Denying the Antecedent: The Fallacy That Never Was, or Sometimes Isn't?

LUÍS DUARTE D'ALMEIDA

EUAN MACDONALD

University of Edinburgh
School of Law
Old College, South Bridge
EH8 9YL Edinburgh
United Kingdom

luis.duarte.almeida@ed.ac.uk
euan.macdonald@ed.ac.uk

Abstract: In this paper we examine two challenges to the orthodox understanding of the fallacy of denying the antecedent. One challenge is to say that passages thought to express the fallacy can usually be given an interpretation on which they express valid arguments, entitling us to query whether the fallacy is commonly, if ever, committed at all. We discuss this claim in Section 1. The second challenge comes from those who think that there are legitimate uses of denying the antecedent that have traditionally been overlooked. In Section 2 we propose a general test for claims of this sort, and assess three versions of this view.

Résumé: Dans cet article, nous examinons deux défis à la compréhension orthodoxe du sophisme de nier l'antécédent. Un défi consiste à dire que généralement notre interprétation d'un passage que nous jugeons fallacieux peut nous donner un argument valide, ce qui nous donne le droit de nous demander si l'erreur est commise souvent, si jamais du tout. Nous discutons cette allégation dans la section 1. Le deuxième défi vient de ceux qui affirment qu'il y a des usages légitimes de nier l'antécédent qui ont traditionnellement été négligés. Dans la section 2, nous vous proposons un test général pour les affirmations de ce genre, et évaluons trois versions de ce point de vue.

Keywords: affirming the consequent; argument reconstruction; charity; denying the antecedent; fallacies

> There is a moralistic argument for belief in God, which was popularized by William James. According to this argument, we ought to believe in God because, if we do not, we shall not behave well . . . It is, in the first place, very doubtful whether belief in God has all the beneficial moral effects that are attributed to it. Many of the best men known to history have been unbelievers. John Stuart

Mill may serve as an instance. And many of the worst men known to history have been believers. Of this there are innumerable instances. Perhaps Henry VIII may serve as typical. (Bertrand Russell, 1997/1952, p. 545-6.)

1. The fallacy that never was?

In a 1994 article, Michael Burke sought to defend the view that the formal fallacy of denying the antecedent is rarely if ever committed, at least "among arguers sufficiently mature to have published writings" (Burke 1994, p. 27). He was neither the first nor the last to argue for a claim of this kind (cf. the discussion and references in Finocchiaro 1981; Adler 1994; Krabbe 1995; Moldovan 2009), but his remains the most sustained and best known defence of this position. It will form the focus of our discussion in this section.

Burke takes himself to have offered reasons of varying strength to support two claims. One is that "the fallacy of denying the antecedent is not common in published writings"; the other, that "published passages cannot often be *charged* with denying the antecedent" (1994, p. 27). We will argue both that Burke has failed to provide us with adequate reasons to warrant these claims, and that there is reason to think that the claims are false.

Burke's argument is premised upon a "principle of fairness", which he styles a "very weak principle of exegetical charity": that we should not prefer a fallacious to a non-fallacious interpretation of an argument "unless the balance of textual, contextual, and other evidence" favours such an interpretation (p. 24). This seemingly plausible condition poses an obstacle "so serious", Burke says, "that I know of no real argument that can justifiably be charged with denying the antecedent" (p. 23). The reason is that it will almost always be possible to offer an alternative and at least equally plausible interpretation on which the argument does not come across as fallacious.

In order to illustrate this point, Burke examines passages commonly offered in textbooks as examples of the fallacy. Some of these examples are "concocted", while some are "real"; but Burke proposes to regard them all "as real utterances", and claims that such passages "cannot be fairly charged with the fallacy of denying the antecedent". The example he discusses in greater detail is one of the "concocted" ones:

If capital punishment deterred murder, it would be justified. Since it doesn't, it isn't.

Burke remarks that it is "at least plausible . . . to take the argument to be an *enthymematic instance of modus ponens* (or of *modus tollens*, depending on the formulation of the unstated conditional)" (p. 24). That is because there is "no adequate reason" to consider the speaker's stated conditional "as part of the argument"; after all, "it is anything but unusual for a passage containing an argument to contain statements that are not part of the argument" (p. 24). In this example, "it is at least plausible to ascribe to the conditional some other role"—the rhetorically and dialectically important role of "clarifying the nature of the arguer's objection to capital punishment, of making clear that the arguer opposes capital punishment *only* because the arguer believes it doesn't deter murder". We can accordingly view the speaker's argument as containing only one stated premise—the premise that capital punishment does not deter murder—together with the "unstated premise: If capital punishment doesn't deter murder, then it isn't justified" (p. 25). But then the argument is valid: it is an instance of *modus ponens*.

This is not an isolated example: the point is meant to hold for all the passages found in textbooks as illustrating the fallacy of denying the antecedent, and indeed more generally. The fairness constraint, Burke says (p. 26), "is satisfied in none of the examples we have considered—and in none of which I am aware":

> It would be unfair to view them as formally fallacious unless there were some special reason for doing so. And in none of the cases known to me does either the context of the passage or any other salient circumstance supply such a reason (p. 27).

In every actual example Burke considers, then, the stated conditional can invariably be understood, "naturally and plausibly", as playing some "non-premisory" role; and unstated premises (and sometimes unstated conclusions) can be identified, warranting—in the absence of "contextual or circumstantial evidence" to the contrary—the interpretation on which the argument is not formally fallacious (p. 25).

We agree that the textbook examples examined by Burke would—if taken as examples of actual arguments—be better understood as formally valid arguments rather than as instances of denying the antecedent.[1] How, then, can we disagree with his

---

[1] It does not follow that we agree that "concocted" examples like the capital punishment one, when found in textbooks, are necessarily poor examples of denying the antecedent. Being concoctions, these examples are presumably being offered by textbook authors as depicting situations in which someone

conclusion? Surely, if textbook examples—presumably chosen precisely because they are exemplary—cannot be justifiably charged with instantiating the fallacy, do we not have good grounds to suppose that this will be true of most, if not all, published arguments?

Burke himself asks the key question: "Are [these] examples . . . in some way unrepresentative" (p. 26)? The basis for our disagreement is that we have no reason to think that they are representative, and some reason to think that they are not.

The first step is to examine the reason Burke gives for thinking that the examples are indeed representative. He begins by noting that "the fallacy of denying the antecedent is instantiable by passages of many different patterns", and proposes that we split these into two groups: those (which he calls "B patterns") that "contain a textual feature indicating . . . that the conditional contained by the passage is a premise and that the passage, therefore, is an instance of denying the antecedent"; and those (the "A patterns") in which there is no such textual indication. He goes on to give some examples of both (p. 26):

A Patterns (relatively common)	**B Patterns** (uncommon)
(a) If A, then B. Not A. So, not B.	(e) Not A. If A, then B. So, not B.
(b) If A, then B. But not A. So, not B.	(f) If A, then B. And not A. So, not B.
(c) If A, then B. Since not A, not B.	(g) Not A. Since B if A, not B.
(d) If A, then B. But since not A, not B.	(h) If A, then B. And since not A, not B.
	(i) Not A. But if A, then B. So, not B.
	(j) Not A. But since B if A, not B.
	(k) Not B, since not A. And if A, then B.

does actually mean to combine a conditional and the denial of the antecedent to infer the denial of the consequent. This is not to say that the examples are plausible, or that such situations are frequent. Yet surely it is the prerogative of those authors to cook up wholly imaginary situations that illustrate precisely what they mean them to illustrate. These are, after all, *their* examples—not Burke's. So Burke can only say that in these examples there is "no reason" to regard the conditionals as premises, or that there is "no evidence" supporting the fallacy charge, by filling up those authors' examples with details that no longer allow them to stand as plausible illustrations of the fallacy of denying the antecedent; and to do that is to change rather than to "interpret" such examples.

B Patterns (uncommon)

(l) Not B, since B if A. And not A.
(m) If A, then B. So not B, since not A.

Burke's argument from here is simple. Passages that instantiate A patterns are relatively common, but fairness almost always requires that we interpret them as formally valid arguments. Argumentative passages that instantiate B patterns, on the other hand, are "almost *surely*" (p. 26) to be interpreted as fallacious—because the text makes it clear that the explicit conditional is indeed a premise of the argument—but we rarely, if ever, find them, and their instantiations sound "unnatural". This provides us, Burke says, with good reason to think that "published passages cannot often be *charged* with denying the antecedent" (p. 27).

This argument, however, seems to us incapable of supporting its conclusion—and for considerations that have nothing to do with the small size of the sample with which Burke is working. ("In examining 81 books, as well as numerous articles on fallacies", Burke "found 9 real passages said to be instances of denying the antecedent", none of which instantiates a *B* pattern: p. 27.) In order to explain why, it will help to be clearer about what it means to "commit" the fallacy of denying the antecedent, and about the role played by "patterns" in identifying it.

