

Material consequence and counter-factuals

DAVID HITCHCOCK

Department of Philosophy
McMaster University
Hamilton, Ontario
Canada L8S 4K1
hitchckd@mcmaster.ca

ABSTRACT: A conclusion is a “material consequence” of reasons if it follows necessarily from them in accordance with a valid form of argument with content. The corresponding universal generalization of the argument’s associated conditional must be true, must be a covering generalization, and must be true of counter-factual instances. But it need not be law-like. Pearl’s structural model semantics is easier to apply to such counter-factual instances than Lewis’s closest-worlds semantics, and gives intuitively correct results.

KEYWORDS: associated conditional, closest-world semantics, consequence, counter-factuals, covering generalization, David Lewis, Judea Pearl, law-like generalization, material consequence, structural model semantics

1. INTRODUCTION

Good arguers support their claims with reasons from which the claim actually follows. To clinch the argument, the claim would have to follow necessarily, in the sense that it is not possible for the reasons to be true and the claim untrue. The claim can follow necessarily in virtue of a contentless form of one’s argument, as when one argues by *modus tollens*:

(1) There is no life on Mars, since its atmosphere is in a static equilibrium and its atmosphere would not be in a static equilibrium if there were life there.

But it can also follow necessarily in virtue of a contentful form of one’s argument, as when one argues more succinctly (and more naturally):

(2) There is no life on Mars, since its atmosphere is in a static equilibrium.

The contentful form of argument in virtue of which the claim now follows is: The atmosphere of planet *x* is in a static equilibrium, so there is no life on planet *x*. This second way in which a claim can follow necessarily from reasons has come to be known, following Sellars (1948, 1953), as *material consequence*. It has been discussed by Bolzano (1972/1837), Peirce (1956/1877), Ryle (1950), Toulmin

(1958), George (1972, 1983), Hitchcock (1985, 1998, 2011), Brandom (1994, 2000), Pinto (2006) and Freeman (2006, 2011).

2. COVERING GENERALIZATIONS

Material consequence evidently requires the truth without exceptions of some contentful generalization of what I shall call the argument's 'associated conditional': the material conditional whose antecedent is the conjunction of the reasons and whose consequent is the claim. Otherwise the argument would not have a valid contentful form in virtue of which the claim follows.

Material consequence also requires that at least one variable in a true generalization of its associated conditional be shared by its antecedent and consequent. Otherwise it would reduce to the *consequentia materialis* of medieval logicians, a relation guaranteed by either the truth of an argument's conclusion or the falsehood of one of its reasons. For, if the conclusion is true, one could construct a true generalization of the argument's associated conditional by generalizing over some content in the reasons that does not occur in the conclusion. Consider for example the argument:

(3) *Mars is a planet, because trees grow.

A generalization of the associated conditional of this argument is that Mars is a planet if things of some kind grow. (For any kind K , if every K grows, then Mars is a planet—which is logically equivalent to the proposition that, if there is some kind K such that every K grows, then Mars is a planet.) And this generalization is true, simply because it has a true consequent: Mars *is* a planet.

Similarly, if a reason is false, one could construct a true generalization of the argument's associated conditional by generalizing over some content in the conclusion that does not occur in the reasons. Consider for example the argument:

(4) *Some cows are reptilian mammals, because Mars is a star.

A generalization of the associated conditional of this argument is that there are reptilian mammals if Mars is a star. (For any kind K , if Mars is a star, then some K s are reptilian mammals—which is logically equivalent to the proposition that there are reptilian mammals if Mars is a star.) And this generalization is true simply because it has a false antecedent: Mars is not a star.

The problem with such examples is that the generalization of the associated conditional is only *trivially* true: it is true either merely because any instance of it has a true consequent or merely because any instance of it has a false antecedent. A satisfactory account of material consequence must require that a generalization of the associated conditional be non-trivially true. Arguments like (3) and (4), where the initial universal quantifiers in the only true generalizations of the associated conditional bind variables that occur either only in the antecedent or only in the consequent, can be rejected as invalid if one adds to the requirement of a true generalization of the associated conditional that at least one initial universal

quantifier in the generalization binds a variable that occurs both in the antecedent and in the consequent. I shall call such a generalization a 'covering generalization' of the argument.

