The impact of free will beliefs on implicit learning

David Wisniewski[1,*], Davide Rigoni[2,3*], Luc Vermeylen[1], Senne Braem[1], Elger Abrahamse[4,5], Marcel Brass[1,6]

1 Department of Experimental Psychology, Ghent University, Ghent, Belgium

2 Hult International Business School, UG Campus, London, United Kingdom

3 Department of Business, Solvay Business School, Free University Brussels, Brussels, Belgium.

4 Department Communication and Cognition, Tilburg University, Tilburg, The Netherlands

5 Department of Educational Sciences, Atlántico Medio University, Spain

6 Berlin School of Mind and Brain/Department of Psychology, Humboldt Universität zu Berlin, Berlin, Germany

* equal contribution

Corresponding author:

David Wisniewski
Department of Experimental Psychology
Ghent University
Henri Dunantlaan 2
9000 Gent, Belgium
Email: david.wisniewski@ugent.be

Abstract

A growing number of studies demonstrate that belief in free will (FWB) is dynamic, and can be reduced experimentally. Most of these studies assume that doing so has beneficial effects on behavior, as FWBs are thought to subdue unwanted automatic processes (e.g. racial stereotypes). However, relying on automatic processes can sometimes be advantageous, for instance during implicit learning (e.g. detecting and exploiting statistical regularities in the environment). In this registered report, we tested whether experimentally reducing FWBs positively affected implicit motor learning. We hypothesized that reducing FWBs would lead to both faster and stronger implicit learning, as measured using the alternating serial reaction time (ASRT) task. While we did show a manipulation effect on free will beliefs, there was no detectable effect on implicit learning processes. This finding adds to the growing body of evidence that free will belief manipulations do not meaningfully affect downstream behavior.

Keywords

Sequence learning, free will, beliefs, registered report

1. Introduction

Most people believe that they possess free will (Wisniewski et al., 2019). That is, they believe they can choose and act in line with their own desires – in principle free from any internal or external constraints (Monroe & Malle, 2010). However, the actual existence of free will has been a hotly debated topic in philosophy for centuries, with more recent contributions from psychology (Aarts & van den Bos, 2011; Baumeister, 2008) and neuroscience (Haggard & Eimer, 1999; Libet et al., 1983; Soon et al., 2008). Regardless of whether one believes that free will exists or not (Fischer et al., 2007; Wegner, 2004), another key question remains: Why is it that so many people believe in free will? Recent estimates show that 82% of US citizens believe in the existence of free will (Wisniewski et al., 2019). One potential explanation is that there are behavioral benefits of believing in free will.

Social psychologists and cognitive neuroscientists alike started to address this question, asking whether believing in free will (or not) has any direct impact on human behavior and cognition. The main hypothesis has been that believing in free will must be beneficial for both the individual and society, because it incites personal responsibility, rational thinking, and self-control, which in turn promote socially desirable behavior (Baumeister et al., 2009). However, some evidence also points to negative consequences of believing in free will, including decreased humility (Earp et al., 2018), increased support for authoritarianism, retributive punishment (Carey & Paulhus, 2013) and potentially economic inequality (Mercier et al., 2020). Still, stronger belief in free will has been associated with improved performance across a variety of domains: correlational studies indicate that people with strong belief in free will perform better in job (Stillman et al., 2010) and academic settings (Feldman et al., 2016). Free will belief even affects behavior at a much lower sensory-motor level, affecting cognitive and neurophysiological markers of self-agency (Lynn et al., 2014;

Rigoni et al., 2011, 2012) and cognitive control (Rigoni et al., 2013, 2015). Recently, much of these findings have been questioned though, as a meta-analysis demonstrated that free will belief manipulations have no detectable downstream effects on e.g. social processes (Genschow et al., 2021).

One key issue with much of the prior literature has been that it focused on tasks in which participants were required to overcome negative automatic processes, such as racial stereotypes (Zhao et al., 2014) or conflicting response tendencies (Rigoni et al., 2012, 2013, 2015). From this perspective, believing in free will mainly seems to foster cognitive control, helping participants to respond in a more agentive manner (Rigoni & Brass, 2014). Yet, there is some evidence conflicting with this view (Earp et al., 2018), and critically, there are many situations in which relying on automatic processes is helpful. For example, (Bocanegra & Hommel, 2014) showed that high levels of cognitive control caused participants to overlook low-level contingencies in stimuli, which would have made task performance much easier if used. Thus, if a strong belief in free will fosters the suppression of automatic processes by way of upregulating cognitive control, we hypothesize that such beliefs will not always be beneficial. They can even be disadvantageous in situations that benefit from the absence of such top-down influences. Thus, even though it seems that disadvantageous effects of reduced free will beliefs are not robust (Genschow et al., 2021), this does not necessarily mean that the same is true for potential advantageous effects of reduced free will beliefs.

Here, we studied the impact of free will belief on automatic processes in learning in a predictable task environment. More precisely, we tested the hypothesis that learning can be hampered by high levels of belief in free will whenever it is based on statistical regularities of the task that are uninstructed. This type of automatic process in learning is typically referred to as implicit or incidental learning, because it describes instances in which learning occurs without the intention to learn and in which the end-product of learning is difficult to express through explicit verbal reports (Cleeremans et al., 1998). To this end, we employed the alternating serial reaction time (ASRT) task (Howard & Howard, 1997). In this task, participants reacted to several different stimuli by pressing one of four different response buttons. Unbeknownst to the participants, responses followed a long sequence (e.g. 1-4-2-3-3-2-4-1-3-4-2-1), with some response sequences occurring more often than others. With time, participants learn the frequent sequences and respond faster to those. Crucially, since learning these sequences seems to rely on implicit learning processes that do not produce easily verbalizable knowledge (Gamble et al., 2014; Howard & Howard, 1997), this task seems ideal to study the benefits of relying on automatic implicit processes to optimize behavior. We hypothesized that reducing free will beliefs would be advantageous for these implicit processes.

Specifically, we hypothesized that implicit learning is malleable to people's belief in free will, and that high levels of belief in free will would lead to reduced implicit learning, as compared to low levels of free will beliefs. One group of participants was subjected to an anti-free will manipulation (i.e. the anti free will group), while another group of participants was subjected to a pro free will manipulation (i.e. the pro free will group), which has been shown to modulate people's beliefs in free will (Rigoni et al., 2011; Seto & Hicks, 2016). Given that implicit learning typically develops over time (Abrahamse et al., 2010), we further expected the two different groups to differ in their implicit learning effect towards the end of practice – when learning was most optimal to observe its modulation.

