



Confounding Ghost Channels and Causality: A New Approach to Causal Information Flows

Nihat Ay^{1,2,3,4}

Received: 29 June 2020 / Accepted: 22 April 2021 / Published online: 8 July 2021
© The Author(s) 2021

Abstract

Information theory provides a fundamental framework for the quantification of information flows through channels, formally Markov kernels. However, quantities such as mutual information and conditional mutual information do not necessarily reflect the causal nature of such flows. We argue that this is often the result of conditioning based on σ -algebras that are not associated with the given channels. We propose a version of the (conditional) mutual information based on families of σ -algebras that are coupled with the underlying channel. This leads to filtrations which allow us to prove a corresponding causal chain rule as a basic requirement within the presented approach.

Keywords Information flow · Causality · Mutual information · Conditional mutual information · Filtration

Mathematics Subject Classification (2010) 62A01 · 62B10 · 94A15 · 94A17 · 94A40

1 Introduction: Information Theory and Causality

Classical information theory [20] is based on the definition of *Shannon entropy*, a measure of uncertainty about the outcome of a variable Z :

$$H(Z) = - \sum_z p(z) \log p(z),$$

where $p(z) = \mathbb{P}(Z = z)$ denotes the distribution of Z . (Throughout this introduction, we consider only variables X , Y , and Z with finite state sets X , Y , and Z , respectively.)

Dedicated to Jürgen Jost on the occasion of his 65th birthday.

✉ Nihat Ay
nihat.ay@tuhh.de

¹ Hamburg University of Technology, Hamburg, Germany

² Max Planck Institute for Mathematics in the Sciences, Leipzig, Germany

³ Leipzig University, Leipzig, Germany

⁴ Santa Fe Institute, Santa Fe, NM, USA

Shannon entropy serves as a building block of further important quantities. The flow of information from a sender X to a receiver Z , for instance, can be quantified as the reduction of uncertainty about the outcome of Z based on the outcome of X . More precisely, we compare two uncertainties here, the uncertainty about the outcome of Z , that is $H(Z)$, with the uncertainty about the outcome of Z *after* knowing the outcome of X , that is

$$H(Z|X) = - \sum_x p(x) \sum_z p(z|x) \log p(z|x), \quad (1)$$

where $p(z|x) = \mathbb{P}(Z = z|X = x)$ denotes the conditional distribution of Z given X . (Note that for x with $p(x) = 0$, conditioning is not well-defined. This ambiguity, however, has no effect on (1), due to the multiplication with $p(x)$.) Naturally, the latter uncertainty, $H(Z|X)$, is smaller than or equal to the first one, $H(Z)$, leading to another fundamental quantity of information theory, the *mutual information*:

$$I(X; Z) = H(Z) - H(Z|X). \quad (2)$$

This difference can also be expressed in geometric terms as the *KL-divergence*

$$I(X; Z) = \sum_x p(x) \sum_z p(z|x) \log \frac{p(z|x)}{p(z)}. \quad (3)$$

The KL-divergence plays an important role in information geometry as a canonical divergence [1–4]. Such a divergence is characterised in terms of natural geometric properties. It is remarkable that this purely geometric approach yields the fundamental information-theoretic quantities which were previously derived from a set of axioms that are formulated in non-geometric terms.

Typically, the conditional distribution $p(z|x)$ is interpreted mechanistically as a channel which receives x as an input and generates z as an output. In this interpretation, the stochasticity of a channel is considered to be the effect of external or hidden disturbances of a deterministic map. This is formalised in terms of a so-called *structural equation*

$$Z = f(X, U), \quad (4)$$

with a deterministic map f and a noise variable U that is independent of X [17]. Integrating out the noise variable U , we obtain a Markov kernel, the formal model of a channel. What do we gain by this construction? Formally, we do not gain much, as a Markov kernel is basically a conditional probability distribution, defined for all x . However, with the representation (4) and the associated Markov kernel, the conditional probability distribution $p(z|x)$ can be interpreted as the result of a (probabilistic) causal effect of X on Z . This interpretation provides the basis for Pearl's influential proposal of a general theory of causality [17]. The mutual information (2) then becomes a measure of the causal information flow from X to Z [6], which is consistent with Shannon's original idea of the amount of information transmitted through a channel [20]. This consistency, however, is apparently violated when dealing with variations or extensions of the sender-receiver setting. In Section 2, we are now going to highlight instances of such inconsistency that will play an important role in this article.

Later in this article, a channel will be formally given in terms of a Markov kernel, with a more explicit notation. In what follows, however, we keep the notation of a conditional probability distribution and state explicitly when we interpret it mechanistically as a channel.

2 Confounding Ghost Channels

The mutual information is symmetric, that is $I(X; Z) = I(Z; X)$. Interpreting it as a measure of causal information flow, this symmetry suggests that we have the same amount of causal information flow in both directions, even though the channel goes from X to Z so that there cannot be any flow of information in the opposite direction. What is wrong here? This apparent problem, let us call it “the symmetry puzzle”, can be resolved quite easily. We *can* revert the direction and compute the conditional distribution $p(x|z) = \frac{p(x)}{p(z)} p(z|x)$, based on elementary rules of probability theory and without reference to any mechanisms. Furthermore, this conditional distribution *can* be mechanistically interpreted and represented in terms of a structural equation (4). (This is always possible for a given conditional distribution.) Such a representation introduces a hypothetical channel for generating the reverted conditional distribution $p(x|z)$, a kind of “ghost” channel that is actually not there. The mutual information then quantifies the causal information flow of this hypothetical channel. The symmetry of the mutual information simply means that the actual causal information flow in forward direction will be equal to the causal information flow of any hypothetical channel in backward direction that is capable of generating the conditional distribution $p(x|z)$. The symmetry puzzle, however, is not the only apparent inconsistency between the (conditional) mutual information and causality. We are now going to highlight another problem, which is closely related to the symmetry puzzle but requires a deeper analysis for its solution.

We now assume that the channel receives x and y as inputs and generates z as an output. With the corresponding conditional distribution $p(z|x, y) = \mathbb{P}(Z = z|X = x, Y = y)$ we have the *conditional mutual information* of Y and Z given X :

$$I(Y; Z|X) = H(Z|X) - H(Z|X, Y) \tag{5}$$

$$= \sum_{x,y} p(x, y) \sum_z p(z|x, y) \log \frac{p(z|x, y)}{p(z|x)}. \tag{6}$$

According to (5), the conditional mutual information compares the uncertainty about z given x , before and after observing the outcome y , reflected by the conditional probabilities $p(z|x)$ and $p(z|x, y)$, respectively. The representation (6) makes this comparison more explicit as a deviation of $p(z|x, y)$ from $p(z|x)$. Together with (3), we obtain the chain rule

$$I(X, Y; Z) = \sum_{x,y} p(x, y) \sum_z p(z|x, y) \log \frac{p(z|x, y)}{p(z)} \tag{7}$$

$$= \sum_{x,y} p(x, y) \sum_z p(z|x, y) \left[\log \frac{p(z|x)}{p(z)} + \log \frac{p(z|x, y)}{p(z|x)} \right] \tag{8}$$

$$= I(X; Z) + I(Y; Z|X). \tag{9}$$

For the computation of both terms, $I(X; Z)$ and $I(Y; Z|X)$, we have to evaluate the “reduced” conditional distribution $p(z|x)$. It is obtained from the original one in the following way:

$$p(z|x) = \sum_y p(y|x) p(z|x, y). \tag{10}$$

This conditional distribution represents a second kind of hypothetical channel, a “ghost channel”, which screens off the actual flow of information. It can be sensitive to information about x that is not necessarily employed by the original channel $p(z|x, y)$. More precisely, given two states x, x' that satisfy $p(z|x, y) = p(z|x', y)$ for all z and all y , we cannot expect

$p(z|x) = p(z|x')$ for all z . This is a consequence of the coupling through $p(y|x)$, on the right-hand side (RHS, for short) of (10). In the most extreme case, y is simply a deterministic map of x , so that the knowledge of y does not provide any additional information about z , that is $p(z|x, y) = p(z|x)$. In the following example we study this case more explicitly and thereby highlight the inconsistency of the terms $I(X; Z)$ and $I(Y; Z|X)$ in (9) with the underlying causal structure. We will argue that the conditional distribution (10) has to be modified in order to allow for a causal interpretation.

Example 1 Consider three variables X, Y, Z with values -1 and $+1$, and assume that Z is obtained as a copy of Y , that is

$$p(z|x, y) = \begin{cases} 1, & \text{if } z = y \\ 0, & \text{otherwise.} \end{cases} \tag{11}$$

We interpret this conditional distribution, which is well-defined for *all* arguments x and y , as a mechanism. This means that all information required for the output Z is contained in Y . Intuitively, we would expect from a measure of information flow to assign zero for the flow from X to Z and a positive value to the flow of information from Y to Z given X . This is however not what we get with the usual definitions of mutual information and conditional mutual information. The reason for that is the stochastic dependence of the inputs X and Y . To be more precise, let us assume that the input distribution is given as

$$p(x, y) = \frac{e^{\beta xy}}{\sum_{x', y' \in \{\pm 1\}} e^{\beta x' y'}}$$

where the parameter β controls the coupling of the inputs. This implies $p(x) = \mathbb{P}(X = x) = 1/2$ and $p(y) = \mathbb{P}(Y = y) = 1/2$ for all $x, y \in \{\pm 1\}$. We can decompose the full mutual information, as a measure of information flow from X and Y together to Z , in the following way:

$$I_\beta(X, Y; Z) = \underbrace{I_\beta(Y; Z)}_{=\log 2} + \underbrace{I_\beta(X; Z|Y)}_{=0} \tag{12}$$

(The subscript β indicates the dependence of the respective information-theoretic quantities on this parameter.) This is consistent with the intuition that Z is receiving all information from Y and no information from X . However, we observe an inconsistency if we decompose the full mutual information in a different way:

$$I_\beta(X, Y; Z) = I_\beta(X; Z) + I_\beta(Y; Z|X) \tag{13}$$

For the two terms on the RHS of (13) we obtain

$$I_\beta(X; Z) = \log(2) - \frac{\log(1 + e^{2\beta})}{1 + e^{2\beta}} - \frac{\log(1 + e^{-2\beta})}{1 + e^{-2\beta}},$$

$$I_\beta(Y; Z|X) = \frac{\log(1 + e^{2\beta})}{1 + e^{2\beta}} + \frac{\log(1 + e^{-2\beta})}{1 + e^{-2\beta}}.$$

These functions are shown in Fig. 1. In the limit $\beta \rightarrow -\infty$ the two inputs become completely anti-correlated with support $(-1, +1)$ and $(+1, -1)$. Correspondingly, for

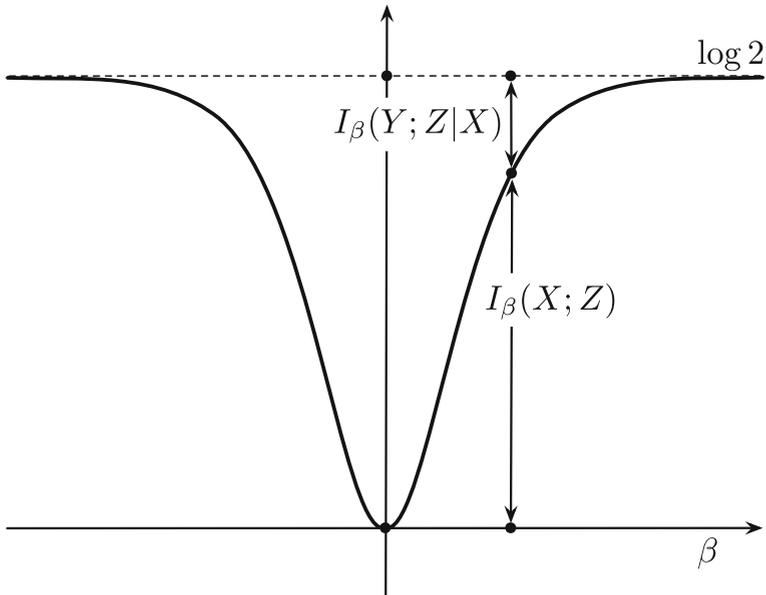


Fig. 1 The mutual information $I_\beta(X; Z)$ and the conditional mutual information $I_\beta(Y; Z|X)$ as functions of β . Even though the channel does not employ any information from X , the mutual information $I_\beta(X; Z)$ converges to the maximal value for $\beta \rightarrow \infty$

$\beta \rightarrow +\infty$ we have complete correlation, and the support is $(-1, -1)$ and $(+1, +1)$. With (13), we obtain the following decomposition:

$$\begin{aligned}
 I(X, Y; Z) &= \lim_{\beta \rightarrow \infty} I_\beta(X, Y; Z) \\
 &= \lim_{\beta \rightarrow \infty} I_\beta(X; Z) + \lim_{\beta \rightarrow \infty} I_\beta(Y; Z|X) \\
 &= \underbrace{I(X; Z)}_{=\log 2} + \underbrace{I(Y; Z|X)}_{=0}.
 \end{aligned}
 \tag{14}$$

The decomposition (14) gives the impression that Z is receiving all information from X and no information from Y . However, we know, by the definition of the mechanism (11), that this is not the case. The actual situation is better reflected by the decomposition (12).

