

A Representation Theorem for Causal Decision Making

Joseph Y. Halpern¹, Evan Piermont²

¹Computer Science Department, Cornell University, Ithaca, USA.

²Department of Economics, Royal Holloway, University of London, UK.

halpern@cs.cornell.edu, evan.piermont@rhul.ac.uk

Abstract

We show that it is possible to understand and identify a decision maker’s subjective causal judgements by observing her preferences over interventions. Following Pearl (2000), we represent causality using *causal models* (also called *structural equations models*), where the world is described by a collection of variables, related by equations. We show that if a preference relation over interventions satisfies certain axioms (related to standard axioms regarding counterfactuals), then we can define (i) a causal model, (ii) a probability capturing the decision-maker’s uncertainty regarding the external factors in the world and (iii) a utility on outcomes such that each intervention is associated with an expected utility and such that intervention A is preferred to B iff the expected utility of A is greater than that of B . In addition, we characterize when the causal model is unique. Thus, our results allow a modeler to test the hypothesis that a decision maker’s preferences are consistent with some causal model and to identify causal judgements from observed behavior.

1 Introduction

Causal judgments play an important role in decision making. When deciding between actions that intervene directly on some aspect of the world, one major source of uncertainty is the indirect effect of such actions via causal interaction. For example, when deciding the interest rate, the Federal Reserve might consider the possibility that a change in the interest rate will cause a change in unemployment, and further that this causal relationship itself might be contingent on other macroeconomic variables.

Uncovering and describing the causal relationship between variables is a task that has led to enormous effort across many different disciplines (see, e.g., (Angrist and Pischke 2009; Cunningham 2021; Hernán and Robins 2020; Morgan and Winship 2007; Parascandola and Weed 2001; Plowright et al. 2008; Pearl 2009; Pearl 2000; Spirtes, Glymour, and Scheines 1993); this list barely scratches the surface). Different decision makers, on account of their private information and personal experience, might hold different beliefs about causal relationships. In this paper, we show that it is possible to understand and identify a decision maker’s subjective causal judgements by observing her preferences over interventions.

A first step to doing this is to decide how to represent causality. Most recent work has focused on using counter-

factuals. In the philosophy community, following Stalnaker (1968) and Lewis (1973), counterfactuals are given semantics using possible worlds equipped with a “closer than” relation; a counterfactual such as “If ϕ were true then ψ would be true” is true at a world ω if ψ is true at the closest world(s) to ω where ϕ is true. Pearl (2000) has championed the use of *causal models* (also called *structural equations models*), graphical models where the world is described by a collection of variables, related by equations. (These are related to models of causality in economics that go back to the work of Haavelmo (1943) and Simon (1953).) The equations model the effect of counterfactual interventions.

We use the latter approach here, although as we shall show, it is closely related to the former approach. Following Pearl, we assume that the world is described by a set of variables, which is split into two sets: the *exogenous* variables, whose values are determined by factors outside the model, and the *endogenous* variables, whose values are determined by the exogenous variables and other endogenous variables. Which variables should be exogenous and which should be endogenous depends on the situation. For example, if the Federal Reserve is deciding whether to change the interest rates, the interest rates should clearly be viewed as endogenous. But for a company trying to decide whether to go ahead with a project that involves borrowing money at the current interest rates, the interest rates are perhaps best viewed as exogenous.

A primitive action is an intervention that sets the value of a particular variable; because the values of the exogenous variables are taken as given, we assume that only endogenous variables can be intervened on. Following Pearl (2000), we use the **do** notation to denote such actions. For example **do**[$Y \leftarrow y$] is the primitive action that sets variable Y to value y . Following Blume, Easley, and Halpern (2021) (BEH from now on), we allow more general actions to be formed from these primitive actions using **if ... then ... else**. Thus, non-primitive actions are conditional interventions of the form **if ϕ then A else B** , where A and B are themselves (possibly primitive) actions and ϕ is a test—a statement regarding the values of variables that is either true or false.

As is standard in decision-theoretic analyses, we assume that the decision maker has a preference relation \succsim over actions. We show that if the preference relation satisfies cer-

tain axioms (that correspond to standard axioms regarding counterfactuals), then we can represent the decision-maker’s preference as the maximization of the expected utility of an actions relative to a causal model equipped with a probability and utility. Specifically, given a preference relation, we can define a causal model, add a probability on contexts (settings of the exogenous variables, which can be viewed as capturing the decision-maker’s uncertainty regarding the external factors in the world), and a utility on outcomes (which are just settings of the variables in the model). In such a causal model, we can determine the expected utility of an action that involves (conditional) interventions. We show that the decision maker prefers action A to B iff the expected utility of A is greater than that of B , given the causal model, probability, and utility.

Example 1. *The decision maker is the U.S. Federal Reserve (the Fed), whose objective is to promote the overall health of the economy, which can be characterized by two variables, inflation rate (X_{inf}) and employment level (X_{emp}). These variables are themselves determined by the following variables: the current interest rate (Y_{rate}), the growth rate of the economy (U_{grow}), and the prior interest rate (U_{prior}). The growth rate and the prior interest rate are exogenous; the remaining variables are endogenous.*

Assume for simplicity that all variables take values in $\{0, 1\}$, where the values 0 and 1 represent the interest rate/inflation/employment level/growth rate being low (respectively, high), and the causal equations are

$$\begin{aligned} X_{inf} &= 1 - Y_{rate} && (F_{X_{inf}}) \\ X_{emp} &= 1 - (Y_{rate} \times (1 - U_{grow})) && (F_{X_{emp}}) \\ Y_{rate} &= U_{prior} && (F_{Y_{rate}}) \end{aligned}$$

That is, both inflation and employment depend on the interest rate, and employment also depends on the current growth rate of the economy, which is exogenous. The current interest rate is set equal the prior rate; that is, it remains unchanged unless intervened on.

The utility of the Fed is completely determined by the inflation rate and employment level, and is given by

$$u(X_{inf}, X_{emp}) = 2X_{emp} - X_{inf}.$$

These preferences over outcomes, and the effect of interventions as determined by the structural equations, determine the Fed’s preferences over interventions.

Consider the case where the prior rate is known but there is uncertainty about the growth rate. The Fed contemplates the value of the two interventions (1) set the interest rate to 1 (i.e., $\text{do}[Y_{rate} \leftarrow 1]$) and (2) set the interest rate to 0 (i.e., $\text{do}[Y_{rate} \leftarrow 0]$). Note that, given the prior interest rate, one of these interventions will be equivalent to doing nothing.

In the case where the growth rate is high (i.e., $U_{grow} = 1$), then $X_{emp} = 1$ and so $\text{do}[Y_{rate} \leftarrow 1]$ yields a utility of 2, whereas $\text{do}[Y_{rate} \leftarrow 0]$ yields a utility of 1. So the Fed prefers to intervene and set the interest rate to 1. Conversely, if the growth rate is low ($U_{grow} = 0$), then $Y_{emp} = 1 - Y_{rate}$, so $\text{do}[Y_{rate} \leftarrow 1]$ yields a utility of 0, whereas $\text{do}[Y_{rate} \leftarrow 0]$ yields a utility of 1. In that case, the Fed prefers to leave the interest rate at 0. Thus, the Fed’s

preference between these two interventions depends on its probabilistic assessment of the growth rate. It would clearly prefer the conditional action

if ($U_{grow} = 1$) then $\text{do}[Y_{rate} \leftarrow 1]$ else $\text{do}[Y_{rate} \leftarrow 0]$

to either constant action.

This example takes the causal model (and the goals, i.e., utilities, of the decision maker) as given, and uses them to infer preferences over interventions/actions. Our representation theorem characterizes when the reverse process is also possible. If and only if our axioms hold, then we can start with preferences over actions and from them deduce a causal structure that is consistent with them. Thus, our results allow a modeler to test the hypothesis that a decision maker’s preferences are consistent with some causal model. In doing so, we provide a definition of *causally sophisticated* behavior as a benchmark of rationality for decision making in the presence of interventions.¹ If a decision maker is causally sophisticated—if her actions can be rationalized by some causal model—our results further determine when the causal model can be uniquely identified. This allows a modeler to explain observed behavior in terms of causal judgments.

To the best of our knowledge, we are the first to examine causal decision making in the context of the structural equations, which is perhaps now the most common approach to representing causality in the social sciences and computer science. Bjorndahl and Halpern (2021) (BH from now on) addressed similar questions in the context of the closest-world approach to counterfactuals of Lewis (1973) and Stalnaker (1968). Interestingly, our technical results use their results (and earlier results of BEH) as the basis for our representation theorem, based on a connection between the two representations of counterfactuals due to Halpern (2013). Our results allow us to relate the two approaches at a decision-theoretic level.

Other approaches to modeling causality have also been considered in the literature. Schenone (2020) and Ellis and Thysenb (2021) take a statistical approach, taking a lack of conditional independence as a definition of causality; Alexander and Gilboa (2023) understand causality through a reduction in the Kolmogorov complexity. As we said, our approach is closer to the approach that is currently the focus of work in the social sciences and computer science. Importantly, by identifying the structural equations, we provide a more detailed insight into the causal mechanisms being considered by the decision maker.

In our representation (like that of BH), both utility and probability are defined on valuations of variables. So, in contrast to the decision-theoretic tradition following Savage (1954) (and, of particular relevance, Schenone (2020)), there is no separation between states (on which probability is defined) and outcomes (on which utilities are defined): our decision maker derives utility directly from the outcome of the intervention. This permits the application of our model to the many economically relevant contexts where the effect of the intervention has direct utilitarian consequences for the

¹This is somewhat analogous to *probabilistic sophistication* in decision making (Machina and Schmeidler 1992).

decision maker (e.g., a government setting policy, or a firm deciding an investment strategy).

