

Does the Consequence Argument Beg the Question?

*I have never been able to get very clear about what a circular argument is—or begging the question either. I am reminded of a remark that Roderick Chisholm once made in response to a charge of having committed one or the other of these offenses: “I seem to have been accused of the fallacy of affirming the antecedent.”*¹

(Peter van Inwagen)

Published in Philosophical Studies 166 (3):575–595 (2013). Please cite published version.

1. Introduction

Various contemporary philosophers have given expression to an argument for the incompatibility of causal determinism and freedom (in the sense that involves access to alternative possibilities or freedom to do otherwise).² This argument—which is parallel in important ways to the ancient argument for the incompatibility of God’s foreknowledge and human freedom³—has been dubbed the “Consequence Argument” by Peter van Inwagen (1983, p. 16):

If determinism is true, then our acts are the consequences of the laws of nature and events in the remote past. But it is not up to us what went on before we were born,

¹ The epigraph comes from van Inwagen (2005, p. 346).

² See, for example, Wiggins (1973); Ginet (1966), (1980), and (1990); van Inwagen (1975) and (1983); and Fischer (1994).

³ Nelson Pike offers a lucid regimentation of the argument in his (1965); for a discussion of the parallels between this argument and the argument for the incompatibility of causal determinism and freedom to do otherwise, see Fischer (1994). There is also a recent debate about whether the Pike-style argument for the incompatibility of God’s foreknowledge and human freedom begs the question, which can be found in Merricks (2009), Fischer and Todd (2011), and Merricks (2011).

and neither is it up to us what the laws of nature are. Therefore, the consequences of these things (including our present acts) are not up to us.

Van Inwagen goes on to say that this argument is “obvious,” in the sense that “it’s one that should occur pretty quickly to any philosopher who asked himself what argument could be found to support incompatibilism.” (1983, p. 16)

The argument has elicited various responses, ranging from acceptance as obviously right to rejection as obviously problematic in one way or another.⁴ Here we wish to focus on one specific response, according to which the Consequence Argument *begs the question*. This is a serious accusation that has not yet been adequately rebutted, and we aim to remedy that in what follows. We begin by giving a formulation of the Consequence Argument. We also offer some tentative proposals about the nature of begging the question. Although the charge of begging the question is frequently made in philosophy, it is surprisingly difficult to pin down the precise nature of this dialectical infelicity (or family of such infelicities). At the risk of seeming impudent, we make some new proposals about the nature of begging the question with an eye to understanding what is going on in central cases in which the charge is legitimately made. We then defend the Consequence Argument against the charge that it begs the question, so construed. Of course, this defense, even if it is successful, does not establish that the argument is obviously right—or right at all. But we hope that our discussion will at least sharpen the issues pertaining to the charge that the argument begs the

⁴ Sympathetic presentations of the Consequence Argument can be found in Fischer (1994), Ginet (1990), Van Inwagen (1983), and Wiggins (1973), among others. For some essentially (although not entirely) friendly critical discussions of the Consequence Argument, and (in some cases) suggestions for adjustments or modifications in the presentation or regimentation of the basic intuitive ideas, see (among others): Fischer (1983) and (1988); O’Connor (2000, pp. 3–18); Warfield (2000); and Warfield and Finch (1998). Perhaps the most visible critical discussion of the Consequence Argument is in Lewis (1981). For recent critiques, see (among many others) Campbell (2007) and (2008); Perry (2004); and Vihvelin (2004) and (2008).

question and thereby make progress in evaluating the Consequence Argument. We contend that, whatever the other liabilities of the argument may be, it does *not* beg the question.

2. Clarifying the Charge of Begging the Question

2.1 *The Consequence Argument*

One difficulty in focusing sharply on the question of whether the Consequence Argument begs the question is that there are so many different formulations of the argument. For example, Peter van Inwagen (1983, pp. 55–105) presents what he calls “three arguments for incompatibilism”—three different versions of the Consequence Argument. (There are also different understandings of “begging the question”; we return to this point below.) To simplify things, and to help focus the debate, we will give a relatively informal version of the Consequence Argument.⁵

We start with this implication of causal determinism: If causal determinism obtains in a possible world w , then any possible world that shares a time-slice with w (i.e., is the same as w with respect to *temporally intrinsic features* at some time) and has the same natural laws as w is identical with w . Alternatively, the idea is that if causal determinism obtains, then any possible world that shares both a time-slice with w (construed as above) and w 's laws of nature shares all time-slices with w . The argument also employs what we will call the Principle of the Fixity of the Past and Laws (PFPL):

⁵ The presentation of the Consequence Argument below (in the text) borrows heavily from Fischer and Ravizza (1998), especially pp. 22–3, which in turn was based on Fischer (1994), pp. 87–98. For a more formal development of this (and related) versions of the Consequence Argument, see Sobel (1998).

PFPL: An agent has it within his power in possible world w to do A only if his doing A can be an extension of the (temporally intrinsic) past in w , holding the natural laws fixed.⁶

The ‘can’ in the formulation of the principle is not meant to be the ‘free will’ can, but, rather, a wider notion of possibility, such as logical possibility. On this interpretation, the principle states that an agent has it in his power (in the ‘free will’ sense) in w to do A only if it is logically possible that his doing A be an extension of the past in w , holding the natural laws fixed. Carl Ginet (1990, pp. 102–3) puts the point nicely:

If I have it open to me now to make the world contain a certain event after now, then I have it open to me now to make the world contain everything that has happened before now plus that event after now. We might call this the principle that *freedom is freedom to add to the given past*

Now our version of the Consequence Argument begins with the assumption that causal determinism obtains in the actual world. Suppose further that someone S does A at time t_3 (in the actual world). It follows from the truth of causal determinism that any world that shares a time-slice with w and has the same natural laws as w is identical with w . So let us suppose that the state of the actual world at t_1 is P , and the set of natural laws in the actual world is L . Now it follows (from the hypothesis of causal determinism) that any possible world whose state is P at t_1 and whose laws are L is identical to the actual world—that is,

⁶ Fischer (1994), p. 88 and Fischer and Ravizza (1998), pp. 21–2. The original formulation of PFPL can be found in Ginet (1990). Fischer (1994), esp. pp. 98–109, argues that accepting PFPL can issue in an illuminating analysis of Newcomb’s Problem. For discussion, see: Carlson (1998) and Fischer (2001a) and (2001b). Kadri Vihvelin (in personal correspondence) has pointed out that PFPL, if intended as a necessary truth, appears to rule out the possibility of time travel. The issue of time travel is unfortunately beyond the scope of this paper, but in future work we hope to explore the relationship between time travel and various fixity principles.

any such world shares all time-slices with the actual world. And thus *no* such world contains *S*'s refraining from doing *A* at *t3*. Hence, *S*'s refraining from doing *A* at *t3* cannot be an extension of the actual past, holding the laws of nature fixed. The Principle of the Fixity of the Past and Laws then implies that *S* cannot at *t2* refrain from doing *A* at *t3* (where *t2* is prior to or contemporaneous with *t3*).⁷

2.2 *The charge of begging the question*

As we said above, there are many different formulations of the Consequence Argument—many different ways of regimenting the intuitive ideas that drive the argument. We take it that the crucial ingredients in all of these specific formulations are captured by PFPL; indeed, we take it that any version of the Consequence Argument must employ—in some way or another—the intuitive ideas of the fixity of the (temporally intrinsic) past and the fixity of the natural laws. As Nelson Pike (1966) once put it, the past is “over-and-done-with” and thus out of our control. Similarly, it is part of common sense that it is neither up to us whether the natural laws apply in a given situation nor what the natural laws are. These ideas are separable, and can be employed in different ways; but for the sake of simplicity in what follows, we will evaluate and defend the combined principle (PFPL), while assuming that the relevant agent does not have access to a world with different laws of nature. This assumption admittedly makes our project less ambitious, because it means that our argument for PFPL will not be sufficient to establish the *soundness* of the Consequence Argument.⁸ But

⁷ In addition to Fischer and Ravizza (1998), there are similar formulations of the Consequence Argument in van Inwagen (1983), pp 83–93; Ginet (1990), pp. 90–123; and Fischer (1994), pp. 87–110.

