

4 The Debate over the Consequence Argument

We turn in this chapter to one of the three most significant influences on contemporary debates about free will, the Consequence Argument for incompatibilism. The Consequence Argument was first introduced by Carl Ginet (1966, cf. 1990), and then developed in different ways by David Wiggins (1973) and Peter van Inwagen (1975, 1983).¹ Since its first appearance, the Consequence Argument has been the single most influential consideration in favor of the thesis that free will understood in terms of leeway freedom is incompatible with determinism. Here is Peter van Inwagen's frequently quoted pithy statement of the argument:

If determinism is true, then our acts are the consequences of the laws of nature and events of the remote past. But it is not up to us what went on before we were born, 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. (1983: 16)

This elegantly cast expression of the argument masks much interesting complexity, complexity which proponents of the argument have addressed in admirable detail. We will turn to that complexity later in the chapter. But before doing so, it will be helpful first to reflect in some detail upon the dialectical status, antecedent to the introduction of the Consequence Argument, of the controversy between classical compatibilists and classical incompatibilists regarding the relation between determinism and leeway freedom.

This chapter will proceed as follows. In Section 4.1, we'll begin by assessing the dispute between classical compatibilists and classical incompatibilists regarding the ability to do otherwise as it stood in the early 1960s, prior to the introduction of the Consequence Argument. Then in Section 4.2 we will give a first pass at setting out a relatively accessible formulation of (a version of) the Consequence Argument. We'll follow that in Section 4.3 by canvassing the main strategies compatibilists have used to resist it. Then, in Section 4.4, we shall set out the Consequence Argument at an advanced level. In Section 4.5 we will critically examine the rule of inference at work in (one version of) the argument. And in Section 4.6 we'll consider some of the most interesting recent disputes regarding its soundness.

4.1. Reflecting on the Classical Controversy over the Ability to Do Otherwise

The classical compatibilists' proposed analysis of the ability to do otherwise can be stated roughly as follows:

An agent is able to do otherwise just in case, if she wanted to do otherwise, she would do otherwise.

As we explained in the previous chapter (Section 3.2), compatibilists offered this proposal as a way of reductively analyzing abilities in terms of counterfactual conditionals. Since the truth of these counterfactuals in general is not threatened by the truth of determinism, this compatibilist analysis was meant to offer an explicit demonstration that the ability to do otherwise is in principle consistent with the truth of determinism. It was thus a major blow to the classical compatibilists' program that their proposed analysis was so decisively refuted.²

Reflecting upon this history, it is tempting to overestimate the dialectical position the classical incompatibilists came to hold. In particular, it is tempting to think that, in the wake of the refutation of the classical compatibilists' proposed analysis (an analysis with its roots dating back at least to Hobbes and Hume), the default presumption was that the preponderance of reason strongly suggested that the ability to do otherwise, and so leeway freedom, was incompatible with determinism. But we think this would be a mistake. How so? To begin, it is worth noting that a refutation, even a decisive refutation, of one argument for the compatibility of determinism and the ability to do otherwise does not itself establish that the ability to do otherwise is *incompatible* with determinism. Discrediting the classical compatibilists' analysis fell far short of showing that they were wrong to contend that the ability to do otherwise is compatible with determinism. Their argument for that thesis might have failed, but this was not proof that their compatibilist thesis itself was false.

At this point, however, an impartial spectator might point out that in the absence of some plausible *positive* account of the compatibility of the ability to do otherwise and determinism, compatibilists were at a transparent disadvantage. Why? Because, on its face, intuition strongly suggests that if an agent's actions are determined, then she cannot do otherwise. If the past and the laws causally ensure what an agent does at a time, then this gives *prima facie* reason to think that at that time the agent is not able to do otherwise. If compatibilists had no special way of accounting for the ability to do otherwise (such as by way of their conditional analysis) to counter the apparent tension between determinism and the ability to do otherwise, then it would seem that the scales were after all tipped rather heavily in favor of the incompatibilists.

