

A Logic For Causal Reasoning

Alexander Bochman

Computer Science Department,
Holon Academic Institute of Technology, Israel
e-mail: bochmana@hait.ac.il

Abstract

We introduce a logical formalism of irreflexive causal production relations that possesses both a standard monotonic semantics, and a natural non-monotonic semantics. The formalism is shown to provide a complete characterization for the causal reasoning behind causal theories from [McCain and Turner, 1997]. It is shown also that any causal relation is reducible to its Horn sub-relation with respect to the nonmonotonic semantics. We describe also a general correspondence between causal relations and abductive systems, which shows, in effect, that causal relations allow to express abductive reasoning. The results of the study seem to suggest causal production relations as a viable general framework for nonmonotonic reasoning.

1 Introduction

Causal theories have been introduced in [McCain and Turner, 1997] as a nonmonotonic formalism that provides a natural solution for both the frame and ramification problem in reasoning about actions (see [Giunchiglia *et al.*, 2001] for a detailed exposition). A causal theory is a set of causal rules $A \Rightarrow B$ that express a kind of a causal relation among propositions. The nonmonotonic semantics of such theories is determined by causally explained interpretations, namely the interpretations that are both closed with respect to the causal rules and such that every fact holding in them is caused.

The above formalism has been defined semantically, and the main aim of our study consists in laying down its logical foundations. As we will show, such foundations can be built in the framework of an inference system for causal rules that we will call *causal production relations*. The logical origins of the latter are in input/output logics [Makinson and der Torre, 2000], but we will supply them with a natural nonmonotonic semantics allowing to represent significant parts of nonmonotonic reasoning. Thus, the main result of the present study is that causal production relations completely characterize the inference rules for causal conditionals that preserve the nonmonotonic semantics of causal theories. It will be shown also that any causal theory is reducible with respect to this semantics to a *determinate* causal theory that contain only Horn causal rules with literals in heads.

The importance of determinate theories lies in the fact that (modulo some mild finiteness restrictions) the explained interpretations of such a theory are precisely the interpretations of its classical *completion* (see [McCain and Turner, 1997]). Consequently, the nonmonotonic consequences of such theories are obtainable by the use of classical inference tools (such as the Causal Calculator, described in [Giunchiglia *et al.*, 2001]).

Finally, we will describe a relationship between causal relations and abductive reasoning. As we will see, the latter will be representable via a special kind of causal relations called *abductive causal relations*. This will serve as yet another justification for our general claim that causal production relations could be used as a general-purpose formalism for nonmonotonic reasoning, a viable alternative to other nonmonotonic formalisms (such as default logic).

2 Causal production relations

We will assume that our language is an ordinary propositional language with the classical connectives and constants $\{\wedge, \vee, \neg, \Rightarrow, \mathbf{t}, \mathbf{f}\}$. In addition, \models and Th will denote, respectively, the classical entailment and the associated logical closure operator.

A *causal rule* is a rule of the form $A \Rightarrow B$, where A and B are classical propositions. Informally, such a rule says that, whenever A holds, B is caused. The following definition describes an inference system for causal rules that will be shown to be adequate (and complete) for causal theories. Actually, many of the postulates below (e.g., And and Or) have already been suggested in the literature (see, [Lifschitz, 1997; Schwind, 1999]).

Definition 2.1. A *causal production relation* (or simply a *causal relation*) is a relation \Rightarrow on the set of propositions satisfying the following conditions:

(Strengthening) If $A \models B$ and $B \Rightarrow C$, then $A \Rightarrow C$;

(Weakening) If $A \Rightarrow B$ and $B \models C$, then $A \Rightarrow C$;

(And) If $A \Rightarrow B$ and $A \Rightarrow C$, then $A \Rightarrow B \wedge C$;

(Or) If $A \Rightarrow C$ and $B \Rightarrow C$, then $A \vee B \Rightarrow C$;

(Cut) If $A \Rightarrow B$ and $A \wedge B \Rightarrow C$, then $A \Rightarrow C$;

(Truth) $\mathbf{t} \Rightarrow \mathbf{t}$;

(Falsity) $\mathbf{f} \Rightarrow \mathbf{f}$.