The problem highlighted in Example 1 can be resolved by an appropriate modification of the conditional probability (10). We are now going to outline this modification, which will provide the main idea of this article. In a first step, let us assume that \bar{y} is fixed as an input to the channel. Which information about x does the channel then use for generating z ? In order to qualitatively describe that information, we lump any two states x and x' together whenever the channel cannot distinguish them, that is

$$p(z|x, \bar{y}) = p(z|x', \bar{y})$$

for all z . This defines a partition $\alpha_{X, \bar{y}}$ of the state set of X that depends on \bar{y} . In a second step, we consider the join of all these partitions, that is their coarsest refinement. More

precisely, we define

$$\alpha_X := \bigvee_{\bar{y} \in Y} \alpha_{X, \bar{y}} := \left\{ \bigcap_{\bar{y} \in Y} A_{\bar{y}} : A_{\bar{y}} \in \alpha_{X, \bar{y}}, \bar{y} \in Y \right\}. \tag{15}$$

The partition α_X represents a qualitative description of the information in X that is used by the channel $p(z|x, y)$. Denote by A_x the set in α_X that contains x . When the channel receives x , in addition to y , then it does not “see” the full x but only the class A_x , and it is easy to verify $p(z|x, y) = p(z|A_x, y)$. Therefore, we replace the conditioning $p(z|x)$ in the above formula (10) by

$$\hat{p}(z|x) := p(z|A_x) = \sum_y p(y|A_x) p(z|A_x, y) = \sum_{x' \in A_x} p(x'|A_x) p(z|x'). \tag{16}$$

This shows that the new conditional distribution $\hat{p}(z|x)$ is obtained by averaging the previous one, $p(z|x)$, according to the information that is actually used by the channel $p(z|x, y)$. Now, replacing in (8) the conditional distribution $p(z|x)$ by $\hat{p}(z|x)$ leads to a corresponding modification of the mutual information and the conditional mutual information:

$$I(X \rightarrow Z) := \sum_x p(x) \sum_z \hat{p}(z|x) \log \frac{\hat{p}(z|x)}{p(z)},$$

$$I(Y \rightarrow Z|X) := \sum_{x,y} p(x, y) \sum_z p(z|x, y) \log \frac{p(z|x, y)}{\hat{p}(z|x)}.$$

It is easy to see that

$$I(X \rightarrow Z) \leq I(X; Z), \quad I(Y \rightarrow Z|X) \geq I(Y; Z|X).$$

However, the sum does not change and we have the chain rule

$$I(X, Y; Z) = I(X \rightarrow Z) + I(Y \rightarrow Z|X).$$

With this new definition, we come back to Example 1. The channel defined by (11) does not use any information from X . Therefore, $\alpha_{X, \bar{y}} = \{X\}$ for all $\bar{y} \in Y$, which implies $\alpha_X = \{X\}$. With formula (16) we obtain $\hat{p}(z|x) = p(z|X) = p(z)$, and therefore

$$I(X \rightarrow Z) = 0 \quad \text{and} \quad I(Y \rightarrow Z|X) = \log 2.$$

If we compare this with (14) we see that the information is shifted from the first to the second term which corresponds to the variable that has the actual causal effect on Z . On the other hand, in both cases the sum of the two contributions equals $\log 2$, the full mutual information $I(X, Y; Z)$.

Causality plays an important role in time series analysis. In this context, *Granger causality* [12, 13] has been the subject of extensive debates which tend to highlight its non-causal nature. Schreiber proposed an information-theoretic quantification of Granger causality, referred to as *transfer entropy*, which is based on conditional mutual information [9, 19]. Even though transfer entropy is an extremely useful and widely applied quantity, it is generally accepted that it has shortcomings as a measure of causal information flow. In particular, it can vanish in cases where the causal effect is the strongest possible. We argue that this is again a result of a ghost channel that is involved in the computation of the classical conditional mutual information and screens off the actual causal information flow. This is demonstrated in the following example which is taken from [6]. Essentially, this example is a reformulation of Example 1, thereby adjusted to the context of time series and stochastic processes.

Example 2 (Transfer entropy) Consider a stochastic process (X_m, Y_m) , $m = 1, 2, \dots$, with state space $\{\pm 1\}^2$ and define $X^m := (X_1, \dots, X_m)$ and $Y^m := (Y_1, \dots, Y_m)$. The transfer entropy at time m is defined as

$$T(Y^{m-1} \rightarrow X_m) := I(Y^{m-1}; X_m | X^{m-1}).$$

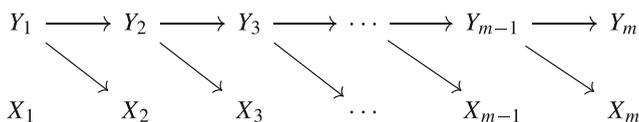
Thus, the transfer entropy quantifies how much information the variables Y_1, \dots, Y_{m-1} contribute to the evaluation of X_m , in addition to the information in X_1, \dots, X_{m-1} . We assume that the process is a Markov chain, given by a transition matrix of the form

$$p(x', y' | x, y) = p(x' | x, y) p(y' | x, y),$$

where

$$p(x' | x, y) = \frac{1}{1 + e^{2\beta x' y}}, \quad p(y' | x, y) = \frac{1}{1 + e^{2\beta y' x}}.$$

The causal structure of the dynamics is represented by the following diagram:



As a stationary distribution we have

$$p(+1, +1) = p(-1, -1) = \frac{1}{2} - ab, \quad p(+1, -1) = p(-1, +1) = ab,$$

where

$$a = \frac{1}{1 + e^{2\beta}}, \quad b = \frac{1}{1 + e^{-2\beta}}.$$

The transfer entropy can be upper bounded as follows (the subscript β indicates the dependence on the coupling parameter β):

$$\begin{aligned} T_\beta(Y^{m-1} \rightarrow X_m) &= I_\beta(Y^{m-1}; X_m | X^{m-1}) \\ &= I_\beta(Y_{m-1}; X_m | X^{m-1}) \\ &= H_\beta(X_m | X^{m-1}) - H_\beta(X_m | X^{m-1}, Y_{m-1}) \\ &= H_\beta(X_m | X^{m-1}) - H_\beta(X_m | X_{m-1}, Y_{m-1}) \\ &\leq H_\beta(X_m | X_{m-1}) - H_\beta(X_m | X_{m-1}, Y_{m-1}) \\ &= I_\beta(Y_{m-1}; X_m | X_{m-1}). \end{aligned}$$

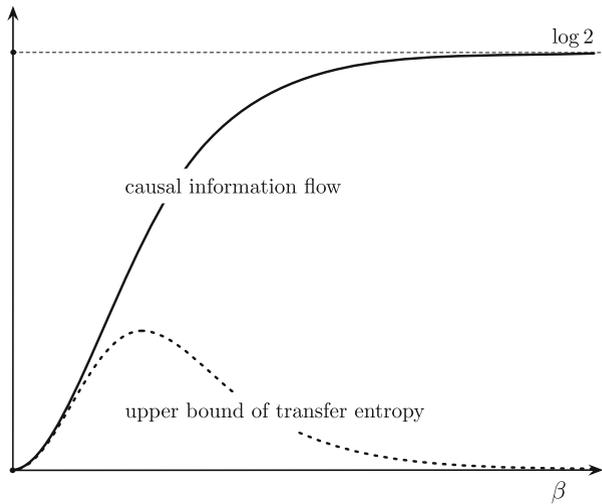
For $\beta = 0$, we have an i.i.d. process with uniform distribution over the states $(+1, +1)$, $(-1, +1)$, $(+1, -1)$, and $(-1, -1)$. For $\beta \rightarrow \infty$, we obtain the deterministic transition

$$(x, y) \mapsto (-y, -y).$$

In this limit, the variables (X_m, Y_m) are completely correlated with $p(+1, +1) = p(-1, -1) = \frac{1}{2}$. In both cases, $\beta = 0$ and $\beta \rightarrow \infty$, the conditional mutual information $I_\beta(Y_{m-1}; X_m | X_{m-1})$, and therefore the transfer entropy $T_\beta(Y^{m-1} \rightarrow X_m)$, vanishes. For $\beta = 0$, this does not represent a problem because any measure of causal information flow should vanish in the i.i.d. case. However, for $\beta \rightarrow \infty$, the variable X_m is causally determined by Y_{m-1} . Therefore, a measure of casual information flow should be maximal in this case. This is not reflected by the transfer entropy. Let us compare this with the information flow measure proposed in this article. Given that X_m only depends on Y_{m-1} , the partition (15) is trivial, that is $\alpha = \{X\}$. Therefore,

$$I(Y^{m-1} \rightarrow X_m | X^{m-1}) = I(Y_{m-1}; X_m).$$

Fig. 2 Dashed line: the conditional mutual information $I_\beta(Y_{m-1}; X_m|X_{m-1})$ as an upper bound of the transfer entropy $T_\beta(Y^{m-1} \rightarrow X_m)$; solid line: the causal information flow $I_\beta(Y^{m-1} \rightarrow X_m|X^{m-1})$ which coincides with the mutual information $I_\beta(Y_{m-1}; X_m)$ in this example



This quantity is converging to the maximal value $\log 2$ for $\beta \rightarrow \infty$. For comparison, both functions are plotted in Fig. 2.

In what follows, we will extend the idea outlined in this section to a more general context of measurable spaces, probability measures, and Markov kernels. In further steps, we will also consider more input nodes.

3 General Information-Theoretic Quantities

In the previous sections, we reviewed fundamental information-theoretic quantities as they are introduced in standard textbooks such as [10]. In this section, we offer an alternative review from a measure-theoretic perspective (see, for instance, [15]). This more abstract setting will allow us to identify natural operations and definitions which are not always visible when dealing with finite state spaces.

Shannon Entropy For a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and a finite measurable partition $\gamma = \{C_1, \dots, C_m\}$, that is $C_i \in \mathcal{F}$, $C_i \cap C_j = \emptyset$ for all $i \neq j$, and $\bigcup_{i=1}^m C_i = \Omega$, the *Shannon entropy* of γ is given by

$$H(\gamma) := - \sum_{C \in \gamma} \mathbb{P}(C) \log \mathbb{P}(C). \tag{17}$$

As a local version of the Shannon entropy, we define

$$h(\gamma) := - \sum_{C \in \gamma} \mathbb{1}_C \log \mathbb{P}(C),$$

where $\mathbb{1}_C$ is the indicator function of C . Denoting by C_ω the set in γ that contains $\omega \in \Omega$, we have $h(\gamma)(\omega) = -\log \mathbb{P}(C_\omega)$. If we integrate the function $h(\gamma)$, we recover the entropy

(17) of the partition γ :

$$\begin{aligned} \int_{\Omega} h(\gamma) d\mathbb{P} &= - \sum_{C \in \gamma} \left\{ \int_{\Omega} \mathbb{1}_C d\mathbb{P} \right\} \log \mathbb{P}(C) \\ &= - \sum_{C \in \gamma} \mathbb{P}(C) \log \mathbb{P}(C) \\ &= H(\gamma). \end{aligned}$$

Conditional Entropy Consider two finite measurable partitions α and γ of Ω , where we assume $\mathbb{P}(A) > 0$ for all $A \in \alpha$. The *conditional entropy* of γ given α is then defined by

$$H(\gamma|\alpha) := - \sum_{A \in \alpha} \mathbb{P}(A) \sum_{C \in \gamma} \mathbb{P}(C|A) \log \mathbb{P}(C|A). \tag{18}$$

As a local version $h(\gamma|\alpha)$ of $H(\gamma|\alpha)$, we define

$$h(\gamma|\alpha) := - \sum_{C \in \gamma} \mathbb{1}_C \log \left(\sum_{A \in \alpha} \mathbb{P}(C|A) \mathbb{1}_A \right). \tag{19}$$

If we evaluate this function for $\omega \in \Omega$ we obtain $h(\gamma|\alpha)(\omega) = -\log \mathbb{P}(C_{\omega}|A_{\omega})$, where A_{ω} and C_{ω} are the atoms in α and γ , respectively, that contain ω . Integrating $h(\gamma|\alpha)$, we recover (18):

$$\begin{aligned} \int_{\Omega} h(\gamma|\alpha) d\mathbb{P} &= - \int_{\Omega} \log \mathbb{P}(C_{\omega}|A_{\omega}) \mathbb{P}(d\omega) \\ &= - \sum_{A \in \alpha} \sum_{C \in \gamma} \int_{A \cap C} \log \mathbb{P}(C_{\omega}|A_{\omega}) \mathbb{P}(d\omega) \\ &= - \sum_{A \in \alpha} \sum_{C \in \gamma} \mathbb{P}(A \cap C) \log \mathbb{P}(C|A) \\ &= H(\gamma|\alpha). \end{aligned}$$

