

Folk Intuitions, Slippery Slopes, and Necessary Fictions: An Essay on Saul Smilansky's Free Will Illusionism

THOMAS NADELHOFFER AND ADAM FELTZ

Humanity is fortunately deceived on the free will issue, and this seems to be a condition of civilized morality and personal sense of value. Illusion and ignorance appear to be conditions for social and personal success.

Saul Smilansky, *Free Will and Illusion*

INTRODUCTION

A number of philosophers have recently become increasingly interested in the potential usefulness of fictitious and illusory beliefs. As a result, a wide variety of fictionalisms and illusionisms have sprung up in areas ranging anywhere from mathematics and modality to morality.¹ In this paper, we focus on the view that Saul Smilansky has dubbed “free will illusionism”—for example, the purportedly descriptive claim that the majority of people have illusory beliefs concerning the existence of libertarian free will, coupled with the normative claim that because dispelling these illusions would produce negative personal and societal consequences, those of us who are unfortunate enough to know the dangerous and gloomy truth about the nonexistence of libertarian free will should simply keep quiet in the name of the common good.

1. See for example, Crimmins (1998), Evans (1982), Field (1980), Field (1989), Fine (1993), Joyce (2001), Joyce (2005), Rosen (1990), Rosen (1994), Van Fraassen (1980), Yablo (2000), and Yablo (2001). For an excellent review of recent developments in moral fictionalism, see Hussain (2004).

In section 1 of this paper, we provide a brief overview of Smilansky's illusionism, paying particular attention to two key empirical assumptions he makes concerning folk beliefs and intuitions. First, we examine the claim that the majority of people believe that humans have libertarian free will and ultimate moral responsibility. Second, we examine Smilansky's claim that if people came to realize that their beliefs about free will and moral responsibility are mistaken, the very moral ties that bind us together would be torn asunder. In section 2 we suggest that these two assumptions do not appear to be supported by the gathering data about folk intuitions, and we discuss why this is particularly problematic for Smilansky's view. We conclude that if future studies cast further doubt on the empirical foundation of Smilansky's illusionism, his view will be very difficult to motivate. After all, if it turns out either that (1) Smilansky is wrong to assume that most people have illusory beliefs about libertarian free will and ultimate responsibility; or (2) he is right that people do have these illusory beliefs, but wrong that they could not dispense with them without losing something very important, then it would no longer make sense to be a free will illusionist. It's not merely that Smilansky happens to use these assumptions as one of several possible ways of motivating his view; these assumptions appear to represent the *only way* of motivating his view.

1. FOLK INTUITIONS, SLIPPERY SLOPES, AND NECESSARY FICTIONS

Smilansky begins with a discussion of what he takes to be our ordinary moral practices—highlighting the central role played by libertarian notions of free will (LFW) and ultimate moral responsibility (UMR). From the outset, he claims to be interested primarily in the “nature of common beliefs” about free will and moral responsibility—beliefs that we are told are tacitly and predominantly libertarian (Smilansky 2000, 26). According to Smilansky, “most people not only believe in actual possibilities and the ability to transcend circumstances, but have distinct and strong beliefs that libertarian free will is a condition for moral responsibility, which is in turn a condition for just reward and punishment” (Smilansky 2000, 27).² For present purposes, we are going to call this closely related group of assumptions about folk intuitions and beliefs concerning the relationship between free will and moral responsibility The Ubiquity of Libertarianism Assumption (ULA).

According to ULA, our ordinary moral practices are built upon the conviction that “blameworthiness requires desert, desert requires responsibility, and responsibility requires control” (Smilansky 2000, 17). Smilansky calls the kind of control that is necessary for robust conceptions of moral blameworthiness and desert “up to usness,” and he suggests that in one very important respect nothing is ultimately up to us. Because he believes (1) that LFW requires that an agent's actions ultimately flow from who she is morally; and (2) that “what a person is,

2. Smilansky calls the principle that underlies this assumption about the relationship between LFW and moral responsibility The Principle of Sole Attribution, which he characterizes in the following way: “Any feature F due to which a person deserves something S in the libertarian free-will dependant sense must, in the normatively relevant respects, be solely attributable to the person or the pertinent aspect A of the person” (Smilansky 2000, 63).

morally, cannot ultimately be under her control” (Smilansky 2001, 74), Smilansky concludes both that LFW is nonexistent and that the notions of UMR which depend on its existence are therefore ungrounded.³

