



ELSEVIER

Contents lists available at ScienceDirect

Consciousness and Cognition

journal homepage: www.elsevier.com/locate/concog

The free will inventory: Measuring beliefs about agency and responsibility


 Thomas Nadelhoffer ^{a,*}, Jason Shepard ^b, Eddy Nahmias ^c, Chandra Sripada ^d, Lisa Thomson Ross ^e
^a Department of Philosophy, College of Charleston, 66 George Street, Charleston, SC 29424, United States

^b Department of Psychology, Emory University, 36 Eagle Row, Atlanta, GA 30322, United States

^c Department of Philosophy, Neuroscience Institute, Georgia State University, P.O. Box 3994, Atlanta, GA 30302-3994, United States

^d Department of Philosophy, The University of Michigan, 2215 Angell Hall, 435 South State Street, Ann Arbor, MI 48109, United States

^e Department of Psychology, College of Charleston, 66 George Street, Charleston, SC 29424, United States

ARTICLE INFO

Article history:

Received 30 October 2013

Keywords:

 Free will
 Determinism
 Dualism
 Soul
 Responsibility
 Punishment
 Psychometrics
 Scale
 Folk intuitions

ABSTRACT

In this paper, we present the results of the construction and validation of a new psychometric tool for measuring beliefs about free will and related concepts: The Free Will Inventory (FWI). In its final form, FWI is a 29-item instrument with two parts. Part 1 consists of three 5-item subscales designed to measure strength of belief in free will, determinism, and dualism. Part 2 consists of a series of fourteen statements designed to further explore the complex network of people's associated beliefs and attitudes about free will, determinism, choice, the soul, predictability, responsibility, and punishment. Having presented the construction and validation of FWI, we discuss several ways that it could be used in future research, highlight some as yet unanswered questions that are ripe for interdisciplinary investigation, and encourage researchers to join us in our efforts to answer these questions.

© 2014 Elsevier Inc. All rights reserved.

1. Introduction

Philosophers have long debated the nature of free will and moral responsibility. Psychologists have more recently joined the fray. While some researchers have tried to shed empirical light on the debate concerning whether we have free will, typically concluding that we do not (e.g., Bargh, 2008; Harris, 2012; Libet, 1999; Wegner, 2002), others have focused instead on exploring our *beliefs* about free will and the role they play in our daily lives (e.g., Alquist, Ainsworth, & Baumeister, 2013; Bergner & Ramon, 2013; Haynes, Rojas, & Viney, 2003; Nettle, 1959; Nettle, 1961; Ogletree & Oberle, 2008; Pronin & Kugler, 2010; Schooler, Nadelhoffer, Nahmias, & Vohs, in press; Stillman & Baumeister, 2010; Stillman et al., 2010; Stroessner & Green, 1990; Viney, Parker-Martin, & Dotten, 1988; Viney, Waldman, & Barchilon, 1982; Waldman, Viney, Bell, Bennett, & Hess, 1983). Each of these two projects is interesting and important in its own right. For present purposes, we focus on the psychology of believing in free will—a topic that has received increasing attention in the wake of gathering evidence that challenging people's beliefs about free will may influence their behavior in surprising and sometimes alarming ways. For example, telling people they do not have free will has been shown to increase cheating (Vohs & Schooler, 2008) decrease helping behavior and increases aggression (Baumeister, Masicampo, & DeWall, 2009), reduce self-control (Rigoni, Kuhn,

* Corresponding author. Fax: +1 843 953 6388.

E-mail addresses: nadelhofferta@cofc.edu (T. Nadelhoffer), jason.s.shepard@emory.edu (J. Shepard), enahmias@gsu.edu (E. Nahmias), sripida@umich.edu (C. Sripada), rossl@cofc.edu (L.T. Ross).

Gaudino, Sartori, & Brass, 2012), and impair detection of errors (Rigoni, Wilquin, Brass, & Burle, 2013). These findings suggest that *believing* in free will may be instrumentally valuable from the standpoints of positive psychology and public morality. This is true regardless of whether or not we *actually* have free will, a related but orthogonal issue that we will not be exploring here.

These recent findings also highlight the importance of having valid and reliable tools for measuring beliefs about free will and related concepts. For while the gathering data suggest that diminishing people's belief in free will may lead to increased cheating, decreased helping, lower punishment judgments, and the like, the validity of these findings depends in part on the validity and reliability of the scales used to measure people's beliefs about free will, determinism, and responsibility. As such, these otherwise exciting findings have highlighted a problem that has hindered empirical research on free will beliefs from the outset—namely, that while researchers have developed several scales during the past thirty years for measuring beliefs about free will (e.g., Paulhus & Carey, 2011; Paulhus & Margesson, 1994; Rakos, Laurene, Skala, & Slane, 2008; Stroessner & Green, 1990; Viney, McIntyre, & Viney, 1984; Viney et al., 1982), each of these tools is problematic. The main goal of our project was to develop a new psychometric instrument that would have greater validity than previous scales while at the same time yielding more fine-grained data concerning how people think about the complex relationships among free will, responsibility, dualism, choice, determinism, and related concepts.

In this paper, we first discuss the extant tools that have been developed for measuring free will beliefs, and we identify some of their limitations. Then, we present the results of the construction and validation of a new psychometric tool for measuring beliefs about free will: The Free Will Inventory (FWI). The FWI is a 29-item instrument with two parts (which can be given together or separately). Part 1 consists of three five-item subscales designed to measure strength of belief in free will, determinism, and dualism. Part 2 consists of a series of 14 statements designed to further explore people's associated beliefs and attitudes about free will, (in)determinism, choice, the soul, predictability, responsibility, and punishment. After presenting the construction and validation of FWI, we (a) discuss several ways that it could be used in future research, (b) highlight some as yet unanswered questions that are ripe for interdisciplinary investigation, and (c) encourage researchers to join us in our efforts to answer these questions.

1.1. Measuring beliefs about free will: An overview

Most of the earliest work on the psychology of believing in free will focused on the relationship between beliefs about free will, determinism, and punishment (e.g., Nettler, 1959; Nettler, 1961; Stroessner & Green, 1990; Viney et al., 1982; Viney et al., 1988). Though groundbreaking, the results were often mixed and hard to interpret. For instance, in one of the first studies on free will beliefs, Nettler (1959) claimed to have found that believing in free will is correlated with cruelty, retribution, and revenge and that believing in determinism is correlated with people being less punitive and treating others with more dignity. In response, Viney et al. (1982) first reported data that seemed to support the opposite conclusion—namely, that people who believe in free will are less rather than more punitive—before later finding “neither reliable correlations between punitiveness and beliefs in free will or determinism nor reliable correlations between rationales for punishment and beliefs in free will or determinism” (Viney et al., 1988, p. 20).

Looking back on the early empirical work on free will beliefs with the benefit of hindsight, there are at least two lessons to learn. First, what researchers find when it comes to beliefs about free will can depend a great deal on how the researchers conceptualize and think about free will. Not only can researchers' own free will beliefs influence how they design their experiments, phrase their questions, etc., but these beliefs may also color how the findings are analyzed and interpreted. Let us call this the *problem of theory contamination*—i.e., when researchers' own theoretical commitments unduly influence or bias their findings. While it may be impossible to keep one's theoretical commitments *entirely* at bay while designing studies and analyzing data, it is important for researchers to make a concerted effort to avoid theory contamination as much as possible. One strategy that we adopted was making sure we had both psychologists and philosophers on our team, as well as people with competing views about free will and its relationship to determinism.

A second lesson to be learned from the early research on free will beliefs is that researchers should have a shared vocabulary (or lexicon) when exploring these types of complex beliefs and attitudes. Instead, much of the early empirical work in this area contains conflicting definitions of key terms (or usages of key terms) such as determinism, libertarian free will, choice, dualism, fatalism, and the like. Moreover, not only are key terms used in ways that differ from one researcher to the next, these terms are often defined in ways that most philosophers who specialize in the free will debate would reject. This generates worries about both content and criterion validity. While philosophical experts on free will should not have *carte blanche* to define contested terminology however they see fit, the long and deep engagement of philosophers with the question of free will should be taken into consideration during empirical investigations of free will beliefs.

At least some of these methodological worries might be alleviated if psychologists and philosophers worked together rather than laboring individually on their own respective sides of the disciplinary divide. For while psychologists do not appear to have a shared lexicon when it comes to free will and related concepts, philosophers have developed a common vocabulary for talking about these issues. Take, for instance, the concept of *determinism*. While the term has a variety of meanings and uses (both in everyday life and amongst psychologists), philosophers typically have one thing in mind when they discuss determinism—namely, the thesis that given the actual past and the laws of nature, there is only one possible future at any moment in time (e.g., van Inwagen, 1983).