The function $h(\gamma|\alpha)$ can be generalised by replacing the partition α by an arbitrary σ -subalgebra \mathcal{A} of \mathcal{F} :

$$h(\gamma|\mathcal{A}) := - \sum_{C \in \gamma} \mathbb{1}_C \log \mathbb{P}(C|\mathcal{A}), \tag{20}$$

where $\mathbb{P}(C|\mathcal{A}) = \mathbb{E}(\mathbb{1}_C|\mathcal{A})$. Note that this function is only \mathbb{P} -almost everywhere defined (abbreviated as \mathbb{P} -a.e.). In the case where the σ -algebra \mathcal{A} is given by a finite partition α with $\mathbb{P}(A) > 0$ for all $A \in \alpha$, we have

$$\mathbb{P}(C|\mathcal{A}) = \sum_{A \in \alpha} \mathbb{P}(C|A) \mathbb{1}_A, \quad \mathbb{P}\text{-a.e.}$$

This shows that the definition (20) is indeed an extension of (19). Correspondingly, integrating (20) yields a generalistaion of (18):

$$\begin{aligned} H(\gamma|\mathcal{A}) &:= \int_{\Omega} h(\gamma|\mathcal{A}) \, d\mathbb{P} \\ &= - \sum_{C \in \gamma} \int_{\Omega} \mathbb{1}_C \log \mathbb{P}(C|\mathcal{A}) \, d\mathbb{P} \\ &= - \sum_{C \in \gamma} \int_{\Omega} \mathbb{P}(C|\mathcal{A}) \log \mathbb{P}(C|\mathcal{A}) \, d\mathbb{P}. \end{aligned}$$

Mutual Information If we subtract from the entropy of a partition γ the conditional entropy of γ given a partition α , we obtain the *mutual information*:

$$I(\alpha; \gamma) := \sum_{A \in \alpha} \mathbb{P}(A) \sum_{C \in \gamma} \mathbb{P}(C|A) \log \frac{\mathbb{P}(C|A)}{\mathbb{P}(C)}. \tag{21}$$

Let us relate this function to the corresponding local functions $h(\gamma)$ and $h(\gamma|\alpha)$. Taking the difference, we obtain

$$\begin{aligned} i(\alpha; \gamma) &:= h(\gamma) - h(\gamma|\alpha) \\ &= - \sum_{C \in \gamma} \mathbb{1}_C \log \mathbb{P}(C) + \sum_{C \in \gamma} \mathbb{1}_C \log \left(\sum_{A \in \alpha} \mathbb{P}(C|A) \mathbb{1}_A \right) \\ &= \sum_{C \in \gamma} \mathbb{1}_C \log \left(\sum_{A \in \alpha} \frac{\mathbb{P}(C|A)}{\mathbb{P}(C)} \mathbb{1}_A \right). \end{aligned}$$

If we evaluate this for $\omega \in \Omega$ we obtain $i(\alpha; \gamma)(\omega) = \log \frac{\mathbb{P}(C_\omega|A_\omega)}{\mathbb{P}(C_\omega)}$, and thus we have

$$I(\alpha; \gamma) = \int_{\Omega} i(\alpha; \gamma) \, d\mathbb{P}.$$

For the general case where the partition α is replaced by a σ -subalgebra \mathcal{A} of \mathcal{F} , we obtain

$$\begin{aligned} i(\mathcal{A}; \gamma) &:= h(\gamma) - h(\gamma|\mathcal{A}) \\ &= - \sum_{C \in \gamma} \mathbb{1}_C \log \mathbb{P}(C) + \sum_{C \in \gamma} \mathbb{1}_C \log \mathbb{P}(C|\mathcal{A}) \\ &= \sum_{C \in \gamma} \mathbb{1}_C \log \frac{\mathbb{P}(C|\mathcal{A})}{\mathbb{P}(C)}. \end{aligned}$$

This leads to a corresponding generalisation of (21):

$$\begin{aligned} I(\mathcal{A}; \gamma) &:= \int_{\Omega} i(\mathcal{A}; \gamma) \, d\mathbb{P} \\ &= \sum_{C \in \gamma} \int_{\Omega} \mathbb{1}_C \log \frac{\mathbb{P}(C|\mathcal{A})}{\mathbb{P}(C)} \, d\mathbb{P} \\ &= \sum_{C \in \gamma} \int_{\Omega} \mathbb{P}(C|\mathcal{A}) \log \frac{\mathbb{P}(C|\mathcal{A})}{\mathbb{P}(C)} \, d\mathbb{P}. \end{aligned} \tag{22}$$

Conditional Mutual Information Finally, we define the *conditional mutual information*. With two σ -subalgebras \mathcal{A} and \mathcal{B} of \mathcal{F} , we define

$$i(\mathcal{B}; \gamma|\mathcal{A}) := h(\gamma|\mathcal{A}) - h(\gamma|\mathcal{A} \vee \mathcal{B}) = \sum_{C \in \gamma} \mathbb{1}_C \log \frac{\mathbb{P}(C|\mathcal{A} \vee \mathcal{B})}{\mathbb{P}(C|\mathcal{A})}.$$

(Here, $\mathcal{A} \vee \mathcal{B}$ denotes the smallest σ -algebra that contains \mathcal{A} and \mathcal{B} .) Integration of this function leads to

$$\begin{aligned} I(\mathcal{B}; \gamma|\mathcal{A}) &= \int_{\Omega} i(\mathcal{B}; \gamma|\mathcal{A}) \, d\mathbb{P} \\ &= \sum_{C \in \gamma} \int_{\Omega} \mathbb{1}_C \log \frac{\mathbb{P}(C|\mathcal{A} \vee \mathcal{B})}{\mathbb{P}(C|\mathcal{A})} \, d\mathbb{P} \\ &= \sum_{C \in \gamma} \int_{\Omega} \mathbb{P}(C|\mathcal{A} \vee \mathcal{B}) \log \frac{\mathbb{P}(C|\mathcal{A} \vee \mathcal{B})}{\mathbb{P}(C|\mathcal{A})} \, d\mathbb{P}. \end{aligned} \tag{23}$$

In a final step, we could further extend $i(\mathcal{B}; \gamma|\mathcal{A})$ and $I(\mathcal{B}; \gamma|\mathcal{A})$ to the case where γ is replaced by a σ -algebra \mathcal{C} , by taking the supremum over all finite partitions γ in \mathcal{C} . However, in this article we restrict attention to a fixed finite partition γ .

4 The Chain Rule as a Guiding Scheme

4.1 Two Inputs

In the introduction, Section 1, we have used the two-input case for discrete random variables in order to highlight the main issue with the classical definitions of the mutual information and the conditional mutual information and to outline the core idea of this article. After having introduced the required information-theoretic quantities for more general variables in Section 3, we now revisit the instructive two-input case and demonstrate how measure-theoretic concepts come into play here very naturally.

Consider measurable spaces (X, \mathcal{X}) , (Y, \mathcal{Y}) , (Z, \mathcal{Z}) , and their product

$$(\Omega, \mathcal{F}) := (X \times Y \times Z, \mathcal{X} \otimes \mathcal{Y} \otimes \mathcal{Z}).$$

In order to ensure the existence of various (regular versions of) conditional distributions, we need to assume that these measurable spaces carry a further structure. Typically, it is sufficient to require that (X, \mathcal{X}) , (Y, \mathcal{Y}) , and (Z, \mathcal{Z}) are Polish spaces (see [11], Theorem 13.1.1), which will be implicitly assumed hereinafter for all measurable spaces.

Now, consider a probability measure μ on $(X \times Y, \mathcal{X} \otimes \mathcal{Y})$ and a Markov kernel

$$\nu : X \times Y \times \mathcal{Z} \rightarrow [0, 1],$$

which models a channel that takes two inputs, $x \in X$ and $y \in Y$, and generates a possibly random output $z \in Z$. This allows us to define a probability measure on the joint space (Ω, \mathcal{F}) , given by

$$\mathbb{P}(A \times B \times C) := \int_{A \times B} \nu(x, y; C) \mu(dx, dy).$$

With the natural projections

$$X : \Omega \rightarrow X, (x, y, z) \mapsto x, \quad Y : \Omega \rightarrow Y, (x, y, z) \mapsto y, \quad Z : \Omega \rightarrow Z, (x, y, z) \mapsto z,$$

we have

$$\mu(A \times B) = \mathbb{P}(X \in A, Y \in B), \tag{24}$$

$$\nu(x, y; C) = \mathbb{P}(Z \in C | X = x, Y = y), \quad \mu\text{-almost everywhere}, \tag{25}$$

where the set of probability one to which “ μ -almost everywhere” refers in (25) is independent of C . Furthermore, we have the marginals

$$\mu_X(A) := \mu(A \times Y) = \mathbb{P}(X \in A),$$

$$\mu_Y(B) := \mu(X \times B) = \mathbb{P}(Y \in B),$$

and, finally, the ν -push-forward measure of μ ,

$$\nu_*(\mu)(C) := \mathbb{P}(Z \in C).$$

Note that the definition of the conditional distribution $\mathbb{P}(Z \in C | X = x, Y = y)$ on the RHS of (25) is quite general and does not exclude cases where $\mathbb{P}(X = x, Y = y) = 0$. It requires a formalism that goes beyond the context of variables with finitely many state sets X, Y , and Z . It is important to outline this formalism in some detail here, which will provide the basis for an appropriate definition of marginal channels. The definition of the conditional distribution

$$\mathbb{P}(Z \in C | X = x, Y = y) \tag{26}$$

involves two steps:

Step 1. We interpret the indicator function $\mathbb{1}_{\{Z \in C\}}$ as an element of the Hilbert space $L^2(\Omega, \mathcal{F}, \mathbb{P})$ and project it onto the (closed) linear subspace of (X, Y) -measurable functions $\Omega \rightarrow \mathbb{R}$. Its projection is referred to as *conditional expectation* and denoted by

$$\mathbb{E}(\mathbb{1}_{\{Z \in C\}} | X, Y). \tag{27}$$

Note that the elements of the Hilbert space $L^2(\Omega, \mathcal{F}, \mathbb{P})$ are equivalence classes of functions where two functions are identified if they coincide on a measurable set of probability one. Therefore, the conditional expectation (27) is only almost surely well defined, where the set of probability one to which “almost surely” refers is dependent on C .

Step 2. Formally, $\mathbb{E}(\mathbb{1}_{\{Z \in C\}} | X, Y)$ is a real-valued function defined on Ω . On the other hand, it is (X, Y) -measurable so that we should be able to interpret it as a function of x and y . Indeed, it follows from the factorisation lemma that there is a unique measurable function $\varphi_C : (X \times Y, \mathcal{X} \otimes \mathcal{Y}) \rightarrow \mathbb{R}$ satisfying $\mathbb{E}(\mathbb{1}_{\{Z \in C\}} | X, Y) = \varphi_C \circ (X, Y)$. The conditional distribution (26) is then simply defined to be the function φ_C , which has x and y as arguments. In the special situation where we start with a Markov kernel ν , we recover it, μ -almost everywhere, in terms of equation (25). It turns out that this equation already describes a quite general situation. Under mild conditions, assuming, for instance, that all measurable spaces are Polish spaces, the conditional distribution (26) can be considered to be a Markov kernel, as a function of x, y , and C .

For the definition of mutual information and conditional mutual information, as generalisations of (3) and (6), respectively, we have to find an appropriate notion of a marginal kernel. We begin with the conditional distribution $\mathbb{P}(Z \in C | X = x)$, as generalisation of $p(z|x)$. For its evaluation we repeat the arguments of the above two steps and consider the conditional expectation

$$\mathbb{E}(\mathbb{1}_{\{Z \in C\}} | X). \tag{28}$$

This is an X -measurable random variable $\Omega \rightarrow \mathbb{R}$. By the factorisation lemma we have a unique measurable function $\nu_X(\cdot; C) : (X, \mathcal{X}) \rightarrow \mathbb{R}$ satisfying $\mathbb{E}(\mathbb{1}_{\{Z \in C\}} | X) = \nu_X(X; C)$, and we set

$$\mathbb{P}(Z \in C | X = x) := \nu_X(x; C).$$

Under mild conditions we can assume that $\nu_X(x; C)$ defines a Markov kernel when considered as a function $\nu_X : X \times \mathcal{Z} \rightarrow [0, 1]$ in x and C .

