

CHAPTER ONE

TWO PROBLEMS FOR CLASSICAL INCOMPATIBILISM*

JOE CAMPBELL

1. Introduction

My main concern in the area of free will research lies with the compatibility problem: Is the *free will thesis*—the claim that some persons have free will—consistent with the thesis of determinism? I'm not interested in attempting to solve traditional problems: Do we have free will? Is determinism true? (van Inwagen 1983) I lean towards epistemological skepticism about metaphysical claims (persons have free will, determinism is true, etc.), so I don't have high hopes for solving traditional problems. My ultimate goal is more modest, for example, to show that there is no reason to doubt that we have free will, and no good argument for *free will skepticism*: the view that no one has free will. Incompatibilism is an essential claim in arguments for free will skepticism. My fascination with arguments for incompatibilism, like the consequence argument, has to do with my fascination with free will skepticism.¹

Ability-to-do-otherwise claims are propositions about actions and *merely possible actions*, actions that one is able to perform yet never in fact performs.² *Classical views* of free will hold that *S* freely does *A* only if *S* is or was able to do otherwise. Analyses of ability-to-do-otherwise claims come in two flavors: *compatibilist accounts* and *incompatibilist accounts*.³ *Classical incompatibilism* is the view that if determinism is true, no one is or ever was able to do otherwise. It is endorsed by incompatibilists of all sorts as well as some *semicompatibilists*, who hold that while determinism might compromise the ability to do otherwise, it does not conflict with moral responsibility (Fischer 1994).

This paper addresses two problems for classical incompatibilism. First, it is shown that some of the best arguments for classical incompatibilism do not really prove the thesis (§ 2; see also Campbell 2007). One recent response to this criticism—weak incompatibilism—is considered and rebutted. After some preliminary comments (§ 3), I then discuss a second problem for classical incompatibilism: the sufficiency problem (§ 4). The *sufficiency problem* is the problem of identifying sufficient conditions for ability-to-do-otherwise claims. I show that incompatibilist accounts suffer from the sufficiency problem. There is a very clear necessary condition for classical incompatibilism: *S* is able to do otherwise only if that *S* does otherwise is consistent with propositions about the past and the laws of nature. Yet there is no clear set of sufficient conditions for the ability to do otherwise, ones that are also clearly incompatible with determinism. Classical compatibilism is often rejected because it fails to provide an adequate analysis of ability-to-do-otherwise claims but classical incompatibilism is no better in this regard.

What is the connection between the two parts of the paper? On the one hand, the best arguments for classical incompatibilism do not establish the thesis. Thus, reasons for endorsing classical incompatibilism have been undercut. It is true that my criticism does not undercut all arguments for classical incompatibilism but until the dust clears and we have time to assess the new arguments and the full impact of my criticism, one should be cautious about saying that there are compelling reasons to adopt classical incompatibilism. On the other hand, the best criticisms of the alternative view—classical compatibilism—focuses on the sufficiency problem. Yet it turns out that classical incompatibilism suffers from the same problem. Putting these two points together is by no means a refutation of classical incompatibilism but it is a nice set of criticisms nonetheless. On balance, when you add up the plusses and the minuses, classical compatibilism is a better theory than classical incompatibilism. I don't establish the last point in this paper but these are two further steps toward that ultimate goal.

2. The Consequence Argument

Classical free will theorists hold that the ability to do otherwise is necessary for free will. Classical incompatibilists further contend that no ability-to-do-otherwise claim is consistent with determinism. One very important argument for classical incompatibilism is the *third argument* (van Inwagen 1983, pp. 93-104; 1989, pp. 404-5; 2000). The third argument uses a modal operator: “ Np ” stands for “ p and no one is or ever

was able to act so as to ensure that $\sim p$ " (van Inwagen 1983, pp. 67-8 and fn. 31, pp. 233-4). It turns out that the original version of the consequence argument is invalid. Specifically, this transfer principle:

(Beta) From Np , $N(p \supset q)$ deduce Nq

is invalid (McKay and Johnson 1996). Here is a short version of the third argument (Widerker 1987) that avoids the interesting but complex issues surrounding principle Beta (for more on this topic, see van Inwagen 2000).

1. $N(P_0 \ \& \ L)$
2. $(P_0 \ \& \ L)$ entails P
3. Therefore, NP

P_0 is a proposition about the *remote* past—that is, a time t_0 before any humans existed. L is the conjunction of the laws of nature. P is any true proposition. There is one transfer principle:

(β') From Np , p entails q deduce Nq .

