

July 31, 2009

How To Think about the Free Will/Determinism Problem

Commonsense says that we have free will. We make choices. And while we are sometimes in a position where we make a choice while being mistaken about what our options are, this is not always, or even usually, the case. Ordinarily, when we make a choice we really do have the choice we think we have. You can stop reading this paper right now, or you can read on. You really can do either of these things. It's up to you. You have a choice.

But, *for all we know*, determinism is true. Determinism is the thesis that the state of the entire universe at any time, together with the laws, is logically sufficient for the state of the universe at any later time.

Commonsense either doesn't know, or doesn't take seriously, the thought of determinism. But as soon as a philosopher explains the thesis of determinism, commonsense sees the problem: the truth of determinism means the absence of free will.

As William James long ago put it:

"Determinism professes that those parts of the universe already laid down absolutely appoint and decree what the other parts will be...necessity on the one hand and impossibility on the other are the sole categories of the real. Possibilities that fail to get realized, are, for determinism, pure illusions; they never were possibilities at all."¹

We can sum up commonsense thinking with the following simple argument.

The No Choice Argument for Incompatibilism:

1. I have free will only if I sometimes have a choice about what to do.

July 31, 2009

2. I have a choice about what to do only if there is more than one thing that I am able to do.
3. If determinism is true, then what I do is (always) the *only* thing that I am able to do. (I am never able to do otherwise.)
4. Therefore, if determinism is true, I never have a choice about what to do.
5. Therefore, if determinism is true, I have no free will.

I think that this way of thinking about the free will/determinism problem is the correct way.² There is something as plain as the nose on our face – the fact that we often have a choice about what to do -- and this apparently obvious fact *seems* to be in direct conflict with determinism. I see the free will/determinism problem as the problem of saying whether this apparent conflict is real.

This may seem obvious – I hope it does seem obvious. But this is not how the literature sees the problem.

1. Through the Lens of Moral Responsibility (How not to Think about Free Will)

In the contemporary literature, the free will/determinism problem is almost invariably viewed through the lens of moral responsibility. It is, of course, widely agreed that having free will is a *necessary* condition of being a morally responsible agent. But most contemporary discussions of the free will/determinism problem forge a much stronger link between questions about moral responsibility and questions about free will.

Let's begin with some particularly striking examples:

...the pessimists reply, all in a rush, that just punishment and moral condemnation implies moral guilt and guilt implies moral responsibility and moral responsibility implies freedom and freedom implies the

July 31, 2009

falsity of determinism...[the optimist replies that] people often decide to do things, really intend to do what they do, know just what they're doing in doing it...But it is here that the lacunae in the optimistic story can be made to show. For the pessimist may be supposed to ask: But why does freedom in this sense justify blame? ³

[P. F. Strawson]

...if people did not regard themselves and one another as responsible beings, life would be unrecognizably different from what it actually is. But the concept of responsibility is a mysterious one which tends, on examination, to become increasingly opaque and to threaten variously to be incoherent or impossible or universally inapplicable... We can express the problem of responsibility in the form of the question "How, if at all, is responsibility possible? And we can express the problem of free will in the form of the question: "What must our relation to our wills be," or better, perhaps, "What kinds of beings must *we* be if we are ever to be responsible for the results of our wills? ⁴

[Susan Wolf]

I believe that libertarians hold the views they do because they believe that any satisfactory conception of freedom of the will must allow us to justify our ascriptions of moral responsibility. They might be persuaded to accept a compatibilist account of freedom of the will if it could be shown that the fact that we are free, in that sense, could be used as the basis for a justification of the claim that we are morally responsible for our actions.⁵

July 31, 2009

[Hilary Bok]

Are we free agents? That depends on what you mean by 'free'. In this book the word 'free' will be used in what I call the ordinary, strong sense of the word. According to which to be a free agent is to be capable of being truly responsible for one's actions.⁶

As I understand it, true moral responsibility is responsibility of such a kind that, if we have it, then it *makes sense*, at least, to suppose that it could be just to punish some of us with (eternal) torment in hell and reward others with (eternal) bliss in heaven.⁷

[Galen Strawson]

These quoted passages, from philosophers whose views about free will and the free will/determinism problem are otherwise very different,⁸ have in common the assumption that if we are not morally responsible, then we are also never free agents; to be a free agent is, according to these philosophers, to have what it takes to be morally responsible for our actions (and for blame to be deserved, or just, or justified).

I think that this way of approaching the free will/determinism problem is a mistake.

If we think about free will in this way, we are imposing a heavy burden on a class of natural facts -- those facts, whatever they are, in virtue of which we have free will. We are saying that a class of natural facts constitute *free will* facts only if they play the role of being *the justifier* of praise, blame, and other practices associated with moral responsibility. Free will is assigned the burden of bridging the 'is-ought' gap, of explaining why moral responsibility is *moral* responsibility.

July 31, 2009

If we think about free will this way, we rule out, by stipulation, the possibility that so far as freedom is concerned, we have what it takes to be morally responsible, but we lack *something else* required for moral responsibility. For instance, insofar as we hold people responsible for what we believe were *their* past actions, it seems plausible that moral responsibility requires that we are numerically the same person through time. If that's right, then a philosophical argument for the conclusion that no one is ever numerically the same person over time is an argument, independent of free will, for the claim that no one is ever morally responsible.⁹ Another example: arguably our practice of holding each other responsible includes the belief that we have the *authority* to demand certain kinds of conduct from other people and to respond adversely -- with sanctions like public acts of blame -- when these demands are not met.¹⁰ If that's right, then a philosophical argument for the conclusion that we never have this kind of authority to make demands on other people is also an argument, independent of free will, for the claim that no one is ever morally responsible.

And if we think about free will this way, questions about what it is to have free will and to act freely (questions in the domain of the philosophy of action and metaphysics) take a back seat to questions about the nature of *just* or *justified* or *deserved* praise and blame (questions in the domain of moral philosophy). We start worrying about the nature of moral responsibility and whether anyone is ever *really* morally responsible, and moral responsibility becomes, as Wolf notes, a mysterious thing.¹¹ It seems urgent, then, to provide an analysis or account of moral responsibility or at least an account of what we do when we hold someone responsible. And in discussing these questions the question of free will either gets left behind or becomes problematic in all the ways that moral responsibility is problematic. And these problems have little or nothing to do with *determinism*.¹²

The problem of moral responsibility - what it is and what it takes for someone to be morally responsible -- is an interesting and important problem but it is a *different problem* than the free will/determinism problem.

July 31, 2009

Not all philosophers who view the free will/determinism problem (and, more generally, questions about free will) through the lens of moral responsibility see the connection in the strong way that Strawson, Wolf, Bok, and Strawson see it. But the following kind of view is very common:

Granted, it's a mistake to conflate free will and moral responsibility, and granted, it is intelligible that someone - a child, for instance - might have free will without having what it takes to be morally responsible for her actions. But we care about freedom only because we care about moral responsibility, and our intuitions about what it is to have free will are clear only insofar as they are linked to moral responsibility. So we should understand the free will/determinism debate as a debate about whether determinism would deprive us of the kind of freedom that is required for moral responsibility.¹³

In my conception, compatibilism is the view that determinism is compatible with whatever sort of freedom is sufficient for moral responsibility, while incompatibilism is the view that determinism is not compatible with this sort of freedom.

