

# Structural Decision Theory

Tung-Ying Wu\*†

---

Judging an act's causal efficacy plays a crucial role in causal decision theory. A recent development appeals to the causal modeling framework with an emphasis on the analysis of intervention based on the causal Bayes net for clarifying what causally depends on our acts. However, few writers have focused on exploring the usefulness of extending structural causal models to decision problems that are not ideal for intervention analysis. The thesis concludes that structural models provide a more general framework for rational decision makers.

---

**1. Introduction.** Decision theories concern an agent's rational choice in a decision problem, where the agent faces different acts to choose from but is uncertain about each act's possible consequences. Suppose she knows about the possible consequences of her different acts, the utility of each consequence, and the probability of each consequence. Then she can acquire the expected utility of each act by multiplying the probability and the utility of each possible consequence of an act and then adding the results of all possible consequences of the act. Philosophers in decision theory contend that a rational choice for an agent is an option that maximizes expected utility.

Causal decision theory (CDT) endorses the principle of expected utility maximization but holds that the agent must take the causal relevance of her acts to their outcomes into consideration. Proponents of CDT share the belief that rational agents should maximize expected utility based on the causal information relevant to their acts but differ in what approach best captures an act's causal efficacy (Lewis 1981, 11; Joyce 1999, 146; Ahmed 2014, 8–9; Weirich 2016).

Interventionist decision theory (IDT) is a form of CDT because IDT also holds that the relevant information that matters to our decision should be

\*To contact the author, please write to: Institute of Philosophy, Chinese Academy of Sciences, Beijing, China; e-mail: tungyingwu@outlook.com.

†I would like to thank Chuang Liu and Malcolm Forster for helpful suggestions and comments.

causal, but IDT approaches an act's causal efficacy through intervention analysis within the framework of causal modeling.<sup>1</sup> IDT holds that an agent should conceive of an act as an intervention that disables all preexisting causes of the act in a decision problem (Meek and Glymour 1994, 1007–8; Pearl 2009, 70, 108–12; Hitchcock 2016, 1158–59; Stern 2017, 4139–42; 2019, 784–85).<sup>2</sup>

More formally, intervention analysis is assessed by the theory of causal Bayes nets. Variables (denoted by uppercase letters) represent tokens of events that serve as relata of (type level) causal relations, and these variables range over possible values (denoted by lowercase letters) that represent these events' occurrence or nonoccurrence or a value if an event is of a quantity. A Bayesian causal model  $M$  is a triple  $\langle G, V, P \rangle$ , where  $V$  is a set that contains variables whose causal relationships we are interested in studying,  $P$  is the probability distribution of each variable, and  $G$  is a directed acyclic graph. Graph  $G$  consists of nodes that represent variables in  $M$  and arrows between nodes that represent causal relations. If the value of a variable  $Y$  depends on  $X$ , then there will be a directed path from  $X$  to  $Y$ . Probability  $P$  satisfies the causal Markov condition if and only if each variable  $X_i$  in  $V$  is independent of all other variables except  $X_i$ 's descendent given  $X_i$ 's parent  $PA_i$ , where " $X_i$ 's descendent" stands for the other variables in  $V$  that are causally downstream from  $X_i$  and " $X_i$ 's parents" stand for  $X_i$ 's immediate causes. More specifically,  $P$  satisfies the causal Markov condition if and only if the following condition holds:  $P(X_1, \dots, X_n) = \prod_i P(X_i | PA(X_i))$ , where  $X_1, \dots, X_n$  are all variables in  $V$ , and  $PA(X_i)$  stands for "parents of  $X_i$ ." An intervention on  $X_j$  removes all its preexisting cause and sets it to a specific value. Hence, the intervention analysis is done by removing  $P(X_j | PA(X_j))$  from the above joint distribution. This amounts to setting  $X_j$  to a specific value and making it no longer dependent on its original parents. Hence, the effect of intervention on  $X_j$  is obtained by the new joint distribution:  $P'(X_1, \dots, X_n) = \prod_{i \neq j} P(X_i | PA(X_i))$ .