The preceding assessment of the dialectic overlooks other credible claims about determinism to which compatibilists were able to point, claims that would help make plausible the thesis that determinism is after all consistent with an agent's possession of the ability (and so the freedom) to do otherwise.

To understand the resources available to the compatibilist, consider what the thesis of determinism would come to as a comprehensive thesis about the natural world. Consider, for instance, a simple instance of a dispositional property, such as the solubility of a piece of solid, dry salt sitting undisturbed in a salt shaker. That substance, salt, is understood in part not just in terms of its chemical composition, and not just in terms of how it is when it remains undissolved. It is also understood in terms of how it *would* behave were it placed in water. In this sense, the salt's solubility is a *modal* property. It is a property about the potentiality or possibility of the item in question under certain conditions. Or consider, for instance, the behavior of a plant, such as a sunflower at night. We understand its nature, even while at rest in the dark, in part in terms of how we would expect it to behave in response to sunlight: it is disposed to follow the sun's arch through the sky.

Turning to the behavior of sophisticated animals, lions, for example, it is part of our understanding of them that they have active powers or dispositions of various sorts—to seek food, to eat, to protect their young, and to evade predators. What the compatibilist can have us note at this point is that these creatures are, in an uncontroversial sense, able to act in ways that at a certain time they might in fact not be acting. If determinism is compatible with a natural world in which ordinary objects have dispositions to behave in various ways under alternative conditions, and more complex items, including various animals, have powers of this general sort, it makes sense to suppose that determinism is compatible with an ability distinctive of persons, the ability to do things that one isn't currently doing (Moore, 1912). Suppose that you have the ability to speak both English and French. You're in a café in Paris, but you order your coffee and croissant in English. Suppose your friend says to you: "You could have ordered in French instead!". In this context it seems that your companion was correct to say what she did no matter what the truth about determinism turns out to be.

Compatibilists might concede at this point that they have no positive thesis about *how* this ability, the free will ability, is to be explained in a way that makes it clear that it is after all compatible with determinism. But their inability to account for this fact about our nature, by way of an analysis of this unique ability, does not provide decisive reason to think that the ability in question is in principle incompatible with the natural world being arranged in such a way that determinism is true. Unless the incompatibilists wished to argue that determinism is incompatible with the wide range of dispositions, powers, and abilities similar to those featured in the rest of the natural world, compatibilists had credible grounds for remaining committed to the thesis that the free will ability, understood as the ability to do otherwise, is also compatible with determinism.

Given the above remarks, it may be best to see the dispute over leeway freedom between compatibilists and incompatibilists as at a relatively evenly balanced impasse. Both sides were able to offer reasonable considerations for their opposing theses. And it seemed that there was little one was able to offer an impartial inquirer to help her adjudicate matters so as to show that one side

rather than the other had claim to a greater preponderance of reason on their side. So consider, in this evenly weighted dispute, the Basic Leeway Argument for Incompatibilism (BLI):

1. If a person acts of her own free will, then she could have done otherwise.
2. If determinism is true, no one can do otherwise than one actually does.
3. Therefore, if determinism is true, no one acts of her own free will.

We set out BLI in Section 3.3 as one of the two arguments capturing different core incompatibilist theses. Given the controversy currently under discussion, the dispute between compatibilists and incompatibilists amounts to a dispute over the second premise of BLI. And it seems that there was little available to move the debate along.

It is now easy for us to explain why the Consequence Argument for incompatibilism had such a powerful influence on the free will debate. Prior to its emergence in the dialectic, incompatibilists had little by way of argument for premise 2 of BLI. What the incompatibilists had to appeal to was a *prima facie*, intuitive judgment that determinism is incompatible with the ability to do otherwise. And while that is after all an initially compelling basis for favoring incompatibilism, even after the fall of the compatibilists' conditional analysis, compatibilists had available to them a plausible conception of natural abilities that underwrites the ability to do otherwise even in a deterministic context. But, as shall soon become clear, the Consequence Argument is in essence an argument—a compelling argument—for the second premise of BLI. As such, it provided a substantial source of support for incompatibilism. To this day it remains one of the most important influences on the free will debate. Many who are incompatibilists today base their commitment to it on the claim that some version of it is after all sound. We turn now to a first pass at setting out the Consequence Argument in a way that moves beyond van Inwagen's formulation of it as we quoted it at the beginning of this chapter.