An arguer commits, with regard to a certain conclusion, the fallacy of denying the antecedent, when her argument for that conclusion can be justifiably reconstructed as a materially invalid instance of one of the following patterns:

> *Denying the antecedent: propositional form*
> (denying the antecedent *stricto sensu*)
>   (1) If p, then q.
>   (2) It is not the case that p.
>   Therefore (from (1) and (2)),
>   (3) It is not the case that q.

> *Denying the antecedent: universal form*
>   (1) For every x: if x is S, then x is P.
>   (2) It is not the case that a is S.
>   Therefore (from (1) and (2)),
>   (3) It is not the case that a is P.

These can be used as patterns of representation: as templates under which the components of the reconstructed argument can

all be made explicit and clearly displayed. But there is no reason to think that when a writer or speaker does give an argument that can be charitably reconstructed as an instance of one of these patterns, her actual utterances—her text or spoken words—will display a grammatical form that is anything like any of the two patterns. In particular, in order for a conditional premise like (1) to be charitably attributed to an arguer, what is required is that she can be taken to be presenting p (or the claim that a is S) as a sufficient condition of q (or of the claim that a is P)—however that is actually articulated by her. Burke's argument is predicated upon the assumption that any argument that can even be considered a candidate for instantiating the fallacy of denying the antecedent will have been put forth in a passage that contains an explicit grammatical conditional in the "If . . . then . . ." form; but that assumption seems to be false.

To say that an arguer commits the fallacy of denying the antecedent is to say that she is committed to the claim that the negation of a sufficient condition is sufficient for the negation of that which was so conditioned. That is a logical error, and one reason that we reconstruct and represent arguments in the manner that we do—with patterns like the ones above—is to bring out this error as clearly as possible.[2] But there is no limit on the ways in which the arguer's logical error could be expressed by her, and no reason at all to suppose that her utterances will have to conform to a particular *grammatical* pattern. Our point, to be clear, is not that we are entitled to "read in" a commitment to a conditional in the absence of strong textual support for such an attribution. Our point is merely that there are many different ways in which an arguer can make it textually clear that she is committed to a conditional proposition—a proposition that we can properly understand and render as an "If . . ., then . . ." claim—that do not involve the actual expression by the arguer of a grammatical "If . . ., then . . ." sentence.

Here is a simple illustration. Suppose someone charged Burke with committing in his paper the fallacy of denying the antecedent, attributing to him the following argument:

> (1) If instances of the listed B patterns were common in published writings, we would have reason to think

---

[2] Reconstruction as a denial of the antecedent is only one way of bringing out this flaw. There will be others: by restating the conditional premise as its logically equivalent contrapositive, for example, an argument that can be charitably reconstructed as a denial of the antecedent can also be charitably reconstructed as an example of the formal fallacy of affirming the consequent (although there may be reasons other than charity for preferring one reconstruction over the other).

        that the fallacy of denying the antecedent can often be justifiably charged.

(2) Instances of the listed B patterns are not common in published writings.
Therefore (from (1) and (2)),

(3) We have no reason to think that the fallacy of denying the antecedent can often be justifiably charged.

Burke might reply that this charge would be uncharitable. But certainly he would not be able to deny that he argues in his paper for the conclusion in (3), and that he does rely on something like premise (2) to support it. Burke is also clearly committed to (1), even though he never states it in this manner. (What he says is that an argument instantiating a B-pattern "almost s*urely* is" an example of denying the antecedent: p. 26, emphasis in original.) And it does not seem plausible to attribute to him the converse of the conditional in (1), which would make the argument valid. (Recall that "B patterns" cover *any* patterns that include a "feature indicating . . . that the conditional contained by the passage is a premise": p. 26.) So Burke's reply would have to be that this particular reconstruction fails to capture his argument for that conclusion. Note, however, that there is no single, discrete passage in his paper—no neatly patterned sequence of sentences—in which that argument is expressed. The charity, or lack thereof, of the reconstruction does not turn on the existence of any such passage. The reconstruction above can be offered as a reconstruction of *an* argument made by Burke *in* his paper, even if the material for the reconstruction is culled from different paragraphs or even different sections. Perhaps this reconstruction *is* uncharitable (we are not fully sure); but that has nothing to do with the grammatical patterns of any fragment of Burke's text.

    There is another, related consideration. Burke has limited his search for examples of denying the antecedent to *precisely those instances in which we would least expect to find them*: passages whose grammatical form already closely resembles the form of the reconstructed (fallacious) argument, and in which any formal fallacy would therefore be at its very clearest. Textbooks may prefer such examples for exactly this reason—which makes textbooks a singularly poor source for the broader claims about argumentation in general that Burke is seeking to defend. Burke's (already small) sample of arguments is rendered unusable by a fatal dose of selection bias.

    If our argument so far is correct, Burke has given us no reason to accept his conclusion that published arguments can rarely if ever be justifiably charged with the fallacy of denying the an-

tecedent. Indeed, if there is no necessary correspondence between the grammatical form of an argument as actually expressed by an arguer, and the form in which we reconstruct and represent that argument, it is difficult to see what—other than a painstaking and comprehensive analysis of a huge number of different texts—*could* provide a reason in support of a strong, general claim like Burke's. While we can certainly accept Burke's insight that many of what may seem, superficially, to be examples of the fallacy of denying the antecedent will turn out, when properly interpreted, to be valid arguments, his interpretative method fails to generate reasons in support of his broader claim.

If, moreover, we could readily find examples of arguments that can be justifiably charged with committing the fallacy—even by Burke's own "fairness" standard—then we would have positive reason to think that Burke is wrong. And given what we just said about selection bias, that reason would be strengthened if we could find such examples in published passages whose grammatical form does track relatively closely one of the formal patterns of denying the antecedent. If this fallacy can be justifiably charged even when at its clearest, then we have, *ceteris paribus*, reason to think that it may be more prevalent when the grammatical structure of the passage does not resemble the formal pattern of reconstruction.

We have already given one such example: the second argument from the passage by Russell that we used as our epigraph. Here is how it can be reconstructed:

(1) For every *x*: if *x* does not believe in God, then *x* does not behave well.
(2) Henry VIII believed in God.
    Therefore (from (1) and (2)),
(3) Henry VIII behaved well.

Does Burke's principle of fairness require that we interpret Russell's passage as expressing a valid argument at this point?

Let us assume there are two constraints here: we must have good reason to suppose that the stated conditional is part of the argument, and not performing some other (e.g. rhetorical) role; and we must have good reason to think that the arguer was not also committed to or relying on the converse of that conditional (or, indeed, on any other conditional that could rescue the argument's validity). Note that these constraints are in fact significantly *stronger* than those proposed by Burke, whose explicit view is that the fact that the stated conditional is part of the argument is itself sufficient grounds to conclude that the argument

"almost surely" is an instance of denying the antecedent, and can be justifiably charged as such (p. 26).

But even this strengthened version of Burke's principle of fairness cannot, in our view, rescue Russell's argument. First, we have very good reason to suppose that here Russell must indeed intend the explicitly stated conditional to be part of the argument: that conditional is the "moralistic" proposition that Russell means to *refute* (but could only do so with this argument if denying the antecedent were a valid form: Russell's point is that since the conclusion is false, and premise (2) true, premise (1) cannot but be false). Secondly, we have equally good reasons—reasons of charity—to reject the idea that Russell is in fact relying on the converse of the explicit conditional or on the corresponding biconditional (either of which would render the argument valid). It is, in the first instance, highly implausible that anyone ever held mere belief in God as a *sufficient* condition of good behaviour, and it would thus have been wholly uncharitable for Russell to have attributed that view to his interlocutors. Perhaps more importantly, Russell is clear that he is here offering a criticism of William James's moralistic argument, which had been set forth in the 1891 essay "The Moral Philosopher and the Moral Life".[3] In that essay, James makes the claim that religious belief brings with it certain psychological benefits that are necessary for what he calls "strenuous" moral living. He never claims that religious belief is *sufficient* to ensure that all believers will live strenuously in this manner, and he is explicit that even those who do live strenuously might not do so in accordance with the correct moral principles ("[E]xactly what the thought of the infinite thinker may be is hidden from us even were we sure of his existence; so that our postulation of him after all serves only to let loose in us the strenuous mood": see James 1891, p. 162). Russell's own stated target, then, cannot be charitably taken to have asserted that belief in God is sufficient for good behaviour. Indeed, on no plausible reading of James's essay is he committed to *any* general proposition to which the behaviour of Henry VIII would be an apt counterexample. Rather than assume that Russell was attacking a straw man, or had made a basic error in his understanding of James's text, it is more charitable—or at least no less charitable—to take *him* to have denied the antecedent.

