On Disjunctive Causal Inference and Indeterminism

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

Abstract

The formalism of causal production relations, introduced in [Bochman, 2003], is extended to causal rules with multiple (disjunctive) heads, purported to provide a direct representation of indeterminate causation. We describe both a monotonic and (stable) nonmonotonic semantics for disjunctive causal theories. It is shown, however, that disjunctive causal rules are reducible to ordinary (singular) causal rules with respect to the stable semantics.

1 Introduction

Causal theories from [McCain and Turner, 1997] are based on causal rules of the form $A \Rightarrow B$ saying that, whenever A holds, B is caused. The formalism and its associated nonmonotonic semantics have suggested a natural solution for the frame and ramification problems in reasoning about actions (see, e.g., [Giunchiglia *et al.*, 2001] for a detailed exposition). In [Bochman, 2003], we have introduced a logical formalism of causal production relations and showed that it constitutes an adequate logical basis for reasoning with such causal rules.

In this report, the formalism of causal relations will be generalized to disjunctive causal rules of the form $A \Rightarrow B_1, \dots, B_n$. Such rules say, roughly, that, whenever A holds, one of B_i is caused. It was suggested already in [Lin, 1996] that rules of this kind can be used for a direct description of causal and action contexts that involve indeterminate effects. We describe both a monotonic and stable nonmonotonic semantics of disjunctive causal theories, and study the properties of this formalism. As our main result, however, we will show that disjunctive causal rules are reducible to ordinary singular causal rules with respect to the stable semantics. Finally, we will discuss the implications of this result for the problem of representing indeterminate effects in causal reasoning.

We will assume that our background language is a classical propositional language with the usual connectives and constants. $\vDash$ will denote the classical entailment. Throughout the paper $a, b, \dots$ will denote finite sets of propositions, while $u, v, \dots$ arbitrary such sets. $\bigwedge a$ ($\bigvee a$) will denote the conjunction (resp. disjunction) of all propositions from a. As a special case, $\bigwedge \emptyset$ ($\bigvee \emptyset$) will denote the constant $\mathbf{t}$ (resp., $\mathbf{f}$). Finally, $\overline{u}$ will denote the complement of u in the set of all propositions, and $\neg u$ the set $\{\neg A \mid A \in u\}$.

1.1 Causal production relations

In order to make the paper self-contained, we will give first a brief description of causal production relations from [Bochman, 2003]. This formalism constitutes, in effect, a logical inference system for reasoning with 'simple' causal rules of the form $A \Rightarrow B$.

Definition 1.1. A *causal (production) relation* is a relation $\Rightarrow$ on the set of classical propositions satisfying the following conditions:

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

(Weakening) If $A \Rightarrow B$ and $B \vDash C$, 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 $A \Rightarrow A$. Still, it should be mentioned that the above formalism of causal relations is thoroughly monotonic, since it satisfies Strengthening. According to the latter, if $A \Rightarrow B$ is an acceptable causal rule, then any causal rule $A \wedge B \Rightarrow C$ will also be acceptable. Nevertheless, due to irreflexivity, the formalism determines also a natural kind of *nonmonotonic* causal reasoning. The latter is based on the requirement that all facts holding in intended models should be caused (explained) by other facts that occur. Accordingly, causal relations give rise to two semantic representations. The first is a monotonic semantics of the causal rules, while the second one is the nonmonotonic semantics that chooses preferred (explained) models from the monotonic semantics.

The monotonic semantics for causal relations is given below.

A *bitheory* is a pair (α, u), where α is a maximal classically consistent set of propositions (a 'world'), while u is a deductively closed set included in α. A *causal semantics* is a set of bitheories.

Remark. Since a deductively closed theory is faithfully representable by a set of worlds, bitheories naturally correspond to possible worlds frames of the form (α, W), where W is a set of worlds, and $\alpha \in W$. Such frames form a semantics of the Universal Causal Logic (UCL), suggested in [Turner, 1999]. Accordingly, the causal semantics, described below, can be seen as just a notational variant of the latter.

Definition 1.2.
- 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 $\mathcal{B}$.

It can be easily verified that the set of causal rules that are valid with respect to a causal semantics satisfies all the above postulates for a causal production relation. Moreover, for a given causal relation, we can construct its *canonical* causal semantics as follows.

