

Research Article

Open Access

Dominik Janzing*

Causal versions of Maximum Entropy and Principle of Insufficient Reason

DOI: DOI, Received ...; revised ...; accepted ..

Abstract: The Principle of Insufficient Reason (PIR) assigns equal probabilities to each alternative of a random experiment whenever there is no reason to prefer one over the other. The Maximum Entropy Principle (MaxEnt) generalizes PIR to the case where statistical information like expectations are given. It is known that both principles result in paradoxical probability updates for joint distributions of cause and effect. This is because constraints on the conditional $P(\text{effect}|\text{cause})$ result in changes of $P(\text{cause})$ that assign higher probability to those values of the cause that offer more options for the effect, suggesting 'intentional behaviour'. Earlier work therefore suggested sequentially maximizing (conditional) entropy according to the causal order, but without further justification apart from plausibility on toy examples. We justify causal modifications of PIR and MaxEnt by separating constraints into restrictions for the cause and restrictions for the mechanism that generates the effect from the cause. We further sketch why Causal PIR also entails 'Information Geometric Causal Inference'. We briefly discuss problems of generalizing the causal version of MaxEnt to arbitrary causal DAGs.

Keywords: cause-effect problem, independence of mechanisms, arrow of time

MSC: 62B10, 62G99

1 Introduction

Understanding asymmetries between cause and effect has attracted researchers from the field of causal discovery particularly since two decades. One challenging problem motivated by the goal of understanding these asymmetries is to distinguish cause and effect from their bivariate distribution. This task cannot be solved by causal discovery methods that rely on conditional independences only [1, 2], but new approaches employ statistical properties other than conditional independences. They rely, for instance, on the additive noise assumption [3–5] or a generalization of the latter [6], or on asymmetries with respect to some notion of description complexity [7–9], or differences regarding regression error [10]. For an overview see also [11] and [12], but also [13] for a critical discussion of some ideas. Although distinction of cause and effect from purely observational data is still challenging, these approaches have stimulated discussions in various directions regarding inferential asymmetries of cause and effect. On the one hand, the relation to the arrow of time in physics has been described by [14] and [15]. On the other hand, it has been argued that the asymmetries entail implications for machine learning for scenarios where the causal direction is known [16, 17].

Here we describe an asymmetry between cause and effect with respect to how we assign priors to a set of possible outcomes of an experiment. Among the most prominent principles to assign priors is the 'Principle of Insufficient Reason' (PIR) and the Principle of Maximum Entropy (MaxEnt) [18]. PIR assigns uniform probabilities to a set of possible outcomes whenever the knowledge about the outcomes is invariant under permutations. MaxEnt, which generalizes PIR, chooses a prior that maximizes entropy subject to the known constraints. For case where the causal direction is known, [19] has argued that MaxEnt can result in implausible distributions and more natural joint distributions result from a sequential maximization: first maximize entropy of the cause subject to all constraints relevant for the latter, and then the conditional entropy of the effect, given the cause, subject to all remaining constraints. However, the arguments of [19] were merely based on intuition without further justification.

*Corresponding Author: Dominik Janzing: Amazon Research, Tübingen, Germany; E-mail: janzind@amazon.com

On a related note, [20] proposes a 'maximum causal entropy principle' for a scenario with two interacting processes X_t, Y_t where X_t is known and Y_t is inferred from its own past and from X_t and its past via a sequential maximization of conditional entropy. [20] justifies the sequential update by arguing that constraints that involve future observations should be ignored *at that respective point in time*. In the appendix we argue that this justification is not sufficient for our purpose.

The goal of this paper is to derive the sequential maximum entropy update rule proposed by [19] from principles that we consider slightly more basic. To this end, Section 2 discusses a simple scenario suggesting that also PIR requires the same modification as MaxEnt. Section 3 tries to justify *Causal PIR* from a deeper principle of independent mechanisms, but also raises questions that remain open in this regard. Section 4 derives the causal version of MaxEnt by [19] from applying Causal PIR to *empirical* distributions. Section 5 describes some problems of generalizing Causal MaxEnt to arbitrary causal DAGs. Section 6 shows that Information Geometric Causal Inference [21] can be derived from Causal PIR similar to Causal MaxEnt.

Proposing new practical inference rules is beyond the scope of this paper. Instead, it aims at better understanding relations between asymmetries of cause versus effect described earlier.

2 Causal Version of PIR

2.1 Standard PIR

The 'Principle of Insufficient Reason (PIR)', also called 'Laplace's Principle of Insufficient Reason' or 'Principle of Indifference' [18], states that in the absence of any relevant evidence, agents should distribute their credence (or 'degrees of belief') equally among all the possible outcomes under consideration. More explicitly, PIR advises to consider all possible alternatives in a random experiment equally likely. For the simple example where we know that one of n urns contains a ball, PIR considers each of the urns as an equally likely location and assigns $P(j) = 1/n$ to each case $j = 1, \dots, n$. For a discussion of justifications of PIR we refer to [22], where also the relation to MaxEnt is discussed in detail.

For our purpose, it is also instructive to rephrase PIR by stating that it advises the uniform prior whenever there is no evidence that breaks the *symmetry* between the alternative outcomes. In a way, PIR then gets a circular structure because any argument against the uniform prior implicitly raises doubts about the symmetry of the problem (obviously, the uniform distribution is the only one that is symmetric under permutation of the alternatives). One insight of our discussion below will be that a reasonable use of PIR is not symmetric with respect to interchanging cause and effect. We are agnostic about whether one should consider this merely as an advice on how to *properly apply PIR in a cause-effect scenario* or as a *causal modification of PIR*.

2.2 Motivating Causal PIR for a simple mechanical device

Consider the mechanical device depicted in Figure 1. It consists of a system with channels having three different entries (top of the figure) and three exits (bottom). The first entry splits into two different channels, while the second and the third entry lead to the same exit. Let us label the three entries with the variable X attaining the values 1, 2, 3, while Y labels the exits 1, 2, 3.

Assume we know that a ball enters one of the three entrances at the top. In absence of any further information, we would consider all three options as equally likely, that is $P(X = 1) = P(X = 2) = P(X = 3) = 1/3$, in agreement with PIR. When rolling through the channel, the ball will take one of the three exits. Whenever it entered at the entrance 2 or 3, it can only take the exit $Y = 1$ due to the topology of the channels. In case it entered at entrance 1, it has the two options later, namely exits $Y = 2$ or $Y = 3$. We now apply PIR for the *conditional* distribution of Y given X and assume that both alternatives are equally likely. The scenario thus yields the joint probabilities shown




Figure 1. A ball enters our mechanical device from the top. Without additional information, we would consider all three options ($X = 1, X = 2, X = 3$) equally likely, that is, assign the probability $1/3$ to them. This results in probability $2/3$ for $Y = 1$ and probability $1/6$ each for $X = 2$ and $X = 3$.

in the table in Figure 2, left. This distribution is clearly asymmetric with respect to X and Y although the topology of the channels is symmetric. Assuming that the ball enters from the bottom, that is, Y labels the entrances and X the exits, thus induces the joint distribution in the table in Figure 2, right, which is obtained by swapping the roles of X and Y .