But if Russell's passage does not strike you as fully clear, here is a more recent example susceptible to a similar analysis—an argument by Nick Cohen writing in *The Guardian*. It shares

---

[3] See Slater's annotation in Russell (1997/1952, p. 768).

with Russell's the characteristic of being an attempted refutation by counterexample (Cohen 2014):

> I'm losing count of and patience with the apologists who tell me there would be no morality without religion. The failure of the serious press and BBC to question this is as shocking as it is depressing. We are almost 150 years on from the moment in 1867 when Matthew Arnold heard the sea of faith's "melancholy, long, withdrawing roar" on Dover Beach. Are religious writers suggesting mid-Victorian Britain was a more moral country in its treatment of women, homosexuals and the poor?

Here, we do not have an explicit, grammatical conditional, and a little more rewriting of the premises is necessary to bring out the form of the argument. But the claim that Cohen is seeking to refute is clearly (as was Russell's) the assertion that religion is a necessary condition of morality. Cohen tries to refute this by discussing countries, rather than individuals, but that does not significantly alter the analysis:

> (1) For every x: if x is not a religious country, then x is not a moral country.
> (2) Mid-Victorian Britain was a religious country.
>   Therefore (from (1) and (2)),
> (3) Mid-Victorian Britain was a moral country.

Again, the claim that religion is a necessary condition of morality is clearly part of the argument: there is no other way to understand it. Likewise, it would be uncharitable to suggest that Cohen was in fact relying—in order to attack it—on the extremely implausible conditional premise that asserted that religion was *sufficient* for morality. As with Russell's passage, the most charitable interpretation here is that Cohen has committed the fallacy of denying the antecedent.

   (Cohen's last sentence might suggest a different interpretation: he could be taken to be attributing to his opponents not merely the view that

> (1*) Religion is a necessary condition of morality,

but also something like the view that

> (2*) If religion is a necessary condition of morality, then (if a country x is more religious than another country y, then country x is more moral than country y).

Cohen could then be taken to be denying the *consequent* of this conditional—the conditional in (2*)—by pointing out, by way of counter-example, that mid-Victorian Britain, though more religious than contemporary Britain, was nonetheless less moral than contemporary Britain; and that would enable him, by *modus tollens*, to reject (1*). Under this interpretation, Cohen's argument would not be formally fallacious; but it would rely on the implausible and highly uncharitable supposition that anyone who believes (1*) is thereby committed to (2*). Defenders of (1*) could simply reply that contemporary Britain, though less religious than mid-Victorian Britain, is nonetheless still a sufficiently religious country—that it meets whatever threshold of religiosity a country must meet in order to be moral—and that (1*) is of course consistent with even the most religious of countries being utterly immoral. This alternative interpretation, then, seems to us not to be differentially justified: our ascription to Cohen of a formally fallacious argument is no less charitable towards him—to say the least—than a reading that acquits him from the fallacy only at the cost of finding *him* guilty of being extremely uncharitable towards his opponents.)

We thus have no reason to think, and some reason to deny, that published arguments seldom commit the fallacy of denying the antecedent.

2. The fallacy that sometimes isn't?

In this section, we consider certain arguments advanced in three recent articles. Prominent in each of these articles is the claim that there are, in certain contexts, "legitimate" instances of the argumentative form of denying the antecedent, and that attending to these instances requires a revision of traditional approaches to the fallacy. It is *this* claim that is of interest to us in what follows.[4] And we propose the following test for the success of any version of such a claim: an instance of the argumentative form of denying the antecedent is legitimate if and only if (a) the argument is cogent, and (b) the cogency of the argument does not depend on the addition of other premises or assumptions, or on the alteration of the conclusion to something other than the negation of the consequent of the conditional. We will take the first constraint to be met if the argument provides *some reason*

---

[4] Each of the articles also makes a range of interesting points in relation to other matters—offering, in particular, analyses of what is actually going on in arguments that might at first be thought to resemble denials of the antecedent—but these are not our concern here.

for its conclusion—even a (weak) defeasible reason for a (weak) probabilistic conclusion. Put otherwise, the argument need not be deductively valid, or even inductively strong; but it must have provided some reason for us to believe in the truth of its conclusion.

It seems plausible to us to view these as minimal constraints on the argumentative legitimacy of any use of denying the antecedent. An argument that fails to provide a reason in the sense outlined above is not a *legitimate* instance of denying the antecedent; and an argument whose cogency depends on the addition of further premises or the alteration of the conclusion is not an instance of denying the antecedent at all. It is worth noting also that one well-noted legitimate use of the form of denying the antecedent—arguments in which the inference from premises to conclusion is semantically, rather than formally, valid—passes this test without difficulty; but that, of course, requires no revision of the traditional approach to the fallacy.

2.1 A legitimate argumentative strategy?

Just over a decade ago, David Godden and Douglas Walton made what sounds like a bold claim: that even "in cases where the conditional employed in [an] argument is properly interpreted as a Philonian (or material) conditional, there are non-fallacious uses of the strategy of denying the antecedent" (Godden and Walton 2004, p. 219). They explicitly state that they are advancing this claim *against* those who maintain that "it [is] always fallacious to argue in the form of denying the antecedent" given a material interpretation of the conditional (p. 225).

If successful, this would indeed be a significant result—one compelling a rethink, or at least a nuancing, of the orthodox position on denying the antecedent as a formal fallacy. But they are not successful. To see why, we can begin by noting what Godden and Walton are *not* denying (p. 230):

> [T]here is a use of DA [denying the antecedent] which is properly interpreted as a deductively fallacious form of argument. Should the move of denying the antecedent occur as a move made by the proponent in an attempt to establish [the negation of][5] the consequent, then the standard account of DA, on which it is a fallacious move in the argument, applies.

---

[5] This is our insertion; we assume that it was omitted from the original in error.

So far, so orthodox. Godden and Walton suggest, however (p. 230), that if the relevant claims occur within an argumentative dialogue, then

> should the move of denying the antecedent occur as a move made by the respondent to an argument, a second usage of denying the antecedent might apply on which the move is not fallacious.

Specifically, if a proponent ("Pro") of a particular conclusion puts forward an argument in the *modus ponens* form—coupling a conditional premise with the assertion of its antecedent—the respondent ("Resp") in the dialogue can respond with a "denial of the antecedent" (p. 230):

> This is the move which concerns us here. In a counter-argument of this sort, the conditional premise of Pro's initial argument is accepted by Resp. But Resp rejects the move, made by Pro, of affirming the antecedent. Instead, Resp denies the antecedent.

Again, Godden and Walton are careful not to claim that the respondent proceeds on this basis to affirm the negation of the consequent. Rather, "the antecedent is denied in an attempt to establish that the consequent [which is the proponent's argument's conclusion] is not acceptable on the grounds expressed by the conditional premise" (pp. 231, 234-5):

> Pro's conclusion . . . is rebutted by Resp's counter-conclusion that [Pro's conclusion] has not been shown to be justified, and that if it is to be accepted as justified it must be done for reasons other than those so far provided by Pro. This counter-conclusion is supported by Resp with an argument which has the form of denying the antecedent.

There is, we think, a significant problem with this account. When Godden and Walton refer to "an argument which has the form of denying the antecedent," they oscillate between using this phrase to refer to an argument with that form, and to refer to a simple statement that denies a sufficient condition. That is, their use of the phrase "argument with the form of denying the antecedent" equivocates between an *argument* with the following form:

> (1) If p, then q.
> (2) It is not the case that p.
> Therefore (from (1) and (2)),
> (3) It is not the case that q.

and a *statement like* (2) made in a dialogical argumentative context in which *someone else* has put forth a statement like (1) as part of a *modus ponens* argument for *q*. Again, on this the authors are clear: in the context of such a dialogue, "denying the antecedent"—in the sense of asserting a claim like (2) in response to the other party's *modus ponens*—cannot legitimately be used to warrant a claim like (3); it can only warrant the claim that the other party has "*failed to . . . establish* the conclusion of her conditional [*modus ponens*] argument", which will "shift the burden of proof" back to the proponent (p. 231). But then Godden and Walton present no challenge whatsoever—differently from what they claim—to what they call the "standard account" (p. 219) of arguments that have the form of denying the antecedent. Their account simply does not address arguments of this form at all. In other words, they fail to meet the second of our two minimal constraints.[6]

If this is correct, then what Godden and Walton have shown is that where, in an argumentative dialogue, one party relies on a premise *p* in support of her conclusion, then taking issue with *p* is a "legitimate argumentative move" for the respondent party who wishes to take issue with the conclusion. But could anyone ever have thought otherwise? This seems to amount to nothing more than the claim that we can legitimately seek to undermine other people's arguments by questioning their premises.

The authors themselves seem at points to acknowledge this result, noting the affinities between their proposed approach and some "widely received" views in the "theory of rebuttal, or counter-argument" (p. 237-8):

---

[6] An anonymous reviewer has suggested to us that Godden and Walton might be interpreted as making the point that instances of the following pattern will always be cogent arguments:

> (1') It has been established that if *p*, then *q*.
> (2') It has not been established that *p*.
> (3') It has not been established that *q*.