What is needed is an argument against the view that the freedoms that soft determinists have advocated are sufficient for moral responsibility, and I develop such an argument here.¹⁴

[Derk Pereboom]

As a theory-neutral term of departure, free will can be defined as the unique ability of persons to exercise control over their conduct in a manner necessary for moral responsibility.¹⁵

July 31, 2009

[Michael McKenna]

Let's reserve the technical term 'morally free' to describe actions that are free in the sense required for moral responsibility. Then we can characterize the Free Will Thesis as the claim that some actions are morally free, and compatibilism as the claim that the Free Will Thesis is compatible with determinism.¹⁶

[Ned Markosian]

This is an improvement insofar as it does not impose upon an account of free will (or free agency or 'moral freedom') the burden of being *the* justifier of blame and insofar as it does not preclude the possibility that something other than lack of free will (or free agency or moral freedom) prevents us from being morally responsible.

But I think that this way of thinking about the free will/determinism problem is, nevertheless, a mistake. Moral responsibility still gets all the attention. Since free will (or free agency or 'moral freedom') is functionally defined as *whatever it is* that plays the role of satisfying the freedom requirement of moral responsibility, we cannot rely on commonsense assumptions about what counts as free will or free agency. If we want to defend the claim that we have the kind of freedom that plays this moral role, or that this kind of freedom is compatible with determinism, we will need to provide an *analysis* or *account* of this "moral freedom". And the only way of testing the adequacy of our account will be by using our intuitions about when agents are praiseworthy, blameworthy, or in some other way morally responsible for what they do.

Derk Pereboom exploits this fact with his notorious "Four Case" argument against compatibilism.¹⁷ He claims that the compatibilist *must* provide an account, or at least sufficient condition, of what it is to satisfy the freedom component of moral responsibility, and he describes four cases and argues that at least one of these cases is a counterexample to each of the best compatibilist accounts, and a challenge to the

July 31, 2009

possibility of *any* compatibilist account.¹⁸ Pereboom's claims about his cases have been criticised. But a more fundamental objection has been overlooked. If we reject the "lens of moral responsibility" approach to questions of free will, a compatibilist can respond to Pereboom's argument by saying: "I don't know about moral responsibility. That is a different and difficult thing. My claim is only that *free will* is compatible with determinism. I don't have an analysis of free will, nor can I even provide a sufficient condition for having free will. But I know it when I see it, and nothing in your argument gives me any reason to believe that determinism would rob me of free will."

Defenders of this way of thinking about the free will/determinism problem would, I suspect, dismiss this objection on the grounds that they are interested only in *moral* freedom and *moral* responsibility. Some of them even go so far as to deny that we have *any* intuitions about free will independent of our intuitions about praiseworthiness, blameworthiness, and moral responsibility.

But I don't think that this is right (and not just because I disagree with the claim about intuitions)¹⁹. There are different kinds of questions we can ask, about free will, about moral responsibility, and about the relevance of determinism to free will and to moral responsibility.

Existential Questions

1. Do we have free will?
2. Are we ever morally responsible for anything?

Compatibility Questions

3. If determinism were true, would we have free will? (More generally, is determinism compatible with free will?)
4. If determinism were true, would we be morally responsible? (More generally, is determinism compatible with moral responsibility?)

Commonsense believes that we have free will and believes that we are morally responsible for at least some of the things that we do. Commonsense is not familiar with

July 31, 2009

the philosophical term of art 'determinism', but once the thesis is explained to someone, that person can give answers to the two Compatibility questions. The most common *first response*, on learning what the thesis of determinism says, is to answer 'no' to both questions; that is, to say that if determinism were true, we would be neither free nor morally responsible.

But the questions are different, and we should not define our terms in a way that precludes the possibility that the answers are different -- for instance, that the answers to the Compatibility Questions are "No; if determinism were true, we would not have free will but, yes, we would still be morally responsible", or "Yes, if determinism were true, we would still have free will, but, no, we would not be morally responsible". Alternatively, the answer to the free will question might be 'don't know' and the answer to the moral responsibility question might be 'no', or 'yes'. Or vice versa.

The problem with the Strawson, Wolf, Bok, and Strawson way of viewing free will through the lens of moral responsibility is that it conflates questions 1 and 2 and thus also 3 and 4, *forcing* the same answers to the two Existential Questions as well as the same answers to the two Compatibility questions.²⁰

The problem with the Pereboom, McKenna, and Markosian way of viewing free will through the lens of moral responsibility is more subtle, since this view permits someone to answer 'yes' to the Existential and Compatibility questions about free will (or to the corresponding questions about Markosian's "Free Will thesis") while answering 'no' to the Existential and Compatibility questions about moral responsibility. But since this view *defines* free will as the kind of freedom that is required for moral responsibility, the following combinations of answers are ruled out, *by stipulation*:

"Yes, we have free will and yes, we are morally responsible; *no*, if determinism were true, we would not have free will; *yes*, if determinism were true, we would still be morally responsible."

July 31, 2009

"Yes, we have free will, and yes, we are morally responsible; *I don't know* if we would still have free will if determinism were true, but *yes*, if determinism were true, we would still be morally responsible."²¹

I do not defend either of these combinations of answers. I believe that we have free will and that we are morally responsible for some of our actions and I believe that determinism is compatible with free will as well as moral responsibility. But these are intelligible positions, and they should not be ruled out, *in advance*, by the terminology we use to discuss the free will/determinism problem.

Objection: Maybe you are right to complain that some philosophers do not take sufficient care to distinguish questions about free will from questions about moral responsibility. But surely you can't deny that free will, *as we ordinarily understand it*, is a necessary (though not sufficient) condition of moral responsibility. After all, if we don't have free will we are always as helpless as a falling rock and blame is always unjust. So we should impose the following constraint on accounts of, or claims about, free will: that they be accounts of, or claims about, the *kind of thing* that is a non-trivially necessary condition of moral responsibility.

Until recently, I would have agreed.²² But I now think that we should not insist, *starting out*, that free will *must* be something that is necessary for moral responsibility. If we insist on this, we are immediately forced into taking a position concerning a philosophical literature that appears to be hopelessly stalemated – the literature concerning Harry Frankfurt's famous argument against the Principle of Alternate Possibilities.²³

Here's the problem. Suppose we have the commonsense view that to have free will is to be able to do otherwise, at least sometimes. And suppose we also say that having free will is a necessary condition for being morally responsible. So far, so good. But Frankfurt and his defenders argue that someone may be morally responsible even if she is *never* able to do otherwise. If they are right, then either free will is not a necessary condition of moral responsibility or we may have free will even if we are *never* able to do otherwise.

July 31, 2009

Now, I think that Frankfurt and his followers are mistaken, and I have written several papers arguing this.²⁴ But I don't want to have this argument every time I write a paper on free will. So let me be explicit. In this paper, I'm not going to assume that free will has *any* conceptual connection with moral responsibility -- not even the apparently innocuous connection of being a necessary but not sufficient condition of moral responsibility. Perhaps free will is necessary for moral responsibility; perhaps it is not. That is not my concern here. My concern here is *only* with the question of whether free will is compatible with determinism.²⁵

I believe that viewing the free will/determinism problem through the lens of moral responsibility has distorted discussion of the problem. Our view of ourselves as free *and* responsible agents is so central to our lives that there is a danger of being lost in one big problem: "How, if at all, is moral responsibility possible?" (Or the only slightly smaller version of this problem: "How, if at all, is the freedom required for moral responsibility possible?") And when we think in these terms, our views tend to polarize in one of two directions: Either free will ("moral freedom") becomes as problematic as moral responsibility²⁶ or moral responsibility and free will ("moral freedom") become so unproblematic that it's hard to see what the fuss was about.²⁷

I believe that if we take away the distorting lens of moral responsibility, we will be in a better position to solve the easier, or at least smaller, or, at least, non-moral or pre-moral, problem of whether determinism would rob us of the everyday kind of free will we unreflectively suppose we have almost all the time -- the kind of free will that licenses talk of alternatives and options, choices, unrealized possibilities, missed opportunities, and wasted abilities.

