Annals of Mathematics and Artificial Intelligence manuscript No.
(will be inserted by the editor)

Qualitative and Quantitative Conditions for the Transitivity of Perceived Causation

Theoretical and Experimental Results

Jean-François Bonnefon · Rui Da Silva Neves · Didier Dubois · Henri Prade

Received: date / Accepted: date

Abstract If A caused B and B caused C, did A cause C? Although laypersons commonly perceive causality as being transitive, some philosophers have questioned this assumption, and models of causality in artificial intelligence are often agnostic with respect to transitivity: They define causation, then check whether the definition makes all, or only some, causal arguments transitive. We consider two formal models of causation that differ in the way they represent uncertainty. The quantitative model uses a crude probabilistic definition, arguably the common core of more sophisticated quantitative definitions; the qualitative model uses a definition based on nonmonotonic consequence relations. Different sufficient conditions for the transitivity of causation are laid bare by the two models: The Markov condition on events for the quantitative model, and a Saliency condition (A is perceived as a typical cause of B) for the qualitative model. We explore the formal and empirical relations between these sufficient conditions, and between the underlying definitions of perceived causation. These connections shed light on the range of applicability of each model. They lead to some speculations about commonsense causal reasoning (which would follow a qualitative model), contrasted with scientific causation

This research was supported by a grant from the Agence Nationale de la Recherche, project number NT05-3-44479.

Jean-François Bonnefon and Rui Da Silva Neves
CLLE (CNRS, UTM, EPHE), Maison de la Recherche
5 al. A. Machado 31058 Toulouse Cedex 9 France
Tel.: +33-5-61503522
Fax: +33-5-61503522
E-mail: bonnefon,neves@univ-tlse2.fr

Didier Dubois and Henri Prade
IRIT (CNRS, UPS), Universit Paul Sabatier
118 Route de Narbonne, F-31062 Toulouse Cedex 9 France
Tel.: +33-5-61556765
Fax: +33-5-61556258
E-mail: dubois,prade@irit.fr

(which would follow a quantitative model). These speculations are supported by a series of three behavioral experiments.

Keywords Causality · Transitivity · Uncertainty · Cognitive Validity

1 Introduction

Making yourself some tea, you put your kettle on the stove. Moments later, the kettle whistles because the water is boiling. The water is boiling because it has been heated to 100 degrees. Is the kettle whistling because the water has been heated to 100 degrees? Most of us agree that it is the case. The kettle example is one where it seems natural to accept that 'A causes C' results from 'A causes B' and 'B causes C.' That is, it makes causal transitivity appear unproblematic.

Although it has always been a strong temptation for laypersons to consider that causality on events is necessarily a transitive relation, philosophers have debated about that assumption. Earlier instances of such a discussion can be found in [32, 40, 41]. E. J. Lowe [40] observed that some transitive causal conclusions were odd, citing as an example the well-known nursery rhyme:

> For want of a nail the shoe was lost,
> For want of a shoe the horse was lost,
> For want of a horse the rider was lost,
> For want of a rider the battle was lost,
> For want of a battle the kingdom was lost,
> And all for the want of a horseshoe nail.

Lowe noted that the transitive conclusion that the want of a nail caused the loss of a kingdom was unplausible, while all the intermediate causal relations were, or could be made, acceptable. J. L. Mackie [41] objected that the transitive conclusion was perfectly acceptable if one interpreted it as meaning that the want of a nail was a partial cause of the loss of a kingdom, even though it was not a subtantial causal contribution. Mackie however conceded that causation might be intransitive if one required the expressed cause to play a substantial role in the production of the effect. Relatedly, G. Hesslow [32] suggested that intransitivity in this example might concern explanation rather than causation itself. That is, that although the transitive causal conclusion was in fact correct, it sounded odd because the want of a nail had weak explanatory relevance in regard to the loss of a kingdom.

More recently, many more examples of intransitive causal arguments were discussed in the field of philosophy [9, 30, 33, 35], strongly suggesting to let go of the constraint that causation should be unconditionally transitive. Although we will not focus our analyses on these examples, we give two of them here as an illustration:

> A man's finger is severed in a factory accident. He is rushed to the hospital, where an expert surgeon reattaches the finger, doing such a splendid job that a year later, it functions as well as if the accident had never happened. The accident causes the surgery, which, in turn, causes the finger to be in a healthy state a year later. But,

> intuitively, the accident does not cause the finger to be in a healthy state a year later [35].

> An assassin plants a bomb under my desk. I find it, and safely remove it. His planting the bomb causes my finding it, which, in turn, causes my continued survival. But, intuitively, his planting it does not cause my continued survival [30].

Note that in the first example where transitivity sounds counterintuitive, there is a notion of *intervention* [44] that breaks transitivity : the surgery is a voluntary act that stops the potential consequences of the finger accident. In the second example, planting the bomb sounds more of a prerequisite for finding it than an actual cause.

In the field of Artificial Intelligence, there has been no strong stance as to whether causation should be transitive (for exceptions, see e.g., [4]). Indeed, models of causality in artificial intelligence such as [31] often take an agnostic stance with respect to transitivity, by defining causation first and then checking whether the definition makes all, or only some, causal arguments transitive. For example, Pearl [44] explains (p. 237) that the transitivity of causation becomes natural if it is understood in terms of indirect influence under a Markovian condition: if A causes B, B causes C *regardless of* A, then A is understood as causing (indirectly) C. We will indeed have more to say on the connection between transitivity and the Markov condition.

In this article, we consider two models of uncertain causation that differ in the way they represent uncertainty. The first model uses a basic probabilistic definition of 'A causes B' as $\Pr(B|A) > \Pr(B)$, or equivalently $\Pr(B \mid A) > \Pr(B \mid \bar{A})$, originally discussed by Good [28,29]. Quite uncontroversially, this definition is too crude, especially from a normative perspective. For example, it would certainly be inadvisable to use this definition to extract causal relations from a body of scientific data. Various augmentations of this definition have accordingly been proposed, most notably by Pearl [44]. We will, however, limit ourselves in this article to this crude definition. Our motivation is twofold. First, our purpose is not to investigate a model of how scientists or other professionals *should* judge causation, but rather to investigate how laypersons might perceive causation. As it turns out, the crude probabilistic definition above provides the basis for one major psychological theory of how laypersons perceive causation [15–17]. Second, we find it advantageous to limit ourselves to the common core of contemporary, more sophisticated probabilistic definitions of causation. Using a richer definition would amount to limiting our analysis to a single probabilistic model of causation. Because all these more sophisticated definitions at least share the idea that $\Pr(B \mid A) > \Pr(B \mid \bar{A})$ when A causes B, we can hope that the conclusions we will reach using this basic idea will have some degree of generality across the various, current probabilistic models of causation.

The second model we will investigate, and that we recently proposed, uses a qualitative representation of uncertainty, based on nonmonotonic consequence relations. As we will see, sufficient conditions for the transitivity of causation laid bare by these two models are different.[1] The quantitative model predicts that causation is transitive as soon as the causal chain is Markovian; and the qualitative model predicts that causation is transitive as soon as the first event in the chain is a salient, normal cause of the middle event. The question then arises of whether the two transitivity conditions are formally related at all; and if so, whether their formal links reflect formal relationships between the two underlying definitions of causation. To answer these questions, we need to:

1. Adapt the Markov condition to a qualitative setting;
   (a) check whether it is a sufficient condition for causal transitivity, in the qualitative sense;
   (b) check whether it is distinct from, stronger than, or weaker than the Saliency condition;
2. Adapt the Saliency condition to a quantitative setting;
   (a) check whether it is a sufficient condition for causal transitivity, in the quantitative sense;
   (b) check whether it is distinct from, stronger than, or weaker than the Markov condition;
3. Compare the notions of causality captured by the qualitative and quantitative conditions; in particular, we will:
   (a) Translate the qualitative definition into a probabilistic setting, and check whether it is stronger or weaker than the standard definition;
   (b) investigate the transitivity conditions of this translated definition.

We eventually point out a gap between the concepts of causation captured by each framework, which may explain the disagreement between them despite the fact that the qualitative framework can be viewed as a mathematical limit of the quantitative one, in terms of extreme, non-standard probabilities [1, 43, 38]. We speculate that the qualitative framework is apt to describe everyday causal thinking by laypersons, and that laypersons are more likely to test for Saliency than to test for the Markov condition when they assess the transitivity of causal arguments. We conclude the article by experimental evidence supporting this conjecture: In three behavioral experiments, we observe that judgments of causal transitivity are predicted by Saliency rather than by the Markov condition.[2]

2 The Quantitative Markov Condition

In this paper, $A, B, C, \dots$ denote events, understood as subsets of situations, exactly one of which occurs. Then, $A \cup B$ denotes the union of such sets

[1] Let us note that these quantitative and qualitative frameworks are both eligible for predicting judgments of causality from reported sequences of events, as discussed in [5]. In particular, their implementation does not raise significant problems from a computational point of view, and is not discussed any further in this article.

