

Frank Veltman's *Logics for Conditionals* Revisited

Richmond H. Thomason

Philosophy Department
University of Michigan
rthomaso@umich.edu

Abstract

Going back to Frank Veltman's dissertation, this note contrasts that approach to conditionals with one that starts with causality and separates epistemology from the interpretation of conditionals.

Contents

1	Introduction	1
2	Conditional semantics	2
3	Interpreting conditionals in specific commonsense domains	3
4	An example	4
5	Conclusion	6

1 Introduction

I wish that Frank Veltman's dissertation, [Vel85] had been more widely read and appreciated. It provides a thoughtful, comprehensive and systematic treatment of logical and philosophical issues in the interpretation of conditionals. Subsequent work would have been improved if authors had studied this work and referred to it; but this happened, I think, far less than it should have. The dissertation remains one of the best sources to be found, for both the philosophy and the mathematical logic of conditionals.

Over the years I have learned many things from Veltman's work, but in this note I will return to the early work on conditionals.

The logic of conditionals, even after many years, enjoys (or perhaps suffers from) fairly wide differences in the theoretical approaches that are endorsed. The approach I've always favored differs from Veltman's. I will discuss the differences with respect to a fundamental issue that is raised in the 1985 dissertation: how to construct a "revision function" for conditionals. This means that Part II of the dissertation, "Possible Worlds Semantics," is especially relevant to my topic. Part I deals mainly with methodological issues. The treatment of conditionals in Part III, "Data Semantics," belongs to an entirely different family of theories of the conditional and I will not discuss it here, although I think that the main points I wish to make could be extended to this case as well.

2 Conditional semantics

I coined the term “revision function” merely to have a temporary term that would include both the selection functions used by Stalnaker and Lewis and the “premise” semantics introduced by Veltman and (later) by Kratzer.

Revision functions are the essential component in models for languages containing conditionals. A conditional involves two syntactic components, which logicians call the antecedent and the consequent. Both the antecedent and the consequent are sentential clauses, and so express propositions. The conditional itself expresses a proposition. Therefore a model must include a conditional interpretation function—a two-place function from propositions to propositions—in order to deliver an interpretation of conditional constructions. In models based on possible worlds (and these are the only models we will consider) such a function must be associated with each possible world.

In selection function semantics¹ a model contains a function s taking a possible world and a proposition into a set of possible worlds (which may well be a unit set, or even empty). The following semantic rule provides a revision function:²

$$(1) w \in \llbracket \phi > \psi \rrbracket \text{ iff } u \in \llbracket \psi \rrbracket \text{ for all } u \in s(w, \llbracket \phi \rrbracket).$$

A conditional formula $\phi > \psi$ is true in w if and only if ψ is true in every world in $s(w, \phi)$.

In premise semantics, a model contains a premise function p from worlds to sets of propositions; the sets that this function returns can be thought of as the set of beliefs of some idealized agent in w . But more needs to be said to obtain an interpretation of conditional constructions from such a premise function.

Veltman discusses how this might be done. A simple version (which assumes the limit assumption) goes like this:

- (i) $P \subseteq \mathcal{P}(\mathcal{W})$ admits ϕ iff $\bigcap P \cap \llbracket \phi \rrbracket \neq \emptyset$.
- (ii) $w \in \llbracket \phi > \psi \rrbracket$ iff for all maximal subsets P of $p(w)$ that admit ϕ , $\bigcap P \cap \llbracket \phi \rrbracket \subseteq \llbracket \psi \rrbracket$.

Premise semantics takes an epistemological approach to the interpretation of conditionals, inspired by the Ramsey rule: to see if you believe a conditional, add the premise to your stock of beliefs, make necessary adjustments, and see if a belief in the consequent is supported by the result of this process. Thus, this approach takes belief revision to form a basis for the interpretation of conditionals.

Although the very large literature subsequent to [AGM85] has told us a great deal about belief revision, as far as I know this has provided no way of showing us how to revise beliefs in well-specified reasoning domains. Suppose, for instance, we are working with a domain for reasoning about mechanical devices of the sort discussed in the AI literature.³ Such a domain will contain a great deal of information about a mechanical device, such as an automobile. The challenge is, given this information, to be able to correctly answer conditionals about the domain, such as this one: *If I step on the brake pedal and turn the ignition key, the engine will start.*

The difficulty with using revision theory to solve such problems is that, in versions of the theory like Veltman’s, too many conditionals that you’d want to be true come out false, whereas in more complicated versions you would need to use a causal theory of the domain to restrict

¹To simplify a bit, and ignore complications having to do with the limit assumption.