3. NON-TRIVIALITY

Requiring that an argument have a true covering generalization is however not enough to exclude cases where generalizations of the associated conditional are true only trivially. For covering generalizations too can be true only because they are trivially true. Consider for example the argument:

(5) *Mars is a planet, because Mars is a star with no mass.

A generalization of the associated conditional of this argument is that stars with no mass are planets. This generalization is true, but only because there are no stars with no mass. The same point can be made about generalizations that are true only because their instances always have a true consequent. Consider the argument:

(6) *Mars has mass, because it is visible from Earth in the night sky.

The generalization that any celestial object visible from Earth in the night sky has mass is true, but only because every celestial object has mass. (In this example, the variable bound by the universal quantifier has been given a range restricted to celestial objects; let us call such a range 'the universe of discourse'. Restriction of the universe of discourse to a well-defined class is clearly legitimate if there is well-supported background knowledge that the subject common to the claim and its supporting reasons belongs to that class, e.g. that Mars is a celestial object. Compare example (2) in the introduction, where the contentful valid form of argument restricts the range of the variable to planets.)

A first response to this problem is to require that an inference-licensing covering generalization not only be true but also have an instance with a true antecedent and an instance with an untrue consequent (Hitchcock, 1998). Alas, it turns out that this requirement is in one way too strict and in another way not strict enough.

As to its being too strict, consider some suppositional reasoning where we assume that some object has a property that we know nothing has—a non-instantiated property. Then consider some other property that according to the laws of physics, say, our imagined object would have if it had the non-instantiated property. Then it seems to follow that our imagined object has that other property. For example, we might suppose that a block of gold has a volume of one cubic metre. Given the density of gold (19,300 kilograms per cubic metre), this block would have a mass of 9.65 metric tonnes. So the following argument would be valid:

(7) This block of gold has a volume of one cubic metre. So its mass is 9.65 metric tonnes.

But the true covering generalization that would license the inference in this argument—the generalization that any block of gold with a volume of one cubic metre has a mass of 9.65 metric tonnes—has, we may suppose, no instance with a true antecedent. However, the absence of such an instance is not the sole reason why the covering generalization is true; another reason why the covering generalization is true is that gold has a density of 19,300 kilograms per cubic metre.

A similar point can be made about a property that every object in the universe of discourse has—an always instantiated property. Suppose that the laws of physics enable one to infer the presence of such an always instantiated property from some property belonging to all instances of some kind. For example, no physical object over the course of its existence both has and lacks mass. Now the laws of physics enable one to infer from the fact that something is an elementary particle of a certain sort what its mass is, and so *a fortiori* that it does not both have and lack mass. So someone might argue:

(8) Photons do not both have and lack mass, since photons are elementary particles.

The conclusion of this argument seems to follow, in virtue of the true covering generalization that no elementary particle both has and lacks mass. But this generalization has no instance with an untrue consequent. Nevertheless, it licenses the inference in our sample argument, because there is another reason why it is true, namely that every elementary particle has a definite mass (which may be zero, as in the case of photons).

The requirement of an instance with a true antecedent and an instance with an untrue consequent is however not only too strict. In another way, it is not strict enough. For, in some arguments with a true covering generalization that meets this requirement, the conclusion intuitively does not follow from the reasons given. Consider for example the argument:

(9) *Napoleon was short, because he ruled France and was exiled to Elba.