2 Causal Models

Let \mathcal{U} and \mathcal{V} denote the set of exogenous and endogenous variables, respectively. Let $\mathcal{R} : \mathcal{U} \cup \mathcal{V} \rightarrow 2^{\mathbb{R}}$ associate to each $Y \in \mathcal{U} \cup \mathcal{V}$ a set $\mathcal{R}(Y)$ of possible values called its range. We assume that the range for each variable is finite. Let $\mathcal{S} = (\mathcal{U}, \mathcal{V}, \mathcal{R})$ denote a *signature*.

A *causal model* is a pair $M = (\mathcal{S}, \mathcal{F})$, where \mathcal{S} is a signature and \mathcal{F} is a collection of structural equations that determine the values of endogenous variables based on the values of other variables. Formally, $\mathcal{F} = \{F_X\}_{X \in \mathcal{V}}$, where

$$F_X : \prod_{Y \in \mathcal{U} \cup (\mathcal{V} - \{X\})} \mathcal{R}(Y) \rightarrow \mathcal{R}(X).$$

A model M is *recursive* if there exists a partial order on \mathcal{V} such that the structural equations for each variable is independent of the variables lower in the order. When we say that X is independent of Y , we mean that F_X does not depend on the value of Y . Given a signature \mathcal{S} , let $\mathcal{M}(\mathcal{S})$ denote the set of all models over the signature \mathcal{S} and $\mathcal{M}^{rec}(\mathcal{S})$ the set of recursive models.

We use the standard vector notation to denote sets of variables or their values. If $\vec{X} = (X_1, \dots, X_n)$ is a vector, we write, in a standard extension of the usual containment relation, $\vec{X} \subseteq \mathcal{U} \cup \mathcal{V}$ to indicate that each $X_i \in \mathcal{U} \cup \mathcal{V}$. Similarly, we write $\vec{x} \in \mathcal{R}(\vec{X})$ if $\vec{x} \in \prod_{i \leq n} \mathcal{R}(X_i)$. For a vector \vec{Y} of variables, $\vec{y} \in \mathcal{R}(\vec{Y})$, and a vector $\vec{X} \subseteq \vec{Y}$, let $\vec{y}|_{\vec{X}}$ denote the restriction of \vec{y} to \vec{X} . In particular, for a single variable X , $\vec{y}|_X$ denotes the X^{th} component.

A *context* is a vector \vec{u} of values for all the exogenous variables \mathcal{U} . Let $\mathcal{C}(\mathcal{S}) = \prod_{Y \in \mathcal{U}} \mathcal{R}(Y)$ consist of all contexts for the signature \mathcal{S} . Call a pair $(M, \vec{u}) \in \mathcal{M}^{rec}(\mathcal{S}) \times \mathcal{C}(\mathcal{S})$ a *situation*. In a recursive model, each situation has a unique solution (i.e., a unique value for each variable that simultaneously satisfies all the structural equations and agrees with the context).

Given a model $M = (\mathcal{S}, \mathcal{F})$, we can construct a new model (over the same signature) that represents the counterfactual situation where some variables are set to specific values. If $\vec{Y} \subseteq \mathcal{V}$ and $\vec{y} \in \mathcal{R}(\vec{Y})$, then $M_{\text{do}[\vec{Y} \leftarrow \vec{y}]} = (\mathcal{S}, \mathcal{F}_{\text{do}[\vec{Y} \leftarrow \vec{y}]})$ denotes the model that is identical to M except that the equation for each variable $X \in \vec{Y}$ in $\mathcal{F}_{\text{do}[\vec{Y} \leftarrow \vec{y}]}$ is replaced by $X = \vec{y}|_X$.² Note that $M_{\text{do}[\vec{Y} \leftarrow \vec{y}]}$ is recursive if M is.

2.1 Syntax and Semantics

Fix a signature \mathcal{S} . For each $X \in \mathcal{U} \cup \mathcal{V}$ and $x \in \mathcal{R}(Y)$, let $X = x$ denote the atomic proposition that says that the variable X takes value x . Let $\mathcal{L}(\mathcal{S})$ denote the language

²It is more standard in the literature to write $M_{\vec{Y} \leftarrow \vec{y}}$ rather than $M_{\text{do}[\vec{Y} \leftarrow \vec{y}]}$. We use the latter notation because it makes it easier to express some later notions.

constructed by starting with these propositions and closing off under negation, conjunction, and disjunction.³

A formula $\phi \in \mathcal{L}(\mathcal{S})$ is either true or false in a situation $(M, \vec{u}) \in \mathcal{M}^{rec}(\mathcal{S}) \times \mathcal{C}(\mathcal{S})$. We write $(M, \vec{u}) \models \psi$ if ψ is true in the situation (M, \vec{u}) . The \models relation is defined inductively.

- $(M, \vec{u}) \models X = x$ iff $X = x$ in the unique solution to the system of equations \mathcal{F} , starting with context \vec{u} .
- $(M, \vec{u}) \models \neg\phi$ iff not $(M, \vec{u}) \models \phi$.
- $(M, \vec{u}) \models \phi \wedge \psi$ iff $(M, \vec{u}) \models \phi$ and $(M, \vec{u}) \models \psi$.
- $(M, \vec{u}) \models \phi \vee \psi$ iff $(M, \vec{u}) \models \phi$ or $(M, \vec{u}) \models \psi$.

An *atom* is a complete description of the values of variables; its truth will completely determine truth of all of formulae in $\mathcal{L}(\mathcal{S})$. Formally, an *atom over \mathcal{S}* is a conjunction of the form $\bigwedge_{Y \in \mathcal{U} \cup \mathcal{V}} Y = y$, where y is some value in $\mathcal{R}(Y)$. We use \vec{a} to denote a generic atom. Let $\mathcal{A}(\mathcal{S})$ denote the set of atoms over \mathcal{S} . Notice that an atom $\vec{a} \in \mathcal{A}(\mathcal{S})$ determines the truth of all formulas in $\mathcal{L}(\mathcal{S})$; we write $\vec{a} \Rightarrow \phi$ if ϕ is true in situations satisfying \vec{a} . Let $\vec{a}_{M, \vec{u}}$ denote unique atom such that $(M, \vec{u}) \models \vec{a}$.

In a slight abuse of notation, identify each context $\vec{u} \in \mathcal{C}$ with the formula characterizing \vec{u} , that is, the formula $\bigwedge_{U \in \mathcal{U}} U = \vec{u}|_U$.

For our later discussion, it is useful to consider an extension of $\mathcal{L}(\mathcal{S})$ that includes formulas of the form $[\vec{Y} \leftarrow \vec{y}]\phi$, where $\phi \in \mathcal{L}(\mathcal{S})$, $\vec{Y} \subseteq \mathcal{V}$, and $\vec{y} \in \mathcal{R}(\vec{Y})$. We can view this formula as saying “after intervening to set the variables in \vec{Y} to \vec{y} , the formula ϕ holds”. Call this extended language $\mathcal{L}^+(\mathcal{S})$. We can extend the semantics that we gave to formulas in $\mathcal{L}(\mathcal{S})$ as follows:

- $(M, \vec{u}) \models [\vec{Y} \leftarrow \vec{y}]\phi$ iff $(M_{\text{do}[\vec{Y} \leftarrow \vec{y}]}, \vec{u}) \models \phi$.

In the sequel, we will also make use of another approach to giving semantics to counterfactuals, due to David Lewis and Robert Stalnaker (Lewis 1973; Stalnaker 1968). This approach is based on the idea of “closest worlds”; roughly speaking, with this approach, $[\vec{Y} \leftarrow \vec{y}]\phi$ is true at a world ω if ϕ is true at all the worlds closest to ω where $\vec{Y} = \vec{y}$. Lewis formalizes the idea of “closest world” using a ternary relation R where, for each $\omega \in \Omega$, $R(\omega, \cdot, \cdot)$ is a partial order on Ω . We can interpret $R(\omega_1, \omega_2, \omega_3)$ as saying that ω_2 is closer to ω_1 than ω_3 is. (In this paper, we focus on the case that, for each world ω , $R(\omega, \cdot, \cdot)$ is a strict linear order.)

In more detail, a *Lewis-style model* is a tuple $M = (\Omega, R, I)$, where R is a ternary relation on Ω as above, and I is an *interpretation* that determines whether each atomic proposition is true or false at each state. Formally, if AP is the set of atomic propositions (as implicitly determined by a set $\mathcal{U} \cup \mathcal{V}$ of exogenous and endogenous variables), then $I : \Omega \times AP \rightarrow \{\text{true}, \text{false}\}$. We can again define a relation \models by induction, by taking

³It is more standard to have the atomic propositions involve only endogenous variables, but for our purposes, it is important to include exogenous variables as well. We explain why when we discuss actions below.

- $(M, \omega) \models X = x$ iff $I(\omega, X = x) = \mathbf{true}$,

defining the semantics of negation, conjunction, and disjunction as above, and for interventions, taking

- $(M, \omega) \models [\vec{Y} \leftarrow \vec{y}] \phi$ iff $(M, \omega') \models \phi$, for all ω' that are minimal worlds according to the order $R(\omega, \cdot, \cdot)$ such that $(M, \omega') \models \vec{Y} = \vec{y}$; that is, the worlds closest to ω for which $\vec{Y} = \vec{y}$ holds.

As shown by Halpern (2013), recursive causal models correspond in a precise sense to a subclass of Lewis-style models; we return to this point in Section 5.

3 Decision Environment

A *primitive action* over \mathcal{S} has the form $\mathbf{do}[Y_1 \leftarrow y_1, \dots, Y_n \leftarrow y_n]$, abbreviated as $\mathbf{do}[\vec{Y} \leftarrow \vec{y}]$, where Y_1, \dots, Y_n is a (possibly) empty list of distinct variables in \mathcal{V} and $y_i \in \mathcal{R}(Y_i)$ for $i = 1, \dots, n$. This action represents an intervention that sets each Y_i to the value y_i . As we said earlier, we allow interventions only on endogenous variables.