The second premise is true given determinism, so the argument appears to show that if determinism is true, then there are no true ability-to-do-otherwise claims. The third argument makes a strong case for classical incompatibilism.

There are issues about the burden of proof in the consequence argument. My view is the burden of proof lies with the classical incompatibilist, for he is the proponent of the argument. The main question when evaluating the argument should be: Does the consequence argument provide a compelling reason for endorsing classical incompatibilism? As a critic, it is not my job to prove classical compatibilism, nor is it my job to convince classical incompatibilists that classical compatibilism is true. I just need to show that the consequence argument does not prove classical incompatibilism; that it does not provide a compelling reason for endorsing classical incompatibilism.

In criticizing the third argument, I do not reject any of the premises or underlying assumptions. Rather, I claim simply that the third argument is not an argument for incompatibilism, for P_0 is not true in all possible worlds. Given the content of Np , classical incompatibilism is not the conclusion of the third argument. This criticism holds whether or not the third argument is sound and it applies to other versions of the consequence argument than the third argument as well as other arguments for classical

incompatibilism that depend on assumptions about the necessity of the past (Bailey 2012).⁴

In a series of papers published in *Analysis* (see especially Campbell 2007 and 2010), I offered two different models to help make my point that the third argument and other arguments fail as arguments for classical incompatibilism. Both models fall under the following description.

Adam*: Suppose W is a determined world such that some adult person exists at every instant of W 's existence. Thus, W has no remote past, no time prior to the existence of human beings. Suppose that at every moment of W 's existence lived Adam, an adult person with all the knowledge, powers, and abilities necessary for moral responsibility. One day Eve is created, and the rest is history. For each of the propositions that comprise W , someone is such that he or she has, or had, a choice about whether that proposition is true.

In the original example *Adam* (2007), the world comes into existence *ex nihilo* and Adam exists from the moment of creation. In the example *Oscillating Adam* (2010), Adam exists in an eternal series of temporal loops. Adam* is ambiguous between these two scenarios.

In reply to my criticism, many have advocated *weak incompatibilism*. According to this view the consequence argument establishes the incompatibility of the free will thesis, the thesis of determinism, and the existence of a remote past. There are several responses to this reply worth noting. First, weak incompatibilism is not a kind of classical incompatibilism. Weak incompatibilists concede that classical incompatibilism is false. The need for the additional contingent assumption of a remote past is proof enough that classical incompatibilism has been abandoned. True, maybe not the spirit of the thesis has been abandoned, but the thesis itself.⁵

What is the relationship between the consequence argument and the third argument, and between the consequence argument and other arguments for classical incompatibilism? Provisionally, formal versions of the consequence argument fall into two types. *Transfer versions* of the consequence argument explicitly endorse a transfer principle—like principle (β')—that transfers the necessity of one set of propositions onto some other set of propositions. *Non-transfer versions* of the consequence argument—like van Inwagen's *second argument* (van Inwagen 1983, 78–93) others (Ginet 1990; Fischer 1994)—do not explicitly endorse such transfer principles. The fundamental question when discussing the relationship between transfer and non-transfer versions of the consequence argument is, do non-transfer versions logically depend upon transfer

principles, even if those principles are not explicitly represented in the argument? According to van Inwagen, “all (logically adequate) arguments for incompatibilism must make some sort of implicit or hidden or covert appeal to Beta” (1994, p. 95). Whether or not something like this thesis is correct is beyond the scope of my paper but that van Inwagen is in support of the thesis should give us pause for thought.

There are other arguments for classical incompatibilism that are not versions of the consequence argument. For a great many of these, the criticism noted above applies as well (Bailey 2012). My criticism applies whenever there are explicit assumptions about the necessity of the past. There are also arguments for classical incompatibilism that are designed to get around criticisms like mine (Warfield 1996; Bailey 2012), though it is debatable whether or not these are versions of the consequence argument. They share some formal similarities with non-transfer versions of the consequence argument. These are complex and contentious topics better suited for other projects.

Furthermore, I don’t want to overstate the conclusion of this section. The conclusion is not that there are no sound arguments for classical incompatibilism. It would be difficult to prove such a claim since it is always open for someone to create a new argument. Rather, the conclusion of this section is more modest: The ability to do otherwise is not obviously incompatible with determinism. This doesn’t establish classical compatibilism but at least it gets the foot in the door.

