, we shall rank our previous possible outcomes – (increased employment) > (*status quo*) > (*inflation*) > (*recession*) – and note that the public's expectation does not depend causally upon what the committee is about to choose, not least since market participants act in advance of the central bank. This provides a dominance argument in favour of expanding the money supply. Now, we have a Newcomb dilemma iff. predictions (as to whether the central bank will expand or not) of the market participants are probabilistically dependent upon the committee decision, that is, whether there is a solid probabilistic interdependence between what the central bank chooses to do and what the public expects. To generate a Newcomb dilemma, we need to assume – in accordance with Broome (1990), Bermúdez (2018), Ahmed (2018) – that $p(\textit{status quo}) > p(\textit{increased employment})$ and that $p(\textit{inflation}) > p(\textit{recession})$.

However, if it can be proven that $p(\textit{increased employment}) = p(\textit{inflation})$ and $p(\textit{status quo}) = p(\textit{recession})$, then the dilemma fails and CDT determines what to do. Ahmed thinks that in a Newcomb situation, the two probabilities (either increased employment and inflation or *status quo* and recession) are the same. He argues that a predictor's (in our case, the public's) evidence-base (ψ) is a subset of the decision-maker's (here, the committee's) evidence-base (ϕ). Namely, $\psi \subseteq \phi$. Decision-makers fully access to their set of pieces of evidence so, we can assume that $p(\phi) = 1$. Therefore, $p(\phi) > p(\psi)$. Now consider this:

$p(\text{Doing } X^{\text{decision-maker}} / \psi \ \& \ \text{Predicting non-}X^{\text{expector}})$ (*increased employment*)
 $p(\text{Doing } X^{\text{decision-maker}} / \psi \ \& \ \text{Predicting } X^{\text{expector}})$ (*inflation*)

and the same is true for the other decision:

$p(\text{Not Doing } X^{\text{decision-maker}} / \psi \ \& \ \text{Predicting non-}X^{\text{expector}})$ (*status quo*)
 $p(\text{Not Doing } X^{\text{decision-maker}} / \psi \ \& \ \text{Predicting } X^{\text{expector}})$ (*recession*)

According to Ahmed, two pairs of probabilities (*increased employment & inflation* and *status quo & recession*) must be the same — very significantly from the subjective perspective of the decision-maker—, if Doing X holds is fully determined by ψ . A rational-decision maker, nevertheless, needs to hold its own actions to be evidentially irrelevant to how the expector forms its beliefs, since its predictions are formed on the basis of ψ (which is only a subset of ϕ). Accordingly, no expector can have more precise information about the decision maker’s choice than the decision maker’s access to its own actions. It results that the probability of expectation (from the expector) must always be lower than what the decision-maker is about to do ($p(\phi) = 1$). If this is so, there is no point to decide in favour of (i.), and in a somewhat real life Newcomb Dilemma, it is always recommended to choose the dominant strategy according to CDT.

5. Newcomb dilemmas in development management

Let us imagine a somewhat fictional case where the chief executive board of Toyota Group is about to decide whether or not to change its research and development (R&D) direction from hybrid cars to electric ones (turning some of its production over to electric cars in the hopes of dominating this new market, which is possible owing to Toyota’s market-leading position in the automotive industry).

Being among the first to enter an emerging market would bring obvious benefits to Toyota, such as enjoying the positive effects of the learning curve, getting to occupy the market segment, creating the impression in customers that the brand in question is the original one (Cohen 2005, 57). However, the chief executive board faces a dilemma since there seems to be a solid probabilistic interdependence between the development of electric cars and the competitors’ expectations of electric car development.

Let us assume a fictional case where the competitors’ expectations have a long and successful record of predicting what developments Toyota is about to make. If Toyota starts developing electric cars, and this is exactly what the competitors have predicted, then the result is going to be only a *slight increase in sales*. This is because Toyota will be able to keep up with the changing competition by utilising its market-leading position and developmental and infrastructural resources, while competitors will also try to occupy this segment.

If, however, Toyota sticks to developing hybrid cars (not developing electric cars), and it is what competitors have predicted, then the *status quo* is the most likely outcome. The research having been conducted to further fine-grain hybrid engines will pay off, stabilising Toyota’s position in the market – at least for a while. Competitors will not need to worry about Toyota’s entering the electric car industry, so the market competition in this particular field will not get enhanced.

However, if Toyota *surprises* its competitors by not getting into electric motor development, Toyota will face a *recession* to a certain degree. This case, where Toyota’s competitors develop electric cars while Toyota does not, will result in Toyota losing its market-leading position while its competitors will gain it.

