This is an early draft to be presented at: The 27th Biennial Meeting of the Philosophy of Science Association, 2020, Baltimore, MD. Please do not cite or quote without permission. Any criticisms or suggestions are welcome. Contact Email: tungyingwu@outlook.com

Structural Decision Theory

Tung-Ying Wu
Fudan University
tungyingwu@outlook.com

Draft as of March 5, 2020

Abstract

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. I found that it is structural models, rather than intervention analysis, serves as a valuable formal tool for a range of realistic decision problems that involve mixed causal mechanisms. 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 (hereafter, 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.[1]

Interventionist decision theory (hereafter, IDT) is a form of CDT because IDT also holds that the relevant information that matters to our decision should be causal, but IDT approaches an act's causal efficacy through intervention analysis within the framework of causal modeling.[2] Specifically, IDT holds that an agent should conceive of an act as an

---

[1] David Lewis, 1981, p. 11; James Joyce, 1999, pp. 146; Ralph Wedgwood, 2013, p. 2644; Arif Ahmed, 2014, pp. 8-9; Paul Weirich, 2016.

[2] Peter Spirtes, Clark Glymour, and Richard Scheines (2000, pp. 47-53) and Judea Pearl (2009, pp. 23-4, 70-4) claim that 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.

More formally, intervention analysis is assessed by the theory of causal Bayes net. 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 non-occurrence, or a value if an event is of a quantity. A Bayesian causal model M is a triple $<G, V, P>$, 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. 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. 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: $P(X_1, \ldots, X_n) = \prod_i P(X_i|PA(X_i))$, where $X_1, \ldots, X_n$ are all variables in V, and "PA_i" stands for "parents of X_i."

An intervention on X_j removes all its pre-existing cause and set 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 set X_j to a specific value and make it no longer depends on its original parents. Hence, the effect of intervention that disables all pre-existing causes of the act in a decision problem.[3,4] This is because causal models 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 importantly, by making use of causal models, one can distinguish causation from correlation.[5]

IDT instructs rational agents to choose an act *x* that maximizes the interventionist expected utility (hereafter, IEU. See below.). 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*.[6] Here is Pearl's definition of IEU:[7]

$$IEU(x) =^{\text{df}} \sum_{y} P(Y = y \mid do(X = x))\, V(Y = y)$$

This definition asserts that *s* should assess the expected utility of an outcome *y* based on evaluating the effect of the intervention to make *s* do *x*.

---

intervention on X_j is obtained by the new joint distribution: $P'(X_1, \ldots, X_n) = \prod_{i \notin j} P(X_i | PA(X_i))$.

[3] See Christopher Meek and Clark Glymour, 1994, pp. 1007-8; Pearl, 2009, p. 70 and pp. 108-112; Christopher Hitchcock, 2016, pp. 1158-9; Reuben Stern, 2017, pp. 4139-42; Stern, 2018, pp. 2-3. 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, p. 1166, and Stern, 2018, pp. 7-8.

[4] Note that the notion of "intervention" in this paper is not the same as James Woodward's (2003, pp. 94-98). In this paper, "intervention analysis" is understood in terms of manipulating the probability distribution in a causal model where the causal Markov condition holds. See footnote 2.

[5] Meek and Glymour, 1994; Pearl, 2009, section 4.1; Hitchcock, 2015, p. 1175; Stern, 2017, p. 4147.

[6] Pearl uses the do-operator to denote "intervention."

[7] Pearl, 2009, p. 108. For similar proposals, see Meek and Glymour, 1994, pp. 1009-10; Hitchcock, 2016, pp. 1162-4.

Nevertheless, Pearl (2017 and forthcoming) recently proposes a new definition of expected utility in terms of structural causal models (hereafter, SCM) as decision-making conditionals. Call the definition of expected utility with an application of SCM "the structural expected utility" (hereafter, SEU):[8]

$$SEU(x) =^{\mathrm{df}} \sum_{y} P(Y_x = y)\, V(Y = y)$$

Pearl entitles $P(Y_x = y)$ as a SCM defined counterfactuals.[9] This definition declares that *s* should evaluate the expected utility of act *x* by using an 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.[10] In contrast, SEU requires delineating functional relations between relevant variables to attain the causal structure.[11] They are nevertheless different methodologies for the agent to approach decision problems.

This paper attempts to assess the scope of SEU and IEU, 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

[8] Pearl, 2017, p. 1 and forthcoming, p. 1. Note that Pearl (2009, p. 108) originally endorsed IEU. Also, Pearl sometimes uses $P(Y = y \mid do(X = x))$ and $P(Y_x = y)$ interchangebly in his writings because the later can be translated and computed by the former under several strong assumptions. Such translation would fail in some examples. See Pearl, 2009, pp.245-7, 289-93 and Pearl et al., 2016, pp. 107-116.