3. “Can” and “Able”

Many philosophers distinguish between *specific* and *general abilities* (Maier 2011; cf. Honoré 1964; Mele 2003b).⁶ Consider this example. Elton John is on a train and there is no piano to play. Is Elton able to play the piano? In some general sense Elton is able to play the piano. He remains a pianist, for instance. Others presume that, since there is no piano to play, there is also a specific sense in which Elton is unable to play the piano. Given this example, we might say that Elton has the *general ability* (or *capacity*) to play the piano but lacks the *specific ability* to play the piano, for there is no piano to play.

Examples of general abilities include Elton John’s ability to play the piano, my ability to ride a bike, and what R. Jay Wallace calls the “powers of reflective self-control” (Wallace 1994, p. 2; see also Vihvelin 2004). Specific abilities are alleged to be abilities that we have at specific times, as opposed to general abilities, which we might attribute to agents without particular occasions for action in mind. This cannot be the difference,

however, for one can have a general ability at a specific time, as well. At one time in my life, I lacked the general ability to ride a bike, though at the present moment I have that ability. All abilities may be indexed to times, so if there is a difference between specific and general abilities, it isn't that the former are temporally indexed whereas the latter are not.

The distinction between specific and general abilities is difficult to specify. Outside of philosophical circles, the presence or absence of a bike is rarely considered when assessing one's ability to ride a bike. It only affects whether or not the ability happens to be manifested. As long as we agree that abilities are the kinds of properties that one can have even though they are never manifested, it shouldn't be a problem to think of all abilities as general abilities. I will not ride a bike now since there is no bike to ride. But I am able to ride a bike, unlike my nephew Logan. A friend wants to go skiing and asks me, "Are you able to ski?" It would be odd to answer "No" just because I've taken my skis to the shop to be waxed. Elton John doesn't lose his ability to play the piano when he boards a train—he does not, upon boarding a train with no piano, cease to be a pianist.⁷

Explications of specific abilities are often vague, though it is usually taken for granted that specific abilities are the kinds of abilities that matter in debates about free will and moral responsibility. Frequently one starts out with an example motivating the distinction between specific and general abilities, like the piano example. Clear explications of specific abilities, though, are hard to come by. For instance, after giving an example about the difference between a man's specific and general ability to serve a tennis ball, John Maier writes: "There is, as it were, nothing standing between him and a serve: every prerequisite for his serving has been met. Such an agent is *in a position to serve*, or has serving as an *option*. Let us say that such an agent has the *specific ability to serve*" (2010, § 1.3). Likewise, Al Mele claims that a general ability "is the kind of ability to *A* that we attribute to agents even though we know that they have no opportunity to *A* at the time of attribution and we have no specific occasion for their *A*-ing in mind" (Mele 2003b, 447). A characterization of specific ability is never provided by Mele but presumably a specific ability to *A* would require the opportunity and occasion to *A* as well as the general ability to *A*. Note that in both of these cases, specific abilities are understood in terms of general abilities. Specific abilities are general abilities *plus* something else, like opportunity or occasion.

Following a suggestion by Randy Clarke, we might say that the specific ability to *A* requires "the opportunity to exercise one's general capacity to *A*" (Clarke 2009, p. 338).⁸ According to Clarke, the ability to

do otherwise—or to believe or desire otherwise—requires more than a general ability, or what Clarke calls a “capacity.” It requires “that one can exercise that capacity” (p. 338). Clarke claims that it is important to distinguish between “having an ability” and “being able to exercise that ability” (p. 338). This suggests that a specific ability is a general ability that one is able to exercise. However, this way of understanding specific abilities is problematic since it includes a cognate of the very term we are trying to understand as part of the analysis. In what sense should I be able to exercise the general ability: in the specific sense or in the general sense? Further, even in the case of general abilities there must be some sense in which it is possible for me to exercise the ability; if there is no sense in which I am able to exercise my general ability to ride a bike, I lack that general ability. It must be possible to exercise general abilities just as well as specific abilities, so we are still left without a difference to ground this distinction.