We will need to say more about what free will is. But let's start with determinism.

2. How not to Think about Determinism

July 31, 2009

Roughly stated, determinism is the conjunction of two claims: i) that we are *no exception* to the laws that govern everything else in the universe; and ii) that the laws state *sufficient*, as opposed to merely necessary, or probabilistic, conditions. More precisely, determinism is the thesis that a complete description of the intrinsic state of the world at any time *t* and a complete statement of the laws of nature together entail every truth about the world at every time later than *t*.

Let's call a possible world "deterministic" iff the thesis of determinism is true at that world; indeterministic iff the thesis of determinism is false at that world. There are different ways a world might be indeterministic. Most radically, a world might be indeterministic by being what I will call a "lawless world"; that is, by being a world where there are no laws at all. A less radical way in which a world might be indeterministic is by being what I'll call a "limited law" world; that is, by being a world where there are laws, but some of the laws are limited in scope or application.²⁸ The least radical possibility, and the one that corresponds most closely to what quantum physics appears to tell us about the actual world, is that a world is indeterministic by being what I will call a "probabilistic world". A probabilistic world is a world where there are laws, and the laws are all-encompassing rather than limited in scope or application, but at least some of the fundamental laws are probabilistic rather than deterministic.

Determinism is not an ontological thesis. Determinism neither entails physicalism nor is entailed by it. There are possible worlds where determinism is true and physicalism false; e.g., worlds where minds are nonphysical things which nevertheless obey strict deterministic laws.²⁹ And there are possible worlds (perhaps our own) where physicalism is true and determinism is false.

It is important to distinguish determinism from some claims about causation:

UC Every event has a cause.

UEC Every event has an event-cause.

UEEC Every event has an event-cause, and no other kind of cause.

July 31, 2009

UC is the weakest claim; it says that every event has a cause, but it doesn't say anything about the entities that are the causes of events; it is consistent with the claim that the causes of events are (or include) events, facts, omissions, persisting objects or agents, or some other kind of entity.³⁰ **UEC** says that every event has at least one event-cause, but it leaves open the possibility that events also have other kinds of causes. **UEEC** makes the strongest claim; it says that every event has a cause, and that the causes of events are *always and only* events.

We should not assume that determinism either entails or is entailed by even the weakest of these claims about causation.

UC - the thesis that every event has a cause - might be true at a world where determinism is false. Consider a world – perhaps our world – where the fundamental laws are probabilistic. At such a world, there are no sufficient causes. But it would be a mistake to conclude that there is no causation. It is now generally agreed, by philosophers who work on causation, that causation may be 'chancy' or probabilistic.³¹ More controversially, **UC** might be true at a world that is neither deterministic nor probabilistic. If there can be causes without laws (if a particular event, object, or person can be a cause without instantiating a law)³², then it may be true, even at a lawless or limited law world, that every event has a cause.

It's less clear whether every deterministic world is a world where **UC** is true. Whether this is so depends on what the correct theory of causation is; in particular, it depends on what the correct theory of causation says about the relation between causation and law. If there can be laws without causation, for instance, there may be deterministic worlds where **UC** is false.

What is clear, however, is that we should *not* make the assumption, almost universally made in the older literature, that determinism is (or is equivalent to) the thesis that every event has a cause. This is an important point, because some of the older arguments in the literature assume that to deny determinism is to claim that there are uncaused events.³³

July 31, 2009

Determinism should not be confused with naturalism. There are different ways of formulating naturalism, but for our purposes we can define naturalism as the conjunction of two claims that say that we are no exception to the causal and nomological workings of the universe.

No Exception (Laws): We are no exception to the laws that govern everything else in the universe.

No Exception (Causation): When we cause things to happen, we do so in a way that differs in complexity but not in kind from the way that natural objects cause things to happen.

Note what these claims say, and what they do not say.

No Exception (Laws) says that the laws apply to us in the same way that they apply to other things; while *leaving it entirely open what the laws are*, and without saying anything about how laws ‘govern’ or ‘apply’ to things. That is, **No Exception (Laws)** is neutral with respect to the claim that the actual world is a deterministic world, a probabilistic world, or a limited law world and it is also neutral with respect to rival philosophical accounts of the nature of laws. **No Exception (Laws)** merely says that *we* are not a special case, so far as the laws are concerned.³⁴

No Exception (Causation) says that our causal interactions with the world are of the same kind as the causal interactions of other things with each other. This leaves it entirely open what these other causal interactions are like, and it leaves it entirely open what the correct theory of causation is. **No Exception (Causation)** merely says that *we* are not a special case, so far as causation is concerned.³⁵

Even though determinism is distinct from naturalism, there is a relation between the two that is worth noting. In the literature, one sometimes hears the claim that ‘determinism is the worst form of naturalism’, and that to be a compatibilist is to be someone who not only believes that we have free will (because she believes that naturalism is true and that free will is compatible with naturalism) but who also believes that *we would have free*

July 31, 2009

will even if science told us that *the 'worst form' of naturalism is true* – that is, if the fundamental laws turn out to be *strict deterministic laws*. This seems right. A libertarian incompatibilist may be a naturalist because she believes that the fundamental laws are probabilistic or limited in the right kinds of ways. But her beliefs about free will are contingent on what science tells us; she draws the line at the “worst case scenario”. If science told us that determinism is true, the libertarian incompatibilist would give up her belief that we have free will.

The upshot of our investigation into determinism and related claims is this: Since we don't know whether determinism is true or false we cannot think of the free will/determinism problem as the problem of reconciling our belief in naturalism with our belief that we have free will. Even if we think that the true threat to free will is naturalism, and that determinism is merely the 'worst case' of naturalism, we need to address the question of whether *determinism*, and not just naturalism, is compatible with free will. If we are incompatibilists, we need to argue that determinism would preclude us from having the free will we might otherwise have. And if we are compatibilists, we need to defend, not only the claim that we *in fact* have free will, but also the claim that we would have free will *even if* determinism turned out to be true.

3. How to Think about the Free Will/Determinism Problem

I propose three constraints on what kind of problem the free will/determinism problem should be:

1. It should be a problem about something obvious and apparently undeniable like ‘there are tables and chairs’ or ‘things continue to exist through time’. Let's call this “the Moorean fact about free will”.³⁶
2. Determinism should be *prima facie* incompatible with the Moorean fact. (Everyone's *first thought*, on grasping what the thesis of determinism says, is: “But if that were true, we wouldn't have free will.”)

July 31, 2009

3. Indeterminism should not be prima facie incompatible with the Moorean fact. (It takes philosophical work to get students to see that indeterminism might also be problematic for free will.)

Free will defined so that it is conceptually linked to moral responsibility does not fit these constraints. At one time, the fact that we have the kind of freedom that justifies blame and punishment might have counted as a Moorean fact. But this is no longer true; these days it is easy to raise doubts about moral responsibility without saying a word about determinism.

Free will defined so that it is conceptually connected to something valuable or ‘worth wanting’ does not fit these constraints. Nor does free will defined in some inherently vague way like “that which makes us dignified and worthy of respect” or ‘that which makes us different from everything else in nature’.

You may wonder: What could free will possibly be, if it is to satisfy these constraints?

I gave my answer at the beginning of this paper. We should understand the free will/determinism problem as a problem about choice. Our core conception of ourselves as agents with free will consists in our belief that we are often in situations where we both *make and also have a choice*.