Using this definition of determinism, the dominant issue in the philosophical debate about free will has traditionally been whether free will and moral responsibility are compatible with determinism. The attempt to address this so-called “compatibility question” has generated a wide variety of positions. *Incompatibilists* claim that free will and determinism cannot coexist. Among incompatibilists, *libertarians* argue that humans have free will, and hence deny the truth of determinism, some arguing that indeterminism helps secure free will (e.g., Kane, 1996), others arguing that free will requires a special causal power, agent-causation (e.g., O’Connor, 2000). Other incompatibilists argue that we lack free will, either because determinism is true (*hard determinists*) or, more commonly, because they argue that free will and moral responsibility are incompatible with indeterminism as well as determinism (*hard incompatibilists* or *skeptics about free will*; e.g., Pereboom, 2001; Strawson, 1994).

However, a long philosophical tradition has defended *compatibilism* about free will and determinism. Compatibilists offer various analyses of the capacities required to be free and responsible agents, such as our abilities to identify with some of our desires over others (Frankfurt, 1988), to act in accord with moral reasons (Wolf, 2000), to be appropriate targets of reactive attitudes such as indignation or approbation (Strawson, 1962), or to be appropriately responsive to reasons (Fischer & Ravizza, 1998). In general, compatibilists argue that free will and moral responsibility do not require the unconditional ability to do otherwise, holding fixed the actual past and laws, nor the seemingly impossible ability to be ultimately responsible for what makes us the way we are (see Strawson, 1994). Instead, compatibilists argue that free and responsible agency requires the capacities involved in self-reflection and practical deliberation; free will is the ability to make choices based on reasons, along with the opportunity to exercise this ability without undue constraints.

Despite the fact that compatibilism is the majority position among philosophers (Bourget & Chalmers, *in press*) and despite the shared vocabulary that philosophers use to discuss free will, psychologists have often operationalized free will and determinism in ways that are not only discordant with the philosophical tradition but which also preclude a combination of beliefs that is not only logically possible but also commonly held (at least by certain groups of people). This has led to some conflicting research programs and unclear results. To see what we have in mind, consider the earliest attempt to develop a scale to measure beliefs about free will, the Free Will and Determinism Scale (FWD) presented in Viney et al. (1982). The general instructions for FWD read as follows:

For centuries, human beings have debated the old philosophical problem of free will and determinism. Free will is the doctrine that assumes that we make real choices which are partially or completely independent of antecedent conditions. Determinism denies this and maintains that causation is operative in all human affairs. Thus, so-called choices are influenced or determined by antecedent conditions. The following materials include statements which represent opinions by those who believe in free will and those who believe in determinism. Please read each of the following statements carefully and then place an x under the statement which most closely corresponds to your opinion on the topic. If your opinion seems to fall somewhere between the statements provided, then mark an x in one of the intermediate spaces.

These instructions highlight the aforementioned worries about theory contamination.

After all, it is clear that Viney et al. (1982) have simply built the assumption of incompatibilism into FWD from the outset, in part by defining ‘determinism’ in contrast with ‘free will’. As such, there is no way to use FWD to tease apart individuals who have incompatibilist intuitions from those who have compatibilist intuitions. Indeed, at some point, Viney and colleagues invite participants to choose between the following two statements: (a) “I believe strongly in free will,” and (b) “I believe strongly in determinism.”

Given the way that FWD is constructed, individuals who do not take determinism to rule out free will have no way to express that belief. Furthermore, if someone expresses belief in free will but also expresses the belief that determinism is false (e.g., because of quantum theory), that does not indicate that she is an *incompatibilist*, since she might think the truth or falsity of determinism is not relevant to whether we have free will (as do most contemporary compatibilist philosophers). The worry is not only theoretical—there is accumulating evidence that compatibilist intuitions are more widespread than philosophers and psychologists have traditionally assumed (e.g., Deery, Bedke, & Nichols, *in preparation*; Monroe & Malle, 2010; Murray & Nahmias, 2012; Nadelhoffer, *in press*; Nadelhoffer & Tocchetto, 2013, chap. 7; Nahmias, Morris, Nadelhoffer, & Turner, 2006; Nichols, 2004; Nichols, 2006; Nichols & Knobe, 2007; Shepard & Reuter, 2012; Shepherd, 2012; Rose & Nichols, 2013). Shedding light on the scope and limits of these compatibilist beliefs and attitudes is at least one of the issues researchers ought to examine. But if the tool we use for measuring people’s free will beliefs is predicated on incompatibilism, we cannot do so.

Unfortunately, the same criticism applies to a more recently developed and commonly used psychometric tool for measuring free will beliefs: the Free Will and Determinism Scale (Filevich et al., 2013; Rakos et al., 2008; Rigoni, Kuhn, Sartori, & Brass, 2011; Rigoni et al., 2012). In the development of their scale, Rakos and colleagues acknowledged that they followed in Viney’s incompatibilist footsteps by conceptualizing free will and determinism “as the opposing poles of an ‘amount of agency’ continuum” (p. 25). Moreover, there are additional worries about the scale designed by Rakos and colleagues. For instance, the validation round involved less than 100 participants, and thus was likely underpowered (Kline, 2011). Moreover, each of the items in their free will subscale contains the term “free will”—which may artificially inflate how well the free will items appear to hang together (Chen & Thissen, 1997). Finally, it is problematic that there is so much overlap between the items in the various subscales. For instance, why is “My decisions are influenced by a higher power” included within the Personal Limitations subscale rather than the Higher Power Control subscale? But for present purposes, the main shortcoming of the scale designed by Rakos and colleagues is the aforementioned worry that it, too, is predicated on incompatibilism.

Paulhus and Carey (2011) recently published a scale that we believe does a better job of avoiding the incompatibilist assumption while at the same time avoiding some of the other problems we have discussed—e.g., small samples, subscale items that all have the same word or term, etc. The Free Will and Determinism Scale (FAD-Plus) is a refinement of an earlier version of the scale (FAD, Paulhus & Margesson, 1994).¹ FAD-Plus is a 27-item scale that consists of four primary subscales: (a) Free Will (e.g., *People have complete free will*); (b) Scientific Determinism (e.g., *As with other animals, human behavior always follows the laws of nature*); (c) Fatalistic Determinism (e.g., *Whether people like it or not, mysterious forces seem to move their lives*); and (d) Unpredictability (e.g., *Chance events seem to be the major cause of human history*).

Contrary to the incompatibilist assumption that influenced the research described above, Paulhus and Carey (2011) found that responses to free will items and responses to scientific determinism items reliably loaded onto largely independent factors. Moreover, over multiple rounds involving several hundred participants, Paulhus and Carey consistently found that aggregate scores on the free will subscale and the scientific determinism subscale tended to be neither positively nor negatively correlated.

In light of these findings, Paulhus and Carey perhaps unsurprisingly conclude that their work on FAD-Plus “supports the conclusion that lay judges see freewill and determinism as quite compatible” (2011, p. 102). However, we prefer to remain agnostic for present purposes. The existing findings on free will beliefs are still too messy and inconclusive to draw any firm conclusions when it comes to the debate about whether most people have incompatibilist beliefs about free will. We need look no further than the data from FAD-Plus for an illustration of just how murky things can be. For instance, scores on their Fatalistic Determinism subscale correlate with scores on the Unpredictability subscale, which we find difficult to explain. More surprisingly, their Fatalistic Determinism subscale does *not* inversely correlate with the Free Will subscale, even though the former includes items such as, “What will be will be—there’s not much you can do about it” and “No matter how hard you try, you can’t change your destiny.” It seems that people who agree with these items, which suggest significant lack of control, would more strongly disagree with items, such as “People have complete control over the decisions they make,” and vice versa. In any case, given the content of the items, it is unclear what exactly the Fatalism and Unpredictability subscales are measuring.

There are additional problems with FAD-Plus as well. For instance, three of the seven Free Will items are about responsibility and blame, such as “Criminals are totally responsible for the bad things they do.” That these items tend to statistically load with items describing control and free will is an interesting finding (and below we will describe similar patterns in our early rounds of data collection), but such items seem to be indirect measures of beliefs about free will and control, and we think it is more useful to have ways to examine the relationships between beliefs about free will and beliefs about responsibility, rather than building them into the same subscale. Furthermore, the Scientific Determinism subscale seems designed to describe the idea that people’s character and actions are caused by their genes and upbringing, rather than describing the philosophical concept of determinism (e.g., the idea that there are sufficient causes for every event, including human decisions). Discovering that beliefs about these types of causation do not correlate with beliefs about free will is important, but it may not inform us about whether people have incompatibilist beliefs. Indeed, even though FAD-Plus does not presuppose incompatibilism as earlier scales did, it does not provide any obvious way to distinguish people who are disposed towards incompatibilism from those who are more inclined toward compatibilism. One could look for people who have high scores on both the Free Will subscale and the Scientific Determinism subscales as a proxy for compatibilism. But this may underestimate compatibilist beliefs since it mistakenly assumes that compatibilists believe strongly in determinism—which certainly need not be the case; and they certainly do not have to agree with statements suggesting, for instance, that our actions are determined by our genes. An indeterminist can still be a compatibilist—which requires merely that one believes that *if* determinism were true, we could be free. Consequently, trying to distinguish compatibilists from incompatibilists by comparing their aggregate subscale scores from FAD-Plus is problematic. Third, FAD-Plus fails to provide us with a way of shedding any light on the relationship between people’s free will beliefs and their beliefs about the soul, the mind, and the brain more generally. Finally, FAD-Plus also does not enable us to explore people’s beliefs about free will, determinism, and dualism and their beliefs about punishment and responsibility.