Here the covering generalization that all rulers of France exiled to Elba were short is true. Further, it meets the additional requirement: it has an instance with a true antecedent (the one concerning Napoleon) and an instance with a false consequent (any instance concerning someone who was not short, such as Giscard d'Estaing). But Napoleon's shortness obviously does not follow from the stated facts of his biography, which are epistemically irrelevant to his height (Hitchcock, 2011). Similar counter-examples can be constructed for any class of individuals that happen to share a property distinct from but not in any way determined by those defining the class. It is said, for example, that the world's largest gold bar weighs 250 kilograms¹, and we may suppose that, in view of the cost of making a gold bar, no gold bar will ever weigh more than 500 kilograms. If so, the generalization that

¹ <http://www.weirdasianews.com/2010/02/16/japan-pours-worlds-largest-gold-bar/>; accessed 2013 04 27.

no gold bar weighs more than 500 kilograms is true, and it meets the additional requirement: there are gold bars, and there are things that weigh more than 500 kilograms (for example, cars). But intuitively the following argument is not valid:

(10) *This block is a gold bar. So it weighs no more than 500 kilograms.

4. COUNTER-FACTUALS

How then can we rule out cases where covering generalizations are true only trivially without ruling out such apparently meritorious arguments as those concerning the weight of a hypothetical cubic-metre block of gold (example 7) and concerning a photon's not both having and lacking a mass (example 8), and without counting in such apparently unmeritorious arguments as those concerning Napoleon's height (example 9) and concerning the weight of a gold bar of unstated volume (example 10)? In examples 7 and 8, the two clauses proposed as means to block trivially true covering generalizations end up blocking covering generalizations that are true non-trivially. So we need to relax or replace those two clauses in such a way as to admit as inference-licensing at least some generalizations that are true non-trivially even though they either have no instances with a true antecedent or have no instances with an untrue consequent. In examples 9 and 10, on the other hand, the two clauses proposed as means to block trivially true covering generalizations failed to block true covering generalizations that intuitively do not license inferences. So we need to tighten up or replace those two clauses in such a way as to rule out, as not able to license an inference, true generalizations that satisfy the two clauses but whose consequent is, we might say, merely accidentally related to their antecedent, as Napoleon's height is to his rule and exile, and as weighing less than 500 kilograms is to being a gold bar.

What seems to be at issue in the counter-examples is whether the generalization's truth-value can be determined independently of knowing the truth-value of its instances. The cubic-metre-block-of-gold argument (7) and the photon argument (8) each have a covering generalization whose truth-value can be determined without taking into account whether it has instances with a true antecedent or whether it has instances with an untrue consequent, and *a fortiori* independently of the fact that the generalization in fact has in the one case no instances with a true antecedent and in the other case no instances with an untrue consequent. On the other hand, the true covering generalization of the Napoleon argument (9) can only be determined to be true by discovering that the one instance with a true antecedent happens to also have, as a matter of separately determined fact, a true consequent. The problem with the gold-bar argument (10) is more difficult to characterize. Perhaps the best account of the problem is that the truth-value of its true covering generalization, that no gold bar weighs more than 500 kilograms, can be determined only by reflecting on the rather extrinsic and accidental reasons for its lacking a counter-example. There is nothing about being a gold bar that precludes it from weighing more than 500 kilograms.

It appears that counter-examples of both sorts can be avoided by shifting to a requirement that an argument have a covering generalization that is true not just of

actual instances but also of hypothetical instances. Thus, in example (7), although there is no block of gold with a volume of one cubic metre, there *could* be such a block, and it *would* have a mass of 9.65 metric tonnes, in view of the density of gold. In example (8), although no physical object both has and lacks mass, there *could* be such an object, and it *would not* be a photon, in view of the fact that every photon has for its entire existence zero mass. In example (9), although there are no *actual* counter-examples to the minimal covering generalization that all rulers of France exiled to Elba were short, there is a *hypothetical* counter-example: Jacques Chirac, the former president of France, *need not*, and indeed *would not*, have been short if he had been exiled to Elba. In example (10), although there are (we might suppose) no *actual* counter-examples to the minimal covering generalization that no gold bars weigh more than 500 kilograms, there is a *hypothetical* counter-example: a gold bar *need not*, and indeed *would not*, weigh no more than 500 kilograms if it were more than twice the volume of the largest gold bar now in existence, which weighs 250 kilograms.

