

IS CAUSAL REASONING HARDER THAN PROBABILISTIC REASONING?

MILAN MOSSÉ

University of California, Berkeley

DULIGUR IBELING

Stanford University
and

THOMAS ICARD

Stanford University

Abstract. Many tasks in statistical and causal inference can be construed as problems of *entailment* in a suitable formal language. We ask whether those problems are more difficult, from a computational perspective, for *causal* probabilistic languages than for pure probabilistic (or “associational”) languages. Despite several senses in which causal reasoning is indeed more complex—both expressively and inferentially—we show that causal entailment (or satisfiability) problems can be systematically and robustly reduced to purely probabilistic problems. Thus there is no jump in computational complexity. Along the way we answer several open problems concerning the complexity of well-known probability logics, in particular demonstrating the $\exists\mathbb{R}$ -completeness of a polynomial probability calculus, as well as a seemingly much simpler system, the logic of comparative conditional probability.

§1. Motivation and preview. There is an uncontroversial sense in which causal reasoning is more difficult than purely probabilistic or statistical reasoning. The latter seems hard enough: estimating probabilities, predicting future events from past observations, determining statistical significance, and adjudicating between statistical hypotheses—these are already formidable tasks, long mired in controversy. No free lunch theorems [6, 41] show that strong assumptions are necessary to gain any inductive purchase on such problems, and there is considerable disagreement about what kinds of assumptions are reasonable in different epistemic and practical circumstances [15]. Problems of causal inference only seem to make our tasks harder. Inferring causal effects, predicting the outcomes of interventions, determining causal direction, and learning a causal model—these problems typically demand statistical reasoning, but they also demand more on the part of the investigator. They may require that we actively interrogate the world through deliberate experimentation rather than passive observation, or that we antecedently accept strong assumptions sufficient to justify the causal conclusions we want to reach, or (very often) both. Indeed, statistical indistinguishability is the norm in causal inference, even with substantive assumptions [43]. As formalized in the causal hierarchy theorem of Bareinboim et al. [5] (see also [26]), it is not only impossible to infer causal information from purely correlational

Received: September 4, 2021.

2020 *Mathematics Subject Classification*: Primary 03A05, Secondary 68Q25.

Keywords and phrases: probability, causation, logic, complexity, ETR.




(or “observational”) data, but also generically impossible to infer counterfactual or explanatory information from purely experimental (or “interventional”) data. From an inferential perspective, probabilistic information vastly underdetermines causal information.

A feature common to both statistical inference and causal inference is that the most prominent approaches to each can be understood, at least in part, as attempts to turn an inductive problem into a deductive one. This is famously true of frequentist methods in the tradition associated with Neyman and Pearson (see [31]), but is arguably true of Bayesian approaches as well. As Gelman and Shalizi [20] suggest, “Statistical models are tools that let us draw inductive inferences on a deductive background,” rendering statistical inferences “deductively guaranteed by probabilistic assumptions” (p. 27). Indeed, one of the benefits of specifying a Bayesian probability model is that it provides an answer to virtually any question about the probability of a hypothesis conditional on data. Given the model and the data, this answer follows as a matter of logic.

Causal underdetermination is likewise confronted with methods for formulating precise inductive assumptions, sometimes allowing answers to causal questions to be derived by mere calculation.

EXAMPLE 1.1 (Do-calculus). *As one prominent example, the do-calculus of Pearl and collaborators (see [33] and Chapter 3 of [34]) establishes systematic correspondences between qualitative (“graphical”) properties of a causal scenario and certain conditional independence statements involving causal quantities. A typical causal quantity of interest is the (average) causal effect, e.g., how likely Y is to take on value y given an intervention setting X to x . In a formal language (introduced in the sequel as $\mathcal{L}_{\text{causal}}$), we write this as $\mathbf{P}([X = x]Y = y)$, or more briefly, $\mathbf{P}([x]y)$.*

Absent assumptions, it is never possible to infer the value of $\mathbf{P}([x]y)$ from observational data [5]. Suppose, however, that we could assume the causal structure has something like the following shape (known in the literature as the front door graph):



For a standard example, we might assume that any causal effect of smoking (X) on cancer (Y) will be mediated by tar deposited in the lungs (Z), and moreover that any unknown sources of variation (U) on X or on Y (or on both), such as a person’s genotype, do not directly influence Z . Under these circumstances, the do-calculus licenses several substantive causal assumptions, which may be rendered precisely in $\mathcal{L}_{\text{causal}}$. Let Γ be the set of equality statements below:

- (i) $\mathbf{P}([x]z) = \mathbf{P}(z|x)$,
- (ii) $\mathbf{P}([z]x) = \mathbf{P}(x)$,
- (iii) $\mathbf{P}([x]y|[x]z) = \mathbf{P}([x, z]y) = \mathbf{P}([z]y)$,
- (iv) $\mathbf{P}([z]y|[z]x) = \mathbf{P}(y|x, z)$.

For instance, (i) says that the causal effect of $X = x$ on $Z = z$ simply coincides with the conditional probability $\mathbf{P}(Z = z|X = x)$. Appealing to a combination of laws of probability and distinctively causal laws involving the “causal-conditional” statements like $[x]y$, it is possible to show that the following equality is in fact entailed by the statements Γ , that is, by (i)–(iv):

$$\mathbf{P}([x]y) = \sum_z \mathbf{P}(z|x) \sum_{x'} \mathbf{P}(y|x', z) \mathbf{P}(x'). \quad (1)$$

In other words, (1) shows that the causal effect of $X = x$ on $Y = y$ can simply be calculated from suitable observational data involving the variables X, Y, Z .

Methods such as these extend beyond the specific problem of estimating causal effects, to include estimation of *counterfactual* quantities as well. For instance, we may want to determine—from experimental data and background assumptions—the joint probability that an individual would survive if *and only if* they are assigned a certain treatment, a quantity we would write as $\mathbf{P}([X = 1]Y = 1 \wedge [X = 0]Y = 0)$. Inferential techniques similar to those in Example 1.1 have been employed in such settings, and have even been automated (e.g., [14]).

More broadly, a number of different approaches to inductive inference, both statistical and causal, can be assimilated to a regiment something like this:

$$\text{Inductive Assumptions} + \text{Data} \models \text{Inferential Conclusion}. \quad (2)$$

In Example 1.1, Γ are the inductive assumptions, the data would be information about $\mathbf{P}(X, Y, Z)$, and the conclusion would be an estimate of the causal effect of $X = x$ on $Y = y$. In a standard Bayesian analysis, the inductive assumption might be a prior probability model for some latent variables (e.g., parameters for a class of probability measures), while the data would be values of some observable variables, and the conclusion might be the posterior values for the hidden variables, or perhaps posterior predictive values for some yet-to-be-observed variables. A critical job of the statistician or data scientist is to identify suitable inductive assumptions that a relevant party judges reasonable (or, ideally if feasible, which are themselves empirically verifiable) and that are sufficiently strong to license meaningful conclusions from the types of data available.

From this vantage point our titular question takes on a new significance. Rather than asking about the difficulty of an inference task in terms of the *strength of assumptions* needed to justify the inference, we could instead ask how difficult it is in general, *computationally speaking*, to reason from inductive assumptions (together with data) to an inferential conclusion, in the strong sense of (2). In other words, we ask how difficult questions like (2) could be across different logical languages for describing relevant assumptions, data, and conclusions.

The contrast of interest in this article is between languages $\mathcal{L}_{\text{prob}}$, suitable for probabilistic reasoning, and languages $\mathcal{L}_{\text{causal}}$, which extend the corresponding probabilistic languages to encompass causal reasoning in addition. In short, $\mathcal{L}_{\text{prob}}$ encompasses “pure” probabilistic reasoning about some set of random variables. In $\mathcal{L}_{\text{causal}}$ we also reason about the probabilities of *causal conditionals*, the causal effect $\mathbf{P}([x]y)$ being a simple example. Such mixed reasoning is crucial for applications like the do-calculus, where causal conclusions depend on distinctively causal assumptions (such as (i)–(iv) in Example 1.1). Some of the emblematic principles of $\mathcal{L}_{\text{causal}}$ reveal

a subtle interplay between the probabilistic and causal-conditional components. For example, the following formula states that if causal interventions which set the values of the variable X thereby affect the values taken the variable Y , then the converse cannot be true:

$$\mathbf{P}([x]_y \wedge [x']_{y'}) > 0 \rightarrow \mathbf{P}([y]_x \wedge [y']_{x'}) = 0. \quad (3)$$

This formula emerges as an instance of a more general scheme in a complete axiomatization of $\mathcal{L}_{\text{causal}}$ (see [25]), implying that X and Y cannot each causally affect the other.

In light of the considerable *empirical* (and expressive) gulf between these two kinds of languages, we might expect to see a parallel jump in computational complexity when moving from $\mathcal{L}_{\text{prob}}$ to $\mathcal{L}_{\text{causal}}$. In a certain respect, $\mathcal{L}_{\text{causal}}$ can be seen as a *combination* of logics, embedding one modal system (a conditional logic) inside another (a probability logic), with non-trivial interactions between the two (such as (3)). It is common wisdom that such combinations may in general drive up complexity, in some cases even resulting in undecidability (see, e.g., [28]). As a famous example, even seemingly innocuous combinations of modalities for knowledge and time (each independently of low complexity) can lead to Π_1^1 -hardness [21]. The present work introduces two main results, which show that this does not happen here: causal reasoning and probabilistic reasoning are, in a precise and robust sense, equally difficult.

The distinction between $\mathcal{L}_{\text{prob}}$ and $\mathcal{L}_{\text{causal}}$ is orthogonal to another distinction, namely *how much arithmetic* we admit in our formal language of probability over a set of probability terms $\mathbf{P}(\delta)$. A wide range of probability logics have been studied in the literature, from pure qualitative comparisons between probability terms (e.g., [19]) to richer fragments capable of reasoning about polynomials over such terms (e.g., [40]). For any such choice $\mathcal{L}_{\text{prob}}$ of probabilistic language we can consider the extension $\mathcal{L}_{\text{causal}}$ to allow not only probability terms, but also causal-probability terms like those introduced above. A strength of our analysis is that we provide a complexity-reflecting reduction from $\mathcal{L}_{\text{causal}}$ to $\mathcal{L}_{\text{prob}}$ in a way that is independent of our choice of probabilistic primitives. Thus, across the landscape of probability logics, we see no increase in complexity. Summarizing, our main result states:

THEOREM 1.2 *Probabilistic reasoning is no harder than causal reasoning. In particular:*

1. *Reasoning about (causal or non-causal) probabilities is as hard as reasoning about sums of (causal or non-causal) probabilities; both are as hard as reasoning about Boolean formulas, or about sums of real numbers.*
2. *Reasoning about (causal or non-causal) conditional probabilities is as hard as reasoning about arbitrary polynomials in (causal or non-causal) probabilities; both are as hard as reasoning about arbitrary polynomials in real numbers.*

The above theorem is stated formally below as Theorem 3.3. While the relationship between probabilistic and causal languages is our main focus, it is worth pointing out that some of our results are of interest beyond the connection with causality. In particular, we find that reasoning in the language of conditional comparative probability is precisely as hard as reasoning in the full existential first-order theory of real numbers ($\exists\mathbb{R}$), thus establishing another notable example of a problem complete for this complexity class. It is also noteworthy that this expressively weak probabilistic language is—from a computational perspective—as complex as the most expressive *causal* languages we consider in the paper (namely, $\mathcal{L}_{\text{causal}}^{\text{poly}}$).

1.1. Relation to previous work. There is a long line of work on probability logic, including a host of results about complexity [1, 18, 32, 42]. As just mentioned, our contribution advances this literature. Concerning causal reasoning, there have been a number of complexity studies for various non-probabilistic causal notions [4, 16]. Most germane to the present study is Halpern’s [22] analysis of the satisfiability problem for deterministic reasoning about causal models, which he shows to be NP-complete (the same as propositional logical reasoning). Eiter and Lukasiewicz [16] studied numerous model-checking queries in a probabilistic setting, including the problem of determining the probability of a specific causal query. They show that this problem is complete for the class #P, the “counting analogue” to NP which also characterizes the problem of determining (approximations for) probabilities of (even very simple) propositional expressions [35].

Our interest in the present contribution is the complexity of reasoning—viz. testing for satisfiability, validity, or entailment, as portrayed in (2)—for probabilistic and causal languages. While this angle has not yet been explored thoroughly in the literature, our study is indebted to, and draws upon, much of this previous work. Theorem 1.2 synthesizes as well as greatly extends a heretofore piecemeal line of results [18, 24, 25]. Moreover, the results just mentioned by Halpern [22] and by Eiter and Lukasiewicz [16]—see also [13]—could be said to lend further support to the claim that causal reasoning is no more difficult (in the sense of computational complexity) than purely probabilistic reasoning.