Lead by our intuition, we have applied PIR twice: first for X , and then for Y , given X . However, the simplicity of the scenario blurs a non-trivial step in this way of reasoning, namely that the experiment is not symmetric with respect to time inversion, or, which is equivalent here, with respect to swapping cause and effect. Here, our asymmetry of reasoning is implicitly based on a belief about the difference between cause vs. effect and past vs. future.

Note that our mechanical toy example does not describe the typical scenario of cause-effect inference since it is uncommon to know the mechanism that relates cause and effect, that is, only the direction is unknown. Typically, we are given observations from X, Y instead of knowledge on the mechanisms. Yet the example is helpful to motivate Causal PIR, which is later used to motivate Causal MaxEnt, which, in turn, is relevant for more realistic inference scenarios.

2.3 Fallback to standard PIR when causal direction is unknown

To elaborate on this, note that the topology of the channel allows 4 different x, y -pairs. Without knowing whether the ball enters from the top or the bottom, PIR lets us assign equal probabilities to each of them since the device is symmetric once the knowledge of the direction of the motion is lost. Obviously, the symmetry of the problem now results in the distribution shown in the table in Figure 2, middle. Note that this distribution may not only be natural when we are *agnostic* about the causal direction, but also if neither of the causal directions is true and the relation between X and Y is due to a common cause. Although the following scenario may seem less natural than the first two ones with X or Y as cause, we mention it to cover also the common cause scenario. Assume that the ball drops from the sky into one of the channels and lies there at some point at rest. If it lies in the regions $X = 2, 3$ or $Y = 2, 3$, its position already defines a unique (x, y) -pair since these values can only occur together with a unique value of Y or X , respectively. In the case where it lies in the regions $X = 1$ or $Y = 1$, we push it towards the branching point to generate the corresponding random value of Y or X , respectively. This way, we have again generated a scenario in which we have no reason to prefer any of the 4 possible (x, y) -pairs over the other. One can argue that the causal structure of this scenario is the DAG shown in Figure 2, middle, where some 'big' unobserved variable Z affects both X and Y , where Z contains position and momentum of the ball and the noise which determines the branching process.




Figure 2. Left: Joint probabilities when the ball enters from the top (X is the cause). Middle: probabilities when the ball enters from the sky (common cause). Right: when the ball enters from the bottom (Y is the cause).

2.4 Paradoxes with standard PIR

As the table in Figure 2, middle, shows, assigning equal probabilities to all 4 possible cases result in higher probabilities to those values of the cause that admit more options for the effect – which suggests ‘intentional behaviour’. Note, however, that the latter interpretation is prone of confusing *ontic* and *epistemic* perspectives: whenever the restriction to these 4 alternatives comes from our knowledge *about the underlying mechanism connecting the cause X with the effect Y* , it is indeed irrational to consider x -values more likely for which there exists a larger number of possible y -values later. However, if we know, for some other reason, that (x, y) is one of the above 4 cases (e.g. because someone told us without telling us (x, y)), there is nothing wrong with updating our *subjective prior* for X in the way resulting from the uniform distribution over the 4 possible pairs. After all, this Bayesian update entails no statement on the underlying causal mechanism. This distinction will be further discussed in Section 3, where we also mention open problems regarding ontic versus epistemic interpretation of constraints.

Similar paradoxes with standard PIR have been described by [23, 24] in a critical discussion of MaxEnt. He described a scenario that he called ‘Pearl’s Puzzle’¹, which we briefly sketch. Assume three individuals A, B, C are invited to a party but don’t know who will be joining. Further assume we consider, a priori, all 8 possible combinations equally likely. In addition, we know that A, B decide independently of each other and of C whether they join, but C will call the host to ask whether both A and B have accepted the invitation and stay at home in this case to avoid seeing both of them together. After accounting for this extra information (excluding the case where A, B, C occur), we are left with 7 remaining combinations, which we would assign equal probabilities to. According to such an update, the joint distribution of A, B has changed after accounting for the information that C ’s decision depends on A and B . Phrasing it in causal terms, the puzzle reads as follows: A, B are the causes and C ’s behaviour their effect. Learning about how C ’s decision depends on A and B actually changes the belief about the mechanism according to which the effect depends on its causes. It is disturbing that an update on this mechanism affects the distribution of the causes (one can also show, which Pearl describes as the main puzzle, that A and B even become dependent by this update).

We will later elaborate on this in the context of the so-called Principle of Independent Mechanisms [11], since Hunter’s and Pearl’s discussions are already lead by such an independence assumption.

To conclude with ‘Pearl’s puzzle’ we briefly sketch how it gets resolved by a sequential use of PIR: since A and B are the causes, we assign a uniform prior over all 4 possible truth values. Afterwards, we assign a uniform prior over all remaining options for C : whenever A and B are coming, C stays at home with probability 1, for all other cases he/she would decide to come with probability 1/2. By construction, whether or not C is coming, is irrelevant for A and B .

¹ Hunter writes: “The example was given by personal communication and has been floating around the uncertain reasoning community for sometime. Pearl informs me that the example was discovered by Norman Dalkey but was first taken as a counterexample to MaxEnt by Pearl”




Figure 3. For the gray path with unknown direction, it is more likely that the ball entered from the top than from the bottom according to Causal PIR.

2.5 General definition of Causal PIR

The way we defined the joint distribution for the mechanical device can be described by the following principle, which also solved the above ‘puzzle’:

Definition 1 (Causal PIR). *Let X and Y be cause and effect with values in finite sets \mathcal{X} and \mathcal{Y} , respectively. If the only knowledge about an observation (x, y) is that it lies in some subset $S \subset \mathcal{X} \times \mathcal{Y}$, Causal PIR assigns uniform distribution to all possible x , for which there exists an y such that $(x, y) \in S$. Then causal PIR assigns the uniform prior over all remaining options for y , given x (that is, all y for which $(x, y) \in S$).*

A priori, we have introduced Causal PIR only as a principle for constructing a prior when the causal direction is known. Conversely, one can certainly use its asymmetry to infer the causal direction by preferring the one with larger likelihood:

Definition 2 (Causal PIR based cause-effect inference). *Given an observation (x, y) generated by either the causal structure $X \rightarrow Y$ or $Y \rightarrow X$. Infer that the true causal direction is the one for which (x, y) has larger likelihood according to the Causal PIR prior.*

Observing, for instance, the path in Figure 3, we thus infer that the ball entered from the top rather than from the bottom: we obtain likelihood $1/3$ for the former versus $1/6$ for the latter.

In a more informal way, we state Causal PIR as follows:

Postulate 1 (informal version of Causal PIR). *Prefer causal models (i.e. directions, DAGs, structural equations) for which*

- *the observed effect admits many values of the cause and*
- *the observed cause admits few values of the effect.*

3 Independent mechanism update

This section and Section 4 repeatedly refer to the Principle of Independent Mechanisms (IM), which we briefly introduce for the special case of a cause-effect pair. A priori, IM is an informal principle stating that, for an unconfounded cause-effect relation, there should be two independent mechanisms in place, one that generates the cause and one that generates the effect from the cause (see [11], Section 2.1, for an overview and discussion of its different aspects). IM has been used as foundational justification for cause-effect inference. One formalization of IM in the literature is the Algorithmic Independence of Conditionals (AIC) by [7] and [25], stating that the shortest description of $P_{X,Y}$

is given by separate descriptions of P_X and $P_{Y|X}$.² This version will be relevant in Section 4, while the following subsection will interpret IM in the sense of decomposing the constraint S according to one that refers to the cause and one referring to the relation between cause and effect.