4.2. A Formulation of the Consequence Argument

The version of the Consequence Argument we'll now consider, a modal version, invokes a compelling pattern of inference applied to modal propositions about what is *power necessary*.³ Power necessity can be understood in terms of what it is not within a person's power to alter. As applied to true propositions, power necessity concerns a person's powerlessness to affect their truth. To say that a person does not have power over a true proposition is to say that she cannot act in such a way that it would be false rather than true. To illustrate, no person has power over the truths of mathematics. That is, no person can act in such a way that the true propositions of mathematics would be false instead.⁴ Hence, the truths of mathematics are, for any person, power necessary. Intuitively, a valid pattern of inference, drawing upon propositions about what is power necessary, unfolds as follows:

If a person has no power over whether a certain fact obtains, and if she also has no power over whether this fact has some other fact as a consequence, then she also has no power over the consequent fact.

Powerlessness, it seems, transfers from one fact to consequences of it. Here is an example:

If poker-playing Diamond Jim, who is holding only two pairs, has no power over the fact that Calamity Sam draws a straight flush, and if a straight flush beats two pairs and we grant Jim also has no power over this, it follows that Jim has no power over the fact that Sam's straight flush beats Jim's two pairs.

This general pattern of inference is applied to the thesis of determinism to yield a powerful argument for incompatibilism. The argument requires the assumption that determinism is true, and that the facts of the past and the laws of nature are fixed. (Given these assumptions, here is a rough and simplified sketch of the argument:

1. No one has power over the facts of the remote past and the laws of nature.
2. No one has power over the fact that the remote past in conjunction with the laws of nature implies that there is only one unique future (that is, no one has power over the fact that determinism is true).
3. Therefore, no one has power over the facts of the future.

According to the Consequence Argument, if determinism is true, no person at any time has any power to alter how her own future will unfold. Assuming free will requires the ability to do otherwise (leeway freedom), then, in light of the Consequence Argument, free will is incompatible with determinism.

We'll now present a more precise version of the argument. Our explanation of the argument involves three steps: first, a presentation of the logical form of the propositions figuring in the argument; second, a demonstration of the inference principle the argument uses; and third, the application of the inference principle to a specific argument for incompatibilism.

First, consider the logical form of the propositions figuring in the argument. The logical form " $N_{s,t}(p)$ " expresses a proposition of the form: "It is power necessary for a person, S, at a time, t, that the proposition p is true," and we will treat this form of expression as synonymous with "p is true, and S is not free at t to act in such a way that, if S were to so act, p would not be true."

Second, the version of the Consequence Argument we'll consider exploits an inference principle (or rule) that is a modalized version of the simple argument form *modus ponens*:

1. p
2. $p \rightarrow q$
3. Therefore, q

Just as *modus ponens* allows us to infer q from the two propositions, p, and p materially implies q ($p \rightarrow q$), so the modalized version allows us to infer that q is power necessary for a person at a time from the two propositions that p is power necessary for a person at a time, and that it is power necessary for a person at a time that q is a consequence of p. The argument form can be represented as follows:

1. $N_{s,t}(p)$
2. $N_{s,t}(p \rightarrow q)$
3. Therefore, $N_{s,t}(q)$

Another common form of notation to represent the preceding expression of this inference rule is as follows:

$$N_{s,t}(p), N_{s,t}(p \rightarrow q) \vdash N_{s,t}(q)$$

In ordinary English, the pattern of inference reads:

1. p obtains, and S is not free at t to act in such a way that, if S were to so act, p would not obtain.
2. p implies q, and S is not free at t to act in such a way that, if S were to so act, p implies q would not obtain.
3. Therefore, q obtains, and S is not free at t to act in such a way that, if S were to so act, q would not obtain.