I’m tempted to say that all abilities are general abilities. What I mean by this is not that specific abilities are uncharacterizable. There are in fact many analyses, examples, and distinctions—such as Kadri Vihvelin’s distinction between *wide* and *narrow* (2013)—that are viable. Rather, my worry is that even when specific abilities are characterized, they eventually lead to the sufficiency problem, or perhaps worse to classical free will skepticism or some other fatalistic outcome.⁹ In as much as specific abilities pick out a modality it appears to be the modality of actuality. We have no problem recognizing our abilities when we are exercising them but examples of unexercised abilities—the ability to do otherwise, for instance—remain contentious. Even when we think we have identified a specific ability, such as my ability to ride a bike given that a bike is available to ride, threats of classical free will skepticism easily surface. There is a bike available to ride but the fact is I don’t want to ride that bike right now; I’d rather work on this paper. If my wants and desires (or their underlying physical states) are considered salient parts of the circumstances, then am I ever able to ride a bike that I don’t ride in anything other than a general sense?

From the above critique of the specific/general ability distinction we cannot derive a more general criticism of classical incompatibilism. Following J.L. Austin (1956), we may distinguish between “can” and “able.” Suppose that the abilities that are relevant to free will and moral responsibility are general abilities. One might still think that while determinism doesn’t impinge on our general abilities, determinism imposes limitations on what one can or cannot do. It might be that “to have freedom of the will over some action ϕ , one must not only have the

ability to do other than ϕ , but also the opportunity to do otherwise” (Franklin 2012, p. 695). Perhaps “can” tracks free will yet “can” is analyzed in terms of general ability plus opportunity. This is Christopher Franklin’s proposal (2012; see also Vihvelin 2000).

Franklin never gives an analysis of “opportunity” but what he says about it is revealing. He takes his lead from Keith Lehrer, who holds that one can do otherwise provided that there is an accessible world W at which one does otherwise and one holds no advantages at W (1976). Lehrer is a compatibilist yet Franklin’s understanding of “advantage” leads him to incompatibilism.¹⁰ Franklin writes: “any difference in the past that is required in order for the agent to perform the action in question constitutes an inadmissible advantage” (2012, p. 702). Following Lehrer, Franklin maintains that differences in the laws of nature constitute an advantage, as well (p. 702). Yet if what one can do is restricted to possible worlds that share the past and laws of the actual world, then given determinism, no one can ever do nor ever could have done otherwise. Franklin provides an argument for classical incompatibilism that endorses a general view of ability.

Clarke critiques the view that being able to act is a matter of having a general ability or capacity, or a set of general abilities or capacities. He writes: “There are several different things that we might be thinking or talking about when we think or say that someone can or is able to do a certain thing” (2009, p. 339). For instance, “ S is able to do A ” might mean “ S has a general capacity to A , and the circumstances are friendly to S ’s exercising that capacity” (p. 336). The comments that Clarke makes about what it is for circumstances to be “friendly” are similar to comments that Franklin makes about “advantages.” For instance, Clarke writes: “whether the circumstances are friendly to the exercise of a capacity can depend, for example, on whether finks or potential masks are present” (p. 339).¹¹ Similarly, when discussing Frankfurt cases, Clarke concedes that an agent might have the general ability or capacity to do otherwise but “the circumstances are not friendly to his exercising that capacity” (p. 340). Perhaps it is compatibilist paranoia but I can’t help but think that so long as propositions about the circumstances, together with the laws of nature, entail that an action will not be performed, then the circumstances will not be friendly, according to Clarke. This strikes me as question begging but that is not the point I want to make.

First pause and consider the spectrum of theories emerging from the distinctions so far. One might think that the relevant sense of “able” is general, or one might think that it is specific; similarly with respect to “can.” One might also think that “can” tracks ability, or that “can” tracks

ability plus opportunity. One might think that the problem of free will and determinism is a problem about ability, or one might think that it is a problem about ability plus opportunity. In attempting to provide an overarching argument against classical incompatibilism, there are a lot of views to consider!

Yet there are important parallels between Franklin's approach and incompatibilist accounts of specific ability that are worth noting. While Clarke's comments about specific abilities suggest an incompatibilist analysis of "able"—one that builds propositions about the circumstances surrounding the action into the account of ability—Franklin is willing to toss the compatibilist an "able"-bone and dig his heels in with an analysis of "can": general ability plus opportunity, where our opportunities may be restricted by propositions about the past together with the laws of nature. From each of these offerings follow incompatibilist accounts where all of the circumstances leading up to an action matter when determining what one can or is able to do.

With so many views to consider, how should a classical compatibilist respond? My approach is to concede (for the purposes of this paper) that compatibilist accounts suffer from the sufficiency problem. Can the classical incompatibilist do any better? For if incompatibilist analyses of "can" and "able" are no better than compatibilist analyses, then incompatibilist accounts of ability have no advantage over compatibilist accounts in this regard. Perhaps *no* analyses of "can" and "able" are forthcoming. That alone cannot be a reason for rejecting the classical free will thesis. We might think, for instance, that basic actions are primitive and unanalyzable but it doesn't follow that there are no basic actions.