The set $\mathbb{A}(\mathcal{S})$ of *actions* over \mathcal{S} is defined recursively, starting with the primitive actions, and closing under **if** ... **then** ... **else**, so that if $A, B \in \mathbb{A}$ and $\phi \in \mathcal{L}(\mathcal{S})$, then

$$\mathbf{if } \phi \mathbf{ then } A \mathbf{ else } B \in \mathbb{A}.$$

We take **if** ϕ **then** A to be an abbreviation of **if** ϕ **then** A **else** $\mathbf{do}[]$ ($\mathbf{do}[]$ is the trivial action that sets no values). Note that although we restrict interventions to interventions to endogenous variables, the tests can involve exogenous variables. For example, even if interests rates are fixed (and hence taken to be exogenous), we may want to say “if the interest rate is 5% then borrow \$1,000,000” (where we take the amount borrowed to be represented by an endogenous variable).

We assume that the decision maker has a preference relation (weak order) $\succsim_{\mathcal{S}}$ over the set of all actions in $\mathbb{A}(\mathcal{S})$. (We often omit the \mathcal{S} if it is clear from context or plays no role in the discussion.) As usual, we write $A \sim B$ if $A \succsim B$ and $B \succsim A$. Call a formula ϕ *null* if its conditional preference is trivial, that is, if $(\mathbf{if } \phi \mathbf{ then } A) \sim (\mathbf{if } \phi \mathbf{ then } B)$ for all actions $A, B \in \mathbb{A}(\mathcal{S})$. Call ϕ *non-null* if it is not null.

Each action $A \in \mathbb{A}(\mathcal{S})$ defines a mapping $h_A : \mathcal{A}(\mathcal{S}) \rightarrow \mathbb{A}^{\text{prim}}(\mathcal{S})$ (where $\mathbb{A}^{\text{prim}}(\mathcal{S})$ is the set of primitive actions). The functions h_A are defined recursively as follows:

$$h_{\mathbf{do}[\vec{Y} \leftarrow \vec{y}]}(\vec{a}) = \mathbf{do}[\vec{Y} \leftarrow \vec{y}]$$

and

$$h_{\mathbf{if } \phi \mathbf{ then } A \mathbf{ else } B}(\vec{a}) = \begin{cases} h_A(\vec{a}) & \text{if } \vec{a} \Rightarrow \phi \\ h_B(\vec{a}) & \text{if } \vec{a} \Rightarrow \neg\phi. \end{cases}$$

3.1 Representation

A *subjective causal expected utility* representation views preferences as the outcomes of maximizing the expected utility of an action, relative to uncertainty regarding the values of exogenous variables. Specifically, the representation is governed by

- M : a model that dictates the causal equations
- \mathbf{p} : a probability distribution on contexts
- \mathbf{u} : a utility function on atoms.

An action (along with the model M) can be thought of as an operation that assigns values to all variables given a context. Specifically given a model M , for each $A \in \mathbb{A}(\mathcal{S})$, define $\beta_A^M : \mathcal{C}(\mathcal{S}) \rightarrow \mathcal{A}(\mathcal{S})$ as

$$\beta_A^M(\vec{u}) = \text{the unique atom } \vec{a} \in \mathcal{A}(\mathcal{S}) \\ \text{such that } (M_{h_A(\vec{a}_M, \vec{u})}, \vec{u}) \models \vec{a}.$$

Unpacking this: if the context is \vec{u} , then the initial assignment of variables is expressed by the atom $\vec{a}_{M, \vec{u}}$, so the primitive action prescribed by A is $h_A(\vec{a}_{M, \vec{u}})$. Thus, the final assignment of variables (i.e., after the intervention $h_A(\vec{a}_{M, \vec{u}})$) is determined by the situation $(M_{h_A(\vec{a}_{M, \vec{u}}), \vec{u}})$: this is what is captured by β .

Definition 1. $(M, \mathbf{p}, \mathbf{u})$ is a subjective causal utility representation of $\succsim_{\mathcal{S}}$, where $M = (\mathcal{S}, \mathcal{F})$ and $\mathcal{S} = (\mathcal{U}, \mathcal{V}, \mathcal{R})$, if \mathbf{p} is a probability on $\mathcal{R}(\mathcal{U})$, \mathbf{u} is a utility function on $\mathcal{A}(\mathcal{S})$, and $A \succsim_{\mathcal{S}} B$ iff

$$\sum_{\vec{u} \in \mathcal{C}(\mathcal{S})} \mathbf{u}(\beta_A^M(\vec{u})) \mathbf{p}(\vec{u}) \geq \sum_{\vec{u} \in \mathcal{C}(\mathcal{S})} \mathbf{u}(\beta_B^M(\vec{u})) \mathbf{p}(\vec{u}). \quad (1)$$

In a representation, we think of contexts as states and atoms as outcomes, so β_A^M can be viewed as a function from states to outcomes. Thus, by associating A with β_A^M , we can think of A as a function from states to outcomes—exactly how Savage views acts.

Notice that the decision maker has no uncertainty about the model M . In particular, should she learn the context, she would know the causal relationships between variables. We discuss this assumption, and its limitations, in Section 6.

4 Axioms

In this section, we discuss the axioms that we need to ensure the existence of a causal expected utility representation.

The first axiom is the cancellation axiom of Blume, Easley, and Halpern (2021). To define it, we need the notion of a *multiset*, which can be thought of as a set that allows for multiple instances of each of its elements; two multisets are equal if they contain the same elements with the same multiplicities. For example, the multiset $\{\{a, a, a, b, b\}\}$ is different from the multiset $\{\{a, a, b, b, b\}\}$: both multisets have five elements, but the multiplicities of a and b differ in the two multisets (assuming $a \neq b$).

Using this notation, we can state the cancellation axiom.

Axiom 1 (Cancellation). *If $A_1 \dots A_n, B_1 \dots B_n \in \mathbb{A}$ and $\{\{h_{A_1}(\vec{a}) \dots h_{A_n}(\vec{a})\}\} = \{\{h_{B_1}(\vec{a}) \dots h_{B_n}(\vec{a})\}\}$ for all $\vec{a} \in \mathcal{A}(\mathcal{S})$, then $A_i \succsim B_i$ for all $i < n$ implies $B_n \succsim A_n$.*

As in the prior literature, cancellation allows us to construct a state space and an additively separable utility representation of $\succsim_{\mathcal{V}}$. The state space that is constructed via the BEH methodology does not have any causal structure. The remaining axioms ensure that we can construct a subjective utility representation that is a causal model.

As the remaining axioms speak directly to the structure of the *do* action, the following notation will be helpful: For each atom $\bar{a} \in \mathcal{A}(S)$, write $\mathbf{do}[\vec{Y} \leftarrow \vec{y}] \rightsquigarrow_{\bar{a}} (Z = z)$ as shorthand for the indifference relation

$$\mathbf{if} \bar{a} \text{ then } \mathbf{do}[\vec{Y} \leftarrow \vec{y}, Z \leftarrow z] \sim \mathbf{if} \bar{a} \text{ then } \mathbf{do}[\vec{Y} \leftarrow \vec{y}]. \quad (2)$$

This says, essentially, that in the context characterized by the atom \bar{a} , setting $\vec{Y} = \vec{y}$ results in $Z = z$. To understand why, notice that intervening to set \vec{Y} to \vec{y} causes Z to equal z if and only if the further intervention setting Z to z has no effect; $\mathbf{do}[\vec{Y} \leftarrow \vec{y}] \rightsquigarrow_{\bar{a}} (Z = z)$ can be thought of as capturing this situation, as the decision maker sees no additional benefit to the additional intervention on Z . Of course, it could be that Z is in fact set to some different value z' , but the decision maker is indifferent between $Z = z$ and $Z = z'$, given the values of the other variables. This indifference is irrelevant as far as the existence of a subjective utility representation goes; we can simply treat the worlds where $Z = z$ and $Z = z'$ identically.

The next axiom ensures that there is only one model in our representation, rather than a distribution over models; that is, the decision maker has no uncertainty about the structural equations. The axiom itself just says that, for each context \bar{u} , there is at most one atom \bar{a} compatible with \bar{u} that is non-null. Intuitively, if M represents \succsim_S , then $(M, \bar{u}) \models \bar{a}$.

Axiom 2 (Model Uniqueness). *For each context \bar{u} , there exists a most one atom $\bar{a} \in \mathcal{A}(S)$ such that $\bar{a} \Rightarrow (U = \bar{u})$ and \bar{a} is non-null.*

The next axiom ensures that there is some value $x \in \mathcal{R}(X)$ such that $X = x$ after \vec{Y} is set to \vec{y} . Galles and Pearl (1998) and Halpern (2000) have analogous axioms.

Axiom 3 (Definiteness). *For each atom \bar{a} , vector of \vec{Y} of endogenous variables, $\vec{y} \in \mathcal{R}(\vec{Y})$, and endogenous variable $X \notin \vec{Y}$, there exists some $x \in \mathcal{R}(X)$ such that $\mathbf{do}[\vec{Y} \leftarrow \vec{y}] \rightsquigarrow_{\bar{a}} (X = x)$.*

The next axiom dictates that intervening to set variables to their actual values does not change the values of other variables; that is, trivial interventions are indeed trivial. It is named after *centering* property considered by Lewis (1973), which ensures that the the closest world to a world ω is ω itself. Let $\bar{a}|_{\vec{Y}}$ denote the restriction of the atom \bar{a} to the conjuncts in \vec{Y} .