1.2. Overview of the paper. In the next two sections (Sections 2 and 3), we introduce the languages and the notions from computational complexity needed to state Theorem 1.2 more formally. The proof of this main result appears in Section 4. Finally, in Section 5 we zoom out to consider what our results show about the relationship between probabilistic and causal reasoning, as well as consider a number of outstanding problems in this domain. In our presentation we assume no prior knowledge of causal modeling, complexity theory, or probability logic. Only elementary logic and probability are presupposed.

§2. Introducing causal and probabilistic languages. In this section, we introduce the syntax and semantics for a series of probabilistic and causal languages. With a precise syntax and semantics in hand, we illustrate that these languages form an expressive hierarchy.

2.1. Syntax. Let \mathbf{V} be a (possibly infinite) collection, representing the (endogenous) random variables under consideration. Informally, these are the variables that we may want to observe, change, query, or otherwise reason about explicitly.

For each variable $V \in \mathbf{V}$, let $\text{Val}(V)$ denote the finite signature (range) of V . For example, for two binary variables we have $\mathbf{V} = \{X, Y\}$ with $\text{Val}(X) = \text{Val}(Y) = \{0, 1\}$. We introduce the following deterministic languages:

$$\begin{aligned} \mathcal{L}_{\text{int}} &:= \top \mid V = v \mid \mathcal{L}_{\text{int}} \wedge \mathcal{L}_{\text{int}}, & V \in \mathbf{V}, v \in \text{Val}(V), \\ \mathcal{L}_{\text{prop}} &:= V = v \mid \neg \mathcal{L}_{\text{prop}} \mid \mathcal{L}_{\text{prop}} \wedge \mathcal{L}_{\text{prop}}, & V \in \mathbf{V}, v \in \text{Val}(V), \\ \mathcal{L}_{\text{full}} &:= [\mathcal{L}_{\text{int}}] \mathcal{L}_{\text{prop}} \mid \neg \mathcal{L}_{\text{full}} \mid \mathcal{L}_{\text{full}} \wedge \mathcal{L}_{\text{full}}. \end{aligned}$$

Choose either $\mathcal{L}_{\text{prop}}$ or $\mathcal{L}_{\text{full}}$ as the *base language* \mathcal{L} . The former is essentially a propositional language with extended ranges, while the latter is a causal conditional

language. The semantics of these formulas will be introduced in Section 2.2, but intuitively we can interpret a formula of $\mathcal{L}_{\text{full}}$, such as $[X = 1]Y = 0$, as expressing a subjunctive conditional: were X to take on value 1, then Y would come to have value 0. We understand the conditional causally, in a sense to be made precise below.

So-called *terms* over the base language are the main ingredient of our probabilistic languages. The most basic term is $\mathbf{P}(\delta)$ for $\delta \in \mathcal{L}$, representing the probability of δ . By varying the composite terms admitted, we can define polynomial, conditional, linear, and comparative languages. Where $\delta, \delta' \in \mathcal{L}$ are formulas of \mathcal{L} :

$$\begin{aligned} T_{\text{poly}}(\mathcal{L}) &\text{ is generated by the grammar } \mathbf{t} := \mathbf{P}(\delta) \mid \mathbf{t} + \mathbf{t}' \mid \mathbf{t} \cdot \mathbf{t}', \\ T_{\text{cond}}(\mathcal{L}) &\text{ is generated by the grammar } \mathbf{t} := \mathbf{P}(\delta \mid \delta'), \\ T_{\text{lin}}(\mathcal{L}) &\text{ is generated by the grammar } \mathbf{t} := \mathbf{P}(\delta) \mid \mathbf{t} + \mathbf{t}', \\ T_{\text{comp}}(\mathcal{L}) &\text{ is generated by the grammar } \mathbf{t} := \mathbf{P}(\delta). \end{aligned}$$

We define for each $* \in \{\text{comp}, \text{lin}, \text{cond}, \text{poly}\}$ the causal and purely probabilistic languages:

$$\begin{aligned} \mathcal{L}_{\text{prob}}^* &:= \mathbf{t} \geq \mathbf{t} \mid \neg \mathcal{L}_{\text{prob}}^* \mid \mathcal{L}_{\text{prob}}^* \wedge \mathcal{L}_{\text{prob}}^*, & \mathbf{t} \in T_*(\mathcal{L}_{\text{prop}}), \\ \mathcal{L}_{\text{causal}}^* &:= \mathbf{t} \geq \mathbf{t} \mid \neg \mathcal{L}_{\text{causal}}^* \mid \mathcal{L}_{\text{causal}}^* \wedge \mathcal{L}_{\text{causal}}^*, & \mathbf{t} \in T_*(\mathcal{L}_{\text{full}}). \end{aligned}$$

Several of these probabilistic languages have appeared in the literature. For instance, $\mathcal{L}_{\text{prob}}^{\text{poly}}$ appeared already in early work by Scott and Krauss [40], while $\mathcal{L}_{\text{prob}}^{\text{lin}}$ was introduced explicitly by Fagin et al. [18]. The language $\mathcal{L}_{\text{causal}}^{\text{poly}}$ was introduced and studied recently in [25] (see also [5, 16]). Many of these languages, however, have not yet received explicit treatment.

2.2. Semantics.

2.2.1. Structural causal models. The semantics for all of these languages will be defined relative to *structural causal models*, which can be understood as a very general framework for encoding *data-generating processes*. In addition to the endogenous variables \mathbf{V} , structural causal models also employ *exogenous variables* \mathbf{U} as a source of random variation among endogenous settings. For extended introductions, see, e.g., [5, 34].

DEFINITION 2.1. A structural causal model (SCM) \mathbf{M} is a tuple $\mathbf{M} = (\mathcal{F}, \mathbb{P}, \mathbf{U}, \mathbf{V})$, with:

- (a) \mathbf{V} a set of endogenous variables, with each $V \in \mathbf{V}$ taking on possible values $\text{Val}(V)$,
- (b) \mathbf{U} a set of exogenous variables, with each $U \in \mathbf{U}$ taking on possible values $\text{Val}(U)$,
- (c) $\mathcal{F} = \{f_V\}_{V \in \mathbf{V}}$ a set of structural functions, such that f_V determines the value of V given the values of the exogenous variables \mathbf{U} and those of the other endogenous variables $V' \in \mathbf{V}$, and
- (d) \mathbb{P} a probability measure on a σ -algebra $\sigma(\mathbf{U})$ on \mathbf{U} .

Here we will assume for convenience that $\text{Val}(V)$ and $\text{Val}(U)$ are all finite.

In addition, we adopt the common assumption that our SCMs are recursive:

DEFINITION 2.2. An SCM \mathbf{M} is recursive if there is a well-order \prec on \mathbf{V} such that \mathcal{F} respects \prec in the following sense: for any $V \in \mathbf{V}$, whenever $\mathbf{v}_1, \mathbf{v}_2 : V \mapsto \text{Val}(V)$ have

the property that $\mathbf{v}_1(V') = \mathbf{v}_2(V')$ for all $V' \prec V$, we are guaranteed that $f_V(\mathbf{v}_1, \mathbf{u}) = f_V(\mathbf{v}_2, \mathbf{u})$.

Intuitively, \mathbf{M} is recursive if for all $V \in \mathbf{V}$, the function f_V ensures that the value of V is determined only by the exogenous random variables $U \in \mathbf{U}$ and endogenous random variables $V' \in \mathbf{V}$ for which $V' \prec V$. Thus in a recursive model \mathbf{M} , the probability measure \mathbb{P} on $\sigma(\mathbf{U})$ induces a joint probability distribution $\mathbb{P}(\mathbf{V})$ over values of the variables $V \in \mathbf{V}$.

Causal interventions represent the result of a *manipulation* to the causal system, and are defined in the standard way (e.g., [34, 43]):

DEFINITION 2.3. *An intervention is a partial function $i : V \mapsto \text{Val}(V)$. It specifies variables $\text{dom}(i) \subseteq \mathbf{V}$ to be held fixed and the values to which they are fixed. An intervention i induces a mapping, also denoted i , of systems of equations $\mathcal{F} = \{f_V\}_{V \in \mathbf{V}}$, such that $i(\mathcal{F})$ is identical to \mathcal{F} , but with f_V replaced by the constant function $f_V(\cdot) = i(V)$ for each $V \in \text{dom}(i)$. Similarly, where \mathbf{M} is a model with equations \mathcal{F} , we write $i(\mathbf{M})$ for the model which is identical to \mathbf{M} but with the equations $i(\mathcal{F})$ in place of \mathcal{F} .*

In order to guarantee that interventions lead to a well-defined semantics, we work with structural causal models which are measurable:

DEFINITION 2.4. *We say that \mathbf{M} is measurable if under every finite intervention i , the joint distribution $\mathbb{P}(\mathbf{V})$ associated with the model $i(\mathbf{M})$ is well-defined.*

For measurable models, one can define a notion of causal influence:

DEFINITION 2.5. *A model \mathbf{M} induces the influence relation $V_i \rightsquigarrow V_j$ when there exist values $v, v' \in \text{Val}(V_j)$ and interventions α, α' differing only in the value they impose upon V_i for which¹*

$$\mathbf{M} \models \mathbf{P}([\alpha]V_j = v \wedge [\alpha']V_j = v') > 0.$$

Given an enumeration of variables V_1, \dots, V_n compatible with a well-order \prec , the model \mathbf{M} is compatible with \prec when it induces no instance $V_i \rightsquigarrow V_j$ with $i > j$.

To illustrate the preceding definitions, we return to the front door graph shown in Example 1.1, and demonstrate an example of an SCM that is compatible with this graph:

EXAMPLE 2.6. *Consider the SCM $\mathbf{M} = (\mathcal{F}, \mathbb{P}, \mathbf{U}, \mathbf{V})$, with the exogenous $\mathbf{U} = \{U, U_X, U_Y, U_Z\}$, each of which has probability $1/2$ of being 1 and probability $1/2$ of being 0, and with three endogenous variables $\mathbf{V} = \{X, Y, Z\}$. The equations $\mathcal{F} = \{f_V\}_{V \in \mathbf{V}}$ are given by*

$$\begin{aligned} f_X(U_X, U) &= U \wedge U_X, \\ f_Z(X, U_Z) &= X \wedge U_Z, \\ f_Y(Z, U_Y, U) &= Z \wedge U \wedge U_Y. \end{aligned}$$

We observe that \mathbf{M} is measurable and recursive with the ordering \prec given by $X \prec Z \prec Y$. Further, $X \rightsquigarrow Z$ and $Z \rightsquigarrow Y$, so that \mathbf{M} indeed realizes the front door graph and is compatible with \prec .

¹ The truth definition for \models is introduced formally below in Section 2.2.2.

2.2.2. *Interpretations of terms and truth definitions.* It suffices to give the semantics for $\mathcal{L}_{\text{causal}}^{\text{poly}}$, since this language includes all of the other languages introduced above. A model is a recursive and measurable SCM $\mathbf{M} = (\mathcal{F}, \mathbb{P}, \mathbf{U}, \mathbf{V})$. For each assignment $\mathbf{u} : U \mapsto \text{Val}(U)$ of values to exogenous variables, each $V \in \mathbf{V}$, and each $v \in \text{Val}(V)$, we define $\mathcal{F}, \mathbf{u} \models V = v$ if the equations \mathcal{F} together with the assignment \mathbf{u} assign the value v to V . Conjunction and negation are defined in the usual way, giving semantics for $\mathcal{F}, \mathbf{u} \models \beta$ for any $\beta \in \mathcal{L}_{\text{prop}}$. If $\mathcal{F}, \mathbf{u} \models \beta$ holds for all \mathbf{u} , then we simply write $\mathcal{F} \models \beta$. When the relation $\mathcal{F}, \mathbf{u} \models \beta$ does not depend on \mathbf{u} at all—that is, we have $\mathcal{F}, \mathbf{u} \models \beta$ iff $\mathcal{F}, \mathbf{u}' \models \beta$ for all \mathbf{u}, \mathbf{u}' and all formulas β —we say that the equations \mathcal{F} are deterministic. For $\beta, \beta' \in \mathcal{L}_{\text{prop}}$, we write $\beta \models \beta'$ when $\mathcal{F} \models \beta \rightarrow \beta'$ for all \mathcal{F} , where material implication is defined in the usual way.