⁸ This assumption (that the relevant agent doesn't have access to a world with different laws) limits the project because our argument, which incorporates a certain sort of backtracking counterfactual, will not support acceptance of PFPL over acceptance of PFP by itself. And the compatibilist could arguably say what should be said about the central case in §3 while endorsing PFP but rejecting PFL. In other words, a full

since our intentions are more modest than a full-blown defense of the soundness of the Consequence Argument—and since without this assumption our argument wouldn't even be sufficient to establish our thesis—we will restrict our considerations to PFPL, given the assumption above (that the relevant agent does not have access to a world with different laws of nature). However: even though our ambitions with respect to the Consequence Argument are modest, we should emphasize that a defense of the Consequence Argument (against the charge of begging the question) is not the only upshot of our efforts here. In addition to defending the Consequence Argument, we will be providing a reason in favor of PFPL—something that's difficult to do, given how basic the principle is. So we view ourselves as shedding light not only on the dialectical issues surrounding the Consequence Argument, but also as motivating a basic principle that has, for the most part, not been well motivated in the literature.⁹

Not surprisingly, given that there are different specific formulations of the Consequence Argument, there are different versions of the claim that the Consequence Argument begs the question. But insofar as all formulations of the Consequence Argument employ (in one way or another) the basic ingredients in PFPL, we can assume that the charge of begging the question pertains to PFPL. And if we can clear PFPL of this charge,

defense of the *soundness* of the Consequence Argument would require an argument that PFPL can do some work that PFP can't. (Thanks to Andrew Bailey and E. J. Coffman for emphasizing these points in personal correspondence.)

⁹ To our knowledge, the only philosopher who has taken up this task (albeit under a different guise) is Warfield, in his (2000). Although Warfield's paper wasn't explicitly written as such (cf. note 15 on pp. 178–9), it can be construed as responding to a challenge offered by Flint in his (1987). (We should also note that Flint does not in that paper accuse the Consequence Argument of begging the question; he merely points out that it rests on a principle [roughly equivalent to what we have dubbed PFPL] that is unpalatable for compatibilists—and thus, we would add, is in need of argumentative support if the Consequence Argument is going to be of any use.)

then our defense should apply, *mutatis mutandis*, to any of the particular variations of the charge.

Thus we can summarize the charge as follows: it seems that if a compatibilist accepts PFPL, the game is up—more specifically, it appears as if the game is up before it has even begun (in earnest). A sense of dialectical unfairness here—a sense that the conclusion is already, as it were, contained in the Principle of the Fixity of the Past and Laws—gives rise to the charge that the argument begs the question.¹⁰

Perhaps the most well-known philosopher who has leveled the charge of begging the question against one version of the Consequence Argument is David Lewis (1981, p. 127, fn 5):

[Van Inwagen’s] definition [of “could have rendered false”] is roughly as follows: an agent could have rendered a proposition false iff he could have arranged things in a certain way, such that his doing so, plus the whole truth about the past, together strictly imply the falsehood of the proposition. On this definition, Premise 6 [“I could not have rendered *L* false.”] simply says that I could not have arranged things in any way such that I was predetermined not to arrange things in that way. It is uninformative to learn that the soft determinist is committed to denying Premise 6 thus understood.

¹⁰ Walton (1989, p. 245) offers the following characterization of begging the question:

According to Hamblin [1970, p. 32], the origin of the term “begging the question” is through the translation of Aristotle’s original Greek phrase *to en arche aitesthai* in turn translated into Latin as *petitio principii*, which means “beg for that which is in the question at issue.” The meaning of this curious phrase becomes clearer in the context of persuasion dialogue based on a conflict of opinion between two parties. In persuasion dialogue, one party may *ask to be granted* certain premises he needs to build up his case to persuade the other party of his thesis (his conclusion to be established in the dispute). The thesis (conclusion) is the *question* that is to be established by this party through his argument. Hence to include this conclusion within the premises asked to be granted is to *beg the question*, that is, to “beg for” the question (conclusion) that is supposed to be proved. In other words the fault is that of “begging for” something that should be earned through the work of argument.

Although Lewis is not addressing PFPL here, and it is not straightforward to interpret his remarks, he appears to be objecting to the same basic idea expressed by PFPL. Indeed, van Inwagen (2005, p. 346) interprets Lewis as objecting to the Consequence Argument as (roughly speaking) begging the question:

These words [i.e., Lewis's words] are harder to understand on the third or fourth reading than they seem to be on the first. If, as I say, I understand Lewis, he is implying that ... my argument for incompatibilism is circular or begs the question against the compatibilist or something of that sort.

Lewis is certainly not the only philosopher who has charged the Consequence Argument with begging the question. Christopher Hill (1992, p. 52) offers a principle that is essentially an instance of PFPL:

If it is within *J*'s power to perform an action, then there is a possible world *w* such that (i) *J* performs the action in *w*, (ii) the state of *w* up to the time at which *J* performs the action is exactly the same as the state of the actual world up to that time, and (iii) the actual laws of nature hold in *w*.

Hill contends that van Inwagen's Consequence Argument relies on this principle, but also that "it is illegitimate to appeal to [this principle] in the present context—to do so is to beg

the question.”¹¹ Also, Hilary Bok (1998, p. 97) says, “For this reason compatibilists can deny van Inwagen’s argument on the grounds that it begs the question against them.”¹²

2.3 *What exactly is begging the question?*

Above we pointed out that it can seem as though accepting PFPL gives the game away too soon; it is as though the conclusion (incompatibilism) is already contained in the premises or that one gets to the conclusion *simply* on the basis of what one has assumed, rather than on the basis of any argumentation. But of course in any good deductive argument, the conclusion will indeed be in some sense already present in the premises. That’s why Chisholm said, in the story presented by van Inwagen in the epigraph above, that he seems to have been accused of committing the fallacy of affirming the antecedent. So what exactly *is* begging the question?

No doubt there is a family (perhaps a dysfunctional family) of related dialectical infelicities that fall under the rubric, “begging the question.” But it seems to us that there is a core notion of begging the question, and it is this core notion that we are attempting to capture with the analysis below. First, though, we need to make a distinction between two different uses of “begging the question.” The first use is a general one: an argument might simply beg the question in general. The second use makes a specific reference to some audience or recipient: an argument might beg the question *against* that audience or recipient.