Axiom 4 (Centeredness). *For each atom \bar{a} , vector of endogenous variables \vec{Y} , and endogenous variable $X \notin \vec{Y}$, we have $\mathbf{do}[\vec{Y} \leftarrow \bar{a}|_{\vec{Y}}] \rightsquigarrow_{\bar{a}} (X = \bar{a}|_X)$.*

Finally, as we are interested in recursive causal models, we require an axiom that forces the dependency order on endogenous variables to be acyclic. Towards this, for $X, Y \in \mathcal{V}$, say that X is *unaffected* by Y (given \bar{a}) if

$$\mathbf{do}[\vec{Z} \leftarrow \vec{z}] \rightsquigarrow_{\bar{a}} (X = x) \text{ iff } \mathbf{do}[\vec{Z} \leftarrow \vec{z}, Y \leftarrow y] \rightsquigarrow_{\bar{a}} (X = x) \quad (3)$$

for all $\vec{Z} \in \mathcal{V} \setminus \{X, Y\}$, $\vec{z} \in \mathcal{R}(\vec{Z})$, $y \in \mathcal{R}(Y)$, and $x \in \mathcal{R}(X)$. So X is unaffected by Y if there is no intervention on Y that changes the decision maker's perception

of X (conditional on atom \bar{a}). If this relation does not hold, then X is *affected* by Y , written $Y \rightsquigarrow_{\bar{a}} X$. Let $\rightsquigarrow = \cup_{\bar{a}} \rightsquigarrow_{\bar{a}}$. Our final axiom states that \rightsquigarrow is acyclic; it is inspired by the corresponding axiom in (Halpern 2000).

Axiom 5 (Recursivity). *\rightsquigarrow is acyclic.*

Note that Axiom A5 implies that $\rightsquigarrow_{\bar{a}}$ is acyclic for each atom \bar{a} .

As we now show, a preference order satisfies these axioms iff it has subjective causal expected utility representation.

Theorem 1. *The preference order \succsim_S satisfies Axioms A1–A5 if and only if it has a subjective causal expected utility representation $(M, \mathbf{p}, \mathbf{u})$. Moreover, if \succsim_S satisfies Axiom A3*, then M is unique over the set of non-null contexts.*

There may, in general, be many different subjective causal expected utility representations of the same preference relations. This is because, when the decision maker is indifferent between distinct atoms, the preference relation cannot distinguish between them, and hence cannot distinguish between structural equations that result in distinct atoms that have the same utility. This can be seen in the following example.

Example 2. *There is one exogenous variable U with $\mathcal{R}(U) = \{0, 1\}$ and one endogenous variable X with $\mathcal{R}(X) = \{0, 1, 2\}$. Consider the following two possible structural models:*

$$\begin{aligned} X &= U & (F_X) \\ X &= 2U & (F'_X), \end{aligned}$$

and assume that the agent believes with probability 1 that $U = 0$. Thus, in both models, she believes that $X = 0$. Further suppose that \mathbf{u} depends only on X , and $\mathbf{u}(X = 0) < \mathbf{u}(X = 1) = \mathbf{u}(X = 2)$. The agent does not care if X gets set to 1 or gets set to 2, so is indifferent between the effect of the interventions F_X and F'_X . Thus, these two models cannot be distinguished.

The following axiom, a strengthening of A3, rules out situations like Example 2.

Axiom* 3 (Strong Definiteness). *For each non-null atom \bar{a} , vector of \vec{Y} of endogenous variables, $\vec{y} \in \mathcal{R}(\vec{Y})$, and endogenous variable $X \notin \vec{Y}$, there exists a unique $x \in \mathcal{R}(X)$ such that $\mathbf{do}[\vec{Y} \leftarrow \vec{y}] \rightsquigarrow_{\bar{a}} (X = x)$.*

Strengthening A3 to A3* is both necessary and sufficient to avoid multiple representations.

Definition 2. *A subjective causal utility representation of a preference relation of \succsim_S is identifiable if, for all representations $(M, \mathbf{p}, \mathbf{u})$ and $(M', \mathbf{p}', \mathbf{u}')$, non-null contexts \bar{u} , and formulas $\phi \in \mathcal{L}^+(\mathcal{S})$,*

$$(M, \bar{u}) \models \phi \quad \text{if and only if} \quad (M', \bar{u}) \models \phi.$$

Theorem 2. *A subjective causal utility representation $(M, \mathbf{p}, \mathbf{u})$ of \succsim_S is identifiable iff \succsim_S satisfies Axiom A3*.*

5 Proofs of Theorems 1 and 2

If the preference relation \succsim_S is trivial (i.e., the decision maker is indifferent between all actions) then it is easy to see that all $(M, \mathbf{p}, \mathbf{u})$ with \mathbf{u} constant represent \succsim_S . Thus, in the proofs that follow, we assume that \succsim_S is nontrivial.

Before getting into the details, we describe the proofs at a higher level. Given a preference order \succsim_S , the first step is to construct a Lewis-style model M_{\succsim_S} that represents \succsim_S . We want M_{\succsim_S} to be recursive in a precise sense. So for each atom $\vec{a} \in \mathcal{A}(S)$, we define a strict linear order $<_{\vec{a}}$ of $\mathcal{A}(S)$, intended to represent the closeness of atoms to \vec{a} . We then show that, for each atom \vec{a} , intervening to set $\vec{Y} \leftarrow \vec{y}$ has the same consequence as intervening to set *all* variables to their values in the $<_{\vec{a}}$ -closest atom to \vec{a} in which $\vec{Y} = \vec{y}$ holds (see Lemma 3). This is the only property that BH used to prove their representation theorem, which used Lewis-style models (see their Lemma 5 and Theorem 1). It follows that we get the desired Lewis-style representation of M_{\succsim_S} . Moreover, in M_{\succsim_S} , $<_{\vec{a}}$ does in fact represent the closeness of atoms to \vec{a} . These observations allow us to appeal to results of Halpern (2013), and convert M_{\succsim_S} to a causal model that also represents \succsim_S .

To begin, we define $<_{\vec{a}}$. Axiom A5 ensures that we can extend \sim (and hence $\sim_{\vec{a}}$ for each atom \vec{a}) to a strict linear order on $\mathcal{U} \cup \mathcal{V}$, denoted \rightsquigarrow , such that for all $U \in \mathcal{U}$ and $X \in \mathcal{V}$ we have $U \rightsquigarrow X$. For atoms $\vec{b}, \vec{c} \in \mathcal{A}(S)$, let $Y_{\vec{b}, \vec{c}} \in \mathcal{U} \cup \mathcal{V}$ be the \rightsquigarrow -minimal variable on which \vec{b} and \vec{c} disagree. For each variable $X \in \mathcal{U} \cup \mathcal{V}$, define $\ll_{\vec{a}}^X$ to be some fixed strict linear order of $\mathcal{R}(X)$ whose minimal element is the value of X in atom \vec{a} . Define $<_{\vec{a}}$ to be, loosely speaking, a lexicographic order over atoms, ordering first over variables, using \rightsquigarrow , and then over values, using $\ll_{\vec{a}}^X$. We discuss the intuition behind this order after giving its definition.

Take $<_{\vec{a}}$ to be a strict linear order over $\mathcal{A}(S)$ such that

OR1. $\vec{b} <_{\vec{a}} \vec{c}$ if $Y_{\vec{a}, \vec{c}} \rightsquigarrow Y_{\vec{a}, \vec{b}}$, and

OR2. If $Y_{\vec{a}, \vec{b}} = Y_{\vec{a}, \vec{c}}$, then let Y denote $Y_{\vec{b}, \vec{c}}$, let \vec{Z} be the set of endogenous variables (strictly) \rightsquigarrow -less than Y , and let $\vec{z} = \vec{b}|_{\vec{Z}} (= \vec{c}|_{\vec{Z}})$. Then $\vec{b} <_{\vec{a}} \vec{c}$ if $\mathbf{do}[\vec{Z} \leftarrow \vec{z}] \rightsquigarrow_{\vec{a}} \vec{b}|_Y$ and either

(i) not $\mathbf{do}[\vec{Z} \leftarrow \vec{z}] \rightsquigarrow_{\vec{a}} \vec{c}|_Y$, or

(ii) $\mathbf{do}[\vec{Z} \leftarrow \vec{z}] \rightsquigarrow_{\vec{a}} \vec{c}|_Y$ and $y_{\vec{b}} \ll_{\vec{a}}^Y y_{\vec{c}}$, where $y_{\vec{b}}$ (resp., $y_{\vec{c}}$) denotes the value of Y in \vec{b} (resp., \vec{c}).

In general, (OR1) and (OR2) do not completely determine the order; there may be atoms for which neither (OR1) nor (OR2) hold. The remainder of the order can be completed arbitrarily.

A few points of intuition regarding this order. In the Lewis-style model M_{\succsim_S} that we construct, the states are represented by pairs of atoms. The operator $\rightsquigarrow_{\vec{a}}$ lets us probe the structural equations through the effect of interventions. When considering which of \vec{b} or \vec{c} is closer to \vec{a} according to $\rightsquigarrow_{\vec{a}}$, we consider two criteria: (i) which of \vec{b} or \vec{c} coincides with \vec{a} longer (where *longer* means “for more variables, starting with the exogenous context and proceeding

via \rightsquigarrow ”), and (ii) if both \vec{b} and \vec{c} deviate from \vec{a} at the same variable Y , does one coincide with the counterfactual assessment of \vec{a} given by $\rightsquigarrow_{\vec{a}}$, and if both do, is one closer to \vec{a} than the other according to $\ll_{\vec{a}}^Y$. These two criteria are captured by (OR1) and (OR2), respectively.