[2] This article is an extended version of the preliminary conference paper [12].

Table 1 A Counterexample to the transitivity of probabilistic causation when $\Pr(C|AB) = \Pr(C|B)$.

	$\bar{A}$		A	
	$\bar{B}$	B	$\bar{B}$	B
$\bar{C}$	2	4	3	4
C	3	4	0	4

expressing the disjunction of events, and their conjunction, usually denoted as a set-intersection $A \cap B$ is shortened into AB for simplicity. The standard definition of causation, from a probabilistic perspective, is that the presence of the cause increases the probability of the effect. That is, 'A causes B' (where A, B, are events, not variables) if and only if $\Pr(B|A) > \Pr(B|\bar{A})$. Note that this expression precludes the cases where either $\Pr(A)$ or $\Pr(B)$ is equal to 0 or 1, and that it can also be written $\Pr(B|A) > \Pr(B)$. Note also that, although it can be written in a symmetrical way expressing positive correlation, $\Pr(AB) > \Pr(A)\Pr(B)$, we will maintain the tradition of calling A and B the the 'cause' and the 'effect,' respectively (provided that the cause temporally precedes the effect). Note that Good's definition implies that $\Pr(A) > 0, \Pr(B) > 0$, i.e. the probability of events such that $\Pr(B|A) > \Pr(B)$ are necessarily positive.

Causation in the crude probabilistic sense is not necessarily transitive. The fact that $\Pr(B|A) > \Pr(B|\bar{A})$, i.e., '$A$ causes B', together with the fact that $\Pr(C|B) > \Pr(C|\bar{B})$, i.e., '$B$ causes C', do not always imply that $\Pr(C|A) > \Pr(C|\bar{A})$, i.e., '$A$ causes C.' A detailed counterexample is presented later on (see Table 1; all tables display integer numbers, corresponding to the number of observations in each cell).

Probabilistic causation, however, is transitive as soon as the two following conditions are jointly satisfied:

$$\Pr(C|AB) = \Pr(C|B) \tag{1}$$
$$\Pr(C|A\bar{B}) = \Pr(C|\bar{B}) \tag{2}$$

Conditions (1) and (2) are the two parts of the Markov condition on events. They express that C is independent of A in the context of B, and in the context of $\bar{B}$, respectively. In the following, we will speak of the 'positive' and 'negative' parts of the Markov condition on events, respectively. Note that this fine-grained decomposition of the Markov condition is possible because our analysis focuses on the relationships between specific values of the variables, rather than on the overall relation between the variables. This fine-grained level of analysis is essential to the results we will provide in this paper.

Remark 1 Both parts are covered by the symmetric expression of the Markov condition $\Pr(abc)\Pr(b) = \Pr(ab)\Pr(bc)$, for Boolean variables a, b, c having values A, $\bar{A}$, etc. This symmetric expression allows for $\Pr(b) = 0$, which is not the case for the conditional, asymmetric formulation, using Kolmogorov conditional probabilities.

Table 2 Postulates and characteristic properties of System P.

	PREMISE 1	**PREMISE 2**	**CONCLUSION**
RIGHT WEAKENING	$E \mid\!\sim F$	$F \subseteq G$	$E \mid\!\sim G$
AND	$E \mid\!\sim F$	$E \mid\!\sim G$	$E \mid\!\sim FG$
OR	$E \mid\!\sim G$	$F \mid\!\sim G$	$E \cup F \mid\!\sim G$
CAUTIOUS MONOTONY	$E \mid\!\sim F$	$E \mid\!\sim G$	$EF \mid\!\sim G$
CUT	$E \mid\!\sim F$	$EF \mid\!\sim G$	$E \mid\!\sim G$

Proposition 1 *If the two parts of the Markov condition on events are satisfied, then it follows from* $\Pr(B|A) > \Pr(B|\bar{A})$ *and* $\Pr(C|B) > \Pr(C|\bar{B})$ *that* $\Pr(C|A) > \Pr(C|\bar{A})$.

Proof First, $\Pr(C|A) = \Pr(C|AB)\Pr(B|A) + \Pr(C|A\bar{B})(1 - \Pr(B|A))$;
Likewise, $\Pr(C|\bar{A}) = \Pr(C|\bar{A}B)\Pr(B|\bar{A}) + \Pr(C|\bar{A}\bar{B})(1 - \Pr(B|\bar{A}))$.
From (1), $\Pr(C|AB) = \Pr(C|\bar{A}B) = \Pr(C|B)$;
From (2), $\Pr(C|A\bar{B}) = \Pr(C|\bar{A}\bar{B}) = \Pr(C|\bar{B})$.
Hence: $\Pr(C|A) - \Pr(C|\bar{A}) = \Pr(C|B)\Pr(B|A) + \Pr(C|\bar{B})\Pr(B|A)$
- $\big(\Pr(C|B)\Pr(B|\bar{A}) + \Pr(C|\bar{B})\Pr(B|\bar{A})\big)$
$= (\Pr(C|B) - \Pr(C|\bar{B}))(\Pr(B|A) - \Pr(B|\bar{A}))$.
This quantity is strictly positive from the definitions of 'A causes B' and 'B causes C.' □

Remark 2 This proof is inspired from Eells and Sober [25], who studied the transitivity of their own version of probabilistic causation.

It will be relevant later on to note that (1) alone is not sufficient to ensure transitivity. Table 1 presents an example where A causes B and B causes C, in the probabilistic sense [3] $\Pr(B|A) = 8/11$ is greater than $\Pr(B) = 16/24$, and $\Pr(C|B) = 8/16$ is greater than $\Pr(C) = 11/24$. Furthermore, the positive part (1) of the Markov condition on events is satisfied since $\Pr(C|BA) = 4/8 = \Pr(C|B\bar{A})$. However, A cannot be said to cause C, since $\Pr(C|A) = 4/11$ is *less* than $\Pr(C) = 11/24$. As it can be expected, the negative part (2) of the Markov condition on events is not satisfied, since $\Pr(C|\bar{B}A) = 0$ is different from $\Pr(C|\bar{B}\bar{A}) = 3/5$.

3 A Qualitative Saliency Condition

Because probabilistic information is not always available for causal inference, Bonnefon and colleagues [5,10,11] offered a qualitative counterpart to probabilistic conceptions of causality. This framework takes advantage of so-called *nonmonotonic consequence relations,* which make it possible to express that the occurrence of B is generally, normally a consequence of the occurrence of

[3] Note that in Table 2, probabilities $\Pr(A), \Pr(B), \Pr(C)$ are positive; it will be so in all subsequent examples of the same vein.

A, but that exceptional situations may arise—without the need for specifying how frequent these exceptions are.

Formally, the relation 'If A is true then normally B is true' is written $A \mid\!\sim B$, read A 'snake' B. The snake operator follows the requirements of System P for nonmonotonic reasoning [34]. That is, $\mid\!\sim$ is reflexive and satisfies the postulates and properties summarized in Table 2. Empirical studies repeatedly demonstrated that the postulates of System P provide an adequate description of the way human reasoners handle exception-laden rules [6, 18, 45].

Based on this representation of background knowledge, perceived causation is defined in the following terms: an event A is perceived to cause another event B in a context K if B was false before A took place (which was normal in context K), and became true afterwards (B is normal as well in context KA):

Definition 1 An agent "perceives that A_{t1} caused B_{t_2} in context K" (denoted $A_{t1} \rhd_K B_{t2}$), if and only if:

- the agent knows of a sequence $\neg B_{t1}$, A_{t1}, B_{t2} ; where $t_2 > t_1$; t_1 and t_2 are the first time points when A and B, respectively, are known to be true by the agent ; and K (the *context*) is the conjunction of all other facts known by the agent at time t_1.
- the agent possesses the pieces of default knowledge that $K \mid\!\sim \neg B$ and $AK \mid\!\sim B$.

In the remainder of this article, $A_{t1} \rhd_K B_{t2}$ will be abridged as $A \rhd_K B$, only when there is no risk of ambiguity about the context or the time points. Likewise, $AK \mid\!\sim B$ will be shortened as $A \mid\!\sim_K B$ when appropriate.