Causal models, together with intervention analysis, represent the causal details relevant to a decision-making context in a rigorous mathematical language. Hence, when engaging with a decision problem, one should use causal models to clarify one's assumptions about the causal structure of the problem, the information that one has available, and the question one is asking. More

1. An intervention  $I$  as an external force sets  $X$  to certain values, and  $I$  neither causes any variable other than  $X$  nor is caused by any other variable in a causal model (Spirtes, Glymour, and Scheines 2000, 47–53; Pearl 2009, 23–24, 70–74).

2. Meek and Glymour (1994) claim that we may conceive of our acts as interventions only when we believe that our actions are not caused by circumstances beyond our control (see Hitchcock 2016, 1166; Stern 2019, 789–90). Note that the notion of "intervention" in this article is not the same as Woodward's (2005, 94–98). In this article, "intervention analysis" is understood in terms of manipulating the probability distribution in a causal model in which the causal Markov condition holds. See below.

importantly, by making use of causal models, one can distinguish causation from correlation (Meek and Glymour 1994; Pearl 2009, sec. 4.1; Hitchcock 2016, 1175; Stern 2017, 4147).

IDT instructs rational agents to choose an act  $x$  that maximizes the interventionist expected utility (IEU). Let  $Y$  be a random variable that ranges over possible outcomes,  $P$  be a rational agent  $s$ 's subjective probability function,  $do(X = x)$  be  $s$ 's intervention to make  $s$  do  $x$ ,  $V(Y = y)$  be the utility of an outcome  $y$ , and  $IEU(x)$  be the interventionist expected utility of act  $x$ . Here is Pearl's (2009, 108) definition of IEU:  $IEU(x) = \text{df} \sum_y P(Y = y) | do(X = x) V(Y = y)$ .<sup>3</sup> This definition asserts that  $s$  should assess the expected utility of an outcome  $y$  on the basis of evaluating the effect of the intervention to make  $s$  do  $x$ .

Nevertheless, Pearl (2017, 2021) recently proposes a new definition of expected utility in terms of structural causal models (SCMs) as decision-making conditionals. Call the definition of expected utility with an application of SCMs the structural expected utility (SEU):  $SEU(x) = \text{df} \sum_y P(Y_x = y) V(Y = y)$ .<sup>4</sup> Pearl entitles  $P(Y_x = y)$  as a SCM-defined counterfactual (2021, 2–6). This definition declares that  $s$  should evaluate the expected utility of act  $x$  by using a SCM analysis of causality.

IEU and SEU are methodologically different approaches. They instruct the agent to use different procedures in evaluating the causal information of decision problems. For instance, IEU tells the agent to obtain the probability distribution and the corresponding causal graph of each variable in a decision problem.<sup>5</sup> In contrast, SEU requires delineating functional relations between relevant variables to attain the causal structure. They are nevertheless different methodologies for the agent to approach decision problems.

This article attempts to assess the scope of SEU and IEU and their effectiveness in making explicit the causal structure of decision problems. Previous work has only focused on IEU's implications for some controversial examples in CDT, such as Newcomb's Problem and the Psychopath Button, or issues of uncertainty about causal dependency (Meek and Glymour 1994, 1008–9; Hitchcock 2016, 1165–69; Stern 2017, 4142; 2019, 797–98). To the best of my knowledge, the distinction between IEU and SEU has not been

3. Pearl uses the do-operator to denote "intervention." For similar proposals, see Meek and Glymour (1994, 1009–10) and Hitchcock (2016, 1162–64).

4. Note that Pearl (2009, 108) originally endorsed IEU. Also, Pearl sometimes uses  $P(Y = y | do(X = x))$  and  $P(Y_x = y)$  interchangeably in his writings because the latter can be translated and computed by the former under several strong assumptions. Such translation would fail in some examples (see Pearl 2009, 245–47, 289–93; Pearl, Glymour, and Jewell 2016, 107–16).

5. See the prior discussion of intervention analysis assessed by causal Bayes nets.

dealt with in depth. The example in the next section demonstrates it is SEU, rather than IEU, that serves as a valuable formal tool for a range of realistic decision problems that involve mixed mechanisms. Therefore, SEU provides a more general framework for rational decision makers.

