Notre Dame Journal of Formal Logic
Volume 41, Number 2, 2000

AN INFORMATION-BASED THEORY OF CONDITIONALS

WAYNE WOBCKE

Abstract We present an approach to combining three areas of research which we claim are all based on information theory: knowledge representation in Artificial Intelligence and Cognitive Science using prototypes, plans, or schemata; formal semantics in natural language, especially the semantics of the 'if-then' conditional construct; and the logic of subjunctive conditionals first developed using a possible worlds semantics by Stalnaker and Lewis. The basic premise of the paper is that both schema-based inference and the semantics of conditionals are based on Dretske's notion of information flow and Barwise and Perry's notion of a constraint in situation semantics. That is, the connection between antecedent A and consequent B of a conditional 'if A were the case then B would be the case' is an informational relation holding with respect to a pragmatically determined utterance situation. The bridge between AI and conditional logic is that a prototype or planning schema represents a situation type, and the background assumptions underlying the application of a schema in a situation correspond to channel conditions on the flow of information. Adapting the work of Stalnaker and Lewis, the semantics of conditionals is modeled by a refinement ordering on situations: a conditional 'if A then B' holds with respect to a situation if all the minimal refinements of the situation that support A also support B. We present new logics of situations, information flow, and subjunctive conditionals based on three-valued partial logic that formalizes our approach, and conclude with a discussion of the resulting theory of conditionals, including the "paradoxes" of conditional implication, the difference between truth conditions and assertability conditions for subjunctive conditionals, and the relationship between subjunctive and indicative conditionals.

1. Information Flow and Schema-Based Inference

The aim of this paper is to combine three strands of research, in artificial intelligence, semantics, and logic. One basic question concerns the semantics of schemata or other

Received November 27, 2000; printed August 30, 2002
2001 Mathematics Subject Classification: Primary, 68T27
Keywords: conditional logic, situation semantics, knowledge representation

similar knowledge structures such as prototypes (Rosch [56]), frames (Minsky [45]), and scripts (Schank and Abelson [59]), representations widely used in artificial intelligence. Just what do schemata represent? To answer this question, we draw on the theory of situation semantics (Barwise and Perry [10]). The main feature of situation semantics that is useful for interpreting schemata is the partiality of situations and situation types. Two distinct senses of partiality are involved: (i) partiality in the sense that a situation need not decide every fact (a schema represents only part of the world), and (ii) partiality in the sense that situation types may contain object and event types (serving as the interpretations of the variables or placeholders in a schema): this kind of partiality arises from abstraction.

A knowledge-based system typically contains not just one schema but a collection of schemata usually organized in a hierarchy. This raises the question of what the hierarchical relationship between schemata represents. To answer this question, we use the notion of a *constraint*. In situation semantics, constraints between event types model meaningful relations: a constraint '*A involves B*' holds if whenever some situation supports A, there must be some related situation that supports B. In this paper, schema-based inference is modeled using constraints. However, constraints are not formalized directly, but rather as arising from a more basic hierarchical relationship between situation types that we call the *refinement* relation. This is meant to capture the idea that a schema is used in reasoning as an idealization or abstraction of a complex state of affairs, and as such, its application relies on a collection of background assumptions. As a reasoner moves through the hierarchy to more and more specific schemata, the set of background assumptions underlying the applicability of the schemata change. These background assumptions are intended to characterize the context of application of the schema: schema-based reasoning is situated reasoning, or reasoning in context. On our approach, the conditionals accepted by the agent in a situation are determined by the minimal refinements of the situation: a conditional 'if A then B' is accepted if all the minimal refinements of the situation that support A also support B. As the reasoner moves about the hierarchy, the refinements of the situation change, and the conditionals accepted also change.

A more detailed intuitive picture is as follows. At any time, the agent is in some situation that it classifies as being an instance of some situation type σ. Any conditionals accepted are determined by *applying* the schema to the present situation: applying a given schema to a given situation means grounding the variables in the schema to particular objects in the situation and then copying part of the schema hierarchy with root at the given schema (with variables grounded) to form a new hierarchy of situations rooted at the given situation. We will say that the situation in this way *inherits* the constraints from the situation type of which it is classified. Furthermore, when the agent learns a new fact A, it should reclassify the situation using one or more refined situation types than σ that support A; as a result, the constraints and hence the conditionals accepted by the agent will change. For example, suppose the agent is trying to start its car and classifies its current situation as being of type σ that supports two event types (and the constraint) between turning an ignition key and the car's starting. The agent accepts the conditional 'if I were to turn this key then this car would start'. Now suppose the agent learns that the car's battery is dead. The agent should reclassify the situation using a more specific type that represents the constraint that turning an ignition key in a car with a dead battery means the car won't start. The agent then accepts conditionals such as 'if I were to turn this key then this car wouldn't start' and 'if I were to charge this battery then this car would start'. The hierarchy of schemata thus implicitly represents a collection of constraints with varying background contexts of applicability.

The situation semantics notion of constraints comes from Dretske's theory of information (Dretske [21]): a constraint '*A involves B*' primarily reflects the property that a signal A carries the information B. In Dretske's theory, information is viewed as a commodity that can be transmitted over channels via signals and extracted from those channels by agents attuned to the information they carry. Dretske's central examples are real physical channels like the channels on a television set that carry information about what occurred prior to transmission. What is important about constraints is that the flow of information is regulated by lawlike connections: accidental correlations cannot convey information. As a simple example, even when it is 4 p.m., a broken watch that always shows the time as being 4 p.m. does not convey the information that it is 4 p.m. because the watch would have displayed this anyway (even when it was not 4 p.m.). That is, the channel of communication has to constitute a reliable indicator of the state of affairs obtaining at the source of the transmission for information flow to occur. Thus the flow of information along a channel is contingent on a set of background "channel conditions" which together ensure the veracity of the information transmitted. This means that information flow, like schema application, is "situated" or dependent on context. Although this seems at first sight to conflict with the condition that constraints are lawlike, constraints are still underwritten by causal laws: it is just that the scheme of such laws is thought of as set against an implicit background.

Real physical channels are far removed from schema application in knowledge representation systems. However, we claim that the relevant features of physical channels do apply to schema-based inference and default reasoning. Let us illustrate this claim with perhaps the two most overused examples of the *bird* object type in a type hierarchy and the *restaurant* script in a planning hierarchy. The schemata are used to infer that Tweety flies or that the customer paid for lunch. Now the distinctive mark of information flow is its lawlike (or nomic) nature, the fact that the signal *must* carry the information. To make sense of these schema-based inferences, we take it that the collection of schemata represent channels consisting of lawlike connections between properties or events. So there is a "channel" connecting the property of being a bird and the property of flying. Although Tweety's being a bird may not cause its flying, the relationship between Tweety's flying and its being a bird is not an accidental one (if any other object were a bird, it would also fly). Similarly, in the case of inference based on plan schemata, although Schank and Abelson [59] explicitly define a script as representing a network of causal relationships between events, these causal relations include relations such as enablement or motivation which are not notions of causality used to describe physical channels. Nevertheless, there is a sense in which paying is related to eating lunch at a restaurant in a lawlike rather than accidental way (the customer *has* to pay at the restaurant in order to fulfill social obligations, and anyone else going to the restaurant would also have to pay).

Inferences made using schemata differ from inferences based on information flow in one very important respect. This is that inferences made using schemata are defeasible, that is, subject to review as further information is obtained: the inference that Tweety flies has to be retracted when it becomes known that Tweety is a penguin. But on Dretske's definition of information, no signal can carry the information that p

unless p is true (no signal can carry the information that Tweety flies unless Tweety actually does fly). Thus the status of an inference as information entails, additionally, the property of nondefeasibility. This means that the logic of information flow is different from the logic of default inference, analogous to the way in which the logic of knowledge differs from the logic of belief. However, an agent is usually in no position to know whether or not it has information, so for an agent any information-based inference is just as provisional as any default inference.[1] It is this property that enables schema-based inference, though defeasible, to be regarded as being supported by constraints. To fix terminology, we call an inference *information-based* if it has the status of information and hence is nondefeasible; we call an inference *constraint-based* if it is based on a signal flowing along a channel that does not necessarily have the status of information. This terminology is justified in part by Dretske's observations on Barwise and Perry's use of constraints to model linguistic conventions: linguistic conventions can be violated whereas informational relations cannot (Dretske [22]).

So far, we have argued that schema-based inference is a special type of defeasible inference based on information flow along channels whose background supporting conditions are unknown to the agent. As noted above, information flow is supported by lawlike regularities, but whether information flows along a channel can only be determined relative to these background channel conditions. Let us now emphasize that even though it initially appears that Dretske's theory of information flow is narrowly applicable only to real physical channels, he clearly intends the theory of information to be applied far more widely: in fact, for Dretske [21], information flow can be relative to agents' beliefs, goals, and purposes, as is required for their application in the kinds of social settings represented by scripts, for example. To explain this, first recall that information flows along a channel only if there is a lawlike relationship between the signal and the information carried. This means that information flows in some situation only if there are no *relevant* alternative situations in which the signal occurs without the presence of the purported information. Fixing the channel conditions amounts to delineating the space of relevant alternative situations, and then the question of which information is conveyed becomes an objective and absolute matter. However, the key point is that the concept of relevance used to fix the background is open to the "interests and purposes of people applying it" ([21], p. 133) which makes the concept of information similarly relative to agents' interests and purposes. The important consequence for our discussion is that a reasoning agent which needs to cope with varying sets of channel conditions, and moreover, does not know whether those channel conditions hold in any given situation, has the same beliefs and dispositions to change beliefs as an agent which attempts to facilitate reasoning by applying a collection of schemata whose context of application it may not know, in situations in which it cannot be certain which schemata are applicable. If the agent itself needs to handle shifts in context, its inferences must be defeasible, regardless of whether or not those inferences are based on information.

Thus, we conclude, schema-based inference is constraint-based inference, or inference based on signals flowing along channels that result in information if the channel conditions hold. This provides a way to understand knowledge representation systems based on hierarchies of schemata as representing constraints as in situation semantics, and it is this connection we explore in the remainder of this paper. In Section 2 we present some initial considerations on the logic of conditionals and conclude that schema-based inferences have the right logical properties to be a basis for understanding subjunctive conditionals, extending the argument in Wobcke [69]. In Section 3 we give an overview of situation semantics before presenting a new formalization of information flow using a three-valued logic which borrows from modal logic, described in Section 4. In brief, a statement expressing an informational relation is treated as a conditional such as 'if the doorbell is ringing then someone is at the door'. The lawlike nature of the connection between signal and information received is part of the analysis of the conditional. We interpret the conditional with respect to a situation which, following Dretske's intuition, comes with a set of relevant alternative situations. The conditional is analyzed as a kind of strict implication: A carries the information B with respect to a situation if there is no relevant alternative in which A holds but B does not. The logic of information flow characterizes the kind of inference that is not defeasible.

In Section 5 we present a conditional logic of constraints that models defeasible reasoning. Defeasible inference is construed as the result of a schema's refinement providing a more accurate world view which disagrees with its less refined parent, or equivalently, as the result of refining the set of channel conditions underlying information flow. Similar shifts in context are an important aspect of the analysis of conditionals presented by Stalnaker [63] and Lewis [39]. In these semantic theories, the context of evaluation of a subjunctive conditional 'if A were the case then B would be the case' is a possible world (the actual world), and the conditional is interpreted as a statement to the effect that B holds in those possible world(s) in which A holds that are most similar to the actual world. The interpretation of conditionals is similarly sensitive to background assumptions which are context dependent, the relativity to context being part of the definition of which worlds are most "similar" to a given world, as argued by Lewis ([39], pp. 91–95). The analogy between conditionals and information flow and the dependence of conditionals on a background context modeled as channel conditions has been used by Barwise [7] to develop a formal theory of conditionals. Barwise's account does not include a logic of conditionals, where by this we mean a characterization of those inferences valid in any context; Cavedon [15] and Restall [54] develop conditional logics based on channel theory. In contrast, we formalize a conditional logic using the refinement ordering on situations in place of Lewis's comparative similarity relations on possible worlds. A conditional 'if A were the case then B would be the case' holds with respect to a situation σ if B holds in all minimal refinements of σ that support A. The axiomatization of the logic is determined by the conditions placed on hierarchies of situations. Here, whereas the situations themselves are partial (in the sense of allowing indeterminacy on the question of whether a fact holds), we shall motivate the condition that all situations are comparable in terms of the refinement relation, so that the hierarchy of situations forms a total preorder. Thus our conditional logic is closely related to Lewis's.

In Section 6 we discuss a number of standard puzzles in conditional reasoning. First, since a hierarchy of schemata determines a hierarchy of situation types, the "paradoxes" of conditional implication (failure of strengthening the antecedent, failure of transitivity, and failure of contraposition) arise naturally from the fact that the hierarchy of schemata utilizes defaults and exceptions. Equivalently, these paradoxes are a consequence of how different channel conditions serve to undermine the flow of information with respect to varying backgrounds. The use of schemata also enables us to give a unified (pragmatic) account of factual and counterfactual conditionals.

Second, we consider a number of examples and counterexamples in conditional reasoning proposed in the literature on nonmonotonic reasoning and belief revision: we argue that our logic can handle the examples correctly. We then discuss the difference between subjunctive and indicative conditionals. Subjunctive conditionals are viewed as the expression of a constraint that holds over situation types, closely connected with information flow, whereas indicative conditionals are viewed as assertions of conditional information relating to specific facts about the world (and which are thus sensitive to those facts and the evidence for those facts in addition to general constraints). We explain the different truth conditions of subjunctives and indicatives in terms of the differences in both the context of evaluation and in the particular hierarchy of situations types used in their evaluation. We explain why subjunctives and indicatives are often interchangeable and suggest that they are just when the truth of the indicative derives only from a general constraint (and not from any additional specific facts about the world). Finally, we discuss some examples from the literature on conditionals: Quine's Bizet-Verdi example concerning the conditional excluded middle and 'might' conditionals.

2. Conditionals and Constraints

In this section we argue that hierarchies of schemata that represent a refinement relation on situation types and hence, implicitly, a collection of constraints, can be taken as underlying subjunctive reasoning. Our work builds on that of Barwise who has argued that conditionals can be modeled using constraints. We have claimed above that schema-based inference can also be treated as a kind of constraint-based inference: in this section, we consider issues specific to the interpretation of conditionals.

Our theory of conditionals is a truth-theoretic account in which a conditional 'if A were to be the case then B would be the case' holds with respect to some situation if that situation inherits a constraint $A \Rightarrow B$ from some situation type of which it is an instance. However, we use a different model of constraints from that of Barwise, who builds into the formalization of the conditional some associated (unspecified) set of background assumptions. Under our approach, a constraint is represented using a hierarchy of situation types (represented by schemata), and the constraints that are applied in a given situation are derived from a "base" situation type in the hierarchy that is assumed to be determined partly pragmatically. The position of the base situation type in the hierarchy implicitly determines the background context against which the conditional is evaluated.

We have asserted that the selection of base type (and hence background assumptions) for interpreting a conditional is determined partly pragmatically but this does not mean that it is completely arbitrary. One criterion commonly used in schema-based inference is a principle of *specificity*, according to which, when two comparable sets of background assumptions are both applicable to a situation, the more refined set of assumptions should be preferred to the less refined set. Specificity says, for example, that when we know Tweety to be a penguin, we should choose the penguin schema as the base type when evaluating a conditional such as 'if Tweety had a three-foot wingspan then it would be able to fly'. Another principle from conditional logic is *minimality*, according to which only those background assumptions that are necessary to the interpretation of the conditional should be countenanced. For example, when evaluating a conditional such as 'if Oscar had wings then he would be able to fly'—when Oscar is a person (say)—a context in which Oscar is a bird should be adopted but no more specific context (it would not do to assume that Oscar was a penguin in evaluating the conditional). Of course, specificity and minimality do not uniquely determine a base type for the evaluation of a conditional, especially in view of that fact that they pull in opposite directions. Consider the Tweety example applied to a fish, that is, the statement ‘if that fish were a bird then it would be able to fly’. It may not be entirely irrational to suppose that if the fish were a bird, it would be a penguin and hence not be able to fly (perhaps the fish is black and white in color and bears a resemblance to a penguin). The point is that different intuitions about particular conditional statements can sometimes be traced to different base types used in their evaluation.

Thus we agree with Stalnaker and Lewis insofar as there must be some pragmatic ambiguity in the interpretation of conditionals, but we differ on precisely where this ambiguity is located. In Lewis’s analysis of conditionals, the context dependence results from the criteria for determining the comparative similarity relation, that is, the relation which defines whether one world is more similar than another to a given world. The point of evaluation of a conditional is always the actual world. So with the Tweety examples above, when Tweety is a penguin in the actual world, this world is the closest world, while presumably a world in which a person, Oscar, is a penguin is less similar to the actual world than a world in which Oscar is a flying bird, although exactly why this should be so is not obvious without an appeal to specificity. The idea of “similarity,” however, should not be taken too literally as it leads to some counterintuitive consequences as discussed by Pollock [49]. In fact, Lewis [40] states that the appropriate comparative similarity relation is determined in part by the known truth of conditionals, so although comparative similarity serves as an analysis of conditionals, it does not ultimately serve as an explanation of them. Our strategy is to reduce the semantics of (a subclass of) conditionals to the application of constraints that hold in the actual world and to confine the pragmatic aspects of conditionals to the determination of the context of evaluation and hence to the determination of those constraints.

One other point that must be clarified is the relationship between general constraints (as represented in a hierarchy of schemata) and specific constraints that are inherited by a particular situation through the application of a schema. The difference can be seen linguistically by rendering a general constraint using a generic statement such as ‘if a bird has wings then it can fly’ (following Cavedon [14]) while rendering a specific application of a constraint using a statement such as ‘if Tweety has wings then Tweety can fly’. In addition, when applying a schema to a situation to create a subhierarchy of situations, we assume that any facts supported by the situation that are consistent with the instantiated schemata remain supported by the situations in the new hierarchy. This allows “irrelevant” facts to be the subject of true conditionals. For example, suppose that in the car-starting situation described above that the car is red. We take it that the car is still red in any more refined situation created by inheriting constraints about starting the engine. This indicates a difference between the truth conditions for specific subjunctive conditionals and those for generics. That is, while the generic ‘if I turn a key in an ignition of a car, the car would be red’ is false with respect to a base situation type, the specific statement ‘if I turn this key in the ignition of this car, this car would be red’ is true in the above situation. In such a way, a subjunctive conditional can be true because of the *absence* of a constraint blocking information flow. This is another aspect of the minimality principle in the

Stalnaker-Lewis semantics of conditionals: the worlds supporting the antecedent of a conditional should differ "minimally" from the initial world.