This definition of causality is closely related to an asymmetric notion of qualitative independence introduced by Dubois *et al.* [21] in the setting of possibility theory. It has a number of formal properties that are explored in [11]. Especially relevant to our current purpose is the fact that perceived causality defined in this way is not generally transitive: it is not always possible to infer $A_{t1} \rhd_K C_{t3}$ from $A_{t1} \rhd_K B_{t2}$ and $B_{t2} \rhd_K C_{t3}$ (where $t_3 > t_2 > t_1$, and K is taken at t_1). However :

Proposition 2 *If $A_{t1} \rhd_K B_{t2}$, $B_{t2} \rhd_K C_{t3}$, and , then $A_{t1} \rhd_K C_{t3}$.*

Proof

- $B \mid\!\sim_K C$ holds from the definition of $B_{t2} \rhd_K C_{t3}$. Together with $B \mid\!\sim_K A$, it gives (by Cautious Monotony) $AB \mid\!\sim_K C$. This together with $A \mid\!\sim_K B$ (from the definition of $A_{t1} \rhd_K B_{t2}$) gives (by Cut) $A \mid\!\sim_K C$.
- $K \mid\!\sim \neg C$ follows directly from the definition of $B_{t2} \rhd_K C_{t3}$.
- Trivially, $t_3 > t_1$; and t_1, t_3 are the first time points where A and C, respectively, were known to be true.

□

We will henceforth call $B \mid\!\sim_K A$ the *Saliency* condition.

Indeed, this condition means that observing B in context K normally leads one to expect that A has occurred. In other terms, A is a very salient cause of B, so salient that it is normally the explanation that one will imagine, by default, to explain the occurrence of B.

4 A Qualitative Markov Condition

In this section, we focus on the qualitative framework by (a) exploring the links between the qualitative Saliency condition and a qualitative rendition of the Markov condition; and (b) exploring the links between this qualitative rendition and the transitivity of qualitative causation. One qualitative rendition of (1) and (2) is:

$$AB \mid\!\sim_K C \text{ if and only if } B \mid\!\sim_K C \tag{3}$$

$$A\bar{B} \mid\!\sim_K C \text{ if and only if } \bar{B} \mid\!\sim_k C \tag{4}$$

This definition encodes the idea that knowing or ignoring the presence of A is irrelevant to the agent's belief in C in the presence of B, and in the presence of its contrary $\bar{B}$. The qualitative Saliency condition implies half of the qualitative Markov condition. More precisely, the Saliency condition (2) implies the positive part (3) but not the negative part (4).

Proposition 3 *If $B \mid\!\sim_K A$, then $AB \mid\!\sim_K C$ if and only if $B \mid\!\sim_K C$.*

Proof The proof simply uses CUT on $B \mid\!\sim_K A$ and $AB \mid\!\sim_K C$ for $\Rightarrow$, and CAUTIOUS MONOTONY on $B \mid\!\sim_K A$ and $B \mid\!\sim_K C$ for $\Leftarrow$. □

However, the saliency condition $B \mid\!\sim_K A$ does not imply the negative part of the Markov condition. It is possible to have $A \rhd_K B, B \rhd_K C$ and $B \mid\!\sim_K A$ without having $A\bar{B} \mid\!\sim_K C$ equivalent to $\bar{B} \mid\!\sim_K C$. For example, it may be that both $B \mid\!\sim_K A$ and $\bar{B} \mid\!\sim_K \bar{A}$ hold; but accepting $\bar{B} \mid\!\sim_K \bar{A}$ stops the derivation of $A\bar{B} \mid\!\sim_K C$ from $\bar{B} \mid\!\sim_K C$ (cautious monotony does not apply). In a qualitative setting, the Saliency condition does not necessarily imply the negative part of the Markov condition on events.

Remark 3 Note that (4) could easily be derived from $\bar{B} \mid\!\sim_K A$. However, if both $B \mid\!\sim_K A$ and $\bar{B} \mid\!\sim_K A$ hold, then (by OR) $(KB) \cup (K\bar{B}) \mid\!\sim A$, and ultimately $K \mid\!\sim A$. As shown in [11], the fact that $K \mid\!\sim A$ precludes that $A \rhd_K B$ for any B: In the qualitative model, normal events cannot be perceived as causes of abnormal events. Thus, adding the condition $\bar{B} \mid\!\sim_K A$ to $B \mid\!\sim_K A$ would allow to derive (3) and (4) instead of just (3), but would at the same time defeat the whole point of reasoning about causes.

Remarkably, and contrary to the quantitative case, the positive part of the qualitative Markov condition is sufficient to give transitivity.

Proposition 4 *If $KA \mid\!\sim B$ and $KB \mid\!\sim C$, then $KA \mid\!\sim C$, provided that ($KAB \mid\!\sim C$ if and only if $KB \mid\!\sim C$).*

Proof From $KB \mid\!\sim C$ and ($KAB \mid\!\sim C$ if and only if $KB \mid\!\sim C$), we get $KAB \mid\!\sim C$. From this relation and $KA \mid\!\sim B$ we arrive at $KA \mid\!\sim C$ by applying CUT. □

Corollary 1 *If $A_{t1} \rhd_K B_{t2}$, $B_{t2} \rhd_K C_{t3}$, and ($KAB \mid\!\sim C$ if and only if $KB \mid\!\sim C$), then $A_{t1} \rhd_K C_{t3}$.*

Proof – $KA \mid\!\sim C$ is given by Proposition 4.
– $K \mid\!\sim \neg C$ follows directly from $B_{t2} \rhd_K C_{t3}$.
– Trivially, $t_3 > t_1$; and t_1, t_3 are the first time points where A and C, respectively, were known to be true.

□

Nevertheless, it can be checked that the positive part of the qualitative Markov condition does not imply the Saliency condition. To show this, we use a model of System P where the former holds but not the latter. Consider a qualitative possibility distribution induced by the well-ordered partition [8]: $(\bar{A}\bar{B}\bar{C}, C, (A \cup B) \setminus C)$. Namely $\forall s \notin A \cup B \cup C, \pi(s) = 1$; $\forall s \in C, \pi(s) = \alpha$; $\forall s \in (A \cup B) \setminus C, \pi(s) = \beta$, with $1 > \alpha > \beta$. Translating $A \mid\!\sim B$ as $\Pi(AB) > \Pi(A\bar{B})$, where $\Pi(A) = \max_{s \in A} \pi(s)$ it is easy to check that:

– $\mid\!\sim \bar{B}$ and $\mid\!\sim \bar{C}$ since $\Pi(B) < 1$ and $\Pi(C) < 1$
– $A \mid\!\sim B, B \mid\!\sim C$ since $\Pi(AB) = \Pi(BC) = \alpha$, assuming $ABC \neq \emptyset$; assuming $A\bar{B}C = \emptyset$, $\Pi(A\bar{B}) = \Pi(B\bar{C}) = \beta$
– $AB \mid\!\sim C$ since $\Pi(ABC) = \alpha > \Pi(AB\bar{C}) = \beta$
– Assuming that $\bar{A}BC \neq \emptyset$, $B \mid\!\sim A$ does not hold since $\Pi(\bar{A}B) = \alpha = \Pi(AB)$.

It may seem surprising that, in the qualitative setting, the positive part of the Markov condition is sufficient for transitivity, although it is weaker than the Saliency condition. We can explain this by exploiting the natural connection between conditional assertions of the form $A \mid\!\sim B$ and conditional probabilities $\Pr(B \mid A)$. Indeed, $A \mid\!\sim B$ can be formally interpreted as a pair of events of the form $(AB, A\bar{B})$, and the logic of conditional assertions has a three-valued semantics based on this representation [24]. $\Pr(B \mid A)$ is entirely determined by $\Pr(AB)$ and $\Pr(A\bar{B})$, and there is a closely related semantics of conditional assertions, whereby the statement of $A \mid\!\sim B$ comes down to an infinitesimal probability statement:

$$A \mid\!\sim B \iff \Pr(B \mid A) > 1 - O(\epsilon),$$

where ϵ is a positive number arbitrarily close to 0, and $O(\epsilon)$ stands for any function that vanishes with ϵ. Inferring $A \mid\!\sim B$ from a set of conditionals comes down to proving that the corresponding infinitesimal probability statements imply that $\Pr(B \mid A) > 1 - O(\epsilon)$, i.e. that when the corresponding probabilities get arbitrarily close to 1, so does the probability $\Pr(B \mid A)$. The properties of reasoning with such extreme probability statements were studied in [1,43,38] and turned out to be the properties of nonmonotonic reasoning laid bare in System P. In fact, the calculus of infinitesimal conditional probabilities of this form is equivalent to System P [38,7]. And degraded forms of inference rules of System P hold with standard conditional probabilities [23,26,37].