You offer me cheesecake or chocolate mousse; I pick the mousse. I’m at a fork in the road and must decide between the winding scenic coastal route, and the speedy but boring freeway; I take the freeway. The speeches are finally over, and the vote has been called: all in favor, raise their hands; I raise my hand. We are in situations like these daily, and in such situations all of the following seem to be true:

July 31, 2009

1. We believe that we have options (alternative courses of action): that is, we believe that there are at least two different things we are able to do. (Eat one dessert or another, take one road or the other, raise our hand or keep it lowered). Call these options 'A' and 'B'.
2. We deliberate (or ponder, ruminate, consult our feelings, etc.) between A and B.
3. We decide to do A. (We make up our minds, make a choice, form an intention.)
4. We do A.
5. The belief that we had before we made up our mind was a true belief. While we were deliberating, we really were able to do A *and* we really were able to do B.

We are, of course, sometimes mistaken about what our options are. We ponder the offerings on the menu, trying to decide between the mousse and the cheesecake, not realizing that the mousse is not available today. We ponder going out for a drive, having forgotten that the car is at the garage. In these situations, 1-4 may be true, but 5 is false.

This shows that we are not entitled to infer, from the fact that we deliberate between A and B, that A and B are both options for us. We may be mistaken about what we are able to do. But it seems incredible to suppose that we are *always* mistaken.

Call the belief that we are often, or at least *sometimes*, in situations where our beliefs about our options are correct and 1-5 are all true: **Choice**.

I take **Choice** to be an uncontroversial part of our commonsense view of ourselves as agents with free will. (The core part of that view.) **Choice** seems as obviously and undeniably true as other facts that only a philosopher would question: the fact that there are tables and chairs, the fact that things continue to exist through time, the fact that I have a pair of hands.

July 31, 2009

I call **Choice** a Moorean fact. Like the fact that there are tables and chairs and continuing things, the fact that we are sometimes able to do more than one thing is more obvious and self-evident than any *analysis* that we can give of it. We do not need to give arguments or provide an analysis in order to defend our belief in a Moorean fact.

In saying this, I do not mean to suggest or imply any of the following: That we should never be convinced, by argument, to reject a Moorean fact;³⁷ that we have a priori knowledge of Moorean facts; that Moore's response to the skeptic is correct;³⁸ that Moorean facts are philosophically unproblematic. In particular, I don't mean to suggest that if something is a Moorean fact there is nothing more to be said about what kind of fact it is. (Compare the debate between perdurantists and endurantists about the two different ways in which it might be true that things continue to exist through time.) But I do claim that we need to be convinced, *by very good argument*, before we should be persuaded to give up our belief that a Moorean fact is a fact.

It would be lonely if I were the only philosopher who saw things this way. Luckily I am not. I offer the following quotes from the philosopher who is largely responsible for how I think about the free will/determinism problem:³⁹

The problem of metaphysical freedom is a problem so abstract that it arises in any imaginable world in which there are beings who make choices. Consider some important choice that confronts you...Consider the two courses of action that confront you. I'll call them simply A and B. Do you really not believe that you are able to do A and able to do B?...It seems clear to me that when I am trying to decide which of two things to do, I commit myself, by the very act of attempting to decide between the two of them, to the thesis that I am able to do each of them.⁴⁰

July 31, 2009

The free-will thesis is the thesis that we are sometimes in the following position with respect to a contemplated future act: we simultaneously have both the following abilities: the ability to perform that act and the ability to refrain from performing that act. ⁴¹

[Peter van Inwagen]

4. How not to think about the Consequence Argument

We need to be forced by argument to give up our belief that a Moorean fact is a fact. Appeal to intuition is not an argument. We already know that taking determinism seriously causes people to doubt the Moorean fact. The philosophical problem is to decide whether these intuitions – this way of responding to the problem – is justified.

There is only one *argument* in the literature for the conclusion that if determinism is true, the Moorean fact of **Choice** does not obtain. ⁴² It is the Consequence argument, due, most famously, to Peter van Inwagen. ⁴³ Here is an informal statement of the argument:

If determinism is true, then our acts are the *logical consequence* of the laws of nature and events in the remote past. But we don't have a choice about what happened before we were born, and we don't have a choice about what the laws of nature are. Therefore, we don't have a choice about the consequences of these things, *including our present acts*.

It can't be denied that this way of stating the argument tends to produce the desired effect. Our attention is directed to two apparently undeniable⁴⁴ facts: that we have no choice about the remote past or the laws of physics, and that at a deterministic world *everything*, including our present acts, can be logically deduced from facts about the remote past and the laws. Thinking in these terms causes us to think: "If determinism is true, then *everything has already been decided*, long before I was born. There is nothing

July 31, 2009

that remains up to me. I *never have a choice about anything*, not even in those situations in which it seems clear as day that I am able to do more than one thing.”

The argument has convinced many. It has, I think, made incompatibilism respectable again. But it has also been widely misunderstood, and it has not received nearly the attention from compatibilists that it should have received. Here is a typical response, from Hilary Bok:

“I have described two conceptions of possibility: the broad compatibilist conception of possibility and possibility tout court...van Inwagen’s argument against compatibilism presupposes rather than establishes, the claim that the sense of possibility relevant to libertarian freedom is possibility tout court.’⁴⁵

By “two conceptions of possibility”, Bok means two conceptions or senses of ‘could have done otherwise’. She is accusing van Inwagen of begging the question by using ‘could have done otherwise’ in a way that entails that a person could have done otherwise only if determinism is false; that is, only if the person's doing otherwise is compossible with the actual past and the laws.

Here is van Inwagen’s response to this kind of compatibilist response:

“Many philosophers, in attempting to spell out the concept of free will, use the phrase ‘could have done otherwise’. I did so myself in *An Essay on Free Will*. Nowadays, however, I very deliberately avoid this phrase. I avoid it because “could have done otherwise” is ambiguous and (experience has shown) its ambiguity has caused much confusion in discussions of free will... A whole chapter of Daniel Dennett’s first book on free will was written to no purpose because he didn’t realize that ‘could have done’ sometimes means ‘might have done’ (and this ‘might’ is ambiguous; it has both an ontological and an epistemic sense) and sometimes ‘was able to do’. “

July 31, 2009

“I want to make what seems to me to be an important point, a point that is, in fact, of central importance if one wishes to think clearly about the freedom of the will: compatibilists and incompatibilists mean the same thing by ‘able’. And what do both compatibilists and incompatibilists mean by ‘able’? Just this, what it means in English, what the word means. And, therefore, ‘free will’, ‘incompatibilist free will’, ‘compatibilist free will’ and ‘libertarian free will’ are four names for one and the same thing. If this thing is a property, they are four names for the property *is on some occasions able to do otherwise*. If this thing is a power or ability, they are four names for the power or ability to do otherwise than one in fact does.”⁴⁶

In these passages, van Inwagen makes two important points -- that it is a strategic mistake to discuss questions of free will in terms of ‘can’ or ‘could’, and that it is a mistake to think that compatibilists and incompatibilists mean different things by ‘able’. I agree with both points. It doesn't follow, however, that whenever we say, in ordinary English, that someone is to do something, we always mean the same thing. Perhaps the locution is ambiguous between two or even three different meanings.⁴⁷ But since we are anchoring our discussion to the Moorean fact of Choice, this should ensure that we are discussing the same thing.⁴⁸

The question before us, then, is whether determinism has the consequence that the ability to do otherwise that we take for granted whenever we make choices is an ersatz ability, an illusion, no ability at all. Could it be that we are *always* mistaken when we believe that we have a choice about what to do?