Constraints capture informational relations: the constraint $A \Rightarrow B$ in situation semantics holds if and only if the signal A carries the information that B, where A and B are fact or event types. For Dretske, an informational relation between A and B is underpinned by a causal relation between A and B although not every causal relation is informational. A favorite example of Barwise and Perry illustrates this point. Suppose a coin is tossed and lands heads. Then there is a causal relation between the tossing of this coin and its landing heads. However, for an informational relation to hold, there must be a generally applicable lawlike condition to the effect that a coin could not have been tossed without its landing heads. Clearly, there is no such general constraint: sometimes a coin lands tails. The correct constraint is that a coin's being tossed carries the information that it lands heads or tails (presumably the background assumptions are held fixed while comparing these possible constraints). According to Dretske, informational relations are determined by considering the network of possible relations between source and receiver, not just single events. A similar point regarding the interpretation of conditionals is made by Pollock who calls the relation between antecedent and consequent required for a conditional to be true one of contingent necessitation.

Conversely, of course, not every informational relation is causal. Thus, again following Dretske's intuitive picture, informational relations can connect any two event types in a causal chain. This means that informational relations can go "backward in time" or relate correlated consequences of a single cause. For example, there are presumably constraints such as 'if a coin lands heads or tails then it was tossed' and 'if a person has lung cancer then she or he has emphysema' (assuming in the latter case that both are the result of smoking, rather than one causing the other).

The essence of our approach is to generalize Dretske's theory of information flow as deriving from causal chains to hierarchies of causal chains under varying background conditions. Causal chains are represented in artificial intelligence by scripts or partial plans: thus it is possible for an agent to represent constraints using a hierarchy of scripts or partial plans. On our account, a constraint $A \Rightarrow B$ holds at some situation type σ if B holds at the minimal refinements of σ that support A. For example, suppose we have causal chains representing the connections between tossing a coin and its landing heads and tossing a coin and its landing tails, that is, situation types supporting *toss(coin)* and *land(heads)* and *toss(coin)* and *land(tails)*. If we assume that neither situation type is strictly less refined than the other, both will be minimal refinements of the original situation type that support *toss(coin)*. Thus the following conditional will be true.

> If a coin is tossed, it will land heads or tails.

However, both of the following conditionals will be false.

> If a coin is tossed, it will land heads.
> If a coin is tossed, it will land tails.

When choosing this least refined context as the context of evaluation for the conditional, it is implicitly assumed that there are no other relevant background facts about the particular coin, for example, it is assumed to be a fair coin.

We now discuss three aspects of the semantics of conditionals which make hierarchies of schemata representing causal chains (and hence hierarchies of situation types) suitable for deriving their truth conditions. First, a conditional is not considered true in isolation of other related true conditionals: a single causal chain may account for the truth of a number of conditionals. Second, where the truth of a conditional relies on an implicit background, a hierarchy of schemata implicitly represents this background context as a collection of (unspecified) assumptions: if such an assumption is incorrect, a conditional that is true in one context may be false in a more refined context. Third, conditionals often come in pairs, one which looks factual and one which looks counterfactual (the first statement with both antecedent and consequent negated): we claim that the same causal chain underlies both conditionals.

The first argument for hierarchies of schemata comes from Hanson's discussion of the explanation of everyday events in relation to causal statements (Hanson [34]). Hanson claims that there are as many true causal statements (and so true conditionals) as there are explanations, and that these explanations are dependent on the point of view adopted. Thus, of a car crash, the following statements are asserted.

> The car crashed because the driver tried to avoid the pedestrian.
> The car crashed because the brakes failed.
> The car crashed because the tires skidded on a patch of ice.

Here it is assumed that in trying to avoid the pedestrian, the driver applied the brakes which failed to grip the icy road. All the conditions and events play some role in the crash, but each causal statement isolates one of them. The point is that each causal statement is true, so the corresponding counterfactuals are all true.

> If the driver hadn't tried to avoid the pedestrian, the car wouldn't have crashed.
> If the brakes hadn't failed, the car wouldn't have crashed.
> If the tires hadn't skidded on a patch of ice, the car wouldn't have crashed.

To handle this phenomenon, we assume that the set of causal statements forming the complex explanation is represented in the schema representing one complex causal chain. All the conditionals are true because the schemata which are used in evaluating the conditionals can be derived from this one schema in a straightforward manner (see the discussion on failure schemata below).

The second argument for hierarchies of schemata is that hierarchies implicitly allow for default assumptions and exceptions. The effect of defaults is realized in the theory of conditionals by the failure of strengthening the antecedent. For example, the first conditional below is true, the second false (in an appropriate context).

> If the battery had been charged, the car would have started.
> If the battery had been charged and disconnected, the car would have started.

We interpret this as follows. In the first example, there is an implicit assumption (a "channel condition") that the battery is connected. This assumption is denied in the antecedent of the second conditional and the causal chain between antecedent and consequent is broken. We assume that there is a schema representing the causal connection between a charged battery and the car starting that is less refined than one representing the causal connection between a charged battery that is disconnected and the car not starting. Thus the least refined situation type that supports the battery being charged also supports the car starting, but the least refined situation type that supports the battery being charged and being disconnected does not support the car starting. This kind of example is also used to explain the failure of transitivity and contraposition.

On the theories of conditionals developed by Stalnaker and Lewis, the next statement follows from the above two statements.

> If the battery had been charged, it would not have been disconnected.

Intuitively, the negation of any condition that blocks a constraint is an explicit channel condition for that constraint. Similar examples arise in the AI literature on default reasoning: the condition that Tweety is not a penguin is one of the assumptions underlying the applicability of the bird schema, implying the truth of the default statement 'a bird is typically not a penguin'. We argue below that this condition is too strong for a proper modeling of constraints. Hence we do not allow this inference, agreeing with authors such as Pollock.

Third, conditionals often come in pairs.

> If the battery had been charged, the car would have started.
> If the battery hadn't been charged, the car wouldn't have started.

Both statements are taken to be true, although we argue that for this to be the case, their contexts of evaluation must be different. In particular, it seems that part of the reason both statements are acceptable is the pragmatic assumption that subjunctive conditionals are normally counterfactual. That is, the first statement is normally uttered in a context in which the battery has not been charged and the car not started, whereas the second in a situation in which the battery has been charged and the car started. However, there is a clear connection between the two statements: they both entail that the dead battery (as opposed to something else) is one reason the car didn't start (the first statement, in addition, implies that it is the only reason).

To represent such pairs using hierarchies of schemata (and the above car accident examples), we must further interpret what a plan schema actually means in cases where one of its represented events fails to occur. In doing so, we follow Mackie [42] who argues that each cause of some event is a necessary component of a collection of events and/or conditions which together are sufficient for the occurrence of the event (but not collectively necessary, for the event could have occurred in other ways). Now for a causal chain containing a dependence of B on A, if A is a necessary component of a collection of actions (all those A such that B depends on A) which together are sufficient for B to occur, when A does not occur, B cannot. We assume that such "failure schemata" appear in the hierarchy of schemata as refinements of their parents: many knowledge-based systems explicitly encode failure schemata in this way. Now to handle the examples, we take it that the context of evaluation for both statements is a causal chain containing a link between a battery being charged and a car starting (among others). The minimal refinement supporting the condition that the battery has been charged is the schema itself (giving the truth of the first statement) whereas the minimal refinement supporting the antecedent of the second conditional is a failure schema (giving the truth of the second statement).

Of course, for the purpose of these examples, we are assuming that charging the battery does play a causal role in the car's starting (so that both statements are true in their respective contexts). If charging the battery does not play this role (the battery's charge was sufficient to begin with), a different schema, one representing an alternative "failed" causal chain, will be the minimal refinement in which the second conditional's antecedent is supported, and the conditional will be false. This raises a very important point. That is, the particular hierarchy of situation types, and in addition, the particular refinement relation representing the truth conditions of subjunctive conditionals, is itself context dependent. It is sensitive to the actual causal relationships holding in the described situation, not only to general causal laws. This is analogous to the context dependence of the comparative similarity relation in Lewis's account of conditionals [39], albeit with a different underlying intuition. Thus there is a difference between the true subjunctive conditionals in a situation and those accepted by an agent.

Our approach can briefly be contrasted with the recent formulation of channel theory by Barwise and Seligman [11] as follows. Most obviously, while both the present theory and channel theory aim to formalize information flow and inference based on constraints, our approach starts with Dretske's intuitive picture of how information flow arises out of a network of causal relations, whereas channel theory starts by assuming the existence of "channels" that realize particular information-carrying signals between senders and receivers. Each channel thus corresponds to a single inference or possibly to a class of single inferences of the form 'if the received signal is of type σ_r then the original signal was of type σ_s. The basic problem in modeling inference in channel theory then is to define theories about relationships between channels, such as when two or more channels can be combined, when one channel subsumes another, and so on. However, rather than relate this question to Dretske's background conditions on information flow, Barwise and Seligman assume that the allowable combinations of channels are derivable from a structure of classifications (of signal tokens into types). But although they recognize the importance of handling exceptions to channel conditions, they do not realize that allowing any two channels that "line up" ([11], p. 91) to be sequentially composed leads to incorrect instantiations of transitivity (if A carries the information B and B carries the information C then A carries the information C). These violations of transitivity occur when the signal type A of the first channel is inconsistent with the background conditions on the channel from B to C. For example, let A be 'the doorbell is short circuited', B be 'the doorbell is ringing', and C be 'someone is at the door'; that B carries the information C is one of Dretske's standard examples. Then, contrary to Barwise and Seligman, it seems intuitively reasonable to accept that A carries the information B, B carries the information C, but A does not carry the information C. Thus transitivity should be rejected by any theory of information flow. Note here that the rule of transitivity, which operates at the level of types, should not be confused with Dretske's Xerox principle [21] which operates at the level of tokens: more precisely, the Xerox principle is that when a token of A carries the information B and that token of B carries the information C, the original token carries the information C. Our example is not a counterexample to this principle because when the token is of type A (and hence of type B), the token of B does not carry the information C since this relies on the background condition $\neg A$.

Channel theory, in addition to being a specific theory of information flow, provides a broad framework in which various logical questions can be considered. It is possible to avoid the specific problem with transitivity in Barwise and Seligman's theory by adopting the approach of Cavedon. Cavedon [15] develops a theory in which informational relations represented as channels are ordered in a hierarchy resulting in a conditional logic similar to, but weaker than, the one provided in this paper. Cavedon's logic is also restricted to nonnested conditionals. Thus the main advantage we claim for our approach is not in formal expressive power, although we do allow nested conditionals, but in representational convenience. We claim that the use of schemata to group together informational relations with common background assumptions leads to more efficient reasoners: this is why such representational schemes are used in AI systems in the first place.

Our basic conclusion from the above discussion is that not only can schema-based inference be taken as constraint-based inference, but that constraints as represented in hierarchies of schemata are appropriate for modeling the semantics of subjunctive conditionals. We can now turn to a logical formalization of this semantics. We first present a theory of situations and information flow, after motivating our theory by consideration of various alternative approaches in situation semantics. We then adapt the Stalnaker-Lewis semantics of conditionals to provide a modeling of hierarchies of situations. A sound and complete axiomatization enables us to compare our logic to Pollock's logic SS and Lewis's logic VC. We then consider some further issues in conditional reasoning and how they might be approached using our framework.

3. An Overview of Situation Semantics

A situation is a portion of space-time consisting of individuals, objects, properties, relations, events, and space-time locations. Using the terminology of Devlin [20], situations resolve *issues*: an issue is an ascription of a property to an object or a claim that objects stand in some relation in some location or throughout some time period. If an issue is resolved in a situation, it is resolved with some *polarity*: either positive or negative. Issues with polarity are called basic *infons* (the basic units of information). Compound infons may be formed as logical combinations of basic and other infons, analogous to complex propositions formed using the standard conjunction, disjunction, negation, and conditional operators. If an issue is resolved positively in a situation, the situation *supports* the corresponding infon; if negatively, the situation *rejects* the infon. The question of whether an infon holds at a situation is an entity that can be true or false and is the analogue of the classical proposition. Such propositions can also be combined using the standard conjunction, disjunction, negation, and conditional operators which, however, should not be confused with the operations on infons. The logic of propositions is usually taken to be the standard propositional calculus.

Situations are contrasted with possible worlds, although some authors, for example, Cresswell [16], have argued that the notion of a possible world is broad enough to include situations. The main difference between situations and possible worlds (as understood by Kripke [36] or Lewis [41]) is that although every proposition must be assigned a truth value in every possible world, a situation may leave some issues unresolved. In Barwise and Perry [10], a distinction is also made between actual situations and abstract situations. Put simply, an actual situation is an actually existing part of the world, while an abstract situation is a set-theoretic object specified by a set of infons which may or may not be supported by any actual situation (more like a model of the propositional calculus than a possible world). So if there is a situation in which John eats an apple (here and now), an abstract situation can be formed in which Mary eats an apple at home on Tuesday, irrespective of whether this actually happens. Barwise and Perry also define the notion of a situation type. A situation type is a set of issues obtained by abstracting from specific space-time locations. Intuitively, it is what two different parts of the world have in common when they resolve the same issues in the same way except for the difference in their locations. For example, there is a situation type of John eating an apple and Mary eating an orange. Situation types, like situations, do not necessarily resolve every issue. Of course, there is no formal reason that only space-time locations can be abstracted in forming situation types from situations: in more recent versions of situation theory (for example, Seligman and Moss [61]), any parameter in a situation may be abstracted—that article also presents a much richer theory of situations than that developed in [10].

Another fundamental notion is that of a constraint. A constraint is a relation between event types. Just as situation types are formed by abstraction from situations, event types are formed by abstraction from events, but importantly, in a more general way. With an event type, in addition to the space-time location, the objects and individuals taking part in the event may also be abstracted, giving rise to a parametrized infon with indeterminates filling the roles in the event. For example, an event such as John eating an apple here and now can be abstracted to form the event type of John eating something, somewhere, some time. A constraint is a relation between event types that is designed to capture the meaningfulness of signs, both natural and nonnatural in the terminology of Grice [32]. Roughly, a constraint $A \Rightarrow B$ is to be interpreted as a condition to the effect that whenever some situation supports an event of type A, some related situation supports an event of type B where corresponding parameters in A and B are anchored to the same objects. For example, there could be a constraint in which A is the event type of the doorbell ringing and B is the event type of the someone being at the door. As this example is meant to illustrate, constraints are what makes information flow possible. More precisely, information about B is carried by A if there is a constraint $A \Rightarrow B$ provided all the background or channel conditions upon which that constraint depends hold.

We are now in a position to give a brief, necessarily selective, summary of the development of situation theory which we will use below to place our approach in the context of this research program. In [10], situations and situation types are modeled as sets of basic infons and constraints are modeled as relations between events and sets of event types. However, this allows only a limited treatment of disjunction: in order to express a constraint of the form 'A carries the information that B or C' (such as the coin-tossing example above), a relation between an event A and a set of events types $\{B, C\}$ is needed. Disjunction is thus implicit in that there is no "disjunctive" piece of information $B \vee C$ and hence no way to combine information acquired using this constraint with other information.

Questions of the correct modeling of constraints and of the admissible ways of compounding infons led Devlin [20], and especially Barwise and Etchemendy [9], to consider the algebraic structure of the class of infons. Devlin simply allows, for any two infons, their conjunction and disjunction to be formed as a compound infon: a situation supports a disjunction if it supports either disjunct and rejects a disjunction if it rejects both disjuncts; a situation supports a conjunction if it supports both conjuncts and rejects a conjunction if it rejects either conjunct. Devlin's approach here borrows from that employed in partial logic (see the survey article by Blamey [12]) and this is the approach we follow below.

Barwise and Etchemendy give a more general treatment of the algebraic structure of infons. In their framework, compound infons can be formed from other infons using conjunction, disjunction, negation, and conditional operations. In classical logic, the set of propositions so constructed forms a Boolean algebra. That is, the class of propositions has a lattice structure (with conjunction the meet operation of the lattice and disjunction the join operation) and a complementation operation corresponding to negation in the logic. The ordering relation $\leq$ on the lattice reflects the entailment of propositions, that is, $A \leq B$ if and only if $A \models B$. The simplest kind of lattice structure suitable for modeling infons is the distributive lattice, which has only the meet and join operations that distribute over each other. This structure is sufficient to model the infon algebras from [20] which we also use.

As Barwise and Etchemendy argue, the lattice of infons is a distributive lattice but not a Boolean algebra. For example, the infon $A \vee \neg A$ is not equal to the $\top$ element of the lattice, because intuitively, the information $A \vee \neg A$ is stronger than no information: it implies, for example, that A is defined. They postulate that the lattice of infons should form a Heyting algebra, defined as follows.

Definition 3.1 A *Heyting algebra* $\langle L, \wedge, \vee, ^{-} \rangle$ is a distributive lattice $\langle L, \wedge, \vee \rangle$ containing elements $\bot$ and $\top$ with a binary pseudo-complementation operation $\rightarrow$ such that the following properties are satisfied for all $a, b \in L$.

(i) $a \wedge (a \rightarrow b) = b$;
(ii) for any $c \in L$, if $a \wedge c \leq b$ then $c \leq a \rightarrow b$.

Note that for any given infons a and b, the infon $a \rightarrow b$ is a new piece of "conditional" information: the weakest piece of information that produces the piece of information b when combined with a. The infon $\neg a$ can be defined as the weakest piece of information which when combined with a gives $\bot$. An important technical point is that the basic infons are not here assumed to be issues with an assigned polarity, that is, an issue with negative polarity is not equal to the negation of the original issue with positive polarity. As a consequence, the infon $\neg\neg a$ is not always equal to the infon a, as it is under Devlin's approach. For example, for an element b not equal to $\top$ or $\bot$, the infon $\neg(b \vee \neg b)$ is equal to $\bot$, but its negation $\neg\neg(b \vee \neg b)$ is equal to $\top$ which is different from $b \vee \neg b$ (in general, $a \leq \neg\neg a$ but not the reverse).

In classical logic, propositions can be modeled formally as sets of possible worlds: a proposition A corresponds to the set of worlds in which A is true. In this way, the algebraic structure of sets of worlds is also a Boolean algebra with the meet, join, and complementation operations as set-theoretic intersection, union, and complementation. Now to mimic this property with sets of situations, Barwise and Etchemendy take a situation to be a closed coherent set of basic infons. A number of possible definitions of coherence are suggested, although the simplest one is sufficient: a situation is coherent if it does not both support and reject the same infon. Closure operations on sets of infons are used to model (deterministic) constraints. A set C of infons is closed if, for any constraint $A \Rightarrow B$, whenever C contains A, C contains B. What results is a Heyting algebra over closed coherent sets of basic infons.