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

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

Now, *a bitheory of* a causal relation is defined as a pair $(\alpha, \mathcal{C}(\alpha))$, where α is a world such that $\mathcal{C}(\alpha) \subseteq \alpha$ (such worlds correspond to interpretations that are closed with respect to the rules of the causal relation). Finally, the canonical semantics of a causal relation will be defined as a set of all its bitheories.

It has been shown in [Bochman, 2003] that the canonical semantics is adequate for the source causal relation, which immediately implies that causal relations are complete for the above causal semantics.

Finally, the canonical semantics determines also a natural *nonmonotonic semantics* for a causal relation. The latter is defined as a set of all bitheories $(\alpha, \mathcal{C}(\alpha))$ of the causal relation, for which $\alpha = \mathcal{C}(\alpha)$. Such bitheories single out worlds that are both closed with respect to the causal rules and such that any proposition that holds in them is caused (explained). The resulting logical system has been shown to be adequate for nonmonotonic reasoning with causal theories of [McCain and Turner, 1997] in the following sense:

- The postulates of causal relations preserve the nonmonotonic semantics of causal theories;
- Causal production relations constitute a maximal logic that preserves the nonmonotonic semantics of causal theories under arbitrary additions of causal rules (see [Bochman, 2003] for details).

2 Disjunctive causal relations

In this section we will describe a generalization of causal relations to disjunctive causal rules that involve sets of propositions in their heads. This generalization will naturally correspond to well-known generalizations of normal logic programs to disjunctive logic programs. An apparent necessity of rules of this kind for representing indeterminate causation could be justified as follows.

Suppose that we have two ordinary causal rules $A \wedge C \Rightarrow B$ and $A \wedge \neg C \Rightarrow \neg B$. The rules can be interpreted as saying that A causes either B or $\neg B$, depending on whether the additional condition C holds. Suppose now that the actual truth-value of C is unknown to us (C is a 'hidden parameter'). Still, the above rules convey a nontrivial information about the situations in question, namely that when A holds, than either B is caused or $\neg B$ is caused. As we will see, this information is expressible using a disjunctive causal rule $A \Rightarrow B, \neg B$. Note in this respect that this information is completely lost in a singular causal rule $A \Rightarrow B \vee \neg B$, which always holds for causal relations (by Truth). This immediately suggests that disjunctive causal rules could be used for a direct representation of indeterminate causation.

It turns out that reasoning with disjunctive causal rules can also be axiomatized in an inference system that will form a natural generalization of causal production relations.

We will consider disjunctive causal rules as rules holding primarily between finite sets of classical propositions: $a \Rightarrow b$ will be taken to mean that if all the propositions in a hold, then one of the propositions in b is caused. As we will see, however, the set of propositions in the premises of the causal rules can always be replaced by their conjunction, so disjunctive causal rules could also be seen as relations between propositions and sets of propositions.

We will use below an ordinary notation for premise and conclusion sets in causal rules. Thus, $a, b \Rightarrow c$ will stand for $a \cup b \Rightarrow c$, and $a, A \Rightarrow$ will mean $a \cup \{A\} \Rightarrow \emptyset$, etc.

Definition 2.1. A *disjunctive causal relation* is a binary relation $\Rightarrow$ on finite sets of classical propositions satisfying the following conditions:

(Left Monotonicity) If $a \Rightarrow b$, then $A, a \Rightarrow b$;

(Right Monotonicity) If $a \Rightarrow b$, then $a \Rightarrow b, A$;

(Cut) If $a \Rightarrow b, A$ and $A, a \Rightarrow b$, then $a \Rightarrow b$;

(Strengthening) If $A \vDash B$ and $a, B \Rightarrow b$, then $a, A \Rightarrow b$;

(Weakening) If $A \vDash B$ and $a \Rightarrow b, A$, then $a \Rightarrow b, B$;

(Left And) If $a, A \wedge B \Rightarrow b$, then $a, A, B \Rightarrow b$;

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

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

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

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

Similarly to ordinary consequence relations, a disjunctive causal relation can be extended to arbitrary sets of propositions in premises and conclusions by requiring *compactness*: for any sets of propositions u, v,

(Compactness) $u \Rightarrow v$ if and only if $a \Rightarrow b$, for some finite $a \subseteq u$, $b \subseteq v$.