For each intervention $\alpha \in \mathcal{L}_{\text{int}}$ and each $\beta \in \mathcal{L}_{\text{prop}}$, we define $\mathcal{F}, \mathbf{u} \models [\alpha]\beta$ iff $i_\alpha(\mathcal{F}), \mathbf{u} \models \beta$, where i_α is the intervention which effects the assignments described by α . We also allow that α may be the trivial intervention \top , in which case we simply write β instead of $[\alpha]\beta$. We define

$$\llbracket \mathbf{P}(\varepsilon) \rrbracket_{\mathbf{M}} = \mathbb{P}(\{\mathbf{u} : \mathcal{F}, \mathbf{u} \models \varepsilon\}).$$

For conditional probability terms we define $\llbracket \mathbf{P}(\delta|\delta') \rrbracket_{\mathbf{M}} = 1$ when $\llbracket \mathbf{P}(\delta') \rrbracket_{\mathbf{M}} = 0$ and using the above definition and the usual ratio definition otherwise. For two terms $\mathbf{t}_1, \mathbf{t}_2$, we define $\mathbf{M} \models \mathbf{t}_1 \geq \mathbf{t}_2$ iff $\llbracket \mathbf{t}_1 \rrbracket_{\mathbf{M}} \geq \llbracket \mathbf{t}_2 \rrbracket_{\mathbf{M}}$. The semantics for negation and conjunction are defined in the usual way, giving a semantics for $\mathbf{M} \models \varphi$ for any $\varphi \in \mathcal{L}_{\text{causal}}^{\text{poly}}$.

With this semantics, probability behaves as expected. For example, we have the following validity for any $\varepsilon_1, \varepsilon_2$:

$$\text{Add. } \mathbf{P}(\varepsilon_1 \wedge \varepsilon_2) + \mathbf{P}(\varepsilon_1 \wedge \neg\varepsilon_2) = \mathbf{P}(\varepsilon_1).$$

Causal interventions behave as expected as well. Indeed, fix any model \mathbf{M} with equations \mathcal{F} , any variable $V \in \mathbf{V}$, and any assignment \mathbf{u} of values to the exogenous variables. Then V takes on at least and at most one value upon the intervention α : this is trivial if α intervenes on V , and it otherwise follows immediately from the fact that once \mathbf{u} is fixed, the values of all variables are determined by the equations $i_\alpha(\mathcal{F})$. In other words, in the language $\mathcal{L}_{\text{causal}}^*$ for any $*$ \in $\{\text{comp, lin, cond, poly}\}$, we have the validity for all \mathbf{M} and \mathbf{u} :

$$\text{Def. } \bigwedge_{\substack{v, v' \in \text{Val}(V) \\ v \neq v'}} \neg[\alpha](V = v \wedge V = v') \wedge \bigvee_{v \in \text{Val}(V)} [\alpha](V = v).$$


More generally, for each $\alpha \in \mathcal{L}_{\text{int}}$, the indexed box $[\alpha]$ can be thought of as a *normal, functional* modal operator.

Having introduced the syntax and semantics for several languages and pointed to some basic validities, we recall in the next subsection various results and examples that illustrate the expressive relationships between these languages.

2.3. A two-dimensional expressive hierarchy.

DEFINITION 2.7. For a formula φ in any of the languages just introduced, let $\text{Mod}(\varphi) = \{\mathbf{M} : \mathbf{M} \models \varphi\}$ be the class of its models. For two languages \mathcal{L}_1 and \mathcal{L}_2 , we say that \mathcal{L}_2 is at least as expressive as \mathcal{L}_1 if for every $\varphi \in \mathcal{L}_1$ there is some $\psi \in \mathcal{L}_2$ such that $\text{Mod}(\varphi) = \text{Mod}(\psi)$. We say \mathcal{L}_2 is strictly more expressive than \mathcal{L}_1 if \mathcal{L}_2 is at least as expressive as \mathcal{L}_1 but not vice versa.

In this section, mostly rehearsing familiar results and examples, we illustrate that the expressivity of the languages \mathcal{L}^* for $\mathcal{L} \in \{\mathcal{L}_{\text{prob}}, \mathcal{L}_{\text{causal}}\}$ and $* \in \{\text{comp}, \text{lin}, \text{cond}, \text{poly}\}$ form an expressive hierarchy along two axes. First, the purely probabilistic language $\mathcal{L}_{\text{prob}}^*$ is always less expressive than the corresponding causal language $\mathcal{L}_{\text{causal}}^*$. Second, $\mathcal{L}_{\text{prob}}^{\text{comp}}$ is less expressive than both \mathcal{L}^{lin} and $\mathcal{L}^{\text{cond}}$, both of which are less expressive than the language $\mathcal{L}^{\text{poly}}$. Where each arrow indicates a strict increase in expressivity, the hierarchy can be shown graphically:²



2.3.1. *First axis: from probabilistic to causal.* To illustrate the expressivity of causal as opposed to purely probabilistic languages, we recall a variation by Bareinboim et al. [5] on an example due to Pearl [34]:

EXAMPLE 2.8 (Causation without correlation). Let $M_1 = (\mathcal{F}, \mathbb{P}, \mathbf{U}, \mathbf{V})$, where \mathbf{U} contains two binary variables U_1, U_2 such that $\mathbb{P}(U_1) = \mathbb{P}(U_2) = 1/2$, and \mathbf{V} contains two variables V_1, V_2 such that $f_{V_1} = U_1$ and $f_{V_2} = U_2$. Then V_1 and V_2 are independent. Having observed this, one could not conclude that V_1 has no causal effect on V_2 ; indeed, consider the model M' , which is like M , except with the mechanisms:

$$f_{V_1} = \mathbf{1}_{U_1=U_2},$$

$$f_{V_2} = U_1 + \mathbf{1}_{V_1=1, U_1=0, U_2=1}.$$

Here $\mathbf{1}_S$ is the indicator function for statement S , equal to 1 if S holds and 0 otherwise. In this case $\mathbb{P}_M(V_1, V_2) = \mathbb{P}_{M'}(V_1, V_2)$, so that the models are indistinguishable in any of the probabilistic languages $\mathcal{L}_{\text{prob}}^*$. However, the models are distinguishable in $\mathcal{L}_{\text{causal}}^{\text{comp}}$, and so in all of the other causal languages. Indeed, note that $\mathbb{P}_M([V_1 = 1]V_2 = 1) = 1/2$ while $\mathbb{P}_{M'}([V_1 = 1]V_2 = 1) = 3/4$. Then, for instance, the following statement

$$\mathbb{P}([V_1 = 1]V_2 = 1) = \mathbb{P}([V_1 = 1]V_2 = 0)$$

belongs to $\mathcal{L}_{\text{causal}}^{\text{comp}}$ and distinguishes M from M' .

As shown in [5] (cf. also [45]), the pattern in Example 2.8 is universal: for any model M it is always possible to find some M' that agrees with M on all of $\mathcal{L}_{\text{prob}}^{\text{poly}}$ but disagrees on $\mathcal{L}_{\text{causal}}^{\text{poly}}$.³

² The arrow in the center of these squares is meant to indicate that $\mathcal{L}_{\text{prob}}^*$ is less expressive than $\mathcal{L}_{\text{causal}}^*$ for any choice of $* \in \{\text{comp}, \text{lin}, \text{cond}, \text{poly}\}$.

³ The Causal Hierarchy Theorem of Bareinboim et al. [5] (refer to [26] for a topological version, enabling the relevant generalization to infinite \mathbf{V}) involves an intermediate language between $\mathcal{L}_{\text{prob}}^{\text{poly}}$ and $\mathcal{L}_{\text{causal}}^{\text{poly}}$, capturing the type of causal information revealed by controlled experiments. Even this three-tiered hierarchy is strict, and in fact one can go further to obtain

THEOREM 2.9. $\mathcal{L}_{\text{causal}}^{\text{poly}}$ is more expressive than $\mathcal{L}_{\text{prob}}^{\text{poly}}$. What is stronger, no $\mathcal{L}_{\text{prob}}^{\text{poly}}$ -theory (i.e., maximally consistent set in this language) uniquely determines a $\mathcal{L}_{\text{causal}}^{\text{poly}}$ -theory.

2.3.2. Second axis: from qualitative to quantitative. Focusing just on probabilistic languages, we will show that $\mathcal{L}_{\text{prob}}^{\text{comp}}$ is less expressive than both $\mathcal{L}_{\text{prob}}^{\text{lin}}$ and $\mathcal{L}_{\text{prob}}^{\text{cond}}$, and that both of these are less expressive than the language $\mathcal{L}_{\text{prob}}^{\text{poly}}$. In each case, it suffices to give two measures $\mathbb{P}_1(\mathbf{V})$ and $\mathbb{P}_2(\mathbf{V})$ which are indistinguishable in the less expressive language but which can be distinguished by some statement in the more expressive one.

Comparative probability. First, we claim that $\mathcal{L}_{\text{prob}}^{\text{comp}}$ is less expressive than $\mathcal{L}_{\text{prob}}^{\text{lin}}$. Suppose we have just a single binary variable X , abbreviating $X = 1$ by q and $X = 0$ by $\neg q$. Then let $\mathbb{P}_1(q) = 2/3$ so that $\mathbb{P}_1(\neg q) = 1/3$, and let $\mathbb{P}_2(q) = 3/5$ so that $\mathbb{P}_2(\neg q) = 2/5$. The qualitative order on the four events $q, \neg q, \top, \perp$ is the same, but, for instance, $\mathbb{P}_1(q) = \mathbb{P}_1(\neg q) + \mathbb{P}_1(\neg q)$, while $\mathbb{P}_2(q) \neq \mathbb{P}_2(\neg q) + \mathbb{P}_2(\neg q)$.

Next, we recall an example due to Luce [29], which shows that $\mathcal{L}_{\text{prob}}^{\text{comp}}$ is less expressive than $\mathcal{L}_{\text{prob}}^{\text{cond}}$. Let p, q, r each be events corresponding to the three possible values taken by a random variable. Consider the measures $\mathbb{P}_1(p) = 5/9, \mathbb{P}_1(q) = 3/9, \mathbb{P}_1(r) = 1/9$ and $\mathbb{P}_2(p) = 6/9, \mathbb{P}_2(q) = 2/9, \mathbb{P}_2(r) = 1/9$. Then the two orders are the same, because for $i \in [2]$,

$$\mathbb{P}_i(\top) > \mathbb{P}_i(p \vee q) > \mathbb{P}_i(p \vee r) > \mathbb{P}_i(p) > \mathbb{P}_i(q \vee r) > \mathbb{P}_i(q) > \mathbb{P}_i(r) > \mathbb{P}_i(\perp).$$

However, the conditional probabilities differ: $\mathbb{P}_1(r|q \vee r) < \mathbb{P}_1(q|p \vee q)$, while $\mathbb{P}_2(r|q \vee r) > \mathbb{P}_2(q|p \vee q)$. In other words, the measures \mathbb{P}_1 and \mathbb{P}_2 are indistinguishable in $\mathcal{L}_{\text{prob}}^{\text{comp}}$ but distinguishable in $\mathcal{L}_{\text{prob}}^{\text{cond}}$.

Polynomials in probabilities. To show that $\mathcal{L}_{\text{prob}}^{\text{lin}}$ is less expressive than $\mathcal{L}_{\text{prob}}^{\text{poly}}$, we simply identify a formula $\varphi \in \mathcal{L}_{\text{prob}}^{\text{poly}}$ such that there is no $\psi \in \mathcal{L}_{\text{prob}}^{\text{lin}}$ with $\text{Mod}(\varphi) = \text{Mod}(\psi)$. For this we can take the example $\mathbb{P}(A \wedge B) = \mathbb{P}(\neg A \vee \neg B) \wedge \mathbb{P}(A|B) = \mathbb{P}(B)$. (This is in fact expressible already in $\mathcal{L}_{\text{prob}}^{\text{cond}}$.) This enforces that $\mathbb{P}(B) = 1/\sqrt{2}$, while Ibeling et al. [27] show that every formula in $\mathcal{L}_{\text{prob}}^{\text{lin}}$ has models in which every probability is rational.