[9] Pearl, forthcoming, pp. 2-6. For the comparison between the causal modeling's and Lewis's accounts of counterfactuals, see Eric Hiddleston, 2005, Woodward, 2003, pp. 133-145, and Pearl, 2009, pp. 238-41, and Pearl, 2017.

[10] See footnote 2.

[11] I will formally expand on this later.

Newcomb's Problem and Psychopath Button, or issues of uncertainty about causal dependency.[12] To the best of my knowledge, the distinction between IEU and SEU has not been dealt with in depth. The example in next section demonstrates that it is SEU, rather than IEU, serves as a valuable formal tool for a range of realistic decision problems that involve mixed causal mechanisms. Therefore, SEU provides a more general framework for rational decision-makers.

The following sections of this paper are 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 SCM. Section 4 employs SCM to analyze the Spinner and shows how SCM 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 (drawn below) and an arrow in the circle. 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 +

[12] Meek and Glymour, 1994, pp. 1008-9; Hitchcock, 2016, pp. 1165-9; Stern, 2017, pp. 4142; Stern, 2018, pp. 15-16.

X units. However, if the arrow stops in area $Z = 2$, Z will be deleterious so the reward will have $X - 2$ units.

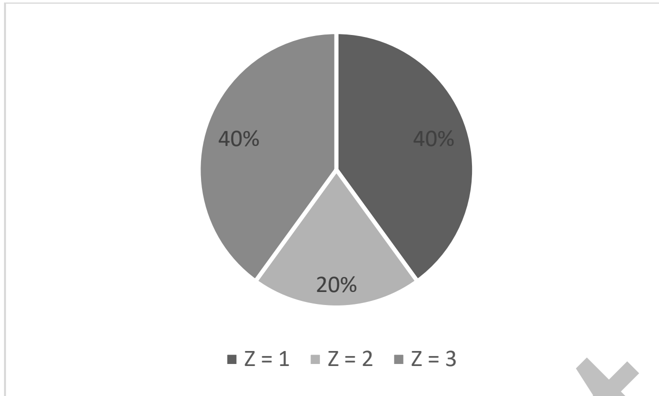

Figure 1. The Spinner

Now, assessing the expected gain of option ADD-X is a complicated task.[13] 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 causal mechanisms of these areas differ from area to area because they exhibit different dispositions that manifest given the presence of the agent's acting on ADD-X.

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 by the theory of causal Bayes net with the assumption that the relevant causal model satisfies the causal Markov condition. However, the procedure of computing the effect of acting on ADD-X as an

[13] This example is a modified case of "additive intervention." Namely, one evaluates the effect of adding some amount from X without removing a pre-existing causal process of X. (In Newcomb's Spinner, I use X as an instrument variable, Y as some amount, Z as a preexisting cause of Y.) See Bill Shipley, 2016, pp. 9-11, 50-4 and Pearl, 2009, section 11.4.4. Pearl et al. (2016, pp. 109-111) confirm that the effect of additive intervention could not be reduced to intervention-expressions alone.

intervention amounts to computing $P(Y = y \mid do(X = q))$, 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. 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 causal mechanisms.[14]

In the Spinner, the agent cannot intervene to fix the amount of the reward. For doing so is an intervention that removes the pre-existing 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. 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

[14] One cannot evaluate the causal efficacy of ADD-X by the analysis of interventions $P(Y = y \mid do(X = x, Z = z))$, $P(Y = y \mid do(X + Z))$, and $P(Y = y \mid do(X - Z))$. As stipulated in the example, it is not possible for the agent to intervene to set Z to a fixed value.

knows each area's causal mechanism, she should evaluate the causal effects of her interventions area by area. Since the issue is predicting the expected gain of ADD-X, the agent should average the causal effects in each area by its proportion to the whole circle to derive the desired quantity.

This paper puts forward a justification for applying SCM 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 SCM to define an act's expected utility in decision theory, namely, SEU. In what follows, I will introduce SCM, which may be of use to a rational agent to accurately predict what is causally downstream from her acts.

3. Structural Causal Models

SCM 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 SCM includes the work of the economist Herbert A. Simon who specialized in decision-making. In his influential papers, Simon 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.[15] I will begin with a brief account of SCM.

[15] Herbert A. Simon, 1957, pp. 10-13.