The remainder of the article is organized as follows. Section 2 presents the example of the spinner and explains why IEU fails to deliver an intuitive result. Section 3 gives a brief overview of SCMs. Section 4 employs SCMs to analyze the spinner and shows how SCMs and SEU, but not intervention analysis and IEU, deliver an intuitive result.

**2. The Spinner.** An agent has a chance to win a prize (called the reward). There is a spinner (see fig. 1) and an arrow in the circle that may be spun on its dial to indicate a result. The agent may choose between two options “SAFE” and “ADD- $X$ .” If the agent plays SAFE, the agent flicks the arrow and gains the value where the arrow stops. Since 40% of the time the arrow stops in area  $Z = 1$ , 20% in area  $Z = 2$ , and 40% in area  $Z = 3$ , the expected average gain for the agent is 2 units of money. In contrast, option ADD- $X$  allows the agent to increase the reward by  $X$  unit(s) of money for a small cost (much smaller than  $X$ ) with the following rule: if the arrow stops in area  $Z = 1$ ,  $Z$  will not be contributive, and the reward will have only  $X$  unit(s) of money. If the arrow stops in area  $Z = 3$ ,  $Z$  will be contributive, so the reward will have  $3 + X$  units. However, if the arrow stops in area  $Z = 2$ ,  $Z$  will be deleterious, so the reward will have  $X - 2$  units.

Now, assessing the expected gain of option ADD- $X$  is a complicated task.<sup>6</sup> The spinner is a mixture of areas of  $Z$  that react differently to the agent’s choosing ADD- $X$ . For example,  $Z$  is contributive to the reward in  $Z = 3$ , not contributive to the reward in area  $Z = 1$ , and deleterious to the reward in area  $Z = 2$ . The mechanisms differ from area to area because they exhibit different dispositions that manifest given the presence of the agent’s acting on ADD- $X$ .<sup>7</sup>

Since the spinner consists of the areas with different dispositional properties, the intervention analysis has difficulty in accurately predicting the causal effect of choosing ADD- $X$ . Simply put, intervention analysis is mostly assessed

6. This example is a modified case of “additive intervention” (Pearl 2009, sec. 11.4.4; Shpitser and Pearl 2009; Shipley 2016, 9–11, 50–54). Namely, one evaluates the effect of adding some amount from  $X$  without removing a preexisting causal process of  $X$ . (In the example, I use  $X$  as an instrument variable,  $Y$  as some amount, and  $Z$  as a preexisting cause of  $Y$ .) Pearl et al. (2016, 109–11) confirm that the effect of additive intervention could not be reduced to intervention expressions alone.

7. Following Bechtel and Abrahamsen (2005, 423), I take mechanism as “a structure performing a function in virtue of its component parts, component operations, and their organization.” The mechanisms here include two parts: one is physical (the spinner), and the other is operational (allocating the reward depending on the outcome of the spin). I thank an anonymous referee for clarifying this point.




Figure 1. Spinner.

by the theory of causal Bayes net. However, the procedure of computing the effect of acting on ADD-X as an intervention amounts to computing  $P(Y = y|do(X = q))$ , which does not fix the level of  $Z$ . Since the level of  $Z$  is not fixed, we may estimate the value of  $Z$  by the expectation ( $E(Z)$ ), and  $E(Z) = 2$  in the spinner. Thus, the intervention analysis implies that the agent should predict that acting on ADD-X as an intervention will always result in the worst case scenario:  $Z$  will be deleterious, so the reward will have  $X - 2$  units.<sup>8</sup> Nevertheless, this is certainly incorrect. For only 20% of the time the value of  $Z$  is deleterious to the reward, but 80% of the time the value of  $Z$  is not deleterious to the reward. It seems that ADD-X does not always lead to the worst causal scenario of the value of  $Z$  being deleterious. Intervention analysis is limited when it is not possible for the agent to intervene on a relevant feature that has a mixture of different mechanisms.

In the spinner, the agent cannot intervene to fix the amount of the reward. For doing so is an intervention that removes the preexisting rule of the spinner, but the agent must flick the arrow, and it is not up to the agent to fix the arrow on the spinner that consists of areas that react differently to adding  $X$ .<sup>9</sup> Thus, it seems that the intervention analysis of choosing ADD-X is unfitting if the intervention analysis is insensitive to the variant causal properties across the circle that is not intervenable. Hence, the agent's intervention analysis of choosing ADD-X is inaccurate, and it remains unclear whether the agent should choose SAFE or ADD-X.