¹¹ Hill (1992, p. 53). For van Inwagen’s reply, see van Inwagen (1992).

¹² In addition to such philosophers as Bok, Lewis, and Hill (who have put the charge in print), various philosophers—including Joseph Keim Campbell and Kadri Vihvelin—have made this charge against the Consequence Argument in some comments on Vihvelin (2010). Campbell, for instance, says that an argument that appeals to PFPL “is not so much an argument as an expression of incompatibilist belief.” According to Campbell, appealing to PFPL begs the question against anyone who accepts the classical view of free will as the ability to do otherwise. Vihvelin says that “If the incompatibilist argues that we lack free will at a deterministic world because we lack the freedom to do otherwise, and then stipulates that the freedom to do otherwise is the power to ‘add to the given past, holding fixed the laws,’ then ... that begs the question.”

These uses are related, but it is possible that they refer to distinct notions. For example, an argument might be perfectly acceptable if directed toward one audience, but would beg the question if directed toward a different audience; this would seem to be a case in which the argument doesn't beg the question in general but does beg the question against certain audiences. And of course the possibility of distinct notions of begging the question implies the possibility of distinct *accusations* of begging the question: the Consequence Argument might be accused of begging the question in general, or it might be accused of begging the question against the compatibilist (as in some of the citations above). Thus we are faced with the question of which notion to analyze, and the question of which accusation to defend against.¹³

Although a comprehensive treatment of begging the question would require giving an account of both notions (or better yet, an account that unifies the notions), we will focus on the general notion, and thus will defend the Consequence Argument against the accusation that it begs the question in general. Three reasons justify this decision. First, we take the general notion to be more basic: an argument that begs the question in general will also beg the question against a particular audience, whereas the converse isn't true. Second, it seems to us that as far as the Consequence Argument is concerned, the allegedly distinct accusations amount to the same thing. And third, even if there is a difference between the two accusations, it seems appropriate to begin with the more general (and thus more serious) accusation. For if the more serious accusation sticks, then the question of whether the audience-relative accusation also sticks becomes less important. Thus we will focus on the general notion in what follows.

¹³ Thanks to E. J. Coffman, Tom Flint, and Aaron Segal for helping us realize the importance of this distinction.

One way to analyze this notion is to say that an argument begs the question just in case no one would accept the relevant premise¹⁴ unless he or she were already convinced of the conclusion.¹⁵ So, for example, it might be thought that no one would accept PFPL without already having accepted incompatibilism about causal determinism and freedom to do otherwise. This idea about the nature of begging the question seems slightly off the mark, however. It makes the issue of whether an argument begs the question hinge on what someone *would accept*, and thus makes what should be a logical or evidential point hinge inappropriately on psychological facts about particular agents. Thus we should adjust the suggestion so that we shift from the question of whether someone *would accept* the relevant premise (under certain circumstances) to the question of whether someone *has a reason* for accepting the premise. More specifically, our suggestion is that an argument begs the question just in case the proponent of the argument has no reasons for accepting the premise that are independent of the conclusion.¹⁶

¹⁴ It's also possible, of course, for an argument to beg the question in virtue of some combination of the premises, but since the offending premises can always be conjoined into one, we'll refer to them in the singular.

¹⁵ Wright (2001, p. 357) offers a similar account: "I have begged the question if I support a conclusion with reasons that would not be accepted in the context by anyone who did not already accept the conclusion."

¹⁶ Neal Tognazzini originally suggested (a conceptual ancestor of) this approach to specifying the nature of begging the question in a comment on Vihvelin (2010). There is a complication, however, that is worth noting. There are two cases in which someone (*S*) might not have any good independent reason to accept the relevant premise. In the first case, there simply aren't any such reasons. In the second case, such reasons exist but *S* isn't aware of them. And this distinction raises a further issue, which has to do with an ambiguity in the term "argument." The Consequence Argument is typically thought of as a set of numbered propositions (i.e., a set of premises, together with a conclusion). But as van Inwagen reminds us (1983, p. 69), a set of numbered propositions is merely a "bookkeeping device" that stands in for an argument. Strictly speaking, an argument (e.g., an argument for incompatibilism) consists of the relevant set of numbered propositions together with a defense of the premises. What this means is that the same bookkeeping device could refer to two different arguments—the first of which begs the question, and the second of which doesn't. This could happen if there exists some independent reason for accepting a particular premise, but the first argument doesn't appeal to that reason, and thus begs the question, whereas the second argument *does* appeal to that reason. (Notice also that when an argument begs the question in virtue of failing to appeal to an independent reason, that could be because the proponent of the argument isn't aware of the reason [as most often happens], or it could be because the proponent of the argument *is* aware of the reason but chooses not to appeal to it. This latter possibility appears to create a problem for our account, because it seems that someone

Given this account of the nature of begging the question, we can sharpen the issue by asking whether there are any reasons for accepting PFPL that are independent of the incompatibilist conclusion. Below (in § III) we will argue that indeed there is a reason to accept PFPL, quite apart from accepting incompatibilism. If we are correct about this point, the Consequence Argument does not beg the question (given our interpretation of the nature of begging the question).

2.4 *Analyzing the notion of independence*

Above we suggested that an argument begs the question insofar as its proponent has no reasons for accepting the premises that are independent of the conclusion. Although we consider this (brief and admittedly preliminary) account of begging the question to show great promise, it does introduce a further issue of interpretation having to do with the notion of *independence*. What exactly does it mean to say that a reason for accepting a proposition is independent of another proposition (given that the first proposition, perhaps together with others, allegedly implies the second)? This is a difficult question—one to which we don't have a fully satisfying answer. But we will consider three possible interpretations, and in so doing illustrate why our suggestion is a step in the right direction.

could offer an argument that is obviously logically circular and yet, according to our analysis, does not beg the question. (For an example of such a situation, cf. the “McCoy” case in Hazlett (2006, p. 346). For the sake of simplicity, and since we are not attempting a full-blown analysis of begging the question, we will ignore this possibility and make the idealizing assumption that anyone offering an argument will offer the best argument available to them, given their evidential state.) If the bookkeeping device doesn't make it clear what independent reasons there may be for the conclusion, then we who encounter the argument are prompted by the principle of charity to at least attempt to come up with them on our own. And if we're unable to do so, given an honest effort, then we are arguably within our rights to level the accusation of begging the question.

So: Although we'll ignore these complications in the main text, what we're arguing is that although the bookkeeping device most often associated with the Consequence Argument doesn't make it clear whether or not the argument itself begs the question, there are reasons for accepting the relevant premise that are independent of the conclusion—and thus the argument itself *need not* beg the question. (Thanks to Joe Campbell, Jacob Domeyer, and Neal Tognazzini for bringing these complications to our attention.)

Perhaps the best place to start is with *non-identity*. On a non-identity interpretation of independence, an argument does *not* beg the question as long as there exists some (good) reason to accept the relevant premise that is not identical to the conclusion.¹⁷ But this interpretation is too permissive. A conjunction of the conclusion and some other true proposition could be a reason to accept the premise, and would not be identical to the conclusion—but of course it would also not be independent of the conclusion in any interesting way.