2. Methods

2.1 Data and code availability

The stage 1 approved protocol can be found here: https://osf.io/m5dwt. All data and code can be found on the Open Science Framework: https://osf.io/6s4yz/.

2.2 Sample characteristics

As preregistered in the stage 1 registered report, we acquired data from n = 275 participants in each of two free will manipulation groups (anti free will, pro free will, for an explanation see below). The sample size was determined through a power analysis ($\alpha = 0.05$ and $1\text{-}\beta = 0.90$), with the goal of detecting three effects reliably: a) an effect of the free will belief manipulation, b) an implicit learning effect in the ASRT task, and c) an effect of the free will belief manipulation on the ASRT learning effects. For the free will belief manipulation effect, we assumed a small effect size (d = 0.25) based on a recent meta-analysis of the field (Genschow et al., 2021). For the free will belief effect on implicit learning, we followed the approach of Vékony et al. (2020), and conservatively chose a smaller than reported effect size of d = 0.50. For the estimation of the overall implicit learning effect, as measured using an ASRT task, we relied on the same publication, and again conservatively chose a smaller than reported effect size of d = 0.50. We then computed the sample size such that the probability of detecting each of the three effects equaled $1\text{-}\beta = 0.90$, which led to the final sample size of n = 275.

Participants were recruited through the Prolific Academic platform, and the inclusion criteria were: age between 18 and 60, first language: English, at least 50 prior submissions on the platform, and at least a 90% approval rate.

Three separate exclusion criteria were then applied. First, we tested for missing data in the free will inventory (FWI, (Nadelhoffer et al., 2014)). Secondly, 3 independent raters screened the responses in the free will manipulation procedure for low-effort responses. If all 3 consistently flagged responses as low effort, that participant was marked for exclusion. Thirdly, if after trial-level removal of data (errors, trials following errors, repetition and trill trials, see below for more details) the remaining number of trials was lower than 3SD below the sample average, that participant was marked for exclusion. Meeting any of these three removal criteria led to a the exclusion of the full data set and that participant was then replaced in the sample. Overall, 64 participants were excluded and replaced this way. The final sample consisted of 276 female, 274 male, and 1 other participant. The average age was 34.85 years (range 19 – 60).

2.3 Procedure

2.3.1 Free will belief manipulation

The experiment was performed online in a web browser. Upon signing up to the experiment and giving their informed consent, participants first underwent a free will belief manipulation. They were randomly assigned to one of two groups: anti free will (FW-) or pro free will (FW+). The free will belief manipulation we used has been successfully employed in past research (Seto & Hicks, 2016), and combines two different manipulations to increase the effect size (reading a text (Vohs & Schooler, 2008), and thinking about different statements (Alquist et al., 2013)). The FW- group read a short text explaining what free will is, and that leading researchers have shown that most of our behavior is determined by environmental factors that are not under our direct control. Afterwards, they were given 10 statements arguing for the non-existence of free will (e.g. “Science has demonstrated that free will is an illusion.”, “Everything a person does is a direct consequence of their environment and genetic makeup.”). They were then asked to think about why these statements are true based on their own experiences and select 3 statements from the list that have proven especially true in their own life. Finally, they were asked to write down how each of the chosen statements is true and provide specific examples from their life with as much detail as possible. The FW+ group also read a short text, explaining what free will is and how leading researchers have shown that our behavior is mostly determined by our own choices. Afterwards they read 10 statements arguing for the existence of free will (e.g. “I demonstrate my free will everyday when I make decisions.”, “I have free will to control my actions and, ultimately, to control my destiny in life.”). They again thought about why these statements are true, select 3 and describe how they have proven true in their own life with as much detail as possible.

2.3.2 Free will inventory

After the manipulation, participants completed the Free Will Inventory (part 1, (Nadelhoffer et al., 2014)). This 15 item questionnaire measures belief in free will (FW), i.e. whether participants think that free will exists. It further measures belief in determinism (DE), i.e. that all events in the physical world are fully determined by prior events, and belief in dualism (DU), i.e. that the mind/soul is a non-physical entity that cannot be reduced to the brain. Each item of the FWI is scored on a 7-point Likert scale (1 = strongly disagree, 7 = strongly agree), with higher scores indicating stronger beliefs. For illustration purposes, each item was rescaled to -1 (strongly disagree) and 1 (strongly agree), with 0 indicating “neither believe nor disbelieve”. Thus, positive values indicate beliefs, while negative values indicate disbelief.

2.3.3 ASRT

Participants then performed the ASRT task ((Gamble et al., 2014; Howard & Howard, 1997), programmed in jsPSych (de Leeuw, 2015), modified from from (Vékony et al., 2020), which is a variant of the SRT task (Nissen & Bullemer, 1987). Four open circles were displayed on a white background at the center of a computer screen (see Figure 1). Each trial started with one of the circles filled in black. Participants were instructed to respond as fast as possible by pressing the corresponding key on the keyboard (“s”, “f”, “j”, and “l”, operated using the left and right index and middle fingers). Each circle remained filled in until the correct response was provided, and after an inter-stimulus interval of 120ms a different circle was filled in. Unknown to the participant, the sequence of the filled in circles followed a pre-defined pattern of 4 locations. The pattern occurred on alternating events. For example, if the locations of the four circles were defined as 1, 2, 3 and 4 from left to right, one pattern that occurred was 1r3r2r4r, such that the first event occurred in location 1, the second event occurred randomly in one of the four locations, the third event occurred in location 3, and so on, according to the pattern. The resulting twenty-four possible permutations were randomized across participants, with each participant encountering only one pattern throughout the task. The first eight trials of each block occurred at random locations, after which the pattern kept repeating until the end of the block, e.g., ‘3r1r2r4r3r1r2r’, and so on. We administered 28 blocks with 64[1] trials each, for a total of 1792 experimental trials.