Now, it is clear that $B \rhd C$ reads $\Pr(C \mid B) > 1 - O(\epsilon)$ and $\Pr(C) < O(\epsilon)$. Hence:

- From $A \rhd B$ and $B \rhd C$ we get $\Pr(AB) > (1 - O(\epsilon)) \Pr(A)$ and $\Pr(BC) > (1 - O(\epsilon)) \Pr(B)$; it yields:
- $\Pr(AB) \Pr(BC) > (1 - O(\epsilon)) \Pr(A) \Pr(B)$.
- This result, along with the positive Markov condition, yields: $\Pr(ABC) \Pr(B) > (1 - O(\epsilon)) \Pr(A) \Pr(B)$.
- Simplifying by $\Pr(B)$: $\Pr(AC) \geq \Pr(ABC) > (1 - O(\epsilon)) \Pr(A)$. Hence $A \rhd C$.

Remark 4 The above setting for plausible inference has been challenged on the ground that it conflicts with the idea of acceptance viewed as high probability: if in a given context $\Pr(A)$ and $\Pr(B)$ are close to 1, then it does not imply that $\Pr(AB)$ is close to 1 at all. It questions the deductive closure of accepted beliefs [47,37]. This is the basis of the lottery paradox of Kyburg [36] (buying a lottery ticket that loses is very likely, but one is sure that there is at least one winner). It explains why Adams, and more recently Pearl, Lehmann and colleagues resort to infinitesimal probabilities to justify the plausible inference rules of system P ($\Pr(AB) \geq 1 - 2O(\epsilon)$ if $\Pr(A) \geq 1 - O(\epsilon)$ and $\Pr(B) \geq 1 - O(\epsilon)$. However, Benferhat et al.[8], have laid bare a non-infinitesimal probabilistic semantics of system P, interpreting conditionals as constraints on so-called big-stepped probabilities (which form super-increasing sequences of numbers $p_1 > p_2 ... > p_n \in [0, 1]$ such that $\forall i = 1, \ldots, n, p_i > \sum_{j=i+1,\ldots,n} p_j$). Noticeably, the set of accepted propositions in the sense of a big-stepped probability is deductively closed in any context. As explained at length in [20], such probabilities are immune to the lottery paradox. Moreover, it is stressed that situations were non-monotonic reasoning makes sense (jumping to conclusions and reasoning with them in standard logic) is when some situations are prominently more likely than others (as expressed by big-stepped probabilities), not when mere chance and pure randomness prevail (as in the lottery paradox example, where all tickets sold are equally likely).

5 A Quantitative Saliency Condition

In this section, we focus on the quantitative case by (a) exploring the links between the quantitative Markov condition and a quantitative rendition of the Saliency condition; and (b) exploring the links between this quantitative rendition and the transitivity of quantitative causation. In agreement with the broadest probabilistic understanding of $\mid\!\sim$, we use the following quantitative rendition of the Saliency condition (2):

$$1 \geq \Pr(A \mid B) \geq k > 0.5. \quad (5)$$

Table 3 A Counterexample to $\Pr(C \mid A\bar{B}) = \Pr(C \mid \bar{B})$ where $\Pr(A \mid B) = 1$.

	$\bar{A}$		A	
	$\bar{B}$	B	$\bar{B}$	B
$\bar{C}$	4	0	1	1
C	2	0	1	1

Table 4 A Counterexample to transitivity where $\Pr(A \mid B) = 1$.

	$\bar{A}$		A	
	$\bar{B}$	B	$\bar{B}$	B
$\bar{C}$	8	0	4	0
C	6	0	1	1

for some appropriate value of the threshold k. We first consider the limit case where $k = 1$. As we will see, this limit case is already informative enough; we will only briefly consider the general case where $0.5 < k < 1$.

The fact that $k = 1$ in (5), i.e., that $\Pr(A \mid B) = 1$, implies the positive part of the quantitative Markov condition. More precisely, $\Pr(A \mid B) = 1$ implies (1) but not (2).

Proposition 5 *If* $\Pr(A \mid B) = 1$ *then* $\Pr(C|AB) = \Pr(C|B)$.

Proof $\Pr(A \mid B) = 1$ implies $\Pr(\bar{A}B) = 0$, and therefore both $\Pr(AB) = \Pr(B)$ and $\Pr(ABC) = \Pr(BC)$. Therefore, $\Pr(ABC)\Pr(B) = \Pr(AB)\Pr(BC)$, which gives (1). □

Just as in the qualitative case, the positive part of the Markov condition does not imply the qualitative Saliency condition, (1) does not imply $\Pr(A \mid B) = 1$. See again Table 1 for a counterexample. $\Pr(A \mid B) = 1$ is thus a stronger condition than the positive part (1) of the Markov condition on events. However, no such relation exists between $\Pr(A \mid B) = 1$ and the negative part of the Markov condition on events: $\Pr(A \mid B) = 1$ does not imply (2). Table 3 displays a simple counterexample where $\Pr(A \mid B) = 1$ but $\Pr(C \mid A\bar{B}) = 1/2$ is different from $\Pr(C \mid \bar{B}) = 3/8$.

Neither does $\Pr(A \mid B) = 1$ imply transitivity. Table 4 displays an example where A causes B and B causes C, in the probabilistic sense: $\Pr(B \mid A) = 1/6$ is greater than $\Pr(B) = 1/20$, and $\Pr(C \mid B) = 1$ is greater than $\Pr(C) = 8/20$. Furthermore, $\Pr(A \mid B) = 1$. But A cannot be said to cause C, as $\Pr(C \mid A) = 2/6$ is *lower* than $\Pr(C) = 8/20$. Unsurprisingly, the negative Markov condition is not satisfied, since $\Pr(C \mid \bar{B}A) = 1/5$ is different from $\Pr(C \mid \bar{B}\bar{A}) = 6/14$.

We have considered so far the limit case where $KB \mathrel{|\!\!\sim} A$ is translated as $\Pr(A \mid B) = 1$, which is the strongest possible rendering of the Saliency condition. Even in this limit case, causation is not necessarily transitive, since $\Pr(C \mid A)$ is not necessarily greater than $\Pr(C)$ when $\Pr(C \mid B) > \Pr(C)$ and

$\Pr(B \mid A) > \Pr(B)$. Now, it can be shown that the optimal lower bound of $\Pr(C \mid A)$ is an *increasing function* of the threshold k with $\Pr(A \mid B) \geq k$.

Indeed, as shown in [22], the optimal lower bound of $\Pr(C \mid A)$ can be written as:

$$\Pr(C \mid A) \geq \Pr(B \mid A) \cdot \max\left(0, 1 - \frac{1 - \Pr(C \mid B)}{\Pr(A \mid B)}\right) \tag{6}$$

It is easily checked from (6) that the optimal lower bound of $\Pr(C \mid A)$ increases with $\Pr(A \mid B)$, and is thus an increasing function of threshold k. Therefore, all other probabilities being equal, if it is not guaranteed that $\Pr(C \mid A) > \Pr(C)$ when $\Pr(A \mid B) = 1$, then it cannot be guaranteed either that $\Pr(C \mid A) > \Pr(C)$ for any value of $k > 0.5$. However, if saliency holds as $\Pr(A \mid B) = 1$, a degraded form of transitivity of conditional probability holds as $\Pr(C \mid A) \geq \Pr(B \mid A) \Pr(C \mid B)$.

6 Two models of causation or two notions of causation?

In summary, the following results have been obtained. In the qualitative setting, using the nonmonotonic Definition 1 of causation, the positive Markov condition on events is sufficient for transitivity to hold: there is no need for the negative Markov condition to hold. Because the Saliency condition implies the positive part of the Markov condition (albeit not its negative part), it is itself a sufficient condition for transitivity.

In the quantitative setting, using the probabilistic definition of causation as $\Pr(B \mid A) > \Pr(B)$, the positive Markov condition on events alone (without the negative part) is not a sufficient condition for transitivity. In this setting, the Saliency condition implies the positive part of the Markov condition, but not its negative part; and it is not itself a sufficient condition for transitivity. In this section, we consider whether or not the pattern of results obtained for the transitivity conditions is a consequence of formal relations between the underlying definitions of causation in the two models.