Finally, the best result – gaining a *market-leading* position in the electric car segment – will come about only if Toyota can surprise its competitors by developing electric cars when no one thought it would. This case seems to be the most straightforward. In this case, Toyota will be able to further dominate the automotive industry because competitors will lag.

| | Competitors predict Toyota’s not developing electric cars | Competitors predict Toyota’s developing electric cars |
|-------------------------|---|---|
| Toyota does not develop | <i>status quo</i> (9) Exp. Ut.: 6.3 (0.7*9) | <i>recession</i> (0) Exp. Ut.: 0 (0.3*0) |
| Toyota develops | <i>market leading</i> (10) Exp. Ut.: 3 (0.3*10) | <i>slight increase in sales</i> (1) Exp. Ut.: 0.7 (0.7*1) |

Toyota’s possible outcomes in this imaginary situation thus sketch a Newcomb-like case because it seems that no matter how Toyota decides, it violates either EDT or CDT.

According to EDT, Toyota is recommended not to develop electric cars. To see why, let us suppose that its competitors can predict Toyota’s developmental strategic moves with 70% precision, and they fail to do so 30% of the time. Given the presented utility values in the table (‘0’, ‘1’, ‘9’, ‘10’), the chief executive board of Toyota Group needs to reason accordingly: If the *status quo* is the case, then the utility value is 9, implying expected utility of 6.3 given Toyota’s competitors’ predictive ability of 70% (0.7 * 9). If Toyota chooses to develop while rivals predict the opposite, then Toyota’s board needs to assume that its rivals will mispredict that it has only a 30% chance. According to this scenario, the utility value is 10, resulting in expected utility of 3 (0.3 * 10). The evidential principle suggests that a *rational agent does what constitutes their best evidence that they will realise their aims*; thus, Toyota needs not to develop electric cars since *the status quo* scenario results in a higher expected utility (6.3) than *the market-leading* case (3).

Alternatively, Toyota's board can apply CDT to decide whether to develop electric cars. The die is cast, and competitors have already made up their minds about what they will do in the electric motor industry. Accordingly, whatever decision Toyota wants to make will have no influence in any respect on its competitors' strategic moves. Therefore, the competitors' prediction – again, just like the meteorologist's forecast and the weather today – has nothing to do with what decision Toyota should make. Competitors are aware only that Toyota's board has to choose between two options: If it decides against developing electric cars, then either *status quo* (9) or *recession* (0) will happen. If it chooses to develop electric cars, it will either grow to dominate the electric car segment, that is, become *market-leading* (10), or it will experience *a slight increase in sales* (1). According to CDT, Toyota needs to choose the dominant decision by developing electric cars.

If this is right, we have found a *development management Newcomb dilemma*.

6. A reply to Ahmed's analysis

If Ahmed is right, the Newcomb dilemma no longer holds in real-life Newcomb situations. Ahmed's analysis might be true for the imaginary choice of the Open Market Committee of the Federal Reserve in the United States, but it does not work for other Newcomb situations. I argue that the analysis cannot account for the presented development management Newcomb dilemma. However, it might be true that the central bank identifies those particular elements of ϕ that account for the set of propositions that completely characterises the expector's evidence-base but it is certainly not true for the competitor's evidence-base.

Let us suppose that one of Toyota's competitors, Tesla (imaginary), entered the R&D field of electric cars earlier and has already tramped over the road and learnt some of the main lessons. Knowing the pitfalls of this research field makes Tesla's evidence-base more extended. Now, let us call 'α' the set of propositions that completely characterises Tesla's evidence-base and let 'β' denote the set of propositions that entirely characterises Toyota's evidence-base. Therefore, we can assume that $\beta \subset \alpha$ which means that every element of β is in α but that α has more. It also allows that although Tesla has gained broader relevant experience, Toyota may have somewhat different approaches. However, in a case like this where an expector (Tesla) has a broader set of propositions, it makes Tesla's ability to predict Toyota's behavior more accurate. Therefore, $p(\alpha) > p(\beta)$. Now consider the following, where 'EC' stands for electric cars.

$p(\text{Develop EC}^{\text{decision-maker}} / \alpha \ \& \ \text{Expecting Not Developing EC}^{\text{expector}})$ *(market leading)*

$p(\text{Develop EC}^{\text{decision-maker}} / \alpha \ \& \ \text{Expecting Developing EC}^{\text{expector}})$ (*slight increase in sales*)

and the same is true for the other decision:

$p(\text{Not Develop EC}^{\text{decision-maker}} / \alpha \ \& \ \text{Expecting Not Developing EC}^{\text{expector}})$ (*status quo*)

$p(\text{Not Develop EC}^{\text{decision-maker}} / \alpha \ \& \ \text{Expecting Developing EC}^{\text{expector}})$ (*recession*)

Given that Tesla knows more, the two pairs of probabilities (*market-leading & slight increase in sales* and *status quo & recession*) cannot be the same (from the decision-maker's subjective perspective), if developing electric cars holds is mostly determined by α . Toyota – assuming that Tesla entered the electric car R&D earlier and has gained broader experience – needs to consider what its competitor predicts since Tesla's expectations are formed based on α (when β is a subset of α). This time the expector (Tesla) has broader information about the decision-maker's (Toyota's) set of propositions grounding its choice, meaning that the probability of expectation (from the expector) must always be higher than what the decision-maker is about to do ($p(\alpha) > p(\beta)$). If I am right, and the two pairs of probabilities (*market-leading & slight increase in sales* and *status quo & recession*) are different, then our imaginary Toyota's board still faces a Newcomb dilemma.