As a result of the alternating pattern, some sequences (Triplets) occurred with a higher probability than others. For example, if the pattern was ‘1r3r2r4r’, the ‘3r2’ triplet was a high probability triplet, since triplets beginning with 3 and ending with 2 occurred every time the pattern was

[1] In the stage 1 protocol, we mistakenly wrote we would administer 58 trials in each block. In fact, we administered 64 trials in each block (8 random trials + 7 repetitions of the 8-trial sequence), which we now corrected in the manuscript.

repeated. Conversely, the sequence ‘2r3’ was a low probability triplet since it only occurred occasionally when 2 and 3 occurred as random events. Implicit sequence learning is therefore quantified as a ‘triplet type’ effect, as the response to the third event in high probability vs. low probability triplets. Repetition (‘222’) and Trill (‘2r2’) triplets were not included in the analysis (Gamble et al., 2014). In addition, incorrect response trials and correct trials following an incorrect response in one of the two preceding trials were excluded. Following Howard & Howard (1997), feedback was displayed at the end of each block. In order to encourage fast and accurate responding, if the accuracy was above 93%, participants were instructed to respond faster. If accuracy was below 92%, they were instructed to respond more accurately. For accuracies between 92% and 93%, they were instructed to continue as before.

2.3.4 Debriefing

After finishing the ASRT task, participants were asked whether they were aware that the sequence of the stimuli followed a pattern by completing a 3-items questionnaire (i.e. “Did you notice anything about how stimuli were presented on the screen, or with respect to how they succeeded each other in time? If yes, please explain”; “Did you notice that there was a pattern in the sequence of the stimuli presented on the screen? If yes, please describe this pattern”; “Have you made an effort to use this pattern to predict the location of the next stimulus? If so, did this help you to perform quicker in the task?”). Participants were also told that the experiment consisted of two groups, how these two groups differed, and to which group they belonged.

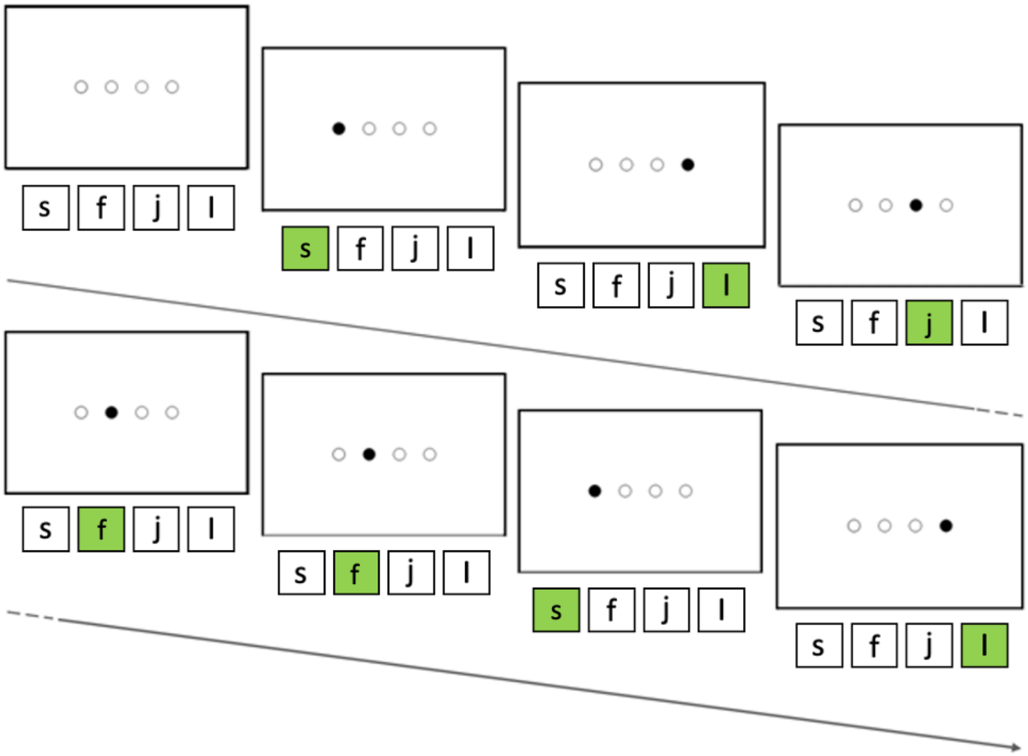

Figure 1. Design. Example of a sequence of trials of the ASRT task. Participants were required to keep their fingers on the 4 green colored keys of the keyboard. In this example, the repeating pattern is '1r3r2r4r'. This procedure resulted in some runs of three events (i.e. triplets) being very frequent, e.g. 3 × 2, where × is any event, and others being very infrequent, e.g. 2 × 3 (i.e. in this example, this would only occur when 2 and 3 occur as random events).

2.4 Planned analyses

All analyses reported here were preregistered in the stage 1 protocol.

2.4.1 Hypothesis 1: Manipulation check

We expected the free will belief manipulation to reduce free will beliefs in the FW- group, as compared to the FW+ group. For this purpose, we first computed the mean score of the free will subscale of the FWI for each participant, and then computed a one-sided, two-sample t-test, in order to test whether free will beliefs decreased in the FW- group as compared to the FW+ group.

2.4.2 Hypothesis 2: Belief effects on ASRT effect strength

In order to test whether implicit learning was stronger in the FW- group, as compared to the FW+ group, we first split all ASRT trials into 4 quartiles of 7 blocks. We expected the ASRT effect to build up across the session, and therefore excluded the first 25% of ASRT trials, where finding a learning effect is unlikely. We summarized reaction time and error rate results by computing the inverse efficiency score (IES) for each participant and condition (IES = reaction time / (1 – error rate). The IES for low frequency triplet trials was then be subtracted from the IES for high frequency triplet trials, to derive a learning score. Values around 0 thus indicated equal performance in high and low frequency triplet trials. Negative values indicated better performance in high frequency than in low frequency triplet trials. These scores were entered into a one-sided two-sample t-test, in order to assess whether learning scores were higher in the FW- group, as compared to the FW+ group.