Trying to flesh out the implications of the nonexistence of LFW and UMR is Smilansky’s central task in *Free Will and Illusion*. On his view, there appear on the surface to be only two responses—namely, compatibilism and hard determinism. And while we are told that neither of these two responses to the nonexistence of LFW and UMR is wholly adequate, Smilansky nevertheless feels that each contains an important grain of truth. First, he points out that compatibilists are surely correct to highlight both the importance and possibility of people having “local reflective control over their actions” (Smilansky 2001, 77). After all, there is a morally significant difference between the kleptomaniac and the common thief—namely, the actions of the latter, unlike the actions of the former, are the result of deliberation and reflection (Smilansky 2001, 77). According to Smilansky, acknowledging the importance of this distinction is necessary if we are to foster a “Community of Responsibility”—for example, a community whereby people’s lives and actions are judged to be “based largely on their choices” (Smilansky 2001, 78).

The members of this compatibilistically grounded moral community are praiseworthy or blameworthy depending on whether they had the ability, capacity, and opportunity to deliberate, reflect upon, and control the consequences of their actions. Smilansky suggests that if we did not sometimes have the capacity to exercise compatibilistic control of our actions, none of our moral practices of blame and praise would be justified. The kleptomaniac and the common thief would end up on the same footing, morally speaking. It is for this reason that Smilansky thinks that compatibilists have contributed an important insight to the free will debate. As he says, “even without libertarian free will, it is reasonable to desire that compatibilist distinctions concerning control affect the way one is treated, and to see this as condition for civilized existence” (Smilansky 2001, 78).

Smilansky nevertheless thinks that compatibilism is hopelessly inadequate when it is judged from what he calls the “ultimate perspective” (Smilansky 2001, 77). On his view, “we can make sense of the notion of autonomy or self-determination on the compatibilist level but, if there is no libertarian free will, no one can be ultimately in control, ultimately responsible, for this self and its determinations” (Smilansky 2001, 75). Hence, even though our punishing an individual for her actions may be justified in compatibilist terms, we are nevertheless punishing her for “what is ultimately her luck, from what follows from who she is—ultimately beyond her control, a state which she had no real opportunity to alter, hence neither her responsibility nor her fault” (Smilansky 2001, 76). At this

3. For present purposes, we need not be concerned with Smilansky’s argument against LFW and UMR. It is worth pointing out that he does not merely suggest that we happen not to have LFW, rather he makes the much stronger claim that the kind of “up-to-usness” one finds in libertarian accounts of free will are incoherent. As he says, “the conditions required by an ethically satisfying sense of libertarian free will, which would give us anything beyond sophisticated formulations of compatibilism, are self-contradictory, and hence cannot be met. This is so irrespective of determinism or causality” (Smilansky 2001, 7). In this respect, Smilansky’s views concerning LFW share several affinities with those of Galen Strawson (1986).

point, we see that notions of compatibilist control, as important as they may be to our moral practices, are nevertheless swallowed up once we view ourselves from the perspective of *sub specie aeternitatis*. So, while we may justifiably be held responsible for our actions in one important sense, in another sense, who we are—and hence what we do—is ultimately little more than dumb luck.

Given that Smilansky thinks that we are not ultimately responsible for our actions, it would seem natural for him to simply follow the hard determinists in denying both the existence of free will and the existence of moral responsibility—a move he resists. For even though he believes that hard determinists are ultimately correct in saying that “we are not, in fact, able to get what we have reason to want” (Smilansky 2000, 2)—namely, LFW and UMR—Smilansky nevertheless believes that hard determinists are too quick to reject or undervalue the importance of compatibilist control for our everyday moral practices. Hence, we are told that if the perspective of the hard determinist were widespread among the masses:

Many people would find it hard to think that the partial compatibilist truth matters, as in fact it ethically does, if they realized the sense in which both the compatibilistically free and the unfree were merely performing according to their mould. And this might lead them to succumb to “pragmatic” consequentialist temptations, or an unprincipled nihilism. The ultimate hard determinist perspective does not leave sufficient moral and psychological “space” for compatibilistically defensible reactive attitudes and moral order. The fragile compatibilist-level plants need to be defended from the chill of the ultimate perspective in the hothouse of illusion. Only if we do not see people from the ultimate perspective can we live in a way which compatibilism affirms—blaming, selectively excusing, respecting, being grateful, and the like. (Smilansky 2001, 89–90)