²This rule covers both Stalnaker’s and a slightly simplified version of Lewis’ semantics.

³See, for instance, [Gel95].

the maximal belief sets used to interpret conditionals. In our automotive example, we can't expect the domain theory to include a material conditional like (2), which tells us that stepping on the brake and turning the ignition key will start the engine.

$$(2) [\text{DO}(\text{Press-brake-pedal}) \wedge \text{DO}(\text{Turn-ignition-key})] \rightarrow \text{NEXT}(\text{Engine-On})$$

This (material) conditional cannot be a domain axiom, because many factors can prevent the engine from starting. Therefore the negation of (2) can figure in maximal belief sets, which will then be causally anomalous. Perhaps that is all right if we're dealing with a causally ignorant agent, but in most useful cases we will be interested in agents who are relatively well-informed about the causal basics of the domain.

Perhaps this problem could be solved by elaborating domain formalizations: including not only a theory of the device, but a theory of the beliefs of the device user. This brings me to my next point: although maybe this could be done, it seems to be redundant, because we should already expect a robust sample of conditionals to be supported by a bare causal theory of the device, without any epistemological additions.

If we have a such a bare causal theory and wish to add to it a theory of what an agent believes, we need first to enlarge the domain states allowed by the bare theory by systematically adding possibilities that violate the constraints of the domain. This larger set of states represents the epistemic possibilities for an agent who is partially ignorant of the domain theory. Subsets of this enlarged set of states will represent an agent's beliefs about the domain, including general beliefs about automobiles and particular beliefs about the situation at hand. According to this approach, then, the theory of belief tracks possible states of the domain (where "possible" is liberally construed).

But we can have more than a bare causal theory; we might wish to add conditionals to it. We might wish to say, for instance, that the conditional

$$(3) [\text{DO}(\text{Press-brake-pedal}) \wedge \text{DO}(\text{Turn-ignition-key})] > \text{NEXT}(\text{Engine-On})$$

is false in a state in which the battery is dead, even if, say, $\text{DO}(\text{Turn-ignition-key})$ is false.

Where we have an enhanced causal theory of this sort, we can proceed to add to it a theory of an agent's beliefs just as we did in the case of the bare domain, by treating beliefs as sets of suitably enlarged set of domain states. If one wants to have both epistemology and conditionals, this is how I think we should do it.

3 Interpreting conditionals in specific commonsense domains

If we follow the modal account of knowledge and belief, as first proposed in [Hin62] (or, for that matter, the account of belief as subjective probability), epistemology recapitulates ontology. We first determine the truth conditions of sentences that do not contain epistemic operators: whether an agent believes such a sentence is then determined by a relation over worlds. For instance, we do not have a "factual" negation and an "epistemic" negation. A negative sentence, such as $\neg p$, even if it's vague or obscure, is true or false in a world w without reference to how agents come to believe $\neg p$. A belief sentence, such as $B_a \neg p$, is true in w iff $\neg p$ is true in all worlds related by R_a to w , where R_a is the belief relation for agent a .

I want to propose that we should treat conditional sentences in much the same way we treat negative ones. If this means we have to discard the Ramsey condition as a guide to conditional semantics, that is a small price to pay for the methodological advances of this procedure.

Consider how this might work out in a specific formalization project.

In formalizing a device theory it is natural, and I would say inevitable, to include causal information as part of the domain.⁴ Ordinarily, except perhaps in tutoring applications, representations of beliefs would be omitted.


Causal information about a domain can provide a method of constructing a selection function, which then provides a semantics for the interpretation of conditionals in the domain. [Tho07] describes how this works in the case of planning domains—domains for reasoning about the consequences of performing actions. These domains require (in the deterministic case) a specification of what momentary state will ensue when an action is performed in an initial state, or (in the nondeterministic case) of what alternative states might result. This involves a causal theory, which might be an axiomatization of the effects of actions in some familiar logic, or which might involve a nonmonotonic logic. [Tho07] shows how to use such a theory to interpret conditionals of the form $\text{DO}(a) > \phi$, where a is an action term and ϕ is a boolean combination of fluent literals. (A fluent is a formula describing aspects of the world that may change.)

In this example, information about actions and their effects automatically induces a specific selection function appropriate to the domain, and so provides a suitable interpretation of conditionals—or at least of certain conditionals, restricted in their syntactic form. I have not attempted to do this for other causal domains that have been formalized by the AI community, but I think the chances are good that it can be done for them as well.⁵

On the other hand, it is not at all clear how to produce an interpretation for conditionals in one of these domains, or in any formalization of a suitable domain, taking the epistemic approach. No matter how much you know about the domain and its causal structure, you would need to begin all over again in taking an epistemic approach to this topic. Our intuitions about beliefs in themselves don't seem to deliver what is needed to make a start. But if we have a domain that already is equipped with conditionals, we have at least one case in which beliefs involving conditionals are clear—the one where the agent understands the causal laws of the domain and is ignorant only of facts about the current situation.