Does this revised account of material consequence rule out arguments like examples 5 and 6, where a covering generalization is true only trivially, i.e. merely because it has no instance with a true antecedent or merely because it has no instance with an untrue consequent? With the argument from Mars' supposedly being a star with no mass to its being a planet (5), the minimal covering generalization has hypothetical counter-examples, even though it has no actual ones: there *could* be a star with no mass, and it *need not*, and indeed *would not*, be a planet. Similar hypothetical counter-examples could be generated for other arguments where the only reason that its minimal covering generalization has no actual counter-examples is that it has no instance with a true antecedent. Consider for example the following parallel to example 5:

(11) *This figure is both a circle with a diameter of non-zero length and a square whose sides are of non-zero length. So it has an area of at least nine square centimetres.

For brevity, let us call a figure that is both a circle with a diameter of non-zero length and a square whose sides are of non-zero length a *squircle*. The minimal covering generalization of argument 11 is that a squircle has an area of at least nine square centimetres. Here again, although there are no actual counter-examples to this generalization, there are hypothetical ones: there *could* be a squircle (if space were different), and it *need not* have an area of at least nine square centimetres; indeed it *would not* have that large an area if, for example, its sides were two centimetres long. In contrast, the following argument is valid:

(12) This figure is a squircle. So it has a non-zero area.

Consider the covering generalization that a figure has a non-zero area if it is a squircle. This generalization is true, because its instances with a hypothetically true antecedent have a consequent that is also true in the hypothetical situation. There

could be a squircle, and it *would* have a non-zero area, namely the square of the length of its sides.

A similar vindication of the requirement that an inference-licensing covering generalization support counter-factual instances comes with a consideration of example 6, the argument from Mars' visibility from Earth in the night sky to its having mass. Any covering generalization of this argument has hypothetical counter-examples: there *could* for example be a celestial object that does not have mass, e.g. a hypothetical star with no mass, and it *need not* be invisible from Earth in the night sky. Here again the reason for finding this argument invalid seems to generalize to all cases where a covering generalization is true only because it has no actual instances with an untrue consequent. Consider the argument:

(13) *Hyenas are carnivores, so they are products of evolution.

Here the minimal covering generalization that all carnivores are products of evolution has no actual counter-examples, simply because all living organisms on Earth, whether carnivores or not, are products of evolution. There *could* however be an organism on Earth that was not a product of evolution, e.g. one created in a laboratory, and it *need not* be a non-carnivore. So the covering generalization, though true of the actual world, does not support counter-factual instances. Again, as with the two arguments 11 and 12 about a squircle, there is a parallel argument to argument 13 that is valid:

(14) Birds are descended from theropod dinosaurs, so they are products of evolution.

Here the minimal covering generalization that all descendants of theropod dinosaurs are products of evolution is true not only of actual instances but also of hypothetical instances. If there were currently living descendants of theropod dinosaurs other than birds, they too would be products of evolution. Also, if there were living organisms on Earth that were not products of evolution, they would not be descended from theropod dinosaurs.

To sum up: A conclusion follows from given reasons if the argument has a covering generalization with neither actual nor hypothetical counter-examples.

5. SEMANTICS

This conception of material consequence might have been arrived at more directly by reflecting on the fact that following *necessarily* requires a form of argument that lacks not only *actual* counter-examples but also *possible* (i.e. hypothetical) counter-examples. With formal consequence, there is no need to consider hypothetical situations, since set-theoretic reasoning can establish that the standard model-theoretic conception will generate all the possibilities against the background assumption of the actual world as it is (Sher, 1996). With material consequence, on the other hand, considering all the possible counter-examples to a contentful form of argument requires attention to hypothetical situations, since substitutions or

interpretations against the background assumption of the world as it is will in general not exhaust the possibilities.