Now Barwise and Etchemendy do not explicitly use their framework to define a logic of constraints and information flow. But it may seem there are two candidates for representing the constraint $A \Rightarrow B$: either as the conditional infon $A \rightarrow B$ or, as Barwise and Etchemendy intend, as the condition that $A \leq B$ in the ordering relation on the lattice (or equivalently, as the condition that $A \rightarrow B$ equals $\top$). The problem with the first definition is that this means there is a constraint relating each pair of infons, giving no way to distinguish the lawlike constraints from the incidental regularities. The problem with the second definition is that this places constraints outside the structure of situations being modeled, so that it is impossible for one situation to satisfy a constraint while another does not, meaning that the varying background conditions that underpin different constraints are not captured.

In addition, as noted by Barwise and Etchemendy ([9], p. 55), the class of Heyting algebras properly contains the class of algebras generated from closed coherent sets of infons, hence the logic of constraints in this framework is at least as strong as intuitionistic logic, but is possibly stronger.

Barwise [8] uses channel theory to formalize the flow of information. In channel theory, each constraint has an associated labeled channel so that a constraint $A \Rightarrow B$ is represented as a formula such as $c : A \rightarrow B$. Intuitively, the channel, which in Dretske's formulation is some lawlike connection between two infons, is itself now construed as a situation c (Barwise uses the term 'information site'). The main property of such channels is, of course, that if a situation s_1 is of type A and s_1 is connected to s_2 by means of a channel c of type $A \rightarrow B$, then s_2 is of type B, here taking infons as types. The logical properties of information flow are captured by rules allowing two (or more) channels to be combined into a single compound channel. Conditional logics arising from this use of channel theory are discussed in Cavedon [15] and Restall [54]. Channel theory is intended for modeling information flow but what seems to have been lost is the way in which information flow relies on lawlike regularities between infons, since the property of a situation being a channel of type $A \rightarrow B$ is represented simply as an atomic formula. If a channel is just one single connection between two information sites, what makes one connection lawlike while another is not? What is missing in this approach is an account of how constraints taken as relations between event types arise from a collection of channels. It is this account that we attempt to formalize using our logics of information flow and constraints presented in the following sections.

4. A Logic of Situations and Information Flow

In this section we present a formal semantics of situations and information flow and define a logic corresponding to this semantics. Our logic, called Infon Calculus IC, is a logic of information flow, that is, it includes formulas representing expressions of the form 'A carries the information B'. Our semantics is motivated by Dretske's discussion of the nomic nature of information flow [21], in particular, the requirement that for A to carry the information that B in some situation, there can be no relevant (nomically possible) alternatives to that situation in which A holds but B does not hold. To formalize this idea, we adapt techniques familiar from modal logic, associating with each situation a collection of possible alternatives. These possible alternatives are situations, not possible worlds, that is, they may leave some issues unresolved. By modeling situations as coherent sets of infons, we are able to adapt techniques from modal and conditional logic. That is, whether a conditional infon holds at a situation is determined by the infons holding in each of a number of possible alternative situations, so is analogous to necessary truth in a Kripke model. Technically speaking, what results is a logic of information flow that is closely related to logics of strict implication and entailment.

Our modeling of situations comes from Devlin [20] and Barwise and Perry [10]. A situation is a set of infons which are themselves issues with an assigned polarity. An issue is some basic property of an object or relation between objects which can either hold (positive polarity) or not hold (negative polarity) at some situation. We do not introduce the complexity of conditional infons that Barwise and Etchemendy use to turn the algebra of infons into a Heyting algebra. Instead, we work only with conjunction, disjunction, and negation of infons and follow Devlin in borrowing from data semantics (Veltman [67]) to define a supports and rejects relation between situations and infons. We also impose a restriction on the logical language which ensures that all legal infons are persistent, that is, guaranteeing that if a legal infon holds at a situation, it holds at all extensions of that situation (Langholm [37] provides a more thorough analysis of persistence).

We also define a logic of propositions, SPC. In situation semantics, a proposition is the question of whether some situation s supports an infon A and is written $s : A$. One way of looking at this is that what classically is the same proposition A becomes different "situated" propositions when coupled with different situations. Each situation therefore provides a context in which the "truth values" of the infons A are determined. That is, a situation generates a modality in the same way that possible worlds generate a modality. As a result, SPC is also a kind of modal logic which, as our name implies, can be thought of as a "situated" version of the propositional calculus.

We begin with the basic definitions assuming a given collection of issues.

Definition 4.1 A *basic infon* is an issue together with a polarity (positive or negative).

Definition 4.2 The *simple infons* are defined as follows. A basic infon is a simple infon, and if A and B are simple infons then so are $\neg A$, $A \wedge B$, and $A \vee B$.

Definition 4.3 The *(persistent) infons* are defined as follows. A simple infon is a (persistent) infon, and if A and B are (persistent) infons then so are $\neg\neg A$, $A \wedge B$, $A \vee B$, and $A \rightarrow B$.

Definition 4.4 The *dual* $A^\perp$ of a basic infon A is the issue of A with opposite polarity. The dual operation can be extended to all simple infons by defining $(A \wedge B)^\perp$ as $A^\perp \vee B^\perp$, $(A \vee B)^\perp$ as $A^\perp \wedge B^\perp$, and $(\neg A)^\perp$ as A (this is justified by Lemma 4.8 below).

Definition 4.5 A *situation* is a set of basic infons.

Definition 4.6 A situation is *coherent* if it does not contain any basic infon and its dual.

Definition 4.7 A coherent situation σ *supports (rejects)* a simple infon A, written $\sigma \models A$ ($\sigma \dashv A$), under the following conditions.

$\sigma \models A$	if $A \in \sigma$ for a basic infon A.
$\sigma \models \neg A$	if $\sigma \dashv A$.
$\sigma \models A \wedge B$	if $\sigma \models A$ and $\sigma \models B$.
$\sigma \models A \vee B$	if $\sigma \models A$ or $\sigma \models B$.
$\sigma \dashv A$	if $A^\perp \in \sigma$ for a basic infon A.
$\sigma \dashv \neg A$	if $\sigma \models A$.
$\sigma \dashv A \wedge B$	if $\sigma \dashv A$ or $\sigma \dashv B$.
$\sigma \dashv A \vee B$	if $\sigma \dashv A$ and $\sigma \dashv B$.

Lemma 4.8 *A coherent situation supports a simple infon A if and only if it rejects its dual $A^\perp$.*

The above definitions are standard in situation semantics and are closely related to those of Devlin in [20]. If the infons are represented in a lattice by the sets of situations that support them, then the class of infons forms a De Morgan lattice, with conjunction corresponding to the set intersection, disjunction corresponding to set union, and negation corresponding to the dual operation, as in the relevant logic of Anderson and Belnap [4]. The connection between information flow and relevant logic has been emphasized by Restall in [54].

Definition 4.9 A *De Morgan lattice* $\langle L, \wedge, \vee, ^- \rangle$ is a distributive lattice $\langle L, \wedge, \vee \rangle$ with a unary complementation operation $^-$ such that the following properties are satisfied for all $a, b \in L$.

(i) $\overline{\overline{a}} = a$,

(ii) $\overline{a \wedge b} = \overline{a} \vee \overline{b}$,

(iii) $\overline{a \vee b} = \overline{a} \wedge \overline{b}$.

However, the above definitions are not enough to define an interesting logic: if a theorem is to correspond to the simple infons which are supported by every situation then there are no theorems at all! Any formula A fails to be supported by the situation which is the empty set of basic infons. A logic of situations can be defined by incorporating a conditional connective corresponding to information flow: the formula $A \rightarrow B$ stands for 'A carries the information B'. Such a formula holds in a situation σ if B is supported by all the relevant alternative situations to σ that support A (Dretske [21]). To model this idea, we associate with each situation σ a collection of alternative situations σ^*: σ supports $A \rightarrow B$ if every situation in σ^* that supports A also supports B (cf. [69]). The definition uses the following notion of an extension of a situation.

Definition 4.10 A situation τ *extends* σ in an information flow model Σ, denoted $\sigma \sqsubseteq \tau$, if and only if $\sigma \models A$ implies $\tau \models A$ for any (persistent) infon A. The set of extensions of a situation σ (with a particular information flow model understood) is denoted by σ^+.

Note that $\sigma \sqsubseteq \tau$ only if $\sigma \subseteq \tau$. Thus our definition extends the standard situation semantics definition.

Definition 4.11 An *information flow model* is a set Σ of coherent situations together with a function * assigning to each coherent situation $\sigma \in \Sigma$ a nonempty set of coherent situations $\sigma^* \in \Sigma$ such that

1. $\sigma^* \subseteq \sigma^+$;
2. $\sigma \in \sigma^*$;
3. whenever $\tau \in \sigma^*$, $\tau^* = \sigma^* \cap \tau^+$.

Conditions 1 and 2 are straightforward. The effect of condition 3 is to ensure that information accrual is cumulative. More precisely, suppose an agent is in some situation σ. As information A is acquired, the situation σ "expands" into an extension τ of σ which by definition supports all the infons supported by σ together with A (and possibly other infons). The condition ensures that the information A is retained as further information is gained in the situation τ. The technical condition to guarantee this is that the alternative situations to τ are exactly those alternatives to σ that are extensions of τ.

Definition 4.12 A coherent situation σ *supports (rejects)* an infon $A \rightarrow B$ under the following conditions. The truth conditions given above are now understood to apply to all (persistent) infons.

$\sigma \models A \to B$ if for any situation $\tau \in \sigma^*$ with $\tau \models A$, $\tau \models B$.
$\sigma \Dashv A \to B$ if for some situation $\tau \in \sigma^*$ with $\tau \models A$, $\tau \Dashv B$.

Definition 4.13 An infon A is *semantically persistent* if for any situation σ in any information flow model, if $\sigma \models A$ then $\tau \models A$ for all situations $\tau \in \sigma^*$.

Lemma 4.14 *A persistent infon is semantically persistent.*

The following scheme defines the axioms of the infon calculus IC. We assume a language in which primitive symbols denote basic infons. $A \leftrightarrow B$ is used as an abbreviation for $(A \to B) \wedge (B \to A)$ and $A \supset B$ as an abbreviation for $\neg A \vee B$. Note that we only allow instantiations of the axiom schemes that result in legal (persistent) infons.

(I1) $A \to (A \vee B)$, $B \to (A \vee B)$
(I2) $(A \wedge B) \to A$, $(A \wedge B) \to B$
(I3) $A \to (B \to (A \wedge B))$
(I4) $A \leftrightarrow \neg\neg A$
(I5) $(A \wedge (B \vee C)) \to ((A \wedge B) \vee (A \wedge C))$
(I6) $((A \vee B) \wedge (A \vee C)) \to (A \vee (B \wedge C))$
(I7) $\neg(A \wedge B) \leftrightarrow (\neg A \vee \neg B)$
(I8) $\neg(A \vee B) \leftrightarrow (\neg A \wedge \neg B)$
(I9) $(A \wedge (A \supset B)) \to B$
(I10) $(A \to B) \supset (A \to B)$
(I11) $((A \to B) \wedge (A \to C)) \to (A \to (B \wedge C))$
(I12) $((A \to C) \wedge (B \to C)) \to ((A \vee B) \to C)$
(I13) $((A \to B) \wedge (B \to C)) \to (A \to C)$
(I14) $(A \to (B \to C)) \leftrightarrow ((A \wedge B) \to C)$
(RI1) From A and $A \to B$ infer B
(RI2) If $\vdash B$ then infer $A \to B$

The following are theorems or derived rules of IC.

$A \to A$
$(A \supset B) \to (A \to B)$
$(A \to C) \to ((A \wedge B) \to C)$
$((A \wedge B) \vee (A \wedge C)) \to (A \wedge (B \vee C))$
$(A \vee (B \wedge C)) \to ((A \vee B) \wedge (A \vee C))$
$(A \to (B \to C)) \to ((A \to B) \to (A \to C))$
From A and $A \supset B$ infer B
From $A \vee \neg A$ infer $(A \to B) \to (A \supset B)$
From $A \vee \neg A$, $B \vee \neg B$ infer $(A \to B) \to (\neg B \to \neg A)$

Lemma 4.15 (Deduction Theorem) *If* $\Gamma \cup \{A\} \vdash B$ *then* $\Gamma \vdash A \to B$.

The system IC captures the basic logical properties of situations and infons. In particular, taken together these are just a statement of the fact that the algebra of sets of situations is a De Morgan lattice (the '$\to$' operator plays the role of the ordering relation $\leq$ on the lattice). The axiom schemes (I1) – (I8) and rule (RI1) represent the De Morgan principles. Axiom scheme (I9) relates to the coherence property of situations: this would not be valid if incoherent situations were admitted. Axiom schemes (I3) and (I14) refer to the special (monotonicity) conditions on the * function and are only valid because of the language restriction allowing only persistent infons. The other axiom schemes and rules come from modal logics of strict implication.

Soundness and completeness of the logic IC is relatively straightforward.

Theorem 4.16 IC *is sound and complete with respect to the class of information flow models.*

Dretske considers the following principles as 'inherent in and essential to the ordinary idea of information' ([21], p. 57). These are both derived inference rules of IC.

(Xerox Principle)	From $A \to B$ and $B \to C$ infer $A \to C$.
(Conjunction Principle)	From $A \to B$, $A \to C$ infer $A \to (B \wedge C)$.

Finally in this section, we present a logic of propositions SPC. Propositions are a coupling of a situation σ and an infon A: atomic SPC formulas are written as $c : A$ for a context symbol c denoting a situation in an information flow model. Atomic formulas may be combined using the standard propositional calculus (PC) connectives $\wedge$, $\vee$, $\neg$, and $\supset$ in the usual way.

Definition 4.17 A *propositional interpretation* is an information flow model Σ together with an assignment of a situation from Σ to each context symbol.

The following scheme defines the axioms of the logic SPC. Unless otherwise stated, A and B range over IC formulas.

(P1)	All truth functional tautologies over the language of SPC
(P2)	$c : A$ for A an axiom of IC
(P3)	$(c : \neg A) \supset \neg(c : A)$
(P4)	$c : (A \supset B) \supset (c : A \supset c : B)$
(P5)	$c : (A \to B) \supset c : (A \supset B)$
(P6)	$c : (A \vee B) \supset (c : A \vee c : B)$
(MP)	From α and $\alpha \supset \beta$ infer β

Axiom (P1), which generates formulas such as $(c : A) \wedge (c : B) \supset (c : A)$, reflects the fact that the propositional connectives apply to situated propositions in the same way as to classical propositions. Axiom (P2) with (MP) ensures that every theorem of IC holds at every situation. Axiom (P3) captures the property of coherence of situations, and axioms (P4) and (P5) ensure that the set of facts holding at each situation is logically closed.

In fact, the above system illustrates a general scheme for turning any logic of situations into a logic of propositions. All that is needed are the axiom schemes of the subsidiary logic encoded as propositions, as in (P2), together with the inference rules encoded using material implication, here following from (P3) – (P5), and each property needed to ensure completeness, such as primeness, encoded here as (P6). Finally, the propositional calculus is embedded in the system by way of the propositional tautologies and modus ponens.

Theorem 4.18 SPC *is a sound and complete inference system with respect to the propositional interpretations.*

5. A Logic of Conditionals and Constraints

The logic IC captures the properties of inference based on information flow, that is, inference which is nondefeasible, whereas schema-based inferences are defeasible. In this section we present a conditional logic of constraints SC extending IC in which the connective $\Rightarrow$ is understood as indicating schema-based defeasible inference. We also lift the restriction limiting the language of IC to persistent infons. Under our formulation of constraints, a constraint $A \Rightarrow B$ holds at a situation type σ if B holds at the minimal refinements of σ that support A (if there is no such situation type, the conditional holds trivially). This definition allows us to adapt standard methods from conditional logic to formalize our approach because of the close analogy between refinement and Lewis's comparative similarity.

Note that although the place of constraints in recent versions of situation semantics is not fully clear (cf. [61]), the notion that constraints hold with respect to a situation type is not standard. The difficulty, as described above, is that either a constraint holds with respect to a situation, in which case the lawlike nature of constraints is not captured, or else constraints lie outside a logical structure of situations (the approach taken in [10]), in which case the logic of constraints is difficult to formulate. As we have discussed, the situation type with respect to which a constraint holds is the denotation of a schema; it is a semantic object that supports a set of parametrized infons in the same way a situation supports a set of ordinary infons. That is, we make an assumption that the question of whether a situation type supports a parametrized infon is primitive in the semantics. Part of the reason this is nonstandard is that in recent versions of situation semantics, it has become customary to conflate a situation type with an abstract parametrized infon that the situation supports. This notion of type is more "extensional" to the extent that the set of types classifying a situation is determined purely by the set of parametrized infons it supports. In contrast, our notion is more "intensional" to the extent that there is a separate class of situation types of which a situation may be an instance, and relationships between such types play a significant role in the formulation of a set of constraints.

The logic of constraints turns out to be identical to the logic of conditionals, although the truth conditions for constraints and conditionals differ. This is because we take the logic of conditionals to be the logic of constraints *as applied to situations.* In this way, a conditional is true with respect to a situation in virtue of inheriting a constraint from a situation type of which it is an instance or, as described in Section 2, in virtue of the absence of a condition blocking the constraint's applicability. In either event, the formal properties of hierarchies of situations are the same as those of the hierarchies of situation types that represent constraints.

Our formal semantics of conditionals is based on total preorders of situations analogous to the possible worlds semantics for conditionals given by Lewis [39]. We briefly describe the motivation behind our definitions. As described above, a conditional $A \Rightarrow B$ holds at a situation σ if B holds at all minimal refinements of σ that support A. The natural requirements on orderings of situations are reflexivity, antisymmetry, and transitivity, so that the ordering forms a partial order. But in addition, we require a property known as almost-connectedness (see below) which, in the context of the other conditions, is similar in power to totality, that is, the condition that all pairs of situations are comparable. Some support for this strong condition follows from the intuition that if a situation supports $A \vee B$, it must support either A or B. To see this, say that B is *exceptional* with respect to A if some minimal A-supporting refinement of σ is strictly less refined than some minimal B-supporting refinement of σ. Without comparability, it would be possible to have a structure of situations in which both A and B were exceptional with respect to $A \vee B$, that is, a scenario in which the minimal refinements of σ that support $A \vee B$ neither contain all the minimal A refinements nor all the minimal B refinements. In such a structure, it is consistent to have $(A \vee B) \Rightarrow C$ but neither $A \Rightarrow C$ nor $B \Rightarrow C$. Yet when the constraint $(A \vee B) \Rightarrow C$ is applied in any actual situation, the minimal refinements of that situation which support $A \vee B$ must support either A or B and, given the truth of the constraint, also C, in which case either $A \Rightarrow C$ or $B \Rightarrow C$ should hold at the original situation. Comparability naturally ensues because, following Stalnaker [63], given any two situations σ_1 and σ_2 that are minimal A and B refinements of σ (respectively), a total preorder of situations $\preceq_\sigma$ with respect to σ can be defined by setting $\sigma_1 \preceq_\sigma \sigma_2$ if and only if σ_1 is a minimal $A \vee B$-supporting refinement of σ.