Finally, we give an example to show that $\mathcal{L}_{\text{prob}}^{\text{cond}}$ is less expressive than $\mathcal{L}_{\text{prob}}^{\text{poly}}$. As above, let p, q, r be events corresponding to possible values taken by a random variable. Define $\mathbb{P}_1(p) = 3/20, \mathbb{P}_1(q) = 4/20$, and $\mathbb{P}_1(r) = 13/20$, while $\mathbb{P}_2(p) = 3/20 - .03, \mathbb{P}_2(q) = 4/20 - .01$, and $\mathbb{P}_2(r) = 13/20 + .04$. One can verify by exhaustion that all comparisons of conditional probabilities agree between \mathbb{P}_1 and \mathbb{P}_2 , and thus they are indistinguishable in $\mathcal{L}_{\text{prob}}^{\text{cond}}$. At the same time, there are statements in $\mathcal{L}_{\text{prob}}^{\text{poly}}$ in which the models differ. For example, $\mathbb{P}_1(r)\mathbb{P}_1(q) < \mathbb{P}_1(p)$, whereas $\mathbb{P}_2(r)\mathbb{P}_2(q) > \mathbb{P}_2(p)$. This shows that $\mathcal{L}_{\text{prob}}^{\text{cond}}$ is less expressive than $\mathcal{L}_{\text{prob}}^{\text{poly}}$. Further, we observe that $\mathbb{P}_1, \mathbb{P}_2$ can be distinguished in

an infinite hierarchy of increasingly expressive causal languages between $\mathcal{L}_{\text{prob}}^{\text{poly}}$ and $\mathcal{L}_{\text{causal}}^{\text{poly}}$. Because we are showing that there is a complexity collapse even from the most expressive to the least expressive systems, we are not concerned in the present work with these intermediate languages.

$\mathcal{L}_{\text{prob}}^{\text{lin}}$: $\mathbb{P}_i(q) \geq 0.2$ for $i = 1$ but not for $i = 2$, and this statement is equivalent to the statement in $\mathcal{L}_{\text{prob}}^{\text{lin}}$ that

$$\underbrace{\mathbb{P}_1(q) + \dots + \mathbb{P}_i(q)}_{10 \text{ times}} \geq \mathbb{P}_i(\top) + \mathbb{P}_i(\top).$$

Together, this observation and the earlier remark that $\mathbb{P}(A \wedge B) = \mathbb{P}(\neg A \vee \neg B) \wedge \mathbb{P}(A|B) = \mathbb{P}(B)$ is expressible in $\mathcal{L}_{\text{prob}}^{\text{cond}}$ show that \mathcal{L}^{lin} and $\mathcal{L}^{\text{cond}}$ are incomparable in expressivity.

Summarizing the results of this section:

THEOREM 2.10. *\mathcal{L}^{lin} and $\mathcal{L}^{\text{cond}}$ are incomparable in expressive power. Both are strictly more expressive than $\mathcal{L}^{\text{comp}}$ and strictly less expressive than $\mathcal{L}^{\text{poly}}$.*

§3. Introducing computational complexity. In this section, we introduce the ideas from complexity theory needed to state our main results. We denote by $\text{SAT}_{\text{prob}}^*$, $\text{SAT}_{\text{causal}}^*$ the satisfiability problems for $\mathcal{L}_{\text{prob}}^*$, $\mathcal{L}_{\text{causal}}^*$, respectively, where $*$ \in {comp, lin, cond, poly}. There are two key definitions:

DEFINITION 3.1. *Say that a map $\varphi \mapsto \psi$ preserves and reflects satisfiability when φ is satisfiable if and only if ψ is satisfiable. Such a map is called a many-one reduction of φ to ψ . Such a map is said to run in polynomial time if it is computable by a Turing machine in a number of time steps that is a polynomial function of the length $|\varphi|$ of the input formula. When the Turing machine is non-deterministic, the map is said to be non-deterministic as well; in this case we say that the reduction is an NP-reduction.*

DEFINITION 3.2. *A decision problem maps an input, represented as a binary string, to an output “yes” or “no.” For example, $\text{SAT}_{\text{prob}}^*$ maps a standard encoding of the formula $\varphi \in \mathcal{L}_{\text{prob}}^*$ to “yes” if it is satisfiable and to “no” otherwise. When each member of a collection \mathcal{C} of decision problems can be reduced via some deterministic, polynomial-time map to a particular decision problem $c \in \mathcal{C}$, one says that the problem c is \mathcal{C} -complete. The class \mathcal{C} of decision problems is called a complexity class.*

When a problem c is complete for some complexity class, this means that the complexity class \mathcal{C} fully characterizes the difficulty of the problem: the problem c is at least as “hard” as any of the problems in \mathcal{C} , and it is itself in \mathcal{C} . Thus any two problems which are complete for a complexity class are equally hard, since each can be reduced in deterministic polynomial time to the other. Complete problems facilitate results relating complexity classes: to show that a class \mathcal{C} is contained in another \mathcal{C}' , it suffices to give deterministic, polynomial-time, many-one reduction from a problem c which is complete for \mathcal{C} to any problem $c' \in \mathcal{C}'$.

Fagin et al. [18] showed that $\text{SAT}_{\text{prob}}^{\text{lin}}$ is complete for the complexity class NP. That $\text{SAT}_{\text{prob}}^{\text{comp}}$ is also NP-complete follows quickly from this result and the Cook–Levin theorem [12], which says that Boolean satisfiability is NP-complete as well. For clarity, we include these known results in the statement of our main result, which gives completeness results for all of the other probabilistic and causal languages defined above:

THEOREM 3.3. *We characterize two sets of tasks:*

1. $\text{SAT}_{\text{prob}}^{\text{comp}}$, $\text{SAT}_{\text{prob}}^{\text{lin}}$, $\text{SAT}_{\text{causal}}^{\text{comp}}$, $\text{SAT}_{\text{causal}}^{\text{lin}}$ are NP-complete.

2. $\text{SAT}_{\text{prob}}^{\text{cond}}, \text{SAT}_{\text{prob}}^{\text{poly}}, \text{SAT}_{\text{causal}}^{\text{cond}}, \text{SAT}_{\text{causal}}^{\text{poly}}$ are $\exists\mathbb{R}$ -complete.

This is the more formal statement of Theorem 1.2.

Since problems that are complete for a class are all equally hard, our main results imply that causal and probabilistic reasoning in these languages do not differ in complexity. In the remainder of this section, we introduce the complexity classes NP and $\exists\mathbb{R}$. We note that the inclusions $\text{NP} \subseteq \exists\mathbb{R} \subseteq \text{PSPACE}$ are known [9], where PSPACE is the set of problems solvable using polynomial space; it is an open problem whether either inclusion is strict. Further, NP and PSPACE are closed under many-one NP-reductions, and ten Cate et al. [11] show that $\exists\mathbb{R}$ is also closed under many-one NP-reductions:

DEFINITION 3.4. *A complexity class \mathcal{C} is closed under many-one NP reductions if to show that a problem is in \mathcal{C} , it suffices to find a polynomial-time NP-reduction of the problem to one that is known to be in \mathcal{C} .*

3.1. The class NP. The class NP contains any problem that can be solved by a non-deterministic Turing machine in a number of steps that grows polynomially in the input size. Equivalently, it contains any problem solvable by a polynomial-time deterministic Turing machine, when the machine is provided with a polynomial-size certificate, which we think of as providing the solution to the problem, or “lucky guesses.” In this case we think of the deterministic Turing machine as a *verifier*, tasked with ensuring that the certificate communicates a valid solution to the problem.

Hundreds of problems are known to be NP-complete. Among them are Boolean satisfiability and the decision problems associated with several natural graph properties, for example possession of a clique of a given size or possession of a Hamiltonian path. See [36] for a survey of such problems and their relations.

3.2. The class $\exists\mathbb{R}$. The Existential Theory of the Reals (ETR) contains all true sentences of the form

$$\text{there exist } x_1, \dots, x_n \in \mathbb{R} \text{ satisfying } \mathcal{S},$$

where \mathcal{S} is a system of equalities and inequalities of arbitrary polynomials in the variables x_1, \dots, x_n . For example, one can state in ETR the existence of the golden ratio, which is the only root of the polynomial $f(x) = x^2 - x - 1$ greater than one, by “there exists $x > 1$ satisfying $f(x) = 0$.” The decision problem of saying whether a given formula $\varphi \in \text{ETR}$ is complete (by definition) for the complexity class $\exists\mathbb{R}$.

The class $\exists\mathbb{R}$ is the real analogue of NP, in two senses. Firstly, the satisfiability problem that is complete for $\exists\mathbb{R}$ features real-valued variables, while the satisfiability problems that are complete for NP typically feature integer- or Boolean-valued variables. Secondly, and more strikingly, Erickson et al. [17] recently showed that while NP is the class of decision problems with answers that can be verified in polynomial time by machines with access to unlimited integer-valued memory, $\exists\mathbb{R}$ is the class of decision problems with answers that can be verified in polynomial time by machines with access to unlimited *real-valued* memory.

As with NP, a myriad of problems are known to be $\exists\mathbb{R}$ -complete. We include some examples that illustrate the diversity of such problems:

- In graph theory, there is the $\exists\mathbb{R}$ -complete problem of deciding whether a given graph can be realized by a straight line drawing [39].

- In game theory, there is the $\exists\mathbb{R}$ -complete problem of deciding whether an (at least) three-player game has a Nash equilibrium with no probability exceeding a fixed threshold [7].
- In geometry, there is the $\exists\mathbb{R}$ -complete “art gallery” problem of finding the smallest number of points from which all points of a given polygon are visible [2].
- In machine learning, there is the $\exists\mathbb{R}$ -complete problem of finding weights for a neural network trained on a given set of data such that the total error of the network falls below a given threshold [3].

For discussions of further $\exists\mathbb{R}$ -complete problems, see [10, 38].

§4. Our results. In this section, we prove our main result (Theorem 1.2). To do this, we first establish that one can reduce satisfiability problems for causal languages to corresponding problems for purely probabilistic languages.

4.1. Reduction.

DEFINITION 4.1. Fix a set OP of operations on \mathbb{R} , and for a given placeholder set S , let $OP(S)$ be the set of terms generated by application of operations in OP to members of S . Define

$$\begin{aligned} \mathcal{L}_{\text{prob}} &= \mathbf{t}_1 \geq \mathbf{t}_2 \mid \neg\mathcal{L}_{\text{prob}} \mid \mathcal{L}_{\text{prob}} \wedge \mathcal{L}_{\text{prob}}, & \mathbf{t}_i &\in OP(\{\mathbf{P}(\varepsilon) : \varepsilon \in \mathcal{L}_{\text{prop}}\}), \\ \mathcal{L}_{\text{causal}} &= \mathbf{t}_1 \geq \mathbf{t}_2 \mid \neg\mathcal{L}_{\text{causal}} \mid \mathcal{L}_{\text{causal}} \wedge \mathcal{L}_{\text{causal}}, & \mathbf{t}_i &\in OP(\{\mathbf{P}(\varepsilon) : \varepsilon \in \mathcal{L}_{\text{full}}\}). \end{aligned}$$

The semantics for these languages are restricted to recursive SEMs.

PROPOSITION 4.2 (Reduction). There exists a many-one NP reduction from $\text{SAT}_{\mathcal{L}_{\text{causal}}}$ to $\text{SAT}_{\mathcal{L}_{\text{prob}}}$.

We first give a prose overview of the main ideas underlying the reduction. Fix $\varphi \in \mathcal{L}_{\text{causal}}$. The key observation is that the reduction is straightforward when every ε with $\mathbf{P}(\varepsilon)$ mentioned in $\varphi \in \mathcal{L}_{\text{causal}}$ is a complete state description, where a complete state description says, for each possible intervention and each variable, what value that variable takes upon that intervention. Indeed, complete state descriptions have three nice properties:

1. *Polynomial-time comparison to ordering.* One can easily check whether a complete state description implies influence relations conflicting with a given order \prec on the variables appearing in it. Indeed, one simply reads which variables influence which variables of the intervention statements appearing in ε .
2. *Existence of model matching probabilities.* If a collection of complete state descriptions does not conflict with an order \prec , then any probability distribution on the descriptions ε has a recursive model that induces it; briefly, one can simply take a distribution over deterministic models for the mutually unsatisfiable descriptions ε .
3. *Small model property.* At most $|\varphi|$ complete state descriptions are mentioned in φ , and so at most that many receive positive probability in any model satisfying φ .

These properties will allow a reduction to go through. Indeed, fix $\varphi \in \mathcal{L}_{\text{causal}}$. Given that φ is satisfiable, one can request as an NP certificate an ordering \prec and (relying

on #3) the small set of complete state descriptions receiving positive probability. One then checks (relying on #1) that these descriptions do not conflict with \prec . Since φ is satisfiable only if there exists a measure satisfying its inequalities, one can safely translate those inequalities into the probabilistic language, giving a satisfiable probabilistic formula ψ . If the probabilistic formula ψ is satisfiable via some measure, one can (relying on #2) infer a corresponding recursive model for the causal formula φ . Thus the map $\varphi \mapsto \psi$ preserves and reflects satisfiability.