How do we evaluate the causal efficacy of an act when the world is a mixture of variant mechanisms in which the act causes different outcomes? Presumably, if the agent knows each area's mechanism, she should evaluate the causal effects of her interventions area by area. Since the issue is predicting

8. A related issue is the condition of unanimity that a cause must raise the probability of its effect in all contexts (Dupré 1984; Eells 1991, 103–4). An in-depth evaluation of this condition would go beyond the scope of this article.

9. Hence, one cannot evaluate the causal efficacy of ADD-X by the analysis of interventions  $P(Y = y|do(X = x, Z = z))$ ,  $P(Y = y|do(X + Z))$ , and  $P(Y = y|do(X - Z))$ .

the expected gain of ADD-X, the agent should average the causal effect in each area by its proportion to the whole circle to derive the desired quantity.

This article puts forward a justification for applying SCMs to evaluate an act's causal efficacy in decision theory. The last question of the above example—how we evaluate the causal efficacy of an act when the world is a mixture of variant mechanisms in which the act causes different outcomes—calls for a SCM analysis. For this purpose, the above example provides an independent reason for employing SCMs to define an act's expected utility in decision theory, namely, SEU. In what follows, I will introduce SCMs, which may be of use to a rational agent to accurately predict what is causally downstream from her acts.

**3. Structural Causal Models.** SCMs can formally represent causal relations in a rigorous mathematical language. They conveniently represent an agent's belief about causal relationships among variables of interest and the causal effect of an intervention. A prior development of SCMs includes the work of the economist Herbert A. Simon who specialized in decision making. In his influential paper, Simon (1957, 10–13) argues that we can define a causal system as some functional relationships in a structure—a specific arrangement of variables and equations in fixing the sequence of computing their solutions. I begin with a brief account of SCMs.

A SCM  $M$  consists of a quadruple  $\langle U, V, f, P \rangle$ , where  $U$  is a set of exogenous (or background) variables and  $V$  is a set of endogenous variables. Exogenous variables represent background factors in  $M$  and are only determined by factors outside the model, and their values do not depend on the other variables in the model. In contrast, endogenous variables are determined only by the other variables in the model. A set of functions,  $f$ , assigns each endogenous variable in  $V$  a value based on the values of the other variables in the model, and  $P$  represents a probability distribution over all variables in  $U$ . Specifically, each function has the form  $X_i = f_i(PA_i, U_i)$ ,  $i = 1, \dots, n$  (Simon 1957, 18–19, 40; Pearl 2009, 202–3; Pearl et al. 2016, 26–27), where  $X_i$  is an endogenous variable in  $V$ ,  $PA_i$  (which stands for “parents of  $X_i$ ”) is a set of variables in  $V$ ,  $U_i$  is an exogenous variable in  $U$ , and  $PA_i$  and  $U_i$  together determine the value of  $X_i$ .

Moreover, by assumption, each variable in  $V$  can only have one distinct equation that determines its value. Hence, each function represents an autonomous mechanism that predicts what value nature would assign to  $X_i$  in response to every possible value combination of  $(PA_i, U_i)$ . They are autonomous in the sense that one function  $f_i$  continues to hold or remains undisrupted by external changes to the other functions in  $f$ . Hence, the causal relations in  $M$  are deterministic given a value assignment of  $U_i$ . Since every  $X_i$  is (partially or wholly) determined by at least one  $U_i$  and every  $U_i$  is not determined by any  $X_i$  in  $V$ , a value assignment of all  $U_i$  in  $U$  determines a unique value

distribution over all  $X_i$  in  $V$  based on  $f$ . If  $P$  is the probability distribution over all exogenous variables, the probability distribution for the endogenous variables is also  $P$  (Simon 1957, 40–43, 54–56; Pearl 2009, 27–32, 203–6; Pearl et al. 2016, 98).<sup>10</sup>

To illustrate SCMs and an explicit representation of the causal relationships in the spinner, in the next section I use SCMs to represent the spinner and demonstrate that the agent can accurately predict what is causally downstream of her acts in SCMs' expressions.