We can now easily extend the classical definitions of mutual information and conditional mutual information to the context of this section. For a finite measurable partition ξ of Z , the sets $Z^{-1}(C)$, $C \in \xi$, form a corresponding finite measurable partition $Z^{-1}(\xi)$ of Ω , and we can use (22) to define the mutual informations

$$\begin{aligned} I_\xi(X, Y; Z) &:= I(\sigma(X) \vee \sigma(Y); Z^{-1}(\xi)) \\ &= \sum_{C \in \xi} \int_{X \times Y} \nu(x, y; C) \log \frac{\nu(x, y; C)}{\nu_*(\mu)(C)} \mu(dx, dy), \end{aligned}$$

and

$$\begin{aligned} I_\xi(X; Z) &:= I(\sigma(X); Z^{-1}(\xi)) \\ &= \sum_{C \in \xi} \int_X \nu_X(x; C) \log \frac{\nu_X(x; C)}{\nu_*(\mu)(C)} \mu_X(dx). \end{aligned} \tag{29}$$

Furthermore, with (23) we define the conditional mutual information

$$\begin{aligned} I_\xi(Y; Z | X) &:= I(\sigma(Y); Z^{-1}(\xi) | \sigma(X)) \\ &= \sum_{C \in \xi} \int_{X \times Y} \nu(x, y; C) \log \frac{\nu(x, y; C)}{\nu_X(x; C)} \mu(dx, dy). \end{aligned} \tag{30}$$

We repeat the computation (8) and decompose the mutual information $I_\xi(X, Y; Z)$ as follows:

$$\begin{aligned} I_\xi(X, Y; Z) &= \sum_{C \in \xi} \int_{X \times Y} \nu(x, y; C) \log \frac{\nu(x, y; C)}{\nu_*(\mu)(C)} \mu(dx, dy) \\ &= \sum_{C \in \xi} \int_{X \times Y} \nu(x, y; C) \left[\log \frac{\nu_X(x; C)}{\nu_*(\mu)(C)} + \log \frac{\nu(x, y; C)}{\nu_X(x; C)} \right] \mu(dx, dy) \tag{31} \\ &= I_\xi(X; Z) + I_\xi(Y; Z | X). \end{aligned} \tag{32}$$

We argue that, in order to have a causal decomposition of the full mutual information $I_\xi(X, Y; Z)$ into two terms similar to $I_\xi(X; Z)$ and $I_\xi(Y; Z | X)$, we have to modify the marginal channel

$$\nu_X(x; C) = \mathbb{P}(Z \in C | X = x)$$

in (31). In this modification, the conditioning with respect to X has to be adjusted to the actual information used by the kernel $\nu(x, y; C)$. To this end, we consider the smallest σ -subalgebra $\mathcal{A}_{X, \bar{y}}$ of \mathcal{X} for which all functions $\nu_{X, \bar{y}}(\cdot; C) := \nu(\cdot, \bar{y}; C)$, $C \in \mathcal{Z}$, are measurable. It corresponds to the partition $\alpha_{X, \bar{y}}$ that appears in (15). Now we generalise the definition (15) of the partition α_X by combining all the σ -algebras $\mathcal{A}_{X, \bar{y}}$:

$$\mathcal{A}_X := \bigvee_{\bar{y} \in Y} \mathcal{A}_{X, \bar{y}} \subseteq \mathcal{X}. \tag{33}$$

Note that, typically, $\mathcal{A}_{X, \bar{y}}$, $\bar{y} \in Y$, as well as \mathcal{A}_X cannot be naturally identified with corresponding σ -subalgebras of the σ -algebra generated by the channel ν , that is the smallest

σ -subalgebra of $\mathcal{X} \otimes \mathcal{Y}$ for which $(x, y) \mapsto \nu(x, y; C)$ is measurable for all $C \in \mathcal{Z}$. This is illustrated by the following example.

Example 3 We consider

$$(\mathbf{X}, \mathcal{X}) = (\mathbf{Y}, \mathcal{Y}) = (\mathbf{Z}, \mathcal{Z}) = (\mathbb{R}, \mathcal{B}(\mathbb{R})),$$

where $\mathcal{B}(\mathbb{R})$ denotes the Borel σ -algebra of \mathbb{R} . We assume that the channel ν is simply given by the addition $(x, y) \mapsto x + y$:

$$\nu(x, y; C) = \mathbb{1}_C(x + y).$$

As $\mathcal{B}(\mathbb{R})$ is generated by the intervals $[r - \varepsilon, r + \varepsilon] \subseteq \mathbb{R}$, the smallest σ -algebra $\mathcal{A} \subseteq \mathcal{X} \otimes \mathcal{Y}$ for which all functions $\nu(\cdot, \cdot; C)$ are measurable is generated by the following sets

$$A(r, \varepsilon) := \left\{ (x, y) \in \mathbb{R}^2 : r - \varepsilon \leq x + y \leq r + \varepsilon \right\}, \quad r \in \mathbb{R}, \varepsilon \in \mathbb{R}_+.$$

Now let us consider $\mathcal{A}_{X, \bar{y}}$, the σ -algebra generated by the kernel

$$\nu_{X, \bar{y}} : \mathbb{R} \times \mathcal{B}(\mathbb{R}) \rightarrow [0, 1], \quad (x, C) \mapsto \nu_{X, \bar{y}}(x; C) := \nu(x, \bar{y}; C).$$

It is easy to see that $\mathcal{A}_{X, \bar{y}}$ is the smallest σ -subalgebra of \mathcal{X} containing the \bar{y} -sections

$$A_{X, \bar{y}}(r, \varepsilon) := \{x \in \mathbb{R} : (x, \bar{y}) \in A(r, \varepsilon)\} = \{x \in \mathbb{R} : r - \bar{y} - \varepsilon \leq x \leq r - \bar{y} + \varepsilon\},$$

that is, $\mathcal{A}_{X, \bar{y}} = \mathcal{B}(\mathbb{R})$. This example shows that the cylinder sets $A \times \mathbb{R}$, $A \in \mathcal{A}_{X, \bar{y}}$, are not necessarily contained in \mathcal{A} (see illustration in Fig. 3).

With the σ -subalgebra \mathcal{A}_X of \mathcal{X} , we can now modify the random variable $X : (\Omega, \mathcal{F}, \mathbb{P}) \rightarrow (\mathbf{X}, \mathcal{X})$ by simply reducing the image σ -algebra to \mathcal{A}_X :

$$\widehat{X} : (\Omega, \mathcal{F}, \mathbb{P}) \rightarrow (\mathbf{X}, \mathcal{A}_X).$$

We will see that this step is crucial here, even though it might appear like a minor technical step at first sight. It allows us to modify (28) by replacing the full σ -algebra of X , \mathcal{X} , by the σ -algebra of \widehat{X} :

$$\mathbb{E}(\mathbb{1}_{\{Z \in C\}} | \widehat{X}).$$

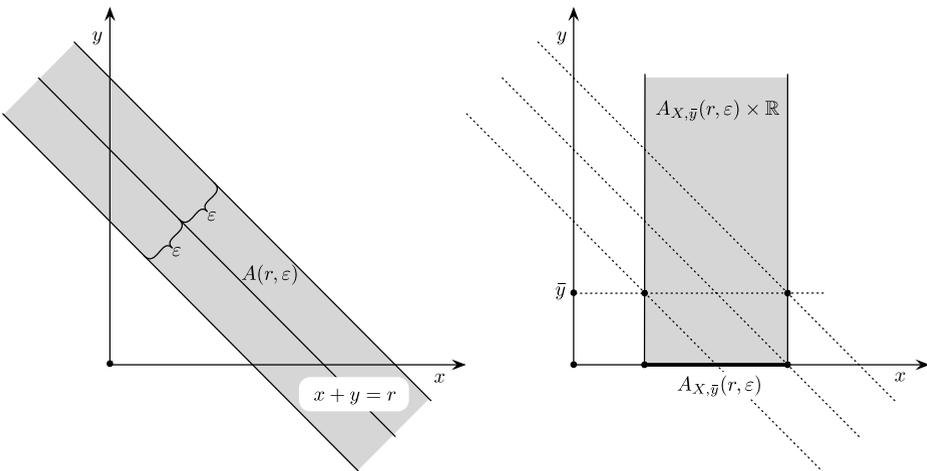


Fig. 3 Illustration of the ν -measurable sets in \mathbb{R}^2 and their sections in \mathbb{R}

This is, by definition, an \widehat{X} -measurable random variable $\Omega \rightarrow \mathbb{R}$. By the factorisation lemma, we can find a unique measurable function $\widehat{\nu}_X(\cdot; C) : (\mathcal{X}, \mathcal{A}_X) \rightarrow \mathbb{R}$ satisfying $\mathbb{E}(\mathbb{1}_{\{Z \in C\}} | \widehat{X}) = \widehat{\nu}_X(\widehat{X}; C)$. This yields a new marginal channel,

$$\mathbb{P}(Z \in C | \widehat{X} = x) := \widehat{\nu}_X(x; C),$$

as a modification of $\nu_X(x; C)$ which appears twice in (31). Note that the kernel $\widehat{\nu}_X(x; C)$ is defined almost surely. However, the definition of a conditional mutual information will be independent of the version of that kernel.

Now we come to the definition of a causal version of the mutual information (29) and the conditional mutual information (30). We simply replace in these definitions $\nu_X(x; C)$ by $\widehat{\nu}_X(x; C)$:

$$\begin{aligned} I_\xi(X \rightarrow Z) &:= I(\sigma(\widehat{X}); Z^{-1}(\xi)) \\ &= \sum_{C \in \xi} \int_{\mathcal{X}} \widehat{\nu}_X(x; C) \log \frac{\widehat{\nu}_X(x; C)}{\nu_*(\mu)(C)} \mu_X(dx), \end{aligned} \tag{34}$$

$$\begin{aligned} I_\xi(Y \rightarrow Z|X) &:= I(\sigma(Y); Z^{-1}(\xi) | \sigma(\widehat{X})) \\ &= \sum_{C \in \xi} \int_{\mathcal{X} \times \mathcal{Y}} \nu(x, y; C) \log \frac{\nu(x, y; C)}{\widehat{\nu}_X(x; C)} \mu(dx, dy). \end{aligned} \tag{35}$$

The following proposition relates the causal quantities (34) and (35) to the corresponding non-causal ones, (29) and (30).

Proposition 4 *We have the chain rule*

$$I_\xi(X, Y; Z) = I_\xi(X \rightarrow Z) + I_\xi(Y \rightarrow Z|X). \tag{36}$$

Furthermore,

$$I_\xi(X \rightarrow Z) \leq I_\xi(X; Z), \quad I_\xi(Y \rightarrow Z|X) \geq I_\xi(Y; Z|X). \tag{37}$$

Proof With C_z denoting the set in ξ that contains z , we have

$$\log \frac{\nu(x, y; C_z)}{\nu_*(\mu)(C_z)} = \log \frac{\nu(x, y; C_z)}{\widehat{\nu}_X(x; C_z)} + \log \frac{\widehat{\nu}_X(x; C_z)}{\nu_*(\mu)(C_z)}.$$

Integrating this with respect to $\nu(x, y; dz)$ we get

$$\begin{aligned} &\int_Z \log \frac{\nu(x, y; C_z)}{\nu_*(\mu)(C_z)} \nu(x, y; dz) \\ &= \int_Z \log \frac{\nu(x, y; C_z)}{\widehat{\nu}_X(x; C_z)} \nu(x, y; dz) + \int_Z \log \frac{\widehat{\nu}_X(x; C_z)}{\nu_*(\mu)(C_z)} \nu(x, y; dz) \\ &= \sum_{C \in \xi} \nu(x, y; C) \log \frac{\nu(x, y; C)}{\widehat{\nu}_X(x; C)} + \sum_{C \in \xi} \nu(x, y; C) \log \frac{\widehat{\nu}_X(x; C)}{\nu_*(\mu)(C)}. \end{aligned} \tag{38}$$

Further integrating the first term of (38) with respect to μ gives us $I_\xi(Y \rightarrow Z|X)$ (see (35)). For the corresponding integration of the second term, we obtain

$$\begin{aligned} & \sum_{C \in \xi} \int_{X \times Y} v(x, y; C) \log \frac{\hat{v}_X(x; C)}{v_*(\mu)(C)} \mu(dx, dy) \\ &= \sum_{C \in \xi} \int_{\Omega} \mathbb{P}(Z \in C|X, Y) \log \frac{\mathbb{P}(Z \in C|\hat{X})}{\mathbb{P}(Z \in C)} d\mathbb{P} \end{aligned} \tag{39}$$

$$= \sum_{C \in \xi} \int_{\Omega} \mathbb{P}(Z \in C|\hat{X}) \log \frac{\mathbb{P}(Z \in C|\hat{X})}{\mathbb{P}(Z \in C)} d\mathbb{P} \tag{40}$$

$$= \sum_{C \in \xi} \int_X \hat{v}_X(x; C) \log \frac{\hat{v}_X(x; C)}{v_*(\mu)(C)} \mu_X(dx). \tag{41}$$

The crucial step (40) follows from the general property of the conditional expectation of a function f with respect to a σ -subalgebra \mathcal{A} :

$$\int_{\Omega} fg \, d\mathbb{P} = \int_{\Omega} \mathbb{E}(f|\mathcal{A})g \, d\mathbb{P}, \quad \text{for all } \mathcal{A}\text{-measurable functions } g.$$

Here, f is given by $\mathbb{P}(Z \in C|X, Y)$, \mathcal{A} is the σ -algebra generated by \hat{X} , and g is given by $\log \frac{\mathbb{P}(Z \in C|\hat{X})}{\mathbb{P}(Z \in C)}$ which is \hat{X} -measurable. The steps (39) and (41) follow directly from the definitions of the Markov kernels, and we finally obtain $I_\xi(X \rightarrow Z)$ (see (34)). This concludes the proof of the chain rule (36).