Another straightforward interpretation involves *entailment*. Perhaps we should say that an argument begs the question when (and only when) there are no reasons for accepting the premise that do not entail the conclusion—i.e., when (and only when) every reason for accepting the premise entails the conclusion. On this interpretation of independence, an argument avoids begging the question only if there is at least one reason for accepting the premise that doesn't entail the conclusion. While this interpretation is definitely an improvement on the non-identity interpretation, it is actually too strict—as we can see by considering the following type of argumentative situation. If an argument contains a premise such that all reasons for accepting that premise entail the premise, and if all of the premises (of the argument in question) are necessarily true, then every reason for accepting the premise will entail the conclusion. According to the entailment interpretation, any argument in this type of situation would beg the question. But clearly we don't want to say that all such

¹⁷ On the identity interpretation, an argument begs the question iff there are no reasons (to accept the relevant premise) that are not identical to the conclusion. Contraposing and canceling a double negation gives us the following: An argument does not beg the question iff there exists a reason to accept the premise that is not identical to the conclusion.

arguments beg the question; saying that would almost be like penalizing the available reasons for strongly supporting the premise.¹⁸

These failures lead us to suspect that it will be difficult, if not impossible, to cash out the notion of independence in purely logical terms. (Begging the question is, after all, supposed to be an informal fallacy.) Thus, we are going to suggest that the best way to interpret the notion of independence (or at least the first step in the right direction) is in terms of the evidential state of the argument's proponent *after* subtracting out his or her acceptance of the conclusion. Our suggestion is that an argument begs the question just in case the proponent of the argument has no reason to accept the relevant premise, apart from a prior acceptance of the conclusion.¹⁹

Let us briefly take stock. We claimed first that psychologizing the notion of begging the question (by characterizing it in terms of whether someone *would accept* a particular premise) is problematic. We suggested a slightly different characterization of the notion, which itself hinges on a notion of independence. Our suggestion shifts the focus from a psychological question (i.e., the question of whether the *recipient* of the argument *would accept*

¹⁸ Entailment in the other direction won't work either. If we say that "independent of the conclusion" should be interpreted as "not entailed by the conclusion," then an argument avoids begging the question if there's at least one reason to accept the premise that isn't entailed by the conclusion. But this interpretation, much like the non-identity interpretation, is too permissive. A conjunction of the conclusion and some other true proposition could be a reason to accept the premise, and would not be entailed by the conclusion, but of course would not be independent of the conclusion in the sense we're looking for. (Thanks to Aaron Segal for his help in highlighting the difficulties in finding a suitable interpretation of the notion of independence.)

¹⁹ It's worth highlighting here that our analysis is in terms of someone's *having* a certain kind of reason, rather than there merely *being* such a reason. If we analyze independence in terms of there *being* such a reason, then we are forced to say that an argument begs the question only if there exist no reasons for the relevant premise that are independent of the conclusion. But consider a case (which we borrow from Aaron Segal in which the relevant premise p is identical to the conclusion c , and there is a reason r for p that entails p —but the proponent of the argument is not aware of this entailment relation (i.e., not aware that there is a reason for p). This argument clearly begs the question (since $p = c$), but there *is* a reason for p (namely, r) that exists apart from a prior acceptance of c . So, given our interpretation of independence, it can't be right to say that an argument begs the question only if there exist no reasons for the relevant premise that are independent of the conclusion.

the relevant premise) to an evidential question (i.e., the question of whether the *proponent* of the argument *has any reasons* for the premise, apart from a prior acceptance of the argument's conclusion). It is, we admit, not obvious how to analyze the crucial notion of independence, and all we have offered is a step in the right direction. But as we see it, our informal analysis threads the needle between two inadequate accounts: the psychologized notion of begging the question and the notion that relies on simple logical analyses of independence. More importantly, our account captures what we take to be the intuitive idea behind the accusation: An argument begs the question just in case the proponent has no reason to accept the relevant premise, absent accepting the conclusion.

Note finally that this account of begging the question does not require a *decisive* independent reason to accept the premise. Having a decisive independent reason would obviously be preferable (from the standpoint of defending the relevant argument), but we will content ourselves with looking for a reason that at least renders it plausible to accept the relevant premise (apart from accepting the conclusion). More specifically, we will identify a set of considerations that render it plausible to accept PFPL but don't rest on incompatibilism about casual determinism and freedom to do otherwise. First, however, we will consider a potential objection to our account.

2.5 What our account says about Moore's anti-skeptical argument

G. E. Moore's argument for an external world, which has often been accused of begging the question,²⁰ will serve as a good test case for our account. Moore's argument runs as follows:

²⁰ Compare Hazlett (2006) and Pryor (2004), both of whom cite examples of the accusation.

- (1) Here are two hands.
- (2) If hands exist, then there is an external world.
- (3) So there is an external world.²¹

Most philosophers would agree that there's something infelicitous about this argument, at least insofar as it is directed toward the skeptic. Some would go further, and say that any viable account of begging the question must deliver the verdict that Moore's argument begs the question.²² Thus it's worth asking what our account has to say about Moore's argument. According to our account, Moore's argument begs the question iff Moore has no reason to accept the relevant premise—"Here are two hands"—apart from a prior acceptance of the conclusion. Moore does antecedently accept the conclusion, but is that his only reason for affirming the premise? Clearly not—for the way in which he affirms the premise (holding up his hands, and waving them around in a certain way) indicates that it is supported by evidence (namely, perceptual and proprioceptive evidence). According to our account, then, Moore's argument doesn't beg the question after all. This verdict is obviously problematic if it is indeed true that any viable account of begging the question should deliver the opposite verdict. But we think there are good reasons to resist this desideratum on accounts of begging the question. We will grant that Moore's argument does exhibit a certain kind of epistemic circularity, or epistemic dependence (between premise and conclusion), but we would argue, as others have, that whereas some kinds of circularity are malignant, the kind of

²¹ The argument comes from Moore (1962), but this particular formulation of the argument comes from Pryor (2004).

²² Thanks to E. J. Coffman for making us aware of this proposed desideratum on accounts of begging the question.

circularity that characterizes Moore's argument is benign.²³ Thus we think that our account actually gives the right result when applied to Moore's argument.

Now of course not everyone has been convinced that some forms of epistemic circularity are benign. To those who remain unconvinced, notice that our account is flexible enough to accommodate the intuition that Moore's argument does in fact beg the question. For we could adjust our account by adding the notion of a *contrastive* reason. One complaint about Moore's argument is that his perceptual or proprioceptive evidence for premise (1) does not undermine the relevant skeptical counterpossibility that is inconsistent with that premise (e.g., the skeptical counterpossibility that he is merely a brain in a vat). In other words, Moore's reason for believing the premise is not a contrastive reason because even though it does count in favor of the premise, it doesn't count *against* the relevant counterpossibility. (Thus we will say—without attempting to provide an analysis of contrastive reasons in general—that some reason *r* is a contrastive reason for some argument premise *p* only if *r* supports *p* while undermining the relevant counterpossibility.²⁴) With this

²³ Both Alston (1986) and Bergmann (2004) have argued that some types of epistemic circularity are benign. Pryor (2004) argues in similar fashion with respect to certain types of epistemic dependence, although he is admittedly more concerned with the *reasoning processes* underlying Moore's argument than he is with the argument itself. In particular, he is primarily concerned with the question of whether the reasoning process represented by Moore's argument is legitimate—i.e., whether that reasoning process can instantiate a genuine justificatory structure. (He thinks it can.) Pryor then points out that even if Moore's reasoning process is legitimate, there may nevertheless be some audience (e.g., a skeptical audience) that would not be convinced by an argument based on that reasoning process. This point presumably generalizes to other contexts in which an argument based on a reasoning process is presented to an audience, which means that even though the Consequence Argument does not beg the question, there may nevertheless be some audiences who would not be convinced by it. This is another issue that would need to be addressed in a full defense of the Consequence Argument.