Attention to hypothetical situations, however, has its own theoretical problems. How is one to determine that, in some instance of a covering generalization, an antecedent that is actually untrue nevertheless *could* be true? If it could, how is one to determine whether the consequent *would* be true in such a hypothetical situation? Similarly, how is one to determine that the actually true consequent of a covering generalization's instance *could* nevertheless be untrue? If it could, how is one to determine whether the antecedent *would* be untrue in such a hypothetical situation? In the cases we examined, it seemed straightforward to make the required determinations. We had no difficulty in thinking of how an actually untrue antecedent *could* be true: a block of gold *could* have a volume of one cubic metre (7), another ruler of France *could* have been exiled to Elba (9), there *could* be a star with no mass (5), a gold bar *could* have a volume more than twice that of the largest gold bar now in existence (10), there *could* be a squircle (11 and 12), a carnivore *could* be created in a laboratory (13), and there *could* be a living descendant of theropod dinosaurs that was not a bird (14). And we had no difficulty in figuring out whether in such a hypothetical situation the consequent *would* be true: the block of gold (7) *would* have a mass of 9.65 metric tonnes, the other ruler of France exiled to Elba (9) *need not* have been short, the star with no mass (5) *need not* and indeed *would not* be a planet, the gold bar (10) *need not* and indeed *would not* weigh no more than 500 kilograms, the squircle *need not* have an area of at least nine square centimetres (11) but *would* have a non-zero area (12), and the laboratory-created carnivore (13) *need not* be a product of evolution but the non-avian living descendant of theropod dinosaurs (14) *would*. Similarly, we had no difficulty in thinking of how an actually true consequent *could* be untrue: a celestial object *might* have no mass (6) and a physical object *might* at different periods in its history have mass and lack mass (8). And we had no difficulty in figuring out whether in such a hypothetical situation the antecedent *would* be untrue: the celestial object with no mass (6) *need not* be invisible from Earth in the night sky, and the physical object that both had and lacked mass (8) *would not* be a photon.

In general, in these cases we are relying on law-like generalizations whose truth-value is determined by the presence or absence of a law of nature, a theorem of geometry, or conventional meanings of terms. Such law-like generalizations are true if and only if they support counter-factual instances. Could one then reduce the concept of material consequence to the existence of a true law-like covering generalization?

It appears not. For we can think of arguments whose conclusion intuitively seems to follow, in virtue of a true covering generalization that supports counter-factual instances, even though the generalization is not law-like. Consider for example the following argument:

(15) President Obama lives in Washington, because he lives in the White House.

One could imagine this argument addressed to someone who thought that Obama merely used the White House as his office, and lived somewhere else, commuting to work. Such a person might not be sure what city Obama lived in. Argument 15 should be a convincing argument for them, as long as they accept the supporting reason on the say-so of its author. For the conclusion obviously follows, in virtue of the covering generalization that whoever lives in the White House lives in Washington. And this covering generalization supports counter-factuals: if Mitt Romney lived in the White House, he would live in Washington. But there is no law of nature or mathematical theorem or conventional definition underpinning this generalization. What makes it true, even in counter-factual instances, is the combination of the purely contingent fact that the plot of ground in which the White House is situated, at 1600 Pennsylvania Avenue, is located in the city of Washington, D.C., with the law-like transitivity of the relation of being within: if x lives in building y and y is located in city z , then x lives in city z .

At the 2011 OSSA conference, in discussion of a presentation subsequently published as (Hitchcock, 2011), Robert Ennis challenged my claim that example 15 had a covering generalization that supported counter-examples, on the ground that, for example, if Vladimir Putin lived in the White House, it would be in Moscow. This challenge raises the vexed question of how in general we determine the truth-value of counter-factuals. If Vladimir Putin did live in the White House, in what city would he be living?