3.1 Constraints on the cause and constraints on functions

We will now describe a principle that justifies Causal PIR in Definition 1. To this end, let $\mathcal{F} := \mathcal{Y}^{\mathcal{X}}$ denote the set of functions $f : \mathcal{X} \rightarrow \mathcal{Y}$ and define the formal random variable F attaining values $f \in \mathcal{F}$. Note that \mathcal{F} can be represented as the k -fold cartesian product of \mathcal{Y} with $k := |\mathcal{X}|$, whose components are indexed by $x \in \mathcal{X}$. In other words, a function f is represented by the k -tuple $(f(x^1), \dots, f(x^k))$ if $\mathcal{X} = \{x^1, \dots, x^k\}$. Accordingly, a distribution on \mathcal{F} is a joint distribution on this cartesian product. Its marginal distribution on component x describes the conditional probability $P_{Y|X=x}$ (see, for instance, Section 3.4 of [11]).

Let $P_{\mathcal{F}}$ be the uniform distribution³ on \mathcal{F} . Fortunately, the uniform distribution has product structure over the components of the cartesian product, which renders $P_{\mathcal{F}}$ particularly easy to deal with. Further, for every input x , every output y is equally likely. In other words, the uniform prior over all functions induces conditional distributions $P_{Y|X=x}$ that are uniform for all x . After assuming also a uniform prior P_X , we have thus obtained a uniform prior $P_{X,Y} = P_X P_{Y|X}$ over all $|\mathcal{X}||\mathcal{Y}|$ combinations, which are $9 = 3 \cdot 3$ in the above example.

After having defined our prior on \mathcal{X} and \mathcal{F} , let us obtain the additional information that the entire device only generates (x, y) -pairs in some set $S \subset \mathcal{X} \times \mathcal{Y}$. For the example above, these are the 4 combinations shown in the tables in Figure 2. We now assume that the constraint S is the result from two independent mechanisms, one for \mathcal{X} and one for \mathcal{F} :

Postulate 2 (separation of constraints). *Given the constraint $(x, y) \in S$ for a cause-effect pair (X, Y) , we assume, by default, that this constraint is enforced by two separate mechanisms. First, there is a mechanism that enforces all x to be in the set*

$$S_X := \{x \in \mathcal{X} \mid (x, y) \in S \text{ for some } y \in \mathcal{Y}\}.$$

Second, there is a mechanism that enforces functions to be in the set

$$S_F := \{f \in \mathcal{F} \mid (x, f(x)) \in S \text{ for all } x \in S_X\}.$$

To better understand the postulate it helps to say what kind of mechanisms it *excludes*: Imagine an agent who chooses functions f that violate $(x, f(x)) \in S$ for *some* inputs $x \in S_X$, but always makes sure that these functions are only combined with inputs x for which the constraints are satisfied. In other words, the agent ensures $(x, f(x)) \in S$ by combining x and f in a smart way. In this case, we would say that the mechanism choosing x and the mechanism choosing f are *dependent*. In Subsection 3.2 we will discuss in what sense this would violate the Principle of Independent Mechanisms.

The above restrictions for X and F together generate the restriction S . Note that the above separation of S into (S_X, S_F) entails *minimal commitment* on both components x and f (while still preserving independence) in the following sense. First, it is obvious that no proper superset $\tilde{S}_X \supset S_X$ guarantees that $(x, y) \in S$, regardless of the constraints for the functions. Second, no larger set $\tilde{S}_F \supset S_F$ guarantees $(x, y) \in S$ unless we require x -values and functions f to respect *joint* constraints.

² Further concrete conceptualizations of IM are: (1) the hypothesis that unlabelled data in semi-supervised learning is only helpful for a so-called ‘anticausal prediction’ scenario (where the cause is predicted from the effect), while it is pointless for ‘causal prediction’ (when the effect is predicted from the cause), see [16]. Further, IM has been formalized as (2) uncorrelatedness of the logarithmic slope of the function connecting X and Y with the density of X in ‘Information Geometric Causal Inference’ [21], and (3) uncorrelatedness of the (absolute squared) transfer function connecting X and Y with the power spectrum of X when X and Y are time series connected by a linear filter (‘Spectral Independence Criterion’ by [26]).

³ Note that Hunter’s solution of ‘Pearl’s puzzle’ [24] also uses an update of a distribution over functions (‘probability measures over counterfactuals’), but is based on the assumption that the constraint is known to refer to the function only, while our scenario describes a constraint for (x, y) for which it is not a priori known what it tells us about the function.

For our mechanical device above, the constraint for X reads that there are 3 possible entries $X = 1, 2, 3$. While this constraint is trivial since our set \mathcal{X} contains only these 3 values, one could also think of a set \mathcal{X} that is a priori larger until our information on S restricts the options for X to the subset S_X consisting of these 3 values. The constraints on F that we conclude from the joint constraint S consists in excluding all functions that map $X = 1$ to y -values other than $Y = 2, 3$ and $X = 2, 3$ to values other than 1. Nevertheless, Postulate 2 is less innocent than our toy example suggests. We will therefore further discuss its justification in Subsection 3.2.

We now obtain the following technically simple result, which we phrase as a theorem since it considers Causal PIR as an implication of the more basic Postulate 2:

Theorem 1 (Causal PIR from independent mechanism update). *Let P_X and $P_{\mathcal{F}}$ be uniform distributions on \mathcal{X} and \mathcal{F} , respectively. Then the conditional distribution of P_X^S is the uniform distribution over all S_X . Further, for every $x \in S_X$, the conditional $P_{Y|X=x}^S$ resulting from the conditional distribution $P_{\mathcal{F}}^S$ is the uniform distribution over all y for which $(x, y) \in S$.*

Proof. The first statement is obvious. To show that the conditional is also uniform over all remaining options, we represent each function f as the k -tuple of y -values

$$(y^1, \dots, y^k) := (f(x^1), f(x^2), \dots, f(x^k)),$$

where x^1, \dots, x^k denote the elements of S_X . The uniform prior over all function thus amounts to the uniform prior over all k -tuples $(y^1, \dots, y^k) \in \mathcal{Y}^k$. Since the uniform distribution over a cartesian product factorizes over its components, we can perform the update independently for each x^j and obtain a uniform distribution over all y for which $(x^j, y) \in S$. \square

3.2 Justifying separation of constraints

The remaining subsection is devoted to the justification of Postulate 2. It needs to be informal because it is a discussion on beliefs about the world rather than insights from statistics or any other branches of mathematics. Further, it can be seen as ‘abstract physics’ in which the ‘hardware’ of the underlying processes is unspecified. We also briefly mention relations to the thermodynamic arrow of time and thus reach a domain that goes beyond the scope of this paper. Accordingly, further justification of Postulate 2 could also raise questions of theoretical physics. However, the main focus on this Subsection is the question *which implicit further assumptions are made* when Postulate 2 is said to be entailed by the Principle of Independence Mechanisms.