But this is not obviously the pattern of denying the antecedent; and as this pattern stands, at any rate, the conjunction of (1') and (2') would by itself give us no reason to think that *q* has not been established (in the relevant context of argumentation)—which is of course also the reason why instances of denying the antecedent are not formally valid arguments. And if, on the other hand, it is only the relevant argumentative context that entitles us to say that a party's failure to establish *p* makes it more likely than not (or even certain) that *q* has not been established, then it is this further premise that will be doing the argumentative work, and the argument would again no longer count as an instance of denying the antecedent.

> In this general framework [of a theory of rebuttal], the legitimate use of denying the antecedent is just a special case of showing that an argument is a bad one by showing that it has a bad premise. Specifically, in a target argument of a *modus ponens* form, while the conditional premise may be acceptable, the premise which asserts the antecedent of the conditional may not be. Denying the antecedent rebuts the initial argument by denying this premise . . . As such, denying the antecedent works as a form of rebuttal by undercutting the initial argument by showing that it has an unacceptable premise . . . Denying the antecedent and showing an argument to be flawed due to an unacceptable premise are analogous in regard to both their argumentative effects and the type of conclusion they licence.

As should be clear, however, "denying the antecedent" in the sense that Godden and Walton intend it here is in no helpful sense "analogous" to showing an argument to be flawed due to an unacceptable premise: the former is simply one instance of the latter. Nor is it a "special case"; it is simply a case, which might *appear* special only as a result of an equivocation between the formally fallacious pattern of argument called "denying the antecedent" and a statement *denying an antecedent* upon whose affirmation one's interlocutor is relying.[7] The authors, aware of this issue, seek to address it in a final footnote; but their remarks do nothing but restate the equivocation in starker terms still (p. 242, n. 24):

> [W]e claim that the project of this paper is . . . worthwhile. While it should be conceded that "*whenever* a respondent denies *any* of [the] proponent's premises, he shifts the burden of proof back to [the] proponent" it does not follow that "there is nothing special about the case here singled out". There is something very special about the case we have singled out: namely, it is typically classified as a formal fallacy of argument!

For the reasons given above, we think this last claim false. There is no challenge to the standard account of denying the antecedent to be found along these lines.

---

[7] Note, also, that the strategy of rejecting the antecedent of the proponent's conditional premise is in no way weakened by the fact that the respondent chooses to deny the proponent's conditional premise *as well*; nor is it strengthened if the respondent accepts (or simply chooses not to take issue with) the conditional premise. The acceptability of *each* of the proponent's premises is a separate question, and there is nothing particularly noteworthy about the case in which the respondent does reject one, but not the other.

2.2 An informational shortcut?

A different attempt to rehabilitate the formal fallacy of denying the antecedent (as well as the formal fallacy of affirming the consequent) is made by Floridi (2009). Floridi too opens with strong claims. He indicts logic for its "ungreen policy" regarding the formal fallacies. Instead of discarding them "as absolutely worthless rubbish", we should be "recycling" the fallacies so as not to "waste their potential contribution to our information processes". For the formal fallacies can be "rather useful . . . ways to gain and manage one's information": they are "not mere mistakes of no value, but informational shortcuts that can be epistemically fruitful if carefully managed" (p. 318). That is what Floridi promises to show; our view is that he too is unsuccessful.

Floridi's proposed "greener" (p. 318) interpretation of the formal fallacies relies on Bayes's theorem. He works with the following example (p. 320):

> Suppose Jill receives many emails, but only a few of them (say 2%) are infected by some software virus. She uses a rather reliable antivirus software, which is successful 95% of the time. The latter does not erase her potentially infected emails, but moves them to a special quarantine folder in the e-mail client software, which Jill can check. Jill wonders how often she should check it for good emails.

The example is clear. Then Floridi writes (p. 320):

> The question [Jill] is implicitly asking is: "what is the probability that A (=the email was infected), given the fact that B (=the email was blocked by the antivirus and placed in the quarantine folder) when, on average, 2% of all the emails I receive are actually infected and my antivirus is successful 95% of the time, that is, it provides only 5% false positives?"

This is a slightly awkward formulation of Jill's question. What she wants to know is what is the probability that *an* email—not "the" email (what email?)—that has been blocked is actually infected. Bayes's theorem enables Jill to calculate the answer: the chance that a blocked email is infected is 28%. "Clearly," as Floridi says (p. 320), "Jill should check her folder regularly."

But what does this have to do with the formal fallacies? Floridi brings Maggie into the picture. He says that Maggie, differently from Jill, is not "a smart girl". She "uses the same antivirus as Jill, and receives roughly the same number of emails", but

she is "astonished" when Jill tells her that "she should check her quarantine folder regularly" (p. 321).

Why is Maggie astonished? Here Floridi's story becomes less clear. He says (p. 321) that Maggie

> thought that, if the email was infected, then the antivirus blocked it, and since the quarantine folder contains only emails blocked by the antivirus, then all the emails in it must be infected.

Why would Maggie think that "if the email was infected, then the antivirus blocked it"? Floridi does not say, but there seems to be only one plausible answer. Maggie must have been *unaware* that the antivirus was only 95% reliable. She must have believed, falsely, that it was 100% reliable.

Floridi continues:

> More formally, [Maggie] reasoned that A → B, B ⊢ A. Jill explains to Maggie that the previous inference is a typical fallacy (AC) [the fallacy of affirming the consequent], but that she should not feel silly at all.

Maybe Jill is not as smart as all that. Maggie's reasoning, at least as Floridi describes it, does not look like the fallacy of affirming the consequent at all. Rather, it looks like this:

(1) All infected emails are blocked.
(2) All emails in the folder are blocked.
       Therefore (from (1) and (2)),
(3) All emails in the folder are infected.

This argument is not valid; but it is not an instance of the fallacy of affirming the consequent. The problem is that Floridi's description of Maggie's reasoning is equivocal: is Maggie reasoning about "all the emails" in the folder, as Floridi writes at one point, or about a particular email ("the email"), as he writes at another point in the same passage? But we can fill in the example to make it fit Floridi's own claims about it. Assume that Maggie was reasoning about some particular email: for example, an email that Jill told her she had sent her the night before, and which (since it was not in Maggie's regular inbox folder) Maggie knew had been blocked. So Maggie—according to Floridi—would be reasoning like this:

(1) All infected emails are blocked.
(2) The email that Jill sent me yesterday is blocked.
       Therefore (from (1) and (2)),
(3) The email that Jill sent me yesterday is infected.

This is an instance of the universal form of the fallacy of affirming the consequent:

(1) For every x: if x is S, then x is P.
(2) a is P.
      Therefore (from (1) and (2)),
(3) a is S.

This seems to be what Floridi has in mind. Yet if indeed this is how Maggie reasoned, then why should she "not feel silly"? Here is Floridi's explanation. He asks us first to "suppose" (counterfactually) that we actually have "perfect, infallible anti-virus software" that generates "no false positives" (and no false negatives) (p. 321). To be clear, this means that it would then be true both that (for the domain of emails received by Maggie)

For every x: if x is blocked, then x is infected,

and that

For every x: if x is infected, then x is blocked.

That means that the probability that a blocked email is *not* infected would be 0; and the probability that an infected email is blocked would be 1. The double implication would therefore be true:

For every x: x is infected if and only if x is blocked.

Floridi says that "perhaps" Maggie "had in mind" the "double implication" (p. 321). Perhaps she did: but then she must have falsely believed that the antivirus was 100% reliable; for otherwise the double implication would not hold. And if she did have the double implication in mind, then her argument, though unsound (since the double implication was actually false), would be deductively valid. The other possibility, Floridi says, is that she did have in mind the formally invalid argument—the universal form of affirming of the consequent—sketched above. But "either way", he claims,

> Maggie was taking a shortcut (she disregarded the probabilities) to focus on the sort of information that she could extract from the fact that those emails were in the quarantine folder. And on the wise advice of being safe rather than sorry, she treated all its content as dangerous. The result is that Maggie is ontologically thrifty (she trusts many less items than Jill) but logically greener (she relies

> on a reasoning that, although formally fallacious, can still be recycled to provide a quick and dirty way of extracting useful information from her environment).

We find this puzzling; we cannot see how this description squares with Floridi's own example. What does it mean to say that Maggie "disregarded" the probabilities? One thing it might mean, as we suggested, is that she was actually *unaware* of the possibility that a blocked email would not be infected. But Floridi's description suggests that this is not the actual issue: it suggests that Maggie somehow *chose* to "focus" on the "sort of information that she could extract from the fact that those emails were in the quarantine folder"—that is, that she was aware of the possibility that not all blocked emails would be infected. But then why should she "treat all its content as dangerous"? That is certainly not a "wise" decision, since she will risk missing quite a few important emails if she does not check the quarantine folder regularly.