²⁴ Notice that we are deliberately avoiding the requirement that a contrastive reason *entail the falsity* of the relevant counterpossibility. This requirement would be too strong. There are some cases in which the relevant counterpossibility is simply the negation of the conclusion, and in these cases a sub-argument for a premise that entails the falsity of the counterpossibility would thereby entail the truth of the conclusion—thus rendering the main argument unnecessary. And if a sub-argument that entails the main argument's conclusion is the only way to support the relevant premise, then the main argument *virtually* begs the question. (The notion of virtually begging the question comes from Klein (1995).) Thus an account of begging the question that required that a contrastive reason entail the falsity of the relevant counterpossibility would be self-defeating, as it would in some cases trade begging the question for virtually begging the question.

notion in hand, we can say that an argument begs the question just in case the proponent has no *contrastive* reason to accept the relevant premise, absent accepting the conclusion.²⁵ In other words, even if the proponent of an argument has an independent reason for the relevant premise, the argument still begs the question unless that reason is a contrastive reason. According to this modified account, Moore's argument would after all beg the question. Thus, while we maintain that our account gives the right results without the notion of a contrastive reason, if there turns out to be a decisive argument against the claim that certain kinds of epistemic circularity are benign, then we can modify our account to incorporate that development without sacrificing our thesis, namely that the Consequence Argument does not beg the question.²⁶

3. A Plausibility Argument for the Principle of the Fixity of the Past and Laws

3.1 A problem for the rejection of PFPL

We now begin to develop an argument for PFPL—an argument that does not rely in any way on a prior acceptance of incompatibilism. PFPL has it that our freedom is the power to add to the given past, holding fixed the laws of nature. This is indeed an intuitive picture. We contend further that if one rejects PFPL, one is in jeopardy of being committed to implausible conclusions about practical reasoning. That is, we will argue that rejecting PFPL could lead to one's having to say that it is rational to choose certain actions in contexts in which it is manifestly irrational to choose those actions. In a nutshell: it is at least plausible that rejecting PFPL would lead to unacceptable consequences for practical reasoning. Thus,

²⁵ Note also that we are not endorsing the more general thesis of contrastivism about reasons (cf. Sinnott-Armstrong (2008)).

²⁶ In note 29 we say a little bit more about why the Consequence Argument doesn't beg the question even if an independent contrastive reason for the premise is required.

there is a plausibility argument for accepting PFPL—an argument that does not depend on incompatibilism.

There are certain examples of which it is plausible to say that backtracking subjunctive conditionals (of a specific kind) are true. Various philosophers have offered examples of this sort, and here is one:

Consider the example of the Icy Patch. Sam saw a boy slip and fall on an icy patch on Sam’s sidewalk on Monday. The boy was seriously injured, and this disturbed Sam deeply. On Tuesday, Sam must decide whether to go ice-skating. Suppose that Sam’s character is such that if he were to decide to go ice-skating at noon on Tuesday, then the boy would not have slipped and hurt himself on Monday.

The situation is puzzling. It seems that Sam is able to decide to go and to go ice-skating on Tuesday.²⁷ And it also appears plausible that if he were to decide to go skating on Tuesday, the terrible accident would not have occurred on Monday. So it appears that Sam *ought* to decide to go ice-skating on Tuesday. And yet, given that Sam knows that the accident did in fact take place on Monday, it also seems *irrational* for Sam to decide to go ice-skating on Tuesday on the basis of a reason flowing from the truth of the backtracker. Nothing *prevents* Sam from deciding to go and from going ice-skating on Tuesday; and if he were to decide to go ice-skating, the accident would not have occurred. And yet it seems inappropriate for Sam to decide to go ice-skating. To do so would seem to exemplify something akin to wishful thinking. (Fischer 1994, p. 95)

Of course, someone might deny that the backtracker (“If Sam were to decide to go ice-skating on Tuesday, the accident would not have occurred on Monday”) is true in the example. After all, if Sam were to decide to go ice-skating on Tuesday, perhaps he would

²⁷ In other words, we can suppose that nothing is preventing Sam from acting out of character by making a decision to go ice-skating. We can be confident that he *won't* act out of character, but it doesn't follow that he *can't*.

have forgotten about the accident or he would have been acting out of character, and so forth. We do not have any knockdown argument that the backtracker is true in the example, but we believe that the story can be filled in so that it is at least *plausible* that the backtracker is indeed true. (Nothing in our argument here depends on defending the claim that the backtracker is obviously true; all we need is the claim that it is at least plausible that, if one fills in the details suitably, the backtracking conditional would be true.)

If we accept PFPL, we can say what we *should* say about practical reasoning in a case such as *Icy Patch*. Intuitively, what Sam can do on Tuesday is to add to the given past (in which the terrible accident did indeed take place on Monday). So any reasons flowing from the non-occurrence of the accident on Monday are just irrelevant to Sam. But how exactly can one embrace this obvious point, if one rejects PFPL? After all, if one rejects PFPL, one is committed to the contention that in some contexts an agent in world *w* has access to possible worlds with *different pasts* from *w*. Then why exactly isn't *Icy Patch* one of these contexts? Having abandoned PFPL, why isn't *Icy Patch* an example in which Sam has access on Tuesday to a possible world in which the accident didn't happen on Monday? More specifically, given a rejection of PFPL, why can't Sam bring it about on Tuesday—simply by deciding to go ice-skating—that the world did not contain the accident on Monday?

Recall (a slightly modified but logically equivalent version of) Carl Ginet's principle that our freedom is the power to add to the given past, holding fixed the laws of nature:

I have it open to me now to make the world contain a certain event after now only if I have it open to me now to make the world contain everything that has happened before now plus that event after now.

(This principle, of course, is equivalent to PFPL, and thus if we reject one then we must reject the other.) In favor of this principle, Ginet offers the following characterization of our deliberative perspective:

My impression at each moment is that *I at the moment*, and nothing prior to that moment, determine which of several open alternatives is the next sort of bodily exertion I voluntarily make. (Ginet 1990, p. 90, emphasis in original)

But notice that if we accept this characterization of our deliberative perspective, while rejecting PFPL, then we are essentially saying that some of those open alternatives are alternatives in which the past contains an event that is not a part of the actual past (or lacks an event that *is* a part of the actual past). In other words, rejecting PFPL commits us to the claim that at least in some contexts, if we have it open to us now to make the world contain a certain event after now, then we have it open to us now to make the world contain some different event in the past plus that event after now. But then why isn't the *Icy Patch* case just such a context? More specifically, on the assumption that PFPL is false, why shouldn't Sam choose to go ice-skating on Tuesday? After all, it is open to him on Tuesday to make the world contain the accident's not happening on Monday plus his decision on Tuesday to go ice-skating. But this is manifestly an unacceptable conclusion.