We now proceed to the technical definitions.

Definition 5.1 The *infons* are defined as follows. A simple infon is an infon, and if A and B are infons then so are $\neg A$, $A \wedge B$, $A \vee B$, $A \rightarrow B$, and $A \Rightarrow B$.

Definition 5.2 An *almost-connected partial order* on a set S is a binary relation $\preceq$ that satisfies the following conditions (for all $\alpha, \beta, \gamma \in S$).

1. reflexive, that is, $\alpha \preceq \alpha$;
2. antisymmetric, that is, $\alpha \preceq \beta$ and $\beta \preceq \alpha$ implies $\alpha = \beta$;
3. transitive, that is, $\alpha \preceq \beta$ and $\beta \preceq \gamma$ implies $\alpha = \gamma$;
4. almost-connected, that is, $\alpha \preceq \beta$ implies $\alpha = \beta, \alpha \preceq \gamma$ or $\gamma \preceq \beta$.

Note that the condition of almost-connectedness implies that for any two elements α and β of S which are incomparable (i.e., $\alpha \not\preceq \beta$ and $\beta \not\preceq \alpha$), the set of elements γ such that $\alpha \preceq \gamma$ is the same as the set of elements γ such that $\beta \preceq \gamma$, and the set of elements γ such that $\gamma \preceq \alpha$ is the same as the set of elements γ such that $\gamma \preceq \beta$. Thus the ordering induces a set of equivalence classes in S: for each element of $s \in S$, there is an equivalence class containing s and the elements with which s is incomparable. This is already very similar to the "sphere" semantics of [39] based on total preorders of possible worlds, and moreover, this connection can be made precise by noting that, given an almost-connected partial order $\preceq$, the ordering $\preceq'$ defined by $\alpha \preceq' \beta$ if and only if $\beta \not\preceq \alpha$ or $\alpha = \beta$ is a total preorder.

The following presents two analogous definitions of constraint models which are equivalent for technical purposes, one understood as applying to situations (for modeling subjunctive conditionals) and the other as applying to situation types (for modeling constraints).

Definition 5.3 A situation σ is *minimal* in a set of situations (situation types) Σ ordered by $\preceq$ if $\sigma \in \Sigma$ and for all $\sigma' \in \Sigma$, $\sigma' \preceq \sigma$ implies $\sigma' = \sigma$.

Definition 5.4 A *constraint model* $\langle \Sigma, V, {}^*, \preceq \rangle$ is a set Σ of coherent situations (situation types) together with (i) a valuation function V that assigns to each element σ of Σ a set of basic infons (parametrized infons) $V(\sigma)$, (ii) an accessibility function * that assigns to each element σ of Σ a set of situations (situation types) $\sigma^* \subseteq \Sigma$, and (iii) an ordering function $\preceq$ that assigns to each element σ of Σ an ordering $\preceq_\sigma$ on σ^* that satisfies the following conditions.

1. $\preceq_\sigma$ is an almost-connected partial order on σ^*;

2. every nonempty subset of σ^* has an element minimal in $\preceq_\sigma$;
3. $\sigma \preceq_\sigma \sigma'$ for all $\sigma' \in \sigma^*$.

Condition 1 represents our standard structural requirements on the ordering of elements in Σ. Condition 2 amounts to a well-ordering principle: there can be no set of elements which form an infinitely descending chain according to the ordering $\prec_\sigma$ (defined as $A \prec_\sigma B$ if and only if $A \preceq_\sigma B$ but not $A = B$). This is used to simplify the semantics. Condition 3 is the "centering" condition of [39]: σ is a minimal element in the ordering and hence by antisymmetry is the unique least element in σ^* under the ordering $\preceq_\sigma$.

We can now define the semantics for a logic of conditionals based on constraint models. The semantic rules for SC are identical to those for IC over their common language. The truth rule for infons $A \Rightarrow B$ is similar to that for $A \rightarrow B$ except that σ^* (the set of relevant alternatives to σ) is replaced by $\min_\sigma(A)$, the set of minimal refinements of σ that support A. The precise definitions are as follows.

Definition 5.5 Given a constraint model $\langle \Sigma, V, {}^*, \preceq \rangle$, let $[A]$ denote the set of situation (types) in Σ that support A. For a situation $\sigma \in \Sigma$, the set of *minimal refinements of* σ *that support* A, $\min_\sigma(A)$, is defined as $\{\sigma' \colon \sigma'$ is minimal in $[A] \cap \sigma^*$ ordered by $\preceq_\sigma\}$.

Definition 5.6 A situation (type) σ supports (rejects) an infon $A \Rightarrow B$ under the following conditions.

$\sigma \models A \Rightarrow B$ if for any situation (type) $\tau \in \min_\sigma(A)$, $\tau \models B$.
$\sigma \,=\!\!\mid A \Rightarrow B$ if for some situation (type) $\tau \in \min_\sigma(A)$, $\tau \,=\!\!\mid B$.

The logic SC is defined by the following axiom schemes and inference rules with corresponding axioms and rules from IC repeated for convenience. Note that because we no longer assume persistence, (I3), (I13), and (I14) of IC are not theorems of SC. This will also mean (see below) that the deduction theorem for SC is not valid.

The axiom schemes for the modified IC are as follows.

(I1) $A \rightarrow (A \vee B)$, $B \rightarrow (A \vee B)$
(I2) $(A \wedge B) \rightarrow A$, $(A \wedge B) \rightarrow B$
(I3′) From A and B infer $A \wedge B$
(I4) $A \leftrightarrow \neg\neg A$
(I5) $(A \wedge (B \vee C)) \rightarrow ((A \wedge B) \vee (A \wedge C))$
(I6) $((A \vee B) \wedge (A \vee C)) \rightarrow (A \vee (B \wedge C))$
(I7) $\neg(A \wedge B) \leftrightarrow (\neg A \vee \neg B)$
(I8) $\neg(A \vee B) \leftrightarrow (\neg A \wedge \neg B)$
(I9) $(A \wedge (A \supset B)) \rightarrow B$
(I10) $(A \rightarrow B) \supset (A \rightarrow B)$
(I11) $((A \rightarrow B) \wedge (A \rightarrow C)) \rightarrow (A \rightarrow (B \wedge C))$
(I12) $((A \rightarrow C) \wedge (B \rightarrow C)) \rightarrow ((A \vee B) \rightarrow C)$
(I13′) $(A \rightarrow (B \rightarrow C)) \rightarrow ((A \rightarrow B) \rightarrow (A \rightarrow C))$
(RI1) From A and $A \rightarrow B$ infer B
(RI2) If $A \vdash B$ then infer $A \rightarrow B$

The axiom schemes relating to the conditional arrow are as follows.

(C1) $(A \rightarrow B) \rightarrow (A \Rightarrow B)$
(C2) $(A \Rightarrow (B \wedge \neg B)) \rightarrow (A \rightarrow (B \wedge \neg B))$

(C3) $(A \wedge (A \Rightarrow B)) \rightarrow B$
(C4) $(A \wedge B) \rightarrow (A \Rightarrow B)$
(C5) $(A \Rightarrow B) \supset (A \Rightarrow B)$
(C6) $((A \Rightarrow B) \wedge (B \rightarrow C)) \rightarrow (A \Rightarrow C)$
(C7) $((A \Rightarrow B) \wedge (A \Rightarrow C)) \rightarrow (A \Rightarrow (B \wedge C))$
(C8) $((A \Rightarrow C) \wedge (B \Rightarrow C)) \rightarrow ((A \vee B) \Rightarrow C)$
(C9) $((A \Rightarrow B) \wedge (A \Rightarrow C)) \rightarrow ((A \wedge B) \Rightarrow C)$
(C10) $((A \Rightarrow B) \wedge ((B \wedge A) \Rightarrow C)) \rightarrow (A \Rightarrow C)$
(C11) $(\neg((A \vee B) \Rightarrow B) \wedge ((A \vee B) \Rightarrow C)) \rightarrow (A \Rightarrow C)$
(RC1) From A and $A \Rightarrow B$ infer B

The following are theorems or derived rules of SC.

$A \Rightarrow A$
$((A \Rightarrow B) \wedge (A \Rightarrow (B \supset C))) \rightarrow (A \Rightarrow C)$
$(A \Rightarrow B) \rightarrow ((A \vee C) \Rightarrow (B \vee C))$
$((A \Rightarrow B) \wedge ((A \vee B) \Rightarrow C)) \rightarrow (B \Rightarrow C)$
$((A \Rightarrow B) \wedge (B \Rightarrow A) \wedge (A \Rightarrow C)) \rightarrow (B \Rightarrow C)$
If $\vdash B \rightarrow C$ then infer $(A \Rightarrow B) \rightarrow (A \Rightarrow C)$
From $A \vee \neg A$ infer $(A \Rightarrow B) \rightarrow (A \rightarrow B)$
From $A \vee \neg A$, $B \vee \neg B$ infer $(A \rightarrow B) \rightarrow (\neg B \rightarrow \neg A)$

Soundness and (weak) completeness of SC can be proven. The completeness proof involves showing that for any infon A, there is a situation (type) σ whose minimal refinements support all and only the infons B such that $A \Rightarrow B$. To prove this, we show that any infon has a unique prime decomposition. Since the prime infons correspond to situations (situation types), it is easy to construct a model for any given prime infon. It is then verified that this gives rise to a constraint model and that all and only the intended infons are supported in this model. The proof, however, works only when the number of distinct infons to be supported by the model is finite (in fact, with infinitely many distinct infons, some sets of infons only have models which violate the well-foundedness condition).

Theorem 5.7 SC *is sound and (weakly) complete with respect to the class of constraint models.*

We now briefly compare SC to well-known conditional logics such as SS [49] and VC [39]. One difficulty in comparing SC to these systems is that both these logics contain the propositional equivalences whereas SC is based on structures of situations. However, as has already been noted, the conditions placed on constraint models effectively mean that the situation (types) form a total preorder so that SC is similar to VC. Thus most of the axiom schemes of VC have analogues in SC. In fact, were PC to be added to SC, the characteristic (CV) axiom would follow from (C11).

(CV) $\neg(A \Rightarrow \neg B) \wedge (A \Rightarrow C) \rightarrow ((A \wedge B) \Rightarrow C)$.

This derivation relies on the equivalence of A and $(A \wedge B) \vee (A \wedge \neg B)$ which does not hold for SC. Indeed (CV) is invalid in SC. However, as we discuss below, (C11) generates some similarly strong consequences that violate some intuitions about simple subjunctive conditionals.

Some other theorems of SS (and hence of VC) that are not theorems of SC also rely on the propositional equivalences (i.e., reflect properties of possible worlds rather than situations). The following axiom scheme is valid on the assumption that B is defined at any minimal A-situation.

$$((A \wedge B) \Rightarrow C) \rightarrow (A \Rightarrow (B \supset C)).$$

In addition, both SS and VC contain the centering axiom (MP1) which is invalid in SC. A plausible alternative, (MP2), is also invalid.

(MP1) $(A \Rightarrow B) \supset (A \supset B)$.
(MP2) $(A \Rightarrow B) \rightarrow (A \rightarrow B)$.

The main reason these are invalid in SC is the modal treatment of information flow. The possible worlds interpretation of (MP1) is that if $A \Rightarrow B$ is true at a world w then if A is also true at w, B must also be true at w. This follows from the fact that if w satisfies A, w is always one of the closest A-worlds to itself. With situations, (MP1) is clearly invalid: it is not supported by the "empty" situation which can support $A \Rightarrow B$ but need not support $A \supset B$. (MP2) is valid if B holds at all situations accessible to all A-situations that are accessible to all $(A \Rightarrow B)$-situations accessible to an initial situation. This is only true if $A \Rightarrow B$ and A are persistent so that they both hold in any situation at which B is required to hold.

However, one limitation of SC is that the deduction theorem is invalid. Recall that for IC the deduction theorem follows from the persistence of infons, that is, the condition that if A holds at a situation then A holds at all situations accessible to that situation. More precisely, the proof of the deduction theorem for IC uses axiom scheme (I14) which is not contained in SC. The problem, however, is (I13) which involves nesting of the conditional operator: a counterexample to the deduction theorem based on (I13) is that $\{A, B\} \vdash A \Rightarrow B$ (for two simple infons A and B), but $\{A\} \nvdash B \rightarrow (A \Rightarrow B)$, and similarly $\nvdash A \rightarrow (B \rightarrow (A \Rightarrow B))$. The main reason for this is the lack of persistence when shifting between contexts in the logic SC, so even though every situation satisfying A and B satisfies $A \Rightarrow B$, there is a situation satisfying A but not $B \rightarrow (A \Rightarrow B)$. For the same reason, the deduction theorem fails if $\Rightarrow$ replaces $\rightarrow$ in the formulation of the theorem.

6. Discussion of the Theory

In this section we elaborate and evaluate our theory by discussing a number of issues in conditional reasoning. These are the "paradoxes" of conditional implication, the technical questions of disjunctive rationality and rational monotony, the difference between truth conditions and assertability conditions for subjunctive conditionals, the relationship between subjunctive and indicative conditionals, and some examples from the literature concerning Quine's competing conditionals and the conditional excluded middle, and 'might' conditionals.

Some of the examples and counterexamples are from the artificial intelligence literature on nonmonotonic reasoning. Nonmonotonic (or default) reasoning variously refers to the problem of inferring conclusions based on incomplete or partial information, or to the problem of predicting the outcomes of plans of actions, and has a close connection to conditional reasoning, particularly to the analysis of indicative conditionals. In fact, the connections go further than this: Gabbay [27] initiated the study of nonmonotonic consequence relations (see also Makinson [43]), and conditional logics were proposed as a formalism for default reasoning by Delgrande ([18] and [19])—this work is extended by Boutilier [13] and Asher and Morreau [6].

The fundamental idea of Gabbay [27] is to characterize commonsense reasoning formalisms using a nonmonotonic consequence relation, that is, a consequence relation, usually denoted by $\mid\!\sim$, that fails to satisfy the property of monotonicity. Monotonicity is defined as the condition that whenever $\Gamma \vdash A$, $\Gamma \cup \Delta \vdash A$ for any sets of formulas Γ and Δ and formula A: note that classical (including conditional) logics do satisfy monotonicity. In the case where Γ and Δ are formulas, Gabbay shows how to define conditions on the consequence relation that correspond closely to the "flat" (i.e., nonnested) formulas of conditional logic. For example, a condition called 'cautious monotony' is that if $A \mid\!\sim B$ and $A \mid\!\sim C$ then $A \wedge B \mid\!\sim C$ which clearly corresponds to our (C9). In artificial intelligence, a set of such properties has been used to characterize systems for nonmonotonic reasoning based on partial orders of states (Kraus et al. [35]) and wellfounded total preorders of states (Lehmann and Magidor [38]) where a 'state' here is analogous to a possible world. It is thus no surprise that the resulting logics correspond to the flat fragments of Pollock's SS and Lewis's VC, respectively. Various authors have formalized translations between such nonmonotonic consequence relations and conditional logics, for example, Arlo-Costa and Shapiro [5], Crocco and Lamarre [17], and Fariñas del Cerro et al. [25].

Now although there is a close formal correspondence between nonmonotonic reasoning systems and conditional logics, the underlying intuitions are different. As a rough general rule, however, stronger inferences are desired in nonmonotonic reasoning systems than in conditional logics, simply because commonsense reasoning requires agents to jump to conclusions that may or may not be supported by valid patterns of conditional reasoning. Thus any general property of conditional logic should be a reasonable condition on a nonmonotonic consequence relation, and conversely, any invalid property of a nonmonotonic consequence relation should correspond to an invalid pattern of conditional reasoning. In this way, counterexamples from the literature on nonmonotonic reasoning provide prima facie counterexamples to purported axioms of conditional logic and we discuss two such counterexamples in this section. The particular conditions on nonmonotonic consequence operations of concern are the rules of disjunctive rationality (DR) and rational monotony (RM), defined as follows [35].

(DR) If $A \not\mid\!\sim C$ and $B \not\mid\!\sim C$ then $A \vee B \not\mid\!\sim C$.
(RM) If $A \not\mid\!\sim \neg B$ and $A \wedge B \not\mid\!\sim C$ then $A \not\mid\!\sim C$.

It is evident that (DR) follows from the analogue of our (C11) while (RM) corresponds to Lewis's (CV) axiom scheme which is more powerful than (C11). That is, (DR) follows from (RM) assuming other reasonable properties of the consequence relation. (RM), and hence (DR), are valid in nonmonotonic reasoning systems such as those of Gärdenfors [29] and Pearl [47] as well as in the system of Lehmann and Magidor [38] mentioned above. However, (RM) and (DR) are not valid in many other formal systems developed for nonmonotonic reasoning such as default logic (Reiter [53]). Thus even in the literature on nonmonotonic reasoning, there is some debate about the acceptability of these rules.

6.1 The paradoxes of conditional implication Stalnaker and Lewis have both argued that strengthening the antecedent (from $A \Rightarrow C$ infer $(A \wedge B) \Rightarrow C$), transitivity (from $A \Rightarrow B$ and $B \Rightarrow C$ infer $A \Rightarrow C$), and contraposition (from $A \Rightarrow B$ infer $\neg B \Rightarrow \neg A$) should all be invalid rules of inference in conditional logic: Lewis calls these the 'paradoxes' of conditional implication, although they are, of course, not paradoxes in the usual sense of that term. We note here that due to our use of ordered structures of situations representing a refinement ordering, the three rules are invalid in SC. As with the earlier systems, this can be illustrated with our favorite simple 'birds fly' example. Suppose we have a situation supporting *bird(tweety)* and *fly(tweety)* with one refinement supporting *penguin(tweety)*, *bird(tweety)*, and ¬*fly(tweety)*. Then the following conditionals are supported in this structure, assuming there are no relevant additional background conditions.

The exceptional "penguin" situation type accounts for the failure of strengthening the antecedent and of transitivity. For strengthening the antecedent, the first conditional is supported while the second is not.

If Tweety were a bird, then it would be able to fly.
If Tweety were a bird and a penguin, then it would be able to fly.

For transitivity, the first two statements are supported while the third is not.