According to this line of reasoning, if people were to lose their faith in the existence in LFW and UMR, not only would society lose its moral compass, but individuals would lose their sense of self-worth and their perceived accountability (Smilansky 2001, 85–86). Consequently, Smilansky claims that “the difficulties caused by the absence of ultimate-level grounding are likely to be great, generating acute psychological discomfort for many people and threatening morality—if, that is, we do not have illusion at our disposal” (Smilansky 2001, 87). For present purposes, we are going to call this The Disutility of Disillusionment Assumption (DDA)—which is essentially the aforementioned ULA coupled both with the denial of LFW and with a prediction that if people were to adopt the perspective of the hard determinist, the moral fabric that holds us together would be torn asunder. On this view, the aforementioned Community of Responsibility could not weather the storm that would result from the wide-scale internalization of the ultimate perspective.

If Smilansky were correct in assuming that (1) the majority of people falsely believe that we have LFW and that LFW is necessary for moral responsibility and that (2) availing people of these false beliefs would produce deleterious social and personal consequences, we would indeed find ourselves in quite a predicament. So

what does Smilansky suggest we do? The short answer is—absolutely nothing. According to Smilansky’s free will illusionism, because the benefits of wide-scale illusory beliefs about the existence of LFW and UMR far outweigh the costs associated with dispelling these beliefs, we should not let the hard determinist cat out of the proverbial bag. As he says, “To put it bluntly: people as a rule ought not to be fully aware of the ultimate inevitability of what they have done, for this will affect the way in which they hold themselves responsible” (Smilansky 2001, 85).

Smilansky’s free will illusionism has both a descriptive and a normative element. Descriptively, he believes that illusory beliefs in LFW and UMR are widespread and that these beliefs have a number of positive effects on society in general as well as on the particular individuals who hold them. Normatively, he believes that given that the closely related network of illusory beliefs concerning LFW and UMR⁴ is crucial for the well-being of both the individual and the community at large, we should leave these beliefs in place even though they are false and confused (Smilansky 2000, 150).

Now that we have a clear idea of what Smilansky’s free will illusionism amounts to, as well as a better understanding of what kinds of empirical assumptions are necessary for motivating his view, we want to turn our attention to some recent studies that give us reason to doubt whether some of his key assumptions about folk intuitions are actually true. And for reasons that we discuss in the upcoming section, we believe that this lack of empirical support spells serious trouble for Smilansky’s entire project. If it turns out that people don’t have illusory beliefs about LFW and UMR, or if they do have illusory beliefs but they could be dispensed with without much ado, then it is unclear why anyone would want to be an illusionist. But we are getting ahead of ourselves. Before we can properly flesh out these kinds of issues, we must first examine the salient data on folk intuitions concerning free will and moral responsibility.

2. HOLDING SMILANSKY’S FEET TO THE EMPIRICAL FIRE

In a series of recent experiments, Eddy Nahmias, Stephen Morris, Thomas Nadelhoffer, and Jason Turner probed folk intuitions concerning free will and moral responsibility with methods borrowed from social psychology (Nahmias et al. 2005, 2006).⁵ One of the stated goals of these studies was to determine in a

4. Smilansky identifies eight distinct but closely related kinds of illusory beliefs about LFW and UMR. The following six are particularly germane to our present discussion: (1) the belief that there is LFW; (2) the belief that there is no threat to LFW; (3) the belief in UMR and full-blown retributivist desert; (4) the belief that we have more than merely compatibilist control over our decisions and actions; (5) the belief that we live in a just world; and (6) the belief that we are living with illusions concerning free will (Smilansky 2000, 215–16).

5. There are a growing number of philosophers in a variety of fields working under the rubric of “experimental philosophy”—one of the stated goals of which is to increase our understanding of folk concepts and intuitions concerning issues that are especially relevant to philosophy. Hence, philosophers have begun engaging in experimental research and contributing to the growing data about folk concepts and intuitions relevant to epistemology (e.g., Nichols, Weinberg, and Stich 2003; Weinberg, Nichols, and Stich 2001), ethics (e.g., Doris and Stich 2003; Nichols 2004a), free will (e.g., Nahmias et al. 2005, 2006; Nichols 2004b), and the philosophy of action (Knobe 2003a, 2003b,