This modal inference rule is known as *Transfer of Powerlessness*, or just *Transfer*. Transfer is supported by applications of it that demonstrate its logical force. For instance, as applied to the example presented above, if, in the midst of his poker game, and holding only two pairs, it is power necessary for Diamond Jim that Calamity Sam draws a straight flush, and if it is power necessary for Diamond Jim that a straight flush beats two pairs, then it is also power necessary for Diamond Jim that Calamity Sam's straight flush beats Jim's two pairs.

Third, we can now apply the Transfer inference principle in the service of incompatibilism. To construct the incompatibilists' Consequence Argument, three premises are needed. One premise involves the principle of the *Fixity of the Past*. It states that no person can change facts of the past relative to her present situation. A second premise involves the principle of the *Fixity of the Laws*. It specifies that no person can alter a law of nature. The third premise invites one to assume that determinism is true—it states that the facts of the past and the laws of nature entail one unique future.

Here, then, is how the Consequence Argument unfolds. Take $(p \& t)$ to stand for the conjunction of two propositions: (p) the facts of the past obtain, and (t) the laws of nature obtain. Given the principles of the Fixity of the Past and the Fixity of the Laws, one can say of their conjunction $(p \& t)$ that they are power necessary for any person. That is, for any person, S, S is not free at any time, t,

to act in such a way that, if S were to so act, the propositions describing the past relative to t_1 or those expressing the laws of nature, would be false ($N_{S_1}(p \& I)$).⁵ This is the first premise of the argument.

Next, $((p \& I) \rightarrow f)$, is an expression of the thesis of determinism, that the propositions describing the past and the laws of nature ($p \& I$) materially imply that there is only one unique future (f). Like the Fixity of the Laws and the Past, the thesis of determinism, supposing it to be true, would also be power necessary for any person at any time.⁶ Hence, the second premise of the Consequence Argument is that for any person, S, S is not free at any time, t_1 to act in such a way that, if S were to so act, the thesis of determinism would not be true [$N_{S_1}((p \& I) \rightarrow f)$].

From these two premises and Transfer, it follows that for any person, S, S is not free at any time, t_1 to act in such a way that, if S were to so act, the actual future, f_1 would not obtain ($N_{S_1}(f)$). Set out formally, the argument can be represented as follows:

1. $N_{S_1}(p \& I)$
2. $N_{S_1}((p \& I) \rightarrow f)$
3. Therefore, $N_{S_1}(f)$

We introduced the modality of power necessity just above by explaining that it is best understood as indexed to persons and times. Strictly speaking, this is correct. But since in the Consequence Argument the claims are meant to apply to any (finite) being, the indexing can safely be dropped as an unproblematic way of simplifying the argument. Hence, we offer this simplification:

1. $N(p \& I)$
2. $N((p \& I) \rightarrow f)$
3. Therefore, $N(f)$

According to the Consequence Argument, given that no person can alter the past or the laws of nature, and assuming that determinism is true, and thus that each person's acts are consequences of the past and the laws, no person can alter the future from the one that is a consequence of the past and the laws given determinism's truth. No one is then able to act other than as she does—no one has leeway freedom.

The Consequence Argument's force is perhaps best captured as expressing a claim about what would have to be true if, at a determined world, a person is able to do otherwise. For that to be true—that is, for the conclusion to the Consequence Argument to be false—a person would have to be able to render false facts about the past or about laws of nature. But, the incompatibilist contends, it is incredible to think that whatever powers of agency a person has, they involve such abilities. Hence, according to the Consequence Argument, if determinism is true no one is free to act other than as she does.