An example may help clarify: Let \vec{a} be the atom $U = 0 \wedge X = 0 \wedge Y = 0 \wedge Z = 0$, and assume that the order \rightsquigarrow on variables is $U \rightsquigarrow X \rightsquigarrow Y \rightsquigarrow Z$. Consider the following atoms:

$$\begin{aligned}\vec{b} &= U = 0 \wedge X = 0 \wedge Y = 0 \wedge Z = 1 \\ \vec{c} &= U = 0 \wedge X = 0 \wedge Y = 1 \wedge Z = 1 \\ \vec{c}' &= U = 0 \wedge X = 0 \wedge Y = 1 \wedge Z = 0.\end{aligned}$$

(OR1) states that if an atom coincides with \vec{a} longer, it is closer to \vec{a} , so \vec{b} is closer to \vec{a} than either \vec{c} or \vec{c}' . Since \vec{c} and \vec{c}' agree up to Z , their closeness to \vec{a} depends on the value(s) of z such that $\mathbf{do}[X \leftarrow 0, Y \leftarrow 1] \rightsquigarrow_{\vec{a}} (Z = z)$. In particular, if we can have $z = 1$ (which is the value of Z according to \vec{c}) and not $z = 0$, then \vec{c} is closer to \vec{a} than \vec{c}' . Intuitively, this is because, according to $\rightsquigarrow_{\vec{a}}$, if we set X to 0 and Y to 1, Z would be 1, so \vec{c} is consistent with the equations we plan to use in the context encoded by \vec{a} , while \vec{c}' is not. Similarly, if we can have $z = 0$ and not $z = 1$, then \vec{c}' is closer. Because of the possibility of indifference, both $z = 0$ and $z = 1$ could be consistent with $\rightsquigarrow_{\vec{a}}$. In this case, the order $\ll_{\vec{a}}^Y$ serves as a consistent method of breaking ties.⁴

Let $\vec{a}_{\mathbf{do}[\vec{Y} \leftarrow \vec{y}]}$ denote the $<_{\vec{a}}$ -minimal atom satisfying $\vec{Y} = \vec{y}$.

Lemma 1. Suppose that \vec{Y} is a set of endogenous variables, $X \in \mathcal{V} \setminus \vec{Y}$, and $\vec{Z}_X \subseteq \mathcal{V}$ is the (possibly empty) set of endogenous variables strictly \rightsquigarrow -less than X . Then we have

M1. $\vec{a}_{\mathbf{do}[\vec{Y} \leftarrow \vec{y}]}|_{\mathcal{U}} = \vec{a}|_{\mathcal{U}}$, and

M2. The value of X in $\vec{a}_{\mathbf{do}[\vec{Y} \leftarrow \vec{y}]}$ is the $\ll_{\vec{a}}^X$ -minimal element of

$$m(\vec{a}, X) =$$

$$\{x \in \mathcal{R}(X) : \mathbf{do}[\vec{Z}_X \leftarrow \vec{a}_{\mathbf{do}[\vec{Y} \leftarrow \vec{y}]}|_{\vec{Z}_X}] \rightsquigarrow_{\vec{a}} (X = x)\}.$$

Proof. If $\vec{a}_{\mathbf{do}[\vec{Y} \leftarrow \vec{y}]}|_{\mathcal{U}} \neq \vec{a}|_{\mathcal{U}}$, then let \vec{b} be such that $\vec{a}_{\mathbf{do}[\vec{Y} \leftarrow \vec{y}]}|_{\mathcal{V}} = \vec{b}|_{\mathcal{V}}$ and $\vec{b}|_{\mathcal{U}} = \vec{a}|_{\mathcal{U}}$. Since \mathcal{U} forms the initial segment of \rightsquigarrow , it follows from (OR1) that $\vec{b} <_{\vec{a}} \vec{a}_{\mathbf{do}[\vec{Y} \leftarrow \vec{y}]}|_{\mathcal{U}}$, yet \vec{b} satisfies $\vec{Y} = \vec{y}$. This is a contradiction; (M1) follows.

To see that (M2) holds, suppose by way of contradiction that it does not hold. Then there exists a variable X such that the value of X in $\vec{a}_{\mathbf{do}[\vec{Y} \leftarrow \vec{y}]}$ is not the $\ll_{\vec{a}}^X$ -minimal element of $m(\vec{a}, X)$ (note that $m(\vec{a}, X)$ is non-empty by Axiom A3).

⁴ $\ll_{\vec{a}}^Y$ is arbitrary except that its initial element is the value of Y in \vec{a} , reflecting the fact that beyond coinciding with \vec{a} if possible, this tie-breaking is arbitrary. Nonetheless, an order still needs to be fixed to ensure tie-breaking is consistent across different interventions.

Let \vec{b} be the atom that coincides with $\vec{a}_{\text{do}[\vec{Y} \leftarrow \vec{y}]}$ for all variables except X , and the value of X in \vec{b} is the $\ll_{\vec{a}}^X$ -minimal element of $m(\vec{a}, X)$. There are two cases:

- $\vec{a}|_{\vec{Z}_X} = \vec{a}_{\text{do}[\vec{Y} \leftarrow \vec{y}]}|_{\vec{Z}_X}$: From Axiom A4, it follows that $\text{do}[\vec{Z}_X \leftarrow \vec{z}_X] \rightsquigarrow_{\vec{a}} \vec{a}|_{\vec{Z}_X}$. By the definition of $\ll_{\vec{a}}^X$, the value of X in \vec{a} is the $\ll_{\vec{a}}^X$ -minimal element of $m(\vec{a}, X)$. By assumption, the value of X in $\vec{a}_{\text{do}[\vec{Y} \leftarrow \vec{y}]}$ is not the $\ll_{\vec{a}}^X$ -minimal element of $m(\vec{a}, X)$, while the value of X in \vec{b} is the $\ll_{\vec{a}}^X$ -minimal element of $m(\vec{a}, X)$. It follows that $\vec{b} <_{\vec{a}} \vec{a}_{\text{do}[\vec{Y} \leftarrow \vec{y}]}$, a contradiction.
- $\vec{a}|_{\vec{Z}_X} \neq \vec{a}_{\text{do}[\vec{Y} \leftarrow \vec{y}]}|_{\vec{Z}_X}$: Since $\vec{a}|_{\vec{Z}_X} \neq \vec{a}_{\text{do}[\vec{Y} \leftarrow \vec{y}]}|_{\vec{Z}_X} = \vec{b}|_{\vec{Z}_X}$, we have $Y_{\vec{a}, \vec{a}, \vec{a}_{\text{do}[\vec{Y} \leftarrow \vec{y}]}} = Y_{\vec{a}, \vec{a}, \vec{b}}$, so by (OR2), $\vec{b} <_{\vec{a}} \vec{a}_{\text{do}[\vec{Y} \leftarrow \vec{y}]}$, and again we have a contradiction. \square

Lemma 2. For all atoms $\vec{a} \in \mathcal{A}(\mathcal{S})$, (disjoint) vectors of endogenous variables \vec{Y} and \vec{Z} , $\vec{y} \in \mathcal{R}(\vec{Y})$, $\vec{z} \in \mathcal{R}(\vec{Z})$, endogenous variables $X \notin \vec{Z} \cup \vec{Y}$ such that for all $Y \in \vec{Y}$, $X \rightsquigarrow Y$, and $x \in \mathcal{R}(X)$, we have

$$\begin{aligned} \text{do}[\vec{Z} \leftarrow \vec{z}] \rightsquigarrow_{\vec{a}} (X = x) \text{ iff} \\ \text{do}[\vec{Z} \leftarrow \vec{z}, \vec{Y} \leftarrow \vec{y}] \rightsquigarrow_{\vec{a}} (X = x). \end{aligned}$$

Proof. By assumption, for all $Y \in \vec{Y}$, we have $Y \not\rightsquigarrow_{\vec{a}} X$. The lemma then follows from a simple induction argument on the variables of \vec{Y} , appealing to (3) in each step. \square

The following lemma shows that, starting with \vec{a} , the effect of setting \vec{Y} to \vec{y} is the same as setting all the endogenous variables to the value that results from setting \vec{Y} to \vec{y} . Intuitively, this is because for a variable $X \notin \vec{Y}$, we are setting it to the value it already has after setting \vec{Y} to \vec{y} , so nothing further changes.

Lemma 3. For all atoms $\vec{a} \in \mathcal{A}(\mathcal{S})$, vectors of endogenous variables \vec{Y} , and $\vec{y} \in \mathcal{R}(\vec{Y})$, we have that

$$\text{if } \vec{a} \text{ then do}[\vec{Y} \leftarrow \vec{y}] \sim \text{if } \vec{a} \text{ then do}[\vee \leftarrow \vec{a}_{\text{do}[\vec{Y} \leftarrow \vec{y}]}|\vee].$$

Proof. We prove this by induction on the variables in \mathcal{V} with respect to the order \rightsquigarrow . Fix some $X \in \mathcal{V} \setminus \vec{Y}$, let $\vec{Z} \subseteq \mathcal{V}$ denote the (possibly empty) vector of endogenous variables strictly \rightsquigarrow -less than X , and let \vec{z} denote the values of the variables \vec{Z} in $\vec{a}_{\text{do}[\vec{Y} \leftarrow \vec{y}]}$. To simplify notation, let x denote the value of X in $\vec{a}_{\text{do}[\vec{Y} \leftarrow \vec{y}]}$, let $\vec{Y}' \subseteq \vec{Y}$ consist of those variables \rightsquigarrow -greater than X , and let \vec{y}' be the restriction of \vec{y} to \vec{Y}' . We now show that if

$$\text{if } \vec{a} \text{ then do}[\vec{Y}' \leftarrow \vec{y}'] \sim \text{if } \vec{a} \text{ then do}[\vec{Z} \leftarrow \vec{z}, \vec{Y}' \leftarrow \vec{y}'] \quad (4)$$

then

$$\begin{aligned} \text{if } \vec{a} \text{ then do}[\vec{Y} \leftarrow \vec{y}] \sim \\ \text{if } \vec{a} \text{ then do}[\vec{Z} \leftarrow \vec{z}, \vec{Y}' \leftarrow \vec{y}', X \leftarrow x]. \end{aligned}$$