As it turns out, the same reduction goes through in the general case, when the ε for which $\mathbf{P}(\varepsilon)$ is mentioned in φ need not be complete state descriptions. Roughly, the strategy is to simply replace every ε such that $\mathbf{P}(\varepsilon)$ is mentioned in $\varphi \in \mathcal{L}_{\text{causal}}$ with an equivalent disjunction of complete state descriptions. The primary complication with this strategy is that there are too many possible interventions, variables, and values those variables could take on; *truly* complete state descriptions are exponentially long, making the reduction computationally intractable. To address this issue, we work with a *restricted class* of state descriptions, which feature only the interventions, variables, and values appearing in the input formula φ :

DEFINITION 4.3. Fix a formula $\varphi \in \mathcal{L}_{\text{prop}} \cup \mathcal{L}_{\text{causal}}$. Let \mathcal{I} contain all interventions appearing in φ and let \mathbf{V}_φ denote all variables appearing in φ . For each variable $V \in \mathbf{V}_\varphi$, let $\text{Assignments}_\varphi(V)$ contain $V = v$ whenever $V = v$ or $V \neq v$ appears in φ , and let it also contain one assignment $V = v^*$ not satisfying either of these conditions. Let Δ_φ contain all possible interventions paired with all possible assignments, where the possibilities are restricted to φ :

$$\Delta_\varphi = \left\{ \bigwedge_{\alpha \in \mathcal{I}} \left([\alpha] \bigwedge_{V \in \mathbf{V}_\varphi} \beta_V^\alpha \right) : \beta_V^\alpha \in \text{Assignments}_\varphi(V) \text{ for } V \in \mathbf{V}_\varphi \right\}.$$

Call $\bigwedge_{V \in \mathbf{V}_\varphi} \beta_V^\alpha$ the results of the intervention α , and β_V^α the result for V of the intervention α . We write $\alpha \in \delta$ when $\delta \in \Delta_\varphi$ as shorthand for $\alpha \in \mathcal{I}$. We write $V \in \alpha$ when α contains some assignment $V = v$.

The following three lemmas confirm that even working with this restricted class of state descriptions, (versions of) the three nice properties outlined above are retained.

DEFINITION 4.4. Fix a formula $\varphi \in \mathcal{L}_{\text{prop}} \cup \mathcal{L}_{\text{causal}}$ and $\Delta' \subseteq \Delta_\varphi$. Fix a well-order \prec on \mathbf{V}_φ . Enumerate the variables V_1, \dots, V_n in \mathbf{V}_φ in a way consistent with \prec . The formula $\delta \in \Delta'$ is compatible with \prec when there exists a model \mathbf{M} that assigns positive probability to δ and that is compatible with \prec . Define Δ_{\prec} to contain all $\delta \in \Delta_\varphi$ compatible with \prec .

LEMMA 4.5 (Polytime comparison to ordering). Fix $\varphi \in \mathcal{L}_{\text{prop}} \cup \mathcal{L}_{\text{causal}}$. Given a set $\Delta' \subseteq |\varphi|$, one can check that $\Delta' \subseteq \Delta_\varphi$ and that each $\delta \in \Delta'$ is compatible with \prec in time polynomial in $|\varphi|$.

This lemma shows that given some statement φ and a set of formulas Δ' , one can efficiently (i.e., in polynomial time) check that the formulas $\delta \in \Delta'$ satisfy two conditions. The first condition is that the formulas δ describe, in the fullest terms possible, the ways that φ could be true (i.e., $\Delta' \subseteq \Delta_\varphi$). The second is that the formulas δ do not rule out the causal influence relations specified by the order \prec , for example

the relations $X \prec Z \prec Y$ induced by the model of smoking's effect on lung cancer discussed in Examples 1.1 and 2.6.

Proof. Checking that $\Delta' \subseteq \Delta_\varphi$ is fast, since one can simply scan φ to make sure that φ mentions precisely interventions mentioned in all $\delta \in \Delta'$; that φ mentions precisely the variables V appearing in the results of every intervention in $\delta \in \Delta'$; and that for each such variable V , at most of one of its assignments $V = v$ in Δ' does not appear as an assignment or a negated assignment in φ .

We now give an algorithm to check whether $\delta \in \Delta'$ is compatible with \prec . We first give prose and formal descriptions of the algorithm and then consider its runtime and correctness.

Order the variables V_1, \dots, V_n in \mathbf{V}_φ in a way consistent with the well-order \prec . For each variable V_i with $i \in [n]$, do the following. First, for each intervention α in δ that mentions V_i , confirm that the intervention leads to satisfiable results: if δ says that upon the intervention α which sets $V_i = v$, the variable V_i takes a value $v' \neq v$, we reject δ , which necessarily has probability 0. Next, for each pair of interventions α, α' in δ which do not intervene on the value assigned to V_i , check whether both interventions result in the same assignments to variables V_j for all $j < i$; we say that such interventions α, α' have *agreement* on all V_j for $j < i$. If this is the case, and yet δ says that these two interventions result in different values for V_i , reject δ ; since V_i can depend only on the values of V_j for $j < i$, when these values are constant, V_i must be constant as well. Here is a formal description of the algorithm. We will write $V \in \alpha$ to denote that the variable V appears (or is mentioned) in the intervention α , i.e., that $V = v$ is a conjunct in α for some value v .

Algorithm 1: Check that $\delta \in \Delta'$ is compatible with \prec .

```

Order the variables  $V_1, \dots, V_n$  in  $\mathbf{V}_\varphi$  according to  $\prec$ 
for  $i$  in  $1, \dots, n$  do
  for intervention  $\alpha$  in  $\delta$  with  $V_i = v$  appearing in  $\alpha$  do
    if  $V_i = v'$  with  $v \neq v'$  appears in the conjunction of assignments
      following  $\alpha$  then
      | return  $\delta$  is unsatisfiable, and so incompatible with  $\prec$ 
    end
  end
  for interventions  $\alpha, \alpha'$  in  $\delta$  agreeing on all  $V_j$  for  $j < i$ , and such that  $V_i \notin \alpha$ 
    and  $V_i \notin \alpha'$  do
    | if  $\alpha$  results in  $V_i = v$  and  $\alpha'$  results in  $V_i = v'$  with  $v \neq v'$  then
    | | return  $\delta$  is incompatible with  $\prec$ 
    | end
  end
end
return  $\delta$  is compatible with  $\prec$ 

```

Below, we show that the above algorithm indeed runs in time $\text{poly}(|\varphi|)$ and is correct, but for clarity, let us step through its execution on some examples. Consider the input $\delta := [V_1 = 0]V_1 = 1$. Then, by the first “if” clause in the algorithm, δ is rejected as

unsatisfiable, since the intervention $[V_1 = 0]$ leads to impossible results. For another example, let δ' be the formula

$$[V_1 = 1 \wedge V_4 = 1](V_1 = 1 \wedge V_2 = 0 \wedge V_3 = 1 \wedge V_4 = 1) \\ \wedge [V_1 = 1 \wedge V_4 = 0](V_1 = 1 \wedge V_2 = 0 \wedge V_3 = 0 \wedge V_4 = 0).$$

Then in the second “if” clause on the third iteration, δ' is rejected as incompatible with \prec , because the interventions $\alpha = [V_1 = 1 \wedge V_4 = 1]$ and $\alpha' = [V_1 = 1 \wedge V_4 = 0]$ do not intervene on V_3 , result in the same values for V_1 and V_2 , and do result in the same value for V_3 , contradicting the fact that V_3 's value must depend only on those assigned to V_1 and V_2 .

It is helpful in considering these examples and the runtime of the algorithm to consider the following table of values:

Results of all interventions in the input formula δ'				
Intervention	V_1	V_2	V_3	V_4
$\alpha = [V_1 = 1 \wedge V_4 = 1]$	$V_1 = 1$	$V_2 = 0$	$V_3 = 1$	$V_4 = 1$
$\alpha' = [V_1 = 1 \wedge V_4 = 0]$	$V_1 = 1$	$V_2 = 0$	$V_3 = 0$	$V_4 = 0$

In effect, the second “for” loop over all interventions α, α' constructs the above table, starting with the leftmost column V_1 and proceeding to the right. The algorithm rejects δ' when two cells in the column V_i and rows α and α' (with $V_i \notin \alpha, \alpha'$) do not assign the same value to V_i but agree on all columns V_j to the left. The restriction that V_i does not appear in α or α' must be included because distinct interventions α, α' can disagree on the values they impose on V_i when intervening on it, regardless of the values assigned to V_j with $j < i$; such disagreement does not constitute a violation of the ordering \prec .

Let us first confirm that this algorithm runs in time $\text{poly}(|\varphi|)$ and then show its correctness. We observe that $\max\{|\delta|, n\} = \text{poly}(|\varphi|)$. The algorithm contains an $O(n)$ loop over V_1, \dots, V_n and two $O(|\delta|^2)$ loops over interventions. The work performed inside of these loops takes time $O(n \cdot |\delta|)$, since we are simply reading δ and checking values for the variables V_j for all $j < i$, which can be stored in a lookup table (like the one above) of size $O(n \cdot |\delta|)$. Thus the runtime of the algorithm is indeed $\text{poly}(|\varphi|)$.

Finally, we confirm that the algorithm is correct. Fix any $\delta \in \Delta'$ and recall that δ is of the form

$$\bigwedge_{\alpha} \left([\alpha] \bigwedge_{V \in V_{\varphi}} \beta_V^{\alpha} \right),$$

where $\beta_V^{\alpha} \in \text{Assignments}_{\varphi}(V)$. First, suppose that the above algorithm declares δ compatible with \prec . We will inductively construct a deterministic model of equations $\mathcal{F} = \{f_{V_i}\}_{i \in [n]}$ and show that $\mathcal{F} \models \delta$ and \mathcal{F} is compatible with \prec . Define f_{V_1} to be the constant function sending all arguments to β_{V_1} , where β_{V_1} is the value of V_1 upon any intervention $\alpha \in \delta$ with $V_1 \notin \alpha$; the second “for” loop in the algorithm ensures that there is at most one such value, and if there is no such value, β_{V_1} can be chosen arbitrarily. Then $f_{V_1} \models \bigwedge_{\alpha} [\alpha] \beta_{V_1}^{\alpha}$. Indeed, this holds by construction for α with $V_1 \notin \alpha$, and it holds trivially for α with $V_1 \in \alpha$, because, by the first

“for” loop, each α is compatible with its results. For the inductive step, define $f_{V_i}(V_1 = \beta_{V_1}, \dots, V_{i-1} = \beta_{V_{i-1}}) = \beta_{V_i}$, where β_{V_i} is the value of V_i upon any intervention $\alpha \in \delta$ for which $V_i \notin \alpha$ and $\beta_{V_j}^\alpha = \beta_{V_j}$ for all $j < i$; by the same reasoning, there is at most one such value, and if there is no such value, β_{V_i} can be chosen arbitrarily. Then by the same reasoning, $f_{V_i} \models \bigwedge_\alpha [\alpha] \beta_{V_i}^\alpha$. Because this holds for all $i \in [n]$, we have $\mathcal{F} \models \delta$. By construction, \mathcal{F} is compatible with \prec , as desired.