6.1 Two views of causation

The connection between nonmonotonic inference and conditional probability, as recalled above, clarifies the difference between probabilistic causality developed after [28], and nonmonotonic causality. The nonmonotonic definition, couched in probabilistic terms, comes down to the following requirements: A causes B if and only if B is little probable per se and very probable in the context where A is true. The weakest quantitative rendition of this definition is:

$$\Pr(B) < \Pr(\bar{B}) \text{ and } \Pr(B \mid A) > \Pr(\bar{B} \mid A). \tag{7}$$

Remark 5 Note that this definition of causation is genuinely asymmetric, contrary to the standard probabilistic definition that amounts to a positive correlation between events.

This weakest rendition already implies $\Pr(B \mid A) > \Pr(B)$ since (7) implies $\Pr(B \mid A) > 0.5 > \Pr(B)$. Hence, the qualitative approach leads to an asymmetric definition of causality that, even in its weakest quantitative counterpart, is clearly stronger than Good's standard probabilistic definition. Causes in the qualitative sense are always causes in the quantitative sense, but there are probabilistic causes that the qualitative setting does not recognize as causes.

There are two classes of such situations. First, situations where, although the presence of the cause does increase the probability of the effect, this latter probability stays lower than .5:

$$\Pr(B) < \Pr(B \mid A) < \frac{1}{2}\,. \tag{8}$$

In these situations, 'A causes B' in the probabilistic but not in the qualitative sense, since the normal course of things is to observe $\bar{B}$ in the context of A.

Remark 6 From the perspective of scientific discovery, it certainly makes sense to talk about the causal role of A in situation (8). And indeed, the probabilistic definition of causality was proposed with a view to capture the nature of scientific explanation in experimental fields (see [48] for a retrospective collection of essays). It may be more debatable whether (8) corresponds to what lay persons would declare as expressing causality. The close inspection of available experimental data [13,39] shows that perceptions of causality in this situation are at best moderate. Note, however, that $\Pr(B)$ and $\Pr(B \mid A)$ always have the same order of magnitude in these experiments.

Although (8) is not a situation recognized as causation by the qualitative model, it may qualify as a situation of 'facilitation.' In [10,11] the facilitation relation is defined as one where B is abnormal, but becomes neither abnormal nor normal when A occurs.

There is another class of situations of probabilistic causation that does not match the requirements for qualitative causation. This class of situations corresponds to cases where the cause increases the probability of an effect that was already highly probable:

$$\Pr(B \mid A) > \Pr(B) > \frac{1}{2}\,. \tag{9}$$

There is no way to express this reinforcement effect in the qualitative setting, due to a lack of expressive power. For A to cause B in the qualitative sense, it is necessary that the perception of B changes from that of an abnormal event, to that of a normal event when A occurs. A cannot be said to be a qualitative cause of B when B already belongs to the normal course of the world even in the absence of A.

Remark 7 A limit situation occurs when B is neither intrinsically normal nor abnormal, but becomes normal in the presence of A. In other terms, one is totally ignorant about whether B is true or not, but starts believing that B is the case when learning that A has occurred. In a qualitative setting, this expresses as [21]:

$$K \not\mid\!\sim B \text{ and } K \not\mid\!\sim \bar{B} \text{ and } KA \mid\!\sim B$$

The probabilistic counterpart to this limit situation may be considered to be that wherein $\Pr(B) = \Pr(\bar{B})$. This representation is debatable, though, as it presupposes that ignorance about B is identified with knowledge about the randomness of B. This situation is not considered to be one of causation in [11]. Rather, following [21], it is considered one of 'justification.' That is, one where, rather than seeing A as the cause of the occurrence of B, agents may consider A to be the justification in their belief that B was going to happen, given that they had not enough information to conclude anything before hearing that A occurred. Here, a justification is understood as an item of information inducing an epistemic change from ignorance to belief, while a cause is viewed as inducing belief reversal.

6.2 Consequences for transitivity

There is a difference of nature between the standard probabilistic definition of causation and the definition inspired by the qualitative model (notwithstanding the fact that the first is based on a symmetric property, whilst the second is asymmetric). Quantitatively, causality is understood in terms of a positive influence of one event on another one, regardless of the prior probability or this last event. Qualitatively, causality is restricted to situations of tendency reversal, whereby an abnormal state of affairs becomes normal upon the occurrence of some event.

Because the quantitative rendition (7) of the qualitative definition of causality is stronger than the standard probabilistic model, the question arises whether Saliency would ensure the transitivity of this stronger notion of causation. The following counterexample shows that this is not the case. Suppose A causes B and B causes C in the sense of (7), and that Saliency strictly holds, i.e., $\Pr(A \mid B) = 1$. Suppose observations as in Table 5. Note that $\Pr(B) < \frac{1}{2}$; $\Pr(C) < \frac{1}{2}$. Moreover $\Pr(B \mid A) = \frac{18}{35} > \frac{1}{2}$; $\Pr(C \mid B) = \frac{10}{18} > \frac{1}{2}$; However, $\Pr(C \mid A) = \frac{14}{35} < \frac{1}{2}$.

One may think of strengthening again the definition of causation by increasing the threshold k and change (7) into what could be called k-causation:

$$\Pr(B) \le 1 - k < 0.5 \text{ and } \Pr(B \mid A) \ge k. \qquad (10)$$

However we can show by means of the generic counterexample in Table 6 that the Saliency condition is not sufficient to ensure transitivity of k-causation, however close is k to 1.

Table 5 Counterexample to transitivity under saliency and strengthened causation.

	$\bar{A}$		A	
	$\bar{B}$	B	$\bar{B}$	B
$\bar{C}$	55	0	13	8
C	10	0	4	10

Table 6 Counterexample to the transitivity of k-causation under saliency.

	$\bar{A}$		A	
	$\bar{B}$	B	$\bar{B}$	B
$\bar{C}$	x	0	w	$2w$
C	y	0	0	$\frac{5wk}{2(1-k)}$

Indeed, the conditions $\Pr(B) \leq 1-k$; $\Pr(C) \leq 1-k$ can be achieved by choosing a sufficiently high value of $x = \Pr(\bar{A}\bar{B}\bar{C})$, tuning $w = \Pr(A\bar{B}\bar{C})$ accordingly for a fixed value of k. Furthermore, when $0.5 < k < 1$, it is clear that $\Pr(B \mid A) = \frac{4+k}{6-k} > k$; $\Pr(C \mid B) = \frac{5k}{4+k} > k$; but $\Pr(C \mid A) = \frac{5k}{6-k} < k$. What it means is that when k increases and the notion of k-causation becomes more demanding, there always exists a small range of situations that do not sanction transitivity. This range shrinks when k increases—and it vanishes when $k = 1$.

Note that replacing Saliency by the two Markov conditions will not do any better at ensuring the transitivity of k-causation. Indeed, suppose $\Pr(C|B) = \Pr(B|A) = k > 0.5$. Then, $\Pr(C|A) = \Pr(C|B)\Pr(B|A) + \Pr(C|\bar{B})(1 - \Pr(B|A))$, using both Markov condition. Therefore: $\Pr(C|A) = k^2 + \Pr(C|\bar{B})(1-k)$. Letting $\Pr(C|\bar{B}) < k$, we arrive at $\Pr(C|A) < k$.

Although, stricto sensu, the transitivity of k-causation can fail under Markov conditions and Saliency, the condition $\Pr(C|A) \geq \Pr(C \mid B)\Pr(B \mid A) > k^2$ is always derivable from $\Pr(B|A) > k$ and $\Pr(C|B) > k$, under the saliency condition $\Pr(A|B) = 1$, from equation (6), with no other assumption. If $\Pr(B|A)$ and $\Pr(C|B)$ are significantly greater than k, then it is possible that $\Pr(C|A) \geq \Pr(C \mid B)\Pr(B \mid A) > k$. Furthermore, from a psychological perspective, it is conceivable that, for some high values of k^2, some individuals may perceive the causal chain as transitive in situations where $k^2 < \Pr(C|A) < k$.

Interestingly, the above discussion suggests a distinction between probabilistic causation and commonsense causality, where the conditions for transitivity differ. We just showed that the qualitative representation of commonsense causality may have a probabilistic rendition that enabled us to compare it to the more traditional notion of probabilistic causation. However, there is still a major difference between the two probabilistic frameworks. Probabilistic causation always assumes probabilistic data is available and that correlation, understood or not as causation, can be read from this data. In the case of commonsense causality, such an assumption cannot be made. Asserting causality and its transitivity are expressed by constraints valid for a family of probability functions. So the probabilistic framework suggested here for commonsense causality belongs to the framework of imprecise probability representations [49], rather than orthodox probability settings.