**4. Additive Intervention.** I use the linear SCM  $M_1 \langle U, V, f, P \rangle$  to represent the causal relationships in the spinner. Let  $X, Y, Z$  be endogenous variables in  $V$ , and let  $I, U_Z$  be exogenous variables in  $U$ .<sup>11</sup> These variables range over possible values (denoted by lowercase letters). Variable  $X$  represents how much value the agent adds to the prize,  $Y$  represents the value of the reward, and  $Z$  represents the value that the arrow points to. The intervention variable  $I$  represents the agent's intervention, and it is an exogenous variable because only outside factors (e.g., the agent's free will) determine its value.

In the spinner,  $X$  can increase the prize  $Y$ , and  $Z$  also causally affects  $Y$ 's value. The following functions represent the causal relations between these variables:  $f_x: X = \{q, 0\}$ ;  $f_z: Z = U_Z$ ;  $f_y: Y = \{Z \text{ if } X = 0; X \text{ if } X > 0 \text{ and } Z < 2; X - Z \text{ if } X > 0 \text{ and } Z = 2; Z + X \text{ if } X > 0 \text{ and } Z > 2\}$ .

Exogenous variable  $U_Z$  determines the value of  $Z$ . The probability distribution of  $U_Z$  is the composition of the spinner:  $P(U_Z = 1) = 0.4$ ,  $P(U_Z = 2) = 0.2$ , and  $P(U_Z = 3) = 0.4$ . Also,  $X = q$  represents the agent's action to add  $q$  to the reward;  $X = 0$  represents the agent's action to add nothing to the reward. Next,  $f_x$  stands for the causal mechanism that specifies how the agent decides to add value to the reward: if she decides to add  $q$  amount of value,  $X$  will be set to  $q$ . If she decides to add nothing,  $X$  will be set to 0.

Function  $f_y$  stands for how  $X$  and  $Z$  determine the amount of the reward: if the agent adds no value ( $X = 0$ ), then the value of  $Y$  will equal  $Z$ . If she adds some value ( $X > 0$ ) but  $Z$  is lower than 2, then the value of  $Y$  will be  $X$ . If she adds some value but  $Z$  equals 2, then the amount of  $Y$  will be  $X - Z$ . If she adds some value, and  $Z$  is larger than 2, then the value of  $Y$  will be  $Z + X$ . This function  $f_y$  demonstrates different mechanisms in which  $Z$  reacts differently to the added value  $X$  in the process of determining the reward  $Y$ .

10. A consequence of a SCM  $M$  is that the probability distribution of every variable in  $M$  satisfies the causal Markov condition. The causal Markov condition holds in SCMs under these further assumptions: (a) there is no causal loop in  $M$ , namely, the associated causal graph is acyclic; (b) the exogenous variables in  $U$  are jointly independent; (c)  $M$  includes every variable that is a cause of two or more other variables; and (d) if any two variables are dependent, then one is a cause of the other or there is a third variable causing both (see Steel 2005, 10; Pearl 2009, 30).

11. For brevity, I omit some exogenous variables.

We turn now to the question of how an agent predicts the overall causal effects of choosing ADD-X. The diverse areas have varied types of mechanisms represented by several levels of  $Z$ . In the spinner, the circle consists of areas with three levels of  $Z$ : 40% is  $Z = 1$ , 20% is  $Z = 2$ , and 40% is  $Z = 3$ . The agent can estimate the results in each level of  $Z$  and averages these effects by the probability distribution of  $Z$ .<sup>12</sup> Specifically, the agent may use  $P(Y_x = y|Z = z)$  to represent the probability that an outcome  $y$  would obtain conditional on the action  $X = x$  in a structural model updated by  $Z = z$ .<sup>13</sup> Given a structural model  $M$  and observed information  $Z = z$ , one can evaluate the conditional  $P(Y_x = y|Z = z)$  in three steps. (1) Abduction: Conditionalize on the evidence  $z$  to determine the value of the variables in  $U$ . (2) Action: Replace the equations corresponding to variables in set  $X$  by the equation  $X = x$ . (3) Prediction: Use the modified model and the updated value of the variables in  $U$  to compute the value of  $Y$  (Galles and Pearl 1998; Pearl 2009, 37, 202–6; Pearl et al. 2016, 92–98).<sup>14</sup>