A structural causal model M consists of a quadruple <U, V, f , P>, where U is a set of exogenous (or background) variables, 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. f is a set of functions that assign each endogenous variable in V a value based on the values of the other variables in the model. P represents a probability distribution over all variables in U. Specifically, each function has the form:[16]

$X_i = f_i\ (PA_i, U_i)$, $i = 1,\ldots,$n

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 causal 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

[16] Simon, 1957, pp. 18-19, 40; Pearl, 2009, pp. 202-3; Pearl et al, 2016, pp. 26-7.

over all exogenous variables, the probability distribution for the endogenous variables is also P.[17, 18]

A structural causal model M corresponds to a causal graph G. If the value of a variable Y depends on X according to the function f_Y, then there will be a directed path from X to Y.[19]

For the sake of illustration of SCM and an explicit representation of the causal relationships in the Spinner, I will use SCM to represent the Spinner, and demonstrate that the agent can accurately predict what is causally downstream of her acts in SCM's expressions in the next section.

4. Additive Intervention

I use the linear SCM M_1: $<U, V, f, P>$ to represent the causal relationships in the Spinner. Let X, Y, Z be endogenous variables in V, and I, U_Z be exogenous variables in U.[20] These variables range over possible values (denoted by lowercase letters). 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 (for example, the agent's free will) determine its value.

[17] Simon, 1957, pp. 40-3, 54-6; Pearl, 2009, pp. 27-32, 205-6; Pearl et al, 2016, p. 98.
[18] A consequence of a structural causal model M is that the probability distribution of every variable in M satisfies the causal Markov condition (CMC). CMC holds in SCM 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 Pearl, 2009, p. 30; Daniel Steel, 2005, p. 10.
[19] Pearl, 2009, p. 203.
[20] For the sake of brevity, I omit some exogeneous variables.

In the Spinner, X can increase the prize Y, and Z also causally affects Y's value. The causal graph G_1 of this model M_1 is figure 2:

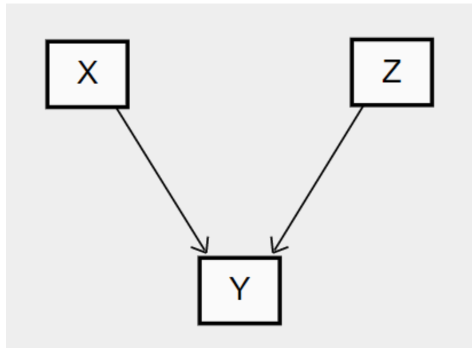

Figure 2. The causal graph G_1 of the Spinner with some exogenous variables omitted.

The following functions represent the causal relations between these variables:

f_X: $X = \{q, 0\}$

f_Z: $Z = U_Z$

$$f_Y: Y = \begin{cases} Z \ if \ X = 0 \\ X \ \ if \ X > 0 \ and \ Z < 2 \\ X - Z \ \ if \ X > 0 \ and \ Z = 2 \\ Z + X \ \ if \ X > 0 \ and \ Z \ > 2 \end{cases}$$

U_Z is an exogenous variable that 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.

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 the agent adds some value ($X > 0$), but Z is lower than 2, then the value of Y will be X. If the agent adds some value, but Z equals

2, then the amount of Y will be $X - Z$. If the agent 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.

Turning now to the question of how an agent predicts the overall causal effects of choosing ADD-X. The diverse areas have varied types of causal 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. One can estimate the results in each level of Z and averages these effects by the probability distribution of Z.[21] I now turn to explain how this sort of prediction is done in M_1.

One 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 $Z = z$.[22] Given a structural model M and observed information $Z = z$, one can evaluate the conditional $P(Y_x = y \mid Z = z)$ in three steps:[23, 24]

(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.

---

[21] In cases where experimental units manifest variant dispositional properties, Spirtes, *et al*. (2000, p.165-7) also cite similar calculations to obtain predictions.
[22] $P(Y_x = y)$ is a subjunctive conditional. "$P(Y_x = y)$" stands for the probability that, had an intervention *do*($X = x$) been performed, an outcome $Y = y$ would obtain.
[23] The following procedure draws from David Galles and Pearl (1998), Pearl (2009, pp. 202-6), and Pearl et al. (2016, pp. 92-8). Joseph Halpern (2000) provide another detailed account of causal inferences in SCM. For the sake of simplicity, I will skip some unnecessary technical details. Note that this is different from Woodward's notion of causality analyzed with counterfactual interventions.
[24] Pearl, 2009, p. 37, 206.

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_1, 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 based on the modified M and the updated values of U.