We now prove the inequalities (37) where we can restrict attention to the first one. We consider the convex function $\phi(r) := r \log \frac{r}{v_*(\mu)(C)}$, for $r > 0$, and $\phi(0) := 0$, and apply Jensen’s inequality for conditional expectations:

$$\mathbb{E} \left(\phi(\mathbb{P}(Z \in C|X)) \middle| \hat{X} \right) \geq \phi \left(\mathbb{E} \left(\mathbb{P}(Z \in C|X) \middle| \hat{X} \right) \right) = \phi(\mathbb{P}(Z \in C|\hat{X})). \tag{42}$$

This implies

$$\begin{aligned} I_\xi(X; Z) &= \sum_{C \in \xi} \int_{\Omega} \mathbb{P}(Z \in C|X) \log \frac{\mathbb{P}(Z \in C|X)}{\mathbb{P}(Z \in C)} d\mathbb{P} \\ &= \sum_{C \in \xi} \mathbb{E}(\phi(\mathbb{P}(Z \in C|X))) \\ &= \sum_{C \in \xi} \mathbb{E} \left(\mathbb{E} \left(\phi(\mathbb{P}(Z \in C|X)) \middle| \hat{X} \right) \right) \\ &\geq \sum_{C \in \xi} \mathbb{E}(\phi(\mathbb{P}(Z \in C|\hat{X}))) \quad (\text{by (42)}) \\ &= \sum_{C \in \xi} \int_{\Omega} \mathbb{P}(Z \in C|\hat{X}) \log \frac{\mathbb{P}(Z \in C|\hat{X})}{\mathbb{P}(Z \in C)} d\mathbb{P} \\ &= I_\xi(X \rightarrow Z). \end{aligned}$$

The second inequality in (37) follows from the first one and the chain rule (36). □

Let us interpret this result. The first inequality in (37) highlights the fact that the stochastic dependence between X and Z , here quantified by the usual mutual information $I_\xi(X; Z)$, cannot be fully attributed to the causal effect of X on Z . Some part of $I_\xi(X; Z)$ is purely

associational, and $I_\xi(X \rightarrow Z)$ constitutes the causal part of it. The second inequality in (37) highlights a different fact. Conditioning on the variable X may “screen off” some part of the causal effect of Y on Z . More precisely, the uncertainty reduction about the outcome of Z through X can be so strong that a further reduction through Y becomes “invisible”. Therefore, the classical conditional mutual information, $I_\xi(Y; Z|X)$, tends to reflect only part of the causal influence of Y on Z given X , $I_\xi(Y \rightarrow Z|X)$. Even though the classical information-theoretic quantities are replaced by their causal versions, the full mutual information can still be decomposed according to the chain rule (36). However, in comparison to the decomposition (32), some amount of it is shifted from one term to the other so that both terms can be interpreted causally.

It turns out that the definitions (34) and (35) require a careful extension if we want to have a general chain rule for more than two input variables. We are now going to highlight this for three input variables.

4.2 Three Inputs

We now consider three input variables. This will reveal that the previous case with two input variables is quite special. An extension to more than two variables requires an adjustment of our definition of causal information flow.

We consider a third input variable (denoted below by W) with values in a measurable space (W, \mathcal{W}) , a probability measure

$$\mu \quad \text{on} \quad (W \times X \times Y, \mathcal{W} \otimes \mathcal{X} \otimes \mathcal{Y}),$$

and an input-output channel, given by a Markov kernel

$$v : W \times X \times Y \times \mathcal{Z} \rightarrow [0, 1].$$

This gives rise to a probability space, consisting of the measurable space

$$(\Omega, \mathcal{F}) := (W \times X \times Y \times Z, \mathcal{W} \otimes \mathcal{X} \otimes \mathcal{Y} \otimes \mathcal{Z}),$$

and the probability measure \mathbb{P} defined by

$$\mathbb{P}(A \times B \times C \times D) := \int_{A \times B \times C} v(w, x, y; D) \mu(dw, dx, dy).$$

Finally, we have the natural projections $W : \Omega \rightarrow W$, $X : \Omega \rightarrow X$, $Y : \Omega \rightarrow Y$, and $Z : \Omega \rightarrow Z$.

The definition of the marginal kernel $\hat{v}_X(x; C)$, as introduced in Section 4.1, is directly applicable to the situation of three input variables. It allows us to define marginal channels by an appropriate grouping of two input variables into one input variable, which formally reduces the three-input case to a two-input case. In particular, we can define the channels $\hat{v}_{W,X}(w, x; C)$ and $\hat{v}_W(w; C)$, by grouping W, X and X, Y , respectively, into one variable. Denoting the set in ξ that contains z by C_z , we then have

$$\log \frac{v(w, x, y; C_z)}{v_*(\mu)(C_z)} = \log \frac{v(w, x, y; C_z)}{\hat{v}_{W,X}(w, x; C_z)} + \log \frac{\hat{v}_{W,X}(w, x; C_z)}{\hat{v}_W(w; C_z)} + \log \frac{\hat{v}_W(w; C_z)}{v_*(\mu)(C_z)}.$$

By integration we obtain

$$\begin{aligned} \int_Z \log \frac{\nu(w, x, y; C_z)}{\nu_*(\mu)(C_z)} \nu(w, x, y; dz) &= \int_Z \log \frac{\nu(w, x, y; C_z)}{\hat{\nu}_{W,X}(w, x; C_z)} \nu(w, x, y; dz) \\ &\quad + \int_Z \log \frac{\hat{\nu}_{W,X}(w, x; C_z)}{\hat{\nu}_W(w; C_z)} \nu(w, x, y; dz) \\ &\quad + \int_Z \log \frac{\hat{\nu}_W(w; C_z)}{\nu_*(\mu)(C_z)} \nu(w, x, y; dz) \\ &= \sum_{C \in \xi} \nu(w, x, y; C) \log \frac{\nu(w, x, y; C)}{\hat{\nu}_{W,X}(w, x; C)} \end{aligned} \tag{43}$$

$$+ \sum_{C \in \xi} \nu(w, x, y; C) \log \frac{\hat{\nu}_{W,X}(w, x; C)}{\hat{\nu}_W(w; C)} \tag{44}$$

$$+ \sum_{C \in \xi} \nu(w, x, y; C) \log \frac{\hat{\nu}_W(w; C)}{\nu_*(\mu)(C)}. \tag{45}$$

An integration of the last term (45) with respect to μ yields, by the same reasoning as in the steps (39), (40), and (41),

$$I_\xi(W \rightarrow Z) = \sum_{C \in \xi} \int_W \hat{\nu}_W(w; C) \log \frac{\hat{\nu}_W(w; C)}{\nu_*(\mu)(C)} \mu_W(dw).$$

A corresponding integration of the first term (43) with respect to μ yields a non-negative quantity that can be interpreted as $I_\xi(Y \rightarrow Z|W, X)$ (see definition (35)). Even though we will have to slightly adjust this first term, the problem we are facing here is most clearly highlighted by the second term, (44). In order to naturally generalise the chain rule (36) we have to interpret the integral of the second term as $I_\xi(X \rightarrow Z|W)$. However, it turns out that, in general,

$$\begin{aligned} I_\xi(X \rightarrow Z|W) &= \sum_{C \in \xi} \int_{W \times X} \hat{\nu}_{W,X}(w, x; C) \log \frac{\hat{\nu}_{W,X}(w, x; C)}{\hat{\nu}_W(w; C)} \mu_{W,X}(dw, dx) \\ &\neq \sum_{C \in \xi} \int_{W \times X \times Y} \nu(w, x, y; C) \log \frac{\hat{\nu}_{W,X}(w, x; C)}{\hat{\nu}_W(w; C)} \mu_{W,X}(dw, dx, dy), \end{aligned} \tag{46}$$

where (46) is the integral of the term (44) with respect to μ . We cannot even ensure that this integral is non-negative. The reason is that the σ -algebra used for the definition of $\hat{\nu}_W(w; C)$ is not necessarily a σ -subalgebra of the one used for the definition of the kernel $\hat{\nu}_{W,X}(w, x; C)$ (the situation is similar to the one of Example 3). Therefore, the reasoning of the steps (39), (40), and (41), cannot be applied here.

The problem highlighted in this section will now be resolved. This will be done by a modification of the involved σ -algebras, which should define a filtration in order to imply a general causal version of the chain rule. In the next section, this modification will be presented for the general case of n input variables.

5 The General Definition of Causal Information Flow

5.1 Filtrations and Information

Let $(X_i, \mathcal{X}_i), i \in N := [n] = \{1, \dots, n\}$, be a family of measurable spaces, the state spaces of the input variables. For each subset M of N , we have the corresponding product space (X_M, \mathcal{X}_M) consisting of $X_M := \times_{i \in M} X_i$ and $\mathcal{X}_M := \otimes_{i \in M} \mathcal{X}_i$. Note that for $M = \emptyset$, the set X_\emptyset consists of one element, the empty sequence ϵ , and $\mathcal{X}_\emptyset = \{\emptyset, \{\epsilon\}\}$ is the trivial σ -algebra with two elements. In addition to the input variables, we consider an output variable with state space (Z, \mathcal{Z}) . The input-output channel is given by a Markov kernel

$$\nu : X_N \times \mathcal{Z} \rightarrow [0, 1].$$

Together with a probability measure μ on (X_N, \mathcal{X}_N) this defines the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ where

$$\Omega := X_N \times Z, \quad \mathcal{F} := \mathcal{X}_N \otimes \mathcal{Z},$$

and

$$\mathbb{P}(A \times C) := \int_A \nu(x; C) \mu(dx), \quad A \in \mathcal{X}_N, C \in \mathcal{Z}.$$

Finally, we have the canonical projections

$$X_M : \Omega \rightarrow X_M, \quad M \subseteq N, \quad \text{and} \quad Z : \Omega \rightarrow Z.$$

We are now going to define the M -marginal of the channel based on a general σ -subalgebra \mathcal{B}_M of \mathcal{X}_M . Below, in Section 5.2, this will allow us to incorporate causal aspects of ν by an appropriate adaptation of \mathcal{B}_M to ν . In order to highlight the flexibility that we have here, let us begin with the usual definition where \mathcal{B}_M equals the largest σ -subalgebra of \mathcal{X}_M , that is \mathcal{X}_M itself. Given a measurable set $C \subseteq Z$, we have the conditional expectation

$$\mathbb{E}(\mathbb{1}_{\{Z \in C\}} | X_M). \tag{47}$$

This is by definition an X_M -measurable function $\Omega \rightarrow \mathbb{R}$. By the factorisation lemma we can represent it as a composition $\mathbb{E}(\mathbb{1}_{\{Z \in C\}} | X_M) = \nu_M(X_M; C)$ with a measurable function $\nu_M(\cdot; C) : (X_M, \mathcal{X}_M) \rightarrow \mathbb{R}$. This allows us to define the conditional distribution

$$\mathbb{P}(Z \in C | X_M = x_M) := \nu_M(x_M; C) \tag{48}$$

which can be interpreted as a channel

$$\nu_M : X_M \times \mathcal{Z} \rightarrow [0, 1], \quad (x_M, C) \mapsto \nu_M(x_M; C).$$

We now modify the outlined marginalisation of ν by reducing the maximal σ -algebra \mathcal{X}_M to the σ -subalgebra \mathcal{B}_M . More precisely, we replace X_M in (47) by

$$\widehat{X}_M : (\Omega, \mathcal{F}, \mathbb{P}) \rightarrow (X_M, \mathcal{B}_M)$$

and consider the conditional expectation

$$\mathbb{E}(\mathbb{1}_{\{Z \in C\}} | \widehat{X}_M).$$

This is now an \widehat{X}_M -measurable function $\Omega \rightarrow \mathbb{R}$, and, by the factorisation lemma, we can represent it as a composition $\mathbb{E}(\mathbb{1}_{\{Z \in C\}} | \widehat{X}_M) = \widehat{\nu}(\widehat{X}_M; C)$ with a measurable function $\widehat{\nu}(\cdot; C) : (X_M, \mathcal{B}_M) \rightarrow \mathbb{R}$. Finally, we have the modification

$$\mathbb{P}(Z \in C | \widehat{X}_M = x_M) := \widehat{\nu}(x_M; C)$$

of the conditional distribution (48), which corresponds to a modified channel

$$\hat{\nu}_M : X_M \times \mathcal{Z} \rightarrow [0, 1], \quad (x_M, C) \mapsto \hat{\nu}_M(x_M; C).$$

By construction, $\hat{\nu}_M$ is \mathcal{B}_M -measurable, that is, it uses only information that is contained in \mathcal{B}_M . For the maximal σ -algebra we recover ν_M . We can also consider the other extreme where \mathcal{B}_M equals the smallest σ -algebra, $\{\emptyset, X_M\}$. In that case, we obtain $\hat{\nu}(x_M; C) = \nu_*(\mu)(C)$. An adjustment of \mathcal{B}_M to the information actually used by ν will allow us to interpret $\hat{\nu}$ causally. In contrast, if we do not have such an adjustment, $\hat{\nu}_M$ will represent a hypothetical channel, a “ghost channel”, based on the σ -algebra of an external observer rather than the σ -algebra of the actual mechanisms of the channel. As we will see in Section 5.2, there are various natural ways to adjust \mathcal{B}_M to ν . However, for the largest subset M of inputs, N , ν will always be measurable with respect to \mathcal{B}_N . This ensures that we recover ν when we condition with respect to \mathcal{B}_N , that is, $\hat{\nu}_N = \nu$.