The first step uses the information  $Z = z$  about the situation to fix the values of the exogenous variables in  $U$ . In particular, each value assignment of variables in  $U$  is the defining characteristic of a single individual or situation. For example, in the model  $M_i$ , a value assignment  $U_i = u_i$  stands for the identity of the agent and the spinner. The second step stands for the minimal modification of the model  $M$  that replaces  $f_x$  with  $X = x$ . The third step predicts the value of  $Y$  on the basis of the modified  $M$  and the updated values of  $U$ .

It is now possible to answer the agent's question of assessing the SEU of choosing ADD-X by SCMs. First, she updates her value assignment of  $U$  from the supposition that  $Z = 1, 2$ , or  $3$  and identifies  $U_z$ . Next, she carries over the updated value of  $U_z$  to the model  $M_1$  modified by  $X = q$ . Finally, she predicts the value of  $Y$  by finding a solution to the following equations:  $f_x: X = q; f_z: Z = U_z; f_y: Y = \{Z \text{ if } X = 0; X \text{ if } X > 0 \text{ and } Z < 2; X - Z \text{ if } X > 0 \text{ and } Z = 2; Z + X \text{ if } X > 0 \text{ and } Z > 2\}$ .

Consequently, she can predict that had she added  $q$  unit(s) to the reward when  $Z = 1$ , the reward would be  $q$ . Equally, she can also predict that had she added  $q$  unit(s) to the reward when  $Z = 2$ , the reward would be  $q - 2$ . Had she added  $q$  unit(s) to the reward when  $Z = 3$ , the reward would be  $q + 3$ . Given that 40% of the time  $Z = 1$ , 20% of the time  $Z = 2$ , and 40% of the time  $Z = 3$ , the SEU of ADD X would be  $q + 0.8$  minus the

12. In cases in which experimental units manifest variant dispositional properties, Spirtes et al. (2000, 165–67) cite similar calculations to obtain predictions.

13. Probability  $P(Y_x = y)$  is a subjunctive conditional, which stands for the probability that, had an intervention *do* ( $X = x$ ) been performed, an outcome  $Y = y$  would obtain.

14. For simplicity, I skip some unnecessary technical details. Note that this is different from Woodward's notion of causality analyzed with counterfactual interventions.



fee that she has to pay. Recall that the expected value of the reward if the agent plays SAFE is invariably 2. Therefore, if option ADD-X allows the agent to pay less than 0.1 unit of money to add  $X > 1.3$  to the reward, she will be quite confident that option ADD-X is preferable to option SAFE.

The implication is that facilitating SCMs and deriving SEU in the spinner and similar situations is more fitting than intervention analysis. As demonstrated in the spinner, the approach of SCMs captures the mixture of variant mechanisms specified by the probability distribution of  $Z$  and the function  $f_z$ , and thereby obtains more accurate characterizations of each area's causal property and the causal efficacy of choosing ADD-X. Hence, in cases in which an agent observes different causal properties that are not intervenable across the population in the real world, the agent might more adequately make statements about her acts' causal efficacy in SCMs' mathematical terms.

The cases of mixed causal properties are realistic but often not ideal for intervention analysis that is appropriate when most members of a population share invariant causal profiles. These cases are common when an act causally affects an extensive system. For example, a socioeconomic policy affects diverse citizens, an educational program affects numerous students, a business decision affects countless customers, and an approved drug affects various patients. These complicated situations are not rare in decision problems.