As can be seen, a disjunctive causal relation forms a subsystem of the classical sequent calculus in which Reflexivity does not hold. Still, it is straightforward to show that the postulates imply the following two general rules describing the relation between causal inference and classical entailment:

(Logical Strengthening) If $c \vDash A$ and $A, a \Rightarrow b$, then $c, a \Rightarrow b$.

(Logical Weakening) If $c \vDash A$ and $a \Rightarrow b, C$, for any $C \in c$, then $a \Rightarrow b, A$.

The above rules imply, in particular, that classically equivalent propositions are interchangeable both in premises and conclusions of the causal rules. The first of the above rules implies also that a finite set of premises in a causal rule can be replaced by their conjunction:

$$a \Rightarrow b \text{ iff } \bigwedge a \Rightarrow b$$

Moreover, the above rules imply, in effect, that the conjunction behaves in a fully classical way in the context of disjunctive causal relations. This is not so, however, for disjunction and negation. Thus, the conclusion sets in disjunctive causal rules cannot be replaced with their classical disjunctions; we have only that $a \Rightarrow b$ implies $a \Rightarrow \bigvee b$, though not vice versa. Also, only one of the two structural rules for negation holds for causal relations, namely

(Reduction) If $a \Rightarrow b, A$, then $a, \neg A \Rightarrow b$.

Any set of disjunctive causal rules Δ determines a unique least disjunctive causal relation that includes Δ; it will be denoted by $\Rightarrow_\Delta$. The latter causal relation consists of all the causal rules that are derivable from Δ by the postulates for disjunctive causal relations.

A *bitheory of* a disjunctive causal relation $\Rightarrow$ will be defined as any bitheory (α, u) such that $\alpha \not\Rightarrow \overline{u}$. As shows the lemma below, such bitheories can be seen as 'closed' with respect to the causal rules of $\Rightarrow$.

Lemma 2.1. *A bitheory (α, u) is a bitheory of a causal relation $\Rightarrow$ if and only if $b \cap u \neq \emptyset$, for any causal rule $a \Rightarrow b$ from $\Rightarrow$ such that $a \subseteq \alpha$.*

If (α, u) is a bitheory of $\Rightarrow$, and v an arbitrary deductively closed set such that $u \subseteq v \subseteq \alpha$, then (α, v) will also be a bitheory of $\Rightarrow$. This indicates that the set of bitheories of a causal relation is determined, in effect, by its inclusion minimal elements[1]. Accordingly, the notion of a minimal bitheory, defined below, will play an important role in what follows.

Definition 2.2. A bitheory (α, u) of $\Rightarrow$ will be called *minimal* if there is no bitheory (α, v) of $\Rightarrow$ such that $v \subset u$.

The following result gives a convenient alternative description of minimal bitheories.

Lemma 2.2. *If α is a causally consistent world, then a pair (α, u) is a minimal bitheory of $\Rightarrow$ if and only if, for any proposition A,*

$$A \in u \text{ iff } \alpha \Rightarrow \overline{u}, A.$$

The importance of minimal bitheories stems from the following fact:

Lemma 2.3. *If $v \not\Rightarrow w$, then there exists a minimal bitheory (α, u) of $\Rightarrow$ such that $v \subseteq \alpha$ and $w \cap u = \emptyset$.*

The above lemma plays the main role in the proof of the completeness theorem stated in the next section.

[1] Such bitheories always exist due to compactness.

2.1 The monotonic semantics of disjunctive causal inference

As we are going to show now, the causal semantics defined earlier for causal production relations turns out to be suitable also for disjunctive causal relations. To this end, we have only to define the notion of validity of disjunctive causal rules with respect to a causal semantics.