4 An example


I'll illustrate this with the marbles-and-boxes example introduced in Chapter 1 of Veltman's dissertation. In this domain there are two boxes, and three marbles of different colors: red, blue, and yellow. Each marble is assigned to only one box, and each box must contain at least one marble. We can represent these worlds with diagrams like the following one, where the red marble is in Box 1 and the others are in Box 2.



There are six possible worlds in the domain:

⁴See, for instance, [Pea87].

⁵In work belonging to a slightly different tradition, Judea Pearl shows how to produce a theory of the probability of (subjunctive) conditionals out of causal graph models. See, for instance, [PG96].



For a conditional language, we need selection functions as well as worlds. I'll assume that selection functions take only unit or empty sets as values. This simplifies things. (It also is the logic I favor for conditionals, but the validity of conditional excluded middle is not really an issue here.) In this case, a selection function f_w for a base world w can be defined by a linear ordering over worlds with least element w : where P is a set of worlds, $f(P) = \{u\}$ if u is the least world in the ordering belonging to P , and is \emptyset if there is no such world.

Therefore a selection function for w_1 in the marbles domain is defined by ordering w_2, w_3, w_4, w_5 , and w_6 , and there are 120 such functions. If we use the number of marbles in different positions to calculate the distance between w_1 and one of these five worlds, w_3 and w_4 are tied in the closest position to w_1 ; w_4 and w_5 are tied in the next position, and w_6 is furthest. There are four selection functions that satisfy these constraints, each defined by a linear ordering of $\{w_2, w_3, w_4, w_5, w_6\}$, as follows.

$$\begin{aligned}
 f_1^1 : & \quad w_2 \prec_1 w_3 \prec_1 w_4 \prec_1 w_5 \prec_1 w_6 \\
 f_2^1 : & \quad w_3 \prec_2 w_2 \prec_2 w_4 \prec_2 w_5 \prec_2 w_6 \\
 f_3^1 : & \quad w_3 \prec_3 w_4 \prec_3 w_5 \prec_3 w_2 \prec_3 w_6 \\
 f_4^1 : & \quad w_4 \prec_4 w_3 \prec_4 w_5 \prec_4 w_2 \prec_4 w_6
 \end{aligned}$$

These selection functions differ only in arbitrary pairwise orderings that are imposed in order to achieve linearity. Arbitrary differences of this sort are best handled using supervaluations, as suggested in [Sta80]. That is, we would say that $f_1^1 \sim f_3^1 \sim f_2^1 \sim f_4^1$, and use this similarity relation to introduce truth-value gaps, so that a conditional like

$$(\neg \text{Red}(1) \vee \neg \text{Blue}(1)) > \text{Red}(1)$$

is truthvalueless, because f_1^1 satisfies it and f_2^1 does not.

I will not bother with supervaluations here, but I will use the relation \sim in constraining belief.

Similarly, we have four selection functions associated with each of the five other worlds: for instance, f_1^2, f_2^2, f_3^2 , and f_4^2 are the four selection functions associated with the base world w_2 .

Now we come to belief. We might wish to deal with believing agents who were ignorant about the domain constraint (at least one marble per box) or about the color and number of marbles, or the number and color of boxes. But to keep things simple, we'll assume that our agent has appropriate beliefs about these things. (This is the simple case that I mentioned at the end of Section 3, above.)

In a domain without conditionals, an agent's uncertainty has simply to do with which world is actual. With conditionals, the uncertainty has also to do with selection functions. An agent may believe that the actual configuration of marbles corresponds to w_1 , but be uncertain about which of the functions $\{f_1^1, f_2^1, f_3^1, f_4^1\}$ is actual. In that case, the agent would believe $\text{Yellow}(2)$, for instance, but would believe neither $\text{Yellow}(1) > \text{Red}(2)$ nor $\text{Yellow}(1) > \neg \text{Red}(2)$. On the other hand, the agent would believe $\neg \text{Red}(1) > \text{Blue}(1)$, since this conditional is true for all four of the selection functions associated with w_1 .