If Tweety were a penguin, then it would be a bird.
If Tweety were a bird, then it would be able to fly.
If Tweety were a penguin, then it would be able to fly.

Contraposition fails because the first statement below is supported but the second is not.

If Tweety were a penguin, then it would be a bird.
If Tweety were not a bird, then it would not be a penguin.

Interestingly, this counterexample to contraposition is not available for SS [49] and stronger logics such as Stalnaker's and Lewis's due to the validity of the following axiom scheme.

$$((A \wedge B) \Rightarrow C) \supset (A \Rightarrow (B \rightarrow C)).$$

The next statement follows from this using the other axioms.

$$(A \Rightarrow C) \wedge ((A \wedge B) \Rightarrow \neg C) \supset (A \Rightarrow \neg B).$$

In default logic terms, this example can be interpreted as stating that if B is a reason for revoking the inference of C from A (so that with B, $\neg C$ follows from A), then $\neg B$ follows by default from A. This is a valid rule of inference in many systems of nonmonotonic reasoning, for example, Pearl's System-Z and its descendants [47]. However, it is not valid in SC because its proof in those systems relies on the equivalence of A and $A \wedge (B \vee \neg B)$. Thus the partiality of situations enables us to avoid one of the well-known problems in nonmonotonic reasoning that arises from repeated use of the original rule to infer that, since any kind of bird is exceptional in at least one respect, birds are, by default, not penguins and not emus and not ostriches, but also not robins and not hawks and not eagles, and so on. The opposite default conclusion, that is, that a bird is typically either a robin or a hawk or an eagle or a penguin, and so on, has some intuitive plausibility. However, SC does not sanction this conclusion either. Intuitively, this conclusion should only follow under the assumption that every bird is classified as an instance of a (known) subtype of birds: since there is no reason to suppose this, we claim SC is correct in not admitting this conclusion.

6.2 Disjunctive rationality In this section we examine a counterexample to disjunctive rationality taken from Rott [57]. Actually, Rott presents the counterexample as one against a postulate of epistemic entrenchment derived from rational monotony (see below), but the example also violates disjunctive rationality. In addition, Rott's conditionals are indicative rather than subjunctive but in this case the difference seems to play no significant role.

The example concerns two possible indicators of a person's health: blood pressure and pulse. Take A to be 'Henry's blood pressure is not all right', B to be 'Henry's pulse is not all right' and C to be 'There is no serious danger to Henry's life'. The usual illnesses whose symptoms are an abnormal blood pressure or pulse rate are due to a deficiency of minerals or proteins, and in such cases, there is no danger to life. Thus we accept $A \vee B \mid\sim C$. But under some circumstances, we may, says Rott, be prepared to reject $A \mid\sim C$ and $B \mid\sim C$. In particular, one possible reason for a high blood pressure is an excess of minerals, in which case we should reject $A \mid\sim C$ (this is supposed to be a life threatening danger). Similarly, one possible reason for a low pulse rate is an excess of proteins, in which case we should reject $B \mid\sim C$ (again supposed life threatening). To make these conclusions more plausible, the symptoms of an excess of minerals (danger) are supposed to be exactly the same as those of a deficiency of proteins (no danger) and vice versa. Thus on the basis of one particular symptom, for example, low pulse rate, the patient could have a deficiency of minerals or an excess of proteins, and so we cannot be sure there is no danger.

The counterexample turns on the intuition that when either A (high blood pressure) or B (low pulse rate) is entertained separately, an excess of minerals or proteins is a possibility serious enough to be considered, but when the possibility of either symptom $A \vee B$ is raised, an excess of either minerals or of proteins is implausible. The support for this intuition is the idea that because, considered separately, an excess of minerals or proteins is more implausible than a deficiency of the same substance, this should not mean that an excess of minerals is as plausible as a deficiency of proteins, nor that an excess of proteins is as plausible as a deficiency of minerals (the fact that they have the same symptoms seems to us irrelevant). This much we can agree with. But what Rott requires is that we agree to both of these statements at the same time: that is, that an excess of minerals could be as plausible as a deficiency of proteins (so that one default inference is blocked), *and at the same time* that an excess of proteins could be as plausible as a deficiency of minerals (so that the other is blocked). Rott ([57], p. 57) seems to imply that we can do this since we don't know anything further to decide the question (according to the example). However, disjunctive rationality allows only one of these statements to be true. The intuition underlying (C11) is one of comparability of situations which means that (i) either both excesses are more implausible than both deficiencies, in which case both $A \mid\sim C$ and $B \mid\sim C$ are accepted rather than rejected, or (ii) an excess of one substance is as plausible as a deficiency of the other, in which case one of $A \mid\sim C$ and $B \mid\sim C$ is accepted and the other rejected. The way the example is described tends to favor (i) as the intended interpretation, that is, given one symptom such as a low pulse rate, on this intuition, it should be inferred that the patient is in no danger since this is probably due to a deficiency of minerals rather than to an excess of proteins. In short, the counterexample seems to rely on a confusion between ignorance and indeterminacy: in situation theoretic terms, ignorance of which constraints hold is elevated to indeterminacy of a system of constraints.[2]

Thus we are inclined to reject the counterexample and conclude that the state of affairs as described is appropriately modeled with ordered structures of situations.

6.3 Rational monotony We have motivated all the axiom schemes and rules of SC. In this section we consider the converse question of whether SC is too strong. We consider, in particular, the rule of rational monotony and the corresponding axiom scheme (CV) in Lewis's conditional logic VC.

One way to understand (RM) in the framework of constraints is that it enforces a kind of independence of the effects of noninterfering background assumptions where B's noninterference with A is represented by the failure of the inference $A \mid\!\sim \neg B$. If, in these circumstances, A allows the default inference of C, that is, $A \mid\!\sim C$, then $A \wedge B$ together must also allow the default inference of C, that is, $A \wedge B \mid\!\sim C$: thus in a weak sense, the default conclusion C is "independent" of B. And conversely, whenever B provides a reason to reject the default conclusion of C from A, $\neg B$ must follow by default from A. Intuitively, this seems too strong a rule since B could be, in Pollock's terms, an undercutting defeater for the connection between A and C even when B does not interfere with A (Pollock [50]). Let us construct a counterexample to (RM). Take A to be 'the flight leaves on time', B to be 'there is a storm', and C to be 'the flight arrives on time'. To make the example work, we have to suppose that $A \mid\!\sim C$ and $A \wedge B \not\mid\!\sim C$ are both true. The first, that is, 'if the plane were to leave on time, it would arrive on time', is not problematic. For the second, suppose that with modern technology, in a storm a plane is equally likely to land normally (on time) as not (delayed), possibly according to other factors such as traffic, state of the runway, and so on, so that $A \wedge B \not\mid\!\sim C$. We claim that it should not follow that $A \mid\!\sim \neg B$, that is, that the absence of a storm is in the normal course of events when planes leave on time. For example, even in a city in the tropics that is frequently subject to storms so that $A \not\mid\!\sim \neg B$, the inferences $A \mid\!\sim C$ and $A \wedge B \not\mid\!\sim C$ can still hold with B remaining a defeater for the default conclusion of C from A.

It remains to check that the example can be modeled using ordered structures of situations. Here we suppose an ordering with two situations, one supporting A and C and a more refined situation supporting A and B. It is easy to verify that $A \mid\!\sim C$, $A \wedge B \not\mid\!\sim C$, and $A \not\mid\!\sim \neg B$ are all true in this model. Thus we conclude that rational monotony is invalid and that Lewis's logic VC is too strong. Note, however, that we have used Lewis's intuitions concerning comparability of states of affairs in motivating SC, so the reason this example cannot be modeled in VC has more to do with that system's use of possible worlds than to its reliance on comparability.

More problematic is a counterexample to (CV) given by Pollock ([49], p. 43). His example starts with three unrelated false statements S (my car is painted black), T (my garbage can blew over), and U (my maple tree died). He then considers the following formula (CV′) which follows from (CV) and which is also valid in SC.

(CV′) $$\neg((S \vee T) \Rightarrow S) \wedge \neg(((S \wedge U) \vee T) \Rightarrow T) \supset \neg((S \vee T) \Rightarrow \neg U).$$

Pollock argues that both antecedents are true but the consequent is false. For the first antecedent, since the disjuncts S and T are unrelated, there is no reason to conclude that if the disjunction $S \vee T$ were true, the disjunct S would be true. Similarly for the second antecedent, as $S \wedge U$ and T are also unrelated, there is no reason to conclude that T would be true if $(S \wedge U) \vee T$ were true. For the consequent, since $\neg U$ is true (by assumption), Pollock claims that $\neg U$ would (still) be true if $S \vee T$ were true, making the consequent false. Now since (CV′) is valid in SC, this example also applies to our logic. However, while we accept that the first antecedent is true and the consequent is false, we think matters are less clear with the second antecedent. In particular, it could be argued that under the notion of minimal change used to interpret the conditionals, the change to the initial situation needed to support $S \wedge U$ is an intuitively larger change than that required to support T, so that the conditional $((S \wedge U) \vee T) \Rightarrow T$ is true in the initial situation. This is the kind of response considered by Nute ([46], p. 69) who does not find it satisfying, presumably because it presupposes that the different changes to the initial situation are comparable—this is the assumption leading to (C11) and thus is exactly the point being disputed. However, the model structures proposed by Pollock violate the analogue of Disjunctive Rationality (DR), a principle for which we have provided some intuitive support. The example does not violate (DR), so one way out of the dilemma would be to define a condition on model structures that validates the analogue of (DR) but not (C11). However, we doubt that such a set of purely structural conditions on constraint (or possible world) models exists (as noted above, following Stalnaker [63], the analogue of (DR) leads naturally, but not necessarily, to comparability).[3]

6.4 Truth conditions and assertability conditions Those authors who argue against truth theoretic accounts of conditionals, especially Adams [2], are perhaps influenced by the fact that conditionals seem to incorporate an inherent ambiguity or indeterminacy which makes one suspect that there can be no simple way to determine the appropriate semantic structures for evaluating them. In place of a semantic theory, Adams attempts to base a theory of conditionals on *assertability conditions*, the conditions under which a rational agent is justified in uttering a particular statement. Roughly, the assertability of a conditional 'if A were the case then B would be the case' is the conditional probability of B on A.

The account we have proposed is truth theoretic and is based on the concept of information. We accept that conditionals are highly ambiguous but we consider this ambiguity to be pragmatic rather than semantic, following Lewis. If conditionals express constraints in context, a major source of ambiguity is the determination of that context and the determination of the background assumptions which can be taken to underwrite the constraint. Once a context is determined, the truth or otherwise of the conditional is fixed by whether or not the expressed constraint holds. Insofar as these factors influence the assertability of a conditional, it is clear that the truth conditions and assertability conditions of conditionals will diverge, and of course, if an agent is ignorant of which constraints actually hold, a true conditional may not be justifiably asserted. In fact, we have already seen an example of this in Rott's putative counterexample to disjunctive rationality. In addition, an assertion of a counterfactual presupposes the falsity of the antecedent, so for a counterfactual to be assertable by an agent, the antecedent's falsity must be believed by the agent.

There is, however, a deeper concern. This is that the notion of information itself, or at least our use of it, is defective. In situation semantics, constraints are informational relations between types of facts or events, and as such, their formulation depends on an agent's conceptual scheme including the agent's power to discriminate properties and abstract types of objects. Recall that Dretske's notion of information is similarly relative to the purposes of agents. Could our conceptual scheme be so deficient that hardly any of the constraints we take to hold actually hold? If so, this would not mean that our semantic account of conditionals was wrong but it would render this account less significant as an analysis of conditionals than one based on assertability conditions. Assertability need not here be equated to conditional probability but could instead be based on the evidence available for a constraint as considered part of a scientific theory. Arguments for this perspective on scientific knowledge are advanced by authors such as Dupré [24] who claims that in the sciences generally, "disorder" is the norm rather than the exception and that there are multiple possible world views acceptable in different contexts. This also squares well with the "pragmatic" view of scientific explanation defended by van Fraassen [66].

This fragmentary nature of an agent's conceptual scheme is taken as an accepted part of situation theory, especially in Seligman's discussion of perspectives [60] and Devlin's of agent schemes [20]. As a consequence, in situation semantics, the idea that sentences unproblematically denote their truth conditions has already been abandoned in favor of a "relational" semantic theory in which the meaning of a sentence is a relation between utterance situations and described situations as detailed in [10]. We have studied the application of these ideas to conditional reasoning in an earlier examination of reasoning about action in artificial intelligence (Wobcke [70]), a problem in which the difference between "traditional" semantic theories and the situation semantics view is clearly apparent. In brief, the problem of reasoning about action as presented by McCarthy and Hayes [44] is to develop a formalism in which the effects of everyday actions can be expressed. The traditional "semantic" view is that an action denotes a function from world states (initial states) to world states (resulting states), and the dilemma is that no agent can be expected to predict the entire world state resulting from the performance of a simple action. Thus a large body of research in this area has concentrated on developing formalisms for representing sample problems based not on any firm semantic foundation but on "epistemological" assumptions, for example, Ginsberg and Smith [30], Shoham [62], Winslett [68]. To give a specific example, a common intuition is that any changes be minimal: in the possible worlds theory, minimality is captured by the use of some ordering on world states, whereas in an epistemological theory, minimality refers to the change in an agent's belief state. While the former is often regarded as impossible to model, the latter seems unmotivated: why should the results of my actions be determined only by what I believe? More recently, authors such as Peppas et al. [48] and Sandewall [58] have sought to establish connections between agent's theories and world states where these states are taken as complete descriptions of parts of the world, the "domain" of the agent's actions. What we argued is that this idea of relativizing reasoning about actions to domains, the motivation of which comes from modeling and simulation and engineering control theory, is analogous to taking a situated view of the semantics of action.[4]

Thus we interpret the semantics of conditionals as being based on a collection of fragmentary world models, views the agent has of the world which may vary from context to context, each of which is dependent on a series of background assumptions which it may be impossible to enumerate. This is a semantic view to the extent that such models do or do not correspond to actual situations in the world, and although this leaves much room for pragmatic ambiguity, we claim that this is because there is this much ambiguity in interpreting conditional statements.

6.5 Subjunctive and indicative conditionals It is widely accepted that subjunctive and indicative conditionals have different acceptance conditions. Adams [1] uses the following pair of examples to illustrate the clear intuitive difference between the two classes of statement.

> If Oswald hadn't killed Kennedy, then someone else would have.
> If Oswald didn't kill Kennedy, then someone else did.

Clearly the second statement is acceptable but the first statement may not be.

According to Stalnaker [64], indicative conditionals differ from subjunctives in that indicatives refer to hypothetical changes in an agent's epistemic state whereas subjunctives hold in virtue of external facts irrespective of an agent's knowledge of those facts. Indicative conditionals are evaluated using a modified Ramsey test, following Ramsey [52]. That is, underlying an indicative conditional is a hypothetical change of belief so as to accept the antecedent of the conditional; the conditional is then acceptable if its consequent holds in the resulting belief state. A common intuition is that changes of belief state should be minimal, and much belief revision research is concerned with giving precise formal conditions governing minimal changes of belief, for example, Gärdenfors [28]. Given only this much, it is easy to see why subjunctives differ from indicatives. For we can suppose that one rational principle of minimality is that no belief held independently of the negation of the conditional's antecedent should be given up in hypothetically accepting the antecedent. For example, given that there is overwhelming evidence for Kennedy's dying when he did, all the alternative belief states considered in evaluating the above indicative conditional should include the belief that Kennedy died, so the conditional will be acceptable.

There are two questions we wish to consider further. First, do indicative conditionals have truth conditions (or only acceptance or assertability conditions), and second, does the logic of indicative conditionals differ from that of subjunctive conditionals? We claim that indicatives have truth conditions, although these truth conditions are different from acceptance conditions since many belief revision functions are equally rational and an agent may not be aware of which one is the "correct" revision function to use to change its beliefs. Such might be the case if an agent is unsure of the independence of evidence supporting different beliefs. Continuing the Kennedy example, suppose there are two possible pieces of evidence for Kennedy's dying of a bullet wound: a bullet wound to the head and a bullet wound to the lungs. Now consider two agents, each of whom believes that Kennedy died of a bullet wound and was shot in the head and the back, but one of whom believes that he died because of the head wound while the other believes it was because of the back wound. Consider now the following conditional.

> If Kennedy wasn't shot in the head, then he wouldn't have died.

The agent who accepts the head wound account accepts this conditional while the other agent does not. But even though both agents are rational and have the same initial set of beliefs, they can't both be right. We know that, in fact, Kennedy died because of a head wound. So we see that the first agent is right, the second is wrong, and that the conditional is true. In the sense that the revision function should respect the informational relations obtaining between the described events, indicative conditionals have truth conditions.

It is a separate question whether the truth conditions of subjunctives and indicatives are the same. It is possible within the belief revision paradigm (the AGM paradigm) of Alchourrón et al. [3] to formulate a "semantic" model in which complete theories play the role of possible worlds, as in Grove [33], and belief states are sets of complete theories. In this way, the AGM belief revision postulates can be seen as analogous to Lewis's conditional logic VC. In a similar way, it is possible to define belief revision functions based on ordered structures of situations, for example, Restall and Slaney [55], although note that in that work, the coherence condition on situations is not required. We believe that just as situations provide an appropriate modeling of parts of the world, they provide an appropriate modeling of belief states, and further, the reasons for insisting on comparability of parts of the world apply also to belief states. As a consequence, the logic of subjunctive and indicative conditionals is SC, even though the truth conditions for subjunctive-indicative pairs are different. The differences between the two classes of conditionals result from the different structures of situations used in evaluating them, the logical similarities from the fact that those structures satisfy the same formal properties.

Subjunctives and indicatives are often interchangeable, and our account suggests that this is so when the truth of the indicative is not dependent on the specific known facts about the case in question so that the same class of alternative situations is considered in each case. This is always the case if the negation of the consequent of the conditional is not a known fact, for example, with future tensed conditionals. Consider the following pair of statements.

> If the battery were to be charged, then the car would start.
> If the battery will be charged, then the car will start.

Suppose the normal background conditions hold except that the battery is not charged. Then both conditionals are true and the structures of situations used in their evaluation are the same. Once it becomes known that the car does not start, for example, with the analogous past tensed conditionals, the difference reappears.

> If the battery had been charged, then the car would have started.
> If the battery was charged, then the car started.

The first statement is presumably true while the second is false.

6.6 Two examples We close this section with discussion of two examples from the literature on conditionals. Consider first Quine's famous pair of "competing" counterfactuals [51].

> If Bizet and Verdi had been compatriots, then Bizet would have been Italian.
> If Bizet and Verdi had been compatriots, then Verdi would have been French.