controlled and systematic manner whether philosophers have been correct to assume that “we come to the table, nearly all of us, as pre-theoretical incompatibilists” (Ekstrom 2002, 310).⁶ After all, if it were true that most nonphilosophers do indeed have what Derk Pereboom has called the “incompatibilist intuition” (Pereboom 2001, 89), then we should expect that if they were given a thought experiment involving an (otherwise ordinary) agent in a deterministic scenario, a majority of them would *not* attribute free will and moral responsibility to the agent. In one of the first attempts to get at the relevant data about folk intuitions concerning free will and moral responsibility, Nahmias et al. ran several studies that presented participants with deterministic scenarios to see whether their intuitions were in line with what philosophers often say about these intuitions.

For present purposes, we are only going to discuss one of their studies that used a vignette involving a Laplacean notion of determinism. Participants in this study were undergraduates who had not studied the free will problem—each of whom read the following scenario and answered two questions about it (then, on the back of the questionnaire, they responded to a manipulation check, were invited to explain their answer, and were offered some demographic information):

Scenario: Imagine that in the next century we discover all the laws of nature, and we build a supercomputer which can deduce from these laws of nature and from the current state of everything in the world exactly what will be happening in the world at any future time. It can look at everything about the way the world is and predict everything about how it will be with one hundred percent accuracy. Suppose that such a supercomputer existed, and it looks at the state of the universe at a certain time on March 25th, 2150 A.D., twenty years before Jeremy Hall is born. The computer then deduces from this information and the laws of nature that Jeremy will definitely rob Fidelity Bank at 6:00 PM on January 26th, 2195. As always, the supercomputer's prediction is correct; Jeremy robs Fidelity Bank at 6:00 PM on January 26th, 2195.

2004; Nadelhoffer 2004a, 2004b, 2004c, 2005, 2006a, 2006b). For a concise introduction to this developing field, see Nichols (2004c).

6. This assumption is particularly common among incompatibilists. In addition to the aforementioned comment from Laura Ekstrom, consider Robert Kane's (1999), claim that “most ordinary people start out as natural incompatibilists . . . [so that] ordinary persons have to be talked out of this natural incompatibilism by the clever arguments of philosophers” (218). Similarly, Galen Strawson (1986) contends that the incompatibilist's libertarian conception of free will, though impossible to satisfy, is precisely “the kind of freedom that most people ordinarily and unreflectively suppose themselves to possess” (30), adding that it is “in our nature to take determinism to pose a serious problem for our notions of responsibility and freedom” (89). And Thomas Pink (2004) tells us that “most of us start off by making an important assumption about freedom. Our freedom of action, we naturally tend to assume, must be incompatible with our actions being determined” (12). In this respect, Smilansky's ULA is simply one example among several where a philosopher assumes that ordinary people are committed—either explicitly or tacitly—to incompatibilism.

One group of participants was asked whether Jeremy “robbed the bank of his own free will” and a second group was asked whether Jeremy was blameworthy for robbing the bank. Surprisingly, a majority of participants in the first group judged that Jeremy robbed the bank of his own free will (seventy-six percent) and a majority in the second group judged that he was blameworthy for having done so (eighty-three percent).⁷

The data from these preliminary studies suggest that ordinary people’s pretheoretical intuitions about free will and responsibility may not be as incompatibilistic as philosophers have assumed. Keep in mind that when participants considered an agent (Jeremy) whose actions were unerringly predicted based on the state of the universe before his birth and the laws of nature, a significant majority nevertheless judged that the agent acted of his own free will and is morally responsible for the action—for example, most participants did *not* recognize an incompatibility between determinism described in these terms and free will or moral responsibility. Hence, the results of the research carried out by Nahmias et al. give us *prima facie* reason to be skeptical concerning the veracity of Smilansky’s aforementioned ULA. After all, Smilansky assumes that most people believe that determinism and free will are incompatible—an assumption that does not settle with the preliminary data.

Of course, one way that Smilansky might respond is to suggest that Nahmias et al. did not make determinism robust or clear enough in their earlier Jeremy studies. In our present attempt to be as charitable to Smilansky as possible, we decided simply to ask participants point blank about the compatibility of free will and determinism. Participants in this new survey were 105 undergraduates at Florida State University who had never studied the free will debate. Each participant received the following question:

Do you think that our actions can be free if all of them are *entirely* determined by our genes, our neurophysiology, and our upbringing?