From (M2), it follows that $\text{do}[\vec{Z} \leftarrow \vec{z}] \rightsquigarrow_{\vec{a}} (X = x)$. Applying Lemma 2, we obtain

$$\text{do}[\vec{Z} \leftarrow \vec{z}, \vec{Y}' \leftarrow \vec{y}'] \rightsquigarrow_{\vec{a}} (X = x). \quad (5)$$

Thus,

$$\begin{aligned} & \text{if } \vec{a} \text{ then do}[\vec{Z} \leftarrow \vec{z}, \vec{Y}' \leftarrow \vec{y}', X \leftarrow x] \\ \sim & \text{if } \vec{a} \text{ then do}[\vec{Z} \leftarrow \vec{z}, \vec{Y}' \leftarrow \vec{y}'] \\ \sim & \text{if } \vec{a} \text{ then do}[\vec{Y} \leftarrow \vec{y}], \end{aligned}$$

where the first indifference comes from from (5) and the definition of $\rightsquigarrow_{\vec{a}}$, and the second from (4). \square

BH's Theorem 1 shows that (in our terminology), given a preference $\succsim_{\mathcal{S}}$ satisfying their Lemma 5 (which is a direct translation of our Lemma 3) and the Cancellation axiom, and a family $<_a$ of linear orders, one for each atom, they can construct a Lewis-style model that represents $\succsim_{\mathcal{S}}$. We outline the construction below, as well as defining formally what it means for a Lewis-style model to represent $\succsim_{\mathcal{S}}$.

The states in the Lewis-style model M are pairs (\vec{a}, \vec{a}') of atoms. Roughly speaking, since we are trying to capture the effect of interventions, we can think of \vec{a} as the current world (the result of performing the intervention) and \vec{a}' as the world before the intervention was performed. The truth of formulas in $\mathcal{L}(\mathcal{S})$ is completely determined by the first component. That is, $(M, (\vec{a}, \vec{a}')) \models Y = y$ if $Y = y$ is a conjunct of \vec{a} ; we extend to negation, conjunction, and negation in the standard way. To extend this semantics to $\mathcal{L}^+(\mathcal{S})$, we need a closeness relation for each pair (\vec{a}, \vec{a}') . Let $\sqsubset_{\vec{a}, \vec{a}'}$ be an arbitrary family of strict linear orders on $\mathcal{A}(\mathcal{S}) \times \mathcal{A}(\mathcal{S})$ such that

$$\langle \vec{b}, \vec{b}' \rangle \sqsubset_{\langle \vec{a}, \vec{a}' \rangle} \langle \vec{c}, \vec{c}' \rangle \text{ whenever } \vec{b} <_{\vec{a}} \vec{c}$$

and

$$\langle \vec{b}, \vec{a} \rangle \sqsubset_{\langle \vec{a}, \vec{a}' \rangle} \langle \vec{b}, \vec{b}' \rangle \text{ for all } \vec{b}' \neq \vec{a}$$

(so that $\langle \vec{a}, \vec{a} \rangle$ is the minimal element of $\sqsubset_{\langle \vec{a}, \vec{a}' \rangle}$).⁵ This completes the description of M .

Let $\text{MIN}_{\langle \vec{a}, \vec{a}' \rangle}(\text{do}[\vec{Y} \leftarrow \vec{y}]) \in \mathcal{A}(\mathcal{S}) \times \mathcal{A}(\mathcal{S})$ denote the $\sqsubset_{\langle \vec{a}, \vec{a}' \rangle}$ -minimal state such that $(M, \text{MIN}_{\langle \vec{a}, \vec{a}' \rangle}(\text{do}[\vec{Y} \leftarrow \vec{y}])) \models \vec{Y} \leftarrow \vec{y}$. Thus, $(M, \langle \vec{a}, \vec{a}' \rangle) \models \text{do}[\vec{Y} \leftarrow \vec{y}] \phi$ iff $(M, \text{MIN}_{\langle \vec{a}, \vec{a}' \rangle}(\text{do}[\vec{Y} \leftarrow \vec{y}])) \models \phi$. It is easy to check that

$$\text{MIN}_{\langle \vec{a}, \vec{a}' \rangle}(\text{do}[\vec{Y} \leftarrow \vec{y}]) = \langle \vec{a}_{\vec{Y} \leftarrow \vec{y}}, \vec{a} \rangle. \quad (6)$$

Thus, the first component of the minimal state that results from performing the intervention $\vec{Y} \leftarrow \vec{y}$ in a state of the form $\langle \vec{a}, \vec{a}' \rangle$ encodes the atom that results when the intervention is performed starting at \vec{a} , while the second component keeps track of the atom we started with.

BH define an analogue of our function h_A , which we denote h'_A , that associates with each of their actions an action

⁵For any strict linear order \prec on $\mathcal{A}(\mathcal{S})$ whose minimal element is \vec{a} , the lexicographic ordering on $\mathcal{A}(\mathcal{S}) \times \mathcal{A}(\mathcal{S})$ with $<_{\vec{a}}$ on the first component and \prec on the second component will suffice.

of the form $\mathbf{do}[\vec{Y} \leftarrow \vec{y}]$.⁶ For a state ω , the analogue of $\beta_A^M(\omega)$ is $\text{MIN}_\omega(h'_A(\omega))$.⁷ BH show that there exists a probability measure \mathbf{p} and utility function \mathbf{u} , both defined on the states of M , such that $A \succ_S B$ iff

$$\begin{aligned} & \sum_{\substack{\langle \vec{a}, \vec{a}' \rangle \in \\ \mathcal{A}(\mathcal{S}) \times \mathcal{A}(\mathcal{S})}} \mathbf{p}(\langle \vec{a}, \vec{a}' \rangle) \cdot \mathbf{u}(\text{MIN}_{\langle \vec{a}, \vec{a}' \rangle}(h'_A(\langle \vec{a}, \vec{a}' \rangle))) \\ \geq & \sum_{\substack{\langle \vec{a}, \vec{a}' \rangle \in \\ \mathcal{A}(\mathcal{S}) \times \mathcal{A}(\mathcal{S})}} \mathbf{p}(\langle \vec{a}, \vec{a}' \rangle) \cdot \mathbf{u}(\text{MIN}_{\langle \vec{a}, \vec{a}' \rangle}(h'_B(\langle \vec{a}, \vec{a}' \rangle))). \end{aligned} \quad (7)$$

This is the sense in which the Lewis-style model M represents \succ_S .

Since BH's Lemma 5 is a direct translation of our Lemma 3, it follows from our axioms as well that this Lewis-style model M represents \succ_S . But we are not done; we need to get a causal model that represents \succ_S . So we want to convert the Lewis-style model constructed by BH to a causal model $M^C = (\mathcal{S}^C, \mathcal{F}^C)$, and construct an appropriate probability \mathbf{p}^C on the contexts of M , and utility \mathbf{u}^C on $\mathcal{A}(\mathcal{S})$.

The first step is easy. We start with a signature \mathcal{S} that determines the language; we take $\mathcal{S}^C = \mathcal{S}$. By Axiom A2, it follows that for each context $\vec{u} \in \mathcal{C}(\mathcal{S})$, there at most one atom such that $\vec{a} \Rightarrow (U = \vec{u})$ and \vec{a} is non-null. Let \mathcal{C}^\dagger denote the set of contexts for which such a non-null atom exists, and for each context $\vec{u} \in \mathcal{C}^\dagger$, let $\vec{a}_{\vec{u}}$ denote this non-null atom. \mathcal{C}^\dagger is non-empty, as we have assumed \succ_S is non-trivial. We define \mathbf{p}^C using the probability \mathbf{p} in the model M provided by BH so that it has support \mathcal{C}^\dagger :

$$\mathbf{p}^C(\vec{u}) = \frac{\sum_{\vec{a}' \in \mathcal{A}(\mathcal{S})} \mathbf{p}(\langle \vec{a}_{\vec{u}}, \vec{a}' \rangle)}{\sum_{\vec{u}' \in \mathcal{C}^\dagger} \sum_{\vec{a}' \in \mathcal{A}(\mathcal{S})} \mathbf{p}(\langle \vec{a}_{\vec{u}'}, \vec{a}' \rangle)}.$$

Observe that, since the truth of formulas in M at state $\langle \vec{a}, \vec{a}' \rangle$ is fully determined by \vec{a} , $h'_A(\langle \vec{a}, \vec{a}' \rangle)$ does not depend on \vec{a}' . In particular, $h'_A(\langle \vec{a}, \vec{a}' \rangle) = h_A(\vec{a})$. Similarly, by (6), $\text{MIN}_{\langle \vec{a}, \vec{a}' \rangle}(\cdot)$ does not depend on \vec{a}' . We can therefore rewrite the BH representation (7) as

$$\begin{aligned} & \sum_{\vec{u} \in \mathcal{C}^\dagger} \mathbf{p}^C(\vec{u}) \cdot \mathbf{u}(\text{MIN}_{\langle \vec{a}_{\vec{u}}, \vec{a}_{\vec{u}} \rangle}(h_A(\vec{a}_{\vec{u}}))) \\ \geq & \sum_{\vec{u} \in \mathcal{C}^\dagger} \mathbf{p}^C(\vec{u}) \cdot \mathbf{u}(\text{MIN}_{\langle \vec{a}_{\vec{u}}, \vec{a}_{\vec{u}} \rangle}(h_B(\vec{a}_{\vec{u}}))). \end{aligned} \quad (8)$$

We next want to show that there exists a set of equations \mathcal{F}^C such that $M^C = (\mathcal{S}^C, \mathcal{F}^C)$ satisfies the same formulas as M . Specifically:

Lemma 4. *There exists a set of equations \mathcal{F}^C such that for all $\phi \in \mathcal{L}(\mathcal{S})$ and all interventions $\vec{Y} \leftarrow \vec{y}$, we have that*

$$\begin{aligned} & (M_{\mathbf{do}[\vec{Y} \leftarrow \vec{y}]}^C, \vec{u}) \models \phi \quad \text{iff} \\ & (M, \text{MIN}_{\langle \vec{a}_{\vec{u}}, \vec{a}_{\vec{u}} \rangle}(\mathbf{do}[\vec{Y} \leftarrow \vec{y}])) \models \phi. \end{aligned} \quad (9)$$

Note that taking $\vec{Y} = \emptyset$ gives a special case of (9): $(M^C, \vec{u}) \models \phi$ iff $(M, \langle \vec{a}_{\vec{u}}, \vec{a}_{\vec{u}} \rangle) \models \phi$.