3.2.1 Constraints from knowledge versus constraints from mechanisms

The first distinction we need to make regarding the constraint $(x, y) \in S$ is whether we assume that there is a *mechanism* that generates pairs in S or whether we know that a particular experiment resulted in an (x, y) -pair in S *by chance* (recall our remarks regarding *ontic* versus *epistemic* perspectives in Subsection 2.4). In the second case, Postulate 2 does not make sense: if $(x, y) \in S$ is not the result of mechanisms that enforces (x, y) to lie in S , it is pointless to postulate *separate* mechanisms. In this case, the further argument resulting in Theorem 1 breaks down: There is, a priori, no reason why *our knowledge* about X and F could not render them dependent,⁴ although the Principle of Independence Mechanisms states that the true mechanisms contain no information about each other. –Note that [7] formalizes *independence* via *algorithmic* information, which is an ontic perspective since it relies on the description length of *known* mechanisms.

One can certainly also justify an *epistemic* ‘Principle of Independence Mechanisms’ stating that our prior about cause-effect pairs should factorize between the mechanism generating the cause and the mechanism generating the

⁴ In general, knowledge about a pair of events can result in a subjective prior that renders them dependent although they are not causally related, as emphasized also by [18].

effect from the cause, but the factorization breaks down after joint observations from cause and effect are available. Although this raises doubts about Postulate 2, we now discuss what kind of inductive bias provides further support.

3.2.2 Bias for ontic interpretation

Let us now describe a scenario in which knowing $(x, y) \in S$ does provide evidence *for the presence of a mechanism* that enforces or at least supports outcomes in S . To this end, assume that the sets \mathcal{X} and \mathcal{Y} are huge (e.g. binary words of length n with $n \geq 100$). Further, assume that S is a *strong* constraint in the sense that it allows only a small fraction of possible outcomes, that is, $|S| \leq (|\mathcal{X}| \times |\mathcal{Y}|)/k$ for some huge number k . Assuming, a priori, a uniform distribution $P_{X,Y}$ on $|\mathcal{X}| \times |\mathcal{Y}|$, that is, we have $P\{(X, Y) \in S\} \leq \frac{1}{k}$. Given some fixed S with this property, we would certainly argue that an observation in S is unlikely without a mechanism that increases the probability of outcomes in S . At first glance, it seems that such a conclusion is only possible if S has been defined *prior* to the experiment. However, it still holds when we can identify a set S *after* observing (x, y) provided that S has *low description length* (here we formalize description length via Kolmogorov complexity [27], that is, let $K(S)$ denote the length of the shortest self-delimiting program that decides whether any pair (x, y) is in S). With these assumptions, it is unlikely to obtain an outcome (x, y) in *any* such tiny set S with low complexity. To see this, let U be the union of all sets S with $K(S) \leq \ell$ and $|S| \leq k$. Since the number of programs of length at most ℓ is at most 2^ℓ [27], the probability for obtaining a result in any of these low complexity sets can be bounded from above as follows:

$$P\{(X, Y) \in U\} \leq \frac{2^\ell}{k}. \quad (1)$$

In case the right hand side of (1) is still significantly smaller than 1 we assume that observing $(x, y) \in S$ indicates the presence of a mechanism that increases the probability of S (compared to the uniform distribution we started with). We phrase this insight as an informal postulate:

Postulate 3 (bias towards mechanisms vs state of knowledge). *Given a system with finite ‘state space’ \mathcal{W} , then the information that the actual state w lies in some set S with low complexity (in the sense that $|S| \cdot 2^{K(S)} \ll |\mathcal{W}|$) is considered as strong evidence for the presence of a mechanism that increases the likelihood of states in S .*

The constraints we will discuss later in the context of MaxEnt will typically be of this type: constraints that describe empirical means of simple functions like polynomials of low order have low complexity (provided that the constants involved have short descriptions), and restrict the combinations of outcomes by huge factors. For the same reasons, typical constraints in thermodynamics also result from *mechanisms*: observing that all particles of a gas are located within a certain volume V can only be explained by a mechanism (e.g. a wall) that confines them to V , rather than being just a coincidence.⁵

3.2.3 Is the constraint S tight?

Together with the bias for an ontic interpretation of constraints, we are now getting slightly closer to deriving Postulate 2 from IM (beyond the few comments made right after stating it). We now assume, for simplicity, that the constraint $(x, y) \in S$ is due to a mechanism that forces all pairs to lie in S (although Postulate 3 is weaker in the sense that it only assumes a mechanism that increases the likelihood of S). The question we are facing is whether S is *tight* in the sense that *all* pairs in S will occur after sufficiently many repetitions of the same experiment. We then need to assume that S originates from separate constraints for X and F because otherwise we would need a mechanism that controls X and F *jointly* by varying them in a way that enforces $(x, f(x)) \in S$, in contradiction to the independence of mechanisms, as sketched after Postulate 2.

⁵ In general, constraints on *macroscopic* variables have negligible description length compared to the typical complexity of the microscopic state of a many-particle system, as also argued by [28].

For the case where S is not tight and the mechanism generates only (x, y) pairs in $S' \subset S$ (but we don't know S') we still choose the update according to Postulate 2 because this is the only possible choice for constraints on X and F that doesn't commit beyond the information we have.

We summarize that assuming that a constraint S arises from independent constraints for \mathcal{X} and \mathcal{F} is our inductive bias, which can be justified under appropriate conditions.

4 From MaxEnt to Causal MaxEnt

4.1 Wallis' argument for MaxEnt

Inferring underdetermined probability distributions by maximizing entropy subject to the available information is a well-established principle in machine learning and statistics, see e.g. [29–31]. The usual formal setting reads:

4.1.1 Accounting for linear constraints

Let us, for simplicity, assume that X is a variable that attains values in some finite set \mathcal{X} . Assume the only information available on P_X is given by the expectations

$$\sum p(x) f_j(X) = c_j, \quad \text{with } c_j \in \mathbb{R}, \quad (2)$$

where f_j are measurable functions. According to MaxEnt we would then choose the unique distribution maximizing the Shannon entropy⁶

$$H(X) := - \sum_x p(x) \log p(x).$$

subject to the constraints (2), which yields

$$p(x) = e^{-\lambda_j f_j(x) - \mu}, \quad (3)$$

with appropriate Lagrange multipliers λ_j, μ .

While distributions that result from MaxEnt often appear intuitively 'natural', or 'smooth'⁷, there is an ongoing debate about how to justify (3) as a *rational* guess [33–35].

[36] stated Postulates that 'consistent' rules for updating a distribution after new information comes in should satisfy, [35] criticized the approach as suffering from hidden implicit assumptions that go beyond what [35] would call 'consistency' requirements. We will therefore prefer the so-called Wallis' derivation (see [18], Section 11.4), which we briefly sketch: Consider an experiment with n draws from the finite probability space $\mathcal{X} = \{x^1, \dots, x^k\}$, and n_1, \dots, n_k with $\sum_j n_j = n$ denotes the number of occurrences of x^j . By elementary combinatorics, the number of combinations for these frequencies reads

$$\#(n_1, \dots, n_k) = \frac{n!}{n_1! n_2! \dots n_k!}. \quad (4)$$

Using Stirling's approximation one can easily show that

$$\frac{1}{n} \log \#(n_1, \dots, n_k) = - \sum_j \frac{n_j}{n} \log \frac{n_j}{n} + O(\log n/n). \quad (5)$$

⁶ For continuous variables, one typically replaces Shannon entropy with differential Shannon entropy [32] $H(X) := - \int p(x) \log p(x) dx$. Since the latter is not invariant with respect to re-parametrization, one should then rather consider minimization of relative entropy to a given prior distribution.