We need to take a closer look at a more precisely formulated version of the Consequence Argument. The version I will discuss is due to David Lewis in his classic “Are We Free to Break the Laws?”⁴⁹, a paper that van Inwagen commends as “the finest essay that has ever been written in defense of compatibilism – possibly the finest essay that has ever been written about any aspect of the free will problem”.⁵⁰

5. How to think about the Consequence Argument

July 31, 2009

Think of the argument as a reductio. A compatibilist is someone who claims that the truth of determinism is compatible with the existence of the kinds of abilities that we assume we have in typical choice situations. Let's call these 'ordinary abilities'. The Consequence argument claims that if we suppose that a deterministic agent has ordinary abilities, we are forced to credit her with incredible abilities as well.

Here's the argument: Suppose, for reductio, that determinism is true, and someone, call her Dana, did not raise her hand but had the ordinary ability to do so. If Dana had exercised her ordinary ability – if she had raised her hand -- then either the remote past or the laws of physics would have been different (would have to have been different). But if that's so, then Dana has at least one of two incredible abilities – the ability to change the remote past or the ability to change the laws. But to suppose that Dana has either of these incredible abilities is absurd. So we must reject the claim that Dana had the ordinary ability to raise her hand.

The first thing to note about this argument is that it relies on a claim about counterfactuals. The argument says that one of these counterfactuals is true:

Different Past: If Dana had raised her hand, the remote past would have been different.

Different Laws: If Dana had raised her hand, the laws would have been different.

Both these counterfactuals seem incredible to us, because we are not used to thinking in terms of determinism, but we are good at evaluating counterfactuals, especially those counterfactuals that we entertain in contexts of choice. And when we contemplate Dana's options, we assume that:

Same Past and Laws: If Dana had raised her hand, the past and the laws would still have been exactly the same.

July 31, 2009

So both **Different Past** and **Different Laws** strike us as false. But that doesn't mean they are false, and if determinism is true, at least one of these counterfactuals *is* true.

Our best theory of counterfactuals – David Lewis's theory – tells us that our commonsense assumption about the past is correct and **Different Past** is false.⁵¹

According to Lewis, if Dana had exercised her ordinary ability to raise her hand, the past would have been exactly the same as it actually was until a time shortly before she chose to raise her hand.⁵² So Dana's ordinary ability to raise her hand does not entail an incredible ability to change the past.

Lewis's theory of counterfactuals does have the consequence that **Different Laws** is true. But **Different Laws** sounds more incredible than it is. If we don't understand Lewis's theory of counterfactuals, we might think that **Different Laws** means:

New Set of Laws: If Dana had raised her hand, our laws would have been replaced by a new and entirely different set of laws.

But this is not how we evaluate counterfactuals. We evaluate counterfactuals by considering the *closest* worlds where the antecedent is true, and the closest worlds are worlds where there are no gratuitous changes from actuality. To suppose a wholly new and different system of laws is to suppose a gratuitous change. Lewis's theory tells us that if Dana had exercised her ordinary ability to raise her hand, the past would have been exactly the same until the occurrence of a small and inconspicuous event – an extra neuron firing in Dana's brain – which marks the divergence from our actual history. Lewis calls this event a 'divergence miracle' (or 'law-breaking event'), since it is an event that is unlawful by the standards of *our* laws. But he tells us that it's a mistake to think that at these non-actual closest worlds one of our laws has been replaced by a *contrary* law; rather, we should think that one of our laws has been replaced by something that might be described either as an "almost-law" or, perhaps, as a *version of the broken law*, "complicated and weakened by a clause to permit the one exception."⁵³

So we should understand **Different Laws** as:

July 31, 2009

Slightly Different Law : If Dana had raised her hand, one of our laws would have been replaced by a *slightly* different law.

You may say: "This hardly helps. If Dana had raised her hand, a law-breaking event -- the divergence miracle -- would have occurred. So the compatibilist remains committed to the claim that Dana has an incredible ability -- the ability to *break* the laws."

But – and this is Lewis’s main point in his reply to van Inwagen – the ability to break a law is the ability to *cause a law-breaking event*. Dana has this ability only if she has the *ability to cause* the divergence miracle. But she doesn’t have this ability. The divergence miracle is over and done with by the time that she acts.⁵⁴

There is more to be said, but I believe that Lewis’s reply succeeds. (I also believe, though I cannot defend this here, that the success of his reply does not require the truth of any particular philosophical account of lawhood.) The Consequence Argument fails. Compatibilists do not need to hide behind different senses or conceptions or kinds of ability, nor do they need to fear that the ability to do otherwise that a deterministic agent has is a “weaker” or "less robust" ability than the ability the libertarian claims we have. In the relevant sense, the *ordinary sense* that we use when we deliberate for the purpose of making a choice, deterministic agents are able to do otherwise.

6. How to Think about Free Will

I would like to close by saying a bit about the nature of the Moorean fact of Choice.

Consider a typical choice situation. I am considering whether to vote ‘yes’ by raising my hand. I think for a moment, and then keep my hand lowered. I refrained from voting ‘yes’, but I was able to vote ‘yes’ because I was able to raise my hand. I had the ordinary ability to do so. Nothing stopped me from doing so. Or so it seems.

July 31, 2009

What are the truth-makers of this fact -- the fact that I was able to vote “yes” by raising my hand?

A simple answer, once widely accepted, but now just as widely rejected is:

Simple Conditional Analysis: “I was able to raise my hand” =(def.) “if I had chosen (or decided, intended, tried, wanted, etc.) to raise my hand, I would have raised my hand”.

More generally, at one time the debate about the free will/determinism problem was regarded as a debate about whether ‘could have done otherwise’ or ‘is able to do otherwise’ can be *analysed* in terms of some kind of counterfactual or subjunctive conditional. Compatibilists argued that ‘can’ is ‘hypothetical’ or ‘conditional’ or “constitutionally iff”; incompatibilists argued that ‘can’ is “categorical” and not analysable. There was a rich literature of cases, and while the discussion did not formulate the problem in terms of what I have been calling the Moorean fact, it was, in my view, an improvement over the current free will literature in this respect: it did not view free will through the lens of moral responsibility. The debate was focused on the meaning of ‘can’ or ‘could’ or ‘is able to’ claims, rather than their moral or normative significance.⁵⁵

It is an unfortunate fact about the history of the free will literature that this debate came to an end sometime in the mid-sixties, shortly before the development of the Lewis/Stalnaker possible worlds semantics for counterfactuals.⁵⁶ The debate came to an end because *everyone* agreed that the attempt to provide a Conditional Analysis of ‘could have done otherwise’ was doomed in principle.⁵⁷ Compatibilists had to look to other strategies for defending compatibilism, and the years since then have seen a wide variety of different strategies: the Strawsonian program⁵⁸; Frankfurt’s argument against the Principle of Alternate Possibilities and the rise of Semi-Compatibilism⁵⁹; accounts of free will in terms of one’s “deep” or “real” self⁶⁰; and, more generally, the increasing

July 31, 2009

tendency to view the free will/determinism problem through the lens of moral responsibility.

But I digress. Here's an example of the kind of case that convinced everyone that the attempt to provide a Conditional Analysis of 'is able to do X' is flawed *in principle*.

Mary is under general anesthesia, which temporarily prevents her from thinking and thus prevents her from deciding or choosing to do anything. If she choose (or decided, intended, etc.) to raise her arm, she *would have to be* conscious, and if she were conscious she would succeed in raising her arm, if she chose. But since she is *in fact* unconscious, she is unable to raise her arm. So it is true that if Mary chose to raise her arm, she would, but false that Mary is able to raise her arm.