⁶BH allow for (but do not require) a richer set of interventions, they allow $\mathbf{do}[\phi]$ for any consistent formula ϕ in the language.

⁷The states for us are contexts, whereas for BH they are pairs of atoms, so using the same symbol ω for states is somewhat of an abuse of notation; we hope that our intention is clear here.

Proof. Let Y_1, \dots, Y_k be the ordering on endogenous variables according to \sim . We define \mathcal{F}^C for the endogenous variables by induction on j , so that (9) holds for all formulas $Y_j = y_j$.

For the base case, note that the value of the variable Y_1 does not depend on the values of any other endogenous variable. It is straightforward to define F_{Y_1} so that $F_{Y_1}(\vec{u}) = y_1$ iff $a_{\vec{u}} \Rightarrow Y_1 = y_1$. It easily follows that $(M^C, \vec{u}) \models Y_1 = y_1$ iff $(M, \langle \vec{a}_{\vec{u}}, \vec{a}_{\vec{u}} \rangle) \models Y_1 = y_1$. Moreover, if $Y_1 \notin \vec{Y}$, then $(M_{\mathbf{do}[\vec{Y} \leftarrow \vec{y}]}^C, \vec{u}) \models Y_1 = y_1$ iff $(M^C, \vec{u}) \models Y_1 = y_1$. By Lemma 2 (taking $\vec{Z} = \emptyset$ and $X = Y_1$) it follows that $(M, \text{MIN}_{\langle \vec{a}_{\vec{u}}, \vec{a}_{\vec{u}} \rangle}(\mathbf{do}[\vec{Y} \leftarrow \vec{y}])) \models Y_1 = y_1$ iff $(M, \langle \vec{a}_{\vec{u}}, \vec{a}_{\vec{u}} \rangle) \models Y_1 = y_1$. On the other hand, if $Y_1 \in \vec{Y}$, then $(M_{\mathbf{do}[\vec{Y} \leftarrow \vec{y}]}^C, \vec{u}) \models Y_1 = y_1$ iff $\vec{y}|_{Y_1} = y_1$ and similarly, $(M, \text{MIN}_{\langle \vec{a}_{\vec{u}}, \vec{a}_{\vec{u}} \rangle}(\mathbf{do}[\vec{Y} \leftarrow \vec{y}])) \models Y_1 = y_1$ iff $\vec{y}|_{Y_1} = y_1$. Thus, we get the desired result in the base case.

For the inductive case, suppose that we have proved the result for Y_1, \dots, Y_j . We can define $F_{Y_{j+1}}$ to be function of the exogenous variables and Y_1, \dots, Y_j such that for all $k \in \{1, \dots, j\}$, and $y_i \in \mathcal{R}(Y_i)$ for $i = 1, \dots, j$, we have that $F_{Y_{j+1}}(\vec{u}, y_1, \dots, y_j) = y_{k+1}$ iff $(M, \text{MIN}_{\langle \vec{a}_{\vec{u}}, \vec{a}_{\vec{u}} \rangle}(\mathbf{do}[Y_1 \leftarrow y_1, \dots, Y_k \leftarrow y_k])) \models Y_{k+1} = y_{k+1}$. Now for arbitrary \vec{Y} , if $Y_{j+1} \notin \vec{Y}$, let $\vec{Y}' = \vec{Y} \cap \{Y_1, \dots, Y_j\}$. Let $\vec{y}' = \vec{y}|_{\vec{Y}'}$. Writing \vec{Y}'' for $\{Y_1, \dots, Y_j\}$, let \vec{y}'' be such that $(M_{\mathbf{do}[\vec{Y}' \leftarrow \vec{y}']}^C, \vec{u}) \models \vec{Y}'' = \vec{y}''$. Then it follows from the induction hypothesis that $(M, \text{MIN}_{\langle \vec{a}_{\vec{u}}, \vec{a}_{\vec{u}} \rangle}(\mathbf{do}[\vec{Y}' \leftarrow \vec{y}'])) \models \vec{Y}'' = \vec{y}''$. Moreover, it easily follows that $(M_{\mathbf{do}[\vec{Y} \leftarrow \vec{y}]}^C, \vec{u}) \models Y_{j+1} = y_{j+1}$ iff $(M_{\mathbf{do}[\vec{Y}' \leftarrow \vec{y}']}^C, \vec{u}) \models Y_{j+1} = y_{j+1}$ and $(M, \text{MIN}_{\langle \vec{a}_{\vec{u}}, \vec{a}_{\vec{u}} \rangle}(\mathbf{do}[\vec{Y} \leftarrow \vec{y}])) \models Y_{j+1} = y_{j+1}$ iff $(M, \text{MIN}_{\langle \vec{a}_{\vec{u}}, \vec{a}_{\vec{u}} \rangle}(\mathbf{do}[\vec{Y}'' \leftarrow \vec{y}''])) \models Y_{j+1} = y_{j+1}$. Since, by the definition of $F_{Y_{j+1}}$, $(M_{\mathbf{do}[\vec{Y}' \leftarrow \vec{y}']}^C, \vec{u}) \models Y_{j+1} = y_{j+1}$ iff $(M, \text{MIN}_{\langle \vec{a}_{\vec{u}}, \vec{a}_{\vec{u}} \rangle}(\mathbf{do}[\vec{Y}'' \leftarrow \vec{y}''])) \models Y_{j+1} = y_{j+1}$, the result easily follows. The argument if $Y_{j+1} \in \vec{Y}$ is the same as in the base case.

The result for arbitrary formulas $\phi \in \mathcal{L}(\mathcal{S})$ is immediate (since conjunction, disjunction, and negation work the same way in both M and M^C). \square

From (6), it follows that $\text{MIN}_{\langle \vec{a}_{\vec{u}}, \vec{a}_{\vec{u}} \rangle}(\mathbf{do}[\vec{Y} \leftarrow \vec{y}]) = \langle \langle \vec{a}_{\vec{u}} \rangle_{\mathbf{do}[\vec{Y} \leftarrow \vec{y}]}, \vec{a} \rangle$; moreover, $(M, \langle \langle \vec{a}_{\vec{u}} \rangle_{\mathbf{do}[\vec{Y} \leftarrow \vec{y}]}, \vec{a} \rangle) \models \langle \vec{a}_{\vec{u}} \rangle_{\mathbf{do}[\vec{Y} \leftarrow \vec{y}]}$. Thus, by Lemma 4, $(M_{\mathbf{do}[\vec{Y} \leftarrow \vec{y}]}^C, \vec{u}) \models \langle \vec{a}_{\vec{u}} \rangle_{\mathbf{do}[\vec{Y} \leftarrow \vec{y}]}$. Moreover, since $(M^c, \vec{u}) \models \vec{a}_{\vec{u}}$, that is, $\vec{a}_{M^c, \vec{u}} = \vec{a}_{\vec{u}}$, we have

$$\beta_A^{M^C}(\vec{u}) = \langle \vec{a}_{\vec{u}} \rangle_{h_A(\vec{a}_{\vec{u}})}. \quad (10)$$

We can now complete the proof that M^C represents the preference order. The utility function \mathbf{u} given by BH is defined on their states, which are pairs of atoms. We define \mathbf{u}^C over individual atoms as

$$\mathbf{u}^C(\vec{a}) = u(\langle \vec{a}, \vec{a}_{\vec{a}|\mathcal{U}} \rangle) \quad (11)$$

if $\vec{a}|\mathcal{U} \in \mathcal{C}^\dagger$ (and define it arbitrarily otherwise). Unpacking this, $\vec{a}|\mathcal{U}$ is the context of the atom \vec{a} , and so, $\vec{a}_{\vec{a}|\mathcal{U}}$ is the

unique non-null atom with the same context as \vec{a} .

$$\begin{aligned} \mathbf{u}(\text{MIN}_{\langle \vec{a}_{\vec{u}}, \vec{a}_{\vec{u}} \rangle}(h_A(\vec{a}_{\vec{u}}))) &= \mathbf{u}(\langle (\vec{a}_{\vec{u}})_{h_A(\vec{a}_{\vec{u}})}, \vec{a}_{\vec{u}} \rangle) && \text{(from (6))} \\ &= \mathbf{u}^C(\langle (\vec{a}_{\vec{u}})_{h_A(\vec{a}_{\vec{u}})} \rangle) && \text{(from (11))} \\ &= \mathbf{u}^C(\beta_A^M(\vec{u})) && \text{(from (10))} \end{aligned}$$

Substituting \mathbf{u}^C for \mathbf{u} in (8) gives the desired representation.