Returning to the question posed in the Spinner, it is now possible to answer the agent's question of assessing SEU of choosing ADD-X by SCM. First, the agent 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: $Z = U_Z$

$$f_Y\text{: } Y = \begin{cases} Z \; if \; X = 0 \\ X \;\; if \; X > 0 \; and \; Z < 2 \\ X - Z \;\; if \; X > 0 \; and \; Z = 2 \\ Z + X \;\; if \; X > 0 \; and \; Z \; > 2 \end{cases}$$

Next, 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 fee that the agent 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 SCM and deriving SEU in the Spinner and similar situations is more fitting than intervention analysis. As demonstrated in the Spinner, the approach of SCM captures the mixture of variant causal mechanisms specified by the probability distribution of Z and the function f_y, and thereby obtains more accurate characterizations of each area's causal property and the causal efficacy of choosing ADD-X. Hence, in cases where 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 SCM's 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; an approved drug affects various patients. It would seem that these complicated situations are not rare in decision problems.

In this paper, I have identified the example of the Spinner which underlines the importance of SCM 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 SCM and SEU is richer than intervention analysis and IEU because SCM and SEU enable the agent to make necessary mathematical statements that relate directly to various causal dispositions in the real world.[25] 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.

References

Ahmed, Arif. 2014. *Evidence, Decision and Causality*. Cambridge University Press.

Dawid, Philip. 2015. "Statistical Causality from a Decision-Theoretic Perspective." *Annual Review of Statistics and Its Application* 2: 273–303.

Egan, Andy. 2007. "Some Counterexamples to Causal Decision Theory." *Philosophical Review* 116 (1): 93–114.

Galles, David, and Judea Pearl. 1998. "An Axiomatic Characterization of Causal Counterfactuals." *Foundations of Science* 3 (1): 151–182.

Gibbard, Allan, and William Harper. 1978. "Counterfactuals and Two Kinds of Expected Utility." In *Foundations and Applications of Decision Theory*, edited by A. Hooker, J. J. Leach, and E. F. McClennen, 125–62. D. Reidel.

Halpern, Joseph Y. 2000. "Axiomatizing Causal Reasoning." *Journal of Artificial Intelligence Research* 12 (1): 317–337.

Hiddleston, Eric. 2005. "A Causal Theory of Counterfactuals." *Noûs* 39 (4): 632–57.

Hitchcock, Christopher. 2013. "What Is the 'Cause' in Causal Decision Theory?" *Erkenntnis (1975-)* 78: 129–46.

———. 2016. "Conditioning, Intervening, and Decision." *Synthese* 193 (4).

Joyce, James M. 1999. *The Foundations of Causal Decision Theory*. Cambridge University Press.

Lewis, David. 1981. "Causal Decision Theory." *Australasian Journal of Philosophy* 59 (1): 5–30.

Meek, Christopher, and Clark Glymour. 1994. "Conditioning and Intervening." *The British Journal for the Philosophy of Science* 45 (4): 1001–21.

Pearl, Judea. 2009. *Causality: Models, Reasoning and Inference*. 2nd edition. Cambridge, U.K. ; New York: Cambridge University Press.

———. 2017. "Physical and Metaphysical Counterfactuals: Evaluating Disjunctive Actions." *Journal of Causal Inference* 5 (2): 1–10.

Pearl, Judea, Madelyn Glymour, and Nicholas P. Jewell. 2016. *Causal Inference in Statistics: A Primer*. 1 edition. Chichester, West Sussex: Wiley.

Shipley, Bill. 2016. *Cause and Correlation in Biology: A User's Guide to Path Analysis, Structural Equations and Causal Inference with R*. 2 edition. Cambridge: Cambridge University Press.

[25] Philip Dawid (2015, pp. 280-2) considers several formal frameworks for analyzing causal processes in decision problems. He agrees that intervention-expressions are not as flexible as the language of SCM.

Shpitser, Ilya, and Judea Pearl. 2009. “Effects of Treatment on the Treated: Identification and Generalization.” In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, 514–521. UAI ’09. Arlington, Virginia, United States: AUAI Press

Simon, Herbert Alexander. 1957. *Models of Man, Social and Rational: Mathematical Essays on Rational Human Behavior in a Society Setting*. New York: John Wiley and Sons.

Spirtes, Peter, Clark Glymour, and Richard Scheines. 2001. *Causation, Prediction, and Search, Second Edition*. Second edition. Cambridge, Mass: A Bradford Book.

Steel, Daniel. 2005. “Indeterminism and the Causal Markov Condition.” *British Journal for the Philosophy of Science* 56 (1): 3–26.

Stern, Reuben. forthcoming. “Decision and Intervention.” *Erkenntnis*, 1–22.

———. 2017. “Interventionist Decision Theory.” *Synthese* 194 (10): 4133–4153.

Wedgwood, Ralph. 2013. “Gandalf’s Solution to the Newcomb Problem.” *Synthese* 190 (14): 2643–75.

Weirich, Paul. 2016. “Causal Decision Theory.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Winter 2016. Metaphysics Research Lab, Stanford University.

Woodward, James. 2005. *Making Things Happen: A Theory of Causal Explanation*. Oxford University Press, USA.