We now consider a family $\mathcal{B} = (\mathcal{B}_M)_{M \subseteq N}$ of σ -algebras. It gives rise to a corresponding family

$$\mathcal{F}_M := X_M^{-1}(\mathcal{B}_M) \subseteq \mathcal{F}, \quad M \subseteq N,$$

of σ -algebras on Ω . We call the family \mathcal{B} *projective*, if the maps

$$\pi_L^M : X_M \rightarrow X_L, \quad x_M = (x_i)_{i \in M} \mapsto x_L = (x_i)_{i \in L}, \quad L \subseteq M \subseteq N,$$

are \mathcal{B}_M - \mathcal{B}_L -measurable. For projective families, we have the following monotonicity:

$$L \subseteq M \quad \Rightarrow \quad \mathcal{F}_L \subseteq \mathcal{F}_M. \tag{49}$$

Given a projective family \mathcal{B} , we now define a corresponding family of information-theoretic quantities which generalise (conditional) mutual information. We begin with a local version, applied to a measurable partition ξ of Z . For $z \in Z$, we denote the set in ξ that contains z by C_z . For $L \subseteq M \subseteq N$, we consider $x_M = (x_L, x_{M \setminus L}) \in X_M$ and define

$$i_{\xi}^{\mathcal{B}}(x_{M \setminus L} \rightarrow z | x_L) := \log \frac{\hat{\nu}_M(x_M; C_z)}{\hat{\nu}_L(x_L; C_z)}.$$

This is a local version of the conditional mutual information. Integration over z yields

$$\int_Z \log \frac{\hat{\nu}_M(x_M; C_z)}{\hat{\nu}_L(x_L; C_z)} \nu(x; dz) = \sum_{C \in \xi} \nu(x; C) \log \frac{\hat{\nu}_M(x_M; C)}{\hat{\nu}_L(x_L; C)}.$$

With a second integration, with respect to μ , we obtain

$$\begin{aligned} & \sum_{C \in \xi} \int_X \nu(x; C) \log \frac{\hat{\nu}_M(x_M; C)}{\hat{\nu}_L(x_L; C)} \mu(dx) \\ &= \sum_{C \in \xi} \int_{\Omega} \mathbb{E}(\mathbb{1}_{\{Z \in C\}} | X) \log \frac{\mathbb{E}(\mathbb{1}_{\{Z \in C\}} | \widehat{X}_M)}{\mathbb{E}(\mathbb{1}_{\{Z \in C\}} | \widehat{X}_L)} d\mathbb{P} \\ &= \sum_{C \in \xi} \int_{\Omega} \mathbb{E}(\mathbb{1}_{\{Z \in C\}} | \widehat{X}_M) \log \frac{\mathbb{E}(\mathbb{1}_{\{Z \in C\}} | \widehat{X}_M)}{\mathbb{E}(\mathbb{1}_{\{Z \in C\}} | \widehat{X}_L)} d\mathbb{P} \quad (\text{by the monotonicity (49)}) \\ &= \sum_{C \in \xi} \int_{X_M} \hat{\nu}_M(x_M; C) \log \frac{\hat{\nu}_M(x_M; C)}{\hat{\nu}_L(x_L; C)} \mu_M(dx_M), \end{aligned}$$

where μ_M denotes the M -marginal of μ . This suggests the following version of the conditional mutual information which we refer to as information flow.

Definition 5 Let $\mathcal{B} = (\mathcal{B}_M)_{M \subseteq N}$ be a projective family of σ -algebras, let ξ be a finite measurable partition of Z , and let $L \subseteq M \subseteq N$. Then we define the *information flow from $X_{M \setminus L}$ to Z given X_L* as

$$I_{\xi}^{\mathcal{B}}(X_{M \setminus L} \rightarrow Z | X_L) := \sum_{C \in \xi} \int_{X_M} \hat{v}_M(x_M; C) \log \frac{\hat{v}_M(x_M; C)}{\hat{v}_L(x_L; C)} \mu_M(dx_M).$$

For $L = \emptyset$ we simplify the notation by $I_{\xi}^{\mathcal{B}}(X_M \rightarrow Z)$ and refer to the *information flow from X_M to Z* .

Given disjoint subsets M_1, M_2, \dots, M_k of N , we use a filtration of σ -algebras for proving a general chain rule for information flows, thereby extending the chain rule in Proposition 4 and resolving the problem highlighted in Section 4.2 for three inputs.

Theorem 6 (General chain rule) *Let $\mathcal{B} = (\mathcal{B}_M)_{M \subseteq N}$ be a projective family of σ -algebras, let ξ be a finite measurable partition of Z , and let M_1, M_2, \dots, M_k be disjoint subsets of N . Then*

$$I_{\xi}^{\mathcal{B}}(X_{M_1}, \dots, X_{M_k} \rightarrow Z) = I_{\xi}^{\mathcal{B}}(X_{M_1} \rightarrow Z) + I_{\xi}^{\mathcal{B}}(X_{M_2} \rightarrow Z | X_{M_1}) + \dots + I_{\xi}^{\mathcal{B}}(X_{M_k} \rightarrow Z | X_{M_1}, X_{M_2}, \dots, X_{M_{k-1}}). \tag{50}$$

Proof Let $M^j := \cup_{i=1}^j M_i$, $j = 0, 1, \dots, k$. The monotonicity (49) implies that the sequence

$$\mathcal{F}_j := \mathcal{F}_{M^j} := X_{M^j}^{-1}(\mathcal{B}_{M^j}), \quad j = 0, \dots, k,$$

is increasing and therefore represents a filtration of σ -algebras. This implies

$$\begin{aligned} & I_{\xi}^{\mathcal{B}}(X_{M_1}, \dots, X_{M_k} \rightarrow Z) \\ &= \sum_{C \in \xi} \int_{X_{M^k}} \hat{v}(x_{M^k}; C) \log \frac{\hat{v}(x_{M^k}; C)}{v_*(\mu)(C)} \mu(dx_{M^k}) \\ &= \sum_{C \in \xi} \int_{X_{M^k}} \hat{v}(x_{M^k}; C) \log \left(\prod_{j=1}^k \frac{\hat{v}_{M^j}(x_{M^j}; C)}{\hat{v}_{M^{j-1}}(x_{M^{j-1}}; C)} \right) \mu(dx_{M^k}) \\ &= \sum_{j=1}^k \sum_{C \in \xi} \int_{X_{M^k}} \hat{v}(x_{M^k}; C) \log \frac{\hat{v}_{M^j}(x_{M^j}; C)}{\hat{v}_{M^{j-1}}(x_{M^{j-1}}; C)} \mu(dx_{M^k}) \\ &= \sum_{j=1}^k \sum_{C \in \xi} \int_{\Omega} \mathbb{P}(Z \in C | \widehat{X}_{M^k}) \log \frac{\mathbb{P}(Z \in C | \widehat{X}_{M^j})}{\mathbb{P}(Z \in C | \widehat{X}_{M^{j-1}})} d\mathbb{P} \\ &= \sum_{j=1}^k \sum_{C \in \xi} \int_{\Omega} \mathbb{P}(Z \in C | \widehat{X}_{M^j}) \log \frac{\mathbb{P}(Z \in C | \widehat{X}_{M^j})}{\mathbb{P}(Z \in C | \widehat{X}_{M^{j-1}})} d\mathbb{P} \\ &= \sum_{j=1}^k \sum_{C \in \xi} \int_{X_{M^j}} \hat{v}_{M^j}(x_{M^j}; C) \log \frac{\hat{v}_{M^j}(x_{M^j}; C)}{\hat{v}_{M^{j-1}}(x_{M^{j-1}}; C)} \mu(dx_{M^j}) \\ &= \sum_{j=1}^k I_{\xi}^{\mathcal{B}}(X_{M_j} \rightarrow Z | X_{M_1}, \dots, X_{M_{j-1}}). \end{aligned}$$

□

We now state basic properties of the information flow. Some of these properties are listed in [14] as natural postulates (P0–P4) for a measure of causal strength.

Proposition 7 (Natural properties) *Let $\mathcal{B} = (\mathcal{B}_M)_{M \subseteq N}$ be a projective family of σ -algebras such that ν is measurable with respect to \mathcal{B}_N , and let ξ be a finite measurable partition of Z . Then the following properties hold:*

- (a) *The information flow from all input variables to the output variable coincides with the mutual information: $I_{\xi}^{\mathcal{B}}(X_N \rightarrow Z) = I_{\xi}(X_N; Z)$.*
- (b) *For a subset M of N , the set of all input variables, the information flow $I_{\xi}^{\mathcal{B}}(X_M \rightarrow Z)$ is smaller than or equal to the mutual information $I_{\xi}(X_M; Z)$.*
- (c) *For a subset M of N , the information flow $I_{\xi}^{\mathcal{B}}(X_M \rightarrow Z|X_{N \setminus M})$ is greater than or equal to the conditional mutual information $I_{\xi}(X_M; Z|X_{N \setminus M})$.*
- (d) *If the information flow $I_{\xi}^{\mathcal{B}}(X_M \rightarrow Z|X_{N \setminus M})$ vanishes then Z is independent of X_M given $X_{N \setminus M}$.*
- (e) *For $L \subseteq M \subseteq N$, we have $I_{\xi}^{\mathcal{B}}(X_L \rightarrow Z) \leq I_{\xi}^{\mathcal{B}}(X_M \rightarrow Z)$. In particular, if $I_{\xi}^{\mathcal{B}}(X_M \rightarrow Z) = 0$ then $I_{\xi}^{\mathcal{B}}(X_L \rightarrow Z) = 0$.*

Proof Statement (a) follows from $\hat{\nu}_N(x_N; C) = \nu(x_N; C)$ and $\hat{\nu}_{\emptyset}(x; C) = \nu_*(\mu)(C)$. The statements (b) and (c) can be proven in the same way as the corresponding inequalities (37) of Proposition 4, thereby using the chain rule

$$I_{\xi}^{\mathcal{B}}(X_N \rightarrow Z) = I_{\xi}^{\mathcal{B}}(X_M \rightarrow Z) + I_{\xi}^{\mathcal{B}}(X_{N \setminus M} \rightarrow Z|X_M)$$

for $M \subseteq N$ (this follows from the general chain rule (50), with $M_1 = M$ and $M_2 = N \setminus M$). In order to prove (d), note that with (c) we have

$$I_{\xi}^{\mathcal{B}}(X_M \rightarrow Z|X_{N \setminus M}) = 0 \implies I_{\xi}(X_M; Z|X_{N \setminus M}) = 0.$$

This implies that X_M is independent of Z given $X_{N \setminus M}$. Finally, (e) follows from the chain rule

$$I_{\xi}^{\mathcal{B}}(X_M \rightarrow Z) = I_{\xi}^{\mathcal{B}}(X_L \rightarrow Z) + I_{\xi}^{\mathcal{B}}(X_{M \setminus L} \rightarrow Z|X_L),$$

by the general chain rule (50), with $M_1 = L$ and $M_2 = M \setminus L$. □

Now let us come back to the two- and three-input cases which we studied in Sections 4.1 and 4.2, respectively. These cases guided our search for the setting that allows us to prove the general chain rule. In the two-input case, we have considered the σ -algebra \mathcal{A}_X , defined by (33). Similarly, we can consider the correspondingly defined σ -algebras \mathcal{A}_Y and $\mathcal{A}_{X,Y}$. The latter is simply the σ -algebra generated by the channel ν . Example 3 then demonstrates that the projections $(x, y) \mapsto x$ and $(x, y) \mapsto y$ do not have to be \mathcal{A}_X, Y - \mathcal{A}_X - and \mathcal{A}_X, Y - \mathcal{A}_Y -measurable, respectively. This simply means that the family, extended by $\mathcal{A}_{\emptyset} = \{\emptyset, \{\epsilon\}\}$, is not necessarily projective. Why does the chain rule still hold this case? The reason is that we can simply extend $\mathcal{A}_{X,Y}$ to the full σ -algebra $\mathcal{X} \otimes \mathcal{Y}$, thereby making the family projective. This extension is trivial in the sense that it has no “visible” effect, simply because conditioning ν with respect to $\mathcal{X} \otimes \mathcal{Y}$ is the same as conditioning it with respect to $\mathcal{A}_{X,Y}$. As a result, we obtain the chain rule in Proposition 4 from Theorem 6. Going from the two-input case to the three-input case, we can define the corresponding family of σ -algebras, which is again not necessarily projective. However, in this case it is not possible to make it projective without visible effect. This demonstrates that the two-input case is quite special. If we want to consider the chain rule as a fundamental property of causal information flows, we have to accept the construction of a projective family of σ -algebras that are associated

with ν as a natural step. The individual causal information flows will then crucially depend on this family. In the following section, we are going to propose two natural ways of such a construction.