While we think FAD-Plus is better, all things considered, than the earlier incompatibilist scales, the need remains for a scale that describes the relevant concepts in a more precise and fine-grained way, ideally lining up more closely with the philosophical tradition, and that examines relationships between a wider range of relevant beliefs. At a minimum, none of the existing scales can shed much light on the aforementioned debate about whether or not most ordinary people have beliefs about free will that support incompatibilism.² In light of the flaws and limitations of the existing psychometric tools, we developed a new instrument for measuring beliefs about free will and related concepts.

¹ Stroessner and Green (1990) were two of the first researchers to appreciate that the relationship between beliefs about free will and determinism might be multi-dimensional. They developed a scale that treated beliefs about free will and determinism as independent (which is a virtue shared by FAD-Plus). Unfortunately, their scale shares several of the shortcomings of the scales we have already discussed—namely, questionable construct validity, small sample sizes, etc. Moreover, since most researchers have recently been using FAD-Plus to explore free will beliefs and associated behaviors (rather than the earlier scale developed by Stroessner and Green), FAD-Plus will be our primary focus here. But Stroessner and Green nevertheless deserve credit for being two of the first researchers to appreciate that beliefs about free will and beliefs about determinism may be largely independent.

² It is worth pointing out that shedding light on the debate about folk incompatibilism was not one of the explicit goals of the researchers who developed the extant scales. As such the researchers cannot be faulted for this particular shortcoming of the tools they developed. This is nevertheless an important limitation for those of us who are interested in this particular debate.

2. Developing and Refining Part 1³ of FWI⁴

Our goal was to adopt an approach to scale construction that was as theory-neutral as possible in the hopes that it would help us minimize the pitfalls of theory contamination. For instance, rather than present participants with a forced incompatibilist choice between free will and determinism, we wanted to present them instead with a variety of statements that would capture a wider range of beliefs and attitudes. Another important decision we made was that we would recruit people from the general population in the US rather than limiting our recruitment to only college students or only white Americans. For instance, FAD-Plus (Paulhus & Carey, 2011) was validated using only white college students (Study 1 and 2) and “residents of the United States with European heritage” (Study 3). Our goal was to make it more likely that our scale was developed and validated using a more representative sample than existing scales both in terms of age and gender and in terms of race and ethnicity. As we will see, we were largely successful on these fronts.

Having settled on a goal and a target audience for our participant pool, we came up with a list of items that we predicted would “hang together” or load onto common psychological factors (or constructs). The topics of the statements included in this initial list ranged widely from free will, responsibility, and punishment to the mind-body relationship, indeterminism, fate, science, and the immaterial soul. At this early stage, we wanted to cast our experimental and conceptual net as widely as possible. Four rounds of data were collected in an effort to develop and refine Part I of FWI. In each development round, FWI was subjected to exploratory factor analyses (EFAs). Items were retained, rejected, or revised, in part, based on the results of the EFAs.

2.1. Exploratory factor analysis overview

EFA is a statistical method used to distill a relatively large number of indicators, or measured variables, into a smaller number of factors, or latent variables. In the context of scale construction, the indicators are the responses to items on the scale, and the factors that emerge from an EFA are presumed to represent underlying constructs. For example, a scale might contain the statements, “People always have the ability to do otherwise,” and “How people’s lives unfold are completely up to them,” among many other statements. In response to these statements participants might be asked to express how much they agree or disagree with the statements using a 7-point scale. These responses are then entered in an EFA, and the EFA produces results that help researchers gain insight into how the various items on the scale cohere, or “hang together,” to form factors. These factors are often interpreted such that items that belong to the same factor express the same underlying construct. For example, maybe the EFA produces a result in which the two example sentences above belong to the same factor. This provides the researcher evidence that these two sentences might be related to each other such that both statements are expressions of the same underlying construct; in this case, maybe they are both expressions of the construct “free will.” However, one of the challenges of EFAs is that they do not always produce results that unambiguously inform the researcher how many factors a set of indicators can appropriately be reduced to or exactly which items should be considered to be a part of what factors. In order for one to effectively use an EFA, one must make decisions about how to interpret EFA results before the EFAs are conducted (see Fabrigar, Wegener, MacCallum, & Strahan, 1999, for a review).

First, one has to decide what *fitting procedure* should be used to estimate the number of factors and the individual factor loadings. *Factor loadings* are regression coefficients between individual items and factors. A higher loading indicates that the item is more strongly associated with the particular factor. There are two commonly used fitting procedures: *maximum likelihood* and *principal components*. Though maximum likelihood has many benefits, it is not robust to violations of assumptions, such as assumptions of normality. The principal components procedure is more robust to violations of assumptions and is thus less likely than maximum likelihood to produce an improper solution. For our EFA, we decided to use principal components because we wanted to use a test that would be robust to violations of assumptions, since we cannot know *a priori* if all assumptions would be met.

Second, one has to decide what are the appropriate number of factors to extract. Two common methods for factor extraction include the *Kaiser criterion* (1960) and *Cattell’s scree method* (1966). According to the Kaiser criterion, one should first compute the eigenvalues associated with each potential factor. Any factors with eigenvalues greater than 1.0 should be included in the model. According to Cattell’s scree test, a scree chart of the eigenvalues should be produced. The scree plot should be visually examined in order to determine the last substantial drop in the magnitude of the eigenvalues. The factors associated with the eigenvalues prior to the last drop are the factors that should be retained in the final model. The scree test has been criticized because deciding where the last substantial drop often involves subjective interpretation on the researcher’s part (Kaiser, 1960). The Kaiser criterion has been criticized as being arbitrary and as often leading to overfactoring, or to the inclusion of more factors than is warranted by the model (Cattell & Jaspers, 1967). However, it is generally deemed less problematic to specify *too many* factors than to specify *too few* factors (Fava & Velicer, 1992; Wood, Tataryn, & Gorsuch, 1996). Since at this exploratory stage we were more worried about missing important factors than worried about limiting factors, we decided to use the Kaiser criterion. However, we do not report on potential factors beyond those that would

³ We will be discussing the development of Part 2 of FWI in Section 4.

⁴ In Sections 2.2–2.5.2, we report brief summaries of Method and Results. For complete list of statements for each round and for factor loadings for each statement for each round, see [Supplementary Material A online](#).

be extracted using the Kaiser criterion because, as noted, the Kaiser criterion is already a liberal criterion that likely leads to overfactoring.⁵

Third, one must decide which factor rotation to apply. EFA models with more than one factor do not have a unique solution. Factor rotations help select the best “simple structure” among potential solutions, which in turn improves the interpretability of the factors. There are two types of factor rotation: orthogonal and oblique. Orthogonal rotations typically provide results that are more easily interpretable than oblique rotations; however, orthogonal rotations assume that factors are uncorrelated, which is typically a problematic assumption in social science research (Fabrigar et al., 1999). Oblique rotations do not assume that factors are uncorrelated, nor do oblique rotations assume that factors are correlated (Fabrigar et al., 1999). In this sense, oblique rotations are much more theory neutral. Since the assumption of uncorrelated factors is typically problematic in social science research, we decided to use an oblique rotation.

Fourth, one must decide how to interpret a factor. When interpreting a factor, one typically looks to see what indicators substantially load onto a factor and then determine what those indicators have in common. Though rules of thumb for what is considered a “substantial loading” vary from text to text, a common rule is that factor loadings above .400 are considered substantial (Hair et al., 1998). Though EFA does not define a factor only by those indicators that load substantially on a factor (a factor is defined by the individual loadings for *all* the indicators on that factor), the goal in using EFA for scale construction is to extract a subset of items that serve as the best indicators of the factor. In order to determine which items loaded substantially onto which factors, we decided to use the Hair et al. “above .400” rule of thumb.

2.2. Round 1

2.2.1. Method

2.2.1.1. Participants. Participants were 293 people recruited online using Qualtrics paid panelist service. Forty-one percent of participants ($n = 121$) reported being male, 46% ($n = 135$) female, and 13% ($n = 37$) did not report gender. Seventy-three percent of participants ($n = 215$) reported being White/Caucasian, 7% ($n = 21$) Black/African-American, 4% ($n = 13$) Asian/Pacific Islander, 2% ($n = 6$) Hispanic/Latino, <1% Native American/American Indian ($n = 1$), and 12% ($n = 37$) did not report ethnicity. Participants' age ranged from 18 to 80, with a mean age of 51.2 (SD = 13.89). Age 43, 52, and 62 marked the 1st, 2nd, and 3rd quartiles, respectively.