4. The Sufficiency Problem

The sufficiency problem is the problem of providing a set of sufficient conditions in an analysis of ability-to-do-otherwise claims. Most analyses are unproblematic when it comes to specifying necessary conditions for the ability to do otherwise. Any incompatibilist account, for instance, is going to hold that ability-to-do-otherwise claims must be consistent with propositions about the past together with the laws of nature. Call this the *incompatibilist condition*: "an agent who did not *A* at *t* was able to *A* at *t* only if in another possible world with the same past and laws, he *A*-s at *t*" (Mele 2003, p. 449). Likewise, even incompatibilists admit that some compatibilist analyses get it right when it comes to necessary conditions for ability-to-do-otherwise claims. Consider the *new dispositionalism*: the view that free will is a set of *dispositional powers*, like flammability,

solubility, and fragility (Vihvelin 2004; see also Clarke 2009; Maier 2010, § 5.2). That dispositional powers exist is consistent with determinism, so the new dispositionalism provides an interesting classical compatibilist account. More importantly, if specific abilities are analyzed in terms of general abilities, then even if specific abilities are the kinds of abilities that are properly identified with the ability to do otherwise, it follows that the ability to do otherwise requires having general abilities. Given this way of understanding things, most philosophers can admit that the new dispositionalism is accurate in articulating a set of necessary conditions for ability-to-do-otherwise claims.

However classical incompatibilism parts ways with the new dispositionalism when it comes to the specification of sufficient conditions for ability-to-do-otherwise claims. Classical incompatibilists contend that the new dispositionalism suffers from the sufficiency problem. In fact, classical incompatibilists extend this criticism to all classical compatibilist analyses of ability-to-do-otherwise claims. In this section, I show that incompatibilist accounts suffer from the sufficiency problem, as well. This cannot be a reason for rejecting compatibilist accounts, in general. For simplicity, our remaining focus is on analyses of “able” as opposed to “can,” though anything said about the former can be said about the latter. Similarly, since it is easier to discuss mere ability claims than it is to discuss ability-to-do-otherwise claims, and since the former can easily be transformed into the latter, analyses of ability claims are discussed below.

What might an adequate incompatibilist analysis of ability claims look like? One approach is to build on the incompatibilist necessary condition. It isn’t enough to merely replace the “only if” in this condition with an “if and only if”: An agent was able to *A* at *t* if and only if in another possible world with the same past and laws, he *A*-s at *t*. To see why this analysis is faulty, consider one of Clarke’s criticisms of the new dispositionalism. “Here is a thought with some intuitive appeal: to be responsible for *A*-ing, one must have possessed a capacity to *A*, and one must have determined, oneself, that one exercise that capacity on that occasion” (2009, p. 342).

Supposing that some sort of self-determination is essential to the kind of ability relevant to free will and moral responsibility, there is no guarantee that it exists merely because there is another possible world *W* with the same propositions about the past and laws of nature as the actual world in which the action is performed. It might be the case that *S*’s *A*-ing at *t* in *W* is a random event and thus not self-determined by *S*.

Thus, one approach that the classical incompatibilist might take is to build on an account of general abilities. Borrowing some suggestions from Clarke, there are many ways in which this might be done.

- S has a general capacity to A , and the circumstances are friendly to S 's exercising that capacity.
- S has a general capacity to A , and it is open to S (at some specified time) to exercise (at some specified time) that capacity.
- S has a general capacity to A , and it is up to S (at some specified time) whether she (at some specified time) exercises that capacity.
- S has a general capacity to A , and S has a choice (at some specified time) about whether she (at some specified time) exercises that capacity. (2009, 336-7)

To these we might add:

- S has a general capacity to A , and A is self-determined by S .

This kind of approach has several problems. First, there are several analyses here, not just one. Why choose one over the other?

Second, supposing that the additional qualifications—self-determination, for instance—are incompatible with determinism is contentious. Suppose, for instance, that the kind of self-determination required for free will and moral responsibility is *agent causation*. Compatibilist accounts of agent causation have been offered (Markosian 2002). This problem is perfectly general. Eventually, we are back to the topic of § 2. There is no reason to expect that incompatibilist analyses structured with Clarke's suggestions in mind will be less controversial than similarly structured compatibilist analyses. Any suggestion that they are rests on a supposed resolution of some version of the compatibility problem.