The moral drawn was that abilities are 'categorical' rather than 'iffy'. Mary's unconscious state has changed certain categorical properties of her brain, the properties that enable her to make choices on the basis of reasoning and deliberation. It is because she lacks these categorical properties that she lacks the ability to *decide or choose whether to raise her arm*. And because she lacks the ability to decide or choose to raise her arm, she also lacks the ability *to raise her arm*.

I think that this way of diagnosing the case is a mistake. I agree that the Simple Conditional Analysis is false. And I agree that this case shows that part of the truth-maker of an ability claim is something that is 'categorical', as the old literature would have said, or, as we would say today, has a *causal basis* in the *intrinsic properties of the person*.⁶¹ It is because general anesthesia has changed the intrinsic properties of Mary's brain that she (temporarily) lacks the ability to make decisions and choices on the basis of deliberation. But I don't think that the ability to raise one's arm requires the ability to deliberate about the pro's and con's of arm-raising; a child might have the first ability without having the second. I think that we should say that Mary has temporarily lost the second ability – the ability to *decide or choose* to raise her arm -- while retaining the ability the child has. After all, general anaesthesia is not a paralytic; it hasn't changed the

July 31, 2009

intrinsic properties of the parts of Mary's brain and body that enable Mary to *move her arm*.

But we don't have to settle this point here. My aim here is not to defend an *analysis* of ability; my aim is to say a bit more about the nature of the Moorean fact of Choice. The lesson to be learned from Mary's case is that *at least part* of the truth-maker of the Moorean fact is constituted by the intrinsically-based abilities that a choice-maker has.⁶² When I refrain from raising my hand, I am different from Mary in the following respect: I am conscious, and I exercise, and therefore possess, the ability to deliberate for the purpose of deciding whether or not to raise my hand. I also have the ability to keep my hand lowered, as well as the ability to raise it; I exercise the first ability; do not exercise the second.⁶³ We may sum this up, somewhat vaguely, by saying that I have 'what it takes' both to *choose* to raise my hand and also to act on my choice by *raising my hand*.

But having the intrinsically-based abilities to *choose* and to *do* both the acts you are contemplating doing does not suffice for the truth of the Moorean fact. I may have the intrinsically-based ability to choose and to do some act that I am contemplating doing without being *able*, in the *sense relevant to Choice*, to do that thing. For suppose that I am contemplating whether to stay home or to go out for a walk. I've got the intrinsically-based ability to go for a walk -- my legs aren't broken, I'm not paralysed. Nor am I disabled in any other way; I don't suffer from agoraphobia, for instance. I decide to stay inside and do so. I made a choice. But, unbeknownst to me, I *didn't have the choice I thought I had*. For I am a prisoner; the door is locked and an armed bodyguard lurks outside, ready to use whatever force it takes to prevent me from going for a walk.

In this situation, I have what we would ordinarily call the ability to go for a walk (the intrinsically-based ability) but I lack what we would ordinarily call the opportunity. I have what it takes to go for a walk, but, unfortunately for me, I am in a situation in which something prevents me from exercising my ability.

July 31, 2009

But this kind of situation is not the norm. Ordinarily, when we deliberate and make a choice between what we take to be two options, A and B, we've got the ability and *also* the opportunity to *choose* and also to *do both* these things. Ordinarily, when I choose to stay home rather than go for a walk, I have the opportunity as well as the ability to go for a walk (and, of course, the opportunity as well as the ability to *decide* whether to stay or to go).

How should we understand the ability/opportunity distinction? I suggest the following: What *abilities* you have is determined by a subset of your *intrinsic properties*; ability facts are, very roughly, facts about what you are like 'beneath your skin'. What *opportunities* you have is determined by a subset of the facts about your *surroundings*. Abilities are shared by duplicates; you and your atom for atom duplicate have exactly the same abilities. Opportunities are a function of your situation; if I swapped places with you, I would also swap opportunities.⁶⁴

This is a partly stipulative (partial) definition; the intention is to provide a mutually exclusive and jointly exhaustive classification of all the ways in which it might be true, of a particular person, on a particular occasion, that she *doesn't have a choice* about something. On my proposal, there are just two ways in which this might be true: Either the fault is in *her*; she lacks, temporarily or permanently, something that would enable her to choose or to do that thing. If so, then I say that she lacks either the ability to choose or the ability to do that thing. Or the fault is in her surroundings; that is, something *extrinsic to her* either prevents or would prevent her from choosing or doing that thing (perhaps by failing to supply her with what she needs to successfully choose or do that thing). If so, then I say that she lacks the opportunity to choose or to do that thing.

The borderline between ability and opportunity may be unclear or disputed insofar as it is unclear, or in dispute, where the boundaries of the person (at the relevant time) are. I suggested a physical boundary ("under the skin"), but this is over-simple since implanted devices controlled by nefarious neurosurgeons are under a person's skin but are not part of *her*. But we can set these borderline and difficult cases aside so far as the free

July 31, 2009

will/determinism problem is concerned. For that problem is to decide whether determinism really has the shocking consequence it appears to have: that no one *ever* has a choice about anything.

Here is an operational test (not a definition) of when someone is able to do something in the sense relevant to **Choice**. Suppose that I have the *ability to do* something – to go for a walk, to ride a bike, to get to the airport in time to catch a plane. And suppose, as we do ordinarily suppose, that I have the ability as well as the opportunity *to choose* whether or not to do that thing. (I'm not asleep, or too drunk to think, or pathologically depressed, or hopelessly distracted by the noise from my neighbors, etc.). If so, then I also have the opportunity *to do* that thing if it is true that *if I now tried to do it, I would succeed*. (Or, perhaps, would have a good, or good enough, chance of succeeding.) The cases where I am locked in my room, lack a bike, or wake up too late to catch my flight are cases where I have the ability but lack the opportunity; the difference lies in the fact that the counterfactual is false. (The Simple Conditional Analysis, though false, was on the right track.)

I propose that this is all that there is to the Moorean fact of Choice. The Moorean fact is the fact that we are often in situations in which we exercise our ability to choose *and* to do something A while having the ability as well as the opportunity to choose *and* to do something else B instead. Any argument for the claim that the truth of determinism is incompatible with the Moorean fact must, then, be an argument for one of two surprising claims: that determinism has the consequence that our abilities go out of existence whenever we don't exercise them or that determinism has the consequence that we lose our opportunities whenever we don't take advantage of them.⁶⁵

Kadri Vihvelin

University of Southern California

July 31, 2009

¹ "The Dilemma of Determinism", *An Address to the Harvard Divinity Students*, 1884, reprinted in *The Will to Believe and other Essays*, New York: Dover Publications, 1956.

² I don't mean to suggest that I endorse the No Choice argument -- I am a compatibilist. But I do think that the argument succinctly sets out the *problem*.

³ "Freedom and Resentment", *Proceedings of the British Academy* 48 (1962), 1-25.

⁴ *Freedom Within Reason*, 1990, Oxford: Oxford University Press, p. 4.

⁵ *Freedom and Responsibility*, Princeton, Princeton University Press, 1998, p. 25.

⁶ *Freedom and Belief*, Oxford: Clarendon Press, 1986, p.1.

⁷ "The Impossibility of Moral Responsibility", *Philosophical Studies* 75 (1994), 5-24.

⁸ P.F. Strawson, Wolf, and Bok are all compatibilists, but their ways of defending compatibilism are very different. Galen Strawson argues that "true moral responsibility" requires a kind of self-making that is impossible for finite, non-godlike creatures like us.