2.2.1.2. Materials. In the initial development round the scale consisted of 30 items. The 30 items were selected to reflect a diverse range of potential folk conceptual categories related to free will, while also being amenable to the finer distinctions and technical notions found in the philosophical discussions of free will. These included statements intended to capture belief in free will (e.g., *People always have the ability to do otherwise*), determinism (e.g., *People's choices and actions have to happen precisely the way they do because of the laws of nature and the way things were at the beginning of the universe*), chance/indeterminism (e.g., *Chance plays a large role in what happens in the universe*), fate (e.g., *Much of what happens to us in life is unavoidable*), moral responsibility (e.g., *People are blameworthy for their bad actions*), and dualism (e.g., *People have non-physical souls that survive after their bodies die*).

2.2.1.3. Procedure. Participants were presented with 30 randomized statements about free will and related constructs of interest and were asked to rate their agreement with each statement using a 7-point Likert scale anchored with *1 Strongly Disagree* and *7 Strongly Agree*. For all four rounds, participants' responses were analyzed using EFA factored with principle components and oblimin rotation. Factors with Eigenvalues above 1.0 were extracted. Following guidelines of Hair et al. (1998), items that had loadings above .400 after oblimin rotation on a factor were considered “substantial loadings.” Any item with a substantial loading on a factor was considered to comprise that factor.

2.2.2. Results

The EFA produced seven factors with eigenvalues above 1.0. Those eigenvalues were 5.951, 3.348, 2.551, 2.028, 1.251, 1.144, and 1.014. The statements that substantially loaded onto Factor 1 included statements intended to reflect determinism and statements intended to reflect fatalism. This suggests one of three possibilities: (1) Although determinism and fatalism are conceptually dissociable, these concepts may not be strongly differentiated in the general population, (2) the concepts may be differentiated in the general population, but the wording of the statements may need to be improved to better capture the distinctions, or (3) the concepts may be differentiated and the wording accurately captures the distinction, but given the nature of the statistics, they are unable to distinguish between the potentially subtle distinctions. Factor 2 included all, and only, statements intended to capture belief in free will. Factor 3 only included statements intended to capture belief in dualism. Factor 4 appears to reflect a “Lack of Control” construct and included some statements about chance and some statements about fatalism. Only two items loaded substantially on Factor 5 and the two items did not appear related.

⁵ If we had used Cattell's scree test as our criterion, we often would have extracted fewer factors and would have never extracted more. In order to confirm that we were not missing any meaningful factors beyond those extracted by the Kaiser criterion, we also conducted a post-hoc exploration of potential factors beyond the factors extracted by the Kaiser criterion. All factors beyond those extracted by the Kaiser criterion were always uninterruptable. For example, many of the factors only had a single item that substantially loaded. In cases, where more than one item substantially loaded, the items did not appear to be conceptually related. These results are typical of potential factors beyond the Kaiser criterion.

Factor 6 included only statements about moral responsibility. Only two items loaded substantially on Factor 7 and both items reflected indeterminism. We expected chance and indeterminism to be conflated, but chance statements loaded alongside some of the fatalism statements in what seemed to be a “Lack of Control” construct. Indeterminism statements loaded on a factor separate from any chance statements (or any other statements). Three items did not highly load onto any of the seven factors.

2.3. Round 2

2.3.1. Method

2.3.1.1. Participants. Participants were 321 people recruited online using Qualtrics paid panelist service. Forty-seven percent of participants ($n = 151$) reported being male, 51% ($n = 163$) female, and 2% ($n = 7$) did not report gender. Eighty-two percent of participants ($n = 264$) reported being White/Caucasian, 7% ($n = 22$) Black/African-American, 3% ($n = 11$) Asian/Pacific Islander, 3% ($n = 11$) Hispanic/Latino, 2% Native American/American Indian ($n = 6$), and 2% ($n = 7$) did not report ethnicity. Participants' age ranged from 18 to 80, with a mean age of 47.51 ($SD = 14.53$). Age 36, 47, and 59 marked the 1st, 2nd, and 3rd quartiles, respectively.

2.3.1.2. Materials. The three items that did not load onto any of the factors in the previous round were dropped. Additional items were either dropped or revised in an effort to capture more coherent conceptual categories of the general population while also adhering to important philosophical categories and distinctions. For example, the fatalism items that loaded onto Factor 1 alongside determinism statements were dropped, while the two fatalism items that loaded onto Factor 4 were retained. After revisions, the scale consisted of 23 items that were intended to capture beliefs in free will, responsibility, determinism, fatalism, dualism, and non-reductionism (e.g., *Humans have both minds and brains and the mind cannot simply be reduced to the brain*).

2.3.1.3. Procedure. Procedure in Round 2 was the same as procedure in Round 1.

2.3.2. Results

The EFA produced five factors with eigenvalues above 1.0. Those eigenvalues were 4.237, 3.776, 2.218, 1.306, 1.041. Factor 1 only included two items, both items expressing belief in non-reductionism. Factor 2 included all and only items related to determinism. No items substantially loaded on Factor 3. Factor 4 consisted of items expressing fatalism. In contrast to Round 1, the EFA results for this round suggest that the general population may make a conceptual distinction between determinism and fatalism. Factor 5 contained five free will and responsibility items. Ten items, including a number of items intended to express beliefs about dualism, did not substantially load onto any factors.

2.4. Round 3

2.4.1. Method

2.4.1.1. Participants. Participants were 317 people recruited online using Qualtrics paid panelist service. Forty-eight percent of participants ($n = 153$) reported being male, 51% ($n = 162$) female, and < 1% ($n = 2$) did not report gender. Eighty-six percent of participants ($n = 271$) reported being White/Caucasian, 7% ($n = 22$) Black/African-American, 4% ($n = 13$) Asian/Pacific Islander, 3% ($n = 8$) Hispanic/Latino, <1% Native American/American Indian ($n = 1$), and < 1% ($n = 2$) did not report ethnicity. Participants' age ranged from 20 to 80, with a mean age of 49.30 ($SD = 13.07$). Age 40, 51, and 59 marked the 1st, 2nd, and 3rd quartiles, respectively.

2.4.1.2. Materials. All of the free will, responsibility, determinism, and fatalism items were retained from the previous round. In the previous round, a dualism construct did not clearly emerge from the retained factors. In an effort to better capture a dualism construct, all the other items from the previous round were heavily scrutinized. Items were revised and rejected, leaving seven items that were intended to capture beliefs in the constructs of dualism (suggesting the existence of a non-physical mind or soul, such as *The fact that we have souls that are distinct from our material bodies is what makes us unique*), or the related but distinct philosophical concept of non-reductionism (suggesting that mental states cannot be understood solely in terms of physical states of the brain, such as *People are greater than the sum of their biological and chemical parts*). The total items for this round were 19.

2.4.1.3. Procedure. Procedure in Round 3 was the same as the procedure in Round 1 and 2.

2.4.2. Round 3 results

The EFA produced four factors with eigenvalues above 1.0. Those eigenvalues were 4.374, 3.793, 1.744, and 1.073. Factor 1 included all and only items that reflected dualism or non-reductionism. Factor 2 included all and only items that reflected determinism or fatalism. Like Round 1 but unlike Round 2, the Round 3 EFA suggested that people tend to conflate determinism and fatalism. Factor 3 included all items that reflected free will and one item that reflected moral responsibility. This moral responsibility item also loaded onto Factor 4. Factor 4 included two items, both reflecting moral responsibility.

2.5. Round 4

2.5.1. Method

2.5.1.1. Participants. Participants were 232 people recruited online using Qualtrics paid panelist service. Fifty-four percent of participants ($n = 126$) reported being male, 43% ($n = 100$) female, and 3% ($n = 6$) did not report gender. Eighty percent of participants ($n = 185$) reported being White/Caucasian, 9% ($n = 21$) Black/African-American, 3% ($n = 8$) Asian/Pacific Islander, 4% ($n = 9$) Hispanic/Latino, 2% Native American/American Indian ($n = 4$), and 2% ($n = 5$) did not report ethnicity. Participants' age ranged from 18 to 80, with a mean age of 47.67 (SD = 14.90). Age 35, 50, and 58 marked the 1st, 2nd, and 3rd quartiles, respectively.

2.5.1.2. Materials. The free will items that loaded together were retained from the previous round, though we added modifiers such as 'completely' and 'always' to them in an attempt to obtain a wider range of responses, and we dropped the few moral responsibility items that sometimes loaded onto factors with free will items. We used seven items to reflect determinism, and dropped the few fatalism items that sometimes loaded onto factors with determinism items. We included 13 dualism/non-reductionism items to find the items that best represented that factor. And we developed some new items to consider whether a 'bypassing' factor would emerge (i.e., items describing whether or not certain mental states play an essential role in behavior).