Third, many of Clarke's additional qualifications include modal concepts like *openness* and *choice* that are related to the concept of *ability*. Hence, such accounts run the risk of circularity. Forgetting about the first and third problems for a moment, the second problem might be avoided by pooling our resources and explicitly including something like the incompatibilist condition into the analysis. Before exploring this option, some preliminary comments are in order.

For several reasons, a *double time-index* account of ability claims is preferred. For instance, many of Clarke's suggestions contain two explicit temporal references and there is no reason to require that the referenced times be the same. Moreover, the incompatibilist condition specifies two temporal references though only one is explicit. For according to the incompatibilist condition, ability claims must be consistent with propositions about the past, and this carries with it an implicit temporal reference, as well. Lastly, we often lose abilities over time. Perhaps I am unable at 2:30pm to catch a flight at 3pm but I was able at 1pm to catch that same flight at 3pm (Lehrer and Taylor 1965). Thus, an incompatibilist

analysis of ability should distinguish between the *time of (possible) action* (3pm in our example) and the *time of ability* (2:30pm vs. 1pm).

In addition, two further concepts need to be introduced. Possible worlds *share the same past* at t if and only if they share the same propositions about the past relative to t . W_1 is an *extension* of W_2 at t if and only if W_1 and W_2 share the same past at t as well as the same laws of nature. Given this understanding of “extension,” every possible world—including the actual world—is an extension of itself. As was noted, Clarke’s suggestions offer a variety of distinct analyses, depending on whether we favor openness, choice, etc. as the salient feature that we should add to an account of general ability. For simplicity, I choose self-determination but my comments should apply to other analyses, as well.

Taking everything said above into account, what might an incompatibilist analysis of specific ability claims look like? Here is a provisional incompatibilist account.

The Incompatibilist Analysis of Specific Ability (IASA): S is able at t_1 to do A at t_2 ($t_1 \leq t_2$) if and only if there is a possible world W that is an extension of the actual world at t_1 where (a) S has the requisite general abilities and (b) A is self-determined by S at t_2 .

Note that given IASA, specific abilities require general abilities. IASA is an incompatibilist analysis, for given it and determinism, no one is or ever was able to do otherwise. If IASA and determinism are true, the only possible world that is an extension of the actual world is the actual world. This holds regardless of which particular times we consider to replace the temporal variables. But if indeterminism is true, depending on how it is manifested in the actual world, we might get radically different accounts of specific ability for any particular person and action, depending on how far back we place the time of ability t_1 . For instance, if we suppose that every event has a 50% chance of occurring and we consider the time of ability to be sometime right after S ’s birth, then S would appear to have a wide range of actions that he was or is able to perform. Yet if we set the time of ability at one millisecond prior to S ’s doing A and the indeterminism of the actual world is more constrained, then S ’s range of actions are similarly limited.

Even if we know that determinism is false, we are very far from ever knowing how indeterminism is manifested in the actual world. Thus, we are equally far from knowing which specific abilities we do or don’t have at any given moment, supposing IASA is true. Consequently, IASA entails a radical *epistemological skepticism about ability*.¹² It is not a useful

account of ability, not one that will help us to decide if some particular psychopath had the ability to control his behavior or even if you were able to brush your teeth this morning given that you brushed them.

Return now to a question raised above: Which time is the time to specify when specifying the time of ability in IASA? In versions of the consequence argument, it seems that propositions about the *remote past*—times prior to the existence of human beings—are what are important. But if our concern is with actions that are an extension of the actual world, then propositions about the past just a millisecond prior to the action should be equally relevant. For this reason, IASA does not appear to offer a single analysis of specific ability. No particular time of ability yields an analysis worth endorsing over the others.

Nonetheless, in understanding what we are specifically able to do, given IASA, which time of ability do we select? In other words, which true propositions about the past matter when considering the proper extensions of the actual world, given the incompatibilist assumptions of IASA? The answer has to be: include *all* the propositions about the past relative to the time of action, even those from just a millisecond prior to the action. After all, if the past is necessary all propositions about the past are on a par. It shouldn't be too hard to see that this eventually leads to classical free will skepticism, or something close. For, holding fixed the propositions about the past up to a millisecond prior to any decision, which decisions would we have made differently? We would need a radical form of indeterminism to avoid free will skepticism and with that would come an equally radical form of the problem of luck as well as the aforementioned epistemological skepticism. Even if we avoided classical free will skepticism, we would be led to *restrictivism*, the view that “rarely, if ever, is anyone able to do otherwise than he in fact does” (van Inwagen 1989, 405).