Now, suppose that δ is compatible with \prec , so that δ is not self-contradictory and there exists some $\mathcal{F} = \{f_{V_i}\}_{i \in [n]}$ compatible with \prec for which $\mathcal{F} \models \delta$. We claim that the above algorithm returns that δ is indeed compatible with \prec . Suppose for a contradiction that on iteration i , the algorithm rejects δ as incompatible with \prec . Since δ is not self-contradictory, it follows by the definition of the algorithm that for some interventions α, α' (with $V_i \notin \alpha, \alpha'$) which agree on all V_j for $j < i$, we have $[\alpha]V_i = v$ and $[\alpha']V_i = v'$ with $v \neq v'$. Let β_j be the value such that the assignments $V_j = \beta_j$ for $j < i$ result from the interventions α and α' . Then

$$\mathcal{F} \models [V_1 = \beta_1, \dots, V_{i-1} = \beta_{i-1}](V_i = v \wedge V_i \neq v'),$$

which is impossible. □

LEMMA 4.6 (Existence of model matching probabilities). *Fix $\varphi \in \mathcal{L}_{\text{prop}} \cup \mathcal{L}_{\text{causal}}$, and suppose \mathbb{P} is a measure on $\Delta_\prec \subseteq \Delta_\varphi$ for some \prec . Then there is a model \mathbf{M} inducing the measure \mathbb{P} on Δ_\prec , i.e., $[\mathbf{P}(\delta)]_{\mathbf{M}} = \mathbb{P}(\delta)$ for all $\delta \in \Delta_\prec$.*

Proof. Let us first define the model \mathbf{M} and then show that it is recursive. Let \mathbf{V}_φ denote all variables appearing in φ . We define $\mathbf{M} = (\mathcal{F}, \mathbb{P}_{\mathbf{M}}, \{U\}, \mathbf{V}_\varphi)$, where $\text{Val}(U) = \Delta_\varphi$ and $\mathbb{P}_{\mathbf{M}}(U = \delta) = \mathbb{P}(\delta)$ for all $\delta \in \Delta_\varphi$. Enumerate the variables $\mathbf{V}_\varphi = \{V_1, \dots, V_n\}$ in a way consistent with \prec . Fix any $\delta \in \Delta_\varphi$. Recall that δ is of the form

$$\bigwedge_\alpha \left([\alpha] \bigwedge_{i \in [n]} \beta_{V_i}^\alpha \right),$$

where $\beta_{V_i}^\alpha \in \text{Assignments}_\varphi(V_i)$. If δ is satisfiable, it has a model, i.e., a deterministic system of equations $\mathcal{F}^\delta = \{f_{V_i}^\delta\}_{i \in [n]}$ such that $\mathcal{F}^\delta \models \delta$. Turning now to define the equations $\mathcal{F} = \{f_{V_i}\}_{i \in [n]}$, for any assignment \mathbf{v} to the variables V_j for $j < i$, put

$$f_{V_i}(\mathbf{v}, U = \delta) = f_{V_i}^\delta(\mathbf{v}).$$

By the above equations and mutual unsatisfiability of $\delta \in \Delta_\prec$, it follows that for all such δ

$$[\mathbf{P}(\delta)]_{\mathbf{M}} = \mathbb{P}_{\mathbf{M}}[U = \delta] = \mathbb{P}(\delta),$$

as required.

It remains for us to confirm that \mathbf{M} is recursive. We claim that the influence relationships $V_i \rightsquigarrow V_j$ induced by the model \mathbf{M} are simply those induced by the deterministic models⁴ \mathcal{F}^δ . This would complete the proof, since by assumption we have $\delta \in \Delta_\prec$, so that \mathcal{F}^δ is compatible with \prec , and therefore $i < j$. Suppose that \mathbf{M} induces the influence relation $V_i \rightsquigarrow V_j$. Then for some interventions α, α' which

⁴ That is, according to Definition 2.5 when each deterministic model is thought of as a probabilistic model in which its respective system of functions is selected with no uncertainty.

disagree only on the value assigned to V_i , some assignment \mathbf{u} to U , and some distinct values v, v' of V_j , we have

$$\mathcal{F}, \mathbf{u} \models [\alpha]V_j = v \wedge [\alpha']V_j = v'.$$

Let δ be the value that \mathbf{u} assigns to U . We claim that

$$\mathcal{F}^\delta \models [\alpha]V_j = v \wedge [\alpha']V_j = v'.$$

Indeed, this follows from the fact that $f_{V_i}(\mathbf{v}, U = \delta) = f_{V_i}^\delta(\mathbf{v})$ for all $i \in [n]$. □

LEMMA 4.7 (Small model property). *Fix $\varphi \in \mathcal{L}_{\text{prop}} \cup \mathcal{L}_{\text{causal}}$. If φ is satisfiable, then φ has a small model, in the sense that the model assigns positive probability to at most $|\varphi|$ elements $\delta \in \Delta_\varphi$.*

Proof. Since φ is satisfiable, it has a recursive model \mathbf{M} with some order \prec . Given the existence of \mathbf{M} , we claim there exists a small model $\mathbf{M}_{\text{small}}$. Indeed, consider the system of equations in the unknowns $\{\mathbb{P}(\delta) : \delta \in \Delta_\prec\}$ given by

$$\begin{aligned} \sum_{\delta \in \Delta_\prec} \mathbb{P}(\delta) &= 1, \\ \sum_{\substack{\delta \in \Delta_\prec \\ \delta \models \varepsilon}} \mathbb{P}(\delta) &= \mathbb{P}_{\mathbf{M}}(\varepsilon), \text{ for each } \varepsilon \text{ such that } \mathbf{P}(\varepsilon) \text{ is mentioned in } \varphi. \end{aligned}$$

There are at most $|\varphi|$ equations, and since for each ε , there exists some $\delta \in \Delta_\prec$ for which $\delta \models \varepsilon$, the equations are non-trivial. Suppose for the moment that $\mathbb{P} = \mathbb{P}_{\mathbf{M}}$ is a solution. Then by a fact of linear algebra (see Lemma 4.8 of [18]), since the at most $|\varphi|$ linear equations have a solution, they have a solution $\mathbb{P} = \mathbb{P}_{\text{small}}$ in which at most $|\varphi|$ of the variables $\mathbb{P}_{\text{small}}(\delta)$ are nonzero. By Lemma 4.6, we then infer from the existence of $\mathbb{P}_{\text{small}}$ that the desired model $\mathbf{M}_{\text{small}}$ exists.

It remains to confirm that $\mathbb{P} = \mathbb{P}_{\mathbf{M}}$ is indeed a solution to the above system of equations. To show this, we must show that for ε with $\mathbf{P}(\varepsilon)$ mentioned in φ ,

$$\mathbb{P}(\varepsilon) = \sum_{\substack{\delta \in \Delta_\prec \\ \delta \models \varepsilon}} \mathbb{P}(\delta).$$

By our choice of \prec , we know that the recursive model \mathbf{M} will assign probability 0 to all $\delta \notin \Delta_\prec$. It thus suffices to show that the above holds when Δ_\prec is replaced with the larger set Δ_φ . To do this, we will put ε into a more manageable form; afterwards, establishing the above equality will be relatively straightforward.

If ε mentions only one intervention α , we claim that $\models \varepsilon \leftrightarrow [\alpha]\beta$, where β is an assignment of variables in φ to various values. Indeed, negation and conjunction distribute over $[\alpha]$, in the sense that $\models \neg[\alpha]\beta \leftrightarrow [\alpha]\neg\beta$ and $\models [\alpha]\beta \wedge [\alpha]\beta' \leftrightarrow [\alpha](\beta \wedge \beta')$, so $[\alpha]$ can be assumed to appear on the outside. Further, since by the validity Def, each variable takes one and only one value upon the intervention α , we can replace β with a disjunction over all assignments to all variables in φ which agree with β . Let us use \mathbf{v}_φ to denote such an assignment $\bigwedge_{V \in \mathbf{v}_\varphi} \beta_V$ where $\beta_V \in \text{Assignments}(V)$, as defined in Definition 4.3. Summing up:

$$\models \varepsilon \leftrightarrow [\alpha] \bigvee_{\mathbf{v}_\varphi \models \beta} \mathbf{v}_\varphi.$$

The exact same ideas apply when ε mentions several interventions $[\alpha_i]\beta_i$ for $i \in [n]$, in which case

$$\models \varepsilon \leftrightarrow \bigwedge_i [\alpha_i] \bigvee_{\mathbf{v}_\varphi \models \beta_i} \mathbf{v}_\varphi \leftrightarrow \bigvee_{\mathbf{v}_\varphi^i \models \beta_i} \bigwedge_i [\alpha_i] \mathbf{v}_\varphi^i.$$

Thus, since all interventions, variables, and assignments appearing in ε are mentioned by the $\delta \in \Delta_\varphi$, and one can always add trivial interventions $[\alpha]\top$, we see that it is a validity that ε is equivalent to a disjunction of formulas $\delta \in \Delta_\varphi$. Finally, we conclude with the observation that since the $\delta \in \Delta_\varphi$ are mutually unsatisfiable, additivity for the measure \mathbb{P} (according to which $\mathbb{P}(\delta \vee \delta') = \mathbb{P}(\delta) + \mathbb{P}(\delta')$ for mutually unsatisfiable δ, δ') tells us that

$$\mathbb{P}(\varepsilon) = \sum_{\substack{\delta \in \Delta_\varphi \\ \delta \models \varepsilon}} \mathbb{P}(\delta),$$

as desired. □

With the lemmas in hand, we now give the desired reduction:

Proof of Proposition 4.2. Fix a $\text{SAT}_{\mathcal{L}_{\text{causal}}}$ instance φ . We first describe the NP certificate and many-one reduction and then prove soundness and completeness. The NP certificate consists of an order \prec on \mathbf{V}_φ and a set of Δ' of size at most $|\varphi|$. The reduction proceeds as follows.

1. Check that $\Delta' \subseteq \Delta_\varphi$ and that each $\delta \in \Delta'$ is compatible with \prec . We note that by Lemma 4.5, this can be done in time polynomial in $|\varphi|$.
2. Replace $\mathbf{P}(\varepsilon)$ appearing in φ with $\mathbf{P}\left(\bigvee_{\substack{\delta \in \Delta' \\ \delta \models \varepsilon}} f(\delta)\right)$, where f is a bijection between Δ' and an arbitrary set of mutually unsatisfiable statements in $\mathcal{L}_{\text{prop}}$. Call the resulting $\mathcal{L}_{\text{prob}}$ formula $\varphi(\Delta')$.

We note that checking $\delta \models \varepsilon$ can be done in polynomial time, since δ is a complete description of the results of all interventions.

First we show completeness. If φ is satisfiable, by Lemma 4.7 it has a small model that assigns positive probability only to some $\Delta' \subseteq \Delta_\prec$ for some ordering \prec , and the probabilities given by this model also solve $\varphi(\Delta')$. So the certificate exists and the reduction succeeds in producing a satisfiable formula.


Now we show soundness. If $\varphi(\Delta')$ is satisfiable, it is solved by some measure \mathbb{P} . This is a measure defined on $f(\Delta')$, and so on $\Delta' \subseteq \Delta_\prec$. Thus since each $\delta \in \Delta'$ is compatible with \prec , by Lemma 4.6 there exists a model \mathbf{M} such that $\llbracket \mathbf{P}(\delta) \rrbracket_{\mathbf{M}} = \mathbb{P}(\delta)$ for $\delta \in \Delta_\varphi$. This \mathbf{M} is a model of the inequalities stated by φ and is recursive, so φ is satisfiable as well. □

4.2. Characterization. Recall that Theorem 3.3 characterizes two sets of tasks:

We characterize two sets of tasks:

1. $\text{SAT}_{\text{prob}}^{\text{comp}}$, $\text{SAT}_{\text{prob}}^{\text{lin}}$, $\text{SAT}_{\text{causal}}^{\text{comp}}$, and $\text{SAT}_{\text{causal}}^{\text{lin}}$ are NP-complete.
2. $\text{SAT}_{\text{prob}}^{\text{cond}}$, $\text{SAT}_{\text{prob}}^{\text{poly}}$, $\text{SAT}_{\text{causal}}^{\text{cond}}$, and $\text{SAT}_{\text{causal}}^{\text{poly}}$ are $\exists\mathbb{R}$ -complete.

We can express these results in a diagram, which holds for $* \in \{\text{prob, causal}\}$:



The line separates $\exists\mathbb{R}$ -complete problems from NP-complete problems, and an arrow from one satisfiability problem to another indicates that any instance of the former problem is an instance of the latter.

We note that these results imply that there exists a many-one, polynomial-time, *deterministic* many-one reduction from $\text{SAT}_*^{\text{causal}}$ to $\text{SAT}_*^{\text{prob}}$, for any $* \in \{\text{comp, lin, cond, poly}\}$, whereas Proposition 4.2 only gives a non-deterministic reduction. To illustrate, recall the model of smoking’s effect on lung cancer discussed in Examples 1.1 and 2.6. Consider again the task of determining whether smoking makes one more likely to possess lung cancer, given one’s causal assumptions Γ and one’s observation of statistical correlation between smoking, tar deposits in the lungs, and lung cancer. In other words, the task is determine whether

$$\Gamma + \text{Correlational data} \models \mathbf{P}([X = 1]Y = 1) > \mathbf{P}([X = 0]Y = 1), \tag{4}$$

where the correlational data include statements such as $\mathbf{P}(Y = 1|X = 1) > \mathbf{P}(Y = 1|X = 0)$ and $0.7 > \mathbf{P}(Y = 1|X = 1) > 0.6$. The above result implies that this task is no more difficult than that of determining whether an analogous entailment

$$\Gamma' + \text{Correlational data} \models \text{Probabilistic conclusion} \tag{5}$$

holds, given *purely probabilistic* assumptions Γ' . Indeed, given Equation (4), one can efficiently (i.e., in polynomial time) construct a probabilistic equation with the form of Equation (5) such that both entailments have the same truth-value; the causal inference goes through if and only if the purely probabilistic inference goes through.