Perhaps it would be helpful to note that in supposing that it is open to Sam on Tuesday to make the world contain the accident's not happening on Monday plus his decision on Tuesday to go ice-skating, we are not thereby supposing that Sam can on Tuesday initiate a backward-flowing causal chain issuing in the accident's not happening on Monday. Rather, we are simply supposing that he can decide on Tuesday to go ice-skating, thereby making the world contain the accident's not happening on Monday plus this

decision on Tuesday. Of course, setting aside the question of whether Sam can on Tuesday cause the accident's not having occurred on Monday, if he can on Tuesday make the world contain the accident's not occurring on Monday, this will no doubt be relevant to his practical reasoning on Tuesday.

So if one rejects PFPL, there will be at least some contexts in which an agent has access to a possible world with a different past. It is possible that some such contexts have the distinctive structure of *Icy Patch*. In such a case, given a rejection of PFPL, it is rational for the agent to choose or decide to do something that is intuitively clearly irrational to choose or decide to do. Thus, it is plausible to claim that the rejection of PFPL leads to unacceptable conclusions for practical reasoning. And if this is correct, then we have a reason to accept the relevant premise of the Consequence Argument that does not depend on the argument's conclusion (that causal determinism and freedom are incompatible).

3.2 Sam and the icy patch: an alternative (but similar) analysis

Intuitively, in *Icy Patch*, Sam can decide on Tuesday to go ice-skating. Also, it seems as if the back-tracking conditional is true: "If Sam were to decide to go ice-skating on Tuesday, then the accident would not have occurred on Monday." Adopting the standard Stalnaker/Lewis semantics for such conditionals, we can say that in the *closest* possible worlds in which Sam decides on Tuesday to go ice-skating, the accident did not occur on Monday. That is, in the world or worlds most similar to the actual world in which Sam decides on Tuesday to go ice-skating, the accident did not occur on Monday.

Consider now the "can-claim," according to which Sam can decide on Tuesday to go ice-skating. It will be helpful to accept, as a working hypothesis, the broad outlines of Keith Lehrer's "possible-worlds analysis" of can-claims. Lehrer (1976, pp. 253–4) says:

... when we say that a person could have done something he did not do, we should not, and I believe do not, thereby affirm that every antecedent necessary condition of his performing the action is fulfilled. It is enough that there be some possible world minimally different from the actual world restricted in an appropriate way so that the person performs the actions and those conditions are fulfilled. We may speak of worlds restricted in the appropriate way as possible worlds that are *accessible* to the agents from the actual world.

So on this sort of possible-worlds account of “can,” an agent can in possible world w perform some action only if he or she does perform the action in a possible world w^* that is accessible to the agent from w . Not all possible worlds are accessible to the agent from w ; the accessible worlds are “suitably restricted.” Of course, the proponent of PFPL will insist that accessibility requires sameness of past (as well as sameness of natural laws); in contrast, a compatibilist such as Lehrer who rejects PFPL will not require sameness of past and natural laws for accessibility.

We contend that it is appropriate to take the reasons that obtain in all worlds accessible from the actual world to the agent at a given time as relevant to the agent’s practical reasoning at the time. Intuitively, the agent has access to those worlds—he can get to those worlds from the actual world. Less metaphorically, he can actualize those worlds. Why not then take the reasons that obtain in those worlds as relevant to the agent’s practical reasoning? If an agent can’t get to a scenario from where he is, then the features of that scenario would seem to be irrelevant to his practical reasoning. But if the agent can indeed get to the scenario from where he is, then why not deem certain features of that scenario relevant to the agent’s practical reasoning? It would seem entirely arbitrary not to allow the agent to take into account reasons that obtain in scenarios genuinely accessible to him

Given a rejection of PFPL, nothing rules it out that in *Icy Patch*, Sam has access on Tuesday to a possible world in which the accident did not occur on Monday. After all, Sam can decide on Tuesday to go ice-skating that afternoon. If this can-claim is rendered true by a possible world w^* accessible to Sam, and on the assumption that we reject PFPL, nothing rules it out that w^* will contain the accident's not happening on Monday. So Sam would have access on Tuesday to a possible world in which the accident did not occur on Monday. (This is just another way of putting the point in the previous section, namely that, given a rejection of PFPL, it is plausible that Sam can on Tuesday make the world contain [among other things] the accident's not happening on Monday). But if this is so, why shouldn't Sam take this as a reason to decide to go ice-skating on Tuesday? If it is appropriate for Sam to take as relevant reasons that obtain in any world genuinely accessible to him at a time, then surely it is (or may well be) rational for him to decide to go ice-skating on Tuesday. But, again, this is a manifestly unacceptable result: it is clearly *irrational* for him so to decide on Tuesday. The problem, once again, is that the rejection of PFPL leads to unacceptable results for practical reasoning. Thus, as above, we have a reason to accept the relevant premise of the Consequence Argument that does not stem from its conclusion.²⁸

²⁸ Note also that the Consequence Argument can be absolved of the charge of begging the question even if an independent contrastive reason is required. The relevant counterpossibility (that's inconsistent with both the relevant premise and the conclusion) is a situation in which determinism is true and S has it within his power to do A even though A cannot be an extension of the temporally intrinsic past in w , holding the natural laws fixed. (In other words, the relevant counterpossibility is simply a situation in which compatibilism is true.) Our defense of PFPL undermines (without entailing the falsity of) this counterpossibility as follows: If it would be irrational when deliberating to consider worlds with a different past, then it would be odd to say that S nonetheless has the power to do otherwise even though doing otherwise cannot be an extension of w 's past. Why would it be irrational to consider (bringing about) a situation that is within one's power to bring about?

4. Gallois's Approach and Diagnosis

In a recent article, Andre Gallois (2009) discusses *Icy Patch* and similar examples. Just so that we have another example of the same sort as *Icy Patch*, it will be useful to have before us an example, due to P.B. Downing, presented by David Lewis. (The example is discussed by Gallois, along with *Icy Patch* and similar cases.) Lewis says:

We know that present conditions have their past causes. We can persuade ourselves, and sometimes do, that if the present were different then these past causes would have to be different. Given such an argument—call it a *backtracking argument*—we willingly grant that if the present were different, the past would be different too.

Jim and Jack quarreled yesterday, and Jack is still hopping mad. We conclude that if Jim asked Jack for help today, Jack would not help him. But wait: Jim is a prideful fellow. He never would ask for help after such a quarrel; if Jim were to ask Jack for help today, there would have been no quarrel yesterday. In that case Jack would be his usual generous self. So if Jim asked Jack for help today, Jack would help him after all.²⁹

Clearly, Jim has it in his power today to ask Jack for help. But it also appears to be true that if Jim were to ask Jack for help today, then the quarrel would not have happened yesterday. Given that it would be much better for Jim if the quarrel hadn't occurred yesterday, then why exactly isn't it rational for Jim to choose to ask Jack for help today?