There are two aspects to such a question. The first is the determination of whether there even *could* be a hypothetical situation in which the false antecedent of the given instance of the covering generalization was true: *could it be* that Vladimir Putin lived in the White House? The second is the determination of whether in such a hypothetical situation the consequent of the given instance of the covering generalization *would have to be* true: *would* Putin live in Washington?

There are currently at least two major candidates for a theoretical analysis of the truth-conditions for counter-factuals: the closest-world semantics of David Lewis (1973) and the structural model semantics of Judea Pearl (2000) and his collaborators. According to Lewis's closest-world semantics, a counter-factual conditional is true if and only if, in every possible world closest to the actual world, the consequent is true if the antecedent is true. In other words, there is no closest possible world where the antecedent is true and the consequent untrue. The difficulty with this semantics is with the construction of a measure of similarity between worlds that would enable us to identify the non-actual worlds that are similar to each other and minimally close to the actual world. In terms of our example, a possible world in which Putin lives in the White House is closer to the actual world if in it Putin lives in Moscow than if in it he lives in Washington, provided that all other things in the two possible worlds are the same. But they would not be the same in those two possible worlds. A world in which Putin lived in the White House and lived in Moscow would require the White House to be in Moscow rather than Washington, assuming that the phrase 'live in' here has the force of picking out the location of a person's principal residence. A world in which Putin lived in the White House and lived in Washington would have the White House still in Washington. Is a world in which Putin lives in the White House and it is in

Moscow closer to the actual world than one in which he lives in the White House and has moved to Washington? Or is it farther away? Or are they equally close? In principle, one possible world is just as close to the actual world as another possible world if each of the two possible worlds has the same number of atomic propositions with a truth-value different from their truth-value in the actual world. And one possible world is closer to the actual world than another if it has fewer atomic propositions than that other world with a truth-value different from their truth-value in the actual world. The difficulty in such examples is to determine which atomic propositions would have a different truth-value in a given possible world. What else would, or might, be the case in a possible world in which Putin lived in the White House? The situation is so different from the actual world that we cannot begin to work out the other changes that would have to occur. Perhaps the result of applying Lewis's semantics to our example would be that there is no closest world in which Putin lives in the White House, and hence that *a fortiori* there is no closest world in which Putin lives in the White House and does not live in Washington. If this is the correct result, then on Lewis's closest-world semantics it is (vacuously) true that Putin would live in Washington if he lived in the White House.

Lewis's closest-world semantics seem to give the intuitively correct result in cases where the counter-factual situation needs relatively few adjustments to our actual situation, so that there is a closest world where the antecedent is true. In a close possible world where Mitt Romney lived in the White House, Romney would have won the U.S. presidential election in November 2012. Other ways in which we could imagine Romney coming to live right now (in May 2013) in the White House would involve far more changes to the truth-value of atomic propositions than those involved in supposing that Romney had won the election instead of Obama. And, if Romney had won the election and so now lived in the White House, he would live in Washington. So, in the closest worlds in which Romney lives in the White House, he also lives in Washington. Hence, on Lewis's closest-world semantics Romney would live in Washington if he lived in the White House.

Pearl's structural model semantics interprets counter-factual conditions in terms of a hypothetical change to equations in a causal model defined by functional causal relationships among variables (Pearl, 2000, p. 205). The change to the equations simulates an external action or spontaneous change that alters the course of history, with minimal change of mechanisms. A causal model of how people come to live in the White House, for example, would include a number of pathways: election as U.S. President and subsequent inauguration, becoming in one way or another part of the immediate co-habiting family of someone who becomes or is U.S. president, joining that part of the White House cleaning and cooking staff that lives in the White House, being invited to stay temporarily as a special guest in the White House. In this causal model, such variables as the geographical location of the White House and the boundaries of the city of Washington, D.C., would be background (exogenous) variables determined by factors outside the model. The counter-factual situation that Vladimir Putin lives in the White House, given that he is the President of Russia, could only come about in the causal model by his staying temporarily as a special guest; we can exclude as not even remote possibilities his entering into a bigamous union with Michelle Obama (or with Barack) or becoming a member of

President Obama's live-in cleaning or cooking staff. But, in any case, the geographical location of the White House and the boundaries of the city of Washington, D.C. are exogenous background variables, whose value remains the same in any hypothetical situation where someone is assumed to live in the White House who actually does not live there. Since the White House is currently located at 1600 Pennsylvania Avenue and that location is within the boundaries of the city of Washington, D.C., then on Pearl's structural model semantics Vladimir Putin would live in Washington if he lived in the White House.