7 Behavioral Results

A case can be made for the idea that the two models we have compared in this article do not address the same phenomenon. The probabilistic definition of causation introduced after Good, and popularized by epistemologists, aims at detecting a positive influence in data observed from natural phenomena. Such causal relations can be read from a careful scrutiny of contingency tables. In contrast, the qualitative model captures a more mundane, commonsense, everyday variant of causation—whereby the change of a state of affairs is understood as caused by the prior occurrence of some abnormal event. This form of causal thinking often comes down to explaining an abnormal fact by some unexpected circumstances. These two kinds of causal thinking are quite distinct, as shown in this article, and perhaps it is no surprise that their transitivity conditions differ.

This observation raises the (widely unaddressed) empirical question of the transitivity of everyday causal thinking. Is everyday causal reasoning, which we take to be qualitative in nature, unconditionally transitive ? And if not, is it sensitive to Saliency, or to Markovian considerations ? The very few behavioral studies that addressed the issue of causal transitivity invariably found that human participants appeared to endow causal relations with full transitivity [2,3,27]. These results, however, must be taken with caution, because they used a very limited sample of causal relations. In particular, no effort was ever made in any of these studies to use causal relations that may defeat transitivity. Overall, we may not want to draw any strong conclusions based on a limited, homogeneous sample of causal relations.

Our first empirical contribution will be to show that transitive causal conclusions can easily be defeated if appropriate causal relations are proposed to human reasoners. Our second question is whether everyday inferences from causal chains are sensitive to Saliency, or to Markovian considerations. If, as we suspect, laypersons adopt a qualitative approach to everyday causal reasoning, then they may perform two different tests before they accept a transitive conclusion. They may consider whether the chain is Markovian (and endorse the transitive conclusion if it is), or they may test whether the first term of the chain is the most salient cause of the middle term (and endorse the transitive conclusion if it is).

Because the qualitative (positive) Markov condition is a weaker condition for transitivity than the Saliency condition, the fact that Saliency holds or not should not affect transitive inferences as long as the chain is Markovian. However, we speculate that the Saliency test is much easier and more natural, cognitively speaking, than the Markov test. In other terms, reasoners may be likely to *substitute* the Markov condition with the easier and more natural Saliency condition. As a consequence, we predict that transitive inferences may be rejected even when the chain is Markovian, as soon as the chain does not meet the Saliency condition.

7.1 Materials for the experiments

The goal of our experiments is to show that human reasoners rely on the Saliency condition, rather than the Markov condition, when they assess the transitivity of a causal chain. To that end, we will construct causal chains that do or do not satisfy the Saliency condition, but that always satisfy the Markov condition. We will then measure whether human reasoners reject the transitivity of causal chains that do not satisfy the Saliency condition, even though they satisfy the Markov condition.

As a preliminary to our experiments, we constructed six causal chains, following a two-step process. The first step consisted of selecting the 'A causes B' relation of the chain: The constraint was that three causal relations among the six should involve many salient possible causes other than A, and that in the three other relations there should have few or no salient possible causes other than A. We took advantage of prior work in the experimental psychology of reasoning that commonly test various relations for the number of alternative causes that they spontaneously suggest to human reasoners [14, 19, 46]. We selected from [14, 19] three relations with one salient cause, and three relations with many possible causes.

The second step was to construct the 'B causes C' relation of each chain, so that the whole chain would be Markovian. We tentatively selected relations that intuitively appeared to be Markovian. To test this intuition prior to using the chains in our series of experiments, we presented 35 adult participants (4 men; mean age = 22.5, square deviation = (SD) = 8.5) with a Markov-assessment task consisting of 10 problems. Each problem was introduced by a sentence that provided a minimal context. Then, the conditional probabilities $\Pr(C|AB)$ and $\Pr(C|\bar{A}B)$ were compared by means of the following question: What is more likely to be true, 'if A and B then C,' or 'if $\bar{A}$ and B then C'? *(The former is more likely, the latter is more likely, they are equally likely.)* Previous research showed that this conditional formulation captures lay perceptions of conditional probability [42]. Here is an example of a complete problem:

> Carole is coming back from a mountain trek. What is more likely to be true: 'if Carole encounters a bear and runs, then she is out of breath,' or 'if Carole does not encounter a bear and runs, she is out of breath'? *(The former is more likely, the latter is more likely, they are equally likely.)*

Six of the 10 problems featured the preselected causal relations, and four problems featured filler items chosen so that they would likely violate the Markov assumption. (Such filler items are commonly inserted in psychological questionnaires alongside items of interest.) The 10 problems appeared in one random order in half the questionnaire, and in the inverse order in the remaining half. For half the participants, the question was phrased: What is more likely to be true, 'if A and B then C,' or 'if $\bar{A}$ and B then C'? For the other half, the order of the clauses was reversed: What is more likely to be true, 'if $\bar{A}$ and B then C,' or 'if A and B then C'?

Half of the participants responded first to the $\Pr(C|AB)$ question, and the other half responded first to the $\Pr(C|\bar{A}B)$ question.

We considered that a chain was Markovian if the most frequent answer to the corresponding problem was *they are equally likely*. Such was the case for all six target problems (51–91% of *they are equally likely* answers), but for none of the four filler items (2–26% of *they are equally likely* answers). We thus made the decision to use these six causal chains in our series of experiments.

The exact phrasing of the causal chains varied in the three experiments, in order to ensure that our findings were robust over the various forms that causality may take in natural language. In Experiment 1, the causal chains were described as 'C occurred because of B, B occurred because of A.' In Experiment 2, the phrasing was 'A caused B, B caused C.' In Experiment 3, the phrasing was 'A generally causes B, B generally causes C.' Couched in the phrasing of Experiment 1, the three (Markovian) chains with one salient cause were:

FINGER Didier was repairing his washing machine. He put a bandage on his finger because it bled. His finger bled because he cut it.

KETTLE Émilie had put her kettle on the stove. The kettle whistled because the water was boiling. The water was boiling because it had been heated to 100 degrees.

WEAPON Fabrice was suspected of murder. He was convicted because his fingerprints were on the weapon. His fingerprints were on the weapon because he held it without wearing gloves.

Using the same phrasing, the three (Markovian) chains with many possible causes were:

APPLE Alice was asleep under an apple tree. She woke up because an apple fell on her. The apple fell on her because it was ripe.

BEAR Carole was coming back from a mountain trek. She was out of breath because she had run. She had run because she had encountered a bear.

COUGH Benjamin was playing in the living room. He was scared because his father coughed in the next room. His father coughed because his food went down the wrong way.

To prevent the potential objection that lesser transitivity of the *Many possible causes* chains may be attributed to their being Markovian to a lesser degree, we note that the chains meant for the *Many possible causes* condition (that we predict to be intransitive) were judged as Markovian to an overall greater extent than the chains meant for the *One salient cause* condition (APPLE 91%, BEAR 77%, COUGH 51%; vs. FINGER 51%, KETTLE 51%, WEAPON 57%).

7.2 Experiment 1

Method Participants were 50 undergraduate students at the University of Toulouse (11 men; mean age = 21.5, SD = 5.3), who volunteered to complete the questionnaire during class time. None of them had previous training in the philosophy or psychology of causation. It took less than 10 minutes to complete the questionnaire.

Participants solved six transitivity problems. All problems followed the same 3-part format: (a) One sentence setting a minimal context; (b) One sentence stating that C happened because of B; (c) One sentence stating that B happened because of A. Participants rated on a 5-point scale the correctness of the conclusion 'C happened because of A' (*totally incorrect, rather incorrect, could be correct as well as incorrect, rather correct, totally correct*; encoded $-2, -1, 0, +1, +2$, respectively). For example, the BEAR problem read:

> Carole was coming back from a mountain trek. She was out of breath because she had run. She had run because she had encountered a bear. Was Carole out of breath because she had encountered a bear?

In three problems (FINGER, KETTLE, WEAPON), the causal relation between A and B had one salient cause. In the other three problems (APPLE, BEAR, COUGH), it had many possible causes. The six problems were presented in random order in one half of the questionnaires, and in reverse order in the other half.

Results Participants' ratings were encoded from -2 to $+2$, with positive ratings indicating acceptance of transitivity. The left part of Fig. 1 displays the global distribution of the 300 responses to problems with one salient cause vs. many possible causes. We averaged each participant's responses to the three problems with one salient cause (on the one hand), and to the three problems with many possible causes (on the other hand). In the three problems with one salient cause, mean transitivity acceptance was $+1.0$ ($SD = 0.8$), whilst it was only -0.6 ($SD = 0.9$) in the three problems with many possible causes. This difference cannot be attributed to chance, $t(49) = 10.0$, $p < .001$, and it is of huge magnitude as per the standards of experimental psychology ($d = 1.9$). This is in fact the case for all the experimental findings we report in this article. Overall, the transitive conclusion was endorsed (i.e., encoded as $+1$ or $+2$) for 80% of the problems with few alternatives, but only for 24% of the problems with many alternatives .