To show the results in the second part of Theorem 1.2, we borrow the following lemma from [2]:

LEMMA 4.8. *Fix variables x_1, \dots, x_n , and set of equations of the form $x_i + x_j = x_k$ or $x_i x_j = 1$, for $i, j, k \in [n]$. Let $\exists\mathbb{R}$ -inverse be the problem of deciding whether there exist reals x_1, \dots, x_n satisfying the equations, subject to the restrictions $x_i \in [1/2, 2]$. This problem is $\exists\mathbb{R}$ -complete.*

Here, for reasons of space we only outline the two steps in the proof of Lemma 4.8. First, one shows that finding a real root of a degree 4 polynomial with rational coefficients is $\exists\mathbb{R}$ -complete, and then one repeatedly performs variable substitutions to get the constraints $x_i + x_j = x_k$ and $x_i x_j = 1$. Second, one shows that any such polynomial has a root within a closed ball about the origin, and then one shifts and scales this ball to contain exactly the range $[1/2, 2]$.

With the lemma in hand, we show the theorem:

Proof of Theorem 1.2. We begin with the first statement. Using the fact that $\text{SAT}_{\text{prob}}^{\text{comp}}$ is NP-hard, it suffices to show that $\text{SAT}_{\text{causal}}^{\text{lin}}$ is inside NP; indeed, since all of the satisfiability problems mentioned in the first statement include $\text{SAT}_{\text{prob}}^{\text{comp}}$ and are included by $\text{SAT}_{\text{causal}}^{\text{lin}}$, they would all then be NP-hard and inside NP, and so would all be NP-complete. It is known both that $\text{SAT}_{\text{prob}}^{\text{lin}}$ is inside NP and that NP is closed under many-one NP reductions; by Proposition 4.2, this places $\text{SAT}_{\text{causal}}^{\text{lin}}$ inside NP, as desired.

We turn now to the second statement. By the same reasoning, it suffices to show that $\text{SAT}_{\text{prob}}^{\text{cond}}$ is $\exists\mathbb{R}$ -hard and that $\text{SAT}_{\text{causal}}^{\text{poly}}$ is inside $\exists\mathbb{R}$. We claim that $\text{SAT}_{\text{prob}}^{\text{poly}}$ is inside $\exists\mathbb{R}$; $\exists\mathbb{R}$ is closed under many-one NP reductions [11], so Proposition 4.2 will place $\text{SAT}_{\text{causal}}^{\text{poly}}$ in $\exists\mathbb{R}$ immediately.

To show that $\text{SAT}_{\text{prob}}^{\text{poly}}$ is inside $\exists\mathbb{R}$, we slightly extend a proof by [25] that the problem is in PSPACE. Suppose that $\varphi \in \mathcal{L}_{\text{prob}}^{\text{poly}}$ is satisfied by some model \mathbb{P} . Again using the fact that $\exists\mathbb{R}$ is closed under NP-reductions, we will provide a reduction of φ to a formula $\psi \in \text{ETR}$. Let E contain all ε such that $\mathbf{P}(\varepsilon)$ appears in φ . Then consider the formula

$$\sum_{\delta \in \Delta_\varphi} \mathbf{P}(\delta) = 1,$$

$$\sum_{\substack{\delta \in \Delta_\varphi \\ \delta \models \varepsilon}} \mathbf{P}(\delta) = \mathbb{P}(\varepsilon) \text{ for } \varepsilon \in E.$$

The measure \mathbb{P} satisfies the above formula, so by Lemma 4.7, the above formula is satisfied by some model $\mathbb{P}_{\text{small}}$ assigning positive probability to a subset $\Delta^+ \subseteq \Delta_\varphi$ of size at most $|E| \leq |\varphi|$. Thus adding to φ the constraint $\sum_{\delta \in \Delta^+} \mathbf{P}(\delta) = 1$ and replacing each $\mathbf{P}(\varepsilon)$ appearing in φ with $\sum_{\delta \in \Delta^+, \delta \models \varepsilon} \mathbf{P}(\delta)$ gives a formula ψ belonging to ETR which has a model, namely $\mathbb{P}_{\text{small}}$ —and conversely, the mutual unsatisfiability of the $\delta \in \Delta^+$, together with the fact that they sum to unity, ensures that any model of ψ is a model of φ . Further, the size constraints on E and Δ^+ ensure that ψ can be formed in polynomial time.

Let us conclude the proof by showing that $\text{SAT}_{\text{prob}}^{\text{cond}}$ is $\exists\mathbb{R}$ -hard. To do this, consider an $\exists\mathbb{R}$ -inverse problem instance φ with variables x_1, \dots, x_n . It suffices to find in polynomial time a $\text{SAT}_{\text{prob}}^{\text{cond}}$ instance ψ preserving and reflecting satisfiability. We first describe the reduction and then show that it preserves and reflects satisfiability.

Corresponding to the variables x_1, \dots, x_n , let us define events $\delta_1, \dots, \delta_n$, with an additional event δ_{n+1} , and with the events $\delta_i \in \mathcal{L}_{\text{prop}}$ mutually unsatisfiable. We introduce further variables $\delta_{i,j}$ for $i, j \in [n]$. We must now translate into $\mathcal{L}_{\text{prob}}^{\text{cond}}$ the constraints of the form (a) $x_i \in [1/2, 2]$, (b) $x_i + x_j = x_k$, and (c) $x_i x_j = 1$. At a high level, the strategies are (a) to require that for any model \mathbb{P} of ψ , we have $\mathbb{P}(\delta_i) \in [1/4n, 1/n]$, (b) to express constraints on sums using probabilities of disjoint unions of events, and (c) to express constraints on products using conditional probabilities. The reduction will proceed so that the map $x_i \mapsto x_i/2n$ sends solutions of φ to solutions of ψ (and so that the map's inverse sends solutions of ψ to solutions of φ).

For each $i \in [n]$, we would like to require that $1/4n \leq \mathbf{P}(\delta_i) \leq 1/n$, adding to these the constraint $\mathbf{P}(\delta_{n+1}) = \mathbf{P}(\neg \bigvee_{i \leq n} \delta_i)$; the former is analogous to the constraint that

$x_i \in [1/2, 2]$, and the latter ensures that the probabilities $\mathbf{P}(\delta_i)$ sum to 1. The only obstacle to including the constraint $1/4n \leq \mathbf{P}(\delta_i) \leq 1/n$ is that $\mathcal{L}_{\text{prop}}^{\text{cond}}$ does not contain constants for $1/4n$ and $1/n$.

This obstacle can be addressed as follows. We replace all instances of $1/n$ with $\mathbf{P}(\bigvee_{\substack{\ell \in [m] \\ b_\ell=1}} \varepsilon_i)$, where:

- The variable b_ℓ denotes the ℓ -th digit in the binary expansion $b_1 \dots b_m$ of $1/n$.
- The events $\varepsilon_\ell \in \mathcal{L}_{\text{prop}}$ for $\ell \in [m]$ are mutually unsatisfiable and mention neither any of the variables mentioned in the $\delta_{i,j}$ for $i, j \in [n]$ nor any of the variables mentioned by δ_k for $k \in [n + 1]$.
- We impose the constraint $\mathbf{P}(\varepsilon_1) = \mathbf{P}(-\varepsilon_1)$ and, for each integer ℓ between 2 and m , impose the two constraints $\mathbf{P}(\varepsilon_\ell) = \mathbf{P}(\varepsilon_{\ell,\ell-1} \wedge \varepsilon_{\ell-1})$ and $\mathbf{P}(\varepsilon_{\ell,\ell-1} | \varepsilon_{\ell-1}) = \mathbf{P}(\varepsilon_{\ell-1})$. Here, the events $\varepsilon_{\ell,\ell-1}$ mention none of the variables mentioned in $\varepsilon_{\ell'}$ for any $\ell' \in [m]$, nor those mentioned in $\delta_{i,j}, \delta_i$ for any $i, j \in [n + 1]$.

These constraints ensure that for any solution \mathbb{P} to ψ , we have $\mathbb{P}(\varepsilon_j) = 2^{-j}$, so that (by mutual unsatisfiability) $\mathbb{P}(\bigvee_{\substack{\ell \in [m] \\ b_\ell=1}} \varepsilon_\ell) = \sum_{\substack{\ell \in [m] \\ b_\ell=1}} \mathbb{P}(\varepsilon_\ell) = 1/n$. The same procedure allows for the expression of arbitrary rational constants, for example $1/4n$ or $1/4n^2$.

Turning next to sums and products, we replace each constraint of the form $x_i + x_j = x_k$ with $\mathbf{P}(\delta_i \vee \delta_j) = \mathbf{P}(\delta_k)$, and we replace each constraint of the form $x_i x_j = 1$ with the two constraints $\mathbf{P}(\delta_{i,j} | \delta_i) = \mathbf{P}(\delta_j)$ and $\mathbf{P}(\delta_{i,j} \wedge \delta_i) = 1/4n^2$. The latter two constraints are included in order to ensure that for any solution \mathbb{P} to ψ , we have $\mathbb{P}(\delta_i) \cdot \mathbb{P}(\delta_j) = 1/4n^2$, giving the desired multiplicative constraint. The events $\delta_{i,j} \in \mathcal{L}_{\text{prop}}$ for $i, j \in [n]$ are to mention none of the variables mentioned in δ_k for any $k \in [n + 1]$.

This completes our description of the reduction. The operations performed are very simple, so to check the runtime of this reduction, it suffices to check the length of ψ . The probabilistic equation corresponding to $x_i + x_j = x_k$ is of the same length and introduces no new variables besides the original δ_i for $i \in [n]$. We next observe that our substitution for the equation $x_i x_j = 1$ introduces two new constraints, each with length polynomial in $|\varphi|$. Finally, we observe that the constraints introduced in the above itemized list number at most $2m$, where $m = \lceil \log 1/4n^2 \rceil$, each being of length polynomial in m , which is in turn polynomial in $|\varphi|$.

Let us now show that the map $\varphi \mapsto \psi$ preserves satisfiability. Given a solution x_1, \dots, x_n to φ , we will construct a solution to ψ . Define $\mathbb{P}(\delta_i) = x_i/2n$ for $i \in [n]$ and $\mathbb{P}(\delta_{n+1}) = 1 - \sum_{i \leq n} \mathbb{P}(\delta_i)$. Then as $1/2 \leq x_i \leq 2$, we have $1/4n \leq \mathbb{P}(\delta_i) \leq 1/n$, and the probabilities $\mathbb{P}(\delta_i)$ sum to 1, as required. Further, since $x + x_j = x_k$, we have (by mutual unsatisfiability) $\mathbb{P}(\delta_i \vee \delta_j) = \mathbb{P}(\delta_i) + \mathbb{P}(\delta_j) = \mathbb{P}(\delta_k)$, as required. Finally, let us check that the constraints $\mathbb{P}(\delta_{i,j} | \delta_i) = \mathbb{P}(\delta_j)$ and $\mathbb{P}(\delta_{i,j} \wedge \delta_i) = 1/4n^2$ are satisfied. Using the definition of conditional probability, these amount to the constraint that

$$\frac{1}{4n^2} = \mathbb{P}(\delta_i) \cdot \mathbb{P}(\delta_j),$$

which is satisfied, since $\mathbb{P}(\delta_i) = x_i/2n, \mathbb{P}(\delta_j) = x_j/2n$, and $x_i x_j = 1$. To prove that $\varphi \mapsto \psi$ reflects satisfiability, we use the same reasoning, this time defining $x_i = 2n \cdot \mathbb{P}(\delta_i)$. \square

§5. Conclusion and outlook. We have shown that questions posed in probabilistic *causal* languages can be systematically reduced to purely probabilistic queries, showing

that the former are—from a computational perspective—no more complex than the latter. In fact, we demonstrated a kind of bifurcation between two classes of languages. On the one hand, languages encompassing at most *addition* enjoy an NP-complete satisfiability problem, whether the language is causal or not. However, as soon as we admit even a modicum of *multiplication* into the language, causal and probabilistic languages become hard for the class $\exists\mathbb{R}$, and even the full language of polynomials over (causal) probability terms is $\exists\mathbb{R}$ -complete. At the low end, this applies to a language with no explicit addition or multiplication, but just inequalities between conditional probability terms, or even simple independence statements for pairs of variables. As clarified in the resulting landscape of formal systems, we have identified an important sense in which causal reasoning is no more difficult than pure probabilistic reasoning. The substantial empirical and expressive gulf between causation and “mere (statistical) association” is evidently not reflected in a complexity gap.