Definition 2.3.
- A rule $a \Rightarrow b$ will be said to *hold* with respect to a bitheory (α, u) if $b \cap u \neq \emptyset$ whenever $a \subseteq \alpha$.
- A rule $a \Rightarrow b$ will be said to be *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 disjunctive 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 disjunctive causal relations, and hence we have

Lemma 2.4. *For any causal semantics $\mathcal{B}$, $\Rightarrow_\mathcal{B}$ is a disjunctive causal relation.*

In order to prove completeness, for any disjunctive causal relation $\Rightarrow$, we consider its *canonical causal semantics*. The latter will be defined as the set $\mathcal{B}_\Rightarrow$ of all minimal bitheories of $\Rightarrow$. Then the following theorem shows that the source disjunctive causal relation is strongly complete for this semantics.

Theorem 2.5. *If $\mathcal{B}_\Rightarrow$ is the canonical semantics for a disjunctive causal relation $\Rightarrow$, then, for any sets of propositions v, w,*

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

The above theorem says that a causal rule $v \Rightarrow w$ belongs to a causal relation $\Rightarrow$ if and only if it is valid in the canonical causal semantics of $\Rightarrow$. Combining this with our previous result, we finally conclude with

Corollary 2.6. *A relation on sets of propositions is a disjunctive causal relation if and only if it is determined by a causal semantics.*

Due to a direct correspondence between causal semantics and the possible worlds semantics from [Turner, 1999], it can be easily seen that our disjunctive causal relations correspond to a subsystem of Turner's UCL. Namely, a causal rule $a \Rightarrow B_1, \ldots, B_n$ corresponds to the UCL formula

$$\bigwedge a \to \mathbf{C}B_1 \vee \cdots \vee \mathbf{C}B_n.$$

Thus, disjunctive causal relations correspond to causal theories of UCL that involve modal formulas with only positive occurrences of the causal operator $\mathbf{C}$ (see also [Lin, 1996] about the use of such rules in representing the indeterminate effects of actions).

3 Singular causal inference

Disjunctive causal rules having a single proposition in their heads will be called *singular* in what follows. As can be seen from the semantic characterization, given in the preceding section, such disjunctive rules have the same meaning as causal production rules from [Bochman, 2003]. In this sense, disjunctive causal relations subsume causal production relations described in the latter paper. In this section we will give a precise characterization of disjunctive causal relations that correspond to causal production relations.

To begin with, the following fact can be easily verified:

Lemma 3.1. *The set of singular rules belonging to a disjunctive causal relation forms a causal production relation.*

In what follows, the above causal production relation will be called the *normal subrelation* of a disjunctive causal relation.

Actually, there is a quite simple and modular recipe how such a subrelation could be obtained from a given set of disjunctive causal rules.

For any set of disjunctive causal rules Δ, let us consider the following set of singular rules:

$$N(\Delta) = \{a, \neg b \Rightarrow \bigvee c \mid a \Rightarrow b, c \in \Delta\}$$

Let us denote by $\Rightarrow^n_{N(\Delta)}$ the least causal production relation that includes $N(\Delta)$. Then the following result shows, in effect, that the set $N(\Delta)$ captures the 'singular content' of Δ.

Theorem 3.2. $\Rightarrow^n_{N(\Delta)}$ *is the normal subrelation of* $\Rightarrow_\Delta$.

The above theorem shows that the set of singular causal rules derivable from a given set of disjunctive causal rules Δ coincides with the set of rules that are derivable from $N(\Delta)$ using only the postulates for causal production relations.

Disjunctive causal relations that are generated by singular causal rules constitute a disjunctive counterpart of causal production relations. Accordingly, we will introduce the following

Definition 3.1. A disjunctive causal relation will be called *singular* if it is a least disjunctive causal relation containing some set of singular causal rules.

Clearly, a disjunctive causal relation $\Rightarrow$ is singular if and only if it coincides with a least disjunctive relation that contain all the singular rules from $\Rightarrow$. In what follows, we are going to give more instructive descriptions of such disjunctive causal relations.

The next corollary gives yet another characterization of singular causal relations. It says that for such relations, worlds produce only determinate effects.

Corollary 3.3. *A disjunctive causal relation is singular if and only if, for any world* α, $\alpha \Rightarrow b, c$ *only if either* $\alpha \Rightarrow b$ *or* $\alpha \Rightarrow c$.

As immediately follows from the above corollary, singular causal relations can also be described in terms of a certain natural constraint on their canonical semantics.