2.5.1.3. Procedure. Procedure in Round 4 was the same as the procedure in Rounds 1, 2, and 3.

2.5.2. Round 4 results

The EFA produced eight factors with eigenvalues above 1.0. However, no factors beyond the third factor had more than one item substantially load, so only the first three factors were examined. Five free will items, along with three mental causation items (non-bypassing items), loaded onto Factor 1. Two other free will items did not load onto any factors. Five determinism items loaded onto Factor 2. No other items loaded onto Factor 2. Two determinism items did not load onto any factors. Four dualism/non-reductionism items, along with two mental causation items loaded onto Factor 3. Other items, including several dualism items and several bypassing items failed to load on any factor.

3. Validating FWI with confirmatory factor analysis

Many lessons were learned from the four development rounds. For example, it is not clear if determinism and fatalism, or if free will and moral responsibility, form clear and distinct conceptual categories for most people. It is possible that most people do distinguish determinism from fatalism, and free will from moral responsibility, but given the close relationships between the concepts within these pairs, it may be difficult for psychometric tools to tease apart these concepts (one goal of Part 2 was to examine some of the conceptual relationships between free will and related concepts). Moreover, we found that with some items, agreement is so high that we had to strengthen the language in order to differentiate people's beliefs. Another lesson learned was that people responded differently to statements using conditionals and counterfactuals than straightforward assertions—which is something that merits closer consideration in future research. Most importantly, we were able to find evidence of clear and distinct categories for the three primary constructs that we sought to measure: free will, determinism, and dualism/anti-reductionism (while anti-reductionism does not entail substance dualism, we were unable to develop distinct constructs for these items, so this subscale includes items representing both concepts). So, after the four aforementioned rounds of preliminary data collection, we settled on the following 15-item scale ((R#) notes in which round the final items originally appeared):

The Free Will Inventory: Part 1⁶

The Free Will Subscale (FW):

1. People always have the ability to do otherwise. (R1)
2. People always have free will. (R1)
3. How people's lives unfold is completely up to them. (R2)
4. People ultimately have complete control over their decisions and their actions. (R2)
5. People have free will even when their choices are completely limited by external circumstances. (R4)

The Determinism Subscale (DE):

1. Everything that has ever happened had to happen precisely as it did, given what happened before. (R4)
2. Every event that has ever occurred, including human decisions and actions, was completely determined by prior events. (R2)

⁶ See [Supplementary Material B online](#) for the complete, standard version of FWI (including directions for use).

3. People's choices and actions must happen precisely the way they do because of the laws of nature and the way things were in the distant past. (R2)
4. A supercomputer that could know everything about the way the universe is now could know everything about the way the universe will be in the future. (R4)
5. Given the way things were at the Big Bang, there is only one way for everything to happen in the universe after that. (R4)

The Dualism/Anti-Reductionism Scale (DU):

1. The fact that we have souls that are distinct from our material bodies is what makes humans unique. (R3)
2. Each person has a non-physical essence that makes that person unique. (R4)
3. The human mind cannot simply be reduced to the brain. (R4)
4. The human mind is more than just a complicated biological machine. (R4)
5. Human action can only be understood in terms of our souls and minds and not just in terms of our brains. (R4)

Having settled on these fifteen items, we proceeded to collect a fifth and final round of data for the purposes of using confirmatory factor analysis (CFA) to formally validate our scale.

3.1. Confirmatory factor analysis overview

CFA statistically examines how well a specified model “fits” the data (Pedhazur & Schmelkin, 1991). It assesses how adequately each observed variable (e.g., Free Will item 1, indicated by v1) represents latent variables (e.g., free will, determinism, dualism) via standardized regression weights. Researchers specify exact linkage patterns between observed variables and underlying factors *a priori*, based on theoretical knowledge or prior research (or both), and then test the goodness of this fit (Byrne, 1994). AMOS version 19 was used to perform the CFA (Arbuckle, 2010). Our model specified 15 regressions (association between the 5 items and the 3 latent variables), 3 covariances (i.e., curved arrows between latent variables), and 15 variances (error terms for each observed variable) for a total of 33 parameters. The recommended minimum criteria of 10 cases per parameter (Schreiber et al., 2006) is met, due to our sample containing 330 cases with no missing data.

There are several measures available to assess the fit of a model. The Chi-square (X^2) test, the Goodness of Fit Index (GFI), and the Comparative Fit Index (CFI) all compare the similarity of the covariance matrix predicted by the model and the observed indicators' covariance matrix. Higher values (ideally above .90) for the GFI and CFI are desirable, as are non-significant X^2 values (Bentler, 1993). The Root Mean Square Error of Approximation (RMSEA) and the Adjusted Goodness of Fit Index (AGFI) are also reported, both of which take into account model complexity. RMSEA values below .05 indicate good fit (Browne & Cudeck, 1993). Values between .05 and .08 indicate fair fit, between .08 and .10 mediocre fit, and above .10 poor fit (Browne & Cudeck, 1993; MacCallum, Browne, & Sugawara, 1996). In addition, AGFI values above .80 are reasonable (Cole, 1987; Pedhazur & Schmelkin, 1991).

3.2. Method

3.2.1. Participants

Participants were 330 people recruited using Qualtrics Paid Panelist service. Nearly half of the participants (48.8%, $n = 161$) reported being male, 48.8% ($n = 161$) were female, and 2% ($n = 8$) did not report gender. The majority of the participants (74.2%, $n = 245$) reported being White/Caucasian, 10% ($n = 33$) were Black/African-American, 5.5% ($n = 18$) were Asian/Pacific Islander, 7.6% ($n = 25$) were Hispanic/Latino, 1.8% were Native American/American Indian ($n = 6$), and less than one percent (.9% ($n = 3$)) did not report ethnicity. Participants' age ranged from 19 to 80, with a mean age of 46.29 ($SD = 14.61$). Age 33, 46, and 58 marked the 1st, 2nd, and 3rd quartiles, respectively.

3.2.2. Materials and procedure

Participants were presented with statements from the final version of the FWI and were asked to rate their agreement with the statement using a 7-point Likert scale anchored with *1 Strongly Disagree* and *7 Strongly Agree*.

3.3. Results

3.3.1. CFA

Observed variables were assigned to their latent factors and the path weights were estimated using the standard maximum-likelihood method for all models. The analysis of the three-factor model yielded acceptable estimates for four of the five indices. The estimates were acceptably high for the GFI (.927) and the CFI (.931) and for the AGFI (.900). The RMSEA was .058, indicating fair fit. Finally, the chi square fit index was statistically significant, $X^2(87) = 183.039$, $p < .001$. Ideally the X^2 statistic should be lower and non-significant, but it is recognized that this fit index is often too stringent in psychometric research (Hopwood & Donnellan, 2010; Paulhus & Carey, 2011; Raykov, 1998).

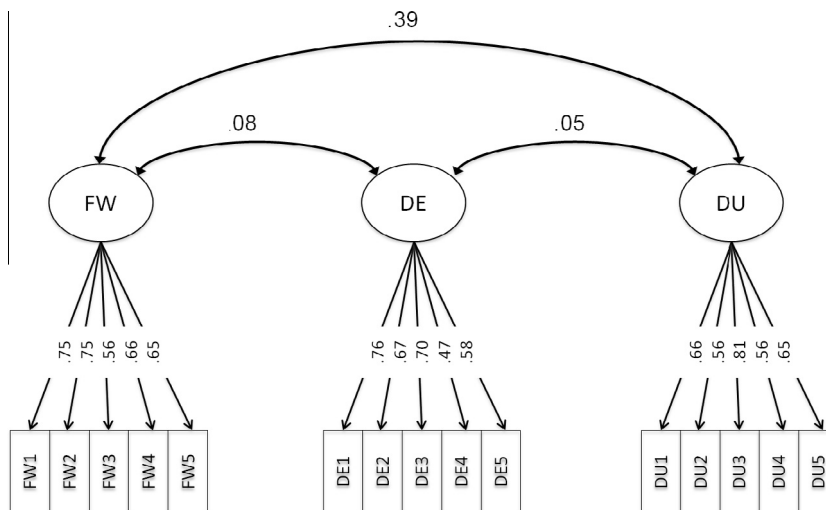


Fig. 1. CFA round regression weights.

The standardized regression weights are shown in Fig. 1. The ranges for items were as follows: Free Will from .56 (fw3) to .75 (fw1 and fw2), Determinism from .47 (de 4) to .76 (de 1), and Dualism from .56 (du2 and du4) to .81 (du3). Looking at the standardized residuals (as suggested by Byrne, 1994), the average off-diagonal standardized residual is .0568, and 86% of the residuals were between $-.10$ and $.10$. The closer the residuals are to zero, the better the fit, although there is no cutoff rule for what is acceptable or unacceptable. The relationship with the strongest degree of misfit (.24) was between fw3 and de1. Overall, the regression weights, the residuals, and the goodness of fit indices, suggest that the model fit is fairly strong.