Even though we used the phrase “*entirely* determined”—something Nahmias et al. explicitly tried to avoid in their earlier studies—forty-four out of the 105 participants (forty-two percent) nevertheless judged that our actions can be free under these conditions.

The first thing worth pointing out is that participants in this study were less likely to give compatibilist-friendly responses to our definitional question than participants in the earlier Nahmias et al. studies. Presumably, the fact that we used the word “determinism,” rather than simply building it tacitly into the vignette as was previously done, partly explains these intuitional differences. There are also likely underlying cognitive differences at work having to do with the difference between affectively “hot” or “cold” scenarios—an interesting likelihood already

7. See Nahmias et al. 2005 and 2006 for further details concerning both the experimental design and the results of their studies.

being explored by Shaun Nichols and Joshua Knobe.⁸ For present purposes, however, we are going to set these kinds of issues aside. Our goal in this paper is not to explain why participants give the answers they do, but rather to flesh out what kind of evidence Smilansky would need to motivate his view. And by our lights, he needs something like ULA to be true. Yet when we explicitly asked participants whether we could be free if our actions were entirely determined, slightly less than half of them were willing to say *yes*.⁹

Minimally, we believe that Smilansky owes an explanation for why fewer than half of the people we asked judged that we could be free even if our actions were entirely determined. And if future studies provide additional confirmation of these results, we will have good reason to conclude that the belief in LFW is not nearly as widespread as Smilansky assumes. At best, the gathering data suggest that Smilansky is only entitled to a much weaker claim—namely, the Bare Majority Libertarian Assumption (BMLA). And given that BMLA is unlikely to be sufficient for motivating Smilansky's free will illusionism, we should reject illusionism until the stronger ULA is established by empirical testing.

For the time being, however, we would like to set aside our worries about the truth of ULA and focus instead on the truth of DDA—which is admittedly a much harder assumption to test since it involves a prediction concerning how people *would* react if they found out that their belief in the existence of LFW were an illusion. One simple way of testing DDA would involve asking people what they thought about Smilansky's worry. For instance, we could ask people who do believe in LFW and UMR what they believe would happen both socially and personally if

8. In a series of interesting experiments, Shaun Nichols and Joshua Knobe (forthcoming) have recently generated similar results. Rather than simply investigating what people's intuitions about free will, determinism, and moral responsibility actually are, Nichols and Knobe investigate the possible cognitive sources of these intuitions. More specifically, they claim that the results of their new studies show that "people have an incompatibilist theory of moral responsibility that is elicited in some contexts but that they also have psychological mechanisms that can lead them to arrive at compatibilist judgments in other contexts." While we find the general descriptivist and cognitivist approach of Nichols and Knobe to be both novel and promising, we believe their experimental design is problematic. For present purposes, it's worth pointing out that even if Nichols and Knobe are correct that people are incompatibilist in theory but compatibilist at heart, this would not help Smilansky. After all, if people end up being compatibilist when it comes to free will and blame attribution in moral cases, then it is unclear why determinism would undermine the so-called community of responsibility that Smilansky claims would be jeopardized if people came to believe determinism were true. But since we think there are some methodological problems with the studies run by Nichols and Knobe in the first place, we aren't sure that their data counts as evidence for ULA or DDA one way or the other. This is an issue that we are currently in the process of exploring.

9. By our lights, the kind of abstract question we asked participants represents a very helpful way of probing folk intuitions concerning determinism, free will, and moral responsibility that does not run the risk of affectively biasing participants' intuitions. So, regardless of whether our wording in terms of "our genes, our neurophysiology, and our upbringing" adequately captures all of the salient underlying issues, we nevertheless believe our general strategy is likely to generate more reliable data than strategies involving morally laden scenarios with bad actors and sometimes horrible consequences. But that is another story for another day. For now, we want to refocus our attention back to the specific issue at hand—namely, whether Smilansky's ULA and DDA enjoy empirical support.

they came to believe that LFW and UMR are actually nonexistent for the reasons Smilansky gives. The goal would be to determine whether people think that if they ceased to believe that they are ultimately free and responsible, they (1) would be more inclined to behave immorally; and (2) would be more inclined to adopt a pessimistic stance toward themselves and their place in the cosmos.¹⁰