One way to make sense of Floridi's example would be to say that Maggie *is* aware that the antivirus is only 95% reliable—and thus aware that possibly, some blocked emails are actually not infected—but that she is also under the erroneous impression that the probability that a blocked email is infected is *very high*. Perhaps that is why Maggie is "astonished" (p. 321) when Jill informs her that the actual probability is around 28%. So maybe this is what Floridi means: Maggie got her probabilities wrong. She falsely supposed that there is a very low probability that she would find any non-infected emails in the quarantine folder; and for that reason she had decided—until Jill came along and set her straight—not to check the folder all that often. So the probabilities she "disregarded" were not the actual probabilities, but *what she took to be the very low probability* that she (falsely) assigned to an email in the folder not being infected.

Still, if this is what Floridi means, what then does it mean to say that Maggie is "taking a shortcut . . . to focus on the sort of information she could extract from the fact that those emails were in the folder" (p. 321)? Or that Maggie was being "logically greener" than Jill, because she "relies on a reasoning that, although fallacious, can still be recycled to provide a quick and dirty way of extracting useful information from her environment" (p. 322)?

If Maggie decides not to check the folder very often because she thinks the probability that it contains non-infected emails is very low, then one thing we can say is that she is *taking a risk* (which she believes is very low) that she might miss some important email. Floridi does at one point say something

like that: he speaks of reasoners like Maggie as making "bets" which may be "risky" (p. 322). But in "betting" that Jill's email, which was blocked, is actually infected, Maggie is not "relying" on any fallacious reasoning whatsoever. She could of course justify her decision in conversation by *saying* simply that in her view, "if an email is in the folder, it is infected"; but what she would thereby mean, and conversationally implicate, would be that if an email is in the folder then it is *almost certainly* infected (or something along these lines). So her view that one should not be checking the quarantine folder all that often would be grounded on the claim that

> For every *x*: if *x* is blocked, then *x* is almost certainly infected.

Her argument would therefore really be:

(1) For every *x*: if *x* is blocked, then *x* is almost certainly infected.
(2) The email that Jill sent me yesterday is blocked.
      Therefore (from (1) and (2)),
(3) The email that Jill sent me yesterday is almost certainly infected.

There is nothing fallacious about this; the argument is deductively valid. If Floridi's point is that Maggie's decision not to check the quarantine folder very often is perfectly rational given her belief that it is very unlikely that the folder contains emails that are not infected, then he is of course right (though the point is not controversial). But there is nothing here that can plausibly be described as "relying" on "fallacious reasoning" *as a way* of "extracting useful information from her environment" (p. 322). The rationality of Maggie's decision is not a matter of what she has "a decent right to conclude" (p. 318); it is a matter of whether she has a decent chance of not missing important emails.

Sensing that his claims may be perceived as "unclear", Floridi goes on to offer another illustration (p. 322):

> Maggie's teacher tells her that, if she doesn't study enough, then she will fail her exam. Unfortunately, Maggie does fail her exam and the teacher reproaches her for not having studied enough. Maggie has learnt her Bayesian lesson, so she knows the teacher's reasoning is fallacious. But she also knows that it is fairly accurate, as a shortcut that gets things right most of the time: on average, students who fail their exams have not studied enough. The teacher should have simply sprinkled her in-

ference and then judgment with some "probably" and "most likely" clauses.

Yet this example, too, is equivocal; and again no plausible reading of a situation like this would support Floridi's claims about the formal fallacies. Once more we are not told whether Maggie's teacher actually realises that it is perfectly possible for a student to have studied enough and still fail the exam; but we may assume that, being a teacher, she does realise that. Let us also assume that the teacher *knows* that the probability that a student who fails did not study enough is relatively high, and in any case higher than 0.5. Now Floridi tells us that the teacher—presumably without first inquiring whether some other factor might explain Maggie's failure—"reproaches" Maggie "for not having studied enough". How then did the teacher reason?

Given the relatively high probability (of which the teacher was aware) that, having failed the exam, Maggie did not study enough, the teacher may have decided to *go with that hypothesis* for practical purposes, and thus to berate Maggie for not having studied enough. This seems to be what Floridi has in mind. But this does not establish that the teacher jumped to—took a "shortcut" to—the conclusion that Maggie did not study enough. The teacher jumped, if anything, to *action*, to an action—berating Maggie—that might turn out not to be justified after all (if it turns out that Maggie did study enough). But this does not mean, and there is no reason to suppose, that the teacher mistook necessity for probability—Floridi says at one point (p. 320) that one possible "explanation" of the formal fallacies is that people "mistake necessity for probability"—and formed the view that necessarily, if Maggie failed the exam, then she did not study enough.

Floridi's claim that the teacher is "using" the formal fallacy of affirming the consequent as a "quick and dirty informational shortcut" is therefore unfounded, as is his more general claim (p. 322) that "[w]hen we use [the formal fallacies of affirming the consequent and denying the antecedent]", we

> bet that $A{\rightarrow}B$, B, $\vdash A$ or that $A{\rightarrow}B$, $\neg A$, $\vdash \neg B$. The bet might be risky (we might be wrong), but it often pays back handsomely in terms of lower amount of information resources needed to reach a conclusion (see the case of the teacher assessing whether Maggie studied enough).

This again is a distorted description of what is going on. When berating Maggie, the teacher is indeed—as was Maggie in the email example—*taking the risk* that her action is not justified.

But this is a risk of which she is perfectly aware. If she is aware of the risk, she is aware of the possibility that Maggie did study enough. And if she is aware of this possibility, then she cannot be reasoning that

(1) If Maggie does not study enough, she will fail the exam.
(2) Maggie failed the exam.
     Therefore (from (1) and (2)),
(3) Maggie did not study enough.

Rather, she reasons that

(1) If Maggie failed the exam, then it is more likely than not that she did not study enough.
(2) Maggie failed the exam.
     Therefore (from (1) and (2)),
(3) It is more likely than not that she did not study enough.

The teacher then decides, on the basis of this likelihood, to berate Maggie, thereby running the risk that she is being unfair. Differently from what Floridi says, then, her reasoning is not fallacious: the only plausible reading of the example is that teacher *did* "sprinkle" her inference with some "probably" and "most likely" clauses. And though it is again true that the teacher is in a sense "disregarding" the possibility that Maggie did study enough, this has nothing to do with her (theoretical) reasoning. Instead, it is in *deciding to act* as she did that she chooses to "disregard" that possibility, "betting", as Floridi says, on the scenario that she thinks more likely to be true. There is no connection between the teacher's "disregarding" that possibility and the formal fallacies.

So nothing supports Floridi's claim that *the formal fallacies* of affirming the consequent and denying the antecedent are "quick and dirty" informational shortcuts (p. 322); or that they can be "interpreted in terms of information-gathering and gain" (p. 324); or that the formal fallacies can "contribut[e] to our information processes" and can be "serviceable" when it comes to "extracting information from whatever is at hand and without much fuss." (p. 318). His examples, properly understood, turn out to be cases in which (a) a person P believes that the probability of a given fact or outcome F is relatively low; (b) F is such that, if F is true, then P has reason not to perform a given action φ; and (c) P decides nonetheless to "bet" that F is actually not the case, and to perform action φ, taking the risk that φ is

unjustified. And there is no sense in which such a person can be said to be taking a "shortcut" to the conclusion that the negation of *F* is true.

Despite his opening promise, then, Floridi's approach too fails to satisfy the second of our minimal constraints. If we are right, he has failed to give us anything that can be called an "interpretation" of the formal fallacies at all, let alone one that is "much more reasonable" (p. 318) than the one that logic traditionally affords.

2.3 An inductively strong argument?

A third author to propose a "significant revision of the view that denying the antecedent is simply a formal fallacy that should be discarded" is Mark Stone (2012, p. 329). He seeks to combine the insights of Godden and Walton on the one hand, and Floridi on the other, and use them to make an even stronger claim. Stone says that he agrees with Godden and Walton that "an argument in which the arguer denies the antecedent can be legitimate as a strategy for responding to a previous argument" (p. 342). But he thinks Godden and Walton "construe the legitimacy of this argument strategy too narrowly" (p. 341). He wants to "go further" (pp. 329, 343):

> Whereas [Godden and Walton] maintain that the force of this type of argument is only that we should *not accept the* [*proponent's*] *conclusion C* for the reasons given in the [proponent's] initial conditional argument, I think its force is that we should *probably accept not C*, that is, we should probably reject C. In other words, I think denying the antecedent has inductive strength. This is the point at which I believe one can make a stronger case for the effectiveness of denying the antecedent in argumentation.

What are his reasons? One of his working examples is Burke's capital punishment case. Stone puts it in the context of an argumentative dialogue, and imagines that someone makes "the following argument in favour of capital punishment" (pp. 343-4), which we will refer to as "Argument 1":

> *Argument 1 (a proponent's argument in favour of capital punishment)*
> (1) If capital punishment deters murder, then it is justified.
> (2) Capital punishment deters murder.