We assume that Jim can today ask Jack for help. On the assumption that PFPL is rejected, it seems that Jim can today make it the case that the world contains the quarrel's not having taken place plus Jim's asking Jack for help later today. Given that Jim can do this—that Jim has genuine access to a world in which the quarrel wouldn't have occurred

²⁹ Lewis (1986, p. 33). Downing gave essentially the same example in his (1959). Gallois (2009) presents the example on p. 236.

yesterday—it seems rational for Jim to choose to ask Jack for help today. But, clearly, it is irrational for Jim to ask Jack for help today, given that the quarrel did in fact take place yesterday; more specifically, it is not appropriate for Jim to take it as a reason for asking Jack for help today that if he were to do so, the quarrel never would have occurred.

Gallois offers the following diagnosis of the examples in question. In *Icy Patch*, Sam actually has reason not to decide to go ice-skating on Tuesday insofar as the terrible accident occurred on Monday. Gallois contends that Sam must keep this reason fixed in his deliberations; this follows from a general principle of “fixity of reasons” defended by Gallois (2009), especially pp. 238–9. Similarly, Jim actually has reason not to ask Jack for help today, given that the quarrel did in fact take place yesterday. Gallois contends that Jim must keep this reason fixed in his deliberations.

Note that Gallois insists that the agents must keep their *actual* reasons fixed, and must not be allowed to consider reasons that obtain in worlds that are *accessible to the agents*. This Principle of Fixity of Reasons does yield the intuitively correct results in the relevant cases: it yields the result that it is irrational for Sam to choose on Tuesday to go ice-skating later that afternoon, and the result that it is irrational for Jim to choose today to ask Jack for help later in the day. But it seems to us that Gallois achieves these results by invoking an *ad hoc* principle. Given that I can make the world contain a certain state of affairs, why shouldn't I take that state of affairs (and any related states of affairs) as relevant to my practical reasoning? Alternatively, given that I have genuine access to a possible world in which a certain state of affairs obtains in the past, why shouldn't I take that state of affairs (and any related states of affairs) as relevant to my practical reasoning? Gallois's invocation of the Principle of the Fixity of Reasons appears to us entirely *ad hoc*; why exactly shouldn't

agents consider the features of accessible scenarios—scenarios they can genuinely get to—in their practical reasoning?

Of course, we agree with Gallois about the irrationality of Sam’s choosing on Tuesday to go ice-skating and Jim’s choosing today to ask Jack for help; but we have a different strategy for getting to these results. As we just pointed out, we find it entirely *ad hoc* and unsatisfying simply to dig in one’s heels and insist that certain reasons in *accessible* scenarios—scenarios that an agent can genuinely “get to”—are irrelevant to practical reasoning. This is the wrong explanation for the correct result. A more attractive approach is to allow consideration in practical reasoning of features of all accessible scenarios, but to restrict accessibility to scenarios with the same past as the actual scenario. That is, we find it much more attractive to accept PFPL, according to which one can only actualize those possible worlds that are extensions of (share the past with) the actual world.

Accepting PFPL allows us to say precisely the right things about cases such as *Icy Patch* without invoking an *ad hoc* principle, such as the Principle of the Fixity of Reasons. It also allows us to say the correct things about (say) Sam’s practical reasoning without insisting that the backtracking conditionals are false. Indeed, one can fully accept the truth of the backtracking subjunctive conditionals compatibly with embracing PFPL.³⁰ Put simply, the can-claim (e.g., “Sam can on Tuesday decide to go ice-skating later that afternoon”) and the backtracking conditional (e.g., “If Sam were to decide on Tuesday to go ice-skating, then the accident would not have taken place on Monday”) are *different* modalities that point us to *different* possible worlds. Whereas the conditional is true in virtue of the *closest* possible worlds in which Sam decides on Tuesday to go ice-skating, the can-claim need not be underwritten

³⁰ For a defense of this contention, see Fischer (1994, pp. 87–110).

by the *closest* such worlds. Therefore, the possible world or worlds in virtue of which the conditional is true need not be the same as the possible world (or worlds) in virtue of which the can-claim is true. Thus, PFPL is entirely compatible with the truth of the backtrackers in cases such as *Icy Patch*.³¹

To recap. Gallois's Principle of the Fixity of Reasons invites the question: Why exactly shouldn't we consider as reasons in our practical reasoning features of scenarios that are genuinely accessible to us—that we can get to from here? Absent an answer to this question, invocation of the principle is *ad hoc*. In contrast, it seems natural and appealing to suppose that our freedom is the power to add to the given past, holding fixed the laws of nature, and further that the reasons relevant to our practical reasoning obtain in scenarios we can get to from where we are. Whereas we agree with Gallois's conclusions, we take a different route, and the ingredients we employ are (unlike Gallois's Principle of the Fixity of Reasons) independently plausible.

5. Conclusion

Various philosophers, including no less a luminary than David Lewis, have contended that the Consequence Argument, in some form or another, begs the question (on some account or implicit understanding of the charge). We have attempted to understand and evaluate this charge. We claim that it is illuminating to think of begging the question as involving the assertion of a premise such that one has no reasons to accept the premise that are independent of accepting the conclusion. We have given an account of the crucial notion of independence that threads the needle between a purely psychological and purely logical

³¹ For a more detailed defense, see Fischer (1994, pp. 87–110) and (2008, pp. 184–6).

interpretation: on our view, an argument begs the question just in case the proponent of the argument has no reason to accept the relevant premise, apart from a prior acceptance of the conclusion. We have argued that, in this sense, the Consequence Argument need not beg the question. There *is* a reason to accept PFPL that does not stem from accepting incompatibilism, namely that rejection of PFPL arguably has implausible results for practical reasoning.

We certainly concede that our argument for the contention that rejecting PFPL will lead to implausible results for practical reasoning is not airtight. One might resist various steps along the way. But we also wish to emphasize that this argument *need not be airtight*. What is required, in order for an argument to escape the charge of begging the question, is a good (but not necessarily decisive) reason to accept the relevant premise that is independent of accepting the conclusion. One needs evidence for the premise that does not stem from the conclusion; but this evidence need not be decisive.

In addition, we have offered something more than a way of conceptualizing the Consequence Argument so that it need not beg the question. We have offered a reason to accept a crucial premise of the Consequence Argument—a premise often thought to be “basic.” It is notoriously difficult to offer arguments or reasons to accept the crucial fixity of the past (and laws) premise in the Consequence Argument. In this paper we have done precisely this: we have shown that the only way to avoid implausible results in practical reasoning is to accept PFPL. Thus, we have offered a strong reason to accept what many have thought must be accepted, without argument, as basic.

The Consequence Argument is highly contentious; we do not here take a stand on whether it is sound. Certainly, reasonable people—and even philosophers!—can disagree about its soundness. Here we hope to have made some progress to the extent that we have

offered reason to think that the Consequence Argument at least does not beg the question. For some this will offer consolation, insofar as the Consequence Argument seems to them so obviously sound. For others it will provide reason to suppose that the Consequence Argument at least rises to a level at which it can be subjected to serious consideration—and perhaps even refutation.³²

John Martin Fischer, University of California, Riverside

Garrett Pendergraft, Pepperdine University

³² Previous versions of this paper were presented to the Metaphysics Reading Group (organized by Michael Rea) at the University of Notre Dame and also the Philosophy Department at the University of Tennessee, Knoxville; we are very grateful for the helpful comments we received on those occasions. We would also like to thank Alex Arnold, Andrew Bailey, Kenneth Boyce, Joseph Keim Campbell, D. Justin Coates, Thomas P. Flint, Michael C. Rea, Bradley Rettler, Aaron Segal, Philip Swenson, Patrick Todd, Neal A. Tognazzini, and Kadri Vihvelin for helpful comments on earlier versions of this paper. (We are particularly grateful to E. J. Coffman for detailed and extremely helpful comments on various drafts.) In addition, we are grateful for comments from an anonymous reviewer for *Philosophical Studies*, which helped shape this paper into its final form.