4.3. Strategies for Resisting the Consequence Argument

The soundness of the Consequence Argument has been contested in various ways: here we focus on the three most influential challenges. A first calls into question the claim that a person is unable to act in such a way that the past would be different than it is. The second calls into question the supposition that a person is unable to act in such a way that the laws would be different than they are. The third objection contends that the inference principle on which the argument relies—the Transfer principle in the version we just set out—is invalid, and thus one cannot draw the desired incompatibilist-friendly conclusion even if the Consequence Argument's premises are all true. Each of these three compatibilist efforts has given rise to some of the most sophisticated contemporary work on the free will problem.⁷

4.3.1. Challenging the Fixity of the Past

As we set it out in the previous section, the Consequence Argument's first premise, $N(p \& I)$, relies on two principles, the Principle of the Fixity of the Past and the Principle of the Fixity of the Laws. Consider first the Principle of the Fixity of the Past, which states that a person cannot alter the past. How could a compatibilist plausibly deny this? It does seem incredible that we might be able to act in the present in such a way that the past would be different. But consider the difference between an agent who has the ability to act in such a way that *she alters the past*, as opposed to an agent who has the ability to act in such a way such that, *if she did so act, the past would have been different*—David Lewis (1981) makes this distinction. The former ability might well be thought outlandish. But the latter ability might be easier to accept.

Let us distinguish in general (and not just as applied to the topic of the Fixity of the Past) between two such notions of ability. The first is a stronger notion of ability, the second weaker. Call them *causal* (CA) and *broad* ability (BA).⁸ Causal ability holds that:

CA: A person has an ability to bring something about, p , just in case there is a course of action such that the person is able to perform such an action, and, if she were to perform it, then she would cause it to be the case that p obtains.

Broad ability holds that:

BA: A person has an ability to bring something about, p , just in case, there is a course of action such that the person is able to perform such an action, and, if she were to perform it, then p would obtain.

These two notions of ability, CA and BA, can be employed to understand how a compatibilist could resist either the Principle of the Fixity of the Past or instead

the Principle of the Fixity of the Laws. The compatibilist position currently under consideration presupposes the notion of BA as it bears on a free agent's relation to the past. When the Principle of the Fixity of the Past is interpreted so that what is at issue is BA, the compatibilist might argue that the principle is false and the ability at issue is not incredible but unproblematic.

To illustrate the difference between CA and BA, consider first what would be required for an agent to act differently (here the claim is not about an agent's *ability* to act differently from how she acted). For example, consider: *If McKenna were dancing on the French Riviera right now, he'd be a lot richer than he is.* Certainly this claim does not mean (at least not given McKenna's dancing skills) that if he goes to the French Riviera to dance, he will *thereby* be made richer. (No one in her right mind would pay to see McKenna dance!) It only means that were he to have gone there to tango, he'd have had to have had a lot more cash *beforehand* in order to finance his escapade.

Now consider an example that is about a claim of ability, one due to John Fischer (1994: 80–2): Each morning at 9:00 the salty old sea dog checks the weather forecast to learn what the weather will be that day. If the weather will be fair at noon, he sets sail. If not, he stays on land. Now one day, at noon, the seadog learned earlier that day, at 9:00, that the weather would be horrible. So he stayed on dry land. But, one might think (setting aside reflections on the free will debate altogether), he is *able* to sail at noon. Why wouldn't he be? There is nothing wrong with him. He is healthy, of sound mind, not hypnotized or deceived. He has what it takes, just then, to go sailing. The fact that he *does not* go sailing, Fischer observes, is not evidence that he *cannot* go sailing—that it is not within his power to do so. But what, we might safely reason, if he were, just then, to exercise his ability to sail and actually sail just then, at noon? Well, a reasonable answer is that, were he to sail just then, we can infer that the forecast at 9:00 a.m. would have been for fair weather.

With Fischer's example in mind, let us carry this over to the Consequence Argument, and consider what resources a compatibilist has in light of these ordinary patterns of inference about actions and exercises of ability in relation to the past. The compatibilist wants to resist the Consequence Argument and to say that at the moment when a free agent acted, even if she was causally determined to act as she did, she just then retained the ability to act differently. The fact that she did not exercise that ability is not itself sufficient reason to conclude that she is unable to exercise it (just like the salty sea dog was able to sail even when he didn't). Had she exercised it, and done otherwise as a result, some feature of the past prior to her so acting would also have been different. Perhaps she would have, just antecedent to her acting, wanted something other than what she actually wanted, or come to believe something different from what she actually believed. But when she acts as she actually does, she nevertheless retains that ability to act differently.