Therefore (from (1) and (2)),
(3) Capital punishment is justified.[8]

This is a valid argument; it is an instance of *modus ponens*. The "corresponding argument that denies the antecedent of this conditional", and which can be "a legitimate argumentative strategy" (p. 343) with which a respondent might seek to undermine *Argument 1*, would be:

> *Argument 2 (a respondent's challenge to Argument 1)*
> (1) If capital punishment deterred murder, then it would be justified.
> (2) Capital punishment does not deter murder.
> Therefore (from (1) and (2)),
> (3) Capital punishment is not justified.

Oddly, Stone appears to take *this*—a respondent's adoption of *Argument 2*—as the kind of "legitimate use" of "denying the antecedent" that Godden and Walton have in mind (cf. Stone 2012, p. 344). Stone misses Godden and Walton's equivocation, which we discussed in Section 2.1; and more importantly, he also seems at points to overlook Godden and Walton's explicit claim that they do not think that a respondent who challenges something like *Argument 1* by denying the antecedent of its conditional premise, is giving an argument *for* the claim that capital punishment is not justified. (In other words, *Argument 2*—which actually *is* an instance of the formal fallacy of denying the antecedent, and whose conclusion is the actual claim that capital punishment is not justified—is not exactly the kind of challenge that Godden and Walton are concerned with.[9])

In any event, Stone begins to motivate his own views by taking issue with those authors who—like Burke, he says—might want to reinterpret arguments like *Argument 2* "in a way that eliminates the apparent fallacy" (p. 328), for example by replacing premise (1) with its inverse. Burke, says Stone, would suggest that we amend *Argument 2* as follows:

---

[8] The way in which we graphically display these arguments differs slightly from Stone's (for example, we number the arguments, and the premises and conclusion of each, for ease of reference); but the content of the arguments remains the same.

[9] "Legitimate employments of denying the antecedent", Godden and Walton say, "cannot be modelled as arguments of the form 'A ⊃ C, ~A ⊨ ~C'"; the respondent "does not seek to establish any claim . . . whatsoever": rather, "the move is made in an attempt to demonstrate that [the proponent's] C[onclusion] has not been established" (cf. Godden and Walton 2004, p. 231).

> *Argument 3 (Burke's revision of Argument 2)*
> (1) If capital punishment does not deter murder, then it is not justified.
> (2) Capital punishment does not deter murder.
>   Therefore (from (1) and (2)),
> (3) Capital punishment is not justified.

Stone makes two claims in connection with this. One is that this revision does not make the argument better since "the same considerations which may be raised to challenge denying the antecedent [that is, to challenge *Argument 2*]" will hold equally against the revised, valid argument (p. 345). Suppose, for example, that someone wants to argue "that capital punishment is justified for retributive reasons" (p. 345):

> With the valid argument [*Argument 3*] this challenge, if good, means that the argument is weak because the assumed premise [premise (1) in *Argument 3*] is false. With denying the antecedent [*Argument 2*], it means that the argument is weak because it has clearly been shown that the premises may be true but the conclusion false.

Stone's second claim is his main thesis that an argument like *Argument 2* is "inductively strong": it "establishes that capital punishment is probably not justified" (p. 345), thus providing us with reasons to *reject* (and not merely reasons not to accept) the conclusion of *Argument 1*.

How is this thesis related to Stone's attempt to broaden Godden and Walton's account of denying the antecedent as a legitimate argumentative strategy? The thought appears to be this. To say that an argument is "inductively strong" is to say that if its premises are true, then the conclusion is probably true. But if the conclusion of *Argument 2* is probably true, then its negation—which is the conclusion of *Argument 1*—is probably false. So *Argument 2* gives us reason, if its premises are true, to reject the conclusion of *Argument 1* (p. 329):

> The premises of an argument that denies the antecedent in [the] context [of its use "to undermine an opponent's position for which plausible reasons have been or may be offered"] can have the logical force of an inductive argument, meaning that if the premises are true the conclusion probably follows. In undermining the opponent's position this form of argument provides reasons for believing that the position is false.

Neither of Stone's claims succeeds, however, and his own example of the retributivist argument for capital punishment will

help us to see why. He is never clear about what that argument would actually look like, but it seems safe to suppose that it would run along the following lines:

> *Argument 4 (the retributivist argument in favour of capital punishment)*
> (1) If capital punishment is an appropriate retributive re sponse to murder, then it is justified.
> (2) Capital punishment is an appropriate retributive re sponse to murder.
>   Therefore (from (1) and (2)),
> (3) Capital punishment is justified.

Stone says that this argument can either be "good" or "fail" (p. 345); and that "if good", the argument holds equally against *Arguments 2* and *3*. Given that *Argument 4* is valid, what does it mean to say that it can either "fail" (or be "unsuccessful": pp. 346, 348) or be "good"? One thing it can mean is that *Argument 4* can be either sound or unsound. So let us suppose that the proponent of *Argument 4* succeeds in showing that the argument is sound—that its premises are both true. Stone says that this would weaken *Argument 2* by showing that its "premises may be true but the conclusion false" (p. 345). But this is misleading. It is true that if *Argument 4* is sound, then the conclusion of *Argument 2* is false. It is also true that the soundness of *Argument 4* is consistent with the truth of both premises of *Argument 2*. But in order to successfully show that in *Argument 2* the truth of the premises is consistent with the falsehood of its conclusion, one does not have to endorse an argument like *Argument 4*. One only has to point out that it is logically possible that there are other justifying factors. In other words, the challenge against denying the antecedent is that the arguer failed to rule out other possibilities, not that those possibilities are actually realised; and this challenge, which suffices to refute *Argument 2*, has no bite against *Argument 3*.

So Stone is wrong that—this was his first claim—the two versions of the deterrence-based argument against capital punishment (that is, the fallacious version in *Argument 2*, and the formally valid version in *Argument 3*) can both be weakened by the "same considerations" (p. 345). It is not the case, differently from what he says, that "the reconstructed valid argument may be criticised for the same reasons as the argument denying the antecedent" (p. 337). He is right, of course, that *Argument 4*, if sound, is also effective against *Argument 3*, since the soundness of *Argument 4* entails that the conditional premise of *Argument 3* is false. But what that shows is that there are considerations

that hold equally against both versions, not that all considerations that hold against one will hold against the other.

What of Stone's second claim—his thesis that an argument like *Argument 2* has "inductive strength" (p. 346), establishing that capital punishment is "probably not justified" (p. 344)? Again, his argument for this claim is never presented in a straightforward way; but he thinks the inductive strength of *Argument 2* "becomes apparent" if we now suppose that the retributivist challenge—*Argument 4*—"fails" (p. 345):

> I would argue that if the challenge fails, the argumentative force of denying the antecedent becomes apparent: it is inductively strong. [*Argument 2*] . . . establishes that capital punishment is probably not justified.

This thought is repeated at several other points; for example (p. 346):

> [S]omeone who believes that deterring murder is a sufficient condition for justifying capital punishment, and who has evidence that it does not deter murder, can make an inductive argument in the form of denying the antecedent [that is, an argument like *Argument 2*] that capital punishment is probably not justified. This argument has inductive strength. How strong it is depends on how successful other arguments to justify capital punishment are. If the retributive argument suggested earlier [that is, *Argument 4*] is unsuccessful, then the inductive argument denying the antecedent is accordingly stronger.

Or more generally (p. 348):

> The conclusion of the argument [*Argument 2*] that capital punishment is not justified is strengthened by the likelihood that other justifications for capital punishment are unsuccessful.

Here is what we take Stone's point to be. Suppose you do believe that deterring murder would suffice to justify capital punishment; but you also believe, on good evidential grounds, that capital punishment does not in fact deter murder. That this one sufficient condition is not met does not of course establish that capital punishment is not justified. There may be other sufficient grounds of justification. Indeed, being an adequate retributive response to murder might be one: that would be the point pressed by anyone who endorsed *Argument 4*. So suppose you also agree that capital punishment would be justified if it were an appropriate retributive response to murder—you agree with

the conditional premise in *Argument* 4—but again you have reason to believe that capital punishment does *not* in fact constitute an appropriate retributive response to murder. All this establishes is that there is *another* sufficient condition for the justification of capital punishment that is not satisfied; the possibility remains that there are other sufficient conditions, conditions you have not yet considered, that would justify capital punishment as well. Still, there are only so many considerations you would be willing to assert as sufficient conditions for the justification of capital punishment. Suppose, then, that for any sufficient condition you can think of—including every one that an opponent might have put to you in an attempt to establish that capital punishment is justified—you have reason to think that that condition is either not satisfied, or unlikely to be satisfied. In that case, since you believe (a) that deterrence is a sufficient condition for justifying capital punishment, (b) that capital punishment does not deter murder—these are the two premises in *Argument 2*—and (c) that it is unlikely that capital punishment can be justified on other grounds, then of course you will be justified in believing (d) that the conclusion that capital punishment is not justified—which is the conclusion in *Argument* 2—is probably false.