References

- Alston, W. P. (1986). Epistemic circularity. *Philosophy and Phenomenological Research*, 47, 1–30.
- Bergmann, M. (2004). Epistemic circularity: Malignant and benign. *Philosophy and Phenomenological Research*, 69, 709–27.
- Bok, H. (1998). *Freedom and responsibility*. Princeton: Princeton University Press. Portions reprinted in G. Watson (Ed.), *Free will*, Second edition (pp. 130–66). New York: Oxford University Press.
- Campbell, J. K. (2007). Free will and the necessity of the past. *Analysis*, 67, 105–11.
- Campbell, J. K. (2008). Reply to Brueckner. *Analysis*, 68, 264–9.
- Carlson, E. (1998). Fischer on backtracking and Newcomb’s problem. *Analysis*, 58, 229–31.
- Downing, P. B. (1959). Subjunctive conditionals, time order, and causation. *Proceedings of the Aristotelian Society*, 59, 125–40.
- Fischer, J. M. (1983). Incompatibilism. *Philosophical Studies*, 43, 127–137.
- Fischer, J. M. (1988). Freedom and miracles. *Noûs*, 22, 235–252.
- Fischer, J. M. (1994). *The metaphysics of free will: An essay on control*. Oxford: Blackwell Publishers.
- Fischer, J. M. (2001a). Newcomb’s problem: A reply to Carlson. *Analysis*, 61, 229–36.
- Fischer, J. M. (2001b). Critical notice of J. Howard Sobel’s *Puzzles of the will*. *Canadian Journal of Philosophy*, 31, 427–444.
- Fischer, J. M. (2008). *My way and life’s highway: Replies to Steward, Smilansky, and Perry*. *Journal of Ethics*, 12, 167–89.
- Fischer, J. M. and Ravizza, M., (1998). *Responsibility and control: A theory of moral responsibility*. New York: Cambridge University Press.
- Fischer J. M. and Todd, P. (2011). The truth about freedom: A reply to Merricks. *The Philosophical Review*, 120, 97–115.
- Flint, T. P. (1987). Compatibilism and the argument from unavoidability. *The Journal of Philosophy*, 84, 423–40.
- Gallois, A. N. (2009). The fixity of reasons. *Philosophical Studies*, 146, 233–48.
- Ginet, C. (1966). Might we have no choice? In K. Lehrer (Ed.), *Freedom and determinism* (pp. 87–104). New York: Random House.
- Ginet, C. (1980). The conditional analysis of freedom. In P. Van Inwagen (Ed.), *Time and cause: Essays presented to Richard Taylor* (pp. 171–86). Dordrecht: D. Reidel.
- Ginet, C. (1990). *On action*. Cambridge: Cambridge University Press.
- Hamblin, C. L. (1970). *Fallacies*. London: Methuen and Company.
- Hazlett, A. (2006). Epistemic conceptions of begging the question. *Erkenntnis*, 65, 343–363.
- Hill, C. (1992). Van Inwagen on the consequence argument. *Analysis*, 52, 49–55.

- Klein, P. (1995). Skepticism and closure: Why the evil genius argument fails. *Philosophical Topics*, 23, 213–36.
- Lehrer, K. (1976). ‘Can’ in theory and practice: A possible-worlds approach. In M. Brand and D. Walton (Eds.), *Action theory* (pp. 241–70). Dordrecht: Reidel.
- Lewis, D. (1981). Are we free to break the laws? *Theoria*, 47, 113–21.
- Lewis, D. (1986). Counterfactual dependence and time’s arrow. In D. Lewis, *Philosophical Papers volume II* (pp. 32–66). New York: Oxford University Press.
- Merricks, T. (2009). Truth and freedom. *The Philosophical Review*, 118, 20–57.
- Merricks, T. (2011). Foreknowledge and freedom. *The Philosophical Review*, 120, 567–86.
- Moore, G. E. (1962). Proof of an external world. In G. E. Moore, *Philosophical papers* (pp. 144–8). New York: Collier Books.
- O’Connor, T. (2000). *Persons and causes: The metaphysics of free will*. New York: Oxford University Press.
- Perry, J. (2004). Compatibilist options. In J.K. Campbell, M. O’Rourke, and D. Shier (Eds.), *Freedom and determinism* (pp. 231–54). Cambridge, MA: MIT Press.
- Pike, N. (1965). Divine omniscience and voluntary action. *The Philosophical Review*, 74, 27–46.
- Pike, N. (1966). Of God and freedom: A rejoinder. *The Philosophical Review*, 75, 369–79.
- Pryor, J. (2004). What’s wrong with Moore’s argument? *Philosophical Issues*, 14, 349–78.
- Sinnott-Armstrong, W. (2008). A contrastivist manifesto. *Social Epistemology*, 22, 257–70.
- Sobel, J. H. (1998). *Puzzles for the will*. Toronto: University of Toronto Press.
- Van Inwagen, P. (1975). The incompatibility of free will and determinism. *Philosophical Studies*, 27, 185–99.
- Van Inwagen, P. (1983). *An essay on free will*. Oxford: Clarendon Press.
- Van Inwagen, P. (1992). Reply to Christopher Hill. *Analysis*, 52, 56–61.
- Van Inwagen, P. (2005). Freedom to break the laws. In P. French and H. Wettstein (Eds.), *Midwest Studies in Philosophy*, 28, 334–50.
- Vihvelin, K. (2004). Free will demystified: A dispositional account. *Philosophical Topics*, 32, 427–50.
- Vihvelin, K. (2008). Compatibilism, incompatibilism, and impossibilism. In J. Hawthorne, T. Sider, and D. Zimmerman (Eds.), *Contemporary debates in metaphysics* (pp. 303–18). Malden, MA: Blackwell Publishing.
- Vihvelin, K. (2010). Clearing the path. Blog post. *Flickers of Freedom*. <http://agencyandresponsibility.typepad.com/flickers-of-freedom/>. Accessed 28 March 2010.
- Walton, D. (1989). *Informal logic: A handbook for critical argumentation*. New York: Cambridge University Press.
- Warfield, T. A. (2000). Causal determinism and human freedom are incompatible: A new argument for incompatibilism. *Philosophical Perspectives*, 14, 167–180.

- Warfield, T. A. and Finch, A. (1998). The *Mind* argument and libertarianism. *Mind*, 107, 515–528.
- Wiggins, D. (1973). Towards a reasonable libertarianism. In T. Honderich (Ed.), *Essays on freedom of action* (pp. 31–62). Boston: Routledge and Kegan Paul.
- Wright, L. (2001). *Critical thinking: An introduction to analytical reading and reasoning*. New York: Oxford University Press.