3.3.2. Psychometric information

The subscales all yielded acceptable internal consistency: the five-item FW subscale Cronbach $\alpha = .803$, the five-item Determinism subscale $\alpha = .772$, and the five-item dualism subscale $\alpha = .771$.

3.3.3. Correlations

The correlation between free will and determinism and between dualism and determinism did not significantly differ from zero, p 's $> .05$. Dualism positively correlated with free will beliefs, $r = .319$, $p < .001$. These correlations were re-run to examine whether there might be gender or ethnicity differences. The two ethnicity groups with the largest number of cases (Caucasian and African American) were the basis of comparison.

The pattern of results was very similar for men, with the only significant relationship being between dualism and free will, $r = .395$, $p < .001$. The pattern was similar for women, though the magnitude of the dualism-free will correlation was weaker, $r = .238$, $p = .002$.

As for ethnicity, the analyses for Caucasians paralleled the full sample; the only significant correlation was between dualism and free will, $r = .252$, $p < .001$. The matrix for African Americans revealed a stronger correlation between determinism and free will ($r = .430$, $p = .013$). In addition, there were significant correlations between dualism and determinism, $r = .399$, $p = .021$, and between free will and dualism, $r = .389$, $p = .025$.

3.3.4. Demographic analyses

We conducted a 2 (gender) by 2 (ethnicity: black versus white) MANOVA on the three beliefs (free will, determinism, and dualism). There were no gender differences or ethnic group differences for the three beliefs. However, there was a significant overall interaction, Pillai's Trace $F(3, 267) = 2.773$, $p = .042$, $\eta^2 = .030$. The interaction was for beliefs regarding dualism, $F(1, 272) = 7.570$, $p = .006$, $\eta^2 = .027$. Among African Americans, men ($M = 5.73$, $SD = .31$) scored higher than women ($M = 4.95$, $SD = .24$). In contrast, Caucasian women ($M = 5.571$, $SD = .10$) scored higher than Caucasian men ($M = 5.21$, $SD = .10$).

4. Developing and refining Part 2 of FWI

One of our primary goals was to design a psychometric instrument for measuring the strength of people's beliefs in free will, determinism, and dualism. In doing so, we also wanted to avoid some of the limitations and problems of the extant tools. After five rounds of data collection and analysis, we believe we were largely successful in accomplishing these goals. Not only were we able to get items that measured the factors we were interested in exploring, but we were able to do so in a way that was more faithful to how philosophers have traditionally understood the key terms in question.

However, from the outset of our project, we realized that it would not be enough to measure only the *strength* of people's beliefs about free will and related concepts—which is all the items from Part 1 of FWI are capable of measuring. We also wanted to be able to explore the relationship between these beliefs. To get a better sense of our worry, keep in mind that a compatibilist and an incompatibilist could both equally agree with the statements from our free will subscale. This highlights the fact that the strength of a person's belief in free will tells us little about the content of this belief. Indeed, it is also possible to have either incompatibilist or compatibilist intuitions but to disagree that humans have free will, since one might think something other than determinism challenges or limits human free will or that free will is impossible (e.g., Strawson, 1994).

In short, we realized early on that if we were going to be able to distinguish compatibilist beliefs from incompatibilist beliefs, or to examine other relationships among various beliefs related to free will, we would need to do more than just measure the strength of people's free will beliefs—that is, we would need to add yet another psychometric instrument to our investigative tool kit. We decided to develop an additional list of statements (Part 2 of FWI) that would better enable us to explore the more fine-grained relationships between people's beliefs about free will and their beliefs and theories about causation, the soul, scientific prediction, choice, responsibility, and punishment. We will now turn our attention to our efforts on this front.

During the process of collecting the data for the purposes of constructing and validating Part 1 of FWI, we also collected two rounds of data on how people think about the nature of free will and moral responsibility (during Round 1 and Round 4, respectively). Between these two rounds of data collection, we revised the items for Part 2 in the hopes that we could use it to shed some additional light on the contours of people's beliefs about free will and related concepts. Items were presented to participants in random order in each study and participants rated their agreement on a 7-point Likert scale that was anchored by *1 Strongly disagree* and *7 Strongly agree*. See Table 1 for the list of items from Part 2 of FWI, along with means and standard deviations from Round 4.

During the process of developing Part 2 of FWI, we found some interesting results. But before we discuss these findings, we should first say a bit more about why we included some of the particular items that we selected. For instance, we included several statements that were specifically intended to help us identify libertarian beliefs about free will and responsibility—namely, items 1, 4, 6, and 8. Each of these items focuses on different aspects of incompatibilism—from indeterminism (1 and 4) and contra-causal powers (6) to so-called “sourcehood” or “ultimate responsibility” (8). By including several different libertarian items, we hoped to maximize the chances that we would be able to distinguish pro-free will compatibilists from pro-free will incompatibilists.⁷

One of the more noteworthy findings from our work on Part 2 of FWI was that most participants endorsed statements that suggest that humans have libertarian free will and the unconditional ability to do otherwise. Consider, for instance, the responses to Part 2 from Round 4. For ease of presentation, we are going to lump together the responses by those who strongly disagree, disagree, and somewhat disagree under the heading “disagree” and we are going to lump together the responses by those who strongly agree, agree, and somewhat agree under the general heading “agree.” Here are the aggregate responses from Round 4 to the libertarian items:

- (1): 79% agree; 6% disagree; 15% neither agree nor disagree
- (4): 74% agree; 10% disagree; 16% neither agree nor disagree
- (6): 47% agree; 17% disagree; 36% neither agree nor disagree
- (8): 78% agree; 11% disagree; 11% neither agree nor disagree

There are several things worth mentioning about these (and related) findings from Part 2. First, the affirmative responses to these items demonstrate that a strong majority of people are willing to endorse statements that explicitly express libertarian views of free will when they are given the chance to do so. Second, responses to the libertarian items correlated in predictable ways with participants' other beliefs and attitudes. For example, item 1 correlated with scores on the free will subscale, $r = .26, p < .01$ and the dualism subscale, $r = .20, p < .01$ in addition to correlating with item 4, $r = .34, p < .01$, item 6, $r = .20, p < .01$, and item 8, $r = .24, p < .01$. Finally, and perhaps most importantly, responses to item 1, unlike aggregate scores on the free will subscale, were negatively correlated with scores on the determinism subscale, $r = -.14, p < .01$. These patterns of correlations are consistent with what one would expect to find if these items were successful at distinguishing libertarian incompatibilists from compatibilists. See Table 2 for correlations between all FWI Part 2 items and FWI Part 1 subscale scores and demographic measures.

Trying to sort out all of these issues is a task for another day. As things stand, we have provided evidence that most people largely agree *both* with explicit statements of compatibilist free will and with explicit statements of incompatibilist free will. Whether participants' responses to the items of our scale can be taken to be indicative of any deep metaphysical folk theories about free will and responsibility is an open question. For now, what we do know is that both those who defend folk

⁷ While we are focusing here on our attempt to distinguish libertarian incompatibilists from compatibilists, we also included several items in Part 2 that were designed to tease apart participants' intuitions about punishment—which was one of the key issues that motivated the earlier work by Viney and colleagues. For instance, items 10, 11, and 12 were specifically designed to distinguish backward looking retributivist intuitions about punishment from forward looking instrumentalist intuitions about punishment. Items 3 and 13 were designed to examine beliefs about the compatibility of free will or responsibility and a law-governed universe. Item 2 was designed to measure agreement with a conditional account of the ability to do otherwise (as opposed to the unconditional account being measured by Item 1). Item 7 was designed to measure agreement with a traditional compatibilist conception of free will. And items 5 and 14 provide information about the people's beliefs about the relationship between dualism and free will. For further discussion of some these other items, see Nadelhoffer (in press).

Table 1

Means and standard deviations for FWI Part 2 statements.