It should be noted that the consequence relation in virtue of which Obama's living in Washington follows from his living in the White House is what medieval logicians called a *consequentia materialis ut nunc*. It holds only "*ut nunc*", i.e. for now. Future changes to the geographical boundaries of Washington, D.C., or construction of a new "White House" outside those geographical boundaries, could bring it about that people who live in the White House do not live in Washington. The conclusion of argument 15 therefore follows not only materially rather than formally, but also for the time being rather than for all time.

Here is another argument where the conclusion appears to follow in accordance with a covering generalization that supports counter-factual instances, even though it is not law-like:

(16) Every human being is mortal, so Socrates is mortal. (Freeman, 2011, p. 183)

The covering generalization that Socrates possesses every property that every human being possesses (for every *F*, if every human being is *F*, then Socrates is *F*) supports counter-factual instances: if every human being were kind, then Socrates would be kind; if every human being had four stomachs, then Socrates would have four stomachs; and so on. But the generalization is not law-like. It is logically equivalent to the proposition that Socrates is a human being, which is a contingent particular fact—contingent because, for example, he might have been an alien. As with the previous example, Pearl's structural model semantics appears to give a better account of why the generalization supports counter-factual instances than does Lewis's closest-world semantics. A causal model of the mechanisms that make Socrates mortal would appeal to various components of his humanity, and ultimately to genetic factors inherited from his parents. A counter-factual instance would alter those aspects of the causal model that affect the variable at issue: the factors responsible for making a human being kind or cruel or indifferent, the mechanisms responsible for the formation of a single stomach in each human being, and so on. The status of Socrates as a human being would not change with such changes, since his species is basic to who he is.

In contrast, if the argument's reason appealed to some causally irrelevant property whose possessors just happened to be mortal, the conclusion would not follow. Consider for example the argument:

(17) *Every two-legged organism is mortal, so Socrates is mortal.

It is true that Socrates possesses every property that every two-legged organism possesses, but this covering generalization (for every F , if every two-legged organism is F , then Socrates is F) does not support counter-factual instances. Suppose for example that every two-legged organism is a reptile. The causal model showing the evolution of species on Earth would then need to be changed to accommodate this counter-factual assumption. Given the direction of the causal mechanisms that have produced both two-legged reptiles (some dinosaurs, birds) and two-legged mammals (human beings, hominids), the change to the causal model to make only reptiles two-legged would involve an evolutionary history in which the ancestors of human beings did not make the shift from being two-legged to being four-legged. In that case, Socrates would be four-legged.

In these three examples, it appears that Pearl's structural model semantics is more easily applicable than Lewis's closest-world semantics to the determination of the truth-value of a singular counter-factual conditional, and that it gives intuitively correct results.

6. SUMMARY

This paper has explored the conditions under which the conclusion of an argument follows materially from the reasons given, where following materially is understood as following in accordance with a contentful valid form of argument. Validity of such a contentful form obviously requires the truth of the corresponding universal generalization of the argument's associated conditional, the material conditional whose antecedent is the conjunction of the argument's premises and whose consequent is the argument's conclusion. This generalization needs to be a covering generalization, in the sense that at least one variable bound by its initial universal quantifiers occurs both in the antecedent and the consequent. But the requirement of a true covering generalization is not enough to rule out as invalid arguments whose true covering generalizations are only trivially true—i.e. true either only because the generalization has no instances with a true antecedent or true only because it has no instances with an untrue consequent. It is tempting to rule out such trivial cases by requiring that an inference-licensing covering generalization have at least one instance with a true antecedent and at least one instance with an untrue consequent. But this restriction both rules out some intuitively valid arguments and fails to rule out some intuitively invalid arguments. The restriction that appears to give just the right results is to require that an inference-licensing generalization supports counter-factual instances. The requirement of support for counter-factual instances can be motivated as not merely *ad hoc* by attending to the *explicandum* of an account of material consequence: the conclusion is to follow *necessarily* from the reasons given.