2.4.3 Hypothesis 3: Belief effects on learning in ASRT task

Lastly, we directly assessed learning rate in the ASRT task. Possibly, decreasing free will beliefs not only led to stronger implicit learning effects, but also accelerated the implicit learning process. Thus, it might be that participants in the FW- group learned faster than participants in the FW+ group. In order to test this hypothesis, we first split the ASRT data into 4 quantiles of 7 blocks, and extracted learning scores (computed just like in Hypothesis 2) within each of the 4 quantiles. Then, we conducted a two-factorial ANOVA, using the factors free will condition (FW-, FW+, between subjects factor), and quantile (1^{st}, 2^{nd}, 3^{rd}, 4^{th} quantile of the ASRT task), using the learning score as the dependent variable. First, we expected to see a learning effect in the ASRT task. This should be reflected in a significant main effect of quantile. Second, the speed of learning should differ between groups, with faster learning in the FW- group, which should be reflected in a bin x group interaction.

2.5 Control analyses

2.5.1 FWI

We also performed a number of control analyses. First, in order to assess whether we replicated previous findings showing a strong belief in free will, we tested whether participants believe in free will or not in this experiment. For this purpose, we performed a one-sided t-test on the free will sub-scale of the FWI, testing whether values were larger than zero (indicating belief in free will). For descriptive purposes, we computed the proportion of participants showing positive values as well.

2.5.2 ASRT

Additionally, we tested whether ASRT performance was comparable across groups, to rule out any confounding influences. Although unlikely, it might be that one group showed fast or more accurate performance in the ASRT, which could complicate interpretation of results across groups. For this purpose, we tested whether the overall IES (across all triplet types) differed between groups, using a two-sided, two-sample t-test. In the unlikely event of a significant difference, overall IES could have been added as a covariate to the planned analyses, in order to correct for any global performance differences between groups.

2.6 Pre-registered exploratory analyses

2.6.1 Correlation analyses

Given recent evidence showing that free will belief manipulation effects are small (Ewusi-Boisvert & Racine, 2018; Genschow et al., 2021), some have used correlational approaches instead of experimental manipulations to study lay views on free will (e.g. (Wisniewski et al., 2019)). To assess the validity of this correlational approach, we performed a number of additional exploratory analyses. We pooled data from both groups into a single large group and then tested our main hypotheses using a correlational approach. Given our large sample size (N = 550), we had enough statistical power to detect even small correlations, $r > 0.07$ ($\alpha = 0.05$, $1-\beta = 0.90$).

We first correlated learning scores (computed like in Hypothesis 2) with the free will sub-scale of the FWI. Positive correlations would indicate evidence for a relation between free will beliefs and implicit learning processes. Similarly, we then assessed the relationship between implicit learning speed and free will beliefs. For this purpose, we first computed learning scores separately for the 4 quantiles of the ASRT task. Then, we fit a linear model to these learning scores, separately for each participant, with quantiles being the predictor. The slope of this model was our estimate of learning speed for each participant. This estimate was then correlated with free will beliefs, and a positive correlation would show evidence for a relation of implicit learning speed and free will beliefs (Hypothesis 3).

2.6.2 Manipulation effects on determinism and dualism beliefs

The FWI contains three different subscales: free will, determinism, and dualism. In the planned analyses, we mainly focus on the free will subscale. We also performed an additional exploratory analysis, to shed more light on the role of determinism and dualism beliefs. The free will belief manipulation we used relies at least partly on arguing for physical determinism ("Everything a person does is a direct consequence of their environment and genetic makeup."). Thus, it might be that participants' FWBs decreased because the manipulation increased their belief in determinism.

A similar argument can be made for dualism as well. From this perspective, determinism and dualism beliefs might have also been affected by the free will belief manipulation. We tested for such possible effects by running the same analysis as for Hypothesis 1, only using determinism and dualism beliefs as dependent variables. This additional analysis provided more insight into how the free will belief manipulation worked, and whether it was specific to free will beliefs or not.

3. Results

3.1 Hypothesis 1

For this and all other Hypotheses, we first report results for the pre-registered analyses. In case we performed additional analyses which were not part of the approved protocol, this is marked as a non-pre-registered analysis and reported in a separate section.

3.1.1 Pre-registered analyses

In order to test whether our free will belief manipulation was successful, we compared free will sub-scale scores (FW) of the FWI between groups (Figure 2). We found a significant effect of group on the free will sub-scale, $t(546.94) = -7.48$, $p < 0.001$, $d = -0.64$.