Though causal relations satisfy most of the rules for classical entailment, their distinctive feature is that they are irreflexive, that is, they do not satisfy the postulate $A \Rightarrow A$. Actually, such relations correspond to a strongest kind of *input-output logics* from [Makinson and der Torre, 2000], basic input-output logics with reusable output, with the only distinction that causal relations satisfy also Falsity. The latter restricts the universe of discourse to consistent theories¹.

As a consequence of Cut, any causal relation will already be transitive, that is, it will satisfy

(Transitivity) If $A \Rightarrow B$ and $B \Rightarrow C$, then $A \Rightarrow C$.

However, transitivity is a weaker property than Cut, since it does not imply the latter (see [Makinson and der Torre, 2000]). In addition, the material implications corresponding to the causal rules can be used as auxiliary assumptions in making derivations. This is a consequence of the following property that holds for causal relations:

(AS) If $A \Rightarrow B$ and $C \wedge (A \rightarrow B) \Rightarrow D$, then $C \Rightarrow D$.

Remark. The notion of a causal production sanctioned by the above postulates is *atemporal*. For example, the rule $p \wedge q \Rightarrow \neg q$ cannot be understood as saying that p and q jointly cause $\neg q$ (afterwards) in a temporal sense; instead, by Cut and Falsity it implies $p \wedge q \Rightarrow \mathbf{f}$, which means, in effect, that $p \wedge q$ cannot hold. Speaking generally, ‘causally consistent’ propositions cannot cause an effect incompatible with them. A representation of temporal and action domains in this formalism can be achieved, however, by adding explicit temporal arguments to propositions, just as in the classical logic (see [Giunchiglia *et al.*, 2001]).

A *constraint* is a causal rule of the form $A \Rightarrow \mathbf{f}$. Such constraints correspond to state constraints in action theories. Now, any causal rule implies the corresponding constraint:

(Reduction) If $A \Rightarrow B$, then $A \wedge \neg B \Rightarrow \mathbf{f}$.

Note, however, that the reverse entailment does not hold. Actually, given the rest of the postulates, Reduction can replace the Cut postulate:

Lemma 2.1. *Cut is equivalent to Reduction.*

Another important fact about causal relations is the following ‘decomposition’ of causal rules:

Lemma 2.2. *Any causal rule $A \Rightarrow B$ is equivalent to a pair of rules $A \wedge \neg B \Rightarrow \mathbf{f}$ and $A \wedge B \Rightarrow B$.*

Causal rules of the form $A \wedge B \Rightarrow B$ are ‘classically trivial’, but they play an important explanatory role in non-reflexive causal reasoning. Namely, they say that, in any causally explained interpretation in which A holds, we can freely accept B , since it is self-explanatory in this context. Accordingly, such rules could be called *conditional abducibles*. We will discuss the role of abducibles later as part of the general correspondence between causal and abductive reasoning².

¹In fact, it has been shown in [Bochman, 2002] that other, weaker input-output logics can also be given a semantics of the kind described in what follows.

²Note also that, under the general correspondence between causal inference and default logic (see [Turner, 1999]), such rules correspond to normal defaults in default logic.

Now the above lemma says that any causal rule can be decomposed into a (non-causal) constraint and a conditional abducible. This decomposition neatly delineates two kinds of information conveyed by causal rules. One is a logical information that constrains the set of admissible models, while the other is an explanatory information describing what propositions are caused (explainable) in such models³.

For a finite set u of propositions, $\bigwedge u$ will denote their conjunction. We will extend causal rules to rules having arbitrary sets of propositions as premises using the following ‘compactness’ recipe: for any set u ,

$$u \Rightarrow A \equiv \bigwedge a \Rightarrow A, \text{ for some finite } a \subseteq u.$$

For a given causal relation \Rightarrow , we will denote by $\mathcal{C}(u)$ the set of propositions caused by u , that is

$$\mathcal{C}(u) = \{A \mid u \Rightarrow A\}.$$

The operator \mathcal{C} will play much the same role as the usual derivability operator for consequence relations. It satisfies the following familiar properties:

$$\text{If } u \subseteq v, \text{ then } \mathcal{C}(u) \subseteq \mathcal{C}(v)$$

$$\mathcal{C}(\mathcal{C}(u)) \subseteq \mathcal{C}(u)$$

Still, \mathcal{C} is not inclusive, that is, $u \subseteq \mathcal{C}(u)$ does not always hold. Also, it is not idempotent, that is, $\mathcal{C}(\mathcal{C}(u))$ can be distinct from $\mathcal{C}(u)$.