⁷ Since any distribution maximizes the entropy subject to appropriate constraints (just choose $f(x) := \log q(x)$ with appropriate constant c), this is certainly a result of the *type of constraints* that typically occur in applications, e.g., if only first and second moments of a distribution are known.

Hence, the number of realizations can be estimated via the entropy of the relative frequencies. Thus, the MaxEnt distribution is the distribution for which the corresponding relative frequencies maximize the number of realizations in the limit of $n \rightarrow \infty$.

Further, one can show that for large enough n , the overwhelming majority of n -tuples satisfying the constraints show empirical distributions that are close to the MaxEnt distribution. Hence, a prior on \mathcal{X}^n that assigns equal probability to each n -tuple, results, after accounting for the constraints, in a posterior that is essentially supported by empirical distributions close to the unique MaxEnt distribution. In this sense, MaxEnt can also be seen as an implication of PIR (when applied to empirical distributions), although MaxEnt is more general from the formal point of view.

4.2 Causal MaxEnt from Causal PIR

We start with motivating Causal MaxEnt in the same way as it is done by [19]. Assume we are given a continuous variable X as cause and a binary variable Y as effect. Let the only information about the joint distribution $P_{X,Y}$ be given by the first and second moments $\mathbf{E}[X]$, $\mathbf{E}[X^2]$, $\mathbf{E}[XY]$, $\mathbf{E}[Y]$, $\mathbf{E}[Y^2]$. One can easily verify that the MaxEnt distribution is a bivariate mixture of Gaussians, where the cases $Y = 0, 1$ correspond to the two mixture components. [19] argues that this distribution would be a plausible distribution if Y was the cause and X the effect, while it is not plausible that the *cause* becomes bimodal just because it has an influence on a binary variable. If one, instead, first maximizes $H(X)$ subject to the constraints $\mathbf{E}[X]$, $\mathbf{E}[X^2]$ and then $H(Y|X)$ subject to the remaining constraints $\mathbf{E}[XY]$, $\mathbf{E}[Y]$, $\mathbf{E}[Y^2]$, the marginal distribution P_X becomes a single Gaussian and $P_{Y=1|X}$ a sigmoid function where the probability for $Y = 1$ smoothly increases or decreases with X , which [19] considers plausible for the causal direction $X \rightarrow Y$. Formally, they have postulated the following principle:

Definition 3 (Causal MaxEnt). *Given some linear constraints for $P_{X,Y}$ for the cause effect pair (X, Y) . Infer the bivariate distribution by first maximizing $H(X)$ subject to all constraints for P_X (entailed by the joint constraints). Then, maximize $H(Y|X)$ subject to the joint constraints.*

[37] show that usual MaxEnt violates the algorithmic independence of P_X and $P_{Y|X}$.

The proof is based on the observation that MaxEnt can result in a joint distribution whose marginal P_X cannot be defined by a separate constraint with simple description. Instead, its simplest description may be ‘the marginal distribution resulting from MaxEnt for the joint constraint’. For the example above with binary X and real-valued Y with second order constraints, P_X is a mixture of two Gaussians, and thus already contains the full information about the joint distribution $P_{X,Y}$. Despite describing this problem of MaxEnt, [37] does not show that Causal MaxEnt is the right replacement of MaxEnt.

We now show that this sequential probability update is a result of Causal PIR when applied to empirical distributions.⁸ Assume we are given ℓ constraints of the form

$$\mathbf{E}[f_j(X, Y)] = c_j \text{ for } j = 1, \dots, \ell. \quad (6)$$

Let us now interpret (6) as constraints for the empirical distribution after n draws. For each pair (\mathbf{x}, \mathbf{y}) of n -tuples $\mathbf{x} := (x_1, \dots, x_n)$ and $\mathbf{y} := (y_1, \dots, y_n)$ we denote by $\mathbf{E}_{(\mathbf{x}, \mathbf{y})}$ the expectation induced by the corresponding empirical distribution of (X, Y) . Finally, we define

$$S := \{(\mathbf{x}, \mathbf{y}) \in \mathcal{X}^n \times \mathcal{Y}^n \text{ with } |\mathbf{E}_{(\mathbf{x}, \mathbf{y})}[f_j(X, Y)] - c_j| \leq \epsilon\}, \quad (7)$$

with some arbitrarily small $\epsilon > 0$, which defines a relaxation of (6) to ensure feasibility for sufficiently large n .

⁸ For readers with interest in physics we note that the independent uniform distributions on \mathcal{X}^n and \mathcal{F}^n our arguments in Subsection 3.1 relied on can be seen as a result of independent mixing processes, the first one mixes the state of the cause and the second one the conditional state of the effect. Accordingly, Ref. [14] has described a physical toy model for a cause-effect relation where this sequential entropy maximization follows from mixing processes that first affect the cause and then the interaction that generates the effect from the cause.

Following our separation of constraints in Postulate 2, we now define $S_{\mathbf{x}}$ as the set of n -tuples \mathbf{x} for which there exists an n -tuple \mathbf{y} such that $(\mathbf{x}, \mathbf{y}) \in S$. Again, Causal PIR tells us to put a uniform prior on $S_{\mathbf{x}}$. Following Subsection 4.1, the overwhelming majority of n -tuples \mathbf{x} in $S_{\mathbf{x}}$ are close to the distribution P_X that maximizes $H(X)$ subject to (6) being feasible for Y . For any $\mathbf{x} \in S_{\mathbf{x}}$ let $S_{\mathbf{x}}$ denote the set of n -tuples \mathbf{y} such that $(\mathbf{x}, \mathbf{y}) \in S$. According to causal PIR, we put a uniform prior on $S_{\mathbf{x}}$. We will again use (5) to derive the conditional empirical distribution that is induced by the majority of the $\mathbf{y} \in S_{\mathbf{x}}$.