According to this first objection, the compatibilist need not say that for a causally determined agent to act otherwise, by her acting otherwise she would *cause* the past of her actual world to be different—which involves ability in the

CA sense. All that the compatibilist needs to invoke is BA, the weaker notion of ability—and thus that if she were to act otherwise, the past would have been different.

4.3.2. Resisting the Fixity of the Laws

Now turn to the Principle of the Fixity of the Laws, which states that an agent cannot alter the laws of nature. The second challenge aims to resist this principle, and to argue that there is a sense in which we have the ability to act so that a law of nature is falsified. Maintaining that an agent could falsify a law of nature seems as counterintuitive as claiming that a person can alter the facts of the past. But by again invoking the distinction between broad and causal ability, compatibilists have aimed to make this challenge plausible.

One way to pursue this strategy is by way of advancing a specific account of the laws of nature. For instance, the compatibilist might argue for a Humean account that holds that a law of nature reduces to regularities among events.⁹ In worlds in which determinism is true, laws that reflect such regularities will be of a sort that they, along with the facts of the past, entail every truth about all later times. Notice that, on such a view, if the history of a world were different from how it is, then different regularities in nature might emerge. Then different laws of nature would be “the” laws of nature. The compatibilist might first point out that no human being has the causal ability to make false an otherwise true law of nature. She cannot *initiate* a variation in the laws of nature by performing an action so that at the time she acts and afterwards, by her so acting, the laws and thus the regularities they reflect would be different from what they are. But an agent might still have the broad ability to act differently from how she does act, so if she were to act in this alternative way, the laws of nature that do in fact obtain would not. Some other regularities would unfold in that alternative causal history. *But she would not be the cause of this change*; her so acting would be an *upshot* of the antecedent fact that the laws were ever so slightly different.

Suppose Captain Ahab chooses in this actual world to eat fish for dinner, but just before choosing to eat the fish, he deliberated about whether to eat steak instead. (Imagine that both were offered to him and he was permitted to choose only one.) Suppose also that he was causally determined to choose to eat fish and not steak. In the actual world, as things unfold, there are regularities involving events that occur prior to his deliberating and his choosing to eat fish. On the Humean account, these regularities will count as laws. However, were Captain Ahab to have chosen to eat steak, the regularities and thus the laws would have been different. In particular, there would be different regularities and thus laws involving events occurring just prior to his deliberating and his choosing to eat steak. Now, in the actual world, as things really did unfold, Captain Ahab does not have the ability to *cause* an actually obtaining law of nature not to obtain. But he does have the broad ability to choose to eat steak and not fish (he's done if in the past), and, were he to have chosen to eat steak, the laws of nature that do obtain in the actual world would not obtain. In particular, a different law

would have governed what happened just prior to his choosing to eat steak. Hence, since he has the ability to choose to eat steak as well in the world in which he does in fact choose to eat fish, Captain Ahab thereby has the broad ability to act in such a way that a law of nature that obtains in the actual world would not obtain.

This type of compatibilist response to the Consequence Argument, first advanced by David Lewis, has come to be known as *local miracle compatibilism* (Lewis, 1981, 1979). Local miracle compatibilism (LMC) does not claim that a free agent at a deterministic world, W1, has the causal ability to break a law of nature, that is, to *cause* a miracle. Local miracle compatibilism only claims that at W1 she has the broad ability to act in a manner such that an actual law of nature would not be a law of nature. If in some other possible world, W2, she acts differently, then *relative to the laws in W1*, a local miracle would occur at W2. That is, a free agent at a deterministic world has the ability to act differently than she does act, because, if she were to act differently, then the laws of nature, just prior to her acting, would be slightly different from the way they are. It's important to emphasize that the alternative course of action in the world W2 would require a miracle *only relative to the laws in the actual world (W1)*. Were she to act on the alternative in W2, some other set of laws would hold in W2. *These she would not violate*.¹⁰ LMC is provocative, but not clearly implausible, and has rightly earned the respect of many serious-minded philosophers.¹¹