Unfortunately, despite the initial allure of running this kind of a study to test Smilansky's DDA, there are several reasons to suspect that simply asking people what they would do if they came to abandon some of their most fundamental beliefs would produce unreliable data. One of the primary shortcomings of this kind of study is that we have good reason to suspect that people will not be good judges of how they would behave if they no longer believed in LFW and UMR. Consider, for instance, the results of one of the more famous studies that have been conducted concerning people's ability to make accurate predictions concerning their future behavior—namely, the shock studies by Stanley Milgram (1963). Prior to these now infamous experiments, Milgram polled psychologists, middle-class adults, and students to see what they thought would be the maximum shock administered by the participants in the experiment. Their predictions ranged anywhere from 195 volts to 300 volts. In Milgram's original experiments, however, sixty-five percent of the participants administered what they believed to be the maximum (and potentially lethal) shock of 450 volts! Indeed, not a single

10. So, that's precisely what we did. In one study, participants were 80 Florida State University undergraduates—each of whom read one of the following two scenarios:

First-person vignette:

Now, for the sake of argument, assume that in the future scientists discover that all of our beliefs, desires, choices, decisions, and actions are *entirely* determined by our genes, our neurophysiology, and our upbringing. Now assume that this scientific discovery leads you to conclude that humans are neither free nor ultimately morally responsible for their actions. Would you be more inclined to behave immorally in light of this knowledge?

Third-person vignette:

Now, for the sake of argument, assume that in the future scientists discover that all of our beliefs, desires, choices, decisions, and actions are *entirely* determined by our genes, our neurophysiology, and our upbringing. Now assume that this scientific discovery leads people to conclude that humans are neither free nor ultimately morally responsible for their actions. Do you think they would be more inclined to behave immorally in light of this knowledge?

The participants' responses to the respective questions in the first- and third-person vignettes were as follows: Whereas only thirty-five percent of the participants who read the first-person case believed they would be more inclined to behave immorally in light of this knowledge, seventy percent of the participants in the third-person vignette believed that others would be more inclined to behave immorally if they came to believe that all of their actions were entirely determined. These results give us reason to doubt the reliability of simply asking someone whether they think they would be more inclined to engage in immoral activity if they no longer believed in LFW. Indeed, the results of the third-person case are likely a more reliable indicator of whether people share Smilansky's belief in DDA. It turns out that people either don't believe—or don't admit—that *they* would be more inclined to behave immorally if *they* no longer believed they had free will, even though they also do believe that *others* would be so inclined. Hence, it turns out that a majority of people may share Smilansky's belief concerning the disutility of disillusionment.

participant stopped before the 300-volt level. These results revealed that the degree of obedience among participants was far higher than people had predicted.

Presumably, had people been asked to predict how *they themselves* would behave rather than being asked how *other people* would behave, their predictions would have been even less reliable. Indeed, subsequent studies have cast further doubt on the reliability of our predictions concerning future behaviors—especially when the circumstances surrounding these behaviors are unfamiliar.¹¹ Hence, we have good reason to suspect that simply asking people to predict how they would act and feel if they discovered that their beliefs concerning free will and moral responsibility are based on little more than illusion would be an insufficient way of testing Smilansky's DDA.

These methodological shortcomings notwithstanding, there is an even more basic reason for not simply asking people what they believe they would do and feel if they no longer believed in LFW and UMR—namely, that this is not the kind of data we would need to determine the merits of Smilansky's DDA. After all, DDA, unlike ULA, is not merely an assumption about what people *believe* would happen in the future if they no longer believed in LFW and UMR, it is a prediction about what would *actually* happen if people no longer had either of these beliefs. In order to test this prediction, we would first need to somehow convince people that they do not have free will and then check to see whether this has the negative effect that Smilansky predicts. Simply asking them what they believe might happen is not going to do the trick.

Arguably the best way to test Smilansky's predictions concerning what would happen if people no longer believed in LFW and UMR would involve identifying and then interviewing people who once believed in LFW and UMR but no longer do. We could then test whether they feel that their lives lack significance. Moreover, because these individuals would already have come to grips with the lack of LFW and UMR, they would not have to speculate concerning how they would react if they came to such a realization—thus, avoiding the major difficulty of asking people in general how they would react if they came to realize that their beliefs in LFW and UMR are illusory. If these people do not experience any serious deleterious effects, then that is even more evidence against Smilansky's claim that such a realization would be seriously debilitating.¹²