To this end, for any $\mathbf{x} \in S_{\mathbf{x}}$ let $n_1^{\mathbf{x}}, \dots, n_k^{\mathbf{x}}$ denote the number of occurrences of the k different elements of \mathcal{X} . Further, for any $\mathbf{y} \in S_{\mathbf{x}}$, let n_j^i be the number of occurrences of the element (i, j) in $\mathcal{X} \times \mathcal{Y}$. For any collection (n_j^i) and any fixed \mathbf{x} , the number of different \mathbf{y} is given by

$$\#(n_1^1, \dots, n_l^k) = \prod_{i=1}^k \frac{n_i^{\mathbf{x}}!}{n_1^i! n_2^i! \dots n_k^i!}, \quad (8)$$

since we need to apply (4) for each element of \mathcal{X} and sample size $n_1^{\mathbf{x}}$. Using the same arguments as for the derivation of (5) we obtain

$$\frac{1}{n} \log \#(n_1^1, \dots, n_l^k) = - \sum_i \frac{n_i^{\mathbf{x}}}{n} \sum_j \frac{n_j^i}{n_i^{\mathbf{x}}} \log \frac{n_j^i}{n_i^{\mathbf{x}}}. \quad (9)$$

$$+O(\log n/n). \quad (10)$$

Recalling that the conditional entropy of Y given X for any probability mass function $p(x, y)$ reads [32]

$$H(Y|X) = - \sum_{x,y} p(x)p(y|x) \log p(y|x),$$

we observe that (9) is the conditional entropy of the empirical distribution. Accordingly, we conclude that, for any fixed \mathbf{x} , the overwhelming majority of n -tuples \mathbf{y} in $S_{\mathbf{x}}$ are those whose empirical distributions are close to the distribution maximizing conditional entropy subject to (6).

The above arguments show that first putting a uniform prior on $S_{\mathbf{x}}$ and then, for fixed \mathbf{x} , a uniform prior on $S_{\mathbf{x}}$ yields a joint distribution on $\mathcal{X}^n \times \mathcal{Y}^n$ that is strongly concentrated on the set of (\mathbf{x}, \mathbf{y}) -pairs whose empirical distribution is given by Causal MaxEnt. In contrast, classical MaxEnt would provide the most likely empirical distribution only if we put uniform prior on S , that is, if we use standard PIR.

5 Generalization of Causal MaxEnt to arbitrary DAGs

Given a causally sufficient set of N variables X_1, \dots, X_N , causally linked by the directed acyclic graph (DAG) G , the causal Markov condition [1, 2] implies that the joint distribution factorizes according to

$$P_{X_1, \dots, X_N} = \prod_{j=1}^N P_{X_j | PA_j}, \quad (11)$$

where $P_{X_j | PA_j}$ denotes the conditional distribution of X_j , given its parents in G . If we are given multivariate constraints of the form

$$\mathbf{E}[f_j(X_1, \dots, X_N)] = c_j, \quad (12)$$

the arguments from Section 4 suggest to obtain the conditionals $P_{X_j | PA_j}$ by sequentially maximizing conditional entropy $H(X_j | PA_j)$ according to an ordering that is consistent with G , a procedure already proposed by [19]. Since we construct the joint distribution as the product of the conditionals $P_{X_j | PA_j}$, it is Markov relative to G by construction. This seems to overcome a problem with classical MaxEnt: maximizing the joint entropy subject to (12) does not necessarily result in an Markovian distribution, while maximizing entropy subject to (12) and (11) is not a convex optimization problem and thus need not have a unique solution (as shown below for a toy example).

Before describing problems with Causal MaxEnt for general DAGs, let us first consider an example where it makes sense. Let X_j be binary variables connected by the causal structure

$$X_1 \rightarrow X_2 \rightarrow \cdots \rightarrow X_N. \quad (13)$$

Assume now we are given a constraint saying ' $X_j = 0$ implies $X_{j+1} = 0$ ' for $j = 1, \dots, N - 1$. Intuitively, this corresponds to a mechanism that appends 0 or 1 to any binary word ending with 1, but it appends only 0 to words ending with 0. In other words, it rules out any binary word (x_1, \dots, x_N) containing the substring 01. Classical MaxEnt would thus result in the uniform distribution over the $N + 1$ binary words $0 \dots 0, 10 \dots 0, 110 \dots 0, \dots, 11 \dots 1$. Causal MaxEnt yields $X_1 = 1$ with probability $1/2$, and all other X_j attain 1 with probability $1/2$ if their predecessor is 1. Thus, the binary words occur with probability $1/2^1, 1/2^2, \dots, 1/2^N, 1/2^N$, a distribution with much lower entropy. In this sense, Causal MaxEnt is more conclusive since it results in smaller uncertainty about the resulting joint distribution after leveraging the causal information.

However, sequential entropy maximization raises the following two problems (ignored by [19]) in case the DAG is not complete:⁹ First, the ordering of nodes is not necessarily unique. Second, sequentially maximizing entropy may render the constraints (12) infeasible, as shown by the following toy example with a DAG G with two variables and no edge. Consider the binary variables X_1, X_2 with values ± 1 . The Markov condition implies the factorization

$$P_{X_1, X_2} = P_{X_1} P_{X_2}. \quad (14)$$

Assume we are given the constraint

$$\mathbf{E}[X_1 X_2] = 1. \quad (15)$$

To implement Causal MaxEnt, let us choose the ordering X_1, X_2 . We observe that (15) entails no restriction for the marginal distributions of X_1 , and thus maximizing $H(X_1)$ yields $P(X_1 = 1) = 1/2$. Then Causal MaxEnt advises us to maximize the entropy of X_2 , given its parents in G (which is the empty set), subject to (15). However, there is no P_{X_2} such that $P_{X_1} P_{X_2}$ satisfies (15), after we have already maximized the entropy of X_1 . To satisfy the constraint, we X_2 need to *depend* on X_1 , which violates the Markov condition. The only joint distributions satisfying constraint (15) and Markov condition (14) are point measures on $(1, 1)$ or $(-1, -1)$, respectively. These are the two solutions of the non-convex problem of maximizing entropy subject to (15) *and* (14). By deciding for one of the solutions we would commit beyond the known constraint (15). If (15) results from independent mechanisms for X_1 and X_2 , it could be that there are either two independent mechanisms generating only the value 1 for both variables, or independent mechanisms generating only -1 for both ones, we just do not know which scenario is the true one. In other words, (15) represents our *knowledge* on the mechanisms, while the mechanisms themselves respect tighter constraints, namely $(X_1, X_2) = (1, 1)$ or $(X_1, X_2) = (-1, -1)$, depending on the scenario. Hence we have an example for the case where constraints are not 'tight' in the sense of our discussion in Subsection 3.2.

More generally speaking, the example shows that decomposing constraints like (12) into *independent* constraints for each of the mechanisms $P_{X_j|P_{A_j}}$ may not be possible. The bivariate example suggests that the requirement to obtain a distribution that factorizes according to the DAG structure prohibits using the constraints entirely, given that we must not commit to any information that is not entailed by the constraints (as we would do by choosing either $(X_1, X_2) = (1, 1)$ or $(X_1, X_2) = (-1, -1)$).

6 Deriving Information Geometric Causal Inference from Causal PIR

Information Geometric Causal Inference (IGCI) [21, 38] is a method for causal discovery, that infers whether two X causes Y for Y causes X from the bivariate distribution $P_{X,Y}$ for the case of an invertible *deterministic* causal relation, i.e., $Y = f(X)$ and $X = f^{-1}(Y)$. Although IGCI is more general, we sketch the idea for variables with

⁹ A DAG is called complete if adding further arrows would result in directed cycles.