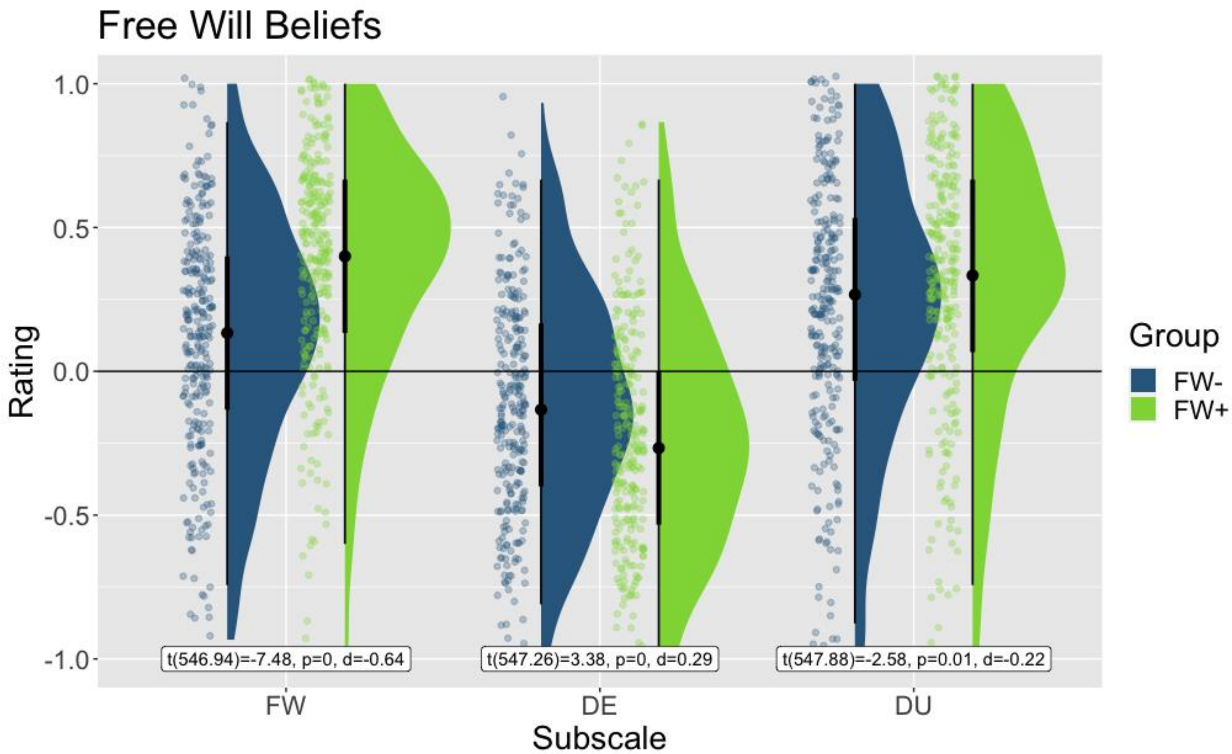

Figure 2: Free will belief scores. This plot illustrates responses on the Free Will Inventory, separately for the free will (FW), determinism (DE), and dualism (DU) sub-scale, and separately for the anti free will (FW-) and control (FW+) groups. Negative values indicate disbelief, positive values indicate belief. Dots represent raw data, distributions are summarized by boxplots (thick black lines and dot) and violin plots. Results of a two-sample t-test comparing both groups are shown at the bottom.

3.2 Hypothesis 2

3.2.1 Pre-registered analyses

In a next step, we tested whether implicit learning was stronger in the FW- group than in the FW+ group. For this purpose, we computed inverse efficiency scores for both high and low frequency triplets. A learning score was then created by computing their difference ($IES_{high\_frequency}$ – $IES_{low\_frequency}$), with negative values indicating stronger implicit learning. A one-sided two-sample t-test was then used to compare learning scores between groups, which yielded no significant difference, t(481.2) = -0.5, p = 0.31, d = -0.04 (Figure 3). We thus found no evidence for a difference in implicit learning strength between the groups.

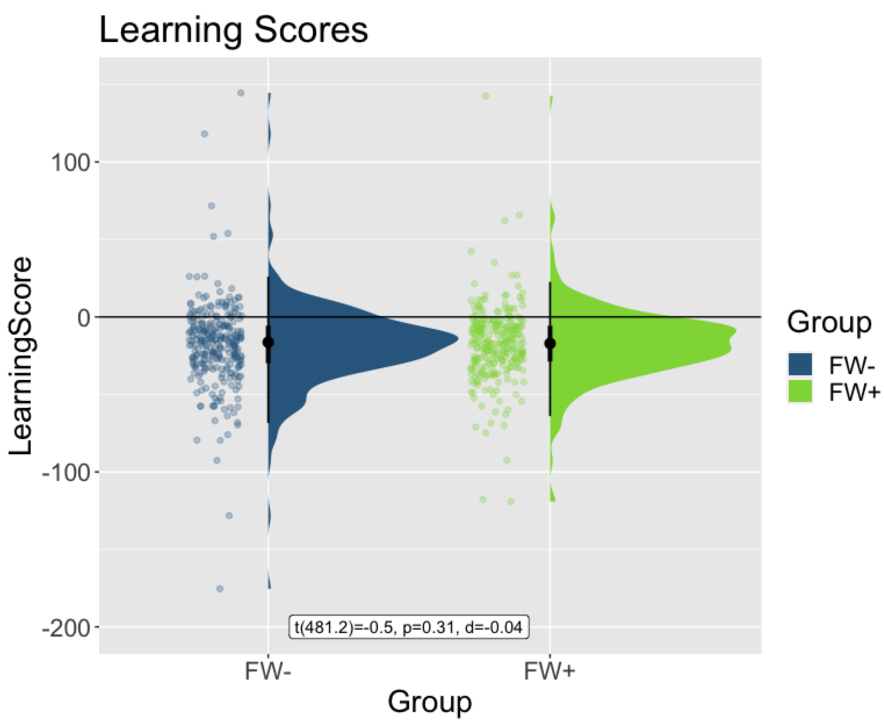

Figure 3: Learning scores, separately for the anti free will (FW-) and control (FW+) groups. Negative values indicate stronger implicit learning. Dots represent raw data, distributions are summarized by boxplots (thick black lines and dot) and violin plots. Results of a two-sample t-test comparing both groups are shown at the bottom.

3.2.2 Non-pre-registered analyses

As can be seen in Figure 3, some participants were clear outliers in the learning score distribution, and we therefore repeated the same analysis after removing outliers. Outliers were defined as values above the sample mean + 1.5 * IQR (inter-quartile-range), or below the sample mean – 1.5 * IQR. After removing outliers, we still observed no significant difference between groups, t(516.53) = -0.16, p = 0.43, d = -0.01.

We also tested whether participants showed implicit learning effects at all in this experiment, by computing a one-sided t-test against 0, separately for the FW- and FW+ group. Both the FW- group, $t(274) = -10.82$, $p < 0.001$, $d = -0.65$, and the FW+ group, t(274) = =6.71, $p < 0.001$, $d = -0.41$, showed clear evidence for implicit learning in the ASRT.

3.3 Hypothesis 3

3.3.1 Pre-registered analyses

We then directly assessed the learning rate in the ASRT, hypothesizing that implicit learning would be faster in the FW- group than in the FW+ group. For this purpose, we plot the ASRT data into 4 quantiles (of 7 experimental blocks each), and extracted learning scores separately for each quantile and group. A two-factorial ANOVA (group x quantile) yielded a significant effect of quantile, $F(2.62,1433.13) = 6.26$, $p < 0.001$, generalized $\eta^2 = 0.008$, reflecting the expected stronger implicit learning effects in later blocks (Figure 4). The main effect of group showed no significant results, $F(1,548) = 0.9$, $p = 0.34$, $\eta^2 = 0.0004$. Should the FW- group exhibit faster implicit learning than the FW+ group, we would expect to see a significant group x quantile interaction, which we did not in the current data set, $F(3,1644) = 0.35$, $p = 0.79$, $\eta^2 = 0.0005$. Therefore, we found no evidence for a difference in implicit learning speed between both groups.