In thinking about how we might design such a study, we suddenly realized that Smilansky's free will illusionism may depend on the truth of some additional questionable assumptions. After all, Smilansky himself is presumably perfectly able

11. For example, Dan Gilbert and his colleagues have shown that in forecasting we often overestimate the effect that negative events will have on our lives. This is probably because we forget to take into account a type of “psychological immune system” that kicks in to soften the effects of the negative events when bad things happen to us. When we forecast, we do not take into account the immune system, and our predictions suffer from “immune neglect.” The immune neglect has the consequence that we overestimate the duration of the negative event's effect; hence we generate a durability bias. Since the immune system only protects the agent against negative effects, it does not augment predictions of positive effects, and therefore does not generate a positive durability bias (Gilbert et al. 2002).

12. We are currently designing some new studies that are intended to help us get at precisely these kind of data.

to live a morally acceptable and personally gratifying life, filled with meaningful choices and loving relationships. So, if disillusionment did not have a deleterious effect on his moral compass and sense of self-worth, why should we assume that everyone else would suffer in a way that Smilansky has not? The same question can be asked regarding most other free will skeptics as well—many of whom are our friends and colleagues. If some philosophers are able to find meaning and purpose in their lives despite losing faith in LFW and UMR, why couldn't everyone else?

Perhaps it is simply because Smilansky believes that the *hoi polloi*, unlike professional philosophers, are not intellectually equipped to appreciate the small grain of truth contained in compatibilistically grounded notions of control and responsibility. Indeed, this is precisely what he seems to have in mind when he says that “reactions and practices which are at least partially valid (have compatibilist grounding) will not be sufficiently adhered to if the absence of libertarian free will is realized” (Smilansky 2001, 87). Hence, because Smilansky believes that the masses would likely throw out the compatibilist baby with the libertarian bathwater, he concludes that those of us who are “in the know” should leave others to their illusory beliefs about free will and determinism. Of course, this too is an empirical claim about people's general inability to appreciate the importance of compatibilist control and responsibility once they have internalized the cold truth about the nonexistence of LFW and UMR.

By our lights, until a concerted effort is made to further test the network of empirical assumptions Smilansky needs in order to be able to motivate his view, we should refrain from adopting free will illusionism. On the one hand, we *cannot* leave people to illusory beliefs they do not have in the first place. On the other hand, even if people do actually have illusory beliefs about things as central to their lives as LFW and UMR, we should only keep the truth hidden from them if we have good reason to suspect that the benefits of doing so outweigh the costs. As we are concerned, simply assuming *tout court* that people are either too unintelligent or weak-willed to handle the truth is not a good enough reason. Of course, if it turns out that future studies show that Smilansky's assumptions are correct after all, then we will be forced to reconsider his view anew. Until then, we believe we have adequate evidence to suspect that illusionism is too problematic to be accepted.¹³

REFERENCES

- Crimmins, M. 1998. “Hesperus and Phosphorus: Sense, Pretense, and Reference.” *Philosophical Review* 107(1): 1–47.
- Doris, J., and Stich, S. 2003. “As a Matter of Fact: Empirical Perspectives on Ethics.” In *The Oxford Handbook of Contemporary Analytic Philosophy*, ed. F. Jackson and M. Smith, 114–54. Oxford: Oxford University Press.
- Ekstrom, L. W. 2002. “Libertarianism and Frankfurt-style Cases.” In *The Oxford Handbook of Free Will*, ed. R. Kane, 309–22. New York: Oxford University Press.
- Evans, G. 1982. *The Varieties of Reference*. Oxford: Clarendon Press.

13. We would like to thank Joshua Knobe, Alfred Mele, Eddy Nahmias, Shaun Nichols, Saul Smilansky, Tamler Sommers, and the other participants at the INPC in 2006 at Washington State University who played along, for their helpful comments and criticisms.