Figure 4. Left: IGCI for a bijective function $f : [0, 1] \rightarrow [0, 1]$. Under certain genericity assumptions, x -values tend to lie in regions with small slope $f'(x)$. Right: Drawing the function f with a fat pen, it induces a relation of possible (x, y) -pairs on the grid (obtained by discretizing X and Y).

range $[0, 1]$ and strictly monotonously increasing f , as shown in Figure 4, left. The intuitive idea is that, for the causal relation $X \rightarrow Y$, 'generic choices' of P_X (independently chosen of f)¹⁰ result in distributions P_Y that tend to have higher density in regions where the derivative $(f^{-1})'(y)$ is large. To exploit this asymmetry for inferring the direction, one infers $X \rightarrow Y$ iff points accumulate in regions of small f' rather than small $f^{-1'}$. Formally, IGCI amounts to inferring the direction $X \rightarrow Y$ iff¹¹

$$\sum_{j=1}^n \log f'(x_j) < 0.$$

IGCI can be obtained as the deterministic and continuous limit of Causal PIR in the following sense. Note that our derivation is close in spirit to the justification of IGCI provided by [39], which relies on counting arguments in the space of discrete functions. However, we want to directly derive it from Causal PIR.

Assume we draw the function f with a fat pen, as shown in Figure 4, right.

Define, after discretizing X and Y to get a grid with $\ell \times \ell$ points, define $R \subset \mathcal{X} \times \mathcal{Y}$ as the points (x, y) lying on the fat stripe. For each x , let $N_X(x)$ denote the number of possible y -values for which $(x, y) \in R$. Define $N_Y(y)$ similarly. For each observed point (x_j, y_j) we have

$$f'(x_j) \approx \frac{N_X(x_j)}{N_Y(y_j)}. \quad (16)$$

Then, $\prod_{j=1}^n N_X(x_j)$ is the number of possible n -tuples \mathbf{y} for the observed \mathbf{x} . Likewise, $\prod_{j=1}^n N_Y(y_j)$, is the number of possible n -tuples \mathbf{x} for the observed \mathbf{y} .

One checks easily that inferring causal direction via Causal PIR based cause-effect inference in Definition 2 thus amounts to comparing $\sum_{j=1}^n \log N_Y(y_j)$ to $\sum_{j=1}^n \log N_X(x_j)$, which, after using (16) amounts to checking the sign of $\sum_{j=1}^n \log f'(x_j)$.

We have thus shown that another non-trivial causal inference method (part from Causal MaxEnt) also follows from applying Causal PIR to the n -fold cartesian product of the underlying probability space.

7 Conclusions

Using a simple mechanical device as toy example, we have argued that our common sense replaces PIR with Causal PIR for bivariate distributions of cause and effect whenever we account for knowledge on the mechanism connecting cause and effect. We have further justified Causal PIR and Causal MaxEnt by assuming that constraints on joint distributions arise from two separate constraints: constraints on the cause and constraints on the cause-effect relation

¹⁰ formalized by the condition $\int_0^1 \log f'(x)p(x)dx \approx \int_0^1 \log f'(x)dx$.

¹¹ note the symmetry $\sum_{j=1}^n \log f'(x_j) = -\sum_{j=1}^n \log f^{-1'}(y_j)$.

(in the sense of possible functions). Earlier work has solved paradoxes with usual MaxEnt by updating priors over functions, too. We have argued, however, that knowledge on the bivariate distribution is not a priori divided into information on the cause and information on the functional relation between cause and effect. We have therefore proposed a way to divide it into these two components in a way that entails minimal commitment for both of them.

Acknowledgements: Many thanks to Armen Allahverdyan and Sergio Hernan Garrido Mejia for helpful comments on relevant literature and Elke Kirschbaum for remarks on an earlier version.

References

- [1] Spirtes P, Glymour C, Scheines R. *Causation, Prediction, and Search*. New York, NY: Springer-Verlag; 1993.
- [2] Pearl J. *Causality*. Cambridge University Press; 2000.
- [3] Kano Y, Shimizu S. Causal inference using nonnormality. In: *Proceedings of the International Symposium on Science of Modeling, the 30th Anniversary of the Information Criterion*. Tokyo, Japan; 2003. p. 261–270.
- [4] Hoyer P, Janzing D, Mooij J, Peters J, Schölkopf B. Nonlinear causal discovery with additive noise models. In: Koller D, Schuurmans D, Bengio Y, Bottou L, editors. *Proceedings of the conference Neural Information Processing Systems (NIPS) 2008*. Vancouver, Canada: MIT Press; 2009. .
- [5] Mooij J, Peters J, Janzing D, Zscheischler J, Schölkopf B. Distinguishing cause from effect using observational data: methods and benchmarks. *Journal of Machine Learning Research*. 2016;17(32):1–102.
- [6] Zhang K, Hyvärinen A. On the Identifiability of the Post-Nonlinear Causal Model. In: *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence*. Montreal, Canada; 2009. .
- [7] Janzing D, Schölkopf B. Causal inference using the algorithmic Markov condition. *IEEE Transactions on Information Theory*. 2010;56(10):5168–5194.
- [8] Marx A, Vreeken J. Telling Cause from Effect Using MDL-Based Local and Global Regression. In: *2017 IEEE International Conference on Data Mining, ICDM 2017, New Orleans, LA, USA, November 18-21, 2017*; 2017. p. 307–316.
- [9] Kocaoglu M, Dimakis AG, Vishwanath S, Hassibi B. Entropic Causal Inference. In: *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence. AAAI’17*. AAAI Press; 2017. p. 1156–1162.
- [10] Blöbaum P, Washio T, Shimizu S. Error Asymmetry in Causal and Anticausal Regression. *Behaviormetrika*. 2017:1–22.
- [11] Peters J, Janzing D, Schölkopf B. *Elements of Causal Inference – Foundations and Learning Algorithms*. MIT Press; 2017.
- [12] Guyon I, Statnikov A, Bakır-Batu B. *Cause Effect Pairs in Machine Learning*. The Springer Series on Challenges in Machine Learning. Berlin & Heidelberg: Springer; 2019.
- [13] Janzing D. The Cause-Effect Problem: Motivation, Ideas, and Popular Misconceptions. In: Guyon I, Statnikov R, Bakır Batu B, editors. *Cause Effect Pairs in Machine Learning*. Springer; 2019. p. 3–26.
- [14] Allahverdyan A, Janzing D. Relating the thermodynamic arrow of time to the causal arrow. *J Stat Mech*. 2008;4:P04001.
- [15] Janzing D, Chaves R, Schölkopf B. Algorithmic independence of initial condition and dynamical law in thermodynamics and causal inference. *New Journal of Physics*. 2016;18(093052):1–13.
- [16] Schölkopf B, Janzing D, Peters J, Sgouritsa E, Zhang K, Mooij J. On causal and anticausal learning. In: J L, Pineau J, editors. *Proceedings of the 29th International Conference on Machine Learning (ICML)*. ACM; 2012. p. 1255–1262.
- [17] Bengio Y, Deleu T, Rahaman N, Ke R, Lachapelle S, Bilaniuk O, et al.. A Meta-Transfer Objective for Learning to Disentangle Causal Mechanisms; 2019. Available from: arXiv:1901.10912.
- [18] Jaynes ET. *Probability theory: the logic of science*. Cambridge, MA: Cambridge University Press; 2003.
- [19] Sun X, Janzing D, Schölkopf B. Causal inference by choosing graphs with most plausible Markov kernels. In: *Proceedings of the 9th International Symposium on Artificial Intelligence and Mathematics*. Fort Lauderdale, FL; 2006. p. 1–11.
- [20] Ziebart B, Bagnell J, Dey A. The Principle of Maximum Causal Entropy for Estimating Interacting Processes. *IEEE Transactions on Information Theory*. 2013 February;59(4):1966–1980.
- [21] Daniusis P, Janzing D, Mooij JM, Zscheischler J, Steudel B, Zhang K, et al. Inferring deterministic causal relations. In: *Proceedings of the 26th Annual Conference on Uncertainty in Artificial Intelligence (UAI)*. AUAI Press; 2010. p. 143–150.
- [22] Uffink J. Can the maximum entropy principle be explained as a consistency requirement? *Studies in History and Philosophy of Science Part B*. 1995;26(3):223 – 261.
- [23] Hunter D. Uncertain Reasoning using Maximum Entropy Inference. In: *Uncertainty in Artificial Intelligence*. vol. 4 of *Machine Intelligence and Pattern Recognition*. North-Holland; 1986. p. 203 – 209.
- [24] Hunter D. Causality and maximum entropy updating. *International Journal of Approximate Reasoning*. 1989;3(1):87 – 114.