(What makes other possibilities unlikely is a matter of either (a) it being unlikely that there are *other* true conditionals setting out sufficient conditions for justifying capital punishment; or (b) if there are other such conditionals, it being unlikely that their antecedents are satisfied.)

Note, however, that the unlikelihood that there are other sufficient conditions that happen to be satisfied has nothing to do with the "strength" of any denying-the-antecedent move; it has nothing to do with the fact that both premises in *Argument 2* may actually be true. Rather, the claim that no other sufficient condition is likely to be satisfied turns on considerations extraneous to that denying-the-antecedent move.[10] So in order for someone's reliance on premises (1) and (2) of *Argument 2* to warrant—as Stone claims it would—the conclusion that capital punishment is probably not justified, the argument would have

---

[10] In a recent paper, Godden and Zenker (2015), who also engage at length with Stone's account, argue that "in whatever way indicative conditionals are interpreted, the [putative] cogency of DA [denying the antecedent] and AC [affirming the consequent] depends on contingent factors that are unasserted by, and remain independent of, the denial of antecedents or the affirmation of consequents" (p. 120); "DA and AC arguments *as stated* are not probative—they fail to provide reasons for their conclusions—unless such assumptions are explicated and met" (p. 106). We agree, and have sought to offer a different argument for this conclusion.

to include a further premise stating that capital punishment is probably not justified unless it deters murder; and the conclusion of the argument would have to be rewritten to state that capital punishment is probably not justified (or something to that effect):

> *Argument 5*
> (1) If capital punishment deterred murder, then it would be justified.
> (2) Capital punishment does not deter murder.
> (3) If capital punishment does not deter murder, it is probably not justified.
> Therefore (from (1) and (2) and (3)),
> (4) Capital punishment is probably not justified.

Premise (3) is a *further* premise in the sense that its negation is consistent with the assertion of the two premises in the original version of the argument (the denial-of-the-antecedent version in *Argument 2*).

Our point, then, is that unless premise (3) is included, the conclusion that capital punishment is probably not justified is not warranted. Whether the conclusion is warranted, therefore, is something that has nothing to do with any "legitimate role" played by denying the antecedent "in establishing that a position is improbable" (Stone 2012, p. 348). In fact, it seems unhelpful and even misleading to discuss this under the "denying the antecedent" heading. After all, as soon as premise (3) is included, premise (1) becomes superfluous: the conclusion in *Argument 5* follows deductively from premises (2) and (3) alone. So what we are really left with is

> *Argument 6*
> (1) Capital punishment does not deter murder.
> (2) If capital punishment does not deter murder, it is probably not justified.
> Therefore (from (1) and (2)),
> (3) Capital punishment is probably not justified.

This is a valid argument, and one in which premise (1) is no longer denying any "antecedent". So Stone's analysis gets things doubly wrong: it leaves the crucial premise out of the argument, and leaves in a needless one. If this is Stone's point, he too fails to satisfy the second of our minimal constraints.

Could Stone object that we are missing his main claim? Is his thesis not that an argument like *Argument 2* has "inductive strength"—that it establishes that "capital punishment is proba-

bly not justified"—*regardless* of how successful other arguments for justifying capital punishment may be? He says that if these other arguments fail, *Argument 2* becomes inductively *stronger*, but does he not think that *Argument 2* has "inductive strength" even if such arguments succeed? Here is the relevant passage again (p. 348):

> [S]omeone who believes that deterring murder is a sufficient condition for justifying capital punishment, and who has evidence that it does not deter murder, can make an inductive argument in the form of denying the antecedent [like *Argument 2*] that capital punishment is probably not justified. This argument has inductive strength. How strong it is depends on how successful other arguments to justify capital punishment are. If the retributive argument suggested earlier [that is, *Argument 4*] is unsuccessful, then the inductive argument denying the antecedent is accordingly *stronger*.

If someone believes that deterring murder is a sufficient condition for justifying capital punishment, then if they have evidence that it *does* deter murder, then obviously they can claim that capital punishment is probably—indeed certainly, since the probability would be 100%—justified. But it does not follow, of course, that evidence that capital punishment *does not* deter murder would warrant the claim that capital punishment is probably not justified. We could grant that if the premises of *Argument 2* are both true, then its conclusion—that capital punishment is not justified—is thereby rendered *more probable*, more likely to be true. But that is not enough to support Stone's claim that *Argument 2*—or denying the antecedent more generally—establishes the opposite conclusion as "probably" false. What Stone's claim requires is not merely that the premises of an argument like *Argument 2* establish the conclusion as more likely to be true than *if the premises were false*. What Stone's claim requires is that the premises of such an argument establish the conclusion as *more likely to be true than false*.

In fact, Stone does not carefully distinguish between (a) the claim that the premises of an argument that denies the antecedent, if true, lend *some* probabilistic support to its conclusion, and (b) the claim that the premises of an argument that denies the antecedent, if true, make the conclusion more likely to be true than false. He uses the same phrases—the adverb "probably", or expressions like "inductive strength"—to make claims of both

kinds.[11] This ambiguity may give his proposal some superficial appeal; if by saying that denying the antecedent has "inductive strength" one means only the former claim, then one might perhaps grant that *Argument 2* has some "inductive strength". But that would not tell us that the argument establishes its conclusion as "probably" true. In order to know *that*, we would need to know something else: something that is not settled by the fact that the premises of *Argument 2* are both true. We would need to know *how likely it is* that the conclusion is true if capital punishment does not deter murder; and this is something that turns, as we saw, on the higher or lower probability that there are *other* sufficient conditions for justifying capital punishment that are satisfied. So we would need to add this information to our premises set; and then—again—the conditional premise could simply be discarded, and we would be left with something like *Argument 6* above.

The flaws in Stone's analyses become particularly vivid when he proposes to consider a couple of "significant examples" which he thinks "clearly fit the argumentative context in which we have established that denying the antecedent has a legitimate role" (p. 350). One example is taken from a book by James Rachels (1990). Stone discusses a passage in which Rachels "argues that Darwinism undermines traditional morality by undermining the idea of human dignity" (Stone 2012, p. 351):

> Rachel's argument that Darwinism undermines the idea of human dignity attacks two main doctrines that are presented as support for the idea that human beings have moral dignity: (1) the view that humans are made in the image of God and (2) the view that human beings are uniquely rational. The way he applies this reasoning to his main argument is worth quoting in full:
>
>> We are now in a position to explain how Darwinism might undermine traditional morality. The claim that Darwinism undermines traditional morality is not the claim that it entails that the doctrine of human dignity is false. It is, instead, the claim that Darwinism provides reason for doubting the truth of the considerations that support the doc-

---

[11] For example, when commenting on a variant of Floridi's case "of a student [Maggie] reasoning about her chances of success on a test"—a case we discussed in Section 2.2—Stone writes: "Since *on average* students who study some pass the test, [Maggie's] reasoning provides support for the conclusion *. . . it is still reasonable for Maggie to draw the conclusion* she does. *In other words, the argument has inductive strength*" (Stone 2014, p. 347, emphasis added).

trine. From a Darwinian perspective, both the image of God thesis and the rationality thesis are suspect. Moreover, there are good Darwinian reasons for thinking it unlikely that any other support for human dignity can be found. Thus, Darwinism furnishes the "new information" that undermines human dignity by taking away its support. (Rachels 1990: 97-98)

Stone says that (p. 351):

[t]he arguments that Rachels attacks depend on two conditionals: "If human beings are created in the image of God, then human beings have moral dignity", and "If human beings are uniquely rational, then human beings have moral dignity". The arguments that Rachels uses to undermine the notion of human dignity involve denying the antecedent of these conditionals.

So he attributes two arguments to Rachels. One is an argument against the image-of-God argument for moral dignity (Stone 2012, p. 351):

(1) If human beings are created in the image of God, then human beings have moral dignity.
(2) But human beings are not created in the image of God (as Darwinism has shown).
     Therefore,
(3) Probably human beings do not have moral dignity.

The other is an argument against the unique rationality argument:

(1) If human beings are uniquely rational, then human beings have moral dignity.
(2) But human beings are not uniquely rational (as Darwinism has shown).
     Therefore,
(3) Probably human beings do not have moral dignity.

But Stone gets Rachels's argument doubly wrong as well. First, he saddles Rachels with a *commitment* to the two conditionals: in Stone's reconstruction, the conditionals feature as an integral part of *Rachels's* arguments. Yet these are Rachels's *opponents'* rather than his own conditionals; and clearly Rachels's strategy

of denying the antecedent of such conditionals does not in the least commit him to accepting the conditionals themselves.[12]

Second, Stone's reconstruction omits a crucial step of Rachels's argument, indeed one that is explicitly adopted in the very excerpt by Rachels that Stone quotes. Rachels does not merely argue that "from a Darwinian perspective, both the image of God thesis and the rationality thesis are suspect". He immediately adds that "moreover, there are good Darwinian reasons for thinking it *unlikely that any other support for human dignity can be found.*" This step is needed for the argument to run; and—again—as soon as we add this premise to the arguments the conditional premises can be dispensed with (unsurprisingly, given that Rachels did not endorse them), and the resulting arguments are deductively valid. Indeed we can reconstruct everything as one single valid argument:

(1) Human beings are not created in the image of God (as Darwinism has shown).