3.3.1 Non-pre-registered analyses

Although we found no group x quantile interaction in our data, we further explored implicit learning effects in different parts of the experiment. It might be that e.g. a group difference would only be visible in the last quantile, when learning effects should be strongest. We found no evidence for this hypothesis however, $t(445.54) = -0.14$, $p = 0.45$, $d = -0.01$ (Figure 4). Furthermore, the FWB manipulation might have sped up implicit learning, but this effect might have been detectable only in the early stages of the experiment – when the learning curve is steepest. To test this hypothesis, we compared both groups in the first quantile only, but this analysis yielded a non-significant result, t(546.83) = -1.59, p = 0.06, d = -0.14.

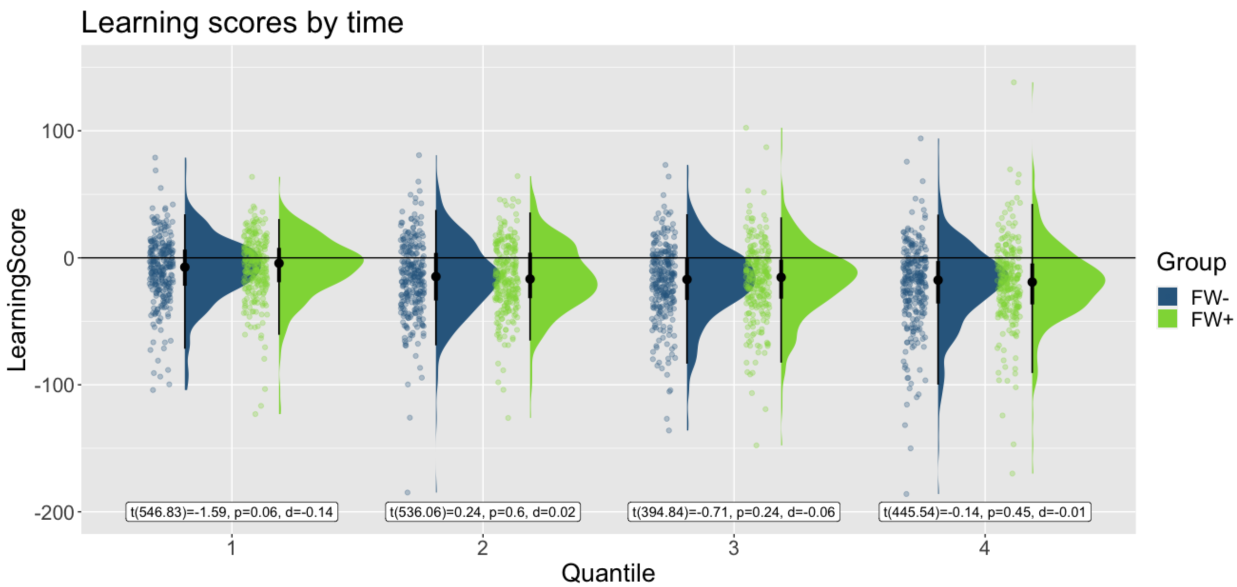

Figure 4: Learning scores by time, separately for the FW- group (afw) and the FW+ group (ctrl), and for each quantile of the experiment. Negative values indicate stronger implicit learning. Dots represent raw data, distributions are summarized by boxplots (thick black lines and dot) and violin plots. Results of two-sample t-tests comparing both groups in each quantile are shown at the bottom.

3.4 Control analyses

3.4.1 FWI

In a first control analysis, we tested whether participants believed in free will or not. We found that 72.18% of the sample believed in free will, and that free will scores were significantly above zero, t(549) = 88.90, p < 0.001, d = 3.79.

3.4.2 ASRT

To rule out potential confounds in ASRT performance between groups, we compared inverse efficiency scores between groups (across both triplet types). We found no significant difference, $t(446.45) = -0.74$, $p = 0.46$, $d = -0.06$, indicating that overall ASRT performance was comparable between both groups.

3.5 Pre-registered exploratory analyses

3.5.1 Correlation analyses

After pooling data from both groups, we correlated learning scores (computed like in Hypothesis 2) with each of the three sub-scales of the FWI. This approach might be more sensitive to detect relations between implicit learning and free will beliefs than group comparisons. We found no significant correlation with free will beliefs, $r = 0.04$, $p = 0.41$, no significant correlation with determinism beliefs, $r = -0.004$, $p = 0.92$, and no significant correlation with dualism beliefs, $r = 0.03$, $p = 0.45$. We also correlated learning speed with each of the three sub-scales of the FWI. However, we found no significant correlation with free will beliefs, $r = 0.04$, $p = 0.09$, no significant correlation with determinism beliefs, $r = -0.001$, $p = 0.94$, and no significant correlation with dualism beliefs, $r = 0.01$, $p = 0.51$. Thus, even when pooling data across groups, we were unable to detect a relationship between free will beliefs and implicit learning.

3.5.2 Manipulation effects on determinism and dualism beliefs

In order to test whether our free will belief manipulation also affected determinism and dualism beliefs, we compared the respective sub-scale scores of the FWI between groups. We found significant effects on determinism beliefs, $t(547.26) = 3.38$, $p < 0.001$, $d = 0.29$, and on dualism beliefs, $t(547.88) = -2.58$, $p = 0.01$, $d = -0.22$. Therefore, our experimental manipulation was not specific to free will beliefs only, and affected related beliefs as well.

4. Discussion

Much of the past research of free will beliefs focused on the potentially negative consequences of low free will beliefs, but effects on e.g. social behavior have not proven robust (Buttrick et al., 2020; Genschow et al., 2021). Here, we questioned the prevalent assumption that a strong belief in free will has mostly positive consequences (Rigoni & Brass, 2014; Vohs & Schooler, 2008). This is often based on the idea that weakening free will beliefs downregulates cognitive control and thus allows unwanted automatic processes to more strongly affect behavior (disinhibition). This might have negative consequences in some cases (Baumeister et al., 2009; Rigoni et al., 2012; Vohs & Schooler, 2008), but might also have positive consequences in other cases (Earp et al., 2018). Here we focused on one such scenario. Implicit learning relies on automatic processes, and is negatively affected by strong cognitive control (Bocanegra & Hommel, 2014). We reasoned that reducing free will beliefs would make implicit learning more effective in an ASRT task (Howard & Howard, 1997). Although we found robust evidence for implicit learning as well as robust evidence for a successful free will belief manipulation, we found no evidence for an effect of free will beliefs on implicit learning. This was true for both the overall strength of the implicit learning effect, as well as the speed with which learning occurs, and was despite the fact that our manipulation effect was unusually strong, $d = -.64$, where effect sizes are usually within the $d = 0.20 – 0.30$ range (Genschow et al., 2021). Even after collapsing data across groups and using a correlational approach, which increased statistical power to detect even very small correlations ($r > 0.07$), we found no evidence for an association between free will beliefs and implicit learning.