In this article, I have identified the example of the spinner that underlines the importance of SCMs and SEU. In that example, the characterization of the causal effect of the act delivered by IEU and the characterization delivered by SEU diverge and the latter—not the former—seems intuitively correct. Moreover, the language of SCMs and SEU is richer than intervention analysis and IEU because SCMs and SEU enable the agent to make necessary mathematical statements that relate directly to various causal dispositions in the real world.<sup>15</sup> The theoretical implication of the spinner is that SEU is recommended in similar situations and that SEU might be a foundation for a more general decision theory.<sup>16</sup>

#### REFERENCES

- Ahmed, Arif. 2014. *Evidence, Decision and Causality*. Cambridge: Cambridge University Press.
- Bareinboim, Elias, Andrew Forney, and Judea Pearl. 2015. "Bandits with Unobserved Confounders: A Causal Approach." In *Advances in Neural Information Processing Systems*, ed. C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, 28:1342–50. Red Hook, NY: Curran.

15. Dawid (2015, 280–82) considers several formal frameworks for analyzing causal processes in decision problems. He agrees that intervention expressions are not as flexible as the language of SCMs.

16. Bareinboim, Forney, and Pearl (2015) argue for a similar conclusion by considering a sequential decision problem. I thank Jiji Zhang for pointing this out.

- Bechtel, William, and Adele Abrahamsen. 2005. "Explanation: A Mechanist Alternative." *Studies in History and Philosophy of Science C* 36 (2): 421–41.
- Dawid, Philip. 2015. "Statistical Causality from a Decision-Theoretic Perspective." *Annual Review of Statistics and Its Application* 2:273–303.
- Dupré, John. 1984. "Probabilistic Causality Emancipated." *Midwest Studies in Philosophy* 9 (1): 169–75.
- Eells, Ellery. 1991. *Probabilistic Causality*. Cambridge Studies in Probability, Induction and Decision Theory. Cambridge: Cambridge University Press.
- Galles, David, and Judea Pearl. 1998. "An Axiomatic Characterization of Causal Counterfactuals." *Foundations of Science* 3 (1): 151–82.
- Hitchcock, Christopher. 2016. "Conditioning, Intervening, and Decision." *Synthese* 193 (4): 1157–76.
- Joyce, James. 1999. *The Foundations of Causal Decision Theory*. Cambridge: Cambridge University Press.
- Lewis, David. 1981. "Causal Decision Theory." *Australasian Journal of Philosophy* 59 (1): 5–30.
- Meeck, Christopher, and Clark Glymour. 1994. "Conditioning and Intervening." *British Journal for the Philosophy of Science* 45 (4): 1001–21.
- Pearl, Judea. 2009. *Causality: Models, Reasoning and Inference*. 2nd ed. Cambridge: Cambridge University Press.
- . 2017. "Physical and Metaphysical Counterfactuals: Evaluating Disjunctive Actions." *Journal of Causal Inference* 5 (2): 1–10.
- . 2021. "Causal and Counterfactual Inference." In *The Handbook of Rationality*, ed. Markus Knauff and Wolfgang Spohn. Cambridge, MA: MIT Press.
- Pearl, Judea, Madelyn Glymour, and Nicholas P. Jewell. 2016. *Causal Inference in Statistics: A Primer*. Chichester: Wiley.
- Shipley, Bill. 2016. *Cause and Correlation in Biology: A User's Guide to Path Analysis, Structural Equations and Causal Inference with R*. 2nd ed. Cambridge: Cambridge University Press.
- Shpitser, Ilya, and Judea Pearl. 2009. "Effects of Treatment on the Treated: Identification and Generalization." In *UAI '09: Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, 514–21. Arlington, VA: AUAI.
- Simon, Herbert Alexander. 1957. *Models of Man, Social and Rational: Mathematical Essays on Rational Human Behavior in a Society Setting*. New York: Wiley.
- Spirtes, Peter, Clark Glymour, and Richard Scheines. 2000. *Causation, Prediction, and Search*. 2nd ed. Cambridge, MA: Bradford.
- Steel, Daniel. 2005. "Indeterminism and the Causal Markov Condition." *British Journal for the Philosophy of Science* 56 (1): 3–26.
- Stern, Reuben. 2017. "Interventionist Decision Theory." *Synthese* 194 (10): 4133–53.
- . 2019. "Decision and Intervention." *Erkenntnis* 84:783–804.
- Weirich, Paul. 2016. "Causal Decision Theory." In *Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta. Stanford, CA: Stanford University. <https://plato.stanford.edu/archives/win2016/entries/decision-causal/>.
- Woodward, James. 2005. *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.