7. Conclusion

Newcomb dilemmas shed light on a discrepancy between the two approaches of our rational reasoning – EDT and CDT. We have examined three versions of the dilemma: the original, highly technical and abstract one plus two more mundane cases of it. It turned out that the general schema of the dilemma may appear in macroeconomic states of affairs, representing real-life Newcomb dilemmas. You might think that even the more everyday versions of the dilemma are too far removed from fully realistic decision situations. I disagree. Even though a clear Newcomb schema is pretty unlikely to occur, the phenomenon of a market participant having broader knowledge of a particular field, making them able to predict what their competitors are about to do, is rather probable. It is also possible that the competitors are very well aware that the other market participant has this special knowledge. If I am right, similar Newcomb situations might be cropping up regularly in development management, leading to seemingly impossible strategic decisions having to be made as to whether to follow EDT or CDT. This may turn out to be an inevitable pitfall of development management.

References

- Ahmed, Arif, ed. "Evidence, Decision and Causality." In Evidence, Decision and Causality, i–ii. Cambridge: Cambridge University Press, 2014. <https://www.cambridge.org/core/books/evidence-decision-and-causality/evidence-decision-and-causality/D58D014B0AE-E742946D0265914047645>.
- Ahmed, Arif, ed. *Newcomb's Problem*. Classic Philosophical Arguments. Cambridge: Cambridge University Press, 2018. <https://doi.org/10.1017/9781316847893>.
- Bermúdez, José Luis. "Prisoner's Dilemma and Newcomb's Problem: Why Lewis's Argument Fails." *Analysis* 73, no. 3 (July 1, 2013): 423–29. <https://doi.org/10.1093/analys/ant034>.
- Bermúdez, José Luis. "Strategic vs. Parametric Choice in Newcomb's Problem and the Prisoner's Dilemma: Reply to Walker." *Philosophia* 43, no. 3 (September 1, 2015): 787–94. <https://doi.org/10.1007/s11406-015-9606-6>.
- Bermúdez, José Luis. "Does *Newcomb's Problem* Actually Exist?" In *Newcomb's Problem*, edited by Arif Ahmed, 19–41. Classic Philosophical Arguments. Cambridge: Cambridge University Press, 2018. <https://doi.org/10.1017/9781316847893.002>.
- Binmore, Kenneth. "Robert Nozick Versus David Lewis." In *Imaginary Philosophical Dialogues: Between Sages down the Ages*, edited by Kenneth Binmore, 181–87. Cham: Springer International Publishing, 2021. https://doi.org/10.1007/978-3-030-65387-3_31.
- Broome, John. "Bolker-Jeffrey Expected Utility Theory and Axiomatic Utilitarianism." *The Review of Economic Studies* 57, no. 3 (July 1, 1990): 477–502. <https://doi.org/10.2307/2298025>.
- Cohen, William A. *The Marketing Plan*. 5th edition. Hoboken, NJ: Wiley, 2005.
- Lewis, David. "Prisoners' Dilemma Is a Newcomb Problem." *Philosophy & Public Affairs* 8, no. 3 (1979): 235–40.
- Lewis, David. "Why Ain'cha Rich?." *Noûs* 15, no. 3 (1981): 377–80. <https://doi.org/10.2307/2215439>.
- Nozick, Robert. "Newcomb's Problem and Two Principles of Choice." In *Essays in Honor of Carl G. Hempel: A Tribute on the Occasion of His Sixty-Fifth Birthday*, edited by Nicholas Rescher, 114–46. Synthese Library. Dordrecht: Springer Netherlands, 1969. https://doi.org/10.1007/978-94-017-1466-2_7.
- Skalse, Joar. "A General Counterexample to Any Decision Theory and Some Responses." *ArXiv E-Prints* 2101 (January 1, 2021): arXiv:2101.00280.
- Walker, Mark. "Rejoinder to Bermúdez on Lewis, Newcomb's Problem and the Prisoner's Dilemma." *Philosophia* 43 (July 20, 2015). <https://doi.org/10.1007/s11406-015-9633-3>.
- Walker, Mark Thomas. "The Real Reason Why the Prisoner's Dilemma Is Not a Newcomb Problem." *Philosophia* 42, no. 3 (September 1, 2014): 841–59. <https://doi.org/10.1007/s11406-014-9516-z>.
- Weber, Thomas A. "A Robust Resolution of Newcomb's Paradox." *Theory and Decision* 81, no. 3 (September 1, 2016): 339–56. <https://doi.org/10.1007/s11238-016-9543-2>.