⁹ Derek Parfit, *Reasons and Persons*, Oxford: Clarendon Press, 1984.

¹⁰ For a discussion of this last requirement on moral responsibility, as well as an interesting discussion of two senses or "faces" of responsibility, see Gary Watson, "Two Faces of Responsibility", in his *Agency and Answerability*, Oxford: Clarendon Press, 2004.

¹¹ In the quoted passage, above.

¹² Galen Strawson is quite clear about the problem being our nature as finite beings, as opposed to determinism. The compatibilists who think of free will this way - Wolf, Bok, and P.F. Strawson - do not see determinism (as opposed to the view that we are part of the natural order of things) as the primary problem either.

¹³ This way of articulating the view is mine but the view expressed is widespread in contemporary discussions of the free will/determinism problem.

¹⁴ "Determinism al Dente" , *Nous* 29 (1995), 21-45.

¹⁵ "Compatibilism", *Stanford Encyclopedia of Philosophy*, 2004.

¹⁶ "A Compatibilist Version of the Theory of Agent Causation", *Pacific Philosophical Quarterly* 80 (1999), 257-277. Markosian argues that if agent-causation is possible and relevant to moral freedom and moral responsibility, then it is compatible with determinism.

¹⁷ Derk Pereboom, "Determinism al Dente", *ibid*, and *Living without Free Will*, New York: Cambridge University Press, 2001.

¹⁸ Pereboom says that compatibilists must provide sufficient conditions for moral responsibility because "it is essential to their case that we can attribute moral responsibility in certain standard conditions." (*Living without Free Will, ibid*, p. 101.)

¹⁹ I don't care much for the recently over-used term 'intuition', but I will be arguing that our intuitions (better: core commonsense beliefs) about free will are based on everyday cases of making and having a choice. The choice may be one for which we hold the person responsible, but it need not be.

²⁰ P.F. Strawson, Wolf, and Bok answer "yes" to the two Existential questions and "yes" to the two Compatibility questions. Galen Strawson answers "no" to the two Existential questions, and, "no" to the two Compatibility questions; however, on his view it makes no difference whether or not determinism is true; either way, we would not have free will or be morally responsible.

²¹ It might look as though Markosian can allow these combinations of answers because he explicitly introduces 'moral freedom' as a technical term and because he issues two disclaimers - he does not assume that moral responsibility and moral freedom require the ability to do otherwise and he acknowledges that moral freedom "may be very different from 'freedom' in some other popular senses of the word." But Markosian in effect appropriates 'free will' as a term for moral freedom by defining the Free Will Thesis as the thesis that some actions are *morally free* and compatibilism as the thesis that the Free Will thesis is compatible with determinism. Someone who wants to respond to the two Compatibility questions by answering 'no' and 'yes' is forced to say something like: 'If determinism were true, there would be some sense of 'free' in which I would be less free than I believe I am, but I would still have the freedom that makes the Free Will Thesis true, so I would still be morally responsible.' ("A Compatibilist Version of the Theory of Agent Causation, *ibid*, p.2)

²² An example of how I recently viewed the problem "through the lens": In my "Free Will Demystified: A Dispositional Account", *Philosophical Topics* 32 (2004), *Agency*, John Fischer, ed., 427-450. " I say, at p. 428: "the 'can' relevant to free will is the 'can' that we have in mind in contexts in which we raise questions about moral responsibility." I was closer to being right in an older paper in which I describe a view that I call Simple Compatibilism and then say: "this search for a more sophisticated Compatibilist theory is a mistake. I think that Simple Compatibilism is basically right, not just as an account of freedom of action, but also as an account of free will. It's not, however, a theory of moral responsibility and much confusion will be avoided once we realize this." ("Stop Me Before I Kill Again", *Philosophical Studies* 75 (1994), 115-148, at p. 141.)

²³ “Alternate Possibilities and Moral Responsibility”, *Journal of Philosophy* 66 (1969), 820-39.

²⁴ “Freedom, Foreknowledge, and the Principle of Alternate Possibilities”, *Canadian Journal of Philosophy* 30 (2000), 1-24; “Free Will Demystified: A Dispositional Account”, *ibid*; and “Foreknowledge, Frankfurt, and Ability to Do Otherwise: A Reply to Fischer”, *Canadian Journal of Philosophy* 38 (2008), 343-372.

²⁵ I would argue (though not here), that although the correct (compatibilist) solution to the free will/determinism problem leaves the moral responsibility/determinism problem unsolved (because it is silent on the question of moral responsibility), it provides us with the following *conditional solution* to the problem: *If moral responsibility is possible for non-godlike creatures like us*, then it is compatible with determinism.

²⁶ This is Galen Strawson's view.

²⁷ Arguably, this is what happens with the "reactive attitudes" account of moral responsibility defended, famously, by P.F. Strawson, in "Freedom and Resentment", *ibid*. Free will drops out of the picture altogether. A hard determinist could accept his view.

²⁸ Cf. O'Connor, *The Metaphysics of Free Will*, Oxford: Oxford University Press, 2000 and Clarke, *Libertarian Accounts of Free Will*, Oxford, 2003.

²⁹ Van Inwagen describes such a world (a world where there are only angels) in his "The Mystery of Metaphysical Freedom", *Metaphysics: The Big Questions*, ed. by van Inwagen and D. Zimmerman, Oxford: Blackwell Publishing, 1998.

³⁰ For a defense of fact-causation and omissions as causes, see D.H. Mellor, "For Facts as Causes and Effects", *Causation and Counterfactuals*, ed. by J. Collins, N. Hall, and L.A. Paul, Cambridge, Mass.: MIT Press, 2004. For a defense of property-instances as causes, see L.A. Paul, "Aspect Causation", *Causation and Counterfactuals*, *ibid*. For defense of states as causes, see Thomson, "Causation: Omissions", *Philosophy and Phenomenological Research* 66 (2003), 81-103. For defense of primitive agent-causation, see O'Connor, *ibid*.

³¹ A world with chancy causes might be a world where some events lack causes altogether. But it need not be. The following seems to be a metaphysical possibility: a world where every event has a chancy cause.

³² G. E.M. Anscombe, *Causality and Determination: An Inaugural Lecture*, Cambridge: Cambridge University Press, 1971.

³³ This is the argument widely known as the 'determined or random' dilemma.

³⁴ This means that a libertarian (or other kind of incompatibilist -- a hard determinist or someone agnostic about the truth or falsity of determinism) does not need to reject No Exception (Laws). Different libertarian accounts differ concerning the kind of laws that are required to accommodate free will. O'Connor (*The Metaphysics of Free Will, ibid*) has some interesting suggestions concerning a kind of limited laws account. Robert Kane, (*The Significance of Free Will*, New York: Oxford University Press, 1998) and Clarke (*Libertarian Accounts of Free Will, ibid*) explore varieties of probabilistic world accounts.

³⁵ Note that a libertarian (or other incompatibilist) need not reject No Exception (Causation). This is most clear in the case of libertarians who explicitly endorse event-causation and the claim that free will is possible at probabilistic worlds (cf. Kane, *ibid*) but even libertarians who defend agent-causation and who insist that free will is possible only at limited law worlds (cf. O'Connor, *ibid*) can endorse No Exception (Causation), provided that they argue that agent-causation is a species of primitive object-causation.

³⁶ In his "Elusive Knowledge", *Australasian Journal of Philosophy* 74 (1996), 549-567, David Lewis characterizes a Moorean fact as "one of those things that we know better than we know the premises of any philosophical argument to the contrary." I like this definition, but it may be stronger than I need. If you prefer a more epistemically neutral characterization, you may think of a Moorean fact as a commonsense platitude or a belief that seems so obviously true that we typically don't bother saying it. I thank Joe Campbell for pressing me on this point.