Quine thinks there is no way to decide between these statements but that they cannot both be true. Lewis points out that the second statement entails the following.

> If Bizet and Verdi had been compatriots, Bizet would not have been Italian.

Hence, with any system that has the law of conditional excluded middle $(A \Rightarrow B) \vee (A \Rightarrow \neg B)$, such as Stalnaker's C2, exactly one of the pair must be true and the other false. Which is which is determined by the possible world the selection function treats as being closest to the actual world, and Lewis sees no way for such a decision to be made on logical grounds. He therefore wants both examples to be false (they are contradictory in Lewis's VC assuming the antecedent to be possible). Stalnaker defends the conditional excluded middle by advocating the use of supervaluations (van Fraassen [65]) in conjunction with his logic. This allows either or both (Stalnaker doesn't decide!) statements of the pair to be semantically indeterminate. The conditional excluded middle is still valid because it holds in all ordinary valuations.

In contrast, we follow [7] in taking both conditionals to be true *in appropriate contexts*, although we do not claim that these contexts can be derived from a collection of schemata. According to Barwise, both conditionals can be used to convey information at different times (e.g., if you knew the nationality of Bizet but not of Verdi, you could acquire that information from the first statement of the pair). Under our theory, the conditionals are viewed as being separable pragmatically, giving rise to two different evaluation contexts, one indicating that Bizet has the nationality Verdi does (and so Bizet is Italian), the other indicating that Verdi has the nationality Bizet does (and so Verdi is French). With respect to each context, one counterfactual is true and the other false. Both statements only appear to be unequivocally true because different contexts are used in evaluating them. What is confusing about the examples is the further context dependence to do with the meaning of 'compatriot'. Clearly, whether two people are compatriots is context dependent if the nationality of the people is part of the context which is allowed to vary.

Finally, consider 'might' conditionals, that is, conditionals of the form 'if A were the case then B might have been the case'. One advantage Lewis claims for his analysis over Stalnaker's is that his account can handle such conditionals while all the obvious definitions for the might connective using Stalnaker's semantics for the standard operator lead to problems. We propose that 'might' conditionals can be defined in terms of the standard subjunctive, as $A \Rightarrow \Diamond B$, where $\Diamond B$ holds with respect to a situation σ if and only if B holds at some situation refining σ. Lewis considers this definition and says it fails on the following example.

> If I had looked in my pocket, I might have found a penny.

This is supposing there to be no penny in my pocket so that the conditional should be false (according to Lewis).

But we claim the example is ambiguous with respect to the type of modality in force: whether it is 'metaphysical' or 'epistemic': Lewis clearly understands the modality to be metaphysical. However, following our discussion of subjunctives and indicatives, we suggest that there are two possible interpretations of the sentence and the structures of situations used to evaluate the statement differ according to which interpretation is considered (the example is presumably false under the 'metaphysical' reading but true under the 'epistemic' reading). Thus the example introduces a new element to the context dependence of conditionals: not just in the determination of the situation with respect to which a constraint is assessed but in the specification of the structures of situations used in their evaluation. Indeed, the whole issue of modality is one which deserves more careful analysis within situation semantics.

7. Conclusion

We have developed an account of a subclass of subjunctive conditionals which, on the one hand, is logically formalized using a three-valued logic of situations, and on the other, is suitable for an agent possessing a collection of schemata to use in evaluating certain kinds of conditionals. The theory is based on treating conditionals as expressing constraints—informational relations between facts and events of the kind that can be modeled using structures of situations. The "paradoxes" of conditional implication are explained as arising from the use of defaults and exceptions in schema hierarchies, which in turn is understood in terms of how, in situation semantics, constraints hold only with respect to a series of background conditions which it may not be possible to make explicit.

The present paper provides a more complex picture of conditional reasoning than the early work of Ramsey [52] and Goodman [31], building on but significantly reworking the previous semantic studies of [63], [39], and [49]. However, much more is required before a complete solution to the problem of conditional reasoning can be implemented because, as we have argued in this paper, information flow is highly context dependent, involving multiple conceptual schemes, and moreover, is subject to pragmatic concerns including the beliefs and motivations of agents. We believe we have specified a framework in which issues in conditional reasoning (for a subclass of conditionals) are aligned with issues in knowledge representation in artificial intelligence and defined a formalism whereby logical questions of conditionals in contexts can be meaningfully discussed. However, it is not at all apparent how the more pragmatic issues surrounding the interpretation of conditionals can begin to be addressed within a computational setting.

Appendix A. Proofs

Lemma 4.8 *A coherent situation supports a simple infon A if and only if it rejects its dual $A^{\perp}$.*

Proof: By structural induction on simple infons: for basic infons, the result follows from the definition of the supports and rejects relations. For infons of the form $A \wedge B$, $(A \wedge B)^{\perp}$ is defined as $A^{\perp} \vee B^{\perp}$, and $\sigma \models A \wedge B$ if and only if $\sigma \models A$ and $\sigma \models B$ if and only if $\sigma \dashv A^{\perp}$ and $\sigma \dashv B^{\perp}$ (by the induction hypothesis), that is, if and only if $\sigma \dashv A^{\perp} \vee B^{\perp}$. For infons of the form $A \vee B$, $(A \vee B)^{\perp}$ is defined as $A^{\perp} \wedge B^{\perp}$, and $\sigma \models A \vee B$ if and only if $\sigma \models A$ or $\sigma \models B$ if and only if $\sigma \dashv A^{\perp}$ or $\sigma \dashv B^{\perp}$ (again by the induction hypothesis), that is, if and only if $\sigma \models A^{\perp} \wedge B^{\perp}$. Finally, for infons of the form $\neg A$, $(\neg A)^{\perp}$ is defined as A, and $\sigma \models \neg A$ if and only if $\sigma \dashv A$ (once again, by the induction hypothesis), that is, if and only if $\sigma \dashv A^{\perp}$. □

Lemma 4.14 *A persistent infon is semantically persistent.*

Proof: By structural induction on infons: the interesting case is that of the conditional connective. Suppose $\sigma \models A \rightarrow B$ in some information flow model (i.e., any situation in σ^* that supports A supports B) and suppose $\tau \in \sigma^*$. Since $\tau^* = \sigma^* \cap \tau^+$, $\tau^* \subseteq \sigma^*$, hence any situation in τ^* that supports A also supports B, that is, $\tau \models A \rightarrow B$. □

Lemma 4.15 (Deduction Theorem) *If $\Gamma \cup \{A\} \vdash B$ then $\Gamma \vdash A \rightarrow B$.*

Proof: Following the standard proof for propositional logic using induction on IC proofs: note first that the derived rule of transitivity, that is, from $A \rightarrow B$ and $B \rightarrow C$ infer $A \rightarrow C$, follows from (I3) and (RI1) twice (to infer $(A \rightarrow B) \wedge (B \rightarrow C)$), and (I13) and (RI1) (to conclude $A \rightarrow C$). Now consider any IC proof. If the line of the proof involves instantiation of an axiom scheme B, $A \rightarrow B$ can be inferred as follows: first $B \rightarrow (A \rightarrow (B \wedge A))$ is an instance of (I3) and $A \rightarrow (B \wedge A)$ follows using (RI1), then since $(B \wedge A) \rightarrow B$ is an instance of (I2), $A \rightarrow B$ follows using transitivity. If the line of the proof involves the application of (RI1), say to B_1 and $B_1 \to B_2$, then by the induction hypothesis, there are proofs from Γ of $A \to B_1$ and $A \to (B_1 \to B_2)$. Using (I11) amd (RI1), $A \to (A \wedge B_1)$ follows from $A \to A$ and $A \to B_1$ ($A \to A$ follows from (I4) and transitivity). But $(A \wedge B_1) \to B_2$ follows from $A \to (B_1 \to B_2)$ using (I14) and (RI1), so $A \to B_2$ follows using transitivity again. Finally, if the line of the proof involves an application of (RI2) to infer $B_1 \to B_2$ when B_2 is an IC theorem, a similar application of (RI2) can be used to infer $(A \wedge B_1) \to B_2$, from which $A \to (B_1 \to B_2)$ follows using (I14) and (RI1). □

The next series of definitions leads up to the proof of completeness for IC. The main intermediate results concern a characterization of the coherent situations using prime closed consistent sets of infons.

Definition A.1 A set of infons Γ is *closed* if $A \in \Gamma$ whenever $\Gamma \vdash A$.

Definition A.2 A set of infons Γ is *prime* if whenever Γ contains $A \vee B$, Γ contains A, or Γ contains B.

Lemma A.3 *The set of infons supported by any situation in any information flow model is prime and closed.*

Proof: The fact that the set is prime follows directly from the definition of the supports relation. The fact that it is closed amounts to the soundness of the system IC. More precisely, let σ be a situation in an information flow model and let Γ be the set of infons supported by σ. Suppose $\Gamma \vdash A$ where A is an infon. We have to show that $\sigma \models A$. This proof is by induction on proofs in IC. First, it is easy to verify for each instance A of an axiom scheme that $\sigma \models A$ (coherence is needed for (I9) and persistence for (I3) and (I14)). Second, when an application of (RI1) and (RI2) produce an infon A, the truth conditions ensure that $\sigma \models A$. □

Lemma A.4 *If Γ is a consistent set of infons then there exists a prime closed consistent set of infons containing Γ.*

Proof: As in the proof of Lindenbaum's lemma, start with an enumeration of the infons $A_1, A_2, \ldots$ in increasing order of complexity. Now given a consistent set of infons Γ, define an infinite sequence of sets of infons $\Gamma_1, \Gamma_2, \ldots$ as follows. First, define $\Gamma_1 = \Gamma$. At step i, if $\Gamma_i \cup \{A_i\} \nvdash$ *false* (where *false* is an arbitrary contradiction), define $\Gamma_{i+1} = Cn(\{\Gamma_i \cup \{A_i\}\})$ (the smallest closed set of sentences containing $\{\Gamma_i \cup \{A_i\})$, otherwise define $\Gamma_{i+1} = \Gamma_i$. We show that Γ_∞ is prime, closed, and consistent. Clearly the construction guarantees that Γ_∞ is closed and consistent for these properties hold at each stage of the process. For primeness, suppose that Γ_∞ contains $B_1 \vee B_2$ but neither B_1 nor B_2. Then at step i when $B_1 \vee B_2$ was considered, $\Gamma_i \cup \{B_1 \vee B_2\} \nvdash$ *false*. But $\Gamma_i \cup \{B_1\} \vdash$ *false* and $\Gamma_i \cup \{B_2\} \vdash$ *false* since B_1 and B_2 were inconsistent with subsets of Γ_i at steps before i and are therefore both inconsistent with Γ_i. So by the deduction theorem, $\Gamma_i \vdash B_1 \to$ *false* and $\Gamma_i \vdash B_2 \to$ *false*. But since Γ_i is closed, (I12) together with these assertions implies that $\Gamma_i \vdash (B_1 \vee B_2) \to$ *false*, so that $\Gamma_i \cup \{B_1 \vee B_2\} \vdash$ *false*, a contradiction. □

Corollary A.5 *If Γ is a consistent set of infons and A is a simple infon such that $\Gamma \nvdash A$ then there exists a prime closed consistent set of infons containing Γ that does not contain A.*

Proof: First note that this does not follow directly from Lemma A.4 since, given such a set Γ and infon A, $\neg A$ may not be consistent with Γ (e.g., in the case where A is $B \vee \neg B$), so Lemma A.4 cannot be applied to the set $\Gamma \cup \{\neg A\}$. However, the construction of Lemma A.4 can be generalized in the following way: where in Lemma A.4 the test at each stage is for whether $\Gamma_i \cup \{A_i\} \nvdash \mathit{false}$, we replace this with the test for whether $\Gamma_i \cup \{A_i\} \nvdash A$. At the end of the process, the set Γ_∞ is clearly closed and does not contain A and, using the same reasoning as in the proof of Lemma A.4, this set is also prime. □

Lemma A.6 *For any prime closed consistent set of infons* Γ*, there is an information flow model containing a coherent situation that supports all and only the infons in* Γ.

Proof: We define a mapping $[\,.\,]$ from prime closed consistent sets of infons to situations. For any such set of infons Γ, define $[\Gamma]$ to be the set of basic infons $\{\langle a, +\rangle : a \in \Gamma\} \cup \{\langle a, -\rangle : \neg a \in \Gamma\}$ where a is a primitive symbol in the language of Γ. Fixing a particular such set Γ, the information flow model consists of the set Σ of all extensions of $[\Gamma]$. For any situation σ contained in Σ, define σ^* to be the union of all $[\Gamma_A^\circ]$ where Γ_A° is a prime closed consistent set of infons containing Γ_A, defined for an infon A as $\{B : A \to B \in \Gamma\}$. We must check that the three conditions on the * function are satisfied. Condition 1, $\sigma^* \subseteq \sigma^+$, follows from the fact that $B \to (A \to B)$ is a theorem of IC (this follows from the deduction theorem), so that whenever $B \in \Gamma$, $B \in \Gamma_A$ for any A. Condition 2, $\sigma \in \sigma^*$, follows from the fact that $\Gamma_A = \Gamma$ when A is an IC theorem, so that Γ is counted as one of the Γ_A (if A is an IC theorem, B follows from $A \to B$ using (RI1), and $A \to B$ follows from B using the deduction theorem, so $B \in \Gamma$ if and only if $A \to B \in \Gamma$). Condition 3 is that whenever $\tau \in \sigma^*$, $\tau^* = \sigma^* \cap \tau^+$. This follows from (I14). More precisely, suppose τ is derived from the set $\Gamma_A = \{C : A \to C \in \Gamma\}$. Then any situation in τ^* is derived from a set $\{C : B \to C \in \Gamma_A\}$. By (I14), any such set is identical to the set $\{C : (A \wedge B) \to C \in \Gamma\}$ so the corresponding situation is contained in σ^*. Similarly, for the converse, any situation υ in σ^* that extends τ derives from a set $\{C\colon B \to C \in \Gamma\}$ that also contains A, that is, $B \to A \in \Gamma$. But then $A \to ((A \wedge B) \to C) \in \Gamma$ if and only if $B \to C \in \Gamma$. Hence υ is in τ^*, being constructed from the set $\{C\colon (A \wedge B) \to C \in \Gamma_A\}$.

Finally, we show by structural induction that for all infons A, $\sigma \models A$ if and only if Γ contains A. Consider the 'if' part of this claim and consider first the conditional connective. If $A \to B \in \Gamma$ then once the induction is established and can be applied to A, $\sigma \models A \to B$ follows by construction, here using the fact that $A \to B$ is persistent (i.e., if $A \to B \in \Gamma$ then for any persistent infon C, $(C \wedge A) \to B \in \Gamma$ by transitivity, so $C \to (A \to B) \in \Gamma$ by (I14), hence $A \to B$ holds at all constructed situations in σ^* because these correspond to sets that contain some such Γ_C).

Conversely, suppose $A \to B \notin \Gamma$ and suppose that by applying (I14) repeatedly, the consequent D of any subformula of B of the form $C \to D$ is a simple infon. We have to construct a situation σ_A that supports the elements of Γ_A but rejects B. We do this by induction on the subformulas of B. More precisely, we define a tree whose nodes are labeled with consistent sets and infons; the root of the tree is labeled Γ_A and B and the infons at the leaves of the tree are all simple infons. The tree is constructed according to the following rules. Let Γ_C and D be the label of some node n in the tree. Then the children of n are defined as follows: if D is of the form $\neg\neg D_1$ then n has one child labeled Γ_C and D_1; if D is of the form $D_1 \wedge D_2$ then n has one or two children labeled Γ_C and D_1 or Γ_C and D_2 depending on whether D_1 and D_2 are not contained in Γ_C (at least one of these must be the case); if D is of the form $D_1 \vee D_2$ then n has two children labeled Γ_C and D_1 and Γ_C and D_2; if D is of the form $D_1 \to D_2$ then n has one child labeled $Cn(\Gamma_C \cup \{D_1\})$ and D_2 when D_1 is consistent with Γ_C and one child labeled Γ_C and D_1 when D_1 is inconsistent with Γ_C. The construction stops at a node when it is labeled with a simple infon. The intuition is that corresponding to each node with label Γ_C and D, there is a situation σ_C supporting Γ_C and rejecting D, and this property we prove by induction on the structure of the tree. The desired situation σ_A to complete the proof is then that corresponding to the root of the tree which supports Γ_A but rejects B. For the proof, consider first the base case of the induction. Note that the leaves of the tree are all labeled with a consistent set Γ_C containing Γ_A and some simple infon D such that D is not contained in Γ_C. Hence by Corollary A.5, for each leaf of the tree, there is some prime consistent set containing Γ_C that does not contain D, hence by construction, some situation supporting Γ_C that rejects D.

We now have to verify that the constructed information flow model supports Γ_A but rejects B. So let n be any nonleaf node in the tree labeled Γ_C and D. The induction hypothesis gives, for each child m of n, a situation σ_m supporting the set of infons and rejecting the infon with which m is labeled. We proceed by case analysis. First, if D is of the form $\neg\neg D_1$ then the situation σ_n is σ_m where m is the child of n. Second, if D is of the form $D_1 \wedge D_2$ then set σ_n to be either of the situations corresponding to the children of n (these are identical). Third, if D is of the form $D_1 \vee D_2$ then set $\sigma_n = \sigma_m$ where m is the child of n. Finally, if D is of the form $D_1 \to D_2$ then when D_1 is inconsistent with Γ_C set σ_n to be σ_m where m is the child of n, and when D_1 is consistent with Γ_C set σ_n to be a new situation corresponding to Γ_C. In all cases but the last, the desired result, that σ_n supports Γ_C and rejects D is clear. For the final case (when D is of the form $D_1 \to D_2$ and D_1 is consistent with Γ_C), it is apparent that σ_m supports Γ_C and D_1 but rejects D_2, and since this situation is accessible to σ_n, we have that σ_n supports Γ_C but rejects $D_1 \to D_2$. Thus the induction step is established and the situation corresponding to the root of the tree provides the desired countermodel to $A \to B$.