3 Monotonic semantics of causal inference

We will describe now a monotonic semantics for causal relations. Actually, it will be just a slight modification of the semantics given in [Turner, 1999].

A fully formal semantic interpretation of causal relations could be given in terms of possible worlds frames of the form (i, W) , where i is a propositional interpretation, while W is a set of propositional interpretations that contains i (see [Turner, 1999]). In order to make our descriptions more transparent, however, we will identify such frames with pairs of theories of the form (α, u) , where α is a world (maximal deductively closed set), while u is a deductively closed set included in α . Such pairs will be called *bitheories*. Clearly, any possible worlds frame (i, W) determines a certain bitheory, and vice versa, so all our definitions will admit a purely semantic reformulation.

By a *causal semantics* we will mean a set of bitheories. The corresponding notion of validity for causal rules is given in the next definition.

Definition 3.1. • A causal rule $A \Rightarrow B$ holds in a bitheory (α, u) if either $A \notin \alpha$, or $B \in u$.

- $A \Rightarrow B$ is *valid* with respect to a causal semantics \mathcal{B} if it holds in all bitheories from \mathcal{B} .

We will denote by $\Rightarrow_{\mathcal{B}}$ the set of all causal rules that are valid in a causal semantics \mathcal{B} . It can be easily verified that this set is closed with respect to the postulates for causal relations, and hence we have

³Cf. a similar decomposition of causal rules in [Thielscher, 1997].

Lemma 3.1. \Rightarrow_B is a causal production relation.

To prove completeness, for any causal relation \Rightarrow we consider its *canonical causal semantics* defined as the set $\mathcal{B}_{\Rightarrow}$ of all bitheories of the form $(\alpha, \mathcal{C}(\alpha))$ such that $\mathcal{C}(\alpha) \subseteq \alpha$. Then the following theorem shows that this semantics determines precisely the source causal relation, including the causal rules with arbitrary sets of premises. In other words, the source causal relation is *strongly complete* for this semantics.

Theorem 3.2. For any set of propositions v ,

$$v \Rightarrow A \text{ iff } A \in u, \text{ for any } (\alpha, u) \in \mathcal{B}_{\Rightarrow} \text{ such that } v \subseteq \alpha.$$

Thus, any causal semantics determines a causal relation and vice versa, any causal relation is generated by some causal semantics. Hence we conclude with the following general completeness result:

Corollary 3.3. A relation on propositions is a causal relation iff it is generated by a causal semantics.

3.1 Possible worlds semantics

Causal relations can also be given a semantic interpretation in terms of standard possible worlds models.

As usual, a *possible worlds model* is a triple $\mathbb{W} = (W, R, V)$, where W is a set of worlds, R is a binary accessibility relation on W , while V is a function assigning each world a propositional interpretation. Intuitively, $\alpha R \beta$ means that α and β are, respectively, an initial state (input) and a possible output state of a causal production based on a given set of causal rules. A possible worlds model is *quasi-reflexive* if its accessibility relation satisfies the condition:

$$\text{If } \alpha R \beta, \text{ then } \alpha R \alpha.$$

The following definition provides the notion of validity for causal rules in such models:

Definition 3.2. $A \Rightarrow B$ will be said to be *valid* in a possible worlds model (W, R) if, for any worlds α, β such that $\alpha R \beta$, if A holds in α , then B holds in β .

Then we have

Theorem 3.4. A relation is causal if and only if it is determined by a quasi-reflexive possible worlds model.

The above semantics immediately suggests a modal translation of causal rules (see [Turner, 1999]). Namely, the validity of $A \Rightarrow B$ in a possible worlds model is equivalent to validity of the formula $A \rightarrow \Box B$, where \Box is a standard modal operator. Consequently, causal rules are representable by such modal formulas.

4 The nonmonotonic semantics

In addition to the monotonic semantics, a causal relation determines also a natural nonmonotonic semantics.

Recall that the canonical semantics of a causal relation \Rightarrow is determined by all worlds α for which $\mathcal{C}(\alpha) \subseteq \alpha$. These are the worlds that are closed with respect to the causal rules of \Rightarrow . Still, due to non-reflexivity of \Rightarrow , only some of them are worlds in which any proposition is also caused by some causal rule. These are the worlds α that are fixed points of \mathcal{C} , that is, the worlds for which $\alpha = \mathcal{C}(\alpha)$. In such worlds any fact is causally explained. It is these worlds that determine the nonmonotonic semantics, defined below.