5.2 Adaptation of the Filtration to the Channel

We are now going to couple the family $(\mathcal{B}_M)_{M \subseteq N}$ to the channel ν so that we can interpret the corresponding marginals $(\hat{\nu}_M)_{M \subseteq N}$ causally. In order to simplify the presentation, we first consider an arbitrary σ -subalgebra \mathcal{A} of \mathcal{X}_N . (Below, \mathcal{A} will be chosen to be the σ -algebra generated by ν .) We begin with information in M in the context of a configuration \bar{x} outside of M , that is $\bar{x} \in X_{N \setminus M}$. Given such an \bar{x} , we define the (M, \bar{x}) -trace of \mathcal{A} as follows: For each $A \in \mathcal{A}$, we consider the (M, \bar{x}) -section of A ,

$$\text{sec}_{M, \bar{x}}(A) := A_{M, \bar{x}} := \{x \in X_M : (x, \bar{x}) \in A\}.$$

These sections then form the (M, \bar{x}) -trace of \mathcal{A} , that is

$$\text{tr}_{M, \bar{x}}(\mathcal{A}) := \mathcal{A}_{M, \bar{x}} := \{A_{M, \bar{x}} : A \in \mathcal{A}\}.$$

Considering all possible contexts $\bar{x} \in X_{N \setminus M}$, we finally define the M -trace of \mathcal{A} as

$$\text{tr}_M(\mathcal{A}) := \mathcal{A}_M := \bigvee_{\bar{x} \in X_{N \setminus M}} \mathcal{A}_{M, \bar{x}}.$$

The (M, \bar{x}) -trace as well as the M -trace of \mathcal{A} are σ -subalgebras of \mathcal{X}_M . Note that in the extreme cases $M = \emptyset$ and $M = N$, we recover $\mathcal{A}_\emptyset = \{\emptyset, \{\epsilon\} = X_\emptyset\}$ (where ϵ denotes the empty sequence), and $\mathcal{A}_N = \mathcal{A}$, respectively.

The family of all M -traces of \mathcal{A} describes how \mathcal{A} is “distributed” over the subsets M of N . However, there is a problem here: The canonical projections π_L^M are not necessarily \mathcal{A}_M - \mathcal{A}_L -measurable. This projectivity property is required for the definition of a measure of causal information flow that satisfies the general chain rule of Theorem 6. We highlighted this problem for the three-input case in Section 4.2. There are two ways to recover the projectivity, first by extending and second by reducing \mathcal{A}_M appropriately. Let us begin with the extension:

$$\overline{\mathcal{A}}_M := \bigvee_{L \subseteq M} \left(\pi_L^M\right)^{-1}(\mathcal{A}_L). \tag{51}$$

We have the following characterisation of the family $\overline{\mathcal{A}}_M, M \subseteq N$, as the smallest projective extension of the family $\mathcal{A}_M, M \subseteq N$.

Proposition 8 (Extension of $\mathcal{A}_M, M \subseteq N$) *The family $\overline{\mathcal{A}}_M, M \subseteq N$, satisfies the following two conditions:*

1. For all $M \subseteq N$, \mathcal{A}_M is contained in $\overline{\mathcal{A}}_M$.
2. For all $L \subseteq M \subseteq N$, the canonical projection π_L^M is $\overline{\mathcal{A}}_M$ - $\overline{\mathcal{A}}_L$ -measurable.

Furthermore, for every family $\mathcal{A}'_M, M \subseteq N$, that satisfies these two conditions (where $\overline{\mathcal{A}}_M$ is replaced by \mathcal{A}'_M), we have

$$\overline{\mathcal{A}}_M \subseteq \mathcal{A}'_M \quad \text{for all } M \subseteq N. \tag{52}$$

Proof The first statement is clear (simply choose on the RHS of (51) $L = M$). For the second statement, we have to show

$$\left(\pi_L^M\right)^{-1}\left(\overline{\mathcal{A}}_L\right)=\overline{\mathcal{A}}_M.$$

Given that

$$\overline{\mathcal{A}}_L=\bigvee_{K \subseteq L}\left(\pi_K^L\right)^{-1}\left(\mathcal{A}_K\right)$$

it is sufficient to verify

$$\left(\pi_L^M\right)^{-1}\left(\left(\pi_K^L\right)^{-1}\left(\mathcal{A}_K\right)\right) \subseteq \overline{\mathcal{A}}_M \quad \text{for all } K \subseteq L. \tag{53}$$

The LHS of (53) reduces to $\left(\pi_K^M\right)^{-1}\left(\mathcal{A}_K\right)$ which is by definition contained in $\overline{\mathcal{A}}_M$.

Finally, we prove the minimality. It is easy to see that any family $\mathcal{A}'_M, M \subseteq N$, that satisfies the two conditions has to contain $\left(\pi_L^M\right)^{-1}\left(\mathcal{A}_L\right), L \subseteq M$. By definition, $\overline{\mathcal{A}}_M$ is the smallest σ -algebra that contains these σ -subalgebras (see (51)). This implies (52). \square

After having defined the smallest extension of the family $\mathcal{A}_M, M \subseteq N$, as one way to recover projectivity, we now come to the alternative way, which is by reduction of that family. More precisely, we define

$$\underline{\mathcal{A}}_M:=\left\{A \in \mathcal{X}_M:\left(\pi_M^N\right)^{-1}(A) \in \mathcal{A}\right\}. \tag{54}$$

We have the following characterisation of this family as the largest projective reduction of $\mathcal{A}_M, M \subseteq N$.

Proposition 9 (Reduction of $\mathcal{A}_M, M \subseteq N$) *The family $\underline{\mathcal{A}}_M, M \subseteq N$, satisfies the following two conditions:*

1. For all $M \subseteq N, \underline{\mathcal{A}}_M$ is contained in \mathcal{A}_M .
2. For all $L \subseteq M \subseteq N$, the canonical projection π_L^M is $\underline{\mathcal{A}}_M$ - $\underline{\mathcal{A}}_L$ -measurable.

Furthermore, for every family $\mathcal{A}'_M, M \subseteq N$, that satisfies these two conditions (where $\underline{\mathcal{A}}_M$ is replaced by \mathcal{A}'_M), we have

$$\mathcal{A}'_M \subseteq \underline{\mathcal{A}}_M \quad \text{for all } M \subseteq N.$$

Proof In order to prove the first statement, let $A \in \underline{\mathcal{A}}_M$. This means that

$$\tilde{A}:=\left(\pi_M^N\right)^{-1}(A)=A \times \mathbf{X}_{N \setminus M} \in \mathcal{A}.$$

For all $\bar{x} \in \mathbf{X}_{N \setminus M}$, we have

$$\text{sec}_{M, \bar{x}}(\tilde{A})=\left\{x \in \mathbf{X}_M:(x, \bar{x}) \in \tilde{A}\right\}=A.$$

This means that $A \in \text{tr}_M(\mathcal{A}) = \mathcal{A}_M$, which concludes the proof of the first statement. Now we come to the measurability of the canonical projection π_L^M . For this, we choose $A \in \underline{\mathcal{A}}_L$ and have to show $\left(\pi_L^M\right)^{-1}(A) \in \underline{\mathcal{A}}_M$:

$$\left(\pi_M^N\right)^{-1}\left(\left(\pi_L^M\right)^{-1}(A)\right)=\left(\pi_L^N\right)^{-1}(A) \in \mathcal{A} \quad \text{(by definition (54)).}$$

Finally, we have to prove the maximality. Let $\mathcal{A}'_M, M \subseteq N$, be a family that satisfies the two conditions. Then

$$\left(\pi_M^N\right)^{-1}\left(\mathcal{A}'_M\right) \subseteq \mathcal{A}'_N \subseteq \mathcal{A}_N = \mathcal{A}.$$

This means that $\mathcal{A}'_M \subseteq \underline{\mathcal{A}}_M$. □

By the definitions, we have

$$\underline{\mathcal{A}}_M \subseteq \mathcal{A} \subseteq \overline{\mathcal{A}}_M, \quad M \subseteq N,$$

where equalities hold for $M = N$ and $M = \emptyset$. More precisely,

$$\underline{\mathcal{A}}_N = \mathcal{A} = \overline{\mathcal{A}}_N, \quad \underline{\mathcal{A}}_\emptyset = \{\emptyset, \{\epsilon\}\} = \overline{\mathcal{A}}_\emptyset.$$

This concludes the constructions for a given σ -algebra \mathcal{A} , without explicit reference to the channel $\nu : X \times \mathcal{Z} \rightarrow [0, 1]$. We now couple the studied σ -algebras with the channel ν and therefore choose \mathcal{A} to be the σ -algebra generated by the channel ν , that is $\sigma(\nu)$. We highlight this coupling by writing \mathcal{A}^ν , as a particular choice of \mathcal{A} , and consider the family $(\mathcal{A}_M^\nu)_{M \subseteq N}$ of its traces, together with the corresponding smallest projective extension $(\overline{\mathcal{A}}_M^\nu)_{M \subseteq N}$ and the largest projective reduction $(\underline{\mathcal{A}}_M^\nu)_{M \subseteq N}$. In the context of a channel, the traces of \mathcal{A}^ν have a natural interpretation. In order to see this, we first consider a configuration $\bar{x} \in X_{N \setminus M}$ and define the “constrained” Markov kernel

$$\nu_{M, \bar{x}} : X_M \times \mathcal{Z} \rightarrow [0, 1], \quad \nu_{M, \bar{x}}(x; C) := \nu(x, \bar{x}; C).$$

We denote the σ -algebra generated by $\nu_{M, \bar{x}}$ by $\sigma_{M, \bar{x}}(\nu)$. Taking all “constraints” \bar{x} into account, we then define

$$\sigma_M(\nu) := \bigvee_{\bar{x} \in X_{N \setminus M}} \sigma_{M, \bar{x}}(\nu).$$

Proposition 10 *Let $\mathcal{A}^\nu \subseteq \mathcal{X}_N$ be the σ -algebra generated by the Markov kernel $\nu : X_N \times \mathcal{Z} \rightarrow [0, 1]$. Then for all $M \subseteq N$ and all $\bar{x} \in X_{N \setminus M}$,*

$$\sigma_{M, \bar{x}}(\nu) = \text{tr}_{M, \bar{x}}(\mathcal{A}^\nu) \quad \text{and} \quad \sigma_M(\nu) = \text{tr}_M(\mathcal{A}^\nu).$$

Proof The σ -algebra \mathcal{A}^ν is the smallest σ -algebra that contains all measurable sets of the form

$$A = \{x \in X_N : \nu(x; C) \in B\}, \tag{55}$$

with some $C \in \mathcal{Z}$ and a Borel set B in $\mathcal{B}([0, 1])$. Now consider the (M, \bar{x}) -section of such a set A :

$$\begin{aligned} \text{sec}_{M, \bar{x}}(A) &= \{x \in X_{N \setminus M} : (x, \bar{x}) \in A\} \\ &= \{x \in X_{N \setminus M} : \nu(x, \bar{x}; C) \in B\} \\ &= \{x \in X_{N \setminus M} : \nu_{M, \bar{x}}(x; C) \in B\}. \end{aligned}$$

This shows that the sections $\text{sec}_{M, \bar{x}}(A)$ of measurable sets A of the form (55) generate $\sigma_{M, \bar{x}}(\nu)$, which proves the first equality. The second equality is a direct implication of the first one. □

The results of the previous section, Theorem 6 and Proposition 7, apply to the information flows, defined for the projective families $(\overline{\mathcal{A}}_M^\nu)_{M \subseteq N}$ and $(\underline{\mathcal{A}}_M^\nu)_{M \subseteq N}$. These families take into account the information that is actually used by the channel ν . Therefore, we can interpret the corresponding marginal channels $\hat{\nu}_M$ causally, where we have to distinguish

two kinds of causality. For the projective family $(\overline{\mathcal{A}}_M^v)_{M \subseteq N}$, the channel $\hat{\nu}_M$ incorporates the information in any input configuration x_K , $K \subseteq M$, that is used by ν in conjunction with a context configuration $\bar{x}_{N \setminus K} = \bar{x}_{N \setminus M} \bar{x}_{M \setminus K}$ outside of K . For the projective family $(\mathcal{A}_M^v)_{M \subseteq N}$, on the other hand, the channel $\hat{\nu}_M$ incorporates the information used by ν that is solely contained in x_M , independent of any context. When comparing a marginal channel $\hat{\nu}_M$ with another marginal channel $\hat{\nu}_L$, where $L \subseteq M$, the corresponding information flows $\bar{I}_\xi(X_{M \setminus L} \rightarrow Z|X_L)$ and $I_\xi(X_{M \setminus L} \rightarrow Z|X_L)$, respectively, quantify the causal effects in $\hat{\nu}_M$ that exceed those in $\hat{\nu}_L$. These measures will capture different causal aspects, where the difference can be large. This is illustrated by the following extension of Example 3.