#	Statements	Means (SD)
1	Free will is the ability to make different choices even if everything leading up to one's choice (e.g., the past, the situation, and their desires, beliefs, etc.) were exactly the same	5.44 (1.30)
2	Free will is the ability to make a choice based on one's beliefs and desires such that, if one had different beliefs or desires, one's choice would have been different as well	5.45 (1.20)
3	People could have free will even if scientists discovered all of the laws that govern all human behavior	5.39 (1.41)
4	To have free will means that a person's decisions and actions could not be perfectly predicted by someone else no matter how much information they had	5.31 (1.43)
5	If it turned out that people lacked non-physical (or immaterial) souls, then they would lack free will	3.82 (1.68)
6	To have free will is to be able to cause things to happen in the world without at the same time being caused to make those things happen	4.53 (1.45)
7	People have free will as long as they are able to do what they want without being coerced or constrained by other people	4.80 (1.71)
8	To be responsible for our present decisions and actions we must also be responsible for all of our prior decisions and actions that led up to the present moment	5.35 (1.45)
9	People deserve to be blamed and punished for bad actions only if they acted of their own free will	4.48 (1.75)
10	People who harm others deserve to be punished even if punishing them will not produce any positive benefits to either the offender or society—e.g., rehabilitation, deterring other would-be offenders, etc.	5.37 (1.46)
11	People who perform harmful actions ought to be rehabilitated so they no longer pose a threat to society	4.95 (1.55)
12	People who perform harmful actions ought to be punished so that other potential offenders are deterred from committing similar harmful actions	5.78 (1.20)
13	People could be morally responsible even if scientists discovered all of the laws that govern human behavior	5.28 (1.46)
14	If it turned out that people lacked non-physical (or immaterial) souls, then they would lack moral responsibility	3.80 (1.77)

Note: $n = 330$. Items 1–7 are items designed to measure people's beliefs about the nature of free will, items 8–14 about moral responsibility.

Table 2

Correlations between FWI Part 2 items and FWI Part 1 subscale scores and demographic measures.

Measure	FW Subscale	DE Subscale	DU Subscale	Age	Political ideology	Religiosity
P2 Item 1	.260**	-.114*	.205**	.074	-.144**	.106
P2 Item 2	.248**	-.116*	.281**	.073	-.068	.046
P2 Item 3	.350**	-.255**	.231**	.035	-.089	.025
P2 Item 4	.245**	-.091	.242**	.104	-.012	.094
P2 Item 5	-.430	.290**	.136*	-.052	.099	.174**
P2 Item 6	.200**	-.006	.046	-.083	-.073	-.048
P2 Item 7	.105	.139*	.078	.022	.017	.073
P2 Item 8	.243**	.098	.224**	.013	.087	.041
P2 Item 9	.088	.141*	.012	-.063	-.017	.018
P2 Item 10	.198**	-.058	.165**	.189**	-.095	.156**
P2 Item 11	.180**	.043	.038	-.091	.136*	.101
P2 Item 12	.261**	-.071	.235**	.048	-.087	.130*
P2 Item 13	.260**	-.186**	.162**	.036	-.021	.065
P2 Item 14	-.028	.280**	.266**	.017	.028	.165**

Note: $n = 330$. Political Ideology Item was a 7-point scale anchored by 1 = *Conservative* 4 = *Moderate* and 7 = *Liberal*. Religiosity item was a 5-point scale anchored by 1 = *Not at All Religious* and 5 = *Very Religious*.

* $p < .05$.

** $p < .01$.

compatibilism and those who defend folk incompatibilism have some explaining to do. This is one reason we have chosen as a team to remain agnostic when it comes to the on-going debate between these two camps. On the one hand, different members of our team have different views about this and related issues. On the other hand, there is too much conflicting data and too many unanswered questions to draw any firm conclusions for now. However, at least one thing is clear—namely, if we are to make further progress when it comes to our understanding of free will beliefs and related concepts, we need to ensure that the psychometric tools we are using for measuring these beliefs are up to the task. Our goal throughout this project was to try to make progress on this front by developing an instrument that had more strengths and fewer weakness than previous models. We humbly believe we accomplished this task.

5. General discussion

Increasingly, scientists argue that their discoveries challenge the existence of free will. Meanwhile, as we saw earlier, several studies indicate that informing people about such discoveries can change the degree to which they believe in free will and subtly alter some of their behaviors. But the psychometric tools used in these studies to measure people's beliefs about free will are arguably problematic (see Section 1). So, we decided a new instrument was needed (especially given the practical importance of the behavioral research). Through a series of refinements, using a representative sample of US participants, we developed and validated Part 1 of FWI to measure strength of belief in free will as accurately as possible, while also measuring strength of belief in determinism and in dualism and non-reductionism about the mind. In each case, one goal was to maximize construct validity, especially by (a) avoiding significant discrepancies between the items and the meaning of the concepts in philosophical and scientific discussions and (b) avoiding problematic connotations among items in each subscale (e.g., conflating determinism with fatalism or free will with sense of responsibility). Using CFA we validated Part 1 of FWI, showing that subscale items successfully measure independent concepts. We also developed Part 2 of FWI to allow researchers to examine more complex relationships among beliefs about free will, choice, responsibility, punishment, the existence of souls, and the possibility of scientific prediction.

We hope that further studies by our group and others will be able to fruitfully use FWI to shed new light on beliefs about free will and related concepts. There are a number of possible avenues that researchers could pursue. For instance, one issue that merits closer consideration is whether intuitions vary between different populations when it comes to free will beliefs. This is something we are currently planning to explore with two different lines of research. First, we are in the process of translating FWI into other languages (e.g., Brazilian Portuguese) to measure potential cross-cultural differences in beliefs about free will and related concepts. Second, because we have used a more representative community participation pool to construct and validate FWI, we are now in the position to administer it to college students to see whether there are any interesting differences. Another issue we plan to explore is the role that elements such as (a) personality type, (b) religiosity, (c) political ideology, and (d) philosophical (or scientific) training play in shaping people's intuitions about agency and responsibility. Finally, we are also developing both a first person and third person version of FWI to see whether we find any self-other differences when it comes to how people think about free will.

In each of these cases, one of our goals will be to further explore the complex relationship between people's beliefs about free will and their beliefs about determinism, dualism, scientific discovery, punishment, and related concepts. But our primary goal will be to try to continue to improve the psychometric tools researchers have at their disposal for undertaking these and related projects. As we have seen, accomplishing this goal is no easy task. Our hope is that others will join us in our efforts. Given the increasing visibility of the free will debate in the popular press and the public sphere, the time is ripe for researchers to develop a better understanding of how people ordinarily think about free will and related concepts. Hopefully, our work on FWI has moved us a few more steps in the right direction.

Acknowledgments

This work was supported by the John D. and Catherine T. MacArthur Foundation, The Law and Neuroscience Project, and The Regents of the University of California (TN), a Big Questions in Free Will (BQFW) grant from the John Templeton Foundation (TN and EN), and a Neuroethics Scholar Program Fellowship from Emory University's Center for Ethics (JS). The opinions expressed in this publication are those of the authors alone and do not necessarily reflect the views of the John D. and Catherine T. MacArthur Foundation, the Law and Neuroscience Project, The Regents of the University of California, the John Templeton Foundation, or Emory University's Center for Ethics. We would like to thank these organizations for their generous support. We would also like to thank the following people for their helpful assistance and suggestions along the way: Lauren Brewer, Richard Gonzalez, Daniela Goya Tocchetto, Trevor Kvaran, Scott Lilienfeld, Alfred Mele, Jonathan Schooler, Kathleen Vohs, audience members at the annual BQFW workshops, and commentators on the blog Flickers of Freedom.

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.concog.2014.01.006>.

References

- Alquist, J. L., Ainsworth, S. E., & Baumeister, R. F. (2013). Determined to conform: Disbelief in free will increases conformity. *Journal of Experimental Social Psychology*, 49, 80–86.
- Arbuckle, J. L. (2010). *IBM SPSS Amos 19.0 User's Guide*. Crawfordville, FL: Amos Development Corporation.
- Bargh, J. (2008). Free will is un-natural. In J. Baer, J. Kaufmann, & R. Baumeister (Eds.), *Are we free? psychology and free will* (pp. 128–154). New York: Oxford University Press.
- Baumeister, R. F., Masicampo, E. J., & DeWall, C. N. (2009). Prosocial benefits of feeling free: Disbelief in free will increases aggression and reduces helpfulness. *Personality and Social Psychology Bulletin*, 35, 260–268.
- Bentler, P. M. (1993). *EQS structural equations program manual*. Los Angeles: BMDP Statistical Software Inc.
- Bergner, R. M., & Ramon, A. (2013). Some implications of beliefs in altruism, free will, and nonreductionism. *The Journal of Social Psychology*, 153, 598–618.
- Bourget, D., & Chalmers, D. (in press). What do philosophers believe? *Philosophical Studies*.