We also explored potential effects in a time-resolved fashion, testing whether free will beliefs affect implicit learning specifically in the beginning of the experiment, where the learning curve should be steepest. Here, we found a small effect in the expected direction ($d = -0.14$), which failed to reach the significance threshold however ($p = 0.06$). This might be interpreted as tentative evidence for free will belief effects on early stages of implicit motor learning, but this would need to be replicated in a future study optimized to detect such small effects and even then, it remains doubtful whether such small effects are meaningful.

Given these null results, we can only speculate as to why we did not observe the expected effects. Since this study was adequately powered to detect even small effects, we will assume that the results represent a true null finding. Effects smaller than those we were able to detect would be too small to interpret meaningfully in any case. One potential explanation would be that high-level beliefs are simply too far removed from low-level motor behavior. Although initially there was evidence for such a link, showing that free will beliefs affected automatic post-error adjustment in behavior (Rigoni et al., 2013), more recently it has proven difficult to replicate these results (Eben et al., 2020). Taken together with the current results, this suggests that automatic motor processes are not affected by free will beliefs. An alternative, less strong interpretation of our results might be that the free will belief manipulation did affect motor behavior, but that these effects were weak and decayed quickly. This would be compatible with the anecdotal evidence we found for an effect in the beginning of the ASRT task. Given this assumption, it might then be that the free will belief manipulation did not fundamentally alter beliefs (even for the duration of the experiment), but rather transiently altered cognitive processes akin to a priming effect. This could be tested in future work by modifying the experimental manipulation procedure. We might repeat the free will belief manipulation several times in between blocks of the ASRT task, refreshing the pro and anti-free will messages, and we would predict transient effects immediately following the manipulation which will decay quickly. If the manipulation method affected beliefs more fundamentally, we would expect to see no such decay however. Clearly, this is speculative at the moment and will require more work in the future.

More broadly, we set out to test whether reducing free will beliefs might have positive effects on behavior. Even in a highly powered, pre-registered study that showed surprisingly strong manipulation effects, we failed to find evidence for an effect on motor behavior which mirrors previous results showing a lack of negative effects on behavior. Future work will have to show whether this is due to inadequate experimental techniques, or due to free will beliefs not affecting downstream behavior after all.

5. References

Aarts, H., & van den Bos, K. (2011). On the Foundations of Beliefs in Free Will: Intentional Binding and Unconscious Priming in Self-Agency. *Psychological Science*, *22*(4), 532–537. https://doi.org/10.1177/0956797611399294

Abrahamse, E. L., Jiménez, L., Verwey, W. B., & Clegg, B. A. (2010). Representing serial action and perception. *Psychonomic Bulletin & Review*, *17*(5), 603–623. https://doi.org/10.3758/PBR.17.5.603

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*(1), 80–86. https://doi.org/10.1016/j.jesp.2012.08.015

Baumeister, R. F. (2008). Free Will in Scientific Psychology. *Perspectives on Psychological Science*, *3*(1), 14–19. https://doi.org/10.1111/j.1745-6916.2008.00057.x

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*(2), 260–268. https://doi.org/10.1177/0146167208327217

Bocanegra, B. R., & Hommel, B. (2014). When Cognitive Control Is Not Adaptive. *Psychological Science*, *25*(6), 1249–1255. https://doi.org/10.1177/0956797614528522

Buttrick, N. R., Aczel, B., Aeschbach, L. F., Bakos, B. E., Brühlmann, F., Claypool, H. M., Hüffmeier, J., Kovacs, M., Schuepfer, K., Szecsi, P., Szuts, A., Szöke, O., Thomae, M., Torka, A.-K., Walker, R. J., & Wood, M. J. (2020). Many Labs 5: Registered Replication of Vohs and Schooler (2008), Experiment 1. *Advances in Methods and Practices in Psychological Science*, *3*(3), Article 3. https://doi.org/10.1177/2515245920917931

Carey, J. M., & Paulhus, D. L. (2013). Worldview Implications of Believing in Free Will and/or Determinism: Politics, Morality, and Punitiveness. *Journal of Personality*, *81*(2), 130–141. https://doi.org/10.1111/j.1467-6494.2012.00799.x

Cleeremans, A., Destrebecqz, A., & Boyer, M. (1998). Implicit learning: News from the front. *Trends in Cognitive Sciences*, *2*(10), 406–416. https://doi.org/10.1016/S1364-6613(98)01232-7

de Leeuw, J. R. (2015). jsPsych: A JavaScript library for creating behavioral experiments in a Web browser. *Behavior Research Methods*, *47*(1), 1–12. https://doi.org/10.3758/s13428-014-0458-y

Earp, B. D., Everett, J. A., Nadelhoffer, T., Caruso, G. D., Shariff, A., & Sinnott-Armstrong, W. (2018). *Determined to Be Humble? Exploring the Relationship Between Belief in Free Will and Humility*. https://doi.org/10.31234/osf.io/3bxra

Eben, C., Chen, Z., Cracco, E., Brass, M., Billieux, J., & Verbruggen, F. (2020). Are post-error adjustments influenced by beliefs in free will? A failure to replicate Rigoni, Wilquin, Brass and Burle, 2013. *Royal Society Open Science*, *7*(11), Article 11. https://doi.org/10.1098/rsos.200664

Ewusi-Boisvert, E., & Racine, E. (2018). A Critical Review of Methodologies and Results in Recent Research on Belief in Free Will. *Neuroethics*, *11*(1), 97–110. https://doi.org/10.1007/s12152-017-9346-3

Feldman, G., Chandrashekar, S. P., & Wong, K. F. E. (2016). The freedom to excel: Belief in free will predicts better academic performance. *Personality and Individual Differences*, *90*, 377–383. https://doi.org/10.1016/j.paid.2015.11.043

Fischer, J. M., Kane, R., Pereboom, D., & Vargas, M. (2007). *Four Views on Free Will*. Wiley-Blackwell.

Gamble, K. R., Lee, J. M., Howard, J. H., & Howard, D. V. (2014). Effects of priming goal pursuit on implicit sequence learning. *Experimental Brain Research*, *232*(11), 3635–3643. https://doi.org/10.1007/s00221-014-4054-2

Genschow, O., Cracco, E., Schneider, J., Protzko, J., Wisniewski, D., Brass, M., & Schooler, J. (2021). *Meta-analysis on belief in free will manipulations* [Preprint]. PsyArXiv. https://doi.org/10.31234/osf.io/quwgr

Haggard, & Eimer, M. (1999). On the relation between brain potentials and the awareness of voluntary movements. *Experimental Brain Research*, *126*(1), 128–133. https://doi.org/10.1007/s002210050722

Howard, J. H., & Howard, D. V. (1997). Age differences in implicit learning of higher order dependencies in serial patterns. *Psychology and Aging*, *12*(4), 634–656. https://doi.org/10.1037//0882-7974.12.4.634

Libet, B., Gleason, C. A., Wright, E. W., & Pearl, D. K. (1983). Time of conscious intention to act in relation to onset of cerebral activity (readiness-potential) the unconscious initiation of a freely voluntary act. *Brain*, *106*(3), 623–642.

Lynn, M. T., Muhle-Karbe, P. S., Aarts, H., & Brass, M. (2014). Priming determinist beliefs diminishes implicit (but not explicit) components of self-agency. *Consciousness Research*, *5*, 1483. https://doi.org/10.3389/fpsyg.2014.01483

Mercier, B., Wiwad, D., Piff, P. K., Aknin, L. B., Robinson, A. R., & Shariff, A. (2020). Does Belief in Free Will Increase Support for Economic Inequality? *Collabra: Psychology*, *6*(1). https://doi.org/10.1525/collabra.303

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*(2), 211–224. https://doi.org/10.1007/s13164-009-0010-7

Nadelhoffer, T., Shepard, J., Nahmias, E., Sripada, C., & Ross, L. T. (2014). The free will inventory: Measuring beliefs about agency and responsibility. *Consciousness and Cognition*, *25*, 27–41. https://doi.org/10.1016/j.concog.2014.01.006

Nissen, M. J., & Bullemer, P. (1987). Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology*, *19*(1), 1–32. https://doi.org/10.1016/0010-0285(87)90002-8

Rigoni, D., & Brass, M. (2014). From Intentions to Neurons: Social and Neural Consequences of Disbelieving in Free Will. *Topoi*, *33*(1), 5–12. https://doi.org/10.1007/s11245-013-9210-y

Rigoni, D., Kühn, S., Gaudino, G., Sartori, G., & Brass, M. (2012). Reducing self-control by weakening belief in free will. *Consciousness and Cognition*, *21*(3), 1482–1490. https://doi.org/10.1016/j.concog.2012.04.004

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. https://doi.org/10.1177/0956797611405680

Rigoni, D., Pourtois, G., & Brass, M. (2015). 'Why should I care?' Challenging free will attenuates neural reaction to errors. *Social Cognitive and Affective Neuroscience*, *10*(2), 262–268. https://doi.org/10.1093/scan/nsu068

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*(2), 264–269. https://doi.org/10.1016/j.cognition.2013.01.009

Seto, E., & Hicks, J. A. (2016). Disassociating the Agent From the Self: Undermining Belief in Free Will Diminishes True Self-Knowledge. *Social Psychological and Personality Science*, *7*(7), 726–734. https://doi.org/10.1177/1948550616653810

Soon, C. S., Brass, M., Heinze, H.-J., & Haynes, J.-D. (2008). Unconscious determinants of free decisions in the human brain. *Nature Neuroscience*, *11*(5), 543–545. https://doi.org/10.1038/nn.2112

Stillman, T. F., Baumeister, R. F., Vohs, K. D., Lambert, N. M., Fincham, F. D., & Brewer, L. E. (2010). Personal Philosophy and Personnel Achievement: Belief in Free Will Predicts Better Job Performance. *Social Psychological and Personality Science*, *1*(1), 43–50. https://doi.org/10.1177/1948550609351600

Vékony, T., Marossy, H., Must, A., Vécsei, L., Janacsek, K., & Nemeth, D. (2020). Speed or Accuracy Instructions During Skill Learning do not Affect the Acquired Knowledge. *Cerebral Cortex Communications*, *1*(1). https://doi.org/10.1093/texcom/tgaa041

Vohs, K. D., & Schooler, J. W. (2008). The Value of Believing in Free Will: Encouraging a Belief in Determinism Increases Cheating. *Psychological Science*, *19*(1), 49–54. https://doi.org/10.1111/j.1467-9280.2008.02045.x

Wegner, D. M. (2004). Precis of the illusion of conscious will. *Behavioral and Brain Sciences*, *27*(05), 649–659.

Wisniewski, D., Deutschländer, R., & Haynes, J.-D. (2019). Free will beliefs are better predicted by dualism than determinism beliefs across different cultures. *PLOS ONE*, *14*(9), e0221617. https://doi.org/10.1371/journal.pone.0221617

Zhao, X., Liu, L., Zhang, X., Shi, J., & Huang, Z. (2014). The Effect of Belief in Free Will on Prejudice. *PLoS ONE*, *9*(3). https://doi.org/10.1371/journal.pone.0091572

6. Data accessibility

The stage 1 approved protocol can be found here: https://osf.io/m5dwt. All data and code can be found on the Open Science ramework: https://osf.io/6s4yz/

7. Competing Interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

8. Acknowledgments

D.W. was supported by Flemish Science Foundation (FWO) grants number FWO.KAN.2019.0023.01 and 665501, by the European Union's Horizon 2020 Research and Innovation Program under the Marie Skłodowska-Curie grant no. 665501, and the Special Research Fund of Ghent University (BOF). L.V. was supported by FWO grant number 11H5619N. S.B. was supported by an ERC Starting grant (European Union's Horizon 2020 research and innovation programme, Grant agreement 852570). M.B. was supported by an Einstein Strategic Professorship (Einstein Foundation Berlin)