To prove Theorem 2, the uniqueness claim, let $(M, \mathbf{p}, \mathbf{u})$ and $(M', \mathbf{p}', \mathbf{u}')$ be two representations of $\succsim_{\mathcal{S}}$. For all $\vec{u} \in \mathcal{C}^\dagger$ and all vectors $\vec{Y}, \vec{y} \in \mathcal{R}(\vec{Y})$, endogenous variables $X \notin \vec{Y}$, and $x \in \mathcal{R}(X)$

$$(M_{\text{do}[\vec{Y} \leftarrow \vec{y}]}, \vec{u}) \models (X = x) \implies \text{do}[\vec{Y} \leftarrow \vec{y}] \sim_{\vec{a}_{\vec{u}}} (X = x); \quad (12)$$

similarly,

$$(M'_{\text{do}[\vec{Y} \leftarrow \vec{y}]}, \vec{u}') \models (X = x) \implies \text{do}[\vec{Y} \leftarrow \vec{y}] \sim_{\vec{a}_{\vec{u}'}} (X = x). \quad (13)$$

Suppose that $\succsim_{\mathcal{S}}$ satisfies A3*. Then there is a unique $x \in \mathcal{R}(X)$ that satisfies the right-hand relations of (12) and (13), so we have

$$(M_{\text{do}[\vec{Y} \leftarrow \vec{y}]}, \vec{u}) \models (X = x) \iff (M'_{\text{do}[\vec{Y} \leftarrow \vec{y}]}, \vec{u}') \models (X = x).$$

It easily follows that M and M' agree on all formulas in $\mathcal{L}^+(\mathcal{S})$, so M is identifiable.

For the converse, suppose that $M \neq M'$ for some $\vec{u} \in \mathcal{C}^\dagger$. This requires that there exist some vector $\vec{Y}, \vec{y} \in \mathcal{R}(\vec{Y})$, and an endogenous variable $X \notin \vec{Y}$ such that $(M_{\text{do}[\vec{Y} \leftarrow \vec{y}]}, \vec{u}) \models (X = x)$ and $(M'_{\text{do}[\vec{Y} \leftarrow \vec{y}]}, \vec{u}') \models (X = x')$ and $x \neq x'$. But then, by (12) and (13), $\succsim_{\mathcal{S}}$ violates A3*.

6 Discussion

We have given a representation theorem in the spirit of Savage (1954) that helps us understand a decision maker's (subjective) causal judgements: If a decision maker's preferences among actions that involve interventions satisfy certain axioms, then we can find a causal model M , a probability over contexts in M , and a utility on states in M such that the decision maker prefers action A to B iff A has higher expected utility than B . Moreover, we have shown if we add another axiom, then M is unique. Our approach builds on earlier work by BH, who proved an analogous representation theorem using Lewis-style models. Interestingly, other than the Cancellation axiom (which is due to BEH), the axioms used by BH are completely different from ours. For example, they do not define an analogue of our $\sim_{\vec{a}}$ relation, which plays a critical role in our definiteness and centeredness axioms. Roughly speaking, the BH axioms build on axioms for counterfactuals used by Lewis, while ours build on axioms for causal models introduced by Galles and Pearl (1998) and Halpern (2001). In future work, we hope to better understand the connection between the axioms (e.g., to what extent could we use the BH axioms as a basis for a representation theorem for causal models).

Our representation assumes that the agent knows the model M . As long as the set of exogenous variables is flexible, this choice is without loss of generality. Uncertainty

about the model can be reduced to uncertainty about a new exogenous variable U^* . In the larger model that includes U^* , U^* acts as a “switch”, so that each of its values selects one of the models the agent considered possible in the uncertain representation. For example, if the Fed is uncertain as to whether the interest rate actually affects inflation, then we could add a variable U^* that governs this relationship. The structural equation

$$X_{inf} = 1 - (Y_{rate} \times U^*) \quad (F_{X_{inf}})$$

then captures the Fed's uncertainty: if $U^* = 1$, the interest rate mitigates inflation; if $U^* = 0$, it has no effect. That said, this technical workaround may be unsuited for applications where the set of variables is fixed beforehand. We conjecture that our axiomatization, without Axiom 2 (Model Uniqueness), will deliver a generalization that allows for model uncertainty, but our proof strategy would have to be changed.

Here we have considered only one-step decisions; that is, the decision-maker performs an intervention, perhaps conditional on some test. It is clearly also of interest to consider sequential decisions. In practice, plans are composed of a sequence of steps; later interventions might depend on earlier interventions. This leads us to consider a richer set of actions such as **if** ϕ_1 **then** A_1 **else** B_1 ; **if** ϕ_2 **then** A_2 **else** B_2 , where the second action (**if** ϕ_2 **then** A_2 **else** B_2) is performed after the first. Getting a representation theorem for Lewis-style models in the presence of sequential actions seems significantly more complicated than in the case of one-step actions (Bjorndahl and Halpern 2023). It would be of interest to see what is involved in getting such a representation theorem using causal models.

Interestingly, we expect there to be significant differences between causal models and Lewis-style once we have sequential decisions. In causal models, a decision maker's preferences regarding sequences of primitive actions are invariant with respect to order (although this is not necessarily the case for conditional interventions). This is because the interventions $\text{do}[\vec{Y} \leftarrow \vec{y}, \vec{Z} \leftarrow \vec{z}]$ and $\text{do}[\vec{Z} \leftarrow \vec{z}, \vec{Y} \leftarrow \vec{y}]$ are identical. By way of contrast, the order in which primitive interventions are performed can have a significant impact in Lewis-style models. In such models, the closest-world operator is local; loosely speaking, the “distance” between two worlds depends on the world at which they are being contemplated. Because of this, intervening to make ϕ true—moving to the closest ϕ -world—also changes the relative closeness of other worlds; the closest $(\phi \wedge \psi)$ -world may bear no relation to the closest ψ -world to the closest ϕ -world.

In future work, we plan to explore the issue of order dependence in both causal models and Lewis-style models, both because we believe that the issue of importance in its own right, and because it will help elucidate the differences between these two approaches to modeling causality. We also plan to explore multi-agent casual reasoning. This introduces many subtleties: reasoning about other agent's models, of course, but also the ability to cause a change in other agent's beliefs (i.e., by providing evidence).

Acknowledgments

Halpern's work is supported in part by NSF grant FMITF-2319186, ARO grant W911NF-17-1-0061, MURI grant W911NF-19-1-0217 from the ARO, and a grant from Open Philanthropy.

References

- Alexander, Y., and Gilboa, I. 2023. Subjective causality. *Revue économique* 74(4):619–633.
- Angrist, J. D., and Pischke, J.-S. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.
- Bjorndahl, A., and Halpern, J. Y. 2021. Language-based decisions. In *Theoretical Aspects of Rationality and Knowledge: Proc. Eighteenth Conference (TARK 2021)*. The proceedings are published as Electronic Proceedings in Theoretical Computer Science.
- Bjorndahl, A., and Halpern, J. Y. 2023. Sequential language-based decisions. In *Theoretical Aspects of Rationality and Knowledge: Proc. Nineteenth Conference (TARK 2023)*. The proceedings are published as Electronic Proceedings in Theoretical Computer Science.
- Blume, L. E.; Easley, D.; and Halpern, J. Y. 2021. Constructive decision theory. *Journal of Economic Theory* 196. An earlier version, entitled “Redoing the Foundations of Decision Theory”, appears in *Principles of Knowledge Representation and Reasoning: Proc. Tenth International Conference (KR '06)*.
- Cunningham, S. 2021. *Causal Inference: The Mixtape*. Yale University Press.
- Ellis, A., and Thysen, H. C. 2021. Subjective causality in choice. Available at <http://arxiv.org/abs/2106.05957>.
- Galles, D., and Pearl, J. 1998. An axiomatic characterization of causal counterfactuals. *Foundation of Science* 3(1):151–182.
- Haavelmo, T. 1943. The statistical implications of a system of simultaneous equations. *Econometrica* 11:1–12.
- Halpern, J. Y. 2000. Axiomatizing causal reasoning. *Journal of A.I. Research* 12:317–337.
- Halpern, J. Y. 2001. Substantive rationality and backward induction. *Games and Economic Behavior* 37:425–435.
- Halpern, J. Y. 2013. From causal models to possible-worlds models of counterfactuals. *Review of Symbolic Logic* 6(2):81–101.
- Hernán, M. A., and Robins, J. M. 2020. *Causal inference: What If*. Chapman and Hall/CRC.
- Lewis, D. 1973. *Counterfactuals*. Cambridge, MA: Harvard University Press.
- Machina, M., and Schmeidler, D. 1992. A more robust definition of subjective probability. *Econometrica* 60(2):745–780.
- Morgan, S., and Winship, C. 2007. *Counterfactuals and Causal Inference*. Cambridge University Press.
- Parascandola, M., and Weed, D. L. 2001. Causation in epidemiology. *Journal of Epidemiology & Community Health* 55(12):905–912.
- Pearl, J. 2000. *Causality: Models, Reasoning, and Inference*. New York: Cambridge University Press.
- Pearl, J. 2009. Causal inference in statistics. *Statistics Surveys* 3:96–146.
- Plowright, R. K.; Sokolow, S. H.; Gorman, M. E.; Daszak, P.; and Foley, J. E. 2008. Causal inference in disease ecology: investigating ecological drivers of disease emergence. *Frontiers in Ecology and the Environment* 6(8):420–429.
- Savage, L. J. 1954. *Foundations of Statistics*. New York: Wiley.
- Schenone, P. 2020. Causality: a decision theoretic foundation. Available at <http://arxiv.org/abs/1812.07414>.
- Simon, H. A. 1953. Causal ordering and identifiability. In Hood, W. C., and Koopmans, T. C., eds., *Studies in Econometric Methods*, Cowles Commission for Research in Economics, Monograph No. 14. New York: Wiley. 49–74.
- Spirites, P.; Glymour, C.; and Scheines, R. 1993. *Causation, Prediction, and Search*. New York: Springer-Verlag.
- Stalnaker, R. C. 1968. A theory of conditionals. In Rescher, N., ed., *Studies in Logical Theory*. Oxford, U.K.: Blackwell. 98–112.