It should be acknowledged that, from the standpoint of inferential practice, questions of the form (2) constitute just one part of a larger methodological pipeline. In some sense this is only a final stage in the process of going from an inductive problem to a deductive conclusion. The formulation of reasonable inductive assumptions can itself be an arduous task, as can translating those assumptions into a language like $\mathcal{L}_{\text{prob}}$ or $\mathcal{L}_{\text{causal}}$ (that is, into the set Γ). Take once again the example of do-calculus (Example 1.1). The idea behind this method is that in many contexts investigators will be in a position to make reasonable *qualitative* (viz. graphical) assumptions, perhaps justified by expert knowledge, to the effect that some variables are *not* causally impacted in a direct way by certain other variables. Even when this method involves nothing more than assuming a specific causal (directed acyclic) graph, it may still take work to determine which causal-probabilistic statements are licensed by the graph. Many subtasks in this connection have been studied. For instance, determining whether three sets of variables in a graph stand in the so-called *d-separation* relation (which in turn guarantees conditional independence) is known to be very easy (it is *linear time*; see, e.g., [37]). Nonetheless, there are certainly other questions related to complexity that one might ask in this and other settings.

Moving beyond statistical and causal inference tasks narrowly construed, the results in this article raise a number of further research questions, both technical and conceptual. For instance, one can easily imagine versions of our causal languages in a multi-agent setting, with a (causal) probability operator \mathbf{P}_a for multiple agents a . As has been widely recognized, strategic interaction routinely involves reasoning about causality and counterfactuals (see, e.g., [44]). Existing formal proposals for capturing these styles of reasoning have been largely qualitative, with counterfactual patterns formalized using models of belief revision rather than structural causal models (see, e.g., [8]). Whereas (“pure”) probability-logical languages have been thoroughly explored in the game theory literature (e.g., [23]), the causal-probability-logical languages studied here would be quite natural to investigate in that context. Echoing our themes in the present article, what happens to computational complexity in this multi-agent setting, and specifically would a reduction to pure (multi-)probability still be possible?

In a more technical vein, there are natural questions about further extensions to even the most expressive languages we considered. To take just one example, much of probabilistic and causal reasoning employs tools from information theory like (conditional) entropy that in turn rely on logarithmic principles, or alternatively (via

inversion), reasoning about *exponentiation*. A major open problem in logic—known as Tarski’s exponential function problem—is to determine whether the first-order theory of the reals with exponentiation is decidable. Short of that, one might hope to show that some of the weaker (causal-)probability languages studied here remain decidable, perhaps even of relatively low complexity, when exponentiation is added. However, for the strongest languages, such as $\mathcal{L}_{\text{prob}}^{\text{poly}}$, this may prove difficult. As Macintyre and Wilkie [30] have shown, decidability of the *existential* theory of the reals with the unary function e^x would already imply a positive answer to Tarski’s problem.

The reader will surely think of further questions and extensions pertaining to our work in this article. We hope that the systems, results, and methods offered here will be useful in these various directions moving forward, and more generally will help to catalyze further research at the fruitful intersection of logic, probability, causality, and complexity.

BIBLIOGRAPHY

- [1] Abadi, M., & Halpern, J. Y. (1994). Decidability and expressiveness for first-order logics of probability. *Information and Computation*, **112**, 1–36.
- [2] Abrahamsen, M., Adamaszek, A., & Miltzow, T. (2018). The art gallery problem is $\exists\mathbb{R}$ -complete. In Diakonikolas, I., and Kempe, D., editors. *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*. New York: Association for Computing Machinery, pp. 65–73.
- [3] Abrahamsen, M., Kleist, L., & Miltzow, T. (2021). Training neural networks is $\exists\mathbb{R}$ -complete. In Ranzato, M., editor. *Proceedings of the Thirty-Fifth Conference on Neural Information Processing Systems (NeurIPS)*. Vancouver: Curran Associates, Inc.
- [4] Aleksandrowicz, G., Chockler, H., Halpern, J. Y., & Ivrii, A. (2017). The computational complexity of structure-based causality. *Journal of Artificial Intelligence Research*, **58**, 431–451.
- [5] Bareinboim, E., Correa, J., Ibeling, D., & Icard, T. (2022). On Pearl’s hierarchy and the foundations of causal inference. In Geffner, H., Dechter, R., and Halpern, J. Y., editors. *Probabilistic and Causal Inference: The Works of Judea Pearl*. New York: ACM Books, pp. 509–556.
- [6] Belot, G. (2020). Absolutely no free lunches! *Theoretical Computer Science*, **845**, 159–180.
- [7] Bilò, V., & Mavronicolas, M. (2017). $\exists\mathbb{R}$ -complete decision problems about symmetric Nash equilibria in symmetric multi-player games. In Vallée, B., and Vollmer, H., editors. *34th Symposium on Theoretical Aspects of Computer Science (STACS 2017)*. Dagstuhl: Schloss Dagstuhl-Leibniz-Zentrum für Informatik.
- [8] Board, O. (2004). Dynamic interactive epistemology. *Games and Economic Behavior*, **49**(1), 49–80.
- [9] Canny, J. (1988). Some algebraic and geometric computations in PSPACE. In Simon, J., editor. *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*. New York: Association for Computing Machinery, pp. 460–467.
- [10] Cardinal, J. (2015). Computational geometry column 62. *ACM SIGACT News*, **46**(4), 69–78.

- [11] ten Cate, B., Kolaitis, P. G., & Othman, W. (2013). Data exchange with arithmetic operations. In Guerrini, G., editor. *Proceedings of the 16th International Conference on Extending Database Technology*. New York: Association for Computing Machinery, pp. 537–548.
- [12] Cook, S. A. (1971). The complexity of theorem-proving procedures. In Harrison, M., Banerji, R., and Ullman, J., editors. *Proceedings of the Third Annual ACM Symposium on Theory of Computing*. New York: Association for Computing Machinery, pp. 151–158.
- [13] Darwiche, A. (2021). Causal inference using tractable circuits. In Ranzato, M., editor. *Proceedings of the Thirty-Fifth Conference on Neural Information Processing Systems (NeurIPS)*. Vancouver: Curran Associates, Inc.
- [14] Duarte, G., Finkelstein, N., Knox, D., Mummolo, J., & Shpitser, I. (2021). An automated approach to causal inference in discrete settings. Preprint, [arXiv:2109.13471v1](https://arxiv.org/abs/2109.13471v1).
- [15] Efron, B. (1978). Controversies in the foundations of statistics. *The American Mathematical Monthly*, **85**(4), 231–246.
- [16] Eiter, T., & Lukasiewicz, T. (2002). Complexity results for structure-based causality. *Artificial Intelligence*, **142**, 53–89.
- [17] Erickson, J., van der Hoog, I., & Miltzow, T. (2020). *Smoothing the gap between NP and $\exists\mathbb{R}$* . In Irani, S., editor. *IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*. Durham, NC: IEEE, pp. 1022–1033.
- [18] Fagin, R., Halpern, J. Y., & Megiddo, N. (1990). A logic for reasoning about probabilities. *Information and Computation*, **87**(1–2), 78–128.
- [19] de Finetti, B. (1937). La prévision: ses lois logiques, ses sources subjectives. *Annales de l'Institut Henri Poincaré*, **7**, 1–68.
- [20] Gelman, A., & Shalizi, C. R. (2013). Philosophy and the practice of Bayesian statistics. *British Journal of Mathematical and Statistical Psychology*, **66**, 8–38.
- [21] Halpern, J., & Vardi, M. (1989). The complexity of reasoning about knowledge and time. *Journal of Computer and System Sciences*, **38**, 195–237.
- [22] Halpern, J. Y. (2000). Axiomatizing causal reasoning. *Journal of Artificial Intelligence Research*, **12**, 317–337.
- [23] Heifetz, A., & Mongin, P. (2001). Probability logic for type spaces. *Games and Economic Behavior*, **35**, 31–53.
- [24] Ibeling, D. (2018). Causal modeling with probabilistic simulation models. In Bellodi E., and Shrijvers, T., editors. *Proceedings of the 5th International Workshop on Probabilistic Logic Programming (PLP)*. Ferrara: Italian Association for Artificial Intelligence, pp. 36–48.
- [25] Ibeling, D., & Icard, T. (2020). Probabilistic reasoning across the causal hierarchy. In Rossi, F., editor. *Proceedings of the 34th AAAI Conference on Artificial Intelligence*; revised as [arXiv:2001.02889v5](https://arxiv.org/abs/2001.02889v5). Palo Alto, CA: Association for the Advancement of Artificial Intelligence.
- [26] ———. (2021). A topological perspective on causal inference. In Ranzato, M., editor. *Proceedings of the Thirty-Fifth Conference on Neural Information Processing Systems (NeurIPS)*. Vancouver: Curran Associates, Inc.
- [27] Ibeling, D., Icard, T., Mierzewski, K., & Mossé, M. (2022). Probing the quantitative–qualitative divide in probabilistic reasoning. Unpublished manuscript.

- [28] Kurucz, A. (2007). Combining modal logics. In van Benthem, J., Blackburn, P., and Wolter, F., editors. *Handbook of Modal Logic*. Amsterdam: Elsevier, pp. 869–924.
- [29] Luce, R. D. (1968). On the numerical representation of qualitative conditional probability. *The Annals of Mathematical Statistics*, **39**(2), 481–491.
- [30] Macintyre, A. & Wilkie, A. J. (1995). On the decidability of the real exponential field. In Odifreddi, P., editor. *Kreiseliana. About and Around Georg Kreisel*. Wellesley, MA: A. K. Peters, pp. 441–467.
- [31] Neyman, J. (1977). Frequentist probability and frequentist statistics. *Synthese*, **36**(1), 97–131.
- [32] Ognjanović, Z., Rašković, M., & Marković, Z. (2016). *Probability Logics: Probability-Based Formalization of Uncertain Reasoning*. Cham: Springer.
- [33] Pearl, J. (1995). Causal diagrams for empirical research. *Biometrika*, **82**(4), 669–710.
- [34] ———. (2009). *Causality*. Cambridge: Cambridge University Press.
- [35] Roth, D. (1996). On the hardness of approximate reasoning. *Artificial Intelligence*, **82**(1), 273–302.
- [36] Ruiz-Vanoye, J. A., Pérez-Ortega, J., Pazos R., R. A., Díaz-Parra, O., Frausto-Solís, J., Huacuja, H. J. F., Cruz-Reyes, L., & Martínez F., J. A. (2011). Survey of polynomial transformations between NP-complete problems. *Journal of Computational and Applied Mathematics*, **235**(16), 4851–4865.
- [37] Schachter, R. D. (1988). Probabilistic inference and influence diagrams. *Operations Research*, **36**, 589–605.
- [38] Schaefer, M. (2009). Complexity of some geometric and topological problems. In Eppstein, D., and Gansner, E. R., editors. *International Symposium on Graph Drawing*. Berlin: Springer, pp. 334–344.
- [39] ———. (2013). Realizability of graphs and linkages. In Pach, J., editor. *Thirty Essays on Geometric Graph Theory*. New York: Springer, pp. 461–482.
- [40] Scott, D., & Krauss, P. (1966). Assigning probabilities to logical formulas. *Studies in Logic and the Foundations of Mathematics*, **43**, 219–264.
- [41] Shalev-Shwartz, S., & Ben-David, S. (2014). *Understanding Machine Learning: From Theory to Algorithms*. New York: Cambridge University Press.
- [42] Speranski, S. O. (2017). Quantifying over events in probability logic: An introduction. *Mathematical Structures in Computer Science*, **27**(8), 1581–1600.
- [43] Spirtes, P., Glymour, C. N., & Scheines, R. (2000). *Causation, Prediction, and Search*. Cambridge, MA: MIT Press.
- [44] Stalnaker, R. C. (1996). Knowledge, belief and counterfactual reasoning in games. *Economics and Philosophy*, **12**, 133–163.
- [45] Suppes, P., & Zanotti, M. (1981). When are probabilistic explanations possible? *Synthese*, **48**, 191–199.

UNIVERSITY OF CALIFORNIA, BERKELEY

BERKELEY, CA 94720-2390, USA

E-mail: milan.mosse@berkeley.edu

STANFORD UNIVERSITY

STANFORD, CA 94305, USA

E-mail: duligur@stanford.edu

E-mail: icard@stanford.edu