The simplest approach to interpreting belief in a possible worlds framework uses a framework R from worlds to sets of worlds. With models of languages containing conditionals and using selection functions, we must generalize this and use functions from worlds and selection functions to worlds and selection function pairs. This, however, can be simplified, if we assume—naturally enough—that an epistemic state depends only on the world. Then, where R is a function from worlds to world/selection-function pairs, belief is interpreted as follows.

$$(4) \quad w \in \llbracket B(\phi) \rrbracket \text{ iff } \langle u, f \rangle \in \llbracket \phi \rrbracket \text{ for all } \langle u, f \rangle \in R(w).$$

It's natural, too, to assume that arbitrary differences between selection functions—differences which have no grounding in anything factual and would have to be resolved using supervaluations—are opaque to attitudes like belief. This leads to the following constraint.

$$(5) \quad \text{If } f \sim g, \text{ where } f \text{ and } g \text{ are selection functions associated with } u, \\ \text{then } \langle u, f \rangle \in R(w) \text{ iff } \langle u, g \rangle \in R(w).$$

We now can evaluate the agent's beliefs in conditionals: formulas having the form $B(\phi > \psi)$. If, for instance, we work with an agent who is totally ignorant of the facts, then in world w_1 $\text{Red}(1) > \text{Blue}(2)$ will be true (because both $\text{Red}(1)$ and $\text{Blue}(2)$ are true), but $B(\text{Red}(1) > \text{Blue}(2))$ will not, because $\langle w_2, f_1^2 \rangle$ is compatible with the agent's beliefs in w_1 .

I am not sure that anything more is required to account for the interaction of conditionals with propositional attitudes like belief. Of course, belief in a conditional is not the same as conditional belief, but the distinction between the two is subtle enough so that it is hard to produce convincing intuitive evidence that anything vital is missing if we have a theory of the former.

I realize that many people believe that subjunctive conditionals are causal and indicative conditionals are epistemic, but in this note I have intentionally steered clear of the distinction between subjunctive and indicative conditionals. Fortunately the main points I want to make here can be made while steering clear of the complications that would entail. With Stalnaker, I believe we should avoid having to make 'if' ambiguous; this means that a pragmatic account of the subjunctive/indicative distinction is to be preferred if it can be made to work.

5 Conclusion

I feel that conditionals are at bottom matters of fact (although they may be vague and context dependent), and that epistemic operators need to be added to conditionals, rather than somehow built into them. Others favor an approach to conditionals that is epistemic from the start. The issue between these views can now be put in the form of a challenge to the epistemologists. I have indicated here how we can produce a logic of belief and conditionals that does justice to the logic of conditionals, uses standard theories of the logic of belief, and, in simple cases, at least, allows us to formalize domains that combine conditionals and an agent's beliefs.

The challenge is for those who believe in an epistemic theory of conditionals to show how their approach could produce something similar.

References

- [AGM85] Carlos E. Alchourrón, Peter Gärdenfors, and David C. Makinson. On the logic of theory change: Partial meet contraction functions and their associated revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.
- [Gel95] Andrew Gelsey. Automated reasoning about machines. *Artificial Intelligence*, 74(1):1–53, 1995.

- [Hin62] Jaakko Hintikka. *Knowledge and Belief*. Cornell University Press, Ithaca, New York, 1962.
- [Pea87] Judea Pearl. Embracing causality in formal reasoning. In Kenneth Forbus and Howard Shrobe, editors, *Proceedings of the Seventh National Conference on Artificial Intelligence*, pages 369–373, Menlo Park, California, 1987. American Association for Artificial Intelligence, AAAI Press.
- [PG96] Judea Pearl and Moisés Goldszmidt. Probabilistic foundations of reasoning with conditionals. In Gerhard Brewka, editor, *Principles of Knowledge Representation*, pages 33–68. CSLI Publications, Stanford, California, 1996.
- [Sta80] Robert C. Stalnaker. A defense of conditional excluded middle. In William L. Harper, Robert Stalnaker, and Glenn Pearce, editors, *Ifs: Conditionals, Belief, Decision, Chance, and Time*, pages 87–104. D. Reidel Publishing Co., Dordrecht, 1980.
- [Tho07] Richmond H. Thomason. Conditionals and action logics. In Eyal Amir, Vladimir Lifschiz, and Rob Miller, editors, *AAAI 2007 Spring Symposium on Commonsense Reasoning*, pages 156–161, Menlo Park, California, 2007. American Association for Artificial Intelligence. Available at http://web.eecs.umich.edu/~rthomaso/documents/lai/Conditionals_and_Action_Logics.pdf.
- [Vel85] Frank Veltman. *Logics for Conditionals*. Ph.D. dissertation, University of Amsterdam, Amsterdam, 1985.