All this establishes the induction step for infons of the form $A \to B$. It remains to consider the other connectives. Here we show by induction on infons A in disjunctive normal form that $\sigma \models A$ if and only if Γ contains A: the atoms in the DNF are the basic infons and the conditional infons and the axiom schemes of IC ensure that such normal forms exist. If A is an atom or the negation of an atom then $A \in \Gamma$ if and only if $\sigma \models A$ by definition or by the above argument (in the case of conditional infons). If A is of the form $A_1 \wedge A_2$ then $A \in \Gamma$ if and only if $A_1 \in \Gamma$ and $A_2 \in \Gamma$ if and only if $\sigma \models A_1$ and $\sigma \models A_2$ if and only if $\sigma \models A$. Finally, if A is of the form $A_1 \vee A_2$ then $A \in \Gamma$ if and only if $A_1 \in \Gamma$ or $A_2 \in \Gamma$ (using primeness) if and only if $\sigma \models A_1$ or $\sigma \models A_2$ if and only if $\sigma \models A$. □

Example A.7 Let $A = (p \to q) \to (q \to r)$. First, by (I14) this is equivalent to $((p \to q) \wedge q) \to r$. The countermodel will have at least three situations: $\sigma_0 \dashv A$, $\sigma_1 \models (p \to q) \wedge q$ but $\sigma_1 \dashv r$, and $\sigma_2 \models p \wedge q$. The model can be pictured as follows, where the basic infons supported and rejected by a situation are shown inside the circle, those rejected written within square brackets for illustrative purposes. Arrows show accessibility under the * function.

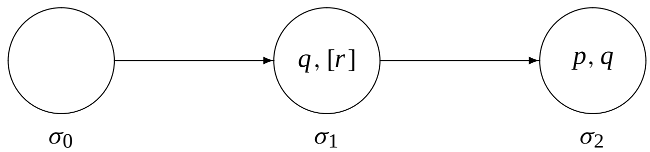

Lemma 4.16 IC *is sound and complete with respect to the class of information flow models.*

Proof: Soundness is easy to check, noting that it relies heavily on the persistence property. For completeness, suppose $\Gamma \nvdash A$; we show that $\Gamma \not\models A$. This follows directly from Lemma A.6, which guarantees the existence of an information flow model with a situation σ such that $\sigma \models B$ if and only if $\Gamma \vdash B$, that is, $\sigma \not\models A$. □

Lemma 4.18 SPC *is a sound and complete inference system with respect to the propositional interpretations.*

Proof: Soundness is easy to check. For completeness, given a consistent SPC theory Γ, for every context symbol c, Lemma A.6 can be used to assign to c an information flow model. This is because (P2)–(P6) guarantee that each set of formulas Γ_c is closed and consistent, hence is contained in a prime closed consistent set of infons, where Γ_c is defined as $\{A\colon (c : A) \in \Gamma\}$. The situation denoted by c is then the situation in the information flow model corresponding to Γ_c. □

The next series of definitions lead up to the proof of soundness and completeness for SC. The initial results concern a characterization of the situations using prime infons.

Definition A.8 Two infons A and B are *provably equivalent*, denoted $A \equiv B$, if and only if $A \vdash B$ and $B \vdash A$ in SC.

Definition A.9 A *conditional infon* is an infon of the form $A \Rightarrow B$ or $A \rightarrow B$.

Definition A.10 A *conditional literal* is an infon of the form A or $\neg A$ where A is a conditional infon.

Definition A.11 An infon p is *prime* over a finite set S of conditional infons if it is a finite conjunction of conditional literals $l_1 \wedge \cdots \wedge l_n$ such that for every infon $A \in S$, either A or $\neg A$ is one of the conjuncts l_i of p, and further, p contains no other conjuncts.

Example A.12 Let $S = \{p \Rightarrow q, q \rightarrow r\}$. Then $(p \Rightarrow q) \wedge \neg(q \rightarrow r)$ is a prime infon over S.

Lemma A.13 *If p and q are consistent prime infons over the same finite set of infons, then $p \vdash q$ only if $p \equiv q$.*

Proof: Suppose p is not equivalent to q. Then p must contain (as one of its conjuncts) some conditional literal r that is not a conjunct of q, that is, $p \vdash r$ and $q \vdash \neg r$ (supposing here the literal r is nonnegated; if not, reverse the roles of p and q). Since $p \vdash q$, $p \vdash \neg r$ by transitivity, so p is inconsistent, contrary to hypothesis. □

Lemma A.14 *If p, q, and r are consistent prime infons over a finite set S of infons, then $p \vdash q \vee r$ implies $p \vdash q$ or $p \vdash r$.*

Proof: Suppose that $p \vdash q \vee r$ but $p \nvdash q$ and $p \nvdash r$. Since $p \vdash q \vee r$, $p \vdash q' \vee r'$ for any literals q' in q and r' in r. But since $p \nvdash q$, there is some literal q_i in q that is not in p, hence $\neg q_i$ is in p and so $p \vdash \neg q_i$. Similarly, there is some literal r_j in r such that $p \vdash \neg r_j$. But since $p \vdash q_i \vee r_j$, p is inconsistent, a contradiction. □

Definition A.15 A *prime decomposition* of an infon A is a finite disjunction of infons $p_1 \vee \cdots \vee p_n$, provably equivalent to A, in which each infon p_i is prime (over the same finite set of infons).

Lemma A.16 *Any infon A is provably equivalent to a prime decomposition of A over some finite set of infons.*

Proof: By induction on infons, following the proof for PC, noting that conditional infons are "propositional" in nature: first replace any basic infon p not within the scope of an '$\Rightarrow$' or '$\rightarrow$' operator by the infon $true \Rightarrow p$ for an arbitrary SC theorem *true*. This leaves an infon built up from conditional infons using the conjunction, disjunction, and negation operators which (C4) guarantees is provably equivalent to A. The IC axioms corresponding to the properties of De Morgan lattices ensure that this is provably equivalent to a disjunction of infons. The set of conditional infons in the prime decomposition is clearly finite. □

Example A.17 Let $A = p \vee (\neg q \wedge (q \Rightarrow r))$. Then $A \equiv (true \Rightarrow p) \vee (\neg(true \Rightarrow q) \wedge (q \Rightarrow r))$. A prime decomposition of A is $((true \Rightarrow p) \wedge \neg(true \Rightarrow q)) \vee ((true \Rightarrow p) \wedge (q \Rightarrow r))$ which is prime over the set $\{true \Rightarrow p, true \Rightarrow q, q \Rightarrow r\}$.

Definition A.18 A prime decomposition $p = p_1 \vee \cdots \vee p_n$ of a infon A is *minimal* if A is not equivalent to $p_1 \vee \cdots [p_i] \cdots \vee p_n$ for any i where this infon is p with the prime p_i omitted.

Corollary A.19 *If a finite set S contains all conditional infons occurring as subformulas of A then a minimal prime decomposition of A over S exists.*

Proof: The construction of Lemma A.16 provides a prime decomposition of A using only the prime infons occurring as subformulas of A. It is easy to turn this into a prime decomposition over S by conjoining the decomposition with all $c \vee \neg c$ contained in S but not occurring in A, then repeatedly applying the distributive law. By ordinary set theory, a minimal prime decomposition therefore exists. □

Lemma A.20 *A minimal prime decomposition of any infon A over a finite set of infons, if it exists, is unique up to provability.*

Proof: Suppose that A has two minimal prime decompositions $P = p_1 \vee \cdots \vee p_n$ and $Q = q_1 \vee \cdots \vee q_m$ over S. Then $p_i \vdash q_1 \vee \cdots \vee q_m$ so $p_i \vdash q_1$ or $\cdots$ or $p_i \vdash q_m$ for each i by Lemma A.14. By suitably reordering the q_j, we can assume that $p_i \vdash q_i$ for each i. Thus $m \leq n$ since otherwise Q is not minimal. Similarly, for each i, $q_i \vdash p_j$ for some j, so $q_i \vdash p_j \vdash q_j$, so $i = j$ (by the minimality of Q) and $n \leq m$. Hence P is provably equivalent to Q. □

We are now in a position to prove completeness for SC. For ease in developing the proof, we define A° to be the set of "nonmonotonic consequences" of A, that is, the set of B such that $A \Rightarrow B$ (relative to a given fixed SC theory) (cf. [35] and especially Freund [26]). We also use this notation to refer to a conjunction of the elements of this set over a finite set of prime infons. The proof uses the prime decomposition of the infon A° to define a model for a consistent SC formula A.

Definition A.21 Given a set Γ of SC sentences, for any infon A, the set A° is defined as $\{B : A \Rightarrow B \in \Gamma\}$.

Lemma A.22 *For a fixed closed set Γ of SC sentences, for any infons A and B, if Γ contains $A \Rightarrow B$ and $B \Rightarrow A$ then A° and B° are equal sets of infons.*

Proof: Suppose Γ contains $A \Rightarrow C$; by (I3′) it follows that Γ contains $(A \Rightarrow B) \wedge (A \Rightarrow C)$, so (C9) and (RI1) give that $(A \wedge B) \Rightarrow C$ is contained in Γ. Also from (I3′), it follows that Γ contains $(B \Rightarrow A) \wedge ((A \wedge B) \Rightarrow C)$, so by (C10) and (RI1), $B \Rightarrow C$ is contained in Γ. That is, we have shown that $A \Rightarrow C \in \Gamma$ implies $B \Rightarrow C \in \Gamma$, so by a similar argument, $B \Rightarrow C \in \Gamma$ implies $A \Rightarrow C \in \Gamma$. Hence the sets corresponding to A° and B° are equal. □

Corollary A.23 *For a fixed closed set Γ of SC sentences, for any infons A and B, if $A \equiv B$ then $A^\circ \equiv B^\circ$ assuming now that A° and B° are prime decompositions over some finite set of infons.*

Proof: If $A \vdash B$ then $\vdash A \rightarrow B$ by (RI2), so $\vdash A \Rightarrow B$ by (C1) and (RI1). Similarly $\vdash B \Rightarrow A$. The result follows from Lemma A.22. □

Lemma A.24 *For a fixed closed set Γ of SC sentences, for any infon A, $A^\circ \equiv A^{\circ\circ}$ assuming both are prime decompositions over some finite set of infons.*

Proof: Since both $A \Rightarrow A^\circ \in \Gamma$ and $A^\circ \Rightarrow A \in \Gamma$, $A^\circ \equiv A^{\circ\circ}$ by Lemma A.22. □

Lemma A.25 *For a fixed closed set Γ of SC sentences, for any infons A and B, $(A \vee B)^\circ \vdash A^\circ \vee B^\circ$, again assuming all are prime decompositions over some finite set of infons.*

Proof: Clearly, $A \Rightarrow A^\circ \in \Gamma$ and $B \Rightarrow B^\circ \in \Gamma$, and since $\vdash A^\circ \rightarrow (A^\circ \vee B^\circ)$ by (I1), $A \Rightarrow (A^\circ \vee B^\circ) \in \Gamma$ using (I3′), (C6), and (RI1). Similarly $B \Rightarrow (A^\circ \vee B^\circ) \in \Gamma$. Hence by (C8) and (RI1), $(A \vee B) \Rightarrow (A^\circ \vee B^\circ) \in \Gamma$. Thus by definition, $A^\circ \vee B^\circ$ is contained in the set $(A \vee B)^\circ$ and hence the infon $A^\circ \vee B^\circ$ follows from $(A \vee B)^\circ$. □

Lemma A.26 *If $p_1 \vee \cdots \vee p_n$ is the minimal prime decomposition of A° over a set S, then $p_i^\circ \equiv p_i$ for each i.*

Proof: Since $A^\circ \equiv ((A^\circ \wedge p_1) \vee \cdots \vee (A^\circ \wedge p_n))$, by Lemmas A.22, A.24, and A.25, $A^\circ \vdash A^{\circ\circ} \vdash ((A^\circ \wedge p_1) \vee \cdots \vee (A^\circ \wedge p_n))^\circ \vdash (A^\circ \wedge p_1)^\circ \vee \cdots \vee (A^\circ \wedge p_n)^\circ \equiv p_1^\circ \vee \cdots \vee p_n^\circ$, that is, $p_1 \vee \cdots \vee p_n \vdash p_1^\circ \vee \cdots \vee p_n^\circ$. So for each i, $p_i \vdash p_1^\circ \vee \cdots \vee p_n^\circ$ and so $p_i \vdash p_j^\circ$ and hence $p_i \vdash p_j$ for some j (since for each j, $p_j^\circ \vdash p_j$). But $p_i \nvdash p_j$ unless $i = j$. This means that for each i, $p_i \vdash p_i^\circ$, and since also $p_i^\circ \vdash p_i$, it follows that $p_i^\circ \equiv p_i$. □

Lemma A.27 *For any consistent infon A, there is a constraint model containing a coherent situation that supports A.*

Proof: We follow the usual Henkin-style construction: for some consistent prime closed theory Γ containing A, we associate a constraint model containing a situation σ that supports all and only the infons in Γ. As usual, it is fairly straightforward to define an ordering on situations from the conditional infons contained in Γ. However, the main complication is the well-orderedness condition. We guarantee well-orderedness by applying the method of filtrations, that is, we identify situations in the model that agree on some set T which contains A and all subformulas of A. This ensures that there are only finitely many distinct situations in the model, so well-orderedness is satisfied.

As in the proof of Lemma A.6, we define an interpretation mapping prime closed consistent sets of infons to situations. The first step, following the filtration method, is to identify such sets that agree over a set T containing A and all subformulas of A. Begin by replacing any basic infon a in A by the infon $true \Rightarrow a$ to give a provably equivalent formula A', then replace A' with its minimal prime decomposition over a set S that contains all conditional infons occurring as subformulas of A': Corollary A.19 guarantees that this exists. A suitable set T is the set of all infons of the form $B \Rightarrow C$ where B and C are disjunctions of prime infons over S. Clearly S and so T are finite and the relation $\sim_T$ on prime closed consistent sets of infons defined by $\Gamma_1 \sim_T \Gamma_2$ if and only if $\Gamma_1 \cap T = \Gamma_2 \cap T$ is an equivalence relation with finitely many equivalence classes.

Let $[\Gamma]$ be an equivalence class under $\sim_T$. We define its interpretation to be a situation whose basic infons are the basic infons contained in all elements of $[\Gamma]$. The constraint model consists of the set Σ of all such interpretations of prime closed consistent sets of infons. For any such situation $\sigma = [\Gamma]$, we define σ^* to consist of the set of all situations σ_p which are the interpretation of an equivalence class $[\Gamma_p]$ as p varies over the prime infons over S and Γ_p varies over the prime closed consistent sets containing $p^\circ \cap S$ (which exist whenever p° is consistent). The situations in σ^* are ordered by setting $\sigma_p \preceq_\sigma \sigma_q$ (for any σ_p, σ_q deriving from some Γ_p, Γ_q) if and only if $p \leq q \in \Gamma$, where $p \leq q$ is the infon $(p \vee q) \Rightarrow p$, as is standard in conditional logic. This is well defined because T contains all infons of the form $p \leq q$ where p and q are prime over S. In addition, σ^* contains situations that reject $B \rightarrow C$ when this infon is not contained in Γ but when $p \Rightarrow C \in \Gamma$ for all primes p in S such that $p \Rightarrow B \in \Gamma$. In such a case, σ^* contains a situation σ' corresponding to an equivalence class derived from a prime closed consistent theory containing B° when B° is consistent. These situations are all ordered maximally in the ordering centered on σ, that is, $\sigma_p \preceq_\sigma \sigma'$ for all the situations σ_p defined above.

We now show that the above construction defines a constraint model. We first show that the ordering $\preceq_\sigma$ is an almost-connected partial order. Reflexivity is straightforward as $p \leq p$ is an axiom of SC. Antisymmetry follows from the fact that when $p \leq q \in \Gamma$ and $q \leq p \in \Gamma$, $p^\circ \equiv (p \vee q)^\circ \equiv q^\circ$ from Lemma A.22. For transitivity, suppose $p \leq q \in \Gamma$ and $q \leq r \in \Gamma$. Since $q \leq r$ means $(q \vee r) \Rightarrow q$, $(p \vee q \vee r) \Rightarrow (p \vee q) \in \Gamma$ by (C6) and (C8), hence $(p \vee q \vee r) \Rightarrow p \in \Gamma$ by (C9), since $p \leq q$ means $(p \vee q) \Rightarrow p$. But $(p \vee q) \Rightarrow p \in \Gamma$ also implies $(p \vee q \vee r) \Rightarrow (p \vee r) \in \Gamma$ by (C6) and (C8), so since $(p \vee q \vee r) \Rightarrow p \in \Gamma$, $(p \vee r) \Rightarrow p \in \Gamma$ by (C9), that is, $p \leq r \in \Gamma$. For almost-connectedness, suppose $p \leq r \in \Gamma$ but $p \not\leq q \in \Gamma$ and $q \not\leq r \in \Gamma$. First, $(p \vee r) \Rightarrow p \in \Gamma$ implies $(p \vee q \vee r) \Rightarrow (p \vee q) \in \Gamma$ by (C6) and (C8). But $\neg((p \vee q) \Rightarrow p) \in \Gamma$ implies $\neg((p \vee q \vee r) \Rightarrow p) \in \Gamma$ by (C6) and (C9), and similarly, $\neg((q \vee r) \Rightarrow q) \in \Gamma$ implies $\neg((p \vee q \vee r) \Rightarrow q) \in \Gamma$. These three conclusions contradict the analogue of (DR) and hence the consistency of Γ. Second, the well-orderedness condition follows from the finiteness of the set of situations Σ. Finally, the fact that σ is the unique minimal element in the ordering $\preceq_\sigma$ follows from (C3) and (C4).

Let σ be a situation associated with the equivalence class deriving from a prime closed consistent set Γ containing A. We show by structural induction on infons that $\sigma \models B$ if and only if $B \in \Gamma$ for all $B \in T$. The only interesting cases are the conditional connectives, for which we must show that $\min_\sigma(B) \models C$ if and only if $B \Rightarrow C \in \Gamma$ and $\sigma \models B \rightarrow C$ if and only if $B \rightarrow C \in \Gamma$. For the first assertion, we show that for any infon B, when the prime decomposition of B° over S is $p_1 \vee \cdots \vee p_n$, $\min_\sigma(A)$ is the collection of all σ_{p_i}. To do this, suppose $\sigma_q \in \Sigma$ and $\sigma_q \models B$. Since $q \Rightarrow B \in \Gamma$, $B \leq q \in \Gamma$, so $B^\circ \leq q \in \Gamma$ (since $B^\circ \leq B \in \Gamma$). Hence $(p_1 \vee \cdots \vee p_n) \leq q \in \Gamma$. By the analogue of (DR), $p_i \leq q \in \Gamma$ for some i, that is, σ_q is not minimal unless it is one of the σ_{p_i}. Note that in the special case that B° is inconsistent, no situation supports B; none of the maximal situations support B since if $B \Rightarrow \mathit{false} \in \Gamma$ then $B \rightarrow \mathit{false} \in \Gamma$ by (C2). Now suppose $B^\circ \vdash C$. Then $p_1 \vee \cdots \vee p_n \vdash C$, so for each i, $p_i \vdash C$. Thus for each i, each $\sigma_{p_i} \models C$, and since $\min_\sigma(B)$ is the collection of all such σ_{p_i}, it follows that $\min_\sigma(B) \models C$, as required. On the other hand, if $B^\circ \nvdash C$, then for some i, $p_i \nvdash C$. So by Lemma A.26, $p_i^\circ \nvdash C$. Hence $\sigma_{p_i} \not\models C$ and so $\min_\sigma(B) \not\models C$, also as required.