4.3.3. Challenging the Transfer Principle

The preceding attempts to prove the Consequence Argument unsound concern different ways to demonstrate the falsity of the first premise of the argument (i.e., no one has any power to alter the past and the laws of nature). A third approach is to set aside debates about the Fixity of the Past and the Laws and instead attempt to prove that the conclusion does not validly follow from the premises. That is, this approach seeks to undermine the inference principle which allows us to infer an inability to act otherwise from an inability to alter the past and the laws.

Michael Stote (1982) proposes a challenge of this kind. His strategy is to show that the Transfer principle

$$N_{S_1}(p), N_{S_1}(p \rightarrow q) \vdash N_{S_1}(q)$$

fails because power necessity is *selective* in the inferences that it licenses. A modality is selective if the inferences it licenses are restricted to certain cases or domains and cannot be generalized as required by valid inference principles that are used to drive valid argument forms.

Consider, by analogy, the case of knowledge. Suppose that René knows that Gassendi is a bachelor, and he knows that, if Gassendi is a bachelor, then Gassendi is unmarried. Given a pattern of inference similar to Transfer, it seems right to conclude that René knows that Gassendi is unmarried. A principle that

might be extrapolated from cases like René's knowledge of Gassendi has been called the Principle of the Closure of Knowledge under Known Implication ("Closure" for short). In a manner structurally similar to Transfer, Closure might be represented as follows:

1. Kp
2. K(p → q)
3. Therefore, Kq

But as plausible as the inference regarding René's knowledge may seem, some philosophers, and Stote is one of them, have contended that inferences of this form are not generally permissible, and thus that Closure is invalid. In particular, they have argued that the following sort of instance is not valid:

1. René knows that he is sitting in a chair.
2. René knows that if he is sitting in a chair, then he is not having a massive hallucination.
3. Therefore, René knows that he is not having a massive hallucination.

Epistemological skeptics invoke a version (the *modus tollens* version) of this pattern of inference (that is, $\sim Kq, K(p \rightarrow q) \vdash \sim Kp$), to reason as follows:

1. René does not know that he is not having a massive hallucination.
2. René knows that if he is sitting in his chair, then he is not having a massive hallucination.
3. Therefore, René does not know that he is sitting in his chair.

But, if this inference is invalid, then it would not follow from René's failure to know that he is not having a massive hallucination that he does not know that he is sitting in his chair. Consequently, while he would not know that he is not having a massive hallucination, this lack of knowledge would not impugn the epistemic status of his external-world belief about his sitting in a chair. The trick, however, is to provide a convincing account of why it is that the sort of inference featured in the skeptic's argument is invalid. (This is not something that we will pursue here.)

In like fashion, applied to Transfer, while sometimes powerlessness over one fact as well as its implications result in powerlessness over an implicated fact, must it always? Or is Transfer selective in the way that Stote thinks Closure is?¹² His key point is that notions like *unavoidability*, or, as we have been discussing, *power necessity*, are sensitive to contexts in a way that only selectively permits the sort of inference we find in the Consequence Argument. Let us work with the idea of *unavoidability*, since that is the notion Stote considers. On his proposal, when we say that something is unavoidable for a person, we have in mind selective contexts in which the facts pertaining to the unavoidability are independent of or bypass that person's agency (Stote, 1982: 19). It is unavoidable for

Barack Obama, for instance, that Caesar crossed the Rubicon. Nothing about his agency—about what he can do—can alter such facts. But when discussing the range of acts Obama is able to perform but does not perform, when he is of sound mind, healthy, uncoerced, and not deceived, these are, for him, in a plain way within his power and so are avoidable. The suggestion is that Transfer-style inferences do not work when they concern aspects of a person's own agency, or at least whenever there is not some special defect or impairment of agency that is in question, such as when an agent is under the grip of an addictive desire.