Mean acceptance ratings for the three problems in the *One salient cause* condition (FINGER, KETTLE, WEAPON) were $+0.8$ ($SD = 1.3$), $+1.1$ ($SD = 1.2$), and $+1.2$ ($SD = 1.0$), respectively. On the contrary, mean acceptance ratings for the three problems in the *Many possible causes* condition (APPLE, BEAR, COUGH) were -0.8 ($SD = 1.1$), -0.6 ($SD = 1.1$), and -0.2 ($SD = 1.2$), respectively. In other terms, participants globally accepted transitivity in the three problems of the *One salient cause* condition, but globally rejected transitivity in the three problems of the *Many possible causes* condition.

7.3 Experiment 2

Method Participants were 51 undergraduate students drawn from the same population as in Experiment 1 (17 men; mean age = 22.4, $SD = 3.1$). Ex-

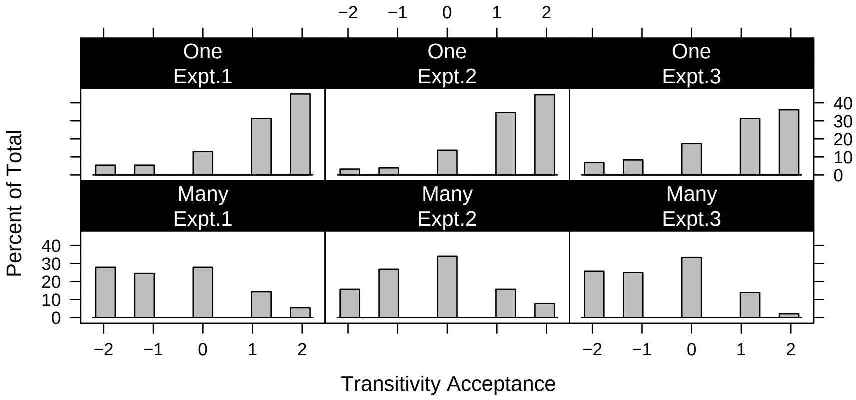

Fig. 1 Distribution of responses to the transitivity question in Experiment 1 (causal explanation), Experiment 2 (particular causal attribution), and Experiment 3 (general causal attribution), for the problems with many possible causes, and for the problems with one salient cause. In all three experiments, transitivity clearly holds for problems with one salient cause, but tends not to hold for problems with many possible causes.

periment 2 was similar to Experiment 1, except that arguments were framed as particular causal attributions rather than causal explanations. The BEAR problem read:

> Carole was coming back from a mountain trek. Encountering a bear had caused her to run. Running had caused her to be out of breath. Did encountering a bear cause Carole to be out of breath?

Results The central part of Fig. 1 displays the global distribution of the 306 responses to problems with one salient cause vs. many possible causes. In the three problems with one salient cause, mean transitivity acceptance was +1.1 (SD = 0.7), whilst it was only −0.3 (SD = 0.9) in the three problems with many possible causes; $t(50) = 10.2$, $p < .001$, $d = 1.7$. Overall, the transitive conclusion was endorsed (i.e., encoded a +1 or a +2) for 76% of the problems with one salient cause, but only for 21% of the problems with many possible causes.

Once more, participants accepted transitivity in the three problems with one salient cause, but rejected transitivity in the three problems with many possible causes. Mean acceptance ratings for the three problems with one salient cause (FINGER, KETTLE, WEAPON) were +1.3 ($SD = 0.9$), +1.2 ($SD = 1.0$), and +0.9 ($SD = 1.0$), respectively. On the contrary, mean acceptance ratings for the three problems with many possible causes (APPLE, BEAR, COUGH) were −0.5 ($SD = 1.1$), −0.2 ($SD = 1.1$), and −0.1 ($SD = 1.2$), respectively.

7.4 Experiment 3

Method Participants were 48 undergraduate students drawn from the same population as in Experiment 1 and 2 (8 men; mean age = 21.2, SD = 4.5). Experiment 3 was similar to Experiments 1 and 2, except that the problems were framed as general causal attributions. The BEAR problem read:

> Encountering a bear during a trek generally causes people to run. Running generally causes people to be out of breath. Can you conclude that encountering a bear during a trek generally causes people to be out of breath?

Philosophers carefully distinguish such generic causation from the instantiated causation featured in Experiments 1 and 2. However, the models we have addressed in this article do not make such a careful distinction. Therefore, it was a natural step to attempt to generalize the findings of Experiments 1 and 2 to generic causation.

Results The right part of Fig. 1 displays the global distribution of the 288 responses to problems with one salient cause vs. many possible causes. In the three problems with one salient cause, mean transitivity acceptance was +0.8 (SD = 0.7), whilst it was only −0.6 (SD = 0.7) in the three problems with many possible causes; $t(47) = 9.5$, $p < .001$, $d = 2.0$. Overall, the transitive conclusion was endorsed (i.e., encoded a +1 or a +2) for 84% of the problems with one salient cause, but only for 33% of the problems with many possible causes.

As previously, participants accepted transitivity in the three problems with one salient cause, but rejected transitivity in the three problems with many possible causes. Mean acceptance ratings for the three problems with one salient cause (FINGER, KETTLE, WEAPON) were +0.9 (SD = 1.0), +1.1 (SD = 1.2), and +0.4 (SD = 1.3), respectively. On the contrary, mean acceptance ratings for the three problems with many possible causes (APPLE, BEAR, COUGH) were −0.9 (SD = 1.1), −0.4 (SD = 1.1), and −0.5 (SD = 1.0), respectively.

7.5 Discussion of behavioral findings

Three behavioral experiments robustly established two main findings. First, lay reasoners do not unconditionally accept the transitivity of causal arguments. This is, to our knowledge, the first time that this fact is clearly established by empirical findings. Second, as expected from the qualitative framework, Saliency is clearly an important factor in the transitivity of ordinary causal transitive judgments, over and above the Markovian character of the causal chains involved. Our results suggest that lay reasoners base their transitive judgments on ‘quick and dirty’ qualitative Saliency considerations, rather than on more sophisticated Markovian considerations, be they qualitative or quantitative.

8 Conclusion

In this article, we have addressed the conditions under which transitive conclusions from causal chains become acceptable. Knowing that A causes B and that B causes C, when are we ready to conclude that A causes C ? We tackled this question through two formal approaches to causation. The first approach is quantitative, and defines causation in terms of probabilistic contrast. The second approach is qualitative, and defines causation in terms of nonmonotonic consequence relations.

We showed that in both approaches, Markovian chains allow transitive conclusions; but also that the qualitative approach features a specific condition for transitivity, which is stronger than the Markov condition. *Saliency* is a sufficient condition for transitivity in the qualitative approach, not in the quantitative approach. Informally speaking, a causal chain "A causes B, B causes C" meets the saliency condition when the occurrence of A is commonly inferred from the occurrence of B; that is, when A can be plausibly inferred from knowing B.

We suggested that the qualitative and quantitative approaches reflected different perspectives on causation, and that the qualitative approach might be closer to the ordinary perspective on causation, adopted in everyday reasoning. Furthermore, we speculated that ordinary reasoners were likely to substitute the weaker Markov condition with the stronger, but cognitively easier, Saliency condition. We thus predicted that the transitivity of ordinary causal inferences would be sensitive to the Saliency condition, even when the Markov condition was otherwise satisfied. We repeatedly validated this speculation through three behavioral experiments, using various linguistic expressions for the denotation of causal relations.

Causality is clearly a multifaceted notion; and one reason for the wide variety of approaches to causality may be the existence of several notions and several contexts having to do with the informal idea of a cause [5]. Few scholars acknowledged this state of facts to the point of suggesting which formal model could account for which variant of causality. We tentatively suggested that the two approaches we studied in this article must not be thought of as competing for providing the best account of causality, but rather as formally tackling two variants of causation.

The scientific variant of causation is essentially concerned with the influence of some variables on other variables; but the commonsense variant of causation rather deals with the way lay persons explain the sudden occurrence of otherwise unexpected events. Because the Markov condition naturally motivates the transitivity of influences among variables, it fits the perspective of the scientific variant, and is accordingly given a prominent role in a literature centered on scientific causation. In the present article, though, we have shown that the commonsense variant of causation may also be tackled formally; and that the difference between the two variants can be laid bare by the existence of different conditions for transitivity. This is an important step toward a complete description of the formal properties of commonsense causation.