Perhaps, in an effort to save the analysis, we should set the time of ability further back, back far enough to avoid the threat of restrictivism. Yet what point in time might that be? If we set the time too far back, then the person won't be specifically able to do anything, for she might lack the general abilities that are necessary for specific ability. Consider the matter of assessing whether or not a particular killer was able to avoid committing a crime. Which point in his history would yield the correct judgment about his level of ability when assessing some particular crime? Depending on which time in his life we choose to fix as the time of ability, we get very different accounts of his range of future options relative to any action. The closer we get to the criminal act, the more the crime seems inevitable, something over which—given the circumstances—he had no

control. Yet as we go further back his range of options increase dramatically.

We might abandon the double time-index account and say that the time of ability is the same as the time of (possible) action. Franklin believes, for instance, that “libertarians must locate indeterminism at the moment of action” (2012, p. 704). Yet it is a short walk from this view to classical free will skepticism. As William Hasker puts it: “by the time something is, it is then too late for it not to be!” (1989, p. 10) There is no reason to suppose that we have more control over the present than over the past. At the moment of action, among the relevant propositions when considering whether (or not) the action is an extension of the actual world must be the claim that the agent did (or did not) perform the action.¹³

Accounts like IASA yield an unrealistic view of ability. How many decisions would we have made otherwise, given that the circumstances immediately prior to the decision are fixed? Certainly if I knew then what I know now, I’d have done many things differently but that is not the matter under consideration, for that is perfectly consistent with the thesis of determinism: same past, same result; different past, (possibly) different result. If in the exact same circumstances, given the exact same set of reasons and information, we would have made a different decision, the question arises: Why? Provided that we could even make sense of this proposal, the problem of luck would return with a vengeance (van Inwagen 1998; 2000), along with the problem of epistemological skepticism about ability.

On the other hand, when comparing the actions of a serial killer and those of average folks, general abilities are very helpful. Did the killer have “the power to grasp and apply moral reasons” or “the power to control [his] behavior in light of such reasons” (Wallace 1994, p.7)? Presumably you and I have those powers. This might provide a significant difference between us and them; one that is not saddled with the inevitable epistemic worries of the libertarian’s specific abilities. Given the level and extent of indeterminism, who knows what options anyone ever had? And who cares? What is significant is whether one has the range of general abilities that give one a chance to lead a fruitful and meaningful life.¹⁴

In considering incompatibilist accounts of specific ability, we’ve encountered several problems. First, if one starts with general abilities and adds additional necessary conditions, there is the issue of having too many different analyses. Which one do we choose and why? Second, even if we were to settle on the relevant necessary condition, whether or not the condition is consistent with determinism is a substantive question. Versions of the consequence argument were recently regarded by

incompatibilists as among the most powerful tools in the incompatibilist toolbox. Now we know that some versions of the argument are faulty, and thus not as good as they might have once looked. Are we really in a position to say that the ability to do otherwise is inconsistent with the claim that an act was self-determined, or some similar claim?

Third, even if we stipulate that an account of specific ability is incompatible with determinism, there is the problem of specifying the relevant temporal indices. For instance, the closer we make the time of ability to the time of (possible) action, the more likely it is that the problems of free will skepticism and restrictivism arise. Preventing this by either introducing a radical indeterminism or considering a longer separation between the time of ability and the time of (possible) action, only brings with it the problem of luck and the related problem of epistemological skepticism about ability.

5. Conclusion

We have examined two problems for classical compatibilism. On the one hand, some of the best arguments for classical incompatibilism have been undermined. Thus, our reason for accepting classical incompatibilism has been undermined. Of course, classical compatibilism is in trouble too for many compatibilist analyses of ability-to-do-otherwise claims suffer from the sufficiency problem. This brings us to the second problem for classical incompatibilism because, as we have seen, incompatibilist accounts suffer from the sufficiency problem, as well. Apparently, specifying the set of sufficient conditions required for the kind of ability relevant to free will and moral responsibility is a difficult matter whether or not one assumes compatibilism or determinism. The sufficiency problem is no reason for endorsing classical incompatibilism over classical compatibilism.