Notice that in the inference invoked in the Consequence Argument unavoidability or power necessity is specified as transferring from a context in which the notion is, as Slotte would have it, appropriately applied, and one in which, in his view, it is not. To illustrate his contention, suppose it is asserted that when Obama freely elects to remain at his desk and keep working, he is free just then to do otherwise and instead take a stroll in the White House gardens with his daughters (and assume that determinism is true). Then, as Slotte might see it, while the following propositions are indeed true:

1. It is unavoidable for Barack Obama that the facts of the past and laws of nature (p&l) are thus and so; and
2. It is unavoidable for Barack Obama that p&l implies that he now remain at his desk and work rather than stroll with his daughters,

the following proposition is false:

3. It is unavoidable for Barack Obama that he remains at his desk and works rather than take a stroll with his daughters.

In the Consequence Argument, the first premise cites considerations that have nothing to do with a person's agency—facts prior to his birth, and the laws of nature, and the second premise cites a fact about what the past and the laws imply, namely, that they imply that he perform the action at issue. It is claimed that these facts are unavoidable for the agent, but from this a conclusion is drawn, relying upon a Transfer-like principle, that the action at issue is unavoidable for him. This, Slotte and other compatibilists (such as Dennett, 1984; Mele, 1995, 2006b) have suggested, is to draw illicitly incompatibilist conclusions about unavoidability from reasonable claims regarding unavoidability.

4.3.4. A Final Challenge

The standard formulation of the core concern exploited by the Consequence Argument as a conflict between an action's being freely willed and the truth of causal determinism has in recent years been challenged by Joseph Campbell (2007, 2008, 2010). Campbell's objection is that this argument relies on the contingent assumption that the agent has a remote past—a past before she existed—and hence does not show that free action is incompatible with

determinism per se. The argument tells us nothing about deterministic worlds in which an agent has no remote past.

We want to note two responses to this objection. Alicia Finch (2013) defends the Consequence Argument against Campbell by arguing that the dynamic of the argument can be restricted to a time within an agent's lifespan, thereby precluding the need for a remote past. Key to her defense is plausibility of the *temporality thesis*, according to which an agent's performing a free action requires that at an earlier time it was up to the agent to perform the action at the later time. Finch's strategy is to argue that in the target cases, the laws and the state of the world at the earlier time are not up to the agent, and if causal determinism is true, this quality of not being up to the agent transfers to the action at the later time.

Carolina Sartorio (2015) responds to Campbell by defending the view that the true incompatibilist concern is not the incompatibility of free will and causal determinism per se, but rather free will and actions being causally determined by factors beyond the agent's control, and Campbell's objection does not address this thesis. We think that Sartorio is right about this. Campbell's objection and these two responses, we believe, are important and insightful.

4.4. The Consequence Argument: A More Precise Formulation

We turn now to a more advanced formulation of the Consequence Argument, the modal version developed by van Inwagen (1983: 93–5).¹³ Our goal in this section is to explain this more precise formulation and comment on details that will prove useful for understanding some of the technical debates about it. In Section 4.5 we will consider one major challenge to the argument so formulated. Introductory students might wish to pass over this section and the next. We include them since some interested in pursuing these topics might profit from a more thorough, technical presentation. But in our view, the most significant objections to the Consequence Argument have already been discussed: they are the challenges to the principles of the Fixity of the Past and the Fixity of the Laws. In Section 4.6 we offer our final assessments of these challenges.

As van Inwagen sets it out, the argument invokes two inference rules (94). The first rule he labels Rule Alpha, α . It specifies that from the fact that p is necessary one can validly infer that no one has, or ever had, a choice about whether p is true—that is, that p is power necessary.¹⁴

α : $\Box p \vdash Np$

The second rule is similar to the one we introduced above, which we called Transfer. Van Inwagen labels it Rule Beta, β , and it can be represented as follows:

β : $Np, N(p \rightarrow q) \vdash Nq$