Example 11 Let

$$(X_i, \mathcal{X}_i) = (\mathbb{R}, \mathcal{B}(\mathbb{R})), \quad i \in \{1, \dots, n\} = N,$$

where $\mathcal{B}(\mathbb{R})$ denotes the Borel σ -algebra of \mathbb{R} . We define the channel simply by the sum of the input states, interpreted as a Markov kernel,

$$\nu(x_1, \dots, x_n; C) := \mathbb{1}_C(x_1 + \dots + x_n).$$

As $\mathcal{B}(\mathbb{R})$ is generated by the intervals $[r - \varepsilon, r + \varepsilon] \subseteq \mathbb{R}$, the smallest σ -algebra \mathcal{A}^ν for which all functions $\nu(\cdot; C)$ are measurable is generated by the following sets

$$A(r, \varepsilon) := \{(x_1, \dots, x_n) \in \mathbb{R}^n : r - \varepsilon \leq x_1 + \dots + x_n \leq r + \varepsilon\}, \quad r \in \mathbb{R}, \varepsilon \in \mathbb{R}_+.$$

For a set $M \subseteq N$ and a context configuration $\bar{x} = (\bar{x}_i)_{i \in N \setminus M} \in \mathbb{R}^{N \setminus M}$, the (M, \bar{x}) -section of $A_{r, \varepsilon}$ is given by

$$\text{sec}_{M, \bar{x}}(A(r, \varepsilon)) = \left\{ x = (x_i)_{i \in M} \in \mathbb{R}^M : r - \sum_{i \in N \setminus M} \bar{x}_i - \varepsilon \leq \sum_{i \in M} x_i \leq r - \sum_{i \in N \setminus M} \bar{x}_i + \varepsilon \right\}.$$

Therefore, the M -trace of \mathcal{A}^ν , \mathcal{A}_M^v , is generated by the halfspaces

$$H_\vartheta := \left\{ x = (x_i)_{i \in M} \in \mathbb{R}^M : \sum_{i \in M} x_i \leq \vartheta \right\}, \quad \vartheta \in \mathbb{R}.$$

For $|M| = 1$, we recover the half lines, so that $\mathcal{A}_{\{i\}}^v = \mathcal{B}(\mathbb{R})$. The projective extension then leads to the largest σ -algebra, the Borel algebra of \mathbb{R}^M :

$$\overline{\mathcal{A}}_M^v = \mathcal{B}(\mathbb{R}^M).$$

Therefore, the marginal channel $\hat{\nu}_M(x; C)$ equals the usual marginal $\nu_M(x; C)$ for the projective extension. For the projective reduction, on the other hand, we obtain the trivial σ -algebra except for $M = N$:

$$\underline{\mathcal{A}}_M^v = \begin{cases} \{\emptyset, \mathbb{R}^M\} & \text{if } M \subsetneq N, \\ \mathcal{A}^\nu & \text{if } M = N. \end{cases}$$

In this case we have $\hat{\nu}_M(x; C) = \nu_*(\mu)(C)$ for $M \neq N$ and $\hat{\nu}_N(x; C) = \nu(x; C)$, where μ is the joint distribution of the input variables.

We now consider the information flows associated with $L \subsetneq M \subseteq N$, for the projective extension as well as for the projective reduction. In both cases these flows coincide with usual (conditional) mutual informations, in an instructive way. More precisely, for the extension we have

$$\bar{I}_\xi(X_{M \setminus L} \rightarrow Z|X_L) = I_\xi(X_{M \setminus L}; Z|X_L).$$

For the reduction, we obtain

$$I_{\xi}(X_{M \setminus L} \rightarrow Z | X_L) = \begin{cases} 0 & \text{if } M \subsetneq N, \\ I_{\xi}(X_N; Z) & \text{if } M = N. \end{cases} \tag{56}$$

Interestingly, (56) does not depend on L . The vanishing of the information flow for $M \neq N$ is due to the fact that the output of the channel, the sum $x_1 + \dots + x_n$, cannot be computed from a proper subset of the inputs. The flow of information only takes place if all inputs are given.

The following example, which is closely related to Example 1, highlights continuity issues of the introduced measures of information flow. They result from the fact that small changes in the mechanisms can lead to large differences of the involved σ -algebras.

Example 12 Consider two input variables, X and Y , and one output variable Z , with corresponding state spaces

$$(X, \mathcal{X}) = (Y, \mathcal{Y}) = (Z, \mathcal{Z}) = (\mathbb{R}, \mathcal{B}(\mathbb{R})).$$

Furthermore, consider two channels, ν and ν' , where ν simply copies the second input, y , and ν' the first one, x . More precisely,

$$\begin{aligned} \nu(x, y; C) &:= \mathbb{1}_C(y), \\ \nu'(x, y; C) &:= \mathbb{1}_C(x). \end{aligned}$$

With $0 \leq \varepsilon \leq 1$, we define the convex combination

$$\nu^{(\varepsilon)} := (1 - \varepsilon)\nu + \varepsilon\nu'.$$

For $\varepsilon = 0$, only the channel ν is acting. Obviously, we have $I_{\xi}(X \rightarrow Z) = \bar{I}_{\xi}(X \rightarrow Z) = 0$, as expected in this case, because ν simply copies y and is not sensitive to x at all. One might intuitively expect that the causal information flow from X to Z stays close to 0 if ε is small but greater than 0. However, this intuition is not reflected by the actual quantities, as defined in this article. It is easy to see that for $\varepsilon \neq 0$, $I_{\xi}(X \rightarrow Z) = \bar{I}_{\xi}(X \rightarrow Z) = I_{\xi}(X; Z)$. Thus, these quantities are not at all sensitive to the parameter ε and behave discontinuously in the limit $\varepsilon \rightarrow 0$.

6 Conclusions

Conditioning is an important operation within the study of causality. The theory of causal networks, pioneered by Pearl [17], introduces interventional conditioning as an operation, the so-called *do-operation*, that is fundamentally different from the classical conditioning based on the general rule $\mathbb{P}(B|A) = \mathbb{P}(A \cap B)/\mathbb{P}(A)$. It models more appropriately experimental setups and avoids confusion with purely associational dependencies. Information theory has been classically used for the quantification of such dependencies, in terms of mutual information and conditional mutual information [20]. Within the original setting of information theory, the mutual information between the input and the output of a channel can be interpreted causally. In the more general context of causal networks, however, confounding effects make a distinction between associations and causal effects more difficult. In such cases, information-theoretic quantities can be misleading as measures of causal effects. In order to overcome this problem, information theory has been coupled with the

interventional calculus of causal networks, and corresponding measures of causal information flow have been proposed [5, 6]. Given that such measures are based on the notion of an experimental intervention, which represents a perturbation of the system, it remains unclear to what extent they quantify causal information flows in the unperturbed system. As another consequence of the interventional conditioning, one cannot expect that causal information flow, as defined in [6], decomposes according to a chain rule. The current article is based on an idea of the author from 2003 which precedes the above-mentioned works on combining the theory of causal networks with information theory. It proposes a way to quantify causal information flows without perturbing the system through intervention. Instead, it is based on classical conditioning in terms of the conditional distribution $\mathbb{P}(B|\mathcal{A})$, where the σ -algebra \mathcal{A} is adjusted to the intrinsic mechanisms of the system. The derived information flow measure satisfies the chain rule and the natural properties of a general measure of causal strength postulated in [14]. The chain rule, together with the generalised Pythagoras relation from information geometry, provide powerful tools within the study of the problem of partial information decomposition [7, 8, 16].

Even though the introduced information flows satisfy natural properties, the aim of the present article is relatively moderate. For instance, the analysis is focussed on a simple network consisting of a number of inputs and one output, which is a strong restriction compared to the setting of [6]. The extension of the present work to more general casual networks remains to be worked out. Furthermore, this article does not address the important problem of causal inference [18]. In addition to these general directions of research, there are various ways to modify and extend the constructions of the present work and thereby potentially highlight further causal aspects of a given channel. The following perspectives are particularly important:

1. In the present article, the information flow has been defined for a fixed finite measurable partition ξ of the state space (Z, \mathcal{Z}) of the output variable Z . A natural further step would be to consider the limit of information flows with respect to an increasing sequence ξ_n , $n = 1, 2, \dots$, so that

$$\bigvee_{n=1}^{\infty} \sigma(\xi_n) = \mathcal{Z}.$$

This limit will be an information flow measure that is independent of a particular partition.

2. Throughout this article, the partition ξ has not been coupled with the σ -algebra of the channel ν . This is the smallest σ -algebra for which all functions $\nu(x; C)$, $C \in \mathcal{Z}$, are measurable. Given that the channel is analysed with respect to the partition ξ , one can restrict attention to the smallest σ -algebra for which the functions $\nu(x; C)$, $C \in \xi$, are measurable. This will be a potentially small σ -subalgebra of the one generated by the channel. We would then have a natural coupling of the partition ξ with the information used by the channel.
3. We started with the family \mathcal{A}_M^ν of M -traces of \mathcal{A}^ν , the σ -algebra generated by ν , as the natural family associated with the channel. However, these traces do not form a projective family of σ -algebras. Such a projectivity is required for the chain rule for corresponding information flows. One can recover projectivity by extension and by reduction, leading to $\overline{\mathcal{A}}_M^\nu$ and $\underline{\mathcal{A}}_M^\nu$, respectively. Example 11 shows that the extension can lead to the largest σ -algebra and the reduction to the trivial one. Given this fact, one might ask whether the extension is too large and the reduction is too small to capture the causal aspects of ν . Even though we argued above that these two projective

families associated with ν capture two different kinds of causal aspects, this question remains to be further pursued. One possible direction would be the analysis of the context-dependent traces of \mathcal{A}^ν , that is the family of $\text{tr}_{M, \bar{x}}(\mathcal{A}^\nu)$, $\bar{x} \in X_{N \setminus M}$. Instead of conditioning with respect to the join

$$\text{tr}_M(\mathcal{A}^\nu) = \bigvee_{\bar{x} \in X_{N \setminus M}} \text{tr}_{M, \bar{x}}(\mathcal{A}^\nu),$$

one could adjust the conditioning to the individual σ -algebras $\text{tr}_{M, \bar{x}}(\mathcal{A}^\nu)$. This would represent an important refinement of the presented theory.

Acknowledgements The author is grateful for valuable comments of two anonymous reviewers. He acknowledges the support of the Deutsche Forschungsgemeinschaft Priority Programme “The Active Self” (SPP 2134).

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Amari, S.-i.: Information Geometry and its Applications. Applied Mathematical Sciences, vol. 194. Springer, Tokyo (2016)
2. Amari, S.-i., Nagaoka, H.: Methods of Information Geometry. Translations of Mathematical Monographs, vol. 191. Oxford University Press, Oxford (2000)
3. Ay, N., Amari, S.-i.: A novel approach to canonical divergences within information geometry. *Entropy* **17**, 8111–8129 (2015)
4. Ay, N., Jost, J., Lê, H.V., Schwachhöfer, L.: Information Geometry. A Series of Modern Surveys in Mathematics, vol. 64. Springer International Publishing, New York (2017)
5. Ay, N., Krakauer, D.C.: Geometric robustness theory and biological networks. *Theory Biosci.* **125**, 93–121 (2007)
6. Ay, N., Polani, D.: Information flows in causal networks. *Adv. Complex Syst.* **11**, 17–41 (2008)
7. Ay, N., Polani, D., Virgo, N.: Information decomposition based on cooperative game theory. *Kybernetika* **56**, 979–1014 (2020)
8. Bertschinger, N., Rauh, J., Olbrich, E., Jost, J., Ay, N.: Quantifying unique information. *Entropy* **16**, 2161–2183 (2014)
9. Bossomaier, T., Barnett, L., Harré, M., Lizier, J.T.: An Introduction to Transfer Entropy. Springer International Publishing (2016)
10. Cover, T.M., Thomas, J.A. Elements of Information Theory, 2nd edn. Wiley-Interscience, Hoboken (2006)
11. Dudley, R.M.: Real Analysis and Probability, 2nd edn. Cambridge Studies in Advanced Mathematics, vol. 74. Cambridge University Press, Cambridge (2002)
12. Granger, C.W.J.: Investigating causal relations by econometric models and cross-spectral methods. *Econometrica* **37**, 424–438 (1969)
13. Granger, C.W.J.: Testing for causality: A personal viewpoint. *J. Econ. Dyn. Control* **2**, 329–352 (1980)
14. Janzing, D., Balduzzi, D., Grosse-Wentrup, M., Schölkopf, B.: Quantifying causal influences. *Ann. Stat.* **41**, 2324–2358 (2013)
15. Kawahara, Y.: Abstract Methods in Information Theory. Multivariate Analysis, vol. 4. World Scientific, New Jersey (1999)

16. Lizier, J., Bertschinger, N., Jost, J., Wibral, M.: Information decomposition of target effects from multi-source interactions: Perspectives on previous, current and future work. *Entropy* **20**, 307 (2018)
17. Pearl, J.: *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge (2000)
18. Peters, J., Janzing, D., Schölkopf, B.: *Elements of Causal Inference: Foundations and Learning Algorithms*. Adaptive Computation and Machine Learning Series. The MIT Press, Cambridge (2017)
19. Schreiber, T.: Measuring information transfer. *Phys. Rev. Lett.* **85**, 461–464 (2000)
20. Shannon, C.E.: A mathematical theory of communication. *Bell Syst. Tech. J.* **27**, 623–656 (1948)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.