Why am I a compatibilist? If you follow the logic of classical incompatibilism eventually you are led to either classical free will skepticism or restrictivism. Thus, the only hope for saving the free will thesis is to endorse classical compatibilism. Further, if we are to make intuitive and helpful distinctions—explain the sense in which I am able to ride a bike and my nephew Logan is not, or explain the sense in which the average person is able to control his behavior yet a psychopath is not—we are going to have to be compatibilists and attend to general abilities. Ultimately, compatibilism is supported by pragmatism, naturalism, and common sense.¹⁵

Notes

¹ My understanding of “incompatibilism” differs from other analyses, like Kadri Vihvelin’s (2008; forthcoming). One difference is that, on my definition, free will skepticism entails incompatibilism. Compatibilism and incompatibilism both entail the free will thesis, according to Vihvelin. My interests in free will skepticism and the consequence argument compel me toward a broader view of incompatibilism.

² For simplicity, I use the term “ability” and its cognates rather than “practical ability.” Practical abilities are those abilities related to action (Mele 2003).

³ Accounts are broader than analyses. All analyses are accounts but not all accounts are analyses.

⁴ The claim is not that propositions about the past are logically or metaphysically necessary. The claim is that no one is or ever was able to render false a proposition about the past.

⁵ Here is a more general criticism of the consequence argument. Claims about the necessity of the past are ambiguous. On one reading, propositions about the past are not necessary at all, for they are just a record of what has happened. What is relevant is whether or not we could have done otherwise *in the past*, not whether we can change the past *now*. On another reading, propositions about the past are necessary in the sense that it is over and done with, out of our reach. However, this necessity does not transfer to claims about the future. Presumably the future is not over and done with, out of our reach.

⁶ Previously I’ve discussed a related distinction between *all-in abilities* and *general abilities* (Campbell 2005; see also Austin 1956).

⁷ I admit that ability-terms might be contextual or even more complex. Perhaps something like interest relative invariantism, or sensitive moderate invariantism, or subject sensitive invariantism is true (Rysiew 2011). Hence, the above discussion is likely an oversimplification. Given that contextualism and these more complex theories offer compatibilist accounts of abilities, these issues shouldn’t worry us (Campbell 2005).

⁸ This is Clarke’s characterization of *all-in ability*. Later he discusses other specific abilities.

⁹ Fatalism is the view that nothing is avoidable (Campbell 2010, 2011; see Markosian 2010). Fatalism entails classical free will skepticism, as we are using these terms.

¹⁰ It is not important for me to provide a further explication of what counts as an advantage (but see Campbell 1997 for my only attempt). What is important is what differentiates compatibilist accounts like Lehrer’s from incompatibilist accounts like Franklin’s.

¹¹ A *finkish disposition* is one that “would vanish immediately, on being put to the test” (Vihvelin 2004, p. 435). In the case of *masked dispositions*, an object is prevented from manifesting its dispositional power without losing the disposition. Unlike finkish dispositions, masked dispositions never vanish.

¹² Free will skepticism is a metaphysical view—the claim that free will does not exist—and not an epistemological skepticism (see Campbell 2011, Ch. 1). Free

will skepticism is (as far as I know) logically unrelated to epistemological skepticism about ability.

¹³ The suggestion is not that Hasker's point is valid. I think there are adequate responses. But even the libertarian eventually resorts to a kind of compatibilism.

¹⁴ In an episode of the TV show *Most Evil* entitled "The Science of Murder," it is claimed that the majority of serial killers suffer from three distinct maladies: brain damage, head injury (which might or might not have caused the brain damage), and extensive abuse and trauma. It is likely that at least some serial killers lack the kinds of general abilities necessary for free will and moral responsibility.

¹⁵ I thank Ryan Hebert, Graham Hubbs, Michael O'Rourke, Kadri Vihvelin, Ben Vilhauer, and Al White for written comments on previous drafts of this paper. Conversations with Vihvelin and White over the last few years have helped shape my views, as well. Earlier versions of this paper were read at the First Annual Graduate Philosophy Conference at the University of Calgary and the Northwest Philosophy Conference. I thank the audiences from those conferences for their contributions. Also, my understanding of the consequence argument and my criticisms of it benefitted from a symposium on my book (Campbell 2011), sponsored by the Freedom Center at the University of Arizona. I am especially indebted to my commentators: Michael McKenna, Carolina Sartorio, and Brandon Warmke. Lastly, I'm thankful for conversations about these topics on the Garden of Forking Paths (now Flickers of Freedom), especially from comments made by John Martin Fischer and Eddy Nahmias.