Law-like generalizations support counter-factual instances. But so, as it turns out, do some generalizations that are not law-like. There are at least two accounts available of the conditions under which a counter-factual singular conditional is true: the closest-worlds semantics of David Lewis (1973) and the structural model semantics of Judea Pearl (2000). It appears from exploration of examples that Pearl's structural model semantics is more easily applied than Lewis's closest-world

semantics to determining the truth-value of counter-factual instances of a covering generalization, and that it gives intuitively correct results.

REFERENCES

- Bolzano, B. (1972). *The Theory of Science: Die Wissenschaftslehre oder Versuch einer Neuen Darstellung der Logik*. Berkeley and Los Angeles: University of California Press. German original first published in 1837.
- Brandom, R. (1994). *Making It Explicit*. Cambridge, MA: Cambridge University Press.
- Brandom, R. (2000). *Articulating Reasons: An Introduction to Inferentialism*. Cambridge, MA: Harvard University Press.
- Freeman, J.B. (2006). Systematizing Toulmin's warrants: an epistemic approach. In D. Hitchcock & B. Verheij (Eds.), *Arguing on the Toulmin Model: New Essays in Argument Analysis and Evaluation* (pp. 87-101, Ch. 7). Dordrecht: Springer.
- Freeman, J.B. (2011). *Argument Structure: Representation and Theory*. Dordrecht: Springer.
- George, R. (1972). Enthymematic consequence, *American Philosophical Quarterly*, 9, 113-116.
- George, R. (1983). Bolzano's consequence, relevance, and enthymemes, *Journal of Philosophical Logic*, 12, 299-318.
- Hitchcock, D. (1985). Enthymematic arguments, *Informal Logic*, 7(2 and 3), 83-97.
- Hitchcock, D. (1998). Does the traditional treatment of enthymemes rest on a mistake?, *Argumentation*, 12, 15-37.
- Hitchcock, D. (2011). Inference claims, *Informal Logic*, 31, 191-228.
- Lewis, D.K. (1973). *Counterfactuals*. Cambridge, MA: Harvard University Press.
- Pearl, J. (2000). *Causality: Models, Reasoning, and Inference*. Cambridge: Cambridge University Press.
- Peirce, C.S. (1955). The fixation of belief. In J. Buchler (Ed.), *Philosophical Writings of Peirce* (pp. 5-22, Ch. 2). New York: Dover. First published in 1877.
- Pinto, R.C. (2006). Evaluating inferences: the nature and role of warrants. In D. Hitchcock & B. Verheij (Eds.), *Arguing on the Toulmin Model: New Essays in Argument Analysis and Evaluation* (pp. 115-143, Ch. 9). Dordrecht: Springer.
- Ryle, G. (1950), 'If', 'So', and 'Because'. In M. Black (Ed.), *Philosophical Analysis* (pp. 323-340), Ithaca: Cornell University Press.
- Sellars, W. (1948). Concepts as involving laws and inconceivable without them, *Philosophy of Science*, 15, 287-315.
- Sellars, W. (1953). Inference and meaning, *Mind*, 62, 313-338.
- Sher, G.Y. (1996). Did Tarski Commit 'Tarski's Fallacy'?, *The Journal of Symbolic Logic*, 61, 653-686.
- Toulmin, S.E. (1958). *The Uses of Argument*. Cambridge: Cambridge University Press.