- Browne, M. W., & Cudeck, R. (1993). Alternative ways of assessing model fit. In K. A. Bollen & J. S. Long (Eds.), *Testing structural equation models* (pp. 136–162). London: Sage Publications.
- Byrne, B. M. (1994). *Structural equation modeling with EQS and EQS/Windows: Basic concepts, applications and programming*. Thousand Oaks, CA: Sage.
- Cattell, R. B. (1966). The scree test for the number of factors. *Multivariate Behavioral Research*, 1, 245–276.
- Cattell, R. B., & Jaspers, J. (1967). A general plasmode (No. 30-10-5-2) for factor analytic exercises and research. *Multivariate Behavioral Research Monographs*, 67.
- Chen, W., & Thissen, D. (1997). Local dependence indexes for item pairs using item response theory. *Journal of Educational and Behavioral Statistics*, 22, 265–289.
- Cole, D. A. (1987). Utility of confirmatory factor analysis in test validation research. *Journal of Consulting and Clinical Psychology*, 55, 584–594.
- Deery, O., Bedke, M., & Nichols, S. (in preparation). Phenomenal abilities: Incompatibilism and the experience of agency.
- Fabrigar, L., Wegener, D., MacCallum, R., & Strahan, E. (1999). Evaluating the use of exploratory factor analysis in psychological research. *Psychological Methods*, 4, 272–299.
- Fava, J. L., & Velicer, W. F. (1992). The effects of over extraction on factor and component analysis. *Multivariate Behavioral Research*, 27, 387–415.
- Filevich, E., Vanneste, P., Brass, M., Fias, W., Haggard, P., & Kuhn, S. (2013). Brain correlates of subjective freedom of choice. *Consciousness and Cognition*, 22, 1271–1284.
- Fischer, J. M., & Ravizza, M. (1998). *Responsibility and control: A theory of moral responsibility*. Cambridge University Press.
- Frankfurt, H. (1988). *The importance of what we care about*. Cambridge University Press.
- Hair, J. F., Jr., Anderson, R. E., Tatham, R. L., & Black, W. C. (1998). *Multivariate data analysis with readings* (5th ed.). Englewood Cliffs, NJ: Prentice Hall.
- Harris, S. (2012). *Free will*. New York: Free Press.
- Haynes, S. D., Rojas, D., & Viney, W. (2003). Free will, determinism, and punishment. *Psychological Reports*, 93, 1013–1021.
- Hopwood, C. J., & Donnellan, M. B. (2010). How should internal structure of personality inventories be evaluated? *Personality and Social Psychology Review*, 14, 244–254.
- Kaiser, H. F. (1960). The application of electronic computers to factor analysis. *Educational and Psychological Measurement*, 20, 141–151.
- Kane, R. (1996). *The significance of free will*. Oxford University Press.
- Kline, R. B. (2011). *Principles and practice of structural equation modeling* (3rd ed.). New York: Guilford Press.
- Libet, B. (1999). Do we have free will? In B. Libet, A. Freeman, & K. Sutherland (Eds.), *The volitional brain* (pp. 47–57). Exeter: Imprint Academic.
- MacCallum, R. C., Browne, M. W., & Sugawara, H. M. (1996). Power analysis and determination of sample size for covariance structure modeling. *Psychological Methods*, 1, 130–149.
- Monroe, A. E., & Malle, B. F. (2010). From uncaused will to conscious choice: The need to study, not speculate about people's folk concept of free will. *Review of Philosophy and Psychology*, 1, 211–224.
- Murray, D., & Nahmias, E. (2012). Explaining away incompatibilist intuitions. *Philosophy and Phenomenological Research*. Online Early Edition 10–20.
- Nadelhoffer, T. (in press). Dualism, libertarianism, and scientific skepticism about free will. In Sinnott-Armstrong, W. (Ed.), *Moral psychology: Neuroscience, free will, and responsibility* (Vol. 4) (pp. 209–216). MIT Press.
- Nadelhoffer, T., & Tochetti, D. (2013). The potential dark side of free will: Some preliminary findings. In Carusso, G. (Ed.), *Exploring the illusion of free will and moral responsibility*. Lexington Books.
- Nahmias, E., Morris, S., Nadelhoffer, T., & Turner, J. (2006). Is incompatibilism intuitive? *Philosophy and Phenomenological Research*, 73, 28–53.
- Nettler, G. (1959). Cruelty, dignity, and determinism. *American Sociological Review*, 24, 375–384.
- Nettler, G. (1961). Good men, bad men, and the perception of reality. *Sociometry*, 3, 279–294.
- Nichols, S. (2004). Folk psychology of free will. *Mind & Language*, 19, 473–502.
- Nichols, S. M., & Oberle, C. D. (2008). The nature, common usage, and implications of free will and determinism. *Behavior and Philosophy*, 36, 97–111.
- Paulhus, D., & Carey, J. (2011). The FAD-Plus: Measuring lay beliefs regarding free will and related constructs. *Journal of Personality and Assessment*, 11(1), 96–104.
- Paulhus, D. L., & Margesson, A. (1994). *Free will and scientific determinism (FAD-4) scale*. Vancouver: University of British Columbia. Unpublished instrument.
- Pedhazur, E. J., & Schmelkin, L. P. (1991). *Measurement, design and analysis: An integrated approach*. Hillsdale, New Jersey: Lawrence Erlbaum Associates, Publishers.
- Pereboom, D. (2001). *Living without free will*. Cambridge University Press.
- Pronin, E., & Kugler, M. B. (2010). People believe they have more free will than others. *Proceedings of the National Academy of the Science*, 107, 22469–22474.
- Rakos, R. F., Laurene, K. R., Skala, S., & Slane, S. (2008). Belief in free will: Measurement and conceptualization innovations. *Behavior and Social Issues*, 17, 20–39.
- Raykov, T. (1998). On the use of confirmatory factor analysis in personality research. *Personality and Individual Differences*, 24, 291–293.
- Rigoni, D., Kuhn, S., Gaudino, G., Sartori, G., & Brass, M. (2012). Reducing self-control by weakening belief in free will. *Consciousness and Cognition*, 21, 1482–1490.
- Rigoni, D., Kuhn, S., Sartori, G., & Brass, M. (2011). Inducing disbelief in free will alters brain correlates of preconscious motor preparation: The brain minds whether we believe in free will or not. *Psychological Science*, 22(5), 613–618.
- Rigoni, D., Wilquin, H., Brass, M., & Burle, B. (2013). When errors do not matter: Weakening belief in intentional control impairs cognitive reaction to errors. *Cognition*, 127, 264–269.
- Rose, D., & Nichols, S. (2013). The lesson of bypassing. *Review of Philosophy and Psychology*, 1–21.
- Schooler, J., Nadelhoffer, T., Nahmias, E., & Vohs, K. (in press). Measuring and manipulating beliefs and behaviors associated with free will: The good, the bad, and the ugly. In Mele, A. (Ed.), *Surrounding free will*. Oxford University Press.
- Schreiber, J. B., Nora, A., Stage, F. K., Barlow, E. A., & King, J. (2006). Reporting structural equation modeling and confirmatory factor analysis results: A review. *The Journal of Educational Research*, 99(6), 323–337.
- Shepard, J., & Reuter, S. (2012). Neuroscience, choice, and the free will debate. *American Journal of Bioethics Neuroscience*, 3, 7–11.
- Shepherd, J. (2012). Free will and consciousness: Experimental studies. *Consciousness and Cognition*, 21, 915–927.
- Stillman, T. F., & Baumeister, R. F. (2010). Guilty, free, and wise: Determinism and psychopathy diminish learning from negative emotions. *Journal of Experimental Social Psychology*, 46, 951–960.
- Stillman, T. F., Baumeister, R. F., Vohs, K. D., Lambert, N. M., Fincham, F. D., & Brewer, L. E. (2010). Personal philosophy and personal achievement: Belief in free will predicts better job performance. *Social Psychological and Personality Science*, 1, 43–50.
- Strawson, P. (1962). Freedom and resentment. *Proceedings of the British Academy*, 48, 1–25.
- Strawson, G. (1994). The impossibility of moral responsibility. *Philosophical Studies*, 75, 5–24.
- Stroessner, S. J., & Green, C. W. (1990). Effects of belief in free will or scientific determinism on attitudes toward punishment and locus of control. *Journal of Social Psychology*, 130, 789–799.
- van Inwagen, P. (1983). *An essay on free will*. Clarendon Press.
- Viney, W., McIntyre, R., & Viney, D. (1984). Validity of a scale designed to measure beliefs in free will and determinism. *Psychological Reports*, 54, 867–872.
- Viney, W., Parker-Martin, P., & Dotten, S. (1988). Beliefs in free will and determinism and lack of relation to punishment rational and magnitude. *Journal of General Psychology*, 115, 15–23.

- Viney, W., Waldman, D., & Barchilon, J. (1982). Attitudes towards punishment in relation to beliefs in free will and determinism. *Human Relations*, 35, 939–950.
- Vohs, K. D., & Schooler, J. W. (2008). The value of believing in free will: Encouraging a belief in determinism increases cheating. *Psychological Science*, 19, 49–54.
- Waldman, D. A., Viney, W., Bell, P. A., Bennett, J. B., & Hess, S. (1983). Internal and external locus of control in relation to beliefs in free will and determinism. *Psychological Reports*, 53, 631–634.
- Wegner, D. (2002). *The illusion of conscious will*. MIT Press.
- Wood, J. M., Tataryn, D. J., & Gorsuch, R. L. (1996). Effects of under- and overextraction in principal axis factor analysis with varimax rotation. *Psychological Methods*, 1, 354–365.