³⁷ Here I disagree with some of the stronger claims about the evidential status of Moorean facts made by recent defenders of commonsense against the skeptic. For good discussion, see Thomas Kelly, "Moorean Facts and Belief Revision, or Can the Skeptic Win?", *Philosophical Perspectives* 19, *Epistemology*, ed. by John Hawthorne, forthcoming and William Lycan, "Moore against the New Skeptics", *Philosophical Studies* 103 (1), 2001, 35-53.

³⁸ Nor am I making *any* claim about knowledge. The free will/determinism problem isn't a problem about whether we *know* that we have free will or whether we *know* that free will is compatible with determinism; it's a problem about whether we *have* free will and whether our having free will is compatible with determinism. (It's a problem within metaphysics, not epistemology.) If the epistemic skeptic is right, then we don't know that we have hands, let alone that we have free will. My intent in calling **Choice** a Moorean fact is only to point out that our belief that we have free will is as firmly embedded within commonsense as our belief that we have hands, that the world didn't begin to exist five minutes ago, that there are tables and chairs and other minds. It has whatever epistemic status these other beliefs have.

³⁹ I offer these quotes as evidence that van Inwagen sees the free will/*determinism* problem in the way that I do -- as a problem about **Choice**. But in his recent paper, "How to think about the Problem of Free Will" *Journal of Ethics*, 2008 12, 337-341, van Inwagen describes a *larger* problem that he calls "the problem of free will" but which I

think is more accurately described as 'the free will-*moral responsibility* problem" because it is the problem of "finding out" which of three "seemingly unanswerable" arguments is "fallacious", and the conjunction of the conclusions of these three arguments is that *moral responsibility* does not exist. The problem that van Inwagen sets out is an important problem. But for the reasons I give in the "lens of moral responsibility" section, I believe that we should not think of it as "the problem of *free will*".

⁴⁰ "The Mystery of Metaphysical Freedom", in *Metaphysics: The Big Questions*, *ibid.*

⁴¹ "How to Think about the Problem of Free Will", *Journal of Ethics*, 2008 12, 337-341.

⁴² I mean that there is only one *serious* argument. There are, of course, lots of bad arguments; in particular, there are lots of arguments based on various fatalist confusions. For criticism of some of these arguments, see my "Compatibilism, Incompatibilism, and Impossibilism", *Contemporary Debates in Metaphysics*, ed. by T. Sider, J. Hawthorne, and D. Zimmerman, Oxford: Blackwell Publishing, 2008

⁴³ *An Essay on Free Will*, Oxford: Clarendon Press, 1983.

⁴⁴ I would argue, however, that these apparently undeniable facts are *not* Moorean facts.

⁴⁵ Bok, *Freedom and Responsibility*, *ibid.*, p. 97.

⁴⁶ "How to Think about the Problem of Free Will", *ibid.*

⁴⁷ Some plausible candidates for the different things we might mean by 'S is able to do X': "S has the skills, competence, or know-how to required to do X"; "S has the skills required to do X *and also* the physical capacity to use these skills (she has what it takes to do X)"; "S has what it takes to do X *and* nothing stands in her way (she's in the right place at the right time; she's got the means and the opportunity; there are no impediments or obstacles, etc).

⁴⁸ This is an oversimplification, but it will do for now. See my final section for more details.

⁴⁹ "Are We Free to Break the Laws?", *Theoria* 47 (1981), 113-121.

⁵⁰ "How to Think about the Problem of Free Will", *ibid.*

⁵¹ "Counterfactual Dependence and Time's Arrow", *Nous* 13 (1979), 455-476. For a good discussion of Lewis's theory of counterfactuals, see Jonathan Bennett, *A Philosophical Guide to Conditionals*, Oxford: Clarendon Press, 2003.

⁵² Why 'shortly before'? Because if Dana were to suddenly raise her hand, out of the blue, without choosing, intending, or even wanting to do so, she would not be exercising

any ability; her hand-raising would be something that 'just happened' to her. What Lewis has in mind is something like the following: Dana thinks of something she didn't actually think of, and this thought gives her a reason (or something she takes to be a reason) to raise her hand.

⁵³ *Counterfactuals*, Cambridge, Mass.; Harvard University Press, 1973, p.75.

⁵⁴ For more details on why Lewis's reply does not commit compatibilists to the claim that free deterministic agents have incredible abilities, see my "Freedom, Causation, and Counterfactuals", *Philosophical Studies* 64 (1991), 161-184, my *SEP* entry, "Arguments for Incompatibilism", and "Compatibilism, Incompatibilism, and Impossibilism", *ibid*.

⁵⁵ For more details about why I think the search for a conditional analysis of ability was abandoned prematurely, see my "Free Will Demystified: A Dispositional Account" *Philosophical Topics* 32 (2004), *Agency*, John Fischer, ed., 427-450.

⁵⁶ Robert Stalnaker, "A Theory of Conditionals", in *Studies in Logical Theory*, N. Rescher, ed', Oxford: Blackwell, 1968 and David Lewis, *Counterfactuals*, *ibid*.

⁵⁷ For an overview of the debate, see Bernard Berofsky, "Ifs, Cans, and Free Will: The Issues", *The Oxford Handbook on Free Will*, Robert Kane, ed., Oxford, 2002.

⁵⁸ See, for instance, R. Jay Wallace, *Responsibility and the Moral Sentiments*, Cambridge, Mass: Harvard University Press, 1994.

⁵⁹ For Frankfurt's argument, see "Alternate Possibilities and Moral Responsibility", *ibid*. This paper has spawned an enormous literature. A good introduction is in *Moral Responsibility and Alternative Possibilities*, ed. by D. Widerker and M. McKenna, Burlington, Vermont: Ashgate, 2003. For Semi-Compatibilism, see John M. Fischer, *The Metaphysics of Free Will*, Oxford: Blackwell, 1994 and John M. Fischer and Mark Ravizza, *Responsibility and Control*, Cambridge: Cambridge University Press, 1998.

⁶⁰ This literature is huge. The classic starting points are Harry Frankfurt, "Freedom of the Will and the Concept of a Person", *Journal of Philosophy* 68 (1971) and Gary Watson, "Free Agency", *Journal of Philosophy* 72 (1975). Wolf criticizes deep self views in her *Freedom within Reason*, *ibid*. Watson responds to Wolf's criticism in "Two Faces of Responsibility", *ibid*.

⁶¹ David Lewis, "Finkish Dispositions", *The Philosophical Quarterly* 47 (1997), 143-158.

⁶² What are these intrinsically-based abilities? An incompatibilist might argue that they are 'agent-causal powers', and some compatibilists would agree. (cf. Ned Markosian, *ibid*). In my "Free Will Demystified: A Dispositional Account", *ibid*, I argue that they are dispositions or bundles of dispositions, differing in complexity but not in kind from intrinsically-based dispositions like fragility.

⁶³ Arguably these are the same ability. I do not assume this here because there are other cases in which the two abilities are clearly different - for instance, cases where we deliberate between going for a walk and going for a bike ride.

⁶⁴ I am assuming, of course, that the laws are held constant.

⁶⁵ I am grateful for comments from Joe Campbell, John Carroll, Randolph Clarke, Janet Levin, Michael McKenna, Rebekah Rice, Terrance Tomkow, Gideon Yaffe, the audience at the INPC conference in Idaho in March 2008, and the audience at the University of Victoria Undergraduate Philosophy Conference in Victoria, British Columbia in March 2009.