References

1. E. Adams. *The logic of conditionals: An application of probability to deductive logic.* Reidel, Dordrecht, 1975.
2. A. K. Barbey and P. Wolff. Causal reasoning from forces. In R. Sun and N. Miyake, editors, *Proceedings of the 28th Annual Conference of the Cognitive Science Society*, page 2439, Mahwah, NJ, 2006. Erlbaum.
3. A. K. Barbey and P. Wolff. Learning causal structure from reasoning. In D. S. McNamara and J. G. Trafton, editors, *Proceedings of the 29th Annual Conference of the Cognitive Science Society*, pages 713–718, Austin, TX, 2007. Cognitive Science Society.
4. J. Bell. Causation as production. In G. Brewka, S. Coradeschi, A. Perini, and P. Traverso, editors, *Proceedings of the 17th European Conference on Artificial Intelligence (ECAI'06)*, pages 327–331, 2006.
5. S. Benferhat, J. F. Bonnefon, P. Chassy, R. M. Da Silva Neves, D. Dubois, F. Dupin de Saint-Cyr, D. Kayser, F. Nouioua, S. Nouioua-Boutouhami, H. Prade, and S. Smaoui. A comparative study of six formal models of causal ascription. *Lecture Notes in Artificial Intelligence*, 5291:47–62, 2008.
6. S. Benferhat, J. F. Bonnefon, and R. M. Da Silva Neves. An overview of possibilistic handling of default reasoning, with experimental studies. *Synthese*, 146:53–70, 2005.
7. S. Benferhat, D. Dubois, and H. Prade. Nonmonotonic reasoning, conditional objects and possibility theory. *Artificial Intelligence*, 92:259–276, 1997.
8. S. Benferhat, D. Dubois, and H. Prade. Possibilistic and standard probabilistic semantics of conditional knowledge bases . *Journal of Logic and Computation*, 9:873–895, 1999.
9. G. Bjornsson. How effects depend on their causes,why causal transitivity fails, and why we care about causation. *Philosophical Studies*, 133:349–390, 2007.
10. J. F. Bonnefon, R. M. Da Silva Neves, D. Dubois, and H. Prade. Background default knowledge and causality ascriptions. In G. Brewka, S. Coradeschi, A. Perini, and P. Traverso, editors, *Proceedings of the 17th European Conference on Artificial Intelligence (ECAI'06)*, pages 11–15, Zurich, 2006. IOS Press.
11. J. F. Bonnefon, R. M. Da Silva Neves, D. Dubois, and H. Prade. Predicting causality ascriptions from background knowledge: Model and experimental validation. *International Journal of Approximate Reasoning*, 48:752–765, 2008.
12. J. F. Bonnefon, D. Dubois, and H. Prade. Transitive observation-based causation, saliency, and the Markov condition. *Lecture Notes in Artificial Intelligence*, 5291:78–91, 2008.
13. M. Buehner, P.W. Cheng, and D. Clifford. From covariation to causation: A test of the assumption of causal power. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29:1119–1140, 2003.
14. N. Campion. Hypothetical and certain inferences from conditional arguments read in text. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32:547–558, 2006.
15. P. W. Cheng. From covariation to causation: A causal power theory. *Psychological Review*, 104:367–405, 1997.
16. P. W. Cheng and L. R. Novick. Causes versus enabling conditions. *Cognition*, 40:83–120, 1991.
17. P. W. Cheng and L. R. Novick. Covariation in natural causal induction. *Psychological Review*, 99:365–382, 1992.
18. R. M. Da Silva Neves, J. F. Bonnefon, and E. Raufaste. An empirical test for patterns of nonmonotonic inference. *Annals of Mathematics and Artificial Intelligence*, 34:107–130, 2002.
19. W. De Neys, W. Schaeken, and G. d'Ydewalle. Causal conditional reasoning and semantic memory retrieval: A test of the semantic memory framework. *Memory and Cognition*, 30:908–920, 2002.
20. D. Dubois, H. Fargier, and H. Prade. Ordinal and probabilistic representations of acceptance. *Journal of Artificial Intelligence Research*, 22:23–56, 2004.

21. D. Dubois, L. Fariñas Del Cerro, A. Herzig, and H. Prade. A roadmap of qualitative independence. In D. Dubois, H. Prade, and E. P. Klement, editors, *Fuzzy Sets, Logics and Reasoning about Knowledge*, Applied Logic (vol. 15), pages 325–350. Kluwer, Dordrecht, 1999.
22. D. Dubois, L. Godo, R. Lopez de Mantaras, and H. Prade. Qualitative reasoning with imprecise probabilities. *Journal of Intelligent Information Systems*, 2:319–363, 1993.
23. D. Dubois and H. Prade. Semantic considerations on order of magnitude reasoning. In M. G. Singh and L. Travé-Massuyès, editors, *Decision Support Systems and Qualitative Reasoning*, pages 223–228. North-Holland, 1991.
24. D. Dubois and H. Prade. Conditional objects as nonmonotonic consequence relations: Main results. In Jon Doyle, Erik Sandewall, and Pietro Torasso, editors, *KR'94: Principles of Knowledge Representation and Reasoning*, pages 170–177. Morgan Kaufmann, San Francisco, CA, 1994.
25. E. Eells and E. Sober. Probabilistic causality and the question of transitivity. *Philosophy of Science*, 50:35–57, 1983.
26. A. Gilio. Probabilistic reasoning under coherence in system P. *Annals of Mathematics and Artificial Intelligence*, 34:5–34, 2002.
27. E. Goldvarg and P. N. Johnson-Laird. Naive causality: a mental model theory of causal meaning and reasoning. *Cognitive Science*, 25:565–610, 2001.
28. I. J. Good. A causal calculus I. *British Journal for the Philosophy of Science*, 11:305–318, 1961.
29. I. J. Good. A causal calculus II. *British Journal for the Philosophy of Science*, 12:43–51, 1962.
30. N. Hall. Causation and the price of transitivity. *Journal of Philosophy*, 97:198–222, 2000.
31. J. Halpern and J. Pearl. Causes and explanations: A structural-model approach — part 1: Causes. *British Journal for the Philosophy of Science*, 56:843–887, 2005.
32. G. Hesslow. The transitivity of causation. *Analysis*, 41:130–133, 1981.
33. C. Hitchcock. The intransitivity of causation revealed in equations and graphs. *Journal of Philosophy*, 98:273–299, 2001.
34. S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44:167–207, 1990.
35. I. Kvart. Transitivity and preemption of causal relevance. *Philosophical Studies*, 64:125–160, 1991.
36. H.E. Kyburg. *Probability and the Logic of Rational Belief.* Wesleyan University Press, Middletown, Ct, 1961.
37. H. E. Kyburg, Jr, C.M. Teng, and G.R. Wheeler. Conditionals and consequences. *Journal of Applied Logic*, 5:638–650, 2007.
38. D. J. Lehmann and M. Magidor. What does a conditional knowledge base entail? *Artif. Intell.*, 55(1):1–60, 1992.
39. K. Lober and D. Shanks. Is causal induction based on causal power? Critique of Cheng (1997). *Psychological Review*, 107:195–212, 2000.
40. E. J. Lowe. For want of a nail. *Analysis*, 40:50–52, 1980.
41. J. L. Mackie. The transitivity of counterfactuals and causation. *Analysis*, 40:53–54, 1980.
42. D. E. Over, C. Hadjichristidis, J. St. B. T Evans, S. J. Handley, and S. A. Sloman. The probability of causal conditionals. *Cognitive Psychology*, 54:62–97, 2007.
43. J. Pearl. *Probabilistic Reasoning in Intelligent Systems.* Morgan Kaufmann, San Mateo, CA, 1988.
44. J. Pearl. *Causality: Models, Reasoning, and Inference.* Cambridge University Press, Cambridge, 2000.
45. N. Pfeifer and G. D. Kleiter. Coherence and nonmonotonicity in human nonmonotonic reasoning. *Synthese*, 146:93–109, 2005.
46. G. Politzer and J. F. Bonnefon. Two varieties of conditionals and two kinds of defeaters help reveal two fundamental types of reasoning. *Mind and Language*, 21:484–503, 2006.
47. D. Poole. The effect of knowledge on belief: Conditioning, specificity and the lottery paradox in default reasoning. *Artificial Intelligence*, 49:281–307, 1991.
48. W. C. Salmon. *Causality and Explanation.* Oxford University Press, New York, 1998.
49. P. Walley. *Statistical reasoning with imprecise Probabilities.* Chapman and Hall, New York, 1991.