Now let $B \rightarrow C \in \Gamma$. For any prime p, if any $\sigma_p \models B$ then $p \Rightarrow B \in \Gamma$ by construction, but since $B \rightarrow C \in \Gamma$, $p \Rightarrow C \in \Gamma$ by (C6), so $\sigma_p \models C$. Conversely, suppose $B \rightarrow C \notin \Gamma$, so B° is consistent by (C2). If $p \Rightarrow C \notin \Gamma$ for some prime p such that $p \Rightarrow B \in \Gamma$, then by the above argument, some situation σ_p supports B but not C, as required. If $p \Rightarrow C \in \Gamma$ for all p such that $p \Rightarrow B \in \Gamma$, some maximally ordered situation supports B but not C. Hence in both cases, $\sigma \not\models B \rightarrow C$, as required. □

Theorem 5.7 SC *is sound and (weakly) complete with respect to the class of constraint models.*

Proof: Soundness is easy to check. For completeness, given any nontheorem A, the infon $\neg(\mathit{true} \Rightarrow A)$ is consistent since otherwise A would be inconsistent (where *true* is an arbitrary SC theorem). Lemma A.27 provides a model for this infon containing a situation at which A is not supported. □

Notes

1. Although we construe schema-based inference as a kind of information extraction, the typical case of information flow involves no reasoning (whether it involves inference is another matter). In Dretske [23], the argument is advanced that certain biological subsystems extract information through being selected for that purpose through evolution. Simple organisms or subsystems are presumably not capable of retracting such information.

2. This is not meant to suggest that Rott is unaware of this potential confusion: semantic concerns are generally not at issue in belief revision research within the AGM paradigm, so there is arguably no need to distinguish ignorance and indeterminacy within that framework.

3. Freund [26] presents one semantic characterization of disjunctive relations, but his main ‘filteredness’ condition refers specifically to the set of A-worlds for each A in a subclass of ‘standard’ models, that is, it is not a condition on model structures themselves.

4. The parallels between reasoning about action and conditional reasoning go further than this, for example, Shoham's theory of reasoning about change [62] is almost identical in details, if not in motivation, to Pollock's semantic account [49] of simple subjunctive conditionals. Moreover, the qualification problem seems to be exactly the problem of delineating the background conditions underlying the truth of a constraint in context (cf. [70]).

References

[1] Adams, E. W., "Subjunctive and indicative conditionals," *Foundations of Language*, vol. 6 (1970), pp. 89–94. 125

[2] Adams, E. W., *The Logic of Conditionals. An Application of Probability to Deductive Logic*, D. Reidel Publishing Co., Dordrecht, 1975. Zbl 0324.02002. MR 58:5043. 123

[3] Alchourrón, C. E., P. Gärdenfors, and D. Makinson, "On the logic of theory change: partial meet contraction and revision functions," *The Journal of Symbolic Logic*, vol. 50 (1985), pp. 510–30. Zbl 0578.03011. MR 87c:03020. 126

[4] Anderson, A. R., and N. D. Belnap, Jr., *Entailment. The Logic of Relevance and Necessity*, vol. 1, Princeton University Press, Princeton, 1975. Zbl 0323.02030. MR 53:10542. 111

[5] Arlo-Costa, H. L., and S. J. Shapiro, "Maps between nonmonotonic and conditional logic," pp. 553–64 in *Principles of Knowledge Representation and Reasoning: Proceedings of the Third International Conference (KR'92)*, edited by B. Nebel, W. R. Swartout, and C. Rich, Morgan Kaufmann, San Mateo, 1992. 119

[6] Asher, N., and M. Morreau, "Commonsense entailment: A modal theory of nonmonotonic reasoning," pp. 387–92 in *IJCAI-91: Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*, edited by J. Mylopoulos and R. Reiter, Morgan Kaufmann, San Mateo, 1991. Zbl 0744.68120. 119

[7] Barwise, J., "Conditionals and conditional information," pp. 21–54 in *On Conditionals*, edited by E .C. Traugott, A. G. B. ter Meulen, J. S. Reilly, and C. A. Ferguson, Cambridge University Press, Cambridge, 1986. 99, 127

[8] Barwise, J., "Constraints, channels, and the flow of information," pp. 3–27 in *Situation Theory and Its Applications*, vol. 3, edited by P. Aczel, D. Israel, Y. Katagiri, and S. Peters, CSLI, Stanford, 1993. 109

[9] Barwise, J., and J. Etchemendy, "Information, infons, and inference," pp. 33–78 in *Situation Theory and its Applications*, vol. 1, edited by R. Cooper, K. Mukai, and J. Perry, CSLI, Stanford, 1990. MR 1 260 726. 107, 109

[10] Barwise, J., and J. Perry, *Situations and Attitudes*, CSLI, Stanford, 1999. Originally published by The MIT Press, Cambridge, 1983. Zbl 0946.03007. MR 2001h:03051. 96, 106, 107, 109, 114, 124

[11] Barwise, J., and J. Seligman, *Information Flow. The Logic of Distributed Systems*, Cambridge University Press, Cambridge, 1997. Zbl 0927.03004. MR 99b:03003. 105

[12] Blamey, S., "Partial logic," pp. 1–70 in *Handbook of Philosophical Logic*, vol. 3, edited by D. M. Gabbay and F. Guenthner, D. Reidel Publishing Co., Dordrecht, 1986. Zbl 0875.03023. 107

[13] Boutilier, C., "Conditional logics of normality as modal systems," pp. 594–99 in *AAAI-90: Proceedings of the Eighth National Conference on Artificial Intelligence*, AAAI Press, Menlo Park, 1990. 119

[14] Cavedon, L., *A Channel-Theoretic Approach to Conditional Reasoning*, Ph.D. thesis, University of Edinburgh, Centre for Cognitive Science, 1995. 101

[15] Cavedon, L., "A channel-theoretic model for conditional logics," pp. 121–36 in *Logic, Language and Computation*, vol. 1, edited by J. Seligman and D. Westerståhl, CSLI, Stanford, 1996. Zbl 0877.03014. MR 1 396 539. 99, 105, 109

[16] Cresswell, M. J., *Semantical Essays*, Kluwer Academic Publishers, Dordrecht, 1988. Zbl 0721.03002. 106

[17] Crocco, G., and P. Lamarre, "On the connection between non-monotonic inference systems and conditional logics," pp. 565–71 in *Principles of Knowledge Representation and Reasoning: Proceedings of the Third International Conference (KR'92)*, edited by B. Nebel, W. R. Swartout, and C. Rich, Morgan Kaufmann, San Mateo, 1992. 119

[18] Delgrande, J. P., "A first-order conditional logic for prototypical properties," *Artificial Intelligence*, vol. 33 (1987), pp. 105–30. Zbl 0654.68106. MR 88g:68099. 119

[19] Delgrande, J. P., "An approach to default reasoning based on a first-order conditional logic: revised report," *Artificial Intelligence*, vol. 36 (1988), pp. 63–90. Zbl 0646.03015. MR 962303. 119

[20] Devlin, K., *Logic and Information*, Cambridge University Press, Cambridge, 1991. Zbl 0732.03001. 106, 107, 108, 109, 110, 124

[21] Dretske, F. I., *Knowledge and the Flow of Information*, Blackwell, Oxford, 1981. 97, 98, 105, 109, 111, 113

[22] Dretske, F. I., "Constraints and meaning," *Linguistics and Philosophy*, vol. 8 (1985), pp. 9–12. 98

[23] Dretske, F. I., *Explaining Behavior*, The MIT Press, Cambridge, 1988. 136

[24] Dupré, J., *The Disorder of Things*, Harvard University Press, Cambridge, 1993. 124

[25] Fariñas del Cerro, L., A. Herzig, and J. Lang, "From ordering-based nonmonotonic reasoning to conditional logics," *Artificial Intelligence*, vol. 66 (1994), pp. 375–93. Zbl 0807.68083. MR 95f:03039. 119

[26] Freund, M., "Injective models and disjunctive relations," *Journal of Logic and Computation*, vol. 3 (1993), pp. 231–47. Zbl 0788.03032. MR 94m:03045. 133, 136

[27] Gabbay, D. M., "Theoretical foundations for nonmonotonic reasoning in expert systems," pp. 439–57 in *Proceedings of the NATO Advanced Study Institute on Logics and Models of Concurrent Systems (La Colle-sur-Loup, 1984)*, edited by K. R. Apt, Springer, Berlin, 1985. Zbl 0581.68068. MR 87d:03071. 118, 119

[28] Gärdenfors, P., *Knowledge in Flux: Modeling the Dynamics of Epistemic States*, The MIT Press, Cambridge, 1988. MR 89k:00015. 125

[29] Gärdenfors, P., "Nonmonotonic inferences based on expectations: A preliminary report," pp. 585–90 in *Principles of Knowledge Representation and Reasoning*, edited by J. F. Allen, R. E. Fikes, and E. J. Sandewall, Morgan Kaufmann, San Mateo, 1991. Zbl 0766.68118. MR 1 142 185. 119

[30] Ginsberg, M. L., and D. E. Smith, "Reasoning about action I: A possible worlds approach," *Artificial Intelligence*, vol. 35 (1988), pp. 165–95. Zbl 0645.68109. MR 89m:68128. 124

[31] Goodman, N., "The problem of counterfactual conditionals," *Journal of Philosophical Logic*, vol. 44 (1947), pp. 113–28. 128

[32] Grice, H. P., "Meaning," *Philosophical Review*, vol. 66 (1957), pp. 377–88. 107

[33] Grove, A., "Two modellings for theory change," *Journal of Philosophical Logic*, vol. 17 (1988), pp. 157–70. Zbl 0639.03025. MR 89e:03027. 126

[34] Hanson, N. R., "Causality," pp. 204–16 in *Philosophical Problems of Causation*, edited by T. L. Beauchamp, Dickenson, Encino, 1961. 103

[35] Kraus, S., D. Lehmann, and M. Magidor, "Nonmonotonic reasoning, preferential models and cumulative logics," *Artificial Intelligence*, vol. 44 (1990), pp. 167–207. Zbl 0782.03012. MR 92e:03033. 119, 133

[36] Kripke, S. A., "Naming and necessity," pp. 253–355 in *Semantics for Natural Language*, edited by D. Davidson and G. Harman, D. Reidel Publishing Co., Dordrecht, 1972. 106

[37] Langholm, T., *Partiality, Truth and Persistence*, CSLI, Stanford, 1988. Zbl 0665.03024. 110

[38] Lehmann, D., and M. Magidor, "What does a conditional knowledge base entail?" *Artificial Intelligence*, vol. 55 (1992), pp. 1–60. Zbl 0762.68057. MR 93i:03035. 119

[39] Lewis, D. K., *Counterfactuals*, Harvard University Press, Cambridge, 1973. MR 54:9979. 99, 105, 114, 115, 116, 117, 128

[40] Lewis, D. K., "Counterfactual dependence and time's arrow," *Nous*, vol. 13 (1979), pp. 455–76. 101

[41] Lewis, D. K., *On the Plurality of Worlds*, Blackwell, Oxford, 1986. 106

[42] Mackie, J. L., "Causes and conditions," pp. 15–38 in *Causation and Conditionals*, edited by E. Sosa, Oxford University Press, Oxford, 1975. 104

[43] Makinson, D., "General patterns in nonmonotonic reasoning," pp. 35–110 in *Handbook of Logic in Artificial Intelligence and Logic Programming*, vol. 3, edited by D. M. Gabbay, C. J. Hogger, and J. A. Robinson, Oxford University Press, New York, 1994. MR 95i:68124. 119

[44] McCarthy, J. M., and P. J. Hayes, "Some philosophical problems from the standpoint of artificial intelligence," pp. 463–502 in *Machine Intelligence*, vol. 4, edited by B. Meltzer and D. Michie, Edinburgh University Press, Edinburgh, 1969. Zbl 0226.68044. 124

[45] Minsky, M., "A framework for representing knowledge," pp. 211–77 in *The Psychology of Computer Vision*, edited by P. H. Winston, McGraw-Hill, New York, 1975. MR 53:9748. 96

[46] Nute, D., *Topics in Conditional Logic*, D. Reidel Publishing Co., Dordrecht, 1980. Zbl 0453.03016. MR 84c:03039. 123

[47] Pearl, J., "System Z: A natural ordering of defaults with tractable applications to nonmonotonic reasoning," pp. 121–35 in *Theoretical Aspects of Reasoning About Knowledge. Proceedings of the Third Conference (TARK 1990)*, edited by R. Parikh, Morgan Kaufmann, San Mateo, 1990. MR 92c:03032. 119, 120

[48] Peppas, P., N. Foo, and W. R. Wobcke, "Events as theory operators," pp. 413–26 in *WOCFAI'91: Proceedings of the First World Conference on the Fundamentals of Artificial Intelligence*, edited by M. De Glas and D. M. Gabbay, Angkor, Paris, 1991. 124

[49] Pollock, J. L., *Subjunctive Reasoning*, D. Reidel Publishing Co., Dordrecht, 1976. Zbl 0345.02016. 101, 117, 120, 122, 128, 137

[50] Pollock, J. L., *How to Build a Person*, The MIT Press, Cambridge, 1989. 122

[51] Quine, W. V. O., *Methods of Logic*, Henry Holt & Company, New York, 1950. Zbl 0038.14811. MR 12,233a. 126

[52] Ramsey, F. P., "General propositions and causality," pp. 237–55 in *The Foundations of Mathematics and Other Logical Essays*, edited by R. B. Braithwaite, Routledge and Kegan Paul, London, 1931. 125, 128

[53] Reiter, R., "A logic for default reasoning," *Artificial Intelligence*, vol. 13 (1980), pp. 81–132. Zbl 0435.68069. MR 83e:68138. 119

[54] Restall, G., "Information flow and relevant logics," pp. 463–77 in *Logic, Language and Computation*, vol. 1, edited by J. Seligman and D. Westerståhl, CSLI, Stanford, 1996. Zbl 0877.03015. MR 97e:03030. 99, 109, 111

[55] Restall, G., and J. K. Slaney, "Realistic belief revision," pp. 367–78 in *WOCFAI'95: Proceedings of the Second World Conference on the Fundamentals of Artificial Intelligence*, edited by M. De Glas and Z. Pawlak, Angkor, Paris, 1995. 126

[56] Rosch, E., "Cognitive representations of semantic categories," *Journal of Experimental Psychology: General*, vol. 104 (1975), pp. 192–233. 96

[57] Rott, H., "Preferential belief change using generalized epistemic entrenchment," *Journal of Logic, Language and Information*, vol. 1 (1992), pp. 45–78. Zbl 0794.03039. MR 95h:03068. 121

[58] Sandewall, E., *Features and Fluents: The Representation of Knowledge About Dynamical Systems*, vol. 1, Oxford University Press, Oxford, 1994. Zbl 0842.68077. MR 96k:68183. 124

[59] Schank, R. C., and R. P. Abelson, *Scripts, Plans, Goals and Understanding*, Lawrence Erlbaum Associates, Hillsdale, 1977. Zbl 0414.68066. 96, 97

[60] Seligman, J., "Perspectives in situation theory," pp. 147–91 in *Situation Theory and its Applications*, vol. 1, edited by R. Cooper, K. Mukai, and J. Perry, CSLI, Stanford, 1990. MR 1 260 731. 124

[61] Seligman, J., and L. S. Moss, "Situation theory," pp. 239–309 in *Handbook of Logic and Language*, edited by J. F. A. K. van Benthem and A G. B. ter Meulen, Elsevier, Amsterdam, 1997. 107, 114

[62] Shoham, Y., *Reasoning About Change*, The MIT Press, Cambridge, 1988. MR 90b:68081. 124, 137

[63] Stalnaker, R. C., "A theory of conditionals," pp. 165–79 in *Causation and Conditionals*, edited by E. Sosa, Oxford University Press, Oxford, 1968. Originally published in *Studies in Logical Theory*, no. 2 in American Philosophical Quarterly Monograph Series, Blackwell, Oxford, 1968, pp. 98–122. 99, 115, 123, 128

[64] Stalnaker, R. C., *Inquiry*, The MIT Press, Cambridge, 1984. 125

[65] van Fraassen, B. C., "Singular terms, truth-value gaps, and free logic," *Journal of Philosophy*, vol. 63 (1966), pp. 481–95. 127

[66] van Fraassen, B. C., *The Scientific Image*, Clarendon Press, Oxford, 1980. 124

[67] Veltman, F., "Data semantics," pp. 541–65 in *Formal Methods in the Study of Language: Part 2*, edited by J. A. G. Groenendijk, T. M. V. Janssen, and M. B. J. Stokhof, Mathematisch Centrum, Amsterdam, 1981. Zbl 0478.03008. MR 84i:03047. 110

[68] Winslett, M., "Reasoning about action using a possible models approach," pp. 89–93 in *AAAI-88: Proceedings, Seventh National Conference on Artificial Intelligence*, edited by T. M. Mitchell and R. G. Smith, Morgan Kaufmann, San Mateo, 1988. 124

[69] Wobcke, W. R., "A schema-based approach to understanding subjunctive conditionals," pp. 1461–6 in *IJCAI-89: Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, edited by N. S. Sridharan, Morgan Kaufmann, San Mateo, 1989. Zbl 0714.68077. 99, 111

[70] Wobcke, W. R., "Reasoning about action from the perspective of situation semantics," pp. 71–76 in *AI'93: Proceedings of the Sixth Australian Joint Conference on Artificial Intelligence*, edited by C. Rowles, H. Liu, N. Y. Foo, World Scientific, Singapore, 1993. 124, 137

Acknowledgments

This work started as Ph.D. research in the Department of Computer Science at the University of Essex: my outlook on these problems has been heavily influenced by my supervisor, Raymond Turner, to whom I am especially grateful. This paper also draws heavily on my work on belief revision carried out in the Knowledge Systems Group at the University of Sydney, and thanks are due to Norman Foo for coordinating this group. Greg Restall explained the connections between (one version of) situation theory and relevant logic, and discussions with David Israel, John Perry, and Lawrence Cavedon have all helped to clarify some of the finer points of situation semantics. Thanks also to Lawrence Cavedon for detailed comments which improved the presentation and content of the paper.

School of Computer Science and Engineering
University of New South Wales
Sydney NSW 2052
AUSTRALIA
wobcke@cse.unsw.edu.au
http://www.cse.unsw.edu.au/~wobcke/