- [25] Lemeire J, Janzing D. Replacing Causal Faithfulness with Algorithmic Independence of Conditionals. *Minds and Machines*. 2012 7;23(2):227–249.
- [26] Shajarisales N, Janzing D, Schölkopf B, Janzing D. Telling cause from effect in deterministic linear dynamical systems. In: *Proceedings of the International Conference on Machine Learning*. Lille; 2015. To appear.
- [27] Li M, Vitányi P. *An Introduction to Kolmogorov Complexity and its Applications*. New York: Springer; 1997.
- [28] Zurek W. Algorithmic randomness and physical entropy. *Phys Rev A*. 1989;40(8):4731–4751.
- [29] Frogner C, Poggio T. Fast and Flexible Inference of Joint Distributions from their Marginals. In: *International Conference on Machine Learning*; 2019. .
- [30] Levy W, Delic H. Maximum entropy aggregation of individual opinions. *IEEE Transactions on Systems, Man, and Cybernetics*. 1994;24(4):606–613.
- [31] Myung IJ, Ramamoorti S, Bailey AD. Maximum Entropy Aggregation of Expert Predictions. *Management Science*. 1996;42(10):1420–1436.
- [32] Cover T, Thomas J. *Elements of Information Theory*. New York: Wileys Series in Telecommunications; 1991.
- [33] Jaynes ET. Information Theory and Statistical Mechanics. II. *Phys Rev*. 1957 Oct;108:171–190.
- [34] Palmieri F, Domenico C. Objective priors from maximum entropy in data classification. *Information Fusion*. 2013;14(2):186 – 198.
- [35] Uffink J. The constraint rule of the maximum entropy principle. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*. 1996;27(1):47 – 79.
- [36] Shore J, Johnson R. AXIOMATIC DERIVATION OF THE PRINCIPLE OF MAXIMUM ENTROPY AND THE PRINCIPLE OF MINIMUM CROSS-ENTROPY. *Naval Research Laboratory, Washington DC*. 1978:1–59.
- [37] Janzing D, Sun X, Schölkopf B. Distinguishing Cause and Effect via Second Order Exponential Models. <http://arxiv.org/abs/09105561>. 2009.
- [38] Janzing D, Mooij J, Zhang K, Lemeire J, Zscheischler J, Daniušis P, et al. Information-geometric approach to inferring causal directions. *Artificial Intelligence*. 2012;182–183:1–31.
- [39] Janzing D, Steudel B, Shajarisales N, Schölkopf B. Justifying information-geometric causal inference. In: *Vovk V, H P, Gammerman A, editors. Measures of Complexity. Festschrift for Alexey Chervonencis. Heidelberg: Springer Verlag; 2015. p. 253–265.*

A Relation to Maximum Causal Entropy

Ref. [20] consider a scenario with two interacting time series $(X_t)_{t \in \mathbb{Z}}$, $(Y_t)_{t \in \mathbb{Z}}$ where the latter (‘predicted process’) is to be inferred from the former (‘known process’). Further, it is assumed that there are constraints (e.g. moment restrictions) on the joint distribution of the bivariate process capturing the relation between the two time series. Maximum Causal Entropy describes a way how an agent optimally accounts for observations from the known process for predicting the other process: While *future* values X_s for $s > t$ can also contain information about the current value Y_t , [20] sequentially maximizes the entropy of Y_t subject to constraints referring to observations of X_s for $s \leq t$, rather than accounting also for constraints that involve the entire process $(X_t)_{t \in \mathbb{Z}}$. The obvious argument is that observations from X_s for $s > t$ are not available at t (which we will criticize below).

Let us first mention an important conceptual difference to Causal MaxEnt. Maximum Causal Entropy a priori restricts the set of joint distributions over which is maximized. This is because it imposes conditional independences since every Y_t is independent of future values X_s ($s > t$), given the past of Y plus the past and present of X . This a priori restriction is not made in Causal MaxEnt.

To explain our problems with the justification of Maximum Causal Entropy, we describe a scenario where Maximum Causal Entropy amounts to our Causal MaxEnt and Causal PIR, but with different justification. Consider the case where the interaction is as in Figure 5. For some fixed t , Y_t is influenced by Y_{t-1} and X_t . Assume that Y_t can attain the values 1, 2, 3 as our variables X, Y in Section 2. Further, let X_t be binary and assume that its influence consists in switching between two different mechanisms for the relation between Y_{t-1} and Y_t : whenever $X_t = 1$, Y_{t-1} and Y_t are related by the mechanical device in Section 2. Whenever $X_t = 0$, Y_t is drawn independently of Y_{t-1} . To infer the joint distribution of Y_{t-1}, Y_t , Maximum Causal Entropy first chooses $P_{Y_{t-1}}$ to be uniform, since the mechanism relating Y_{t-1} and Y_t is not known at that time. Then, after observing X_t , it constructs $P_{Y_t|Y_{t-1}, X_t}$ as the uniform distribution over all Y_t allowed by the mechanism determined by X_t . Here, the resulting joint distribution coincides with the one constructed via Causal PIR and Causal MaxEnt for both cases $X_t = 0, 1$.




Figure 5. Interaction between two time series in which X_t controls the mechanism relating Y_{t-1} and Y_t , but all other observations are independent.

However, we do not believe that [20] answers the question of *why* one should account for the known constraints in this sequential way. The fact that X_t is *not known* when Y_{t-1} is inferred, does not justify to entirely ignore the knowledge on the mechanisms. After all, we know that there are $13 = 4 + 9$ possible combinations for the triple (y_{t-1}, y_t, x_t) . Assigning a uniform prior over them would result in a marginal distribution $P_{Y_{t-1}}$ that still slightly prefers the value Y_{t-1} because it offers more options for Y_t in case X_t attains 1 (and equally many options otherwise). In other words, from the point of view of optimally using all available information, we could, in any step, also account for constraints that refer to variables whose values are not known at that point in time and take their uncertainty in account. We believe that arguments similar to the ones in the present paper are required to complement the justification by [20].