- Field, H. 1980. *Science Without Numbers: A Defense of Nominalism*. Princeton, NJ: Princeton University Press.
- . 1989. *Realism, Mathematics, and Modality*. Oxford: Blackwell.
- Fine, A. 1993. "Fictionalism." *Midwest Studies in Philosophy* 18: 1–18.
- Gilbert, D. T., Pinel, E. C., Wilson, T. D., Blumberg, S. J., and Wheatley, T. P. 2002. "Durability Bias in Affective Forecasting." In *Heuristics and Biases: The Psychology of Intuitive Judgment*, ed. Thomas Gilovich, Dale Griffin, and Daniel Kahneman, 292–312. Cambridge: Cambridge University Press.
- Hussain, N. 2004. "The Return of Moral Fictionalism." *Philosophical Perspectives* 18: 150–87.
- Joyce, R. 2001. *The Myth of Morality*. New York: Cambridge University Press.
- . 2005. "Moral Fictionalism." In *Fictionalist Approaches to Metaphysics*, ed. Mark Kalderon, 287–313. Oxford: Oxford University Press.
- Kane, R. 1999. "Responsibility, Luck, and Chance: Reflections on Free Will and Indeterminism." *Journal of Philosophy* 96(5): 217–40.
- Knobe, J. 2003a. "Intentional Action and Side-Effects in Ordinary Language." *Analysis* 63: 190–93.
- . 2003b. "Intentional Action in Folk Psychology: An Experimental Investigation." *Philosophical Psychology* 16: 309–24.
- . 2004. "Intention, Intentional Action and Moral Considerations." *Analysis* 64: 181–87.
- Milgram, S. 1963. "Behavioral Study of Obedience." *Journal of Abnormal and Social Psychology* 67: 371–78.
- Nadelhoffer, T. 2004a. "The Butler Problem Revisited." *Analysis* 64: 277–84.
- . 2004b. "Praise, Side Effects, and Intentional Action." *The Journal of Theoretical and Philosophical Psychology* 24: 196–213.
- . 2004c. "Blame, Badness, and Intentional Action: A Reply to Knobe and Mendlow." *The Journal of Theoretical and Philosophical Psychology* 24: 259–69.
- . 2005. "Skill, Luck, and Intentional Action." *Philosophical Psychology* 18(3): 343–54.
- . 2006a. "Bad Acts, Blameworthy Agents, and Intentional Actions: Some Problems for Jury Impartiality." *Philosophical Explorations* 9(2): 203–20.
- . 2006b. "On Trying to Save the Simple View." *Mind & Language* 21(5): 565–86.
- Nahmias, E., Morris, S., Nadelhoffer, T., and Turner, J. 2005. "Surveying Free Will." *Philosophical Psychology* 18(5): 561–84.
- . 2006. "Is Incompatibilism Intuitive?" *Philosophy and Phenomenological Research* 73(1): 28–53.
- Nichols, Shaun. 2004a. "After Objectivity." *Philosophical Psychology* 17: 5–28.
- . 2004b. "The Folk Psychology of Free Will: Fits and Starts." *Mind & Language* 19: 473–502.
- . 2004c. "Folk Concepts and Intuitions: From Philosophy to Cognitive Science." *Trends in Cognitive Science* 8: 514–18.
- Nichols, S., and Knobe, J. Forthcoming. "Moral Responsibility and Determinism: The Cognitive Science of Folk Intuition." *Nous*.
- Nichols, S., Weinberg, J., and Stich, S. 2003. "Metaskepticism: Meditations in Ethno-Epistemology." In *The Sceptics*, ed. S. Luper, 227–47. Burlington, VT: Ashgate.
- Pereboom, Derk. 2001. *Living Without Free Will*. New York: Cambridge University Press.
- Pink, T. 2004. *Free Will: A Very Short Introduction*. New York: Oxford University Press.
- Rosen, G. 1990. "Modal Fictionalism." *Mind* 99: 327–54.
- . 1994. "What Is Constructive Empiricism?" *Philosophical Studies* 74(2): 143–78.
- Smilansky, S. 2000. *Free Will and Illusion*. Oxford: Oxford University Press.
- . 2001. "Free Will: From Nature to Illusion." *Proceedings of the Aristotelian Society* 101: 71–95.
- Strawson, G. 1986. *Freedom and Belief*. Oxford: Oxford University Press.
- Van Fraassen, B. 1980. *The Scientific Image*. Oxford: Clarendon Press.
- Weinberg, J., Nichols, S., and Stich, S. 2001. "Normativity and Epistemic Intuitions." *Philosophical Topics* 29: 429–60.
- Yablo, S. 2000. "A Paradox of Existence." In *Empty Names, Fiction, and the Puzzles of Non-Existence*, ed. A. Everett and T. Hofweber, 275–312. Stanford, CA: CSLI Publications.
- . 2001. "Go Figure: A Path Through Fictionalism." *Midwest Studies in Philosophy* 25: 72–102.