People Judge Others to Have More Voluntary Control Over Beliefs Than They Themselves Do

Corey Cusimano
Princeton University

Geoffrey P. Goodwin
University of Pennsylvania

People think other individuals have considerable control over what they believe. However, no work to date has investigated how people judge their own belief control, nor whether such judgments diverge from their judgments of others. We addressed this gap in 7 studies and found that people judge others to be more able to voluntarily change what they believe than they themselves are. This occurs when people judge others who disagree with them (Study 1) as well as others who agree with them (Studies 2–5, 7), and it occurs when people judge strangers (Studies 1, 2, 4, and 5) as well as close others (Studies 3 and 7). It appears not to be explained by impression management or self-enhancement motives (Study 3). Rather, there is a discrepancy between the evidentiary constraints on belief change that people access via introspection, and their default assumptions about the ease of voluntary belief revision. That is, people tend spontaneously to think about the evidence that supports their beliefs, which leads them to judge their beliefs as outside their control. But they apparently fail to generalize this sense of constraint to others, and similarly fail to incorporate it into their generic model of beliefs (Studies 4–7). We discuss the implications of our findings for theories of ideology-based conflict, actor–observer biases, naïve realism, and ongoing debates regarding people's actual capacity to voluntarily change what they believe.

Keywords: actor observer, attribution, belief, control, theory of mind

Supplemental materials: http://dx.doi.org/10.1037/pspa0000198.supp

The language of belief is infused with attributions of control. People talk about what they and others *can* believe ("you can believe what you want, but if you ignore the rocks you'll be badly hurt"), what they *choose* to believe ("I choose to believe in the inherent intelligence and good sense of the average Malaysian voter," "One can choose to believe or not believe in God"), and what they *intend* or *decide* to believe ("He said he didn't know and I intend to believe him," "People are going to decide to believe what they want to believe").[1] These locutions express attributions of control that play an important role in how people evaluate and react to others' beliefs. For instance, recent studies have shown that people commonly attribute a high degree of intentional control to others over what they believe (Cusimano & Goodwin, 2019). That is, people incline toward judging that others (a) intentionally

choose what they believe, (b) have control over what they believe, and (c) can choose to stop holding specific beliefs should they want to. Furthermore, just as with behavior, people appear to rely on these attributions of control when they evaluate belief holders. Individuals who attribute more intentional control to others over what they believe are more likely to blame those others for holding immoral or unjustified beliefs (Cusimano & Goodwin, 2019). Thus, in keeping with the fundamental role that attributions of control play in determining how we explain and judge other people's behavior (Heider, 1958; Malle, Guglielmo, & Monroe, 2014; Weiner, 1995), control attributions appear to occupy a similarly fundamental role when it comes to beliefs.

In the present studies, we investigate attributions of voluntary control over beliefs (also called *doxastic control*) for the self as compared with others, which offers a pertinent test of rival theories. There are plausible theoretical reasons to predict that people would attribute more control to themselves over their beliefs than they attribute to others (over theirs); but there are also plausible reasons to predict the opposite pattern. Resolving this question therefore has important theoretical implications. On a more practical level, discrepant self–other judgments about doxastic control have the potential to exacerbate real-world disagreements over discordant beliefs. For instance, people often feel personally affronted when others do not share their beliefs (e.g., Golman, Loewenstein, Moene, & Zarri, 2016). If they also judge themselves

Corey Cusimano, University Center for Human Values and Program in Cognitive Science, Princeton University; Geoffrey P. Goodwin, Department of Psychology, University of Pennsylvania.

Correspondence concerning this article should be addressed to Corey Cusimano, University Center for Human Values and Program in Cognitive Science, Princeton University, Peretsman Scully Hall, Princeton, NJ 08540. E-mail: cusimano@princeton.edu

[1] These examples were obtained from *The Corpus of News on the Web* (Davies, 2008).

to have less (or more) control over changing their beliefs than others judge them to have, then this may lead to discrepant expectations about which party can choose to change their minds, thus potentially amplifying the original conflict. Both factors point to the relevance of understanding whether, and why, people judge that they and others have different levels of control over what they judge to be true.

As noted above, two divergent predictions emerge from past research. One line of research predicts that people should tend to judge themselves as having more control over their own beliefs than others have over theirs because, to most people, control is desirable, and people often self-enhance desirable properties. People have a strong preference to feel and exert control and react negatively to feeling a loss of it (e.g., Brehm, 1966; Burger & Cooper, 1979; Kelley, 1971; Seligman, 1974, 1975; Wortman & Brehm, 1975). Indeed, the desire for control has been described as one of the strongest human motivations (Gebhardt & Brosschot, 2002; see also Bandura, 1977; Deci & Ryan, 1985; White, 1959). As a consequence, people tend to overattribute control to themselves, inflating how much control they think they have over many things in their life. Indeed, past work suggests that attributions of control are readily biased by motivational concerns (e.g., Alicke, 2000; Burger, 1986; Clark et al., 2014; Mazzocco, Alicke, & Davis, 2004; Miller & Norman, 1975). For instance, the so-called *illusion of control*—whereby people attribute control to themselves over things they in fact have no control over (Langer, 1975; but see Gino, Sharek, & Moore, 2011)—appears especially pronounced in individuals who have a strong desire for control (Burger, 1986).

Because the desire for control pertains to the self and not to others, we would therefore expect people to inflate self-directed, but not other-directed, attributions. Consistent with this reasoning, several studies have found that people attribute to themselves greater control over their own actions than they grant to others (Pronin & Kugler, 2010). In particular, people regard their own behavior as driven more by their own intentions and desires than the same behavior performed by their roommates (Pronin & Kugler, 2010, Study 4; see also Miller & Norman, 1975). Therefore, if people reason about their beliefs in the same way that they reason about their behavior, they should grant themselves more volitional control over their beliefs than they grant to others (over theirs).

Yet there is also reason to postulate precisely the opposite prediction, namely, that people will attribute *less* belief control to themselves than they attribute to others. This is the prediction we made, and our argument proceeds as follows: Unlike actions, beliefs are typically experienced as uncontrollable on account of internal, psychological constraints on belief change. However, because people tend to have difficulty reasoning about, and fully accounting for, the hidden, psychological constraints operating on others, they should routinely fail to account for these constraints when attributing control to others. Instead, we argue that they rely on a default, unreflective judgment that beliefs, like behaviors, are generally controllable, which results in their attributing to others a high degree of control over their beliefs. The combination of these factors should yield a self–other discrepancy, such that believers attribute to themselves less control over their beliefs than they attribute to others. We motivate this line of reasoning below with two key premises.

Our first premise is that beliefs—more so than actions—are subject to psychological constraints that limit people's ability to change them

voluntarily (James, 1937). One major source of constraint is the perceived evidence that impinges upon people's beliefs. Consider this passage in William James's essay, *The Will to Believe*, in which he reflects on this evidentiary constraint on belief:

> Can we, by any effort of our will, or by any strength of wish that it were true, believe ourselves well and about when we are roaring with rheumatism in bed, or feel certain that the sum of the two one-dollar bills in our pocket must be a hundred dollars? We can say any of these things, but we are absolutely impotent to believe them. (p. 5)

In this passage, James asserts that when people possess strong evidence in favor of a belief (e.g., that they have two dollars in their pocket), they cannot simply choose to form a contrary belief.

Existing work that directly investigates belief formation and change appears to corroborate this picture—suggesting that beliefs are partially outside people's voluntary control, precisely because of these evidentiary constraints. In essence, although people can indirectly influence the quality of their beliefs, including how rational and justified those beliefs are (e.g., by exposing themselves to new information, or by deliberating in specific ways; Baron, 2008; Haran, Ritov, & Mellers, 2013; Stanovich & West, 1997; Webster & Kruglanski, 1994), they cannot simply adopt whatever belief they want to (Epley & Gilovich, 2016; Sloman, Fernbach, & Hagmayer, 2010). For instance, when presented with strong arguments in favor of a proposition, people tend to change their beliefs, even when they would prefer not to (Petty & Cacioppo, 1986; Wood & Porter, 2019).[2]

Furthermore, anecdotal evidence suggests that people often experience their beliefs as constrained by evidence, which leads them to view themselves as having low control over their beliefs (e.g., Alston, 1988; Epley & Gilovich, 2016; Pascal, 1852; see Turri, Rose, & Buckwalter, 2018 for a review). For instance, Alston (1988) considers whether he could choose to believe that the United States is still a colony of Britain just by deciding that it is so. He concludes that he "has no such power" (p. 263) and could not change his mind even if offered a huge sum of money. Thus, it appears that confrontation with evidence limits the capacity for voluntary belief change and does so in a way that gives rise to an experience of belief constraint that is discoverable through introspection.

People may also apprehend evidentiary constraints on their beliefs even when there is little objective evidence supporting those beliefs. A core postulate of the well-known theory of naïve realism is that people assume that they "see entities and events as they are in objective reality" and that their "social attitudes, beliefs, preferences, priorities, and the like follow from a relatively dispassionate, unbiased, and essentially 'unmediated' apprehension of the information or evidence at hand" (Ross & Ward, 1996,

---

[2] This conclusion is not undermined by the phenomenon of motivated reasoning, as it might seem to be at first. Although there is widespread agreement that people sometimes reason in motivated (i.e., biased) ways (Baumeister & Newman, 1994; Kunda, 1990), the existence of this phenomenon does not imply that people have conscious, volitional control over their beliefs. In fact, motivated reasoning is likely to work best when it bypasses the will, with the relevant motivations affecting the kinds of information that people consider, rather than operating directly via the will to control final belief states (Baumeister & Newman, 1994; Epley & Gilovich, 2016). We discuss the relevance of motivated reasoning, as well as the claim that beliefs in fact are uncontrollable, in the General Discussion.

p. 110; see also Griffin & Ross, 1991). This tendency to regard one's beliefs as well (i.e., objectively) founded—separate from the actual state of the evidence with regard to any given belief—could therefore exacerbate the felt sense of belief constraint.

Finally, the sense of constraint may also arise from nonevidentiary factors. Many beliefs people hold are the product of unconscious or arational processes (Nisbett & Wilson, 1977; Zajonc, 1980). For instance, repeated exposure to a stimulus can cause people to believe that it is preferable, safe, or of high quality (Zajonc, 1980). Without having access to how such beliefs came about in the first place, people would likely have little ability to change them (Wilson & Brekke, 1994); they would also likely *experience* little ability to exert such control. In sum, a mixture of both evidentiary and nonevidentiary factors could jointly contribute to the experience people have that their beliefs are constrained. To date, however, no work has examined whether lay people actually experience their own beliefs as outside of their control, which was a major purpose of our investigations.

Our second premise is that people will often fail to appreciate that others suffer this same sense of constraint over their beliefs. This idea derives from a broader difficulty people have in appreciating others' inner experiences (Pronin, 2009), stemming from the fact that people do not directly experience others' mental states but have to infer them indirectly (Jones & Nisbett, 1972). Consequently, people judge others to have less complex mental experiences than they themselves do, fail to appreciate the subjective importance of others' experiences, and judge others' emotions as less intense than their own (Johnson, 1987; McFarland & Miller, 1990; Miller & McFarland, 1987; Pronin, Fleming, & Steffel, 2008; Pronin, Kruger, Savitsky, & Ross, 2001). This failure to appreciate others' inner experiences leads people to mis-predict and misunderstand others' behavior (e.g., Bierbrauer, 1979; Jones & Harris, 1967; see Gilbert & Malone, 1995, for a review). Furthermore, people often fail to account for psychological constraints operating on *themselves* if they are not directly experiencing them in the moment, leading to similar prediction errors (e.g., Gilbert, Gill, & Wilson, 2002; Loewenstein, 1996; Van Boven & Loewenstein, 2003). Based on this background research, we predicted that people will insufficiently appreciate others' felt experience of low belief control, attributing to them more control than those others attribute to themselves.

To summarize, if people's ability to alter their beliefs is genuinely constrained—as it appears to be—and if it is constrained by forces that are not directly observable in others, such as the evidence perceived in favor of a particular proposition—as it also appears to be—then we should expect observers to be less likely to incorporate these internal constraints when judging others' control over their beliefs. Actors, however, should be more likely to incorporate the constraints they encounter while introspecting on their own beliefs, and should therefore attribute lower control to themselves over their own beliefs. As a consequence, we should expect people routinely to judge themselves as having less control over their own beliefs than others have over theirs.

Overview of Studies

The present set of studies sought to test these predictions. No work that we are aware of has measured whether lay people judge that they have control over their own beliefs; similarly, no work has

compared self and other-directed ratings of belief control. Our studies address these questions, thereby enabling a test of the two competing theories described above. Based on the reasoning outlined above, our prediction was that when considering the ability to control specific, concrete beliefs, people would attribute to themselves less control than they would attribute to others. However, we remained open to the possibility that the alternative prediction (more belief control attributed to self than other) would instead prove correct, and the studies were capable of revealing this.

We conducted seven studies to address these issues. In Study 1, we find that for opposing beliefs on important social issues, people reliably judge themselves to have less ability to change their beliefs than others have. In Study 2, we find that this effect generalizes to a case in which self and other hold the same belief (rather than opposing beliefs). In Study 3, we compare immoral beliefs and actions to examine an alternative account of our findings, namely that subjects think that it would be bad (or look bad) to say that they can change their beliefs, and so they therefore lower ratings of their own control relative to their ratings of others (as might be predicted by some theories of self-enhancement or self-presentation). Thus, Studies 1–3 establish the main finding and rule out salient alternative explanations.

Studies 4 and 5 directly test the process model implicated by the line of reasoning presented above. In Study 4, we test whether the self–other control discrepancy is explained by differences in what information people spontaneously think about while judging their own and others' belief control. Consistent with our theorizing, we find that people are more likely to think about psychological constraints—such as the evidence in support of a given belief—when judging their own belief control than when judging others' belief control. By contrast, people are more likely to think about their generic concept of beliefs—as highly controllable, or not—when judging others. In Study 5, we manipulate subjects' attention to either the supporting evidence that they or another person has over a given belief, or to their own default, generic conception of belief controllability. Consistent with findings from Study 4, we find that subjects who write about their generic conception of belief provided higher attributions of control compared with subjects who wrote about evidence.

In Studies 6 and 7 we examine whether the self–other discrepancy is attenuated when people do not consider evidence for a specific belief, such as when they reason about their own and others' doxastic control in general terms. In Study 6, we find that people attribute to themselves more control when considering their belief control in general than when considering specific beliefs that they hold. And in Study 7, we find that the self–other difference occurs only when people consider specific beliefs; it is fully attenuated when people consider their own and others' control over beliefs in general. These latter findings are directly predicted by the theoretical reasoning outlined previously, which posits that it is introspective access to psychological constraints (chiefly, perceived evidence) that drives down self-directed attributions of control relative to other-directed attributions.

Transparent Reporting

All sample sizes and statistical analyses were preregistered (preregistration information can be found in the online supplemental materials). The only exceptions are the posttests associated with Study 3. All analyses are reported as either planned (i.e., preregistered) or exploratory. Only one study used exclusionary criteria,

which were also preregistered (see Study 3 for details). Otherwise, all sample sizes represent the number of subjects who completed the study, as no subjects or data were removed prior to analysis. Data and annotated R analysis scripts for each study are located at the first author's OSF page: https://osf.io/z8ctb/.

In all studies, we attempted to maximize power by using within-subject manipulations and, in Studies 1 and 2, multiple trials per subject. However, all studies manipulate the target of the attribution (self vs. other) between subjects. In pretests, we observed the predicted effects when manipulating the target of attribution within-subjects, but these designs also tended to produce strong effects of condition order (data available upon request). Accordingly, we ran each reported study with a between-subjects design. Lastly, all studies reported in this article were approved by the Office of Research Ethics at the University of Pennsylvania.

Study 1

Study 1 investigated whether people judge their own ability to change a belief about an important social issue differently from another person's ability to change their opposing belief on the same issue. We examined subjects' beliefs on four topics: (a) whether God exists (God), (b) whether genetically modified foods should be prohibited (GMF), (c) whether government regulation is the best way to address global climate change (Climate), and (d) whether social media has had a negative overall impact on dating (Social Media). We selected these topics because they reflect timely and important social issues over which people frequently disagree (Pew Research Center, 2015a, 2015b, 2016a, 2016b).

Our primary prediction was that judgments of belief control for the self would be lower than corresponding judgments of belief control for the other person. The study also contained an exploratory component. We were interested in whether this predicted discrepancy would apply across two distinct judgments of belief control. The two judgments concerned (a) whether the agent (self or other) *could choose to change the belief if they wanted to* and (b) whether the agent (self or other) *intentionally chose* to hold a particular belief. Judging whether one could choose to change a belief involves considering one's current ability to intentionally bring about belief change.[3] It captures the notion of control in past scholars' introspective accounts of low belief control (see the William James example provided earlier), and it also accords with the notion of control in metacognition research (e.g., Nelson & Narens, 1990). By contrast, judging that one intentionally chose a belief is a retrospective judgment that involves recalling one's history of arriving at that belief, including whether one had a desire to adopt the belief the in the first place (Malle & Knobe, 1997a). We included it because we considered it possible that such retrospective judgments of intentional choice would similarly yield a self–other discrepancy, with ratings of the self's intentionality lower than those of others' intentionality.

Method

Participants. We recruited 394 people (mean age = 37, 184 reported female) from Amazon's Mechanical Turk system to participate in the experiment.

Design. We used a 2 × 2 between-subjects design, investigating attribution target (self vs. other) and control measure (chose vs. change). Subjects were randomly assigned to make intentional choice and voluntary change judgments (within-subjects) about either their own or others' belief control (between-subjects). Each subject did this for four beliefs that they held (God, GMF, Climate, and Social Media).

Procedure. At the beginning of the study, subjects reported their current belief on four topics: God, GMF, Climate, and Social Media. They did so by choosing which of two opposing statements they agreed with on each issue (e.g., God exists vs. God does *not* exist, genetic modification should be prohibited vs. genetic modification should *not* be prohibited; see Appendix A for full text of stimuli). Subjects were then randomly assigned to respond to follow-up questions about either their own control over these beliefs (self), or alternatively, another person's control (other) over beliefs opposite to those held by the self. Subjects indicated whether they (or the other person) deliberately chose to hold each of the four beliefs (chose), as well as whether they (or the other) could choose to hold the opposite belief (change).

Which statements subjects were presented with varied depending on what beliefs they indicated at the beginning of the study. In the self condition, if a subject initially indicated that they believed that genetically modified foods should be prohibited, they would then have rated their agreement with the following chose statement, "I deliberately chose to believe that genetically modified foods should be prohibited," and with the following change statement, "If I wanted to, I could choose to believe that genetically modified foods should *not* be prohibited." However, if the subject originally indicated that they believed that genetically modified foods should not be prohibited, then they would instead have rated their agreement with, "I deliberately chose to believe that genetically modified foods should *not* be prohibited," and "If I wanted to, I could choose to believe that genetically modified foods should be prohibited."

In the other condition, before rating control, participants were given the following instructions (paragraph breaks indicated by "//"):

> We are now going to ask you a series of questions about other people who, in a prior study we conducted, indicated what they believed about each of these four topics. // We are keeping it confidential who they were, just as all data we collect is kept confidential, so try to imagine another Mechanical Turk worker, similar to yourself, who holds the belief we describe. // In each case, you will be reading about a person similar to you but who holds an attitude that you do not hold.

For each of the four beliefs, participants then saw a statement like the following (bold in original text):

> A participant from a previous experiment indicated that he or she "**believes that genetically modified foods should *not* be prohibited.**" You indicated that you believe the opposite. // Please indicate your agreement with the following statements. (Even if you are not certain of the answer, please indicate what you think is most likely).

---

[3] This measure comprised our primary dependent variable in the full set of studies we report, although at the time we ran Studies 1 and 2 we were not sure whether, or how, it would differ from the intentional choice measure.

The specific content was matched to be the opposite of the subject's own belief. The chose and change questions used wordings similar to those used in the self condition (see above), modified as needed for the other condition. For instance, in the case above, subjects rated their agreement with, "This person deliberately chose to believe that genetically modified food should *not* be prohibited," and "If this person wanted to, he/she could choose to believe that genetically modified foods should be prohibited."

All ratings were made on 1–7 rating scales with 1 labeled *completely disagree* and 7 labeled *completely agree*. Subjects in both the self and other conditions responded to control questions for all four beliefs, which were shown on separate screens in an order randomly set for each participant. At the end of the study subjects indicated their sex, age, and were debriefed. No other measures were collected.

Results

As planned, we ran a linear mixed-effect model regressing agreement ratings on attribution target (self vs. other), control measure (chose vs. change), and their interaction. The model also included random by-subject and by-belief content intercepts as well as random by-subject and by-belief content slopes for the effect of control measure.[4] We observed significant effects of target, $b = -0.53$, $SE = 0.12$, $df = 392$, $t = -4.33$, $p < .001$, $R^2_{(m)} = 0.02$, and control measure, $b = 0.81$, $SE = 0.18$, $df = 4.41$, $t = 4.60$, $p < .001$, $R^2_{(m)} = 0.04$, as well as their significant interaction, $b = 1.42$, $SE = 0.17$, $df = 392$, $t = 8.20$, $p < .001$, $R^2_{(m)} = 0.04$ (see Figure 1).

Tests of simple effects revealed that, as predicted, subjects' judgments of their own ability to voluntarily change their beliefs ($M = 3.83$, $SD = 2.12$) were significantly lower than corresponding judgments of others' ability to do so ($M = 5.08$, $SD = 1.73$), $b = -1.24$, $SE = 0.15$, $df = 392$, $t = -8.00$, $p < .001$, $R^2_{(m)} = 0.05$. However, subjects did not report that they chose their own beliefs any more or less than they reported others chose their beliefs, $b = 0.18$, $SE = 0.14$, $df = 392$, $t = 1.23$, $p = .218$, $R^2_{(m)} < 0.01$. When analyzed separately, all four beliefs revealed the same pattern of results (see Table 1 for means and standard deviations; see Figure 2 for voluntary change ratings).
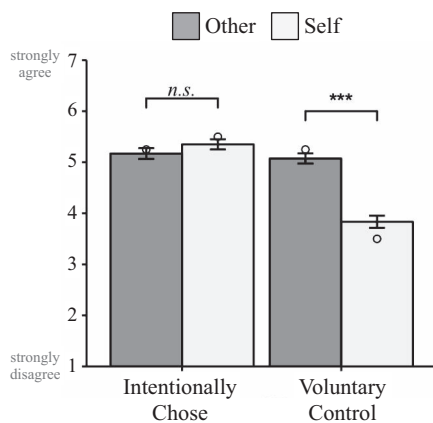

Figure 1. Means and standard errors from participant responses in Study 1. Circles represent median values. *** $p < .001$.

Discussion

Study 1 showed that people judged themselves less capable than others of changing their beliefs. This finding replicated across all four belief statements investigated in this study. This finding points to the possibility that there are *two* sources of disagreement in cases of everyday belief conflict: Individuals with opposing beliefs disagree not only about the matter at hand, but also about who between them could choose to change their mind. We address the implications of this finding in the General Discussion.

In contrast with change ratings, we observed no difference between judgments of whether the self or the other person intentionally chose to have the beliefs in question. This null effect was not directly predicted, nor was the difference between change and choose ratings. However, as we noted in the Introduction to this study, the *change* measure of control accords best with the operationalizations of control in the metacognition literature, as well as with prior anecdotal reports of low belief control (e.g., William Alston's account). One post hoc explanation for the observed difference is that, because the change measure is focused in the present (or the immediate future), it directly confronts people with the limits they face when trying to control a given belief, whereas the intentional choice measure, being retrospective (did I choose this belief in the past?), evokes these limits much more indirectly (as well as being subject to memory loss). Given that the theory predicting lowered self-ratings of control hinges on whether subjects directly experience the limits to their own control, this difference in the focus of the two questions may account for the difference in the pattern of ratings. Consistent with this idea, control ratings were lower overall for the change measure than for the choice measure, as a function of the lower ratings in the self-change condition, specifically. Regardless of the explanation, in Study 2 we examined whether this difference between the two measures replicated.

The self–other discrepancy for change judgments provides initial support for the theory that people view their own beliefs as less controllable than others', while also highlighting a potentially important dynamic between people who hold opposing attitudes. However, the fact that subjects only judged someone who held an opposing belief leaves open the possibility that the discrepancy is limited to cases of disagreement, rather than reflecting a more general self–other difference. In particular, people might judge that the disagreeing other has an incorrect belief, and that incorrect beliefs are more changeable than correct ones, not that other people generally have more control over their beliefs. If the self–other discrepancy is truly general, it would need to replicate in cases where the other person holds the same belief as the self. We therefore tested this in Study 2.

---

[4] In Studies 1 and 2 we computed linear mixed-effect models using the lme4 (Bates, Mächler, Bolker, & Walker, 2015) package in the R computing environment. For effect sizes, we calculated partial-R^2 ($R^2_{[m]}$) for each fixed effect using the r2glmm package (Jaeger, 2017), which implements the approach suggested by Nakagawa and Schielzeth (2013). Finally, p values were generated using Kenward–Roger's method via the pbkrtest package (Halekoh & Højsgaard, 2014).

CUSIMANO AND GOODWIN

Table 1

Means (and Standard Deviations) for Judgments About Each Belief in Studies 1 and 2

Study	Control measure	Attribution target	Climate	GMF	God	Social media
Study 1	Chose	Other	5.11 (1.81)	5.35 (1.72)	5.21 (1.93)	5.03 (1.69)
		Self	5.36 (1.68)	5.41 (1.70)	5.41 (2.01)	5.23 (1.77)
	Change	Other	5.15 (1.64)	5.22 (1.67)	4.93 (1.95)	5.01 (1.63)
		Self	3.76 (2.03)	4.09 (2.02)	3.21 (2.31)	4.29 (1.94)
Study 2	Chose	Other	5.06 (1.86)	4.99 (1.95)	5.38 (1.83)	4.62 (1.97)
		Self	5.37 (1.72)	4.93 (2.03)	5.49 (2.14)	4.90 (1.90)
	Change	Other	4.48 (2.04)	4.66 (2.08)	4.29 (2.26)	4.93 (1.82)
		Self	3.97 (2.21)	4.32 (2.14)	3.57 (2.39)	4.53 (1.97)

Note. Ratings were made on 1–7 scales (1 = *completely disagree*; 7 = *completely agree*). GMF = genetically modified foods.

Study 2

Method

Participants. One hundred ninety-eight people (mean age = 39; 112 reported female) recruited from Amazon's Mechanical Turk platform participated in the experiment.

Design and procedures. We replicated the design of Study 1 in Study 2, crossing attribution target (self vs. other) and control measure (chose vs. change). After reporting what they believe, subjects were randomly assigned to respond either to questions about their own, or another person's beliefs. As in Study 1, subjects answered both control questions for four different beliefs that they held.

Subjects first rated their agreement with four statements that were adapted from the topics used in Study 1 (e.g., "God exists"; see Appendix B for the full text of all items). The four statements were presented in a new random order for each subject, and each rating was made on a 6-point rating scale with the following options, in order: *strongly disagree*, *disagree*, *somewhat disagree*, *somewhat agree*, *agree*, and *strongly agree*. Subjects were next randomly assigned to either the self or the other condition and answered follow-up control questions about each of the four beliefs.
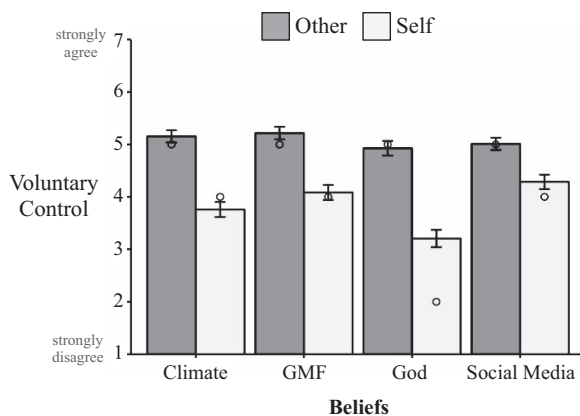

Figure 2. Means (and standard errors) for subjects' agreement ratings in response to the control question, across self and other conditions, for each of the four beliefs in Study 1. Circles represent median values.

In the self condition, subjects were first reminded of what they had just reported believing and were then asked to indicate their agreement with statements that they "deliberately chose" and "could choose to believe / to not believe" the earlier statements they had endorsed. For instance, those who indicated that they "strongly agreed" with the statement, "genetically modified foods should be prohibited" were reminded, "You indicated that you **strongly agree**" with the statement "**Genetically modified foods should be prohibited.**" // Please indicate your agreement with the statements below. (bold original). The two statements pertained to deliberate choice (e.g., "I **deliberately chose to believe** that genetically modified foods should be prohibited") and the ability to choose *not* to believe the statement (e.g., "If I wanted to, **I could choose *not* to believe** that genetically modified foods should be prohibited") in that order. For cases in which subjects had first indicated disagreement, the statement about choice referred to whether the subject "deliberately chose not to believe," and the statement about voluntary change referred to believing the proposition.

In the other condition, subjects were provided instructions indicating that they would answer questions about a person who believed the same thing that they did:

> We are now going to ask you a series of questions about other people who hold similar beliefs to you—specifically people who responded the same way to these questions in earlier studies. // In each case, we are reporting someone's belief that we measured in a prior study we conducted (though of course we are keeping it confidential who they were, just as all data we collect is kept confidential). // For each of the following questions, try to imagine another Mechanical Turk worker, similar to yourself, who holds the belief we describe.

To illustrate, subjects who indicated that they *strongly agree* with the GMF statement were presented with the following prompt, "Another mechanical turk worker, from a prior study we conducted, indicated that they **strongly agree**" with the statement "**Genetically modified foods should be prohibited**." These subjects then indicated their agreement with the chose and change questions: "This person **deliberately chose to believe** that genetically modified foods should be prohibited," and "If this person wanted to, **he/she could choose *not* to believe** that genetically modified foods should be prohibited," respectively. As in the self condition, the prompts were modified to match subjects' initial agreement or disagreement with each statement.

Subjects rated their agreement with the chose and change questions on 7-point scales with 1 indicating *completely disagree* and 7 indicating *completely agree*. The four items (each consisting of a pair of questions) were presented on separate pages and in a random order for each subject. At the end of the study, participants reported their sex and age before being debriefed. No other measures were collected.

Results

As planned, we followed the same analysis procedure from Study 1, regressing agreement ratings on control type, attribution target, and the interaction of control type and attribution target, using a linear mixed-effects model with random intercepts for subject and belief content, and random by-subject and by-belief slopes for the effect of control type. These analyses revealed no significant effect of control type, such that chose ratings were only marginally higher ($M = 5.09$, $SD = 1.51$) than change ratings ($M = 4.34$, $SD = 1.73$), $b = 0.75$, $SE = 0.34$, $df = 3.84$, $t = 2.20$, $p = .095$, $R^2_{(m)} = 0.03$. Similarly, there was no main effect of target (self: $M = 4.80$, $SD = 1.64$; other: $M = 4.64$, $SD = 1.69$), $b = -0.17$, $SE = 0.19$, $df = 196$, $t = -0.90$, $p = .371$, $R^2_{(m)} < 0.01$, but there was a significant interaction between control type and target, just as there had been in Study 1, $b = 0.66$, $SE = 0.27$, $df = 196$, $t = 2.41$, $p = .016$, $R^2_{(m)} = 0.01$.

A test of simple effects confirmed our prediction that ratings of change were lower for self ($M = 4.10$, $SD = 1.72$) than for other ($M = 4.59$, $SD = 1.72$), $b = -0.49$, $SE = 0.24$, $df = 196$, $t = -2.02$, $p = .044$, $R^2_{(m)} = 0.01$, with no corresponding difference between self ($M = 5.17$, $SD = 1.48$) and other ($M = 5.01$, $SD = 1.54$) for chose ratings, $b = 0.16$, $SE = 0.22$, $df = 196$, $t = 0.75$, $p = .454$, $R^2_{(m)} < 0.01$ (see Figure 3). Examined another way, change ratings were lower than chose ratings in the self condition, $b = 1.08$, $SE = 0.37$, $df = 5.15$, $t = 2.94$, $p = .03$, $R^2_{(m)} = 0.03$, but not the other condition, $b = 0.42$, $SE = 0.37$, $df = 5.15$, $t = 1.15$, $p = .301$, $R^2_{(m)} = 0.01$. As planned, we repeated the analysis

Figure 3. Mean ratings (and standard errors) across conditions in Study 2. Circles represent median values. $^*p < .05$.

from Study 1 investigating differences between self and other for each belief individually. The results were less consistent than those observed in Study 1. Self-directed ratings were lower than other ratings for God, $t(195.42) = -2.20$, $p = .029$. However, for the other items, while the means all trended in the expected direction (other higher than self), voluntary change ratings were not significantly different for Media, $t(194.77) = -1.71$, $p = .089$, GMF, $t(195.81) = -1.11$, $p = .268$, or Climate, $t(194.76) = -1.50$, $p = .136$.

Discussion

Study 2 replicated Study 1's findings that people judge others to have a greater ability to voluntarily change their beliefs than they themselves do. In Study 2, this result occurred even though subjects judged another person's ability to stop believing a mutually shared belief rather than an opposing belief. This finding suggests that there may be a general self–other discrepancy in attributions of control over beliefs rather than the difference being limited only to cases of disagreement. As in Study 1, this difference occurred only for judgments of the voluntary ability to change one's beliefs and did not occur for judgments of intentional choice; thus, we have further evidence that the self–other difference is specific to judgments of voluntary change. For this reason, we focused only on judgments of voluntary change in the subsequent studies and return briefly to this issue in the General Discussion.

The findings thus far are consistent with the idea that people's unique introspective access to the constraints on their own beliefs causes them to rate their own belief control lower than that of others. But there are some alternative explanations for this discrepancy that need to be addressed. One in particular is that people may regard voluntarily changing their beliefs (especially without exposure to new, justifying information) as wrong or counternormative, which in turn affects ratings of their own belief control. This idea is encapsulated by William James, who writes, "the talk of believing by our volition . . . is worse than silly, it is vile" (p. 7, 1937; see also discussion in Clifford, 1877). Corroborating this perspective, recent research has indeed shown that some people regard adhering to the norms of rationality as a moral issue (Ståhl, Zaal, & Skitka, 2016). And, because people tend to regard the majority of their beliefs as reasonable and justified (Pronin, Gilovich, & Ross, 2004; Ross & Ward, 1996), voluntarily changing these beliefs may therefore strike many people as unjustified or irresponsible, and therefore, as morally questionable. Accordingly, this might explain why our subjects were reluctant to grant themselves the capacity to exert voluntary control over their beliefs.

The putative badness of voluntary belief change raises two distinct alternative mechanisms for the findings so far. First, people may privately judge that they are less able to voluntarily change their beliefs on the grounds that, as generally good people, they are less capable than others of immoral behavior. Supporting this idea, prior findings suggest that people generally hold a more favorable moral view of themselves than they hold of others (e.g., Allison, Messick, & Goethals, 1989; Alicke, 1985). Second, subjects' judgments may reflect their desire to present themselves in a good light to the experimenter. On this account, subjects judge it as reputation enhancing to say that they could not perform some unvirtuous behavior—a motivation that would depress reports of their own, but not others' control over beliefs. If either of these
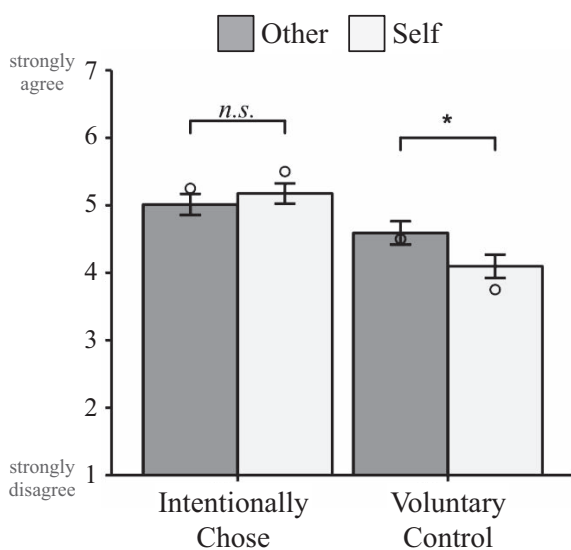
alternatives account for the discrepancy reported above, then the results would reflect existing and well-established biases.

Accordingly, Study 3 tested whether the tendency to regard one's own beliefs as less controllable than others' beliefs is explained by general self-presentational or self-enhancement concerns. We compared people's judgments of their ability to voluntarily change a belief with their judgments of their ability to voluntarily perform a hypothetical immoral behavior. According to the self-enhancement or self-presentation explanation, self-ratings of belief control are depressed relative to others because people regard belief change as bad. If this is the case, then subjects should similarly depress (relative to others) reports that they could choose to perform an (equivalently bad or worse) immoral behavior.

However, our theory does not predict this. Instead, it predicts that the self–other difference for belief control should be significantly attenuated when people judge hypothetical bad behavior. The reason for this is that the constraints on belief change (including, for instance, the supporting evidence for a belief; see Introduction) are salient as soon as a person starts thinking about a given belief. People can access their beliefs very easily—for example, while sitting down at a computer participating in an experiment on belief control—and, as a consequence, they ought to be able to access the constraints on their beliefs very easily, too. Indeed, at a moment's notice, a person can voluntarily attempt to change any given belief, which should expose any latent, previously unconsidered constraints on their ability to do so. By contrast, it is more difficult to think about future hypothetical behavior in a way that evokes the situational pressures or visceral states that would limit or compromise one's control (e.g., Loewenstein, 1996; Gilbert & Malone, 1995). Indeed, people often fail to consider relevant constraints on hypothetical behavior, including situational forces as well as subjective experiences, meaning that their judgments for both themselves and others are often similarly negligent of such constraints (e.g., Pronin et al., 2008; Pronin & Ross, 2006).

Based on this reasoning, we predicted that whereas the self–other difference should occur for belief change, it should be much smaller or even nonexistent for judgments of control over immoral behavior. For behavior, we expected judgments of control for self and other to be quite high and much more similar to one another than in the case of belief change. We therefore predicted a statistical interaction between self–other and belief-behavior.

Study 3

In Study 3, subjects reported how much control either they or a close other would have either to believe that a prototypically immoral act was not immoral, or to perform that very same immoral act. In posttests, performing the immoral act was rated as worse than the holding the immoral belief (see below). The study therefore represents a conservative test, because the alternative explanations under consideration hinge on the idea that people are unwilling to report voluntary control over beliefs because they judge that the exercise of such control is wrong or socially undesirable. On this account, given the results of the posttests, the self–other discrepancy should be no larger for belief control than it is for behavior control. But, if the earlier results instead reflect a belief-specific mechanism (at least in part), then the self–other discrepancy should be larger for belief control than it is for behavior control.

Study 3 departed from Studies 1 and 2 in another important way. Whereas in Studies 1 and 2 subjects rated a distant other's belief control, in Study 3 we prompted subjects to make judgments about someone very close to them. Past research has shown that people more readily project their own mental states to close others rather than distant others, and that they are also more inclined to adopt the perspective of liked versus disliked others (Epley & Waytz, 2009). Thus, Study 3 represents a conservative test in this way as well, because a self–other discrepancy should be less likely to occur for close others.

Method

Participants. To make meaningful self–other comparisons (and also meaningful comparisons between beliefs and behavior), subjects had to report both that they thought the action in question (see below for details) was wrong, and that their nominated close other also believed that it was wrong. Subjects who did not do so were excluded, and we preregistered this exclusion plan.

Data collection occurred in two phases. In the first phase, we recruited 549 people from Amazon's Mechanical Turk. Fifty-five of these subjects (10%) did not qualify for the study because they failed one of the two exclusion criteria above. To reach our recruitment target of 500, we recruited another nine people (also from Amazon's Mechanical Turk). Of these, one subject failed one of the criteria for inclusion, yielding a final sample of 502 subjects (mean age = 36.2, 258 reported female).

Design. The study had a 2 × 2 mixed between-within design, with the target of control attributions (self vs. other) manipulated between-subjects and the type of behavior (belief vs. act) manipulated within-subjects. Subjects were randomly assigned to either the self or other condition, and in each case, responded to questions about their ability to change a moral belief and to perform a corresponding immoral behavior (in random order).

Procedure. At the beginning of the study, subjects read a short description of a prototypical immoral behavior:

> Sometimes people take advantage of another person's costly mistake. Specifically, sometimes a person will see that another person walking ahead of them has dropped $20 on the ground but, instead of returning the money, the person behind will just keep it.

After reading about this behavior, subjects reported whether they agree or disagree with the statement, "I believe that in this situation it is wrong to keep the $20 instead of returning it to its original owner." Next, participants were asked to think of someone close to them, such as a best friend, romantic partner, or spouse, and to type out that person's initials. Once they had done so, they were asked whether they agree that, "The person whose initials I typed above believes that in this situation it is wrong to keep the $20 instead of returning it to its owner." Subjects could select either, "Agree, this person believes that in this situation it is wrong to keep the $20 instead of returning it to its owner," or "Disagree, this person believes that in this situation it is not wrong to keep the $20 instead of returning it to its owner."

If subjects indicated either that they believed that keeping the $20 is not wrong, or that the close other they nominated did not believe that it is wrong, then they were not selected to continue in the study. Instead, they were redirected to the short demographics questionnaire (described below), then debriefed, and paid in full

for participating. We excluded these subjects so that all subjects were making control judgments about a belief they shared with their close other. This ensured that we did not reintroduce the possible confound that we had eliminated in Study 2 (i.e., the possibility that people might think that a belief opposite their own is less justifiable, and therefore more changeable).

Subjects who indicated that both they and their close other believed the target act was wrong were randomly assigned to answer follow-up questions about either themselves or their close other. In the self condition, subjects indicated their agreement with a statement about whether they could choose to perform the immoral act, "In this situation, if I wanted to, I could choose to keep the $20 instead of returning it to its original owner," as well as a statement about their ability to believe otherwise about the moral status of this act, "If I wanted to, I could choose to *believe* that keeping the $20 instead of returning it to its owner is *not* wrong." The order of these statements was counterbalanced, and agreement was assessed using 7-point rating scales (1 = *Completely disagree*; 7 = *Completely agree*). In the other conditions, the statements were altered by dynamically inserting the initials of the person the subject had nominated. For instance, if the second author had nominated the first author, the statements would have read, "In this situation, if CC wanted to, CC could choose to keep the $20 instead of returning it to its original owner," and "If CC wanted to, CC could choose to *believe* that keeping the $20 instead of returning it to its owner is *not* wrong," for the act and belief conditions, respectively. For each statement, subjects were instructed to provide their answer without considering how much they (or the close other) actually would want to do or believe the specified thing. Subjects answered each question on a different screen.

At the end of the study, subjects indicated their age and sex, and then were debriefed. No other measures were collected.

Results

As planned, we conducted a mixed between-within ANOVA on the agreement ratings with behavior type (action vs. belief), the attribution target (self vs. other), and their interaction as the predictor variables. Overall, there was a main effect of behavior type, such that subjects agreed that they and others could choose to keep the $20 ($M = 5.01$, $SD = 2.21$), more so than they could choose to believe that keeping the $20 was not wrong ($M = 3.72$, $SD = 2.21$), $F(1, 500) = 167.41$, $p < .001$, $\eta_G^2 = 0.08$. There was no main effect of target, as subjects did not judge themselves ($M = 4.28$, $SD = 2.35$) to have more control overall than others ($M = 4.45$, $SD = 2.25$), $F(1, 500) = 1.03$, $p = .311$, $\eta_G^2 < 0.001$. However, as predicted, we observed a significant interaction such that there was a discrepancy between self and other judgments for beliefs but not actions, $F(1, 500) = 5.05$, $p = .025$, $\eta_G^2 = 0.003$; see Figure 4, below. Because we observed non-normality in the action conditions, we conducted a nonpreregistered Kruskal–Wallis rank sum test, analyzing whether the action-belief difference scores differed between self and other. This test also revealed a significant effect, such that the action-belief difference was wider for the self than for others, thereby corroborating the interaction nonparametrically, $\chi^2(df = 1) = 6.96$, $p = .008$. Follow-up independent-samples t tests revealed that subjects rated themselves less able to choose to believe that keeping the $20 is wrong ($M =$
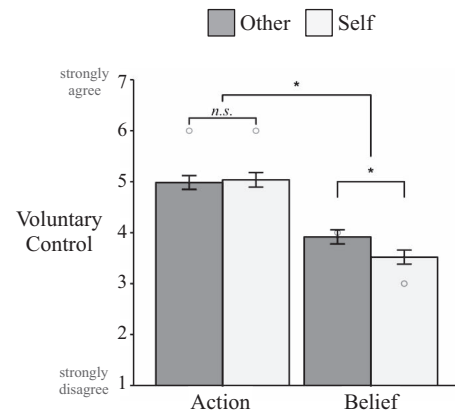

Figure 4. Means (and standard errors) of agreement ratings across conditions in Study 3. Circles represent median responses for each condition. * $p < .05$.

3.52, $SD = 2.20$) than their close other ($M = 3.92$, $SD = 2.21$), $t(499.99) = -2.02$, $p = .044$, $d = -0.18$. But there was no difference between self ($M = 5.04$, $SD = 2.25$) and other ($M = 4.98$, $SD = 2.18$) when comparing actions, $t(499.26) = 0.26$, $p = .793$, $d = 0.02$.[5]

Discussion

Subjects reported that they were less able than a close other to change a belief that an immoral behavior was wrong, yet they reported no difference in their respective capacities to carry out that very same immoral behavior (ratings of action capacity were generally high for both self and other). These findings cast doubt on the idea that general self-presentational or self-enhancement concerns underlie the self–other asymmetry observed in Studies 1 and 2. Rather, they suggest that the true causal mechanism at play in our experiments is more specific to judgments about beliefs (or belief-like mental states).

A strength of Study 3 is that the contents of the belief and the corresponding behavior were closely matched. However, separate

---

[5] One possibility is that the smaller difference in the action condition relative to the belief condition is attributable to a ceiling effect in the action condition. Consistent with this possibility, even though the average ratings for actions were approximately 5 on the scale, the majority of values (55.3%) were a 6 or 7 on the 1–7 scale (see Figure 4). However, although we cannot definitively rule this out, there are two reasons to suspect that it is not the most likely explanation. First, compared with our prior research (Cusimano & Goodwin, 2019), the average control rating of 5 observed in the present study is unusually low for an intentional behavior. In other studies with the same dependent variable, we have found that people's responses tightly cluster around a mean of 6 on a 7-point scale when judging others' control over their intentional behaviors. Thus, people could have used higher regions of the scale, and often do.

Another method for diagnosing a possible ceiling (or floor) effect is to examine the cumulative response distributions to determine whether there is divergence between conditions for values further from the ceiling/floor (see e.g., Simonsohn, Simmons, & Nelson, 2014). When we investigated the cumulative response distributions in the action condition, we observed a uniform absence of differences along every region of the scale. This suggests that the observed interaction is not the product of a ceiling effect in the action condition.

from content, the respective valences of these two stimuli also need to be considered—that is, how bad people thought it would be either to change the moral belief in question, or to carry out the described behavior. If subjects thought that changing their beliefs was worse than carrying out the corresponding behavior, then this might explain their reluctance to indicate a capacity for such belief change, and an alternative explanation based in self-presentation would thereby gain credibility. To check this possibility, we carried out three posttests that compared people's moral judgments of the act described in Study 3 (picking up money that someone dropped and keeping it rather than returning it) with their moral judgments of holding (or professing) the corresponding belief (i.e., that such an act is not wrong).

The three posttests all called for subjects to make moral character judgments. In each case, subjects were initially presented with a description of the target behavior and indicated whether they agreed or disagreed that it was wrong. Only those subjects who agreed that the behavior was wrong continued with the remainder of the study. Posttest 1 compared performing an immoral act with judging (privately) that the act is not immoral. Posttest 2 compared performing the immoral act with professing (publicly) that the act is not immoral. Posttest 3 compared professing that one could not perform the immoral act with professing that one could not choose to believe that the act is immoral. Below is a brief description of these studies; see Appendix C for full text of the vignettes.

Posttest 1. In our first test, we compared subjects' impressions of someone who acts immorally with their impressions of someone who believes that the very same immoral act is not wrong. Fifty-six subjects (of 61) reported that the target behavior was morally wrong. These subjects then read a story involving three characters, Jones, Smith, and Peters. In the story, Jones accidentally and unknowingly dropped a $20 bill on the sidewalk. Smith, who was walking a little way behind him and saw this happen, picked up the $20 and pocketed it. Peters, who witnessed this scene from across the street, privately judged Smith's action as not wrong. Subjects reported their impression of each person. They judged Smith, who kept the $20 ($M = -1.53$, $SD = 1.60$) more negatively than Peters, who merely thought that this act was not wrong ($M = -0.15$, $SD = 1.84$), $t(55) = -5.60$, $p < .001$.

Posttest 2. Our second test compared subjects' impressions of someone who publicly states that a wrong act is permissible with their impressions of someone who performs that same act. One hundred sixty-four subjects (of 175) initially reported that the target behavior was morally wrong. These subjects then read a vignette in which they imagined themselves in conversation with the same third party (Peters) as in posttest 1. After witnessing Smith take the money, the subject turns to Peters and describes what they witnessed. Peters responds that he does not think the act was wrong (see Appendix C for full text). Subjects then made judgments of both Smith, who took the $20, and Peters, who proclaimed that this act was not wrong. Subjects reported having a worse impression of Smith ($M = -1.80$, $SD = 1.35$) than they did of Peters ($M = -1.52$, $SD = 1.29$), $t(163) = -2.99$, $p = .003$, 95% CI [-0.48, -0.1], $d = -0.22$. Thus, in both of posttests 1 and 2, performing an immoral act was judged more negatively than believing (or proclaiming) that the same immoral act is not wrong.

Posttest 3. Our third test was more complex. It compared subjects' impressions of someone who publicly states that he could

not perform an immoral act with their impressions of someone who publicly states that he could not choose to believe that the same act is not wrong. This comparison was designed to alleviate a concern that subjects may report a lack of control over immoral beliefs because there is a distinctive reputational advantage from doing so, as compared with stating that one could not perform an immoral behavior. The claim that one could perform an immoral action may seem universally and uncontroversially true, particularly if it is interpreted in physical and not psychological terms (since it clearly is physically possible to perform the act). Claiming otherwise may therefore seem false and unwarranted, thereby conferring little reputational advantage. It may therefore seem less diagnostic of a person's character than the claim that one could choose to hold an immoral belief—which seems less universally true, more tied to idiosyncratic psychological factors, and therefore more diagnostic of a person's underlying character.[6]

To perform a complete test of this alternative explanation, we also asked subjects to judge two additional individuals—one who states that he could perform the act (though would not), and one who states that he could choose to believe that the act is not wrong (even though he does not believe this currently). Judgments of these individuals constituted a baseline against which to compare the original *could not* judgments, allowing a test of the relative boosts gained by denying the ability to act or believe immorally, respectively.

One hundred nine subjects (of 118) initially reported that the target behavior was morally wrong and continued to the remainder of the study. They read the same basic vignette as in the previous studies, describing Jones accidentally dropping $20 and Smith picking it up and keeping it for himself (there was no Peters in this vignette). Subjects were asked to imagine that they were a witness to this event, seated at a nearby table with several other people. The group then engages in a conversation about the event. The first critical comparison showed that a person who states that, if he were in a similar situation, he could not choose to keep the money ($M = 2.14$, $SD = 1.61$), was judged more positively compared with a person who states that he could not choose to believe that taking the money is not wrong ($M = 0.96$, $SD = 1.65$), $t(107) = 6.11$, $p < .001$. We then compared both of these means against the two baseline conditions, in which the target individual stated that he could perform the act ($M = -0.02$, $SD = 1.43$), or could choose to hold the immoral belief ($M = 0.00$, $SD = 1.49$). These comparisons showed that the relative boost that an individual gains from asserting the lack of ability to hold an immoral belief ($M = -0.96$, $SD = 2.24$) is no greater, and is in fact significantly smaller than the boost gained from asserting the lack of ability to perform an immoral action ($M = -2.16$, $SD = 2.11$), $t(107) = 5.16$, $p < .001$.

In sum, these three posttests indicate the following: When an immoral act and an immoral belief are matched in content, (a) merely holding the belief is seen as less wrong than performing the act, (b) professing the belief is seen as less wrong than performing the act, and (c) professing an inability to hold the belief is judged less favorably than professing the inability to perform the act, and provides a smaller boost in impressions (when compared against

the denial of these abilities). Instead, there appears to be a stronger reputational advantage gained by claiming that one could not perform the relevant immoral behavior. On balance then, these posttests indicate that there is no special reputational advantage that accrues from denying the ability to choose to hold immoral beliefs. Accordingly, the main finding from Study 3—less control attributed to the self than to others over holding immoral beliefs, but no difference for immoral actions—seems unlikely to have resulted from self-enhancement or impression management motivations.

We argue instead that the results are explained by people's greater sensitivity to the internal constraints on beliefs when judging their own belief control as opposed to others'. When people consider their own ability to change a particular belief, they introspect on that belief and, in doing so, confront a sense of low control. This sense of low control drives down people's judgments of their own control over specific beliefs. However, this appreciation of low control is not generalized to others. What then, does account for how people judge others' control over their beliefs?

So far, we have said little about this process, and so we elaborate upon it here. We first postulate that, when people judge how much control they have over a given belief that they hold, they tend to consider factors that are specific to that belief and its surrounding psychological context, such as their reasons for holding the belief, how confident they are in it, and how it fits in with their other beliefs and attitudes. Some of these considerations, such as the supporting evidence for their belief, invoke a sense of constraint, resulting in attributions of low control (see Introduction). Because these factors are often only accessible through introspection, people are less likely to consider them when judging others' belief control. Instead, when judging others, we hypothesize that people are more likely to rely on a generic conception of belief according to which beliefs are highly controllable. Just as people hold generic conceptions of intelligence (Dweck & Leggett, 1988) and emotion (e.g., Ford & Gross, 2019; Tamir, John, Srivastava, & Gross, 2007), they also hold a generic conception of belief (e.g., Akhtar & Wheeler, 2016). One aspect of this generic conception is that beliefs are highly controllable (Cusimano & Goodwin, 2019).[7] However, we argue that this conception of belief as highly controllable is discrepant with the sense of low control induced when people consider changing their own concrete beliefs.

We therefore hypothesize two distinct mechanisms that jointly produce the self–other discrepancy. These mechanisms can be represented diagrammatically as shown in Figure 5. This model generates several predictions beyond the basic self–other discrepancy including (a) what people spontaneously think about while making control judgments about the self or others, (b) what spontaneous thoughts are associated with judgments of low or high belief control, (c) when people will make discrepant belief control judgments for the self versus others, (d) when the self–other discrepancy in control attribution will attenuate, and (e) when people's judgments for the self will be relatively high or low. We describe and test these predictions in Studies 4–7.

Study 4

In Study 4, subjects were asked to attribute control either to themselves or to another person who held an identical belief, and
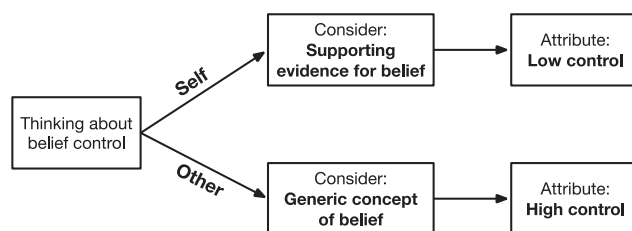

Figure 5. Two routes to attributing belief control when thinking about a belief shared with another person. When people consider their own voluntary belief control, they access the supporting evidence for their belief, which elicits a sense of low control. By contrast, when people consider others' belief control, they access a preexisting, generic conception of beliefs as controllable.

then to write down what they thought about while making their control judgment. Two independent and blind coders then coded these written statements for the presence of either (a) generic reasoning about beliefs or (b) supporting evidence for the belief (described in more detail below). This allowed us to test four predictions that derive from our model of belief control attribution described above. First, we predicted that people often attribute belief control by relying on a generic concept of beliefs, and that reliance on this generic concept is typically associated high belief control ratings. Second, we predicted that people will also sometimes consider specific information about the belief under consideration, such as what evidence or reasons the believer has for it. Thinking about evidence in favor of the belief should be associated with lower attributions of belief control (than reliance on the generic concept). Third, we predicted that people will tend to think about and weigh these two considerations differently when judging their own and others' belief control. Specifically, people will tend to think about beliefs more generically when attributing control to others but will think about the reasons underlying their beliefs when attributing control to themselves. Fourth, differences in what subjects focus on will in turn predict lowered judgments of belief control for the self as compared with others.

Method

Participants. Three hundred ninety-nine subjects (mean age = 37; 195 reported female, 202 reported male) recruited from Amazon's Mechanical Turk platform participated in the experiment.

Design and procedures. At the beginning of the study, subjects reported their current belief about anthropogenic climate

---

[7] The following supplemental study lends additional support to this claim. Two hundred sixty-seven University of Pennsylvania undergraduates a college were asked to rate their agreement (1: *strongly disagree*, 7: *strongly agree*) with four statements about how generally controllable beliefs are. Subjects indicated strong agreement with all four statements, including (a) "People can decide to believe something even when they have good reasons to believe the opposite." ($M = 5.32$, $SD = 1.22$); (b) "If someone really wants to believe that something is true, they can choose to believe it." ($M = 5.35$, $SD = 1.28$); (c) "People can make themselves believe whatever they want to." ($M = 5.07$, $SD = 1.49$); and (d) "No matter what, people can voluntarily choose to believe something if it benefits them to do so." ($M = 5.22$, $SD = 1.42$). These ratings were all significantly above the mid-points of the scales ($ts > 11.2$, $ps < .001$).

change by indicating which of two contradictory statements best reflects their stance on the issue ("I believe that Earth's climate is warming as a result of human activity" or "I believe that Earth's climate is not warming as a result of human activity").

Subjects then reported how much control either they ("self" condition) or another person with an identical belief ("other" condition) has over the belief the subject just reported. Subjects did this by reporting their agreement or disagreement with a statement affirming voluntary control over the belief, similar to the one used in Study 3. The exact text depended on what the subject had just reported. For instance, in the self condition, if the subject had reported that they believed that, "Earth's climate is warming as a result of human behavior," they would then have seen the following prompt:

> You reported that you believe that Earth's climate is warming as a result of human behavior. Now please rate to what extent you agree with the statement: I voluntarily hold this belief. If I wanted to, I could choose to believe that Earth's climate is *not* warming as a result of human behavior.

In the other condition, subjects were told about another Mechanical Turk worker (male, also living in the United States) who reported holding the same belief that they did. Subjects then reported their agreement with the same control statement used in the self condition (minimally modified to refer to this other individual; e.g., "This other person voluntarily holds this belief. If he wanted to, he could choose to believe that Earth's climate is *not* warming as a result of human behavior."). Subjects indicated their agreement or disagreement on a 7-point rating scale anchored at 1 (*completely disagree*) and 7 (*completely agree*).

On the next screen, subjects then wrote down what they thought about while deciding whether or not they or the other person has voluntary control over the belief. In both the self and other conditions, subjects were reminded of their earlier answer to the control question and were prompted to write down what they thought about while giving that answer. For instance, in the self condition, subjects saw the following prompt:

> You just made a judgment about how much voluntary control you have over one of your beliefs. Specifically, on a 1 (*strongly disagree*) to 7 (*strongly agree*) scale, you responded with "[3]" to the statement that you could choose to believe that [Earth's climate is *not* warming as a result of human behavior] if you wanted. We are now interested in what information you considered and weighed when you made that judgment. In the box below, please write down whatever came to mind while you were trying to decide whether you have voluntary control over this belief.

The exact text of the prompts varied depending on which climate change belief subjects reported holding (e.g., belief or denial), and on their prior responses to the control question (e.g., "[3]" in the example above). In the other condition, the prompt was modified to refer to the thoughts that subjects had while judging the other person's belief control.

Finally, subjects completed a demographics questionnaire and were then debriefed about the study.

Coding. Two independent coders, blind to our hypotheses, coded subjects' written descriptions. Specifically, coders noted whether each response contained either (a) a generic statement about belief control or (b) any supporting evidence or arguments

for the specific belief in question (responses could contain both). To complete this coding task, coders were provided with each subject's open-ended response, and the belief in question, but were not provided information about the experimental condition (self vs. other) or the subjects' response on the control question.[8] Coders were instructed to follow the following guidelines:

> Generic Conception of Belief Control: If the subject writes something which states or presupposes a general statement about whether (a) beliefs (in general) are controllable, (b) beliefs (in general) are uncontrollable, (c) beliefs (in general) are chosen, *or* (d) beliefs (in general) are not chosen, indicate "1." Otherwise indicate "0." Subjects need not explicitly state that they are making a general statement about beliefs, but this should be implicit in their response.

> Reasons/Evidence in Support of the Belief: If the subject writes something which articulates (or alludes to the existence of) reasons, evidence, or arguments in support of the specific belief in question, indicate "1." Otherwise indicate "0". It should be implicit in the subject's response that the believer is presently consciously aware of the reasons and evidence being described, but this need not be mentioned explicitly.

These general guidelines were supplemented with additional instructions for how to code boundary cases, which we uncovered while reviewing open-ended responses for a pilot version of this study. For example, some subjects simply reiterated that they did or did not have control over the belief in question (e.g., "I have/do not have control over this belief"). This sort of statement was not counted as citing a generic notion of belief control because it simply states the subject's final conclusion for the belief in question without providing any insight into their underlying reasoning. We therefore regarded it as too narrow in scope to count as a generic statement of belief control. Similarly, we stipulated that to count as "supporting evidence," a response would need to cite evidence or arguments in *support of* the currently held belief, rather than in opposition to it. Our theory explicitly states that only evidence in support of a given belief constitutes a psychological constraint which drives down perceived control. Indeed, evidence or arguments that oppose a belief should facilitate rather than constrain belief control. Accordingly, this restriction was necessary to properly test the theory. Complete coding instructions, which include clarificatory guidance for boundary cases, are available in the online supplemental materials.

Example responses are shown in Table 2. Intercoder reliability was high (84% for generic reasoning about belief control, and 85% for reasons in support of the belief). Disagreements were resolved through discussion until coders agreed on 100% of cases. Subjects' responses and final coding decisions are available on the OSF page for this project, https://osf.io/z8ctb/.

Results

As in prior studies, subjects attributed greater voluntary control over the belief to the other person ($M = 5.40$, $SD = 1.89$) than they did to themselves ($M = 4.54$; $SD = 2.33$), $t(373.37) = 4.06$, $p <$

---

[8] We could not, however, maintain strict blindness to experimental condition or reported control. For some trials there were cues in the subjects' response that revealed the condition (e.g., first person pronouns) as well as the degree of control they had reported (see, e.g., Table 2).

Table 2

Example Statements by Subjects When Prompted to Report What They Considered When Attributing Control Either to Themselves or to Another Person

Consideration type	Example
Generic statements about belief control	"I have voluntary control over my beliefs."
	"You can choose to believe whatever you want, really."
	"Because people are allowed to believe what they wish and they have control over what they believe."
	"Anyone can choose to believe anything."
	"He could say it, but I don't think he could just change his beliefs just like that."
	"I think people can think and believe whatever they want to believe even if it is not true."
	"I think that anyone can choose to believe in what they want to. Even with a large amount of scientific evidence against them, there are often people who choose to believe something different."
Supporting evidence or arguments	"There are too many facts that point to humans causing global warming that it would be difficult for me to believe the opposite viewpoint. So, I am not sure I have voluntary control because it would be hard for me to believe the opposite viewpoint."
	"I thought about factual evidence and research. I can't deny facts. I can't just pretend these facts don't exist."
	"The population of the world and change in the climate increased significantly in the last century. Therefore, there is a positive correlation between increased human activity and Earth climate."
	"I can't voluntary hold a belief to something that has been scientifically proven. To compare, it's like me saying I'm not 140 lbs., even though every measure tool I use shows that I'm 140 lbs."
	"In this particular situation, if the other person chose to believe that Earth's climate is not warming due to human activity, it would be very difficult, if not impossible, to truly believe that, given all scientific evidence proving the opposite."

Note. Responses have been lightly copyedited (e.g., capitalization and punctuation).

.001, 95% CI [0.45, 1.28], $d = .41$ (Figure 6A). We therefore next turned to what subjects considered while making these control judgments.

The majority of subjects ($n = 251$; 63%) reported thinking about either the generic controllability of beliefs ($n = 157$; 39%) or the evidence in favor of the belief in question ($n = 119$; 30%). However, rates of considering these factors differed across condition (Figure 6B). Subjects were more likely to think about whether beliefs are in general controllable when considering another person's belief control ($n = 105$; 51%) than they were when thinking about their own belief control ($n = 52$; 27%), $\chi^2(1, N = 399) = 24.67$, $p < .001$. By contrast, subjects were more likely to think about what evidence or reasons backup the belief when considering their own belief control ($n = 83$; 43%) than when considering others' ($n = 36$; 18%), $\chi^2(1, N = 399) = 28.4$, $p < .001$.

We hypothesized that these considerations would be associated with different attributions of belief control. As expected, subjects who reported thinking generically about belief control attributed high control on average ($M = 5.68$; $SD = 1.84$), whereas subjects who considered reasons for holding the belief (their own or others') attributed significantly lower control ($M = 3.89$; $SD = 2.39$), $t(214.99) = 6.80$, $p < .001$, 95% CI [1.27, 2.31], $d = 0.86$ (Figure 6C). Control ratings associated with generic considerations were significantly above the midpoint of the agreement scale, $t(156) = 11.47$, $p < .001$, 95% CI [5.39, 5.97], whereas the control ratings associated with specific reasons were not significantly different from the midpoint, $t(118) = -0.50$, $p = .618$, 95% CI [3.46, 4.32].

If making judgments about the self or another person causes people to consider different information, which in turn causes different attributions of control for self and other, then differences in subjects' considerations across condition should account for a significant portion of the condition difference in their control judgments. To test this, we conducted a multiple-mediation analysis with the presence of generic belief statements, and the presence of supporting evidence statements, as independent mediators

(Figure 6D). This analysis revealed evidence consistent with both the presence of generic belief control statements ($a_1 = -0.25$, $p < .001$, $b_1 = 0.78$, $p < .001$, $ab_1 = -0.19$, 95% CI [-0.32, -0.07][9]) and the presence of supporting evidence statements ($a_2 = 0.25$, $p < .001$, $b_2 = -1.24$, $p < .001$, $ab_2 = -0.31$, 95% CI [-0.47, -0.15]) independently mediating the self–other control difference, but in opposite directions. Generic statements about belief control were more prevalent in the other condition and predicted higher control judgments, whereas supporting evidence statements were more prevalent in the self condition and predicted lower control judgments. Together, the indirect effects of thinking about generic belief control and supporting evidence accounted for 58% of the total effect on control attributions ($c = -0.86$, $p < .001$; $c' = -0.36$, $p = .087$).

Discussion

Study 4 replicated prior results showing that people attribute greater voluntary belief control to others than to themselves. Just as in Studies 2 and 3, we observed this even when subjects attributed control to someone who held the exact same belief that they did. However, Study 4 went beyond these prior studies by asking subjects to report what they thought about when attributing control to self and other. Consistent with our theorizing, when judging others, people were more likely to think about belief control in generic terms, and less likely to consider the specific support that the person had for the belief. These two different

---

[9] We follow recommendations from Yzerbyt, Muller, Batailler, and Judd (2018) for reporting mediation analyses. We first present results from the joint significance tests of (a) the estimated effect of the independent variable on the mediator, and (b) the estimated effect of the mediator on the outcome, and conclude that there is evidence consistent with mediation when both a and b are significant. We then report the bootstrapped estimated size of the indirect effect (ab) and its 95% confidence interval. Analyses were conducted using the *lavaan* package (Rosseel, 2012).
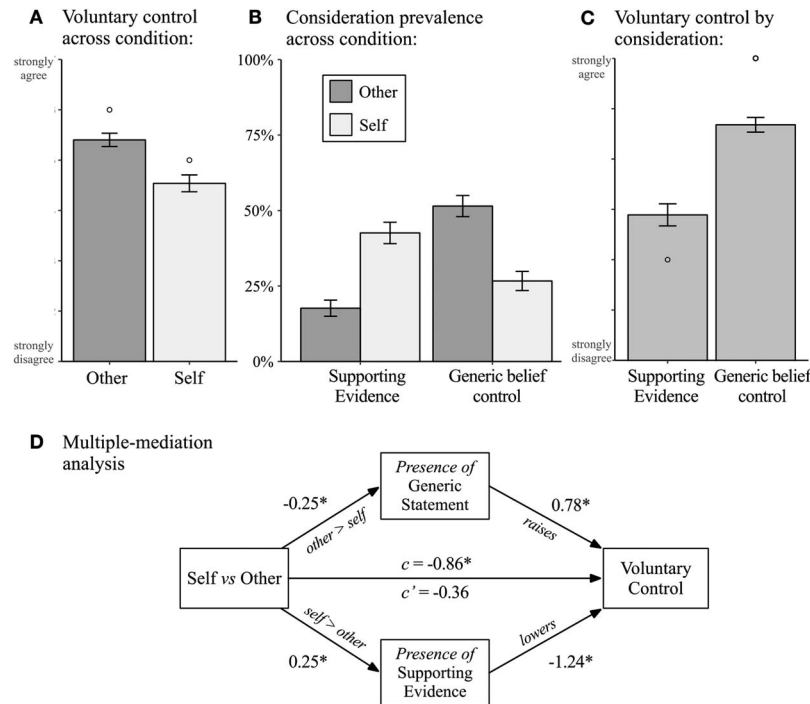
Figure 6. Main results from Study 4. (A) Mean voluntary control judgments across conditions. Error bars represent standard errors. Circles represent median values. (B) The proportion of trials in which subjects reported thinking about belief control in generic terms, or about evidence for the belief, across condition. Error bars represent standard errors. (C) Control judgments split by whether the subject reported thinking generically about belief control or about specific reasons for the belief. Circles represent median values. (D) Multiple mediation analysis testing the mediating role of generic statements and supporting evidence on attributions of voluntary control. $^{*} p < .05$.

considerations were, in turn, associated with different reports of belief control. Specifically, when subjects thought about belief control in generic terms, they tended to attribute a high degree of control. By contrast, when subjects thought about what evidence they or the other person had for the belief, they tended to attribute lower control. This finding is consistent with anecdotal reports of low belief control (see Introduction), in which people view evidence in favor of their belief as constraining voluntary belief change. Finally, mediation analysis indicated that these differences in attributions may underlie the self–other difference in control judgments.

However, although we observed causal evidence that the target of appraisal—self versus other—affects what information people spontaneously consider, this study can only provide correlational evidence that these considerations are associated with subsequent attributions of belief control. Study 5 addresses this directly by manipulating what information subjects consider prior to attributing belief control—either their generic concept of belief control, or the specific evidence supporting a given belief. If our theory is correct, then people should tend to attribute greater belief control when thinking generically than when considering specific evidence for a given belief.

Study 5

In Study 5, we crossed the self–other manipulation with another manipulation that was designed to influence what subjects thought

about while considering belief control. This manipulation comprised three levels: one group of subjects wrote either about belief control in general terms, thereby drawing upon their implicit, generic conception of belief control, another group wrote about the evidence and arguments supporting a given belief, while a third group wrote about belief control without any accompanying instructions.

Our primary prediction was that, relative to the no task condition, subjects' ratings of belief control would be lower when directed to attend to the supporting evidence for a given belief, but higher when directed to attend to their generic conception of belief control. This would provide causal evidence for the two mechanisms that feature in our account, namely that considering evidence reduces people's sense of belief control, while considering belief control generically increases it.

The findings from Study 4 also suggest a more specific pattern of results that we also tested for. Study 4 showed that, when judging their own belief control, subjects by default tend to consider supporting evidence. Accordingly, we predicted that, when judging their own beliefs, subjects' control ratings in the no task condition would match the ratings in the supporting evidence condition (since evidence should be considered in both conditions), but be lower than the ratings in the generic condition. Study 4 also showed that, when judging others, subjects tend to consider belief control in more general terms—thus, we predicted that, when judging others' beliefs, control ratings in the no task condi-

tion would match the ratings in the generic condition but be higher than the ratings in the supporting evidence condition. Put another way, whereas supporting evidence and generic belief concepts are not equally salient for self and other when people are unprompted (see Study 4), the writing prompts ought to make these considerations equally salient for self and other, thereby attenuating differences in control (and resulting in a significant interaction between target and writing task).

Method

Participants. One thousand, one hundred eighty-five subjects (mean age = 38; 691 reported female) recruited from Amazon's Mechanical Turk platform participated in the experiment.

Design. The study comprised a 2 (Attribution Target) × 3 (Preattribution Writing Task) between-subjects design. Subjects were randomly assigned to attribute belief control either to themselves ("self" condition) or to an anonymous Mechanical Turk worker who shared the same belief ("other" condition). Independent of this manipulation, we also randomly assigned subjects to complete one of three writing tasks prior to making a control judgment. In the "generic" writing task condition, subjects wrote about whether beliefs generally are the kinds of things that are freely controllable or not. In the "supporting evidence" writing task condition, subjects wrote about what evidence or arguments they have (or the other person has) in support of the specific belief in question. And finally, in the "no task" condition, subjects attributed voluntary control without any prompt, similar to the procedure described in Studies 3 and 4.

Procedure. At the beginning of the study, subjects reported their belief about anthropogenic climate change by indicating which of two contradictory statements best reflects their belief on the issue ("I believe that Earth's climate is warming as a result of human activity" or "I believe that Earth's climate is not warming as a result of human activity"). On the next screen, subjects were reminded of the belief that they had just reported. Subjects assigned to the "other" condition were then also told, "Someone from a prior study we conducted on mechanical turk (male, also living the United States) reported holding the same belief that you hold. Specifically, he also believes that Earth's climate is warming as a result of human behavior." Similar to Study 4, the text of the anonymous other person's belief matched the belief that subject had just reported holding. At the bottom of the screen, we told subjects that we were interested in how much voluntary control they thought that either they (self condition) or the other person (other condition) had over the belief and instructed them to continue to the next screen.

On the next screen, subjects who were assigned to the "generic" or the "supporting evidence" conditions were given a short writing prompt to complete. Subjects assigned to the "no task" condition did not complete any writing task, but instead immediately made a control judgment (described below). In the "supporting evidence" writing task, subjects read the following prompt:

> Please take a moment to think to about the reasons you have to believe that Earth's climate is warming as a result of human activity. In 1–2 sentences, please describe the evidence and arguments you have to suggest that this belief is true.

This prompt was minimally modified to match the belief that subjects reported at the beginning of the study, and, in the other condition, to refer to the evidence and arguments the other person might have, rather than those possessed by the self.

Subjects assigned to the "generic" writing condition were instead prompted as follows:

> Before you make a judgment about any specific belief, please take a moment to consider whether beliefs are, in general, the sorts of things people can freely choose and change (or not). In 1–2 sentences, please write about whether beliefs are freely chosen and changeable (or not) as well as why you think that.

The generic reasoning prompt was identical across the self and other attribution target conditions. As the prompt makes clear, subjects were free to write anything they wanted about belief control (e.g., that beliefs generally are highly controllable, or not controllable at all, etc.). Based on prior studies (e.g., Study 4 and the supplemental study reported in footnote 7), we expected that subjects would typically write about beliefs as controllable, and that directing their attention toward this generic conception would increase their attributions of belief control (even when considering a specific belief).

Subjects then attributed control either to themselves or to the other person regarding their shared climate change belief. As in prior studies, subjects reported their agreement with a passage asserting that they, or the other person, could voluntarily change their belief. For instance, in the self condition, if the subject affirmed a belief in anthropogenic climate change, they would read, "Now please rate to what extent you agree with the statement: I voluntarily hold my belief about climate change. If I wanted to, I could choose instead to believe that Earth's climate is *not* warming as a result of human behavior." In the other condition, for the same belief, subjects read, "Now please rate to what extent you agree with the statement: This other person voluntarily holds his belief about climate change. If he wanted to, he could choose instead to believe that Earth's climate is *not* warming as a result of human behavior." The final sentence in this passage always described the opposite of what the subject had reported believing at the beginning of the study. As in prior studies, subjects reported their agreement on a 1–7 rating scale (1 = *completely disagree*, 7 = *completely agree*).

At the end of the study, subjects indicated their age and sex, and were then debriefed. No other measures were collected.

Results

Mean control ratings for all conditions are presented in Table 3, and Figure 7, below.

We regressed control judgments on our target manipulation (self vs. other), our writing task manipulation (no task vs. generic vs. supporting evidence), and the interaction between these two factors. The writing task manipulation was dummy coded with "no task" as the reference level so that statistical comparisons reflect deviations from mean responses to the no task condition. The target manipulation (self vs. other) was mean-centered. This analysis revealed that, in the no task condition, subjects attributed to themselves less control ($M = 3.90$, $SD = 2.19$) than they attributed to others ($M = 4.89$, $SD = 2.04$), $b = -0.99$, $SE = 0.21$,

$t = -4.79$, $p < .001$, replicating our main findings from Studies 1–4.

This analysis also confirmed our predictions that writing about generic belief control, or supporting evidence for a belief, would increase, and decrease control judgments, respectively. Averaged across the self and other conditions, subjects in the generic condition ($M = 4.80$, $SD = 2.13$) reported higher control compared with the no task condition ($M = 4.40$, $SD = 2.17$), $b = 0.41$, $SE = 0.15$, $t = 2.79$, $p = .005$. Additionally, subjects who were assigned to think about the supporting evidence for the belief ($M = 4.01$, $SD = 2.16$) attributed lower control on average than did subjects in the no task condition, $b = -0.40$, $SE = 0.15$, $t = -2.63$, $p = .009$.

Follow-up analyses revealed that subjects attributed greater control in the generic condition relative to the supporting evidence condition, $b = -0.81$, $SE = 0.15$, $t = -5.27$, $p < .001$, and planned t tests showed that these differences replicated within the self and other conditions ($ps < .001$). Thus, consistent with our proposal, control judgments for self and other appear to be a product of whether people are thinking about belief control generically, or about supporting evidence for the belief.

Contrary to our expectations, the effects of our writing manipulations relative to the no task condition were not affected by whether subjects attributed control to themselves or another person. We observed no significant interaction between target (self vs. other) and the impact of the generic belief control task relative to the no task condition, $b = 0.25$, $SE = 0.30$, $t = 0.85$, $p = .397$. Likewise, the impact of considering supporting evidence (relative to the no task condition) was not moderated by whether subjects attributed control to themselves or another person, $b = 0.27$, $SE = 0.30$, $t = 0.91$, $p = .364$. In sum, we successfully manipulated attributions of control as a function of the writing task, such that thinking generically about beliefs increased control and thinking about supporting evidence decreased control; but the absence of any further interactions means that we did not observe support for the subtler predictions described in the Introduction to this study.[10]

Discussion

The results from Study 5 replicated prior findings and supported additional predictions made by our account of the self–other discrepancy. As in prior studies, subjects continued to attribute less voluntary control to themselves over a given belief than to another person who held the same belief.

Table 3

Means (and Standard Deviations) for Control Judgments Across Each of the Conditions in Study 5

	Attribution target	
Writing condition	Self	Other
No task	3.90 (2.19)	4.89 (2.04)
Generic	4.44 (2.21)	5.18 (1.99)
Supporting evidence	3.64 (2.19)	4.35 (2.07)

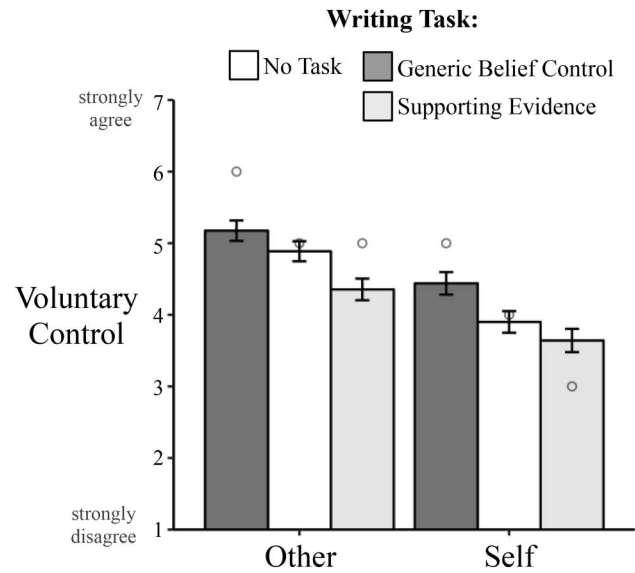*Note.* Judgments were made on 1–7 rating scales. Higher numbers indicate higher attributions of control.

Figure 7. Results from Study 5. Mean attributions of control (with standard errors) across each of the six conditions. Circles represent median values. Higher numeric values indicate stronger agreement with the passage asserting belief control.

However, Study 5 went beyond the prior studies by providing causal evidence for the underlying mechanisms we hypothesized. When subjects briefly considered and described the evidence that they or another person had to support a given belief, reports of voluntary control for both self and other decreased (compared with when subjects did not perform any writing task). This finding provides causal evidence for the correlational findings reported in Study 4. When people think of supporting evidence for a belief, which they are more prone to do for their own beliefs, this causes them to report lower ratings of belief control.

Similarly, Study 5 also provided causal evidence for the control-increasing impact of thinking about one's generic conception of belief. When subjects' attention was directed to their generic conception of belief, they attributed higher belief control (than when given no prior task). This finding is notable in that subjects were not instructed to write anything in particular about belief control. Rather, the manipulation affected people's

---

[10] Although the predicted interactions were not statistically significant, a series of exploratory analyses show that mean ratings of control trended in the pattern we predicted at the outset of this study. For instance, in the self condition, we observed no significant difference between subjects in the no task condition and the supporting evidence condition, $t(380.65) = 1.17$, $p = .244$, 95% CI [−0.18, 0.70], $d = 0.12$, though we did observe a significant increase in control between the no task and generic conditions, $t(404.12) = -2.47$, $p = .014$, 95% CI [−0.97, −0.11], $d = 0.26$. This is consistent with our supposition that when people are judging themselves, the no task condition is similar to the supporting evidence condition. The reverse was observed in the other-directed condition: Relative to the no task condition, subjects attributed significantly lower control in the supporting evidence condition, $t(392.72) = 2.59$, $p = .01$, 95% CI [0.13, 0.94], $d = 0.24$, but there was no difference between the no task and generic conditions, $t(403.20) = -1.44$, $p = .15$, 95% CI [−0.68, 0.10], $d = 0.14$. This is consistent with our supposition that when people judge others, the no task condition is similar to generic condition.

judgments merely by prompting them to recall their existing generic conception of belief, whatever it happened to be (e.g., low control, high control, etc.). This provides further evidence for the claim that people's default conception of beliefs is that they are controllable (see also footnote 7). And it provides complementary causal evidence for the correlational evidence reported in Study 4: When people think about their generic conception of belief, which they are more prone to do when thinking about others, they attribute greater control.

One pattern of results emerged in Study 5 which failed to support our theoretical account of the self–other discrepancy. Our theory predicts that if either the supporting evidence for a belief, or the generic conception of belief, are equally salient for both self and other, then the self–other discrepancy in control ratings should attenuate. However, despite selectively manipulating subjects' attention to each of these different facets of belief, we did not observe a significant attenuation of the original self–other discrepancy in either case. One possible explanation for this is that our manipulations were not strong enough to override people's default attentional focus in each case. For instance, even when directed to think about the supporting evidence that another person has for a given belief, a subject may not be able to do so as vividly as they can for their own belief. Similarly, even when directed to consider a generic conception of belief while answering a control question for the self, subjects may still spontaneously represent the supporting evidence for their own belief, thus somewhat weakening the overall effect of the manipulation. Nonetheless, despite not finding evidence for the most idealized pattern of results that our theory might predict, Study 5 clearly establishes two of its fundamental assertions, namely that thinking about supporting evidence lowers reported belief control, while thinking generically about beliefs increases reported control.

Because it probes the underlying mechanisms of the self–other discrepancy, Study 5 yields a deeper understanding of this discrepancy. It suggests that the self-other discrepancy arises because people focus on different kinds of information when thinking about the self as opposed to others. As a consequence, if we can manipulate what information people focus on when thinking about the self (say), then we should be able produce an analogous discrepancy within the self.

Study 6

In Study 6, to more powerfully manipulate what information people focus on when thinking about belief control, we manipulated the scope of the control question subjects answered. One group of subjects indicated how much control they have over a set of specific beliefs that they currently hold, similar to the procedure used in Studies 1–5. Another group of subjects indicated how much control they have over their beliefs "in general." We hypothesized that, when reasoning about specific beliefs that they hold, subjects would spontaneously access the supporting evidence they have for those beliefs, which should cause them to report relatively lower control. But, when attributing to themselves belief control in general, subjects would tend not to consider such supporting evidence, relying instead on a more generic conception of belief control, which should cause them to report relatively higher control.

Method

Participants. Three hundred two people (mean age = 37.7; 172 reported female, two unreported) were recruited from Amazon's Mechanical Turk (AMT) to participate in this study.

Design and procedure. Subjects were told that we were conducting a study about people's general assessments of their life. At the start of the study, subjects were randomly assigned either to judge the degree of control they have over their beliefs in general (beliefs in general condition) or to judge their control over a set of specific beliefs they reported (specific beliefs condition).

In the "beliefs in general" condition, subjects rated their agreement (1 = *completely disagree*; 7 = *completely agree*) with five passages probing to what extent they thought they had control over different parts of their life, including where they live, their habits, and their job. Four of these passages were filler items, and embedded within them was the single target passage, "My current beliefs are ones that I voluntarily hold. Specifically, I could change what I believe if I wanted to even if this means I was being wrong or immoral by doing so." The four filler items were included to best match the task demands of the "specific beliefs" condition (see below). To reduce acquiescence bias (given that control is generally seen as positive), all the passages were framed so as to highlight a possible downside of exercising control. For instance, in the work question (a filler item) subjects saw the statement "My current job is one that I voluntarily hold. Specifically, I could change where I work if I wanted to even if this meant having a worse job." The passages were each shown on a separate page and in a random order (see Appendix D for complete wordings).

In the specific beliefs condition, subjects wrote down a series of beliefs they currently held. Then, for each one, they indicated whether they held and could change that belief voluntarily. This happened in three stages. In the first stage, subjects were instructed to respond to the unconstrained prompt, "I believe that . . ." with the first belief that came to their mind. In the second stage, subjects responded similarly to a series of prompts intended to solicit beliefs about specific topics. These included beliefs about the subject's work ("I believe that my work . . ."), their family ("I believe that my family . . ."), and themselves ("I believe that I . . ."), as well as a moral belief ("I believe that it is wrong to . . ."). These prompts were shown on separate pages and in a random order. In the third stage, subjects in the specific beliefs condition then rated how voluntarily controllable each of the earlier beliefs they produced was. For instance, if a subject filled in the first prompt with, "most people are good," then they would respond to the following statement (bolding in the original):

> Earlier you wrote that you **believe that . . . most people are good**. Please indicate whether you agree or disagree with the following statement: **I voluntarily hold this belief. If I wanted to, I could choose to *not* believe that most people are good**.

Subjects indicated their agreement on 7-point rating scales (1 = *completely disagree*; 7 = *completely agree*) for each of the five concrete beliefs they had earlier produced. Each belief was shown on a separate page and in a random order.

At the end of the study, subjects from both conditions filled out a demographics form that asked them for their sex and age, and were then debriefed.

Results

Subjects assigned to the general condition on average agreed that, in general, they have voluntary control over their beliefs ($M = 4.99$, $SD = 1.92$). As predicted, this was significantly higher than subjects' average voluntary control ratings in the specific belief condition ($M = 3.96$, $SD = 1.44$), $t(299.96) = 5.34$, $p < .001$, 95% CI [0.65, 1.41], $d = 0.60$. Furthermore, each specific belief prompt in the specific belief condition produced voluntary control ratings that were, on average, lower than the ratings in the general condition. Beliefs generated by the unconstrained belief prompt ($M = 4.45$, $SD = 2.11$) were rated as less voluntarily controllable than beliefs in general ($M = 4.99$, $SD = 1.92$), $t(263.59) = 2.32$, $p = .021$, 95% CI [0.08, 1.01], $d = 0.27$, as were beliefs about work ($M = 4.38$, $SD = 1.87$), $t(281.67) = 2.80$, $p = .005$, 95% CI [0.18, 1.05], $d = 0.32$, the subject themselves ($M = 4.13$, $SD = 2.10$), $t(264.07) = 3.66$, $p < .001$, 95% CI [0.40, 1.33], $d = 0.43$, their family ($M = 3.49$, $SD = 2.14$), $t(260.97) = 6.30$, $p < .001$, 95% CI [1.03, 1.97], $d = 0.74$, and finally, morality ($M = 3.36$, $SD = 2.13$), $t(261.89) = 6.87$, $p < .001$, 95% CI [1.17, 2.10], $d = 0.81$ (see Figure 8).

Discussion

These results provide support for the hypothesis that the accessibility of supporting evidence modulates people's self-attributions of belief control. When asked generally whether they have control over what they currently believe, subjects tended to indicate that they do have such control. Yet, when asked about specific beliefs that they held, subjects attributed to themselves a lower degree of control. We observed this specific-general difference across every specific belief category we assessed, including beliefs about the subjects' themselves, their lives, and their morality. Most strikingly, even when subjects considered the very first specific belief that came their mind, with no constraints on its content, they attributed lower control to themselves over this specific belief than when they considered their beliefs "in general"; thereby suggesting that it was not the specific belief prompts we used that gave rise to the observed differences. In sum, people's typical default assumption about their own control over their beliefs does not match their concrete experience of belief control.

These results are consistent with and add further detail to the picture that has emerged from the prior studies. The high general ratings of control observed in the present study are consistent with the high ratings of belief control subjects gave when asked about people in general (footnote 7), with the high ratings of belief control given for others in Studies 1–5, and with the link between thinking generically about beliefs and ratings of high control in Studies 4 and 5. Similarly, the lower ratings of control for specific beliefs broadly match the ratings given for the self (over specific beliefs) in Studies 1–5.

In Study 7, we tested a final prediction that combines the self–other and specific–general differences observed in the earlier studies. Our theoretical account explains these two phenomena in the same way. When people consider the amount of control they have over specific beliefs that they hold, they are confronted with the psychological constraints on such control, which drives down judgments of their own belief control. But, when people attribute control to others (either over specific beliefs or in general), or when they attribute control to themselves in general (without considering their specific beliefs), they rely on a more generic conception according to which beliefs are quite controllable. Accordingly, both the self–other discrepancy and the specific-general discrepancy arise because of a failure to integrate the experience of trying to control or change one's own beliefs with a more general model of belief controllability.

If this reasoning is correct, then the self–other discrepancy we observed in Studies 1–5 should be most pronounced when people are asked to consider specific beliefs, and should be attenuated or eliminated when people consider belief control in general. We tested this prediction in Study 7.

Study 7

Method

Participants. We recruited 597 subjects (mean age = 35; 365 reported female, 231 reported male, 1 unreported) from Amazon's Mechanical Turk.

Design and procedure. This study used a 2×2 between-subjects design. Subjects were randomly assigned to one of four conditions created by crossing belief condition (specific belief vs. beliefs in general) and target condition (self vs. other).

At the beginning of the study, subjects were asked to think of a person close to them, write down that person's initials in a text box, and indicate their relationship to the person (they could indicate best friend, close family member, romantic partner, or spouse, or write in how they would describe the relationship). For ease of exposition, we will refer to this person as the "Close Other," or CO.

On the next page, subjects in all conditions entered four beliefs that they shared with the CO. They did so by completing four sentence fragments of the form, "We both believe that . . .". Subjects were instructed that the beliefs could be about anything, but had to be the first ones that came to mind.
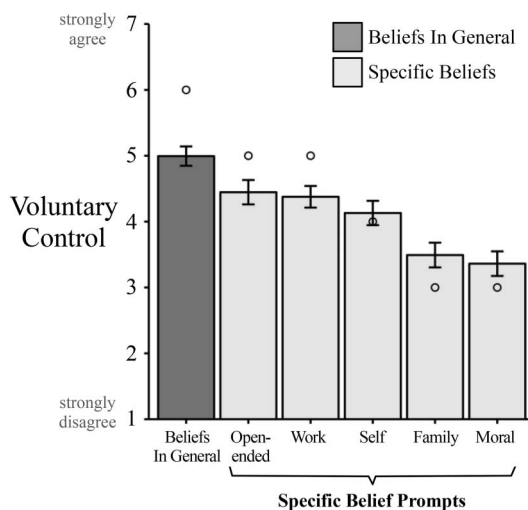

Figure 8. Means (and standard errors) of subjects' agreement with statements that they could choose to change their beliefs in general (dark bar), or that they could choose to change a particular belief (light bars); Study 6. Circles represent median values.

In the specific belief condition, subjects were then presented with each of the four beliefs they had just entered (which they jointly believed with their CO). For each belief, they reported their agreement with statements that they or their CO had voluntary control over the belief. The specific contents of each belief were dynamically inserted within the corresponding belief statement. For instance, in the Self condition, subjects would see the following:

> You wrote that you believe that . . . [BELIEF]. Please indicate whether you agree or disagree with the following statement: I voluntarily hold this belief. If I wanted to, I could choose to not believe that [BELIEF].

In the other condition, the statements were the same except that the initials of the CO were inserted into the statement:

> You wrote that [INITIALS] believes that . . . [BELIEF]. Please indicate whether you agree or disagree with the following statement: [INITIALS] voluntarily holds this belief. If [INITIALS] wanted to, [INITIALS] could choose to not believe that [BELIEF].

The order of presentation of the four beliefs was randomly determined for each participant.

In the "beliefs in general" condition, subjects did not return to the specific beliefs they had reported. Instead, they indicated their agreement with four general passages about either their own, or the CO's life, similar to those used in Study 5 (see Appendix E for full text). Only one of these passages—the belief passage—was relevant to our interests. The other passages were fillers that were included to match the specific beliefs condition in length. The key belief passage corresponded closely to the wording of the statements in the specific-belief conditions. For instance, in the self condition, subjects indicated their agreement with the passage, "My current beliefs are ones that I voluntarily hold. Specifically, I could choose to hold different beliefs if I wanted to even if this meant being wrong or immoral." In the other condition, the initials of the CO were dynamically inserted into the passage as follows, "[INITIALS]'s current beliefs are ones that he/she voluntarily holds. Specifically, [INITIALS] could choose to hold different beliefs if he/she wanted to even if this meant being wrong or immoral." The distractor passages were about the subject's (or CO's) behavior, work, and home (e.g., "My current home is one that I voluntarily live in. Specifically, I could change where I live

if I wanted to even if it meant changing many other parts of my life."). The order of the four passages was randomly determined for each subject.

All ratings were made on a 7-point rating scale (1 = *completely disagree*; 7 = *completely agree*). At the end of the study, all subjects reported their age and sex, and then were debriefed and paid.

Results

Examples of subjects' shared beliefs are provided in Table 4, below. As planned, we averaged together subjects' four agreement ratings in the specific belief condition ($\alpha = .88$). We then conducted an ANOVA regressing agreement ratings on attribution target (self vs. other), belief condition (general vs. specific), and their interaction. We observed a main effect of target such that subjects rated their own control ($M = 4.86$, $SD = 2.04$) lower than that of others ($M = 5.19$, $SD = 1.85$), $F(1, 593) = 4.27$, $p = .039$, $\eta_G^2 = 0.07$. We also observed a main effect of belief condition such that control ratings for specific beliefs ($M = 4.78$, $SD = 2.01$) were lower than control ratings for beliefs in general ($M = 5.27$, $SD = 1.86$), $F(1, 593) = 9.84$, $p = .002$, $\eta_G^2 = 0.016$. And, as predicted, these effects were qualified by a significant interaction, $F(1, 593) = 4.67$, $p = .031$, $\eta_G^2 = 0.008$ (see Figure 9).

As predicted, in the specific belief condition, subjects' ratings of their own control ($M = 4.44$, $SD = 2.08$) were significantly lower than their ratings of others' control ($M = 5.11$, $SD = 1.90$), $t(292.4) = 2.89$, $p = .004$, $d = 0.34$. However, there was no such self–other difference in the general belief condition (self, $M = 5.28$, $SD = 1.92$; other, $M = 5.27$, $SD = 1.80$), $t(296.68) = 0.062$, $p = .951$, $d < 0.01$. Examined another way, subjects rated their control over their nominated specific beliefs ($M = 4.44$, $SD = 2.08$) significantly lower than they rated their control over beliefs in general ($M = 5.28$, $SD = 1.92$), $t(293.61) = 3.61$, $p < .001$, $d = 0.42$, but there was no corresponding difference between ratings of others' specific ($M = 5.11$, $SD = 1.90$) and general ($M = 5.27$, $SD = 1.80$) belief control, $t(295.83) = 0.72$, $p = .472$, $d = 0.08$.

Discussion

The results from Study 7 replicated the findings from Studies 1–5 showing that, when considering control over specific beliefs,

Table 4

Examples of Specific Beliefs Submitted by Subjects in Study 7

	Beliefs shared with close other			
Subject	1	2	3	4
A	Lifting is healthy.	Video games are fun.	Money is an important measure of success.	Intelligence is important.
B	Science is important.	Racism is bad.	Free healthcare should be a human right.	The world is doomed.
C	We are soulmates.	Abortion is wrong.	Our kids come first no matter what.	We have to make changes to improve our life.
D	We live comfortably.	We have the best children.	Parenting is tiring.	We have a bright future.
E	Cannabis is more-or-less harmless.	Social norms are too repressive.	Skateboarding is the best sport.	Video games are great fun.
F	There is a God.	Fat food makes us feel bad.	Our son is amazing.	A smaller house is better.
G	Cheating is bad.	Religion is dumb.	[Person] is dumb.	The earth is round.
H	Paul George will go to the Lakers.	Lebron is the best basketball player in the league.	"Impractical Jokers" is hilarious.	Going to the gym has benefits beyond enhancing physical appearance.

Note. Subjects were instructed to write about four beliefs they shared with a close other whom they had nominated on the previous screen. Each belief was elicited with the sentence fragment, "We believe that. . . ." Responses have been lightly copyedited (e.g., capitalization and punctuation).
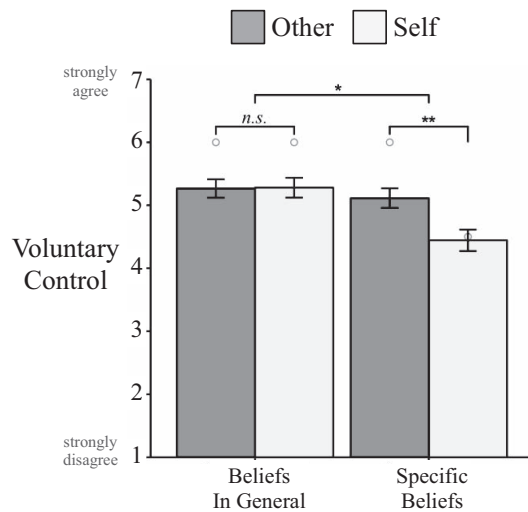
Figure 9. Means (and standard errors) for subjects' agreement ratings in Study 7 across conditions. Circles represent median values. * $p < .05$. ** $p < .01$.

subjects judged themselves to have less control than others. As in Studies 3, 4, and 5, we observed this even when subjects rated beliefs they shared with the other person. Going beyond the earlier studies, we observed this discrepancy even though subjects were no longer constrained to consider a specific belief supplied by the experimenter, but were able to consider whichever beliefs first came to mind. We also replicated Study 6's finding that people judge themselves as having high control over their beliefs in general, but less control over specifically identified beliefs. This finding is noteworthy given that subjects wrote down four specific beliefs they held before indicating their control over beliefs in general—thereby showing that people maintain a general impression of belief control unless asked very directly about their control over specific beliefs. Most importantly, however, we found that the discrepancy between self and other control ratings attenuated in the general beliefs condition. That is, when people judged their own and others' control over beliefs in general, the differences in control ratings we observed in Studies 1–5 were no longer apparent. These findings are consistent with the theory we have articulated; namely, that self–other differences in belief control arise from the experience of psychological constraints on one's own ability to control specific beliefs, and from the subsequent failure to generalize this experiential understanding to others, as well as to integrate it within one's own general theory of belief control.

General Discussion

Although control is a ubiquitous and important judgment in social cognition, determining behaviors such as blame and credit, and reward and punishment, it has not been extensively explored in relation to people's beliefs. In this paper we investigated how much control people attribute to themselves as compared with others. A body of past research suggests that people might inflate self-directed attributions of doxastic control. After all, control is highly desirable, and attributions of one's own control are sometimes inflated as a consequence (e.g., Burger, 1986; Miller &

Norman, 1975). Recent research has even indicated that people tend to attribute to themselves greater free will over their own actions than they attribute to others, consistent with well-established self-enhancement biases (Pronin & Kugler, 2010). Despite these findings, we hypothesized that things would be different in the realm of beliefs.

Our starting premise was that people are constrained in their ability to pick and choose what they believe. Specifically, when people perceive good reasons to believe something, these reasons limit their ability to choose to believe otherwise (James, 1937, see Introduction). However, we also know that people often fail to appreciate the constraints that other people operate under, and that this failure can be especially pronounced for psychological constraints (Gilbert & Malone, 1995). That is, people often fail to appreciate that others' behavior is influenced by psychological constraints such as emotion or stress, which is representative of a broader failure to appreciate the complexity of others' inner lives (Johnson, 1987; McFarland & Miller, 1990; Miller & McFarland, 1987; Pronin, 2009). Accordingly, when it comes to belief change, people may also fail to appreciate the psychological constraints limiting others' ability to change their beliefs. This line of reasoning therefore generates a prediction that contrasts with what might be predicted on the basis of self-enhancement research, namely that people will tend to attribute to themselves less control over their beliefs than they attribute to others, at least insofar as they are considering concrete instances of belief.

We conducted a series of studies to adjudicate between these two possibilities. These studies consistently demonstrated that, when reasoning concretely about beliefs, subjects tended to attribute less control to themselves than to others. This self–other discrepancy occurred both when subjects held beliefs different from those held by another person (Study 1), as well as when they each held identical beliefs (Studies 2, 3, 4, 5, and 7). It arose not only when subjects considered strangers (Studies 1, 2, 4, and 5), but also when they considered close others (Studies 3 and 7). And it arose not only for beliefs supplied by us as the experimenters (Studies 1–5) but also for beliefs that subjects themselves supplied (Study 7). Thus, it appears to be robust.

To further explain how this discrepancy comes about, we hypothesized that belief control judgments are the product of (at least) two distinct sources of information. On the one hand, people hold a generic conception of beliefs as highly controllable. Absent other overriding information, people rely on this conception when attributing belief control to others, as well as when they attribute belief control (in general) to themselves. On the other hand, when they have access to it, people also incorporate information about the evidence that supports a given belief. In contrast to the default conception of belief, consideration of the evidence and reasons in support of a belief tends to lower judgments of voluntary belief control (because this evidentiary information often encapsulates genuine psychological constraints on belief change; cf., Alston, 1988; James, 1937). However, although people can apply their default conception of belief to themselves and others with equal ease, they more easily incorporate evidentiary information in their self-directed attributions of control. This is because they can introspect on the evidentiary basis of their own beliefs but lack such direct access to the evidentiary basis of others' beliefs. Accordingly, because people do not readily integrate evidentiary information into their control judgments of others, they tend to

attribute greater belief control to others than to themselves. In line with this reasoning, the divergent control attributions for self and other should reflect differences in people's attention to these two sources of information. We found strong support for this theorizing in Studies 4–7.

In Study 4, we asked subjects to report what they thought about when attributing belief control to themselves or another person regarding a given belief, coding both for generic statements about belief controllability as well as mentions of evidence or reasons in support of the given belief. As expected, for both self and other, we found that when subjects thought generically, they tended to attribute high control. But, when they thought about the evidence they or the other person had for the belief, they attributed lower control. Furthermore, subjects were much more likely to think about beliefs generically when thinking about others, whereas they were much more likely to think about supporting evidence when thinking about themselves. Thus, differences in what information people spontaneously consider when thinking about the self versus another person appear to underlie the discrepant ratings of belief control.

Study 5 provided causal evidence that directing people's attention to either their own generic conception of belief control, or to the supporting evidence for their own or another's belief, affects their attributions of control. When subjects' default conception of belief control was salient, they attributed high control toward themselves and others. But, when the supporting evidence for their own (or the other's) belief was made salient, subjects' attributions of control were significantly lower. Taken together, Studies 4 and 5 provide causal evidence that reasoning about the self versus another person causes different information to become salient (evidence vs. a more generic conception of beliefs, respectively), which in turn impacts attributions of control.

Studies 6 and 7 produced additional support for this account by testing two further predictions—one concerning when people will grant themselves high versus low control over beliefs, and the other concerning when the previously observed self–other discrepancy will attenuate. Our model predicts that if people can be led to think about their own belief control without accessing the factors that constrain it, then they should rely on their default conception of beliefs, and in so doing, attribute to themselves high control. Supporting this idea, when we asked subjects whether "in general" (i.e., without thinking about any specific belief that they hold) they have control over what they believe, they tended to grant themselves high control—significantly higher than the control they granted themselves when they thought about specific beliefs that they held (Studies 5 and 6). Furthermore, supporting the second prediction, when one group of subjects judged belief control "in general," the self–other difference in control observed in all prior studies completely attenuated (Study 7). Thus, our proposal not only accounts for the original self–other discrepancy, but it also accurately predicts when this discrepancy will disappear, as well as when an isomorphic discrepancy can be reproduced within the self.

Attributions of Intentionality and Voluntary Control

In Studies 1 and 2, the self–other discrepancy consistently emerged for judgments regarding the ability to voluntarily change beliefs but not for judgments regarding whether those beliefs were intentionally chosen in the first place. At first blush, this result

seems out of step with our theorizing that perceptions of constraint drive down self-directed attributions of control relative to other-directed attributions. And indeed, the difference across these two measures was not one that we had initially predicted.

However, in hindsight, the lack of effect on the choice measure can be reconciled with the similar lack of a self–other difference when beliefs were considered in a generic fashion. When an individual considers whether a particular belief was "intentionally chosen," he or she must make a retrospective judgment about an episode of belief formation in the past. Arguably, under these circumstances, the individual is psychologically distant from the actual experience of forming or changing a belief. Similarly, when individuals judge the controllability of their beliefs "in general," they are also psychologically removed from any immediate effort to control their beliefs, and therefore, do not feel the tug of low control that would typically arise if they had a specific belief in mind. In both cases, this psychological distance may result in people relying on a default conception of belief which posits high control. If this reasoning is correct, then a self–other discrepancy for attributions of intentional choice should occur when the judger is psychologically close to the constraints on belief formation— that is, during the process of initial belief formation; future research should explore this possibility.

Alternative Explanations

One alternative explanation for the observed self–other discrepancy is that it reflects a form of self-enhancement (or impression management). Voluntarily controlling a belief merely because one wants to may be seen as violating or disrespecting the norms of belief (Clifford, 1877; James, 1937; Ståhl et al., 2016). Because of the negative connotations associated with this sort of control, subjects may be motivated either (a) not to perceive it in themselves (self-enhancement), or (b) not to report it to an experimenter, even when they do perceive themselves as having it (impression management). However, if either one of these alternatives explained the tendency to attribute less volitional control to the self than to others, then we should have observed a similar discrepancy for judgments about the ability to *act* poorly, because the ability to act in norm-violating ways has similar negative connotations. Yet, we observed no such self–other difference in attributions of the ability to carry out immoral acts that were matched in content to the respective beliefs (Study 3). Thus, this alternative explanation is not well supported by the evidence.

However, perhaps a subtler form of this alternative explanation can account for the results. What if people believe they have less control over their beliefs than do others because they consider themselves to be more objective than others? Indeed, prior work on the phenomenon of naïve realism has shown that people often assume they are objective perceivers of the world (Pronin, Lin, & Ross, 2002; Pronin et al., 2004; Ross & Ward, 1996). A natural corollary of assuming that one perceives the world "directly," and forms objective beliefs on that basis, is the idea that one's beliefs are outside one's control. After all, if your beliefs are dictated by the way the world actually is, and not by some extra input on your part, then what you believe must be limited to what is objectively true. Perhaps then, even when people share a belief with a close other, they hold a prior assumption that the other person is less

objective than they are, and therefore more able to voluntarily choose to believe false or unjustified propositions.

One reading of this alternative explanation is that people hold a generalized assumption that their own beliefs are more grounded in objective processes than are others' (Ehrlinger, Gilovich, & Ross, 2005). This straightforward interpretation could potentially explain some of our results, but it is incompatible with the full set of our findings. A default belief in one's own superior objectivity should apply just as readily to specific beliefs as it does to beliefs considered in general. But, in that case, this account cannot explain the results of Study 6, which revealed a self–other difference only for specific beliefs, and not at all for beliefs considered in general. Accordingly, this rendering of an alternative explanation inspired by naïve realism seems inadequate.

However, perhaps the belief in one's superior objectivity is triggered only when people consider trying to change a specific belief that they hold. On this account, considering one's own belief concretely prompts people to judge that they are especially objective (which does not happen otherwise), and this judgment in turn leads people to infer that they lack voluntary control (i.e., because of their objectivity). As a consequence, people attribute to themselves low control over the specific belief in question. This alternative account posits that self-attributed objectivity is the pivotal mediating judgment in the chain of inference underlying our results, with belief constraint (that is, a lack of control over one's beliefs) merely being an inferential byproduct of this assumption. In contrast, our account posits that directly experienced constraint is the key factor.

This alternative could potentially explain our findings, and we cannot rule it out definitively. However, three observations diminish its plausibility. First, it seems somewhat implausible that people only believe they are objective when they consider concrete beliefs—indeed, related research shows that people attribute objectivity to themselves in both specific and general judgment contexts (see, e.g., Ehrlinger et al., 2005). Second, it is hard to see how an attribution of one's own objectivity should lead directly to an inference of belief constraint without some additional consideration of how relevant evidence acts as a constraining factor. But in that case, this alternative becomes less parsimonious than our own—it posits more intermediate psychological steps en route to a judgment of the self's lower belief control. We propose, instead, that people need only recognize psychological constraints, such as the evidence or arguments in support of a given belief, without any need for an additional inference about their own or others' objectivity. Third, an explanation for our finding couched in the framework of naïve realism assumes that ordinary people believe that objectivity entails constraint. However, we are not sure that people hold this view. Instead, they may think that "being objective" is, in principle, a free choice that they and others have. In this case, attributions of objectivity would be largely orthogonal to attributions of control: Whether or not someone was objective in their belief formation, people may judge that, in principle, they could choose to be biased (or not) if they wanted to be. Our own data cannot directly speak to this possibility, but we think such coincident judgments of high objectivity and high (general) control could be borne out in future research. This would in turn undermine the assumption that underlies a naive realist explanation of our data—namely that greater perceived objectivity *necessitates* greater perceived belief constraint.

In sum, there are several considerations that incline us away from interpreting the present findings solely within a naïve realist theoretical framework. Although it may be that one source of the sense of constraint over one's beliefs stems from an assumption of the self's greater objectivity, this source alone seems insufficient—an explanation of this sort cannot parsimoniously explain the attenuation of the self–other discrepancy for general beliefs, and it relies on an ad hoc assumption about when people attribute objectivity to themselves. Nor can it easily accommodate (possible) conjoint judgments of objectivity and choice. For this reason, we see the findings presented here as complementing theories of objectivity and bias attribution, rather than being subsumed by them. That said, better understanding the nuanced relationship between judgments of belief control and attributions of objectivity or bias is an important goal for future research.

Relationship to Other Actor-Observer Discrepancies

As noted in the Introduction, the discrepancy between self-directed and other-directed attributions of belief control is caused, in part, by the fact that "actors" have different kinds and amounts of information about themselves than "observers" do. This informational difference plays an important role in many other "actor–observer" discrepancies. For instance, it helps explain why people attend to mental and behavioral events differently for self and other (Malle & Pearce, 2001), why they are inclined to explain different behaviors for self and other (Malle & Knobe, 1997b), and finally, why they provide different explanations for the behavior of self and other (e.g., Jones & Nisbett, 1972; Malle, Knobe, & Nelson, 2007; Storms, 1973). In a similar way, a difference in informational access underpins the present self–other discrepancy: It is based in part on the fact that actors are keenly aware of the evidentiary constraints that limit their own capacity to voluntarily change their beliefs, whereas observers are less aware of these constraints. Thus, our findings contribute to a growing list of ways that informational differences between actors and observers generate discrepancies in how they evaluate and understand behavior, and so can be interpreted as demonstrating a new kind of actor–observer bias.

However, this broad difference in information access does not on its own explain the self–other discrepancy we document. There are several additional, quite specific factors that are unique to the present demonstration. It relies on the fact that beliefs are typically formed in response to evidence. It relies on the fact that apprehending supporting evidence constrains a person's sense of their voluntary ability to change a given belief. It relies on the fact that people by default hold a conception of beliefs as highly controllable. And it relies on the fact that this default conception of high controllability is often at odds with the sense of belief control that emerges as a result of apprehending pertinent supporting evidence for a belief. To the best of our knowledge, none of these specific factors are present in any of the existing actor–observer demonstrations. Thus, although the present self–other discrepancy shares the same fundamental starting point as several other actor–observer biases (a difference in information access), the more nuanced factors that drive it are novel and independent of any existing biases within the literature.

Implications for Belief-Based Conflict

Control is a precondition for blame and criticism in everyday life (Markman & Tetlock, 2000; Weiner, 1995). To blame or criticize someone, the critic must be able to show that their criticism is deserved (e.g., that the target had control over what she is being blamed for), otherwise they risk castigation from others (Malle et al., 2014). Victims of blame who feel that they have been blamed unfairly (e.g., for something that they had no control over) may subsequently defend themselves and blame the critic in return (e.g., Dersley & Wootton, 2000; see Malle et al., 2014, for discussion). Thus, a divergence of opinion regarding who had control over what can lead to intractable conflict between parties that causes long-term damage to relationships. Although most work has documented the relationship between control and responsibility in the domain of behavior (e.g., Weiner, 1995), a tight association between control and responsibility has also recently been documented in the domain of belief (Cusimano & Goodwin, 2019). Accordingly, self–other discrepancies in attributions of belief control may similarly exacerbate conflicts in which one person holds another person responsible for their beliefs.

We have demonstrated that when two parties disagree about what to believe, they will likely disagree not only about who is wrong, but also about who is even capable of changing their mind. This secondary disagreement may make compromise especially difficult to achieve, because each party believes that they are literally incapable of voluntarily changing their own mind, while simultaneously believing that the other party can change theirs. For instance, if one party to a disagreement believes that they are incapable of adopting the other party's view, then they will not be motivated to compromise their view. If, at the same time, the other party expects the first person to compromise, and believes that doing so is within that person's control, then the fact that the first person does not do so should lead to judgments that she is culpably intransigent, and therefore deserving of (perhaps additional) blame or punishment. And because the first person believes that she cannot simply choose to change her mind, he will resent being blamed or punished for not doing so and may criticize the other person in return. And so on.

Well into this research, we discovered a real-world example that illustrates elements of this speculative drama. On April 10th, 2018, Megan McArdle, a conservative columnist, published an article voicing a widely held thesis in conservative circles: that liberals are biased against conservatives on the basis of their beliefs, holding derisive attitudes about them that are unfair in the same way that prejudice directed toward other minorities is unfair (McArdle, 2018). The next day, Hamilton Nolan published a rebuttal to McArdle, offering a defense of the treatment of conservatives at the hands of liberals (Nolan, 2018). Tellingly, Nolan's article was titled "Ideology is Choice," and his central argument went as follows: "unlike race and gender and sexual persuasion, it [being conservative] is an intellectual choice. It can be changed at any time." Thus, McArdle chastised liberals for holding conservatives responsible for their beliefs (on the grounds that beliefs are uncontrollable)[11], whereas Nolan defended the practice (on the grounds that beliefs are controllable). These authors, much like the subjects in our studies, appear to conceptualize each other's agency over their own beliefs differently, which in turn appears to lead to very different perspectives on their respective moral responsibilities.

The Relationship Between *Experienced* and *Actual* Control

How do lay judgments of control correspond to *actual* control? Are believers, who attribute relatively less control to themselves, or observers, who attribute relatively more control to those believers, more accurate?

The dominant view among psychologists and philosophers is that people do not have direct control over what they believe (Alston, 1988; Epley & Gilovich, 2016; Gilbert, 1991, 1993; James, 1937; Wegner, 1994; Wilson & Brekke, 1994). According to these scholars, people cannot simply believe whatever they want to because beliefs are spontaneous, automatic responses to information, akin to perceptual processes. If people have any control over what they believe, the story often goes, it is highly indirect (e.g., exerted only by controlling the information one is exposed to; e.g., Alston, 1988; Epley & Gilovich, 2016), or extraordinarily rare (e.g., choosing what to believe is possible only under conditions of extreme uncertainty or ambivalence; James, 1937; Sloman et al., 2010; though see Steup, 2017, for a dissenting view).

However, there is a dearth of empirical work on whether people can exert deliberate, voluntary control over their beliefs. In fact, as far as we know, there is no direct evidence testing whether this is possible. Considerable work on the phenomenon of motivated reasoning and self-deception indicates that people's desires and preferences can impact their beliefs, but this research does not speak directly to the question of voluntary control. That said, the evidence from these studies is consistent with the view that people lack voluntary control over their beliefs (e.g., Klein & Kunda, 1992; Kruglanski & Webster, 1996; and Kunda, 1990, for a review). For instance, evidence to date suggests that desires influence beliefs indirectly and unconsciously—by influencing how incoming information is interpreted, what information is stored in memory, and what information is subsequently retrieved (Epley & Gilovich, 2016; Pronin et al., 2002; see Kunda, 1990, for an extended discussion). When psychologists give people direct evidence of their biases, remove ambiguity, or prevent biased retrieval, motivated reasoning and self-deception seem to all but disappear (e.g., Bar-hillel & Budescu, 1995; Sloman et al., 2010). Moreover, when people are given strong evidence in favor of some conclusion, they often heed this evidence, even when the conclusion is upsetting or undesired (see Wood & Porter, 2019, for recent experimental evidence; see Petty & Cacioppo, 1986, for a review).

Most philosophical arguments for the putative uncontrollability of beliefs are grounded in self-reports of low control (see Introduction). Our findings suggest that these self-reports are only partially shared by lay people (as compared with philosophers). The moderate ratings of control in our studies suggest that lay

---

[11] Evidence that McArdle regards ideology as uncontrollable comes from her public discussion of her piece at the time of its publication, available at the following link: https://twitter.com/asymmetricinfo/status/984162049440800768. Here, she writes in defense of her claim that conservative ideology is the target of *unjust* prejudice that, "you can't just decide to believe something different about a fact you can see" as well as, "you cannot decide to unbelieve something you hold to be deeply true." See the online supplemental materials for archived screenshots.

people appear to attribute to themselves considerably more control than do philosophers (although reported control drops significantly when people call to mind the evidence they have for a given belief; Studies 4 and 5). However, we should be wary of the idea that the experience of control is diagnostic of actual control. Introspection can be a poor guide to our mental processes or capacities, and prior work shows that people are notoriously poor at reasoning about the origin and quality of their beliefs (Davison, 1983; Hauser, Cushman, Young, Kang-Xing, & Mikhail, 2007; Nisbett & Wilson, 1977; Pronin et al., 2002). There is also reason to think that people are poor judges of what they do and do not have control over (e.g., Buehler, Griffin, & Ross, 1994; Koehler & Poon, 2006). Recent evidence suggests that morality constrains people's sense of choice, such that choosing between moral options feels more constrained than choosing between nonmoral options (Kouchaki, Smith, & Savani, 2018). Yet, people who judge that they "could never" harm someone at Time 1, may end up doing so at Time 2 when incentivized the right way. For instance, people who are sexually aroused indicate a greater willingness to engage in morally questionable behaviors than their unaroused selves predict (e.g., falsely telling a partner that they love them just to increase the chance of having sex with that person, Ariely & Loewenstein, 2006). For this reason, external observers may sometimes be better judges of actors than actors are of themselves (e.g., Bass & Yammarino, 1991; MacDonald & Ross, 1999; Risucci, Tortolani, & Ward, 1989). Perhaps the people around us, who see our epistemic foibles more clearly than we do, are better informed about our capacity to choose our beliefs capriciously than we are.

In sum, our finding that people experience their beliefs as partially outside of their voluntary control provides prima facie evidence that this is indeed the case. However, in our view, the question is far from settled, as one's internal sense of control is not necessarily diagnostic of one's actual capacity. The bottom line is that without direct information about people's voluntary abilities to change their beliefs, we cannot know whether or how people are erring in their judgments in the present studies.

Limitations and Future Directions

Subjects in our studies (with the exception of the study reported in footnote 7) were exclusively recruited through Amazon's Mechanical Turk. Although samples recruited from AMT are more representative of the United States than typical university student samples, individuals on AMT tend to be less religious, wealthier, and better educated than the average person in the United States (Paolacci & Chandler, 2014). Additionally, our entire sample consisted of people living in the United States who, like other so-called WEIRD populations, are wealthier and better educated than most people in the world, and are predominately Christian (Henrich, Heine, & Norenzayan, 2010). Cross cultural work has revealed striking differences in how different groups think about individuals' agency. Of particular note, people in some non-U.S. cultures appear to attribute less agency to individuals than do people in the United States (e.g., Iyengar & Lepper, 1999; Kitayama, Snibbe, Markus, & Suzuki, 2004; Miller, Das, & Chakravarthy, 2011; Morris & Peng, 1994; Savani, Markus, Naidu, Kumar, & Berlia, 2010; Spector et al., 2004). For instance, compared with children in the United States, Nepalese children are more inclined to view some behaviors as constrained by social rules and

therefore outside of their control, with this cultural gap widening with age (Chernyak, Kushnir, Sullivan, & Wang, 2013). In a similar vein, Indian adults appear to be less likely than U.S. adults to construe everyday behaviors as choices (Savani et al., 2010). Of clearest relevance to the present studies, some work suggests that Christians tend to attribute more control to others over deviant mental states (e.g., consciously entertaining thoughts of having an affair) than do Jews, thus showing evidence for cultural moderation with respect to mental states in particular (Cohen & Rozin, 2001). In light of this sort of evidence, we should not automatically assume that the results from our studies will replicate across different cultural or religious contexts.

Although we are uncertain as to whether our findings will generalize to all cultures, our findings do suggest an important direction for cross-cultural work. Specifically, future work measuring attributions of belief control should measure both lay theories of belief control and the introspective experience of belief control. One advantage of measuring both is that doing so could reveal different amounts of variation between these two measures of control across cultures. For instance, if we assume that beliefs are indeed uncontrollable to a significant degree (see above), we should expect that the felt experience of low control will vary little from culture to culture. By contrast, the lay theory of belief, which may be influenced by highly variable norms (e.g., religious norms, Cohen & Rozin, 2001), or folk theories of agency (see paragraph above), may be more likely to vary across cultures. For this reason, we speculate that self–other differences in belief control are most likely to arise in cultures where the lay theory of belief posits high control, as it is in these cultures where this lay theory will most likely diverge from the felt experience of belief.

Another limitation of our studies regards the limited range of beliefs that we sampled. The beliefs in Studies 1–5 were highly abstract, complex, or value-laden (e.g., belief in God, the correct policy for genetically modified foods, the wrongness of not returning money to its rightful owner). We addressed this in Studies 6 and 7 by using beliefs that subjects themselves provided—specifically, the first beliefs that came to mind. This yielded a considerably wider sampling of belief contents (see Table 4 for a list of examples). Yet, it still leaves open the question of how people judge their own control relative to that of others for very simple, concrete beliefs (e.g., "there is a two thirds chance of pulling a marble out of the bucket," "there is a quarter in my pocket," "it is raining"). We are ambivalent about whether to expect the same discrepancy in cases such as these. It may be that the self–other difference is attenuated or eliminated given that the relevant constraints on belief change are far more apparent for beliefs of this sort. Continuing to delimit the bounds of the self–other discrepancy remains a valuable goal for future research.

Although the present article focuses only on the constraints on belief change, it may be that other mental states, including desires, evaluative attitudes, and emotions, are subject to similar constraints. If they are, then we might expect similar self–other discrepancies in perceived control—particularly in light of past work showing that people generally attribute high control to others over many mental states (Cusimano & Goodwin, 2019). Indeed, there is already one reason to expect the self–other discrepancy to extend to other mental states, namely, that a person's beliefs often play a pivotal role in determining his or her other mental states. For instance, if someone is depressed because she believes she will not recover from a severe

illness, an observer may think she is more capable of cheering up than she herself does, precisely because the observer judges her as more able to change her belief about her prognosis than she herself does. However, whether such self–other differences do in fact extend to other mental states awaits empirical testing.

Finally, the present findings raise an important question; namely, why is it that people tend by default to regard beliefs as quite controllable, shifting away from this only when they consider beliefs concretely? In other words, why would people hold an intuitive theory of belief that diverges from their experience (and possibly the reality) of belief? Precisely where this general theory of belief control comes from is uncertain. One possibility is that the general desirability of having control leads people unreflectively to attribute high control. Another possibility is that people conflate control with the capacity to be rational, such that they intuitively associate the process of weighing reasons to form beliefs with the ability to exert voluntary control over those beliefs. Then, because they assume that people (including themselves) are generally rational when forming beliefs, they conclude that people also have voluntary control over their beliefs. Lastly, perhaps people recognize the phenomenon of motivated reasoning (i.e., they recognize the influence that desires can have over beliefs), and mistake this phenomenon as providing evidence for voluntary control. The present results cannot locate the source of the default assumption that beliefs are controllable, but they suggest that it is somewhat unreflective, and can be unseated once people consider more deeply what is involved in controlling their beliefs.

Conclusion

The present article uncovers an important discrepancy in how people think about their own and others' beliefs; namely, that people judge that others have a greater capacity to voluntarily change their beliefs than they, themselves do. Put succinctly, when someone says, "You can choose to believe in God, or you can choose not to believe in God," they may often mean that *you* can choose but *they* cannot. We have argued that this discrepancy derives from two distinct ways people reason about belief control: either by consulting their default theory of belief, or by introspecting and reporting what they feel when they consider voluntarily changing a belief. When people consider their own belief control, they rely on these two distinct processes depending on the nature of the judgment. When judging their control over beliefs in general, they apply their default theory which leads to relatively high judgments of control. But when judging their control over a specific belief, people introspect on their ability to change that belief, which often reveals obstacles, thereby lowering their ratings of control. In contrast, when people consider others, they simply apply their default theory of belief, regardless of whether they are judging specific beliefs or beliefs in general. Accordingly, in both cases, they judge that others have considerable control over what they believe. Discrepant attributions of control for self and other emerge as a result of these different processes. Such discrepancies may in turn have important downstream effects on people's behavior during disagreements. More work is needed to explore these downstream effects, as well as to understand how much control people actually have over what they believe. Predictably, we find the results from these studies compelling, but we acknowledge that other readers may believe whatever they please.

References

Akhtar, O., & Wheeler, S. C. (2016). Belief in the immutability of attitudes both increases and decreases advocacy. *Journal of Personality and Social Psychology, 111,* 475–492. http://dx.doi.org/10.1037/pspa 0000060

Allison, S. T., Messick, D. M., & Goethals, G. R. (1989). On being better but not smarter than others: The Muhammad Ali effect. *Social Cognition, 7,* 275–295.

Alicke, M. D. (1985). Global self-evaluation as determined by the desirability and controllability of trait adjectives. *Journal of Personality and Social Psychology, 49,* 1621–1630. http://dx.doi.org/10.1037/0022-3514 .49.6.1621

Alicke, M. D. (2000). Culpable control and the psychology of blame. *Psychological Bulletin, 126,* 556–574. http://dx.doi.org/10.1037/0033-2909.126.4.556

Alston, W. P. (1988). The deontological conception of epistemic justification. *Philosophical Perspectives, 2,* 257–299. http://dx.doi.org/10.2307/2214077

Ariely, D., & Loewenstein, G. (2006). The heat of the moment: The effect of sexual arousal on sexual decision making. *Journal of Behavioral Decision Making, 19,* 87–98. http://dx.doi.org/10.1002/bdm.501

Bandura, A. (1977). *Social learning theory.* Englewood Cliffs, NJ: Prentice Hall.

Bar-hillel, M., & Budescu, D. (1995). The elusive wishful thinking effect. *Thinking & Reasoning, 1,* 71–103. http://dx.doi.org/10.1080/135 46789508256906

Baron, J. (2008). *Thinking and deciding* (4th ed.). New York, NY: Cambridge University Press.

Bass, B., & Yammarino, F. (1991). Congruence of self and others' leadership ratings of naval officers for understanding successful performance. *Applied Psychology, 40,* 437–454. http://dx.doi.org/10.1111/j .1464-0597.1991.tb01002.x

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software, 67,* 1–48. http://dx.doi.org/10.18637/jss.v067.i01

Baumeister, R. F., & Newman, L. S. (1994). Self-regulation of cognitive inference and decision processes. *Personality and Social Psychology Bulletin, 20,* 3–19. http://dx.doi.org/10.1177/0146167294201001

Bierbrauer, G. (1979). Why did he do it? Attribution of obedience and the phenomenon of dispositional bias. *European Journal of Social Psychology, 9,* 67–84. http://dx.doi.org/10.1002/ejsp.2420090106

Brehm, J. W. (1966). *A theory of psychological reactance.* New York, NY: Academic Press.

Buehler, R., Griffin, D., & Ross, M. (1994). Exploring the "planning fallacy": Why people underestimate their task completion times. *Journal of Personality and Social Psychology, 67,* 366–381. http://dx.doi.org/10 .1037/0022-3514.67.3.366

Burger, J. M. (1986). Desire for control and the illusion of control: The effects of familiarity and sequence of outcomes. *Journal of Research in Personality, 20,* 66–76. http://dx.doi.org/10.1016/0092-6566(86) 90110-8

Burger, J., & Cooper, H. (1979). The desirability of control. *Motivation and Emotion, 3,* 381–393. http://dx.doi.org/10.1007/BF00994052

Chernyak, N., Kushnir, T., Sullivan, K. M., & Wang, Q. (2013). A comparison of American and Nepalese children's concepts of freedom of choice and social constraint. *Cognitive Science, 37,* 1343–1355. http://dx.doi.org/10.1111/cogs.12046

Clark, C. J., Luguri, J. B., Ditto, P. H., Knobe, J., Shariff, A. F., & Baumeister, R. F. (2014). Free to punish: A motivated account of free will belief. *Journal of Personality and Social Psychology, 106,* 501–513. http://dx.doi.org/10.1037/a0035880

Clifford, W. K. (1877). The ethics of belief. In T. Madigan (Ed.), *The ethics of belief and other essays* (pp. 70–96). Amherst, MA: Prometheus.

Cohen, A. B., & Rozin, P. (2001). Religion and the morality of mentality. *Journal of Personality and Social Psychology, 81,* 697–710. http://dx.doi.org/10.1037/0022-3514.81.4.697

Cusimano, C., & Goodwin, G. P. (2019). Lay beliefs about the controllability of everyday mental states. *Journal of Experimental Psychology: General, 148,* 1701–1732. http://dx.doi.org/10.1037/xge0000547

Davies, M. (2008) *The corpus of contemporary American English (COCA): 560 million words, 1990-present.* Retrieved from https://www.english-corpora.org/coca/

Davison, W. P. (1983). The third-person effect in communication. *Public Opinion Quarterly, 47,* 1. http://dx.doi.org/10.1086/268763

Deci, E. L., & Ryan, R. M. (1985). *Intrinsic motivation and self-determination in human behavior.* New York, NY: Plenum Press. http://dx.doi.org/10.1007/978-1-4899-2271-7

Dersley, I., & Wootton, A. (2000). Complaint sequences within antagonistic argument. *Research on Language and Social Interaction, 33,* 375–406. http://dx.doi.org/10.1207/S15327973RLSI3304_02

Dweck, C. S., & Leggett, E. L. (1988). A social-cognitive approach to motivation and personality. *Psychological Review, 95,* 256–273. http://dx.doi.org/10.1037/0033-295X.95.2.256

Ehrlinger, J., Gilovich, T., & Ross, L. (2005). Peering into the bias blind spot: People's assessments of bias in themselves and others. *Personality and Social Psychology Bulletin, 31,* 680–692. http://dx.doi.org/10.1177/0146167204271570

Epley, N., & Gilovich, T. (2016). The mechanics of motivated reasoning. *The Journal of Economic Perspectives, 30,* 133–140. http://dx.doi.org/10.1257/jep.30.3.133

Epley, N., & Waytz, A. (2009). Mind perception. In S. T. Fiske, D. T. Gilbert, & G. Lindzey (Eds.), *The handbook of social psychology* (5th ed., pp. 498–541). New York, NY: Wiley.

Ford, B. Q., & Gross, J. J. (2019). Why beliefs about emotion matter: An emotion-regulation perspective. *Current Directions in Psychological Science, 28*(1), 74–81. http://dx.doi.org/10.1177/0963721418806697

Gebhardt, W. A., & Brosschot, J. F. (2002). Desirability of control: Psychometric properties and relationships with locus of control, personality, coping, and mental and somatic complaints in three Dutch samples. *European Journal of Personality, 16,* 423–438. http://dx.doi.org/10.1002/per.463

Gilbert, D. T. (1991). How mental systems believe. *American Psychologist, 46,* 107–119. http://dx.doi.org/10.1037/0003-066X.46.2.107

Gilbert, D. T. (1993). The assent of man: Mental representation and the control of belief. In D. M. Wegner & J. W. Pennebaker (Eds.), *Century psychology series. Handbook of mental control* (pp. 57–87). Englewood Cliffs, NJ: Prentice Hall.

Gilbert, D. T., Gill, M. J., & Wilson, T. D. (2002). The future is now: Temporal correction in affective forecasting. *Organizational Behavior and Human Decision Processes, 88,* 430–444. http://dx.doi.org/10.1006/obhd.2001.2982

Gilbert, D. T., & Malone, P. S. (1995). The correspondence bias. *Psychological Bulletin, 117,* 21–38. http://dx.doi.org/10.1037/0033-2909.117.1.21

Gino, F., Sharek, Z., & Moore, D. A. (2011). Keeping the illusion of control under control: Ceilings, floors, and imperfect calibration. *Organizational Behavior and Human Decision Processes, 114,* 104–114. http://dx.doi.org/10.1016/j.obhdp.2010.10.002

Golman, R., Loewenstein, G., Moene, K. O., & Zarri, L. (2016). The preference for belief consonance. *The Journal of Economic Perspectives, 30,* 165–188. http://dx.doi.org/10.1257/jep.30.3.165

Griffin, D. W., & Ross, L. (1991). Subjective construal, social inference, and human misunderstanding. *Advances in Experimental Social Psychology, 24,* 319–359. http://dx.doi.org/10.1016/S0065-2601(08)60333-0

Halekoh, U., & Højsgaard, S. (2014). A kenward-roger approximation and parametric bootstrap methods for tests in linear mixed models–the R package pbkrtest. *Journal of Statistical Software, 59,* 1–30.

Haran, U., Ritov, I., & Mellers, B. A. (2013). The role of actively open-minded thinking in information acquisition, accuracy, and calibration. *Judgment and Decision Making, 8,* 188–201.

Hauser, M., Cushman, F., Young, L., Kang-Xing, J. R., & Mikhail, J. (2007). A dissociation between moral judgments and justifications. *Mind & Language, 22,* 1–21. http://dx.doi.org/10.1111/j.1468-0017.2006.00297.x

Heider, F. (1958). *The psychology of interpersonal relations.* Hoboken, NJ: Wiley. http://dx.doi.org/10.1037/10628-000

Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences, 33,* 61–83. http://dx.doi.org/10.1017/S0140525X0999152X

Iyengar, S. S. S., & Lepper, M. R. (1999). Rethinking the value of choice: A cultural perspective on intrinsic motivation. *Journal of Personality and Social Psychology, 76,* 349–366. http://dx.doi.org/10.1037/0022-3514.76.3.349

Jaeger, J. (2017). r2glmm: Computes R squared for mixed (multilevel) models. R package Version 0.1.2. Retrieved from https://CRAN.R-project.org/package=r2glmm

James, W. (1937). The will to believe. In *The will to believe, and other essays in popular philosophy* (pp. 1–31). London, UK: Longmans, Green and Co.

Johnson, J. T. (1987). The heart on the sleeve and the secret self: Estimations of hidden emotion in self and acquaintances. *Journal of Personality, 55,* 563–582. http://dx.doi.org/10.1111/j.1467-6494.1987.tb00452.x

Jones, E. E., & Harris, V. A. (1967). The attribution of attitudes. *Journal of Experimental Social Psychology, 3,* 1–24. http://dx.doi.org/10.1016/0022-1031(67)90034-0

Jones, E. E., & Nisbett, R. E. (1972). The actor and the observer: Divergent perceptions of the causes of behavior. In E. E. Jones, D. E. Kanouse, H. H. Kelley, R. E. Nisbett, S. Valins, & B. Weiner (Eds.), *Attribution: Perceiving the causes of behavior* (pp. 79–94). Morristown, NJ: General Learning Press.

Kelley, H. H. (1971). *Attribution in social interaction.* Morristown, NJ: General Learning Press.

Kitayama, S., Snibbe, A. C., Markus, H. R., & Suzuki, T. (2004). Is there any "free" choice? Self and dissonance in two cultures. *Psychological Science, 15,* 527–533. http://dx.doi.org/10.1111/j.0956-7976.2004.00714.x

Klein, W. M., & Kunda, Z. (1992). Motivated person perception: Constructing justifications for desired beliefs. *Journal of Experimental Social Psychology, 28,* 145–168. http://dx.doi.org/10.1016/0022-1031(92)90036-J

Koehler, D. J., & Poon, C. S. K. (2006). Self-predictions overweight strength of current intentions. *Journal of Experimental Social Psychology, 42,* 517–524. http://dx.doi.org/10.1016/j.jesp.2005.08.003

Kouchaki, M., Smith, I. H., & Savani, K. (2018). Does deciding among morally relevant options feel like making a choice? How morality constrains people's sense of choice. *Journal of Personality and Social Psychology, 115,* 788–804. http://dx.doi.org/10.1037/pspa0000128

Kruglanski, A. W., & Webster, D. M. (1996). Motivated closing of the mind: "seizing" and "freezing". *Psychological Review, 103,* 263–283. http://dx.doi.org/10.1037/0033-295X.103.2.263

Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin, 108,* 480–498. http://dx.doi.org/10.1037/0033-2909.108.3.480

Langer, E. J. (1975). The illusion of control. *Journal of Personality and Social Psychology, 32,* 311–328. http://dx.doi.org/10.1037/0022-3514.32.2.311

Loewenstein, G. (1996). Out of control: Visceral influences on behavior. *Organizational Behavior and Human Decision Processes, 65,* 272–292. http://dx.doi.org/10.1006/obhd.1996.0028

MacDonald, T., & Ross, M. (1999). Assessing the accuracy of predictions about dating relationships: How and why do lovers' predictions differ from those made by observers? *Personality and Social Psychology Bulletin, 25,* 1417–1429. http://dx.doi.org/10.1177/0146167299259007

Malle, B. F., Guglielmo, S., & Monroe, A. E. (2014). A theory of blame. *Psychological Inquiry, 25,* 147–186. http://dx.doi.org/10.1080/1047840X.2014.877340

Malle, B. F., & Knobe, J. (1997a). The folk concept of intentionality. *Journal of Experimental Social Psychology, 33,* 101–121. http://dx.doi.org/10.1006/jesp.1996.1314

Malle, B. F., & Knobe, J. (1997b). Which behaviors do people explain? A basic actor-observer asymmetry. *Journal of Personality and Social Psychology, 72,* 288–304. http://dx.doi.org/10.1037/0022-3514.72.2.288

Malle, B. F., Knobe, J. M., & Nelson, S. E. (2007). Actor-observer asymmetries in explanations of behavior: New answers to an old question. *Journal of Personality and Social Psychology, 93,* 491–514. http://dx.doi.org/10.1037/0022-3514.93.4.491

Malle, B. F., & Pearce, G. E. (2001). Attention to behavioral events during interaction: Two actor-observer gaps and three attempts to close them. *Journal of Personality and Social Psychology, 81,* 278–294. http://dx.doi.org/10.1037/0022-3514.81.2.278

Markman, K. D., & Tetlock, P. E. (2000). 'I couldn't have known': Accountability, foreseeability and counterfactual denials of responsibility. *British Journal of Social Psychology, 39,* 313–325. http://dx.doi.org/10.1348/014466600164499

Mazzocco, P. J., Alicke, M. D., & Davis, T. L. (2004). On the robustness of outcome bias: No constraint by prior culpability. *Basic and Applied Social Psychology, 26,* 131–146. http://dx.doi.org/10.1080/01973533.2004.9646401

McArdle, M. (2018). Bias against conservatives works like any other prejudice. *Wall Street Journal.* Retrieved from https://www.washingtonpost.com/opinions/bias-against-conservatives-works-like-any-other-prejudice/2018/04/10/17fa1838-3c40-11e8-974f-aacd97698cef_story.html

McFarland, C., & Miller, D. T. (1990). Judgments of self-other similarity. *Personality and Social Psychology Bulletin, 16,* 475–484. http://dx.doi.org/10.1177/0146167290163006

Miller, D. T., & McFarland, C. (1987). Pluralistic ignorance: When similarity is interpreted as dissimilarity. *Journal of Personality and Social Psychology, 53,* 298–305. http://dx.doi.org/10.1037/0022-3514.53.2.298

Miller, D. T., & Norman, S. A. (1975). Actor-observer differences in perceptions of effective control. *Journal of Personality and Social Psychology, 31,* 503–515. http://dx.doi.org/10.1037/h0076485

Miller, J. G., Das, R., & Chakravarthy, S. (2011). Culture and the role of choice in agency. *Journal of Personality and Social Psychology, 101,* 46–61. http://dx.doi.org/10.1037/a0023330

Morris, M. W., & Peng, K. (1994). Culture and cause: American and Chinese attributions for social and physical events. *Journal of Personality and Social Psychology, 67,* 949–971. http://dx.doi.org/10.1037/0022-3514.67.6.949

Nakagawa, S., & Schielzeth, H. (2013). A general and simple method for obtaining R^2 from generalized linear mixed-effects models. *Methods in Ecology and Evolution, 4,* 133–142. http://dx.doi.org/10.1111/j.2041-210x.2012.00261.x

Nelson, T. O., & Narens, L. (1990). Metamemory: A theoretical framework and new findings. In G. H. Bower (Ed.), *Psychology of learning and motivation* (Vol. 26, pp. 125–173). Amsterdam, the Netherlands: Elsevier.

Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review, 84,* 231–259. http://dx.doi.org/10.1037/0033-295X.84.3.231

Nolan, H. (2018). Ideology is a choice. *Splinter.* Retrieved from https://splinternews.com/ideology-is-a-choice-1825172619

Paolacci, G., & Chandler, J. (2014). Inside the turk. *Current Directions in Psychological Science, 23,* 184–188. http://dx.doi.org/10.1177/0963721414531598

Pascal, B. (1852). *Pensées.* Paris, France: Dezobry et E. Magdeleine.

Petty, R. E., & Cacioppo, J. T. (1986). The elaboration likelihood model of persuasion. *Communication and persuasion* (pp. 1–24). New York, NY: Springer.

Pew Research Center. (2015a). *Teen, social media and technology.* Washington, DC: Author.

Pew Research Center. (2015b). *U.S. public becoming less religious.* Washington, DC: Author.

Pew Research Center. (2016a). *The new food fights: U.S. public divides over food science.* Washington, DC: Author.

Pew Research Center. (2016b). *The politics of climate.* Washington, DC: Author.

Pronin, E. (2009). The introspection illusion. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 8, pp. 1–67). San Diego, CA: Elsevier Academic Press. http://dx.doi.org/10.1016/S0065-2601(08)00401-2

Pronin, E., Fleming, J. J., & Steffel, M. (2008). Value revelations: Disclosure is in the eye of the beholder. *Journal of Personality and Social Psychology, 95*(4), 795–809. http://dx.doi.org/10.1037/a0012710

Pronin, E., Gilovich, T., & Ross, L. (2004). Objectivity in the eye of the beholder: Divergent perceptions of bias in self versus others. *Psychological Review, 111,* 781–799. http://dx.doi.org/10.1037/0033-295X.111.3.781

Pronin, E., Kruger, J., Savitsky, K., & Ross, L. (2001). You don't know me, but I know you: The illusion of asymmetric insight. *Journal of Personality and Social Psychology, 81,* 639–656. http://dx.doi.org/10.1037/0022-3514.81.4.639

Pronin, E., & Kugler, M. B. (2010). People believe they have more free will than others. *Proceedings of the National Academy of Sciences of the United States of America, 107,* 22469–22474. http://dx.doi.org/10.1073/pnas.1012046108

Pronin, E., Lin, D. Y., & Ross, L. (2002). The bias blind spot: Perceptions of bias in self versus others. *Personality and Social Psychology Bulletin, 28,* 369–381. http://dx.doi.org/10.1177/0146167202286008

Pronin, E., Olivola, C. Y., & Kennedy, K. A. (2008). Doing unto future selves as you would do unto others: Psychological distance and decision making. *Personality and Social Psychology Bulletin, 34,* 224–236. http://dx.doi.org/10.1177/0146167207310023

Pronin, E., & Ross, L. (2006). Temporal differences in trait self-ascription: When the self is seen as an other. *Journal of Personality and Social Psychology, 90,* 197–209. http://dx.doi.org/10.1037/0022-3514.90.2.197

Risucci, D. A., Tortolani, A. J., & Ward, R. J. (1989). Ratings of surgical residents by self, supervisors and peers. *Surgery, Gynecology & Obstetrics, 169,* 519–526.

Ross, L., & Ward, A. (1996). Naive realism in everyday life: Implications for social conflict and misunderstanding. In E. S. Reed, E. Turiel, & T. Brown (Eds.), *The Jean Piaget symposium series: Values and knowledge* (pp. 103–135). Hillsdale, NJ: Erlbaum.

Rosseel, Y. (2012). lavaan: An R package for structural equation modeling. *Journal of Statistical Software, 48,* 1–36. http://dx.doi.org/10.18637/jss.v048.i02

Savani, K., Markus, H. R., Naidu, N. V., Kumar, S., & Berlia, N. (2010). What counts as a choice? U.S. Americans are more likely than Indians to construe actions as choices. *Psychological Science, 21,* 391–398. http://dx.doi.org/10.1177/0956797609359908

Seligman, M. E. P. (1974). Depression and learned helplessness. In R. J. Friedman & M. M. Katz (Eds.), *The psychology of depression: Contemporary theory and research* (pp. 83–113). Washington, DC: Winston-Wiley.

Seligman, M. E. P. (1975). *Helplessness: On depression, development, and death*. San Francisco, CA: Freeman.

Simonsohn, U., Simmons, J. P., & Nelson, L. D. (2014). Anchoring is not a false-positive: Maniadis, Tufano, and List's (2014) 'Failure-to-Replicate' is actually entirely consistent with the original. *SSRN Electronic Journal*.

Sloman, S. A., Fernbach, P. M., & Hagmayer, Y. (2010). Self-deception requires vagueness. *Cognition, 115,* 268–281. http://dx.doi.org/10.1016/j.cognition.2009.12.017

Spector, P. E., Sanchez, J. I., Siu, O. L., Salgado, J., & Ma, J. (2004). Eastern versus western control beliefs at work: An investigation of secondary control, socioinstrumental control, and work locus of control in china and the U.S. *Applied Psychology, 53,* 38–60. http://dx.doi.org/10.1111/j.1464-0597.2004.00160.x

Ståhl, T., Zaal, M. P., & Skitka, L. J. (2016). Moralized rationality: Relying on logic and evidence in the formation and evaluation of belief can be seen as a moral issue. *PLoS ONE, 11,* e0166332. http://dx.doi.org/10.1371/journal.pone.0166332

Stanovich, K. E., & West, R. F. (1997). Reasoning independently of prior belief and individual differences in actively open-minded thinking. *Journal of Educational Psychology, 89,* 342–357. http://dx.doi.org/10.1037/0022-0663.89.2.342

Steup, M. (2017). Believing intentionally. *Synthese, 194,* 2673–2694. http://dx.doi.org/10.1007/s11229-015-0780-7

Storms, M. D. (1973). Videotape and the attribution process: Reversing actors' and observers' points of view. *Journal of Personality and Social Psychology, 27,* 165–175. http://dx.doi.org/10.1037/h0034782

Tamir, M., John, O. P., Srivastava, S., & Gross, J. J. (2007). Implicit theories of emotion: Affective and social outcomes across a major life transition. *Journal of Personality and Social Psychology, 92,* 731–744. http://dx.doi.org/10.1037/0022-3514.92.4.731

Turri, J., Rose, D., & Buckwalter, W. (2018). Choosing and refusing: Doxastic voluntarism and folk psychology. *Philosophical Studies, 175,* 2507–2537. http://dx.doi.org/10.1007/s11098-017-0970-x

Van Boven, L., & Loewenstein, G. (2003). Social projection of transient drive states. *Personality and Social Psychology Bulletin, 29,* 1159–1168. http://dx.doi.org/10.1177/0146167203254597

Webster, D. M., & Kruglanski, A. W. (1994). Individual differences in need for cognitive closure. *Journal of Personality and Social Psychology, 67,* 1049–1062. http://dx.doi.org/10.1037/0022-3514.67.6.1049

Wegner, D. M. (1994). Ironic processes of mental control. *Psychological Review, 101,* 34–52. http://dx.doi.org/10.1037/0033-295X.101.1.34

Weiner, B. (1995). *Judgments of responsibility: A foundation for a theory of social conduct*. New York, NY: Guilford Press.

White, R. W. (1959). Motivation reconsidered: The concept of competence. *Psychological Review, 66,* 297–333. http://dx.doi.org/10.1037/h0040934

Wilson, T. D., & Brekke, N. (1994). Mental contamination and mental correction: Unwanted influences on judgments and evaluations. *Psychological Bulletin, 116,* 117–142. http://dx.doi.org/10.1037/0033-2909.116.1.117

Wood, T., & Porter, E. (2019). The elusive backfire effect: Mass attitudes' steadfast factual adherence. *Political Behavior, 41,* 135–163.

Wortman, C. B., & Brehm, J. W. (1975). Responses to uncontrollable outcomes: An integration of reactance theory and the learned helplessness model. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 8, pp. 277–336). San Diego