A causal semantics $\mathcal{B}$ will be called *functional*, if for any world α there is no more than one theory u such that $(\alpha, u) \in \mathcal{B}$. Then we have

Theorem 3.4. *A disjunctive causal relation is singular if and only if its canonical causal semantics is functional.*

Unfortunately, the above claim cannot be extended to arbitrary causal semantics, since functional causal semantics do not always produce singular causal relations.

In order to give yet another important description of singular causal relations, let us consider first the following example:

Example. Two singular rules $A \wedge C \Rightarrow B$ and $A \wedge \neg C \Rightarrow \neg B$ imply $A \Rightarrow B, \neg B$ (by Right Monotonicity and Or), though they imply neither $A \Rightarrow B$, nor $A \Rightarrow \neg B$.

The above example shows that singular causal rules can generate nontrivial disjunctive rules that cannot be decomposed directly into singular rules. As we are going to show, however, this example describes, in a sense, the only way of producing nontrivial disjunctive rules from singular ones. Thus, the following theorem shows that, for singular causal relations, disjunctive effects are always 'separable' by adding some further assumptions to the premises.

Theorem 3.5. *A disjunctive causal relation is singular if and only if it satisfies the following condition:*

> *If* $a \Rightarrow b, c$, *then* $a, A \Rightarrow b$ *and* $a, \neg A \Rightarrow c$, *for some proposition* A.

The above theorem amounts to saying that indeterminate effects arise in singular disjunctive causal relations only due to 'forgetting' of some relevant parameters. This characteristic property is even more vivid in the following corollary (which has actually been used in the proof of the theorem):

Corollary 3.6. *A disjunctive causal relation is singular if and only if it satisfies the following condition:*

> $a \Rightarrow B_1, \dots, B_n$ *iff there are pairwise incompatible propositions* $A_1, \dots, A_n$ *such that their disjunction is a tautology, and* $a, A_i \Rightarrow B_i$, *for any* $i \leq n$.

In order to obtain a more comprehensive picture of the role of disjunctive causal rules, it is also instructive to consider cases when such rules are *logically* reducible to singular rules with respect to disjunctive causal relations. For example, the rule $A \Rightarrow B, \neg B$, discussed earlier, is equivalent to a pair of singular causal rules $A, B \Rightarrow B$ and $A, \neg B \Rightarrow \neg B$. This is a consequence of the following general fact:

Lemma 3.7. *A rule* $a \Rightarrow B_1, \dots, B_n$ *is reducible to the set of singular rules* $\{a, B_i \Rightarrow B_i \mid i = 1, \dots, n\}$ *and a constraint* $a, \neg B_1, \dots, \neg B_n \Rightarrow \mathbf{f}$ *if and only if* $a, B_i, B_j \Rightarrow B_i$, *for any* $i \neq j$.

Note that the above characteristic condition for reduction holds, in particular, when all B_i are incompatible, given a, that is, when $a, B_i, B_j \Rightarrow \mathbf{f}$. For example (cf. [Lin, 1996]), a disjunctive rule

$$a \Rightarrow B \wedge C, B \wedge \neg C, \neg B \wedge C$$

is reducible to the following set of singular rules:

$$\begin{array}{cc} a, B \Rightarrow B & a, \neg B \Rightarrow \neg B \\ a, C \Rightarrow C & a, \neg C \Rightarrow \neg C \\ \multicolumn{2}{c}{a, \neg B, \neg C \Rightarrow \mathbf{f}} \end{array}$$

The above results show, in effect, that irreducibly disjunctive rules can arise only in cases when their heads contain mutually compatible propositions. For such rules, we need a different strategy.

4 The stable nonmonotonic semantics

A nonmonotonic semantics for arbitrary theories of universal causal logic has been suggested in [Turner, 1999]. Since disjunctive causal relations form a subsystem of UCL, this semantics can be immediately translated into our framework.

Definition 4.1. A world α will be said to be *causally explained* with respect to a disjunctive causal relation $\Rightarrow$ if (α, α) is a minimal bitheory of $\Rightarrow$. The set of all causally explained worlds will be said to form a *stable nonmonotonic semantics* of $\Rightarrow$.