(2) Human beings are not uniquely rational (as Darwinism has shown).

(3) If human beings are neither created in the image of God nor uniquely rational, then probably human beings do not have moral dignity (for Darwinism has shown that it is unlikely that any other support for moral dignity can be found).
Therefore (from (1), (2) and (3)),

(4) Probably human beings do not have moral dignity.

The other example Stone discusses is an argument by John Locke, in *An Essay Concerning Human Understanding*, against "the claim that some ideas or principles are innate" (Stone 2014, p. 353):

> The main argument that Locke attacks is that universal consent about principles proves them innate. "There is nothing more commonly taken for granted", Locke explains,
>
>> than that there are certain Principles both Speculative and Practical (for they speak of both) universally agreed upon by all Mankind: which therefore, they argue, must needs be the constant Impressions, which the Souls of Men receive in their first Beings, and which they bring into the World with

---

[12] See also note 7 above.

them, as necessarily and really as they do any of
their inherent faculties. (Locke 1975: 49)

Part of Locke's argument against these innate principles
aims to undermine the position that there are innate prin-
ciples by taking away its support. Locke declares that
"this Argument of Universal Consent, which is made use
of, to prove innate Principles, seems to me a Demonstra-
tion that there are none such: Because there are none to
which all Mankind give an Universal Assent" (Locke
1975/1689: 49).

Stone thinks that Locke's "argument is correctly represented as
an example of denying the antecedent" (Stone 2014, p. 353):

(1) If there are principles that have universal agreement,
then these principles must be innate.
(2) But there are no principles that have universal agree
ment.
Therefore,
(3) These principles are not innate.

Stone downplays Locke's explicit claim that he takes this to be a
"demonstration" that there are no innate principles, and thinks
this should be read as simply an attempt to "undermine" his op-
ponents' arguments. This is highly unpersuasive; Locke does
appear to be denying the antecedent of the conditional in order
to (fallaciously) derive the negation of its consequent (in much
the same way as Russell and Cohen in the arguments we dis-
cussed at the end in Section 1). Is there a way of avoiding attrib-
uting the formal fallacy to Locke? In his book *Locke on Human
Understanding*, E. J. Lowe suggests one. "On the face of it",
Lowe writes, Locke's argument "is a blatant example of denying
the antecedent" (Lowe 1995, p. 24):

The innatist, Locke has suggested, makes the following
claim:

1. If any principle is universally assented to, then it is in-
nate.

The innatist then allegedly conjoins 1 with the claim that
certain principles *are* universally assented to, and validly
draws the conclusion that those principles are innate. But
Locke himself now *denies* the antecedent of 1 by assert-
ing:

2. No principle is universally assented to;

and concludes thence

3. No principle is innate.

But 1 and 2 do *not* entail 3. (To suppose that they do is precisely to commit the fallacy of "denying the antecedent".) What is needed in conjunction with 2 to entail 3 is, rather,
4. If any principle is innate, then it is universally assented to.

"However", as Lowe goes on to say (p. 25), "a justification of Locke's strategy may be forthcoming if we take him to suppose that the innatist has and can have *nothing other* than universal assent to offer in support of the existence of innate principles". Under this reading, then, Locke's argument is not fallacious; it is to be reconstructed instead as:

(1) No principle is universally assented to.
(2) There is no other possible basis (that is, no basis other than universal assent) for the existence of innate principles.
Therefore (from (1) and (2)),
(3) No principle is innate.

This is the same kind of reconstruction we offered both for the capital punishment example (*Argument 6* above) and for Rachels's Darwinist argument. (The only difference is that here the second premise affirms that it is certain, rather than just likely, that there is no other basis on which the opposite conclusion could be supported.) Whether or not Lowe's reading of the argument is warranted is a question of Lockean exegesis that we do not have to adjudicate. What matters is that of the two competing reconstructions of Locke's argument, none matches Stone's analysis. For the reconstruction above is a deductively valid argument having nothing to do with denying the antecedent; and the alternative reading is that Locke did deny the antecedent (of his opponents' conditional) seeking to derive the negation of the consequent: but then he did commit the formal fallacy.

3. Conclusion

This paper offered discussions of several recent treatments of the fallacy of denying the antecedent. In Section 1, we argued that Burke's interpretative approach does not, contrary to what

he claims, give us reason to think that the fallacy is seldom if ever committed in print. Our discussion was not targeted at the interpretative approach itself, on which we have endorsed no strong views. Rather, our argument was that the interpretative approach should not be tied to the assumption that passages that can charitably be said to commit the fallacy must instantiate what Burke calls a "B pattern". Once that questionable assumption is discarded, Burke's claim that the fallacy is rarely committed is no longer supported. There are, moreover, we also argued, positive reasons to think that Burke's claim is actually false.

In Section 2, which was mostly critical in both purpose and tone, we discussed three recent attempts to rehabilitate the fallacy of denying the antecedent: Godden and Walton's claim to have identified non-fallacious uses of "what is typically classified as formal fallacy of argument!"; Floridi's claim that the formal fallacies are "informational shortcuts that can be epistemically fruitful"; and Stone's claim that denying the antecedent has "inductive strength" and is not "simply a formal fallacy that should be discarded". We argued that none of these claims is warranted. Again, our strategy was not to offer comprehensive indictments of these authors' broader accounts of fallacies, argumentative dialogues, burdens of proof, and so on. Indeed, on many aspects of such accounts we have simply remained agnostic. What we did target was the attention-grabbing claim, made on different grounds by each of those authors, that their accounts pose an interesting or illuminating challenge to the orthodox view of the formal fallacies. In our view, they do not.

After all, in order to establish that there are legitimate instances of "what is typically classified as a formal fallacy of argument" (to use Godden and Walton's phrase again), those authors would have to have shown that there can be arguments—apart from those that happen to be materially valid by virtue of the meaning of the terms occurring in the premises—that (a) are cogent, and (b) do not rely for their cogency on further assumptions or premises (or changes to the conclusion) beyond those of the denying the antecedent pattern. Stone's examples fail to meet condition (a), while Godden and Walton's and Floridi's fail to meet condition (b). But arguments that fail to meet the former are not aptly called "legitimate", and those that fail to meet the latter are not instances of denying the antecedent.

References

Adler, Jonathan E. 1994. Fallacies and alternative interpretations. *Australasian Journal of Philosophy* 72: 271-282.

Burke, Michael B. 1994. Denying the antecedent: A common fallacy? *Informal Logic* 16: 23-30.

Cohen, Nick. 2014. The phantom menace of militant atheism. *The Guardian* (7th September, 2014), http://www.theguardian.com/commentisfree/2014/sep/07/militant-atheism-religious-apologists-intellectuals. Accessed 16 October 2015.

Finocchiaro, Maurice A. 1981. Fallacies and the evaluation of reasoning. *American Philosophical Quarterly* 18: 13-22.

Floridi, Luciano. 2009. Logical fallacies as informational shortcuts. *Synthese* 167: 317-25.

Godden, David M. and Douglas Walton. 2004. Denying the antecedent as a legitimate argumentative strategy: A dialectical model. *Informal Logic* 24: 219-43.

Godden, David M. and Frank Zenker. 2015. Denying antecedents and affirming consequents: The state of the art. *Informal Logic* 35: 88-134.

James, William. 1979/1891. The moral philosopher and the moral life. In *The works of William James: Electronic edition. Volume 6: The will to believe and other essays in popular philosophy*, eds. Frederick H. Burkhardt et al. Cambridge, MA: Harvard University Press.

Krabbe, Erik C. W. 1995. Can we ever pin one down to a formal fallacy? *Poznań Studies in the Philosophy of the Sciences and the Humanities* 45: 151-164.

Locke, John. 1975/1689. *An essay concerning human understanding*, ed. Peter H. Nidditch. Oxford: Clarendon Press.

Lowe, E. J. 1995. *Locke on human understanding*. London and New York: Routledge.

Moldovan, Andrei. 2009. Pragmatic considerations in the interpretation of denying the antecedent. *Informal Logic* 29: 309-26.

Rachels, James. 1990. *Created from animals: The moral implications of Darwinism*. Oxford and New York: Oxford University Press.

Russell, Bertrand. 1997/1952. Is there a God? In *The collected papers of Bertrand Russell. Volume 11: Last philosophical testament 1943-68*, ed. John G. Slater. London: Routledge.

Stone, Mark. 2012. Denying the antecedent: Its effective use in argumentation. *Informal Logic* 32: 327-56.