Definition 4.1. A *causally explained world* for a causal relation \Rightarrow is a world α such that $\alpha = \mathcal{C}(\alpha)$. The *nonmonotonic semantics* of \Rightarrow is the set of all causally explained worlds.

Propositions that hold in all causally explained worlds are considered as the nonmonotonic consequences determined by the causal relation. This semantics is indeed nonmonotonic, since the set of such consequences changes nonmonotonically with changes in the underlying set of causal rules.

Now we will establish a correspondence between the above semantics and the nonmonotonic semantics of causal theories from [McCain and Turner, 1997].

A *causal theory* is an arbitrary set of causal rules. Given a causal theory Δ and a set u of propositions, we define the following set:

$$\Delta(u) = \{A \mid B \Rightarrow A \in \Delta, \text{ for some } B \in u\}$$

According to [McCain and Turner, 1997], a world α is *causally explained* with respect to Δ if it is the only world containing $\Delta(\alpha)$. Or, equivalently, when $\alpha = \text{Th}(\Delta(\alpha))$.

Now, causal relations can also be considered as (rather big) causal theories. Moreover, it can be easily verified that for the latter the above definition is equivalent to the definition of causally explained worlds, given earlier. Note also that for any causal theory Δ there exists a least causal relation that includes Δ . We will denote it by \Rightarrow_{Δ} . Actually, \Rightarrow_{Δ} is the set of all causal rules that can be derived from Δ using the postulates for causal relations.

Let us introduce now the following definition:

Definition 4.2. Two causal theories will be called (non-monotonically) *equivalent* if they determine the same set of causally explained worlds.

As a first part of our general result, the next theorem shows that the postulates of causal inference preserve the above non-monotonic semantics.

Theorem 4.1. \Rightarrow_{Δ} is equivalent to Δ .

The above theorem shows, in effect, that the logic of causal relations is adequate for the nonmonotonic reasoning in causal theories.

Let us say that a causal rule is *Horn* one if it has the form $A \Rightarrow l$, where l is a literal or the falsity constant \mathbf{f} ⁴. Our next result will show that any causal theory is equivalent to a Horn theory. But first one more distinction.

Definition 4.3. • A causal theory is *determinate* if it contains only Horn rules. A causal relation will be called *determinate* if it is generated by a determinate causal theory.

- A causal theory Δ will be said to be *locally finite* if, for any propositional atom p , there is only a finite number of causal rules in Δ that contain p in their heads.
- A causal theory will be called *definite* if it is determinate and locally finite.

Clearly, any finite causal theory will be locally finite, though not vice versa. Similarly, any finite determinate theory will be definite. As can be easily verified, a determinate

⁴Thus, any constraint $A \Rightarrow \mathbf{f}$ will be a Horn rule in this sense.

causal theory is definite if and only if, for any literal l , there is no more than a finite number of rules with the head l . Consequently, our definition of definite theories turns out to be equivalent to that given in [McCain and Turner, 1997].

By a *determinate subrelation* of a causal relation \Rightarrow we will mean the least causal relation generated by the set of all Horn rules from \Rightarrow . Then we have

Theorem 4.2. *Any causal relation is nonmonotonically equivalent to its determinate subrelation.*

By Theorem 4.1 we can conclude now that any causal theory is equivalent to some determinate theory.

Example. The causal theory

$$\{p \Rightarrow p, \neg p \Rightarrow \neg p, \neg q \Rightarrow \neg q, \mathbf{t} \Rightarrow p \vee q\}$$

has a single causally explained interpretation $\{p, \neg q\}$. The last rule $\mathbf{t} \Rightarrow p \vee q$ is not Horn, but it implies a constraint $\neg(p \vee q) \Rightarrow \mathbf{f}$ (by Reduction), and it can be verified that the determinate causal theory obtained by replacing the rule with this constraint will be equivalent to the original theory.

Actually, there is a simple algorithm of transforming any causal theory Δ into an equivalent determinate theory Δ_d . It consists in two steps:

1. For any causal rule $A \Rightarrow B$ from Δ , a constraint $A \wedge \neg B \Rightarrow \mathbf{f}$ is added to Δ_d ;
2. If $\{A_i \Rightarrow B_i\}$ is a (minimal) set of rules from Δ such that $\bigwedge B_i \models l$, for some literal l , then a rule $\bigwedge A_i \Rightarrow l$ is added to Δ_d .

Then the following result can be proved:

Lemma 4.3. Δ is equivalent to Δ_d .

Unfortunately, the above algorithm is not modular. Moreover, it does not preserve local finiteness: there are locally finite causal theories such that their determinate counterparts are not locally finite. Still, in many simple cases it gives a convenient recipe for transforming an arbitrary causal theory into an equivalent determinate theory.

Now we are going to show the second part of our main result, namely that causal relations constitute a maximal logic suitable for reasoning in causal theories. To begin with, we introduce

Definition 4.4. Causal theories Δ and Γ are *causally equivalent* if \Rightarrow_Δ coincides with \Rightarrow_Γ .

Two theories are causally equivalent if each can be obtained from the other using the postulates of causal relations. In other words, they are equivalent in the ‘causal logic’ associated with the latter. For example, a theory $\{A \vee B \Rightarrow C\}$ is causally equivalent to the theory $\{A \Rightarrow C, B \Rightarrow C\}$ (by Strengthening and Or). Speaking generally, any causal theory is causally equivalent to a theory with rules of the form $\bigwedge l_i \Rightarrow \bigvee l_j$, where l_i, l_j are literals.

Now, by Theorem 4.1, we immediately obtain

Corollary 4.4. *Causally equivalent theories are nonmonotonically equivalent.*

The reverse implication in the above corollary does not hold, and a deep reason for this is that the causal equivalence

is a monotonic (logical) notion, and hence, unlike the non-monotonic equivalence, it is preserved under addition of new causal rules. For example, though any causal theory is equivalent to a determinate one, they may give different results if we add some causal rules to them.

Example. (continued) Let us add a rule $\mathbf{t} \Rightarrow \neg p \vee q$ both to the causal theory in the preceding example, and to its determinate counterpart. In the first case we obtain a theory which implies $\mathbf{t} \Rightarrow q$ (by And), and hence it will have two explained interpretations $\{p, q\}$ and $\{\neg p, q\}$. For the reduced determinate theory, however, the addition of this rule will give $\{p, q\}$ as the only explained interpretation.

What we need is a stronger, monotonic counterpart of the notion of equivalence that would be preserved under addition of new causal rules. This immediately suggests the following definition.

Definition 4.5. Two causal theories Δ and Γ will be said to be *strongly equivalent* if, for any set Φ of causal rules, $\Delta \cup \Phi$ is nonmonotonically equivalent to $\Gamma \cup \Phi$.

Strongly equivalent theories are ‘equivalent forever’, that is, they are interchangeable in any larger causal theory without changing the nonmonotonic semantics. Consequently, strong equivalence can be seen as an equivalence with respect to the monotonic logic of causal theories. And the next result shows that this logic is precisely the logic of causal relations.

Theorem 4.5. *Two causal theories are strongly equivalent if and only if they are causally equivalent.*

The above result (and its proof) has the same pattern as the corresponding results about strong equivalence of logic programs and default theories (see [Lifschitz *et al.*, 2001; Turner, 2001]).

The above result implies, in effect, that causal relations are maximal inference relations that are adequate for reasoning with causal theories: any postulate that is not valid for causal relations can be ‘falsified’ by finding a suitable extension of two causal theories that would determine different causally explained interpretations, and hence would produce different nonmonotonic conclusions.

5 Causation versus abduction

Causal inference has numerous connections with other non-monotonic formalisms. Thus, [Turner, 1999] describes the relations with circumscription, autoepistemic and default logic. In this section we will describe a correspondence between causal inference and abductive reasoning. In addition to specific results, this will give us a broader perspective on the representation capabilities of causal relations.

A general *abductive framework* can be defined as a pair $\mathbb{A} = (\text{Cn}, \mathcal{A})$, where Cn is a supraclassical consequence relation, while \mathcal{A} is a distinguished set of *abducibles*. A proposition A is *explainable* in an abductive framework \mathbb{A} if there exists a consistent set of abducibles $a \subseteq \mathcal{A}$ such that $A \in \text{Cn}(a)$.

A general correspondence between abductive frameworks and input-output logics has been established in [Bochman, 2002], based on the identification between the above notion

of explainability and input-output derivability. By this correspondence, abducibles are representable by ‘reflexive’ propositions satisfying input-output rules $A \Rightarrow A$, while abductive frameworks themselves correspond exactly to input-output logics satisfying an additional postulate of Abduction (see below). By these results, input-output logics allow to give a syntax-independent representation of abductive reasoning.

5.1 The abductive subrelation

We will begin with the following definitions:

Definition 5.1. • A proposition A is an *abducible* in a causal relation \Rightarrow if $A \Rightarrow A$;

- A causal relation is *abductive* if it satisfies

(Abduction) If $A \Rightarrow B$, then $A \Rightarrow C \Rightarrow B$, for some abducible C .

Causal inference in abductive causal relations is always mediated by abducibles. Consequently, propositions caused by worlds are caused, in effect, by the abducibles that hold in these worlds:

$$\mathcal{C}(\alpha) = \mathcal{C}(\alpha \cap \mathcal{A}),$$

where \mathcal{A} is the set of abducibles. As we mentioned, abductive causal relations correspond to a certain class of abductive systems. Moreover, it has been shown in [Bochman, 2002] that the nonmonotonic semantics of such causal relations describes, in effect, the *explanatory closure*, or *completion*, in associated abductive frameworks (see [Console *et al.*, 1991]).

It turns out that any causal relation \Rightarrow includes a greatest abductive subrelation. The latter can be defined as follows:

$$A \Rightarrow^a B \equiv (\exists C)(A \Rightarrow C \Rightarrow C \Rightarrow B)$$

Theorem 5.1. *If \Rightarrow is a causal relation, then \Rightarrow^a is the greatest abductive causal relation included in \Rightarrow .*

The abductive subrelation of a causal relation preserves many properties of the latter. For example, both have the same constraints and abducibles.

Recall that in causally explained interpretations any proposition is caused by other propositions, the latter are also caused by accepted propositions, and so on. Clearly, if our ‘causal resources’ are limited, such a causal chain should stop somewhere. More exactly, it should reach abducible (self-explanatory) propositions. This indicates that in many cases the nonmonotonic semantics associated with a causal theory should be determined by the corresponding abductive subrelation. Below we will make this claim precise.

Definition 5.2. A causal relation will be called *weakly abductive* if it is nonmonotonically equivalent to its abductive subrelation.

The next definition will give us an important sufficient condition for weak abductivity.

Definition 5.3. A causal relation \Rightarrow is *well-founded* if any infinite sequence $\{A_0, A_1, A_2, \dots\}$ of propositions such that $A_{n+1} \Rightarrow A_n$, for every $n \geq 0$, contains an abducible.

The above definition describes a variant of a standard notion of well-foundedness with respect to the (transitive) causal order. It should be clear that any causal relation in

a finite language should be well-founded. Moreover, let us say that a causal relation is *finitary* if it is generated by some finite set of causal rules. Then we have

Lemma 5.2. *Any finitary causal relation is well-founded.*

Finally, the next result shows that all such causal relations will be weakly abductive.

Theorem 5.3. *Any well-founded causal relation is weakly abductive.*

It turns out that well-foundedness is not the only condition that is sufficient for weak abductivity. Thus, adequately acyclic causal theories (see [McCain and Turner, 1998] for a definition) are not in general well-founded, but they also satisfy this property.

Theorem 5.4. *Any adequately acyclic causal relation is weakly abductive.*

The above results show that in many cases of interest, the nonmonotonic semantics of causal theories can also be computed using abduction. Still, as a general conclusion, we can say that causal inference constitutes a proper generalization of abductive reasoning beyond the situations where facts are explainable by reference to a fixed set of self-explanatory abducibles.

6 Conclusions and perspectives

Summing up the main results, we can argue that causal relations constitute a powerful formalism for nonmonotonic reasoning, especially suitable for representing action domains. It has been shown, in particular, that it gives an exact description of the logic underlying causal theories. We have seen also that it allows to give a syntax-independent representation of abductive reasoning. If we add to this also the natural correspondences with other nonmonotonic formalisms, such as default logic and circumscription (see [Turner, 1999]), we can safely conclude that causal inference covers a significant part of nonmonotonic reasoning.

Viewed from this perspective, there is still much to be done in order to realize the opportunities created by the formalism. There are two kinds of problems that need to be resolved in this respect.

The nonmonotonic semantics for causal theories is based on the principle of universal causation which is obviously very strong. The principle implies, for example, that if we have no causal rules for a certain proposition, it should be false in all explainable interpretations. As a result, adding a new propositional atom to the language makes any causal theory inconsistent, since we have no causal rules for it. This makes causal theories extremely sensitive to the underlying language in which they are formulated. One way out has been suggested in [Lifschitz, 1997]; it amounts to restricting the principle of universal causation to a particular subset of *explainable* propositions. This approach, however, is purely syntactical and hence retains language dependence.

More subtle, yet perceptible difficulties arise also in representing indeterminate situations in causal theories. Thus, since any causal theory is reducible to a determinate theory, causal rules with disjunctive heads are ignored in the nonmonotonic semantics; more exactly, they are informative only

to the extent that they imply corresponding non-causal constraints or Horn rules. This does not mean, however, that we cannot represent indeterminate information in causal theories. Actually, one of the main contributions of [McCain and Turner, 1997] consisted in showing how we can do this quite naturally in common cases (see also [Lin, 1996]). Still, there is yet no systematic understanding whether and how an indeterminate information can be represented by Horn causal rules.

A more general problem concerns the role and place of causal inference in general nonmonotonic reasoning. Though the former covers many areas of nonmonotonic reasoning, it does not cover them all. Thus, it does not seem suitable for solving the qualification problem in representing actions. Speaking generally, causal reasoning appears to be independent of the kind of nonmonotonicity described by preferential inference from [Kraus *et al.*, 1990]. This naturally suggests that the two kinds of nonmonotonic reasoning with conditionals could be combined into a single formalism, a grand uniform theory of nonmonotonic reasoning. Actually, this idea is not new; it has been explored more than ten years ago in [Geffner, 1992]. It remains to be seen whether current studies of causal reasoning can contribute to viability of such a general theory.

Acknowledgments. The author is grateful to Vladimir Lifschitz, Hudson Turner, David Makinson and Leendert van der Torre, as well as to the anonymous referees, for their instructive comments and suggestions.

References

- [Bochman, 2002] A. Bochman. Input-output logics: Production inference, nonmonotonicity, and abduction, 2002. (Draft).
- [Console *et al.*, 1991] L. Console, D. Theseider Dupre, and P. Torasso. On the relationship between abduction and deduction. *Journal of Logic and Computation*, 1:661–690, 1991.
- [Geffner, 1992] H. Geffner. *Default Reasoning. Causal and Conditional Theories*. MIT Press, 1992.
- [Giunchiglia *et al.*, 2001] E. Giunchiglia, J. Lee, V. Lifschitz, N. McCain, and H. Turner. Nonmonotonic causal theories, 2001. *Artificial Intelligence* (to appear).
- [Kraus *et al.*, 1990] S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44:167–207, 1990.
- [Lifschitz *et al.*, 2001] V. Lifschitz, D. Pearce, and A. Valverde. Strongly equivalent logic programs. *ACM Transactions on Computational Logic*, 2:526–541, 2001.
- [Lifschitz, 1997] V. Lifschitz. On the logic of causal explanation. *Artificial Intelligence*, 96:451–465, 1997.
- [Lin, 1996] F. Lin. Embracing causality in specifying the indeterminate effects of actions. In *Proceedings AAAI-96*, pages 670–676, 1996.
- [Makinson and der Torre, 2000] D. Makinson and L. Van der Torre. Input/Output logics. *Journal of Philosophical Logic*, 29:383–408, 2000.
- [McCain and Turner, 1997] N. McCain and H. Turner. Causal theories of action and change. In *Proceedings AAAI-97*, pages 460–465, 1997.
- [McCain and Turner, 1998] N. McCain and H. Turner. Satisfiability planning with causal theories. In *Proceedings KR-98*, Trento, 1998.
- [Schwind, 1999] C. Schwind. Causality in action theories. *Linköping Electronic Articles in Computer and Information Science*, 4(4), 1999.
- [Thielscher, 1997] M. Thielscher. Ramification and causality. *Artificial Intelligence*, 89:317–364, 1997.
- [Turner, 1999] H. Turner. A logic of universal causation. *Artificial Intelligence*, 113:87–123, 1999.
- [Turner, 2001] H. Turner. Strong equivalence for logic programs and default theories (made easy). In T. Eiter, W. Faber, and M. Truszczyński, editors, *Proceedings LP-NMR'01*, pages 81–92, Vienna, 2001. Springer. LNAI 2173.