As in the singular case, causally explained worlds of the stable semantics can be seen as worlds that are closed with respect to the causal rules and, in addition, any proposition that holds in them is among the conclusions of some causal rule that is 'active' in the world. Accordingly, the stable semantics appears as a natural, even almost inevitable, extension of the nonmonotonic semantics for singular causal theories to the disjunctive case. In addition, in the general correspondence between UCL and disjunctive default logic [Gelfond *et al.*, 1991], established in [Turner, 1999], causally explained interpretations correspond to extensions of disjunctive default theories. Moreover, by the same correspondence, simple disjunctive theories (with sets of literals as premises and conclusions of causal rules) can be translated into disjunctive logic programs, and then causally explained interpretations will correspond to answer sets of such programs. Summing up, there is a tight correspondence between the above semantics and respectable semantics of nonmonotonic formalisms and disjunctive logic programming.

Nevertheless, we are going to show now that the stable semantics imposes quite a radical view on the role of disjunctive causal rules. Namely, such rules can be considered as an inessential 'syntactic sugar' that can be uniformly reduced to ordinary, non-disjunctive causal rules. To be more exact, the following quite surprising result can be shown:

Theorem 4.1. *The stable nonmonotonic semantics of a disjunctive causal relation coincides with the nonmonotonic semantics of its normal subrelation.*

Recall that the the 'singular content' of a disjunctive theory Δ is fully described by the set of singular rules $N(\Delta)$, defined earlier. Accordingly, as an immediate consequence of the above result, we obtain the following

Corollary 4.2. *The stable nonmonotonic semantics of a disjunctive causal theory Δ coincides with the nonmonotonic semantics of a singular causal theory $N(\Delta)$.*

In other words, from the perspective of the stable semantics, any disjunctive causal theory is reducible to its singular counterpart.

5 Discussion

A general conclusion that could be made from the above results is that, from the point of view of a stable nonmonotonic semantics, any disjunctive causal rule is logically equivalent to a set of singular causal rules. In other words, disjunctive causal rules do not extend our representation capabilities in dealing with causation, even an indeterminate one.

It is a good moment to recall now that the stable semantics of disjunctive logic programs (see [Gelfond and Lifschitz, 1991]), a close relative of the stable causal semantics, is problematic in its treatment of indeterminate information. Namely, due to minimization that is involved in the definition of stable sets, the resulting stable semantics tends to give an exclusive interpretation to disjunctions. It turns out that the same shortcoming is preserved by the above stable semantics for disjunctive causal theories. The following, now well-known, example has been suggested by Ray Reiter:

Example. Suppose that we drop a pin on a board painted black and white. As a result, the pin lands on the board in such a way that it touches either black or white area, *or both*. Moreover, if the pin is large (or black and white areas are small), then a most probable effect is that the pin touches both black and white areas.

An apparently natural representation of this situation could be given using the main disjunctive causal rule

$$Drop \Rightarrow Black, White$$

and a couple of auxiliary causal assertions that need not bother us here. Unfortunately, the stable semantics of the resulting causal theory does not explain the world in which the pin touches both black and white areas, though it readily justifies worlds in which the pin touches only one of them. This shortcoming can now be made more vivid using our results. Thus, the above disjunctive rule is reducible to the following three singular rules:

$$Drop, \neg Black \Rightarrow White$$
$$Drop, \neg White \Rightarrow Black$$
$$Drop \Rightarrow Black \vee White$$

As has been shown in [Bochman, 2003], the nonmonotonic semantics of a causal production relation is determined only by Horn (determinate) causal rules that belong to it. In our case, the last rule is not determinate, and hence only the first two rules participate in determining the nonmonotonic semantics. Note, however, that these two rules cannot be applied simultaneously, due to the constraint $Drop \wedge \neg Black \wedge \neg White \Rightarrow \mathbf{f}$ that is implied by each of them. As a result, such rules cannot give an explanation for a quite possible fact $Black \wedge White$.

There are two main options that are conceivable in dealing with this situation; these options correspond to the two main components of our causal formalism. First, we have a logical (monotonic) formalism of disjunctive causal inference that determines not only the internal logic of causal rules, but also, and most importantly, an apparently plausible knowledge representation framework for describing indeterminate causal information. We could challenge this part of the formalism and claim that disjunctive causal rules do not give a direct representation of nondeterminate effects.

Example. (continued) Note first that some seemingly plausible improvements of the above representation still do not achieve the intended effect. For example, in order to allow for a possibility of $Black \wedge White$ output, we could attempt to use the following disjunctive rule[2]:

$$Drop \Rightarrow Black, White, Black \land White.$$

Unfortunately, the latter rule is logically reducible to the original rule with respect to disjunctive causal relations. Still, a disjunctive rule with an intended behavior can be given for this situation (see [Lin, 1996]); it amounts to the following:

$$Drop \Rightarrow Black \land \neg White, White \land \neg Black, Black \land White.$$

But in this case also we have that the rule is logically equivalent to a set of singular rules (see the example following Lemma 3.7).

Continuing this line of thought, an even stronger claim could be made, namely that we do not need disjunctive causal rules at all in order to represent indeterminate causal effects. Actually, one of the main contributions of [McCain and Turner, 1997] consisted in showing that we usually can represent an indeterminate causal information using ordinary, singular causal rules. Still, there is yet no systematic understanding whether we can always represent such an information in this way.

There exists, however, an entirely different way of dealing with the above problem. If we take for granted that disjunctive causal rules provide an adequate and natural representation of indeterminate causation, we have no choice but to question the adequacy of the stable nonmonotonic semantics for disjunctive causal theories. Actually, in the short history of disjunctive logic programming there have been many attempts to define a semantics that would preserve the usual inclusive understanding of logical disjunctions, such as Minker's Weak GCWA (see [Lobo *et al.*, 1992]) and Sakama's possible model semantics [Sakama, 1989]. A related idea has been pursued in [Eiter *et al.*, 1993], where the authors have suggested a semantics for circumscription in which the disjunction is interpreted inclusively. It remains to be seen whether a plausible nonmonotonic semantics of this kind could be given also to disjunctive causal rules. These are the questions that deserve further study.

Acknowledgment. Many thanks to Hudson Turner for extensive comments and discussions on the topics of this study, as well as to the anonymous referee for instructive suggestions.

References

[Bochman, 2003] A. Bochman. A logic for causal reasoning. In *Proceedings IJCAI'03*, 2003.

[Eiter *et al.*, 1993] T. Eiter, G. Gottlob, and Y. Gurevich. Curb your theory! a circumscriptive approach for inclusive interpretation of disjunctive information. In R. Bajcsy, editor, *Proceedings IJCAI-93*, pages 634–639. Morgan Kaufman, 1993.

[Gelfond and Lifschitz, 1991] M. Gelfond and V. Lifschitz. Classical negation and disjunctive databases. *New Generation Computing*, 9:365–385, 1991.

[Gelfond *et al.*, 1991] M. Gelfond, V. Lifschitz, H. Przymusińska, and M. Truszczyński. Disjunctive defaults. In *Proc. Second Int. Conf. on Principles of Knowledge Representation and Reasoning, KR'91*, pages 230–237, Cambridge, Mass., 1991.

[Giunchiglia *et al.*, 2001] E. Giunchiglia, J. Lee, V. Lifschitz, N. McCain, and H. Turner. Nonmonotonic causal theories, 2001. *Artificial Intelligence* (to appear).

[Lin, 1996] F. Lin. Embracing causality in specifying the indeterminate effects of actions. In *Proceedings AAAI-96*, pages 670–676, 1996.

[Lobo *et al.*, 1992] J. Lobo, J. Minker, and A. Rajasekar. *Foundations of Disjunctive Logic Programming*. MIT Press, Cambridge, Mass., 1992.

[McCain and Turner, 1997] N. McCain and H. Turner. Causal theories of action and change. In *Proceedings AAAI-97*, pages 460–465, 1997.

[Sakama, 1989] C. Sakama. Possible model semantics for disjunctive databases. In W. Kim, J.-M. Nicolas, and S. Nishio, editors, *Proc. 1st Int. Conference on Deductive and Oject-Oriented Databases*, pages 369–383, N.Y., 1989. Elsevier.

[Turner, 1999] H. Turner. A logic of universal causation. *Artificial Intelligence*, 113:87–123, 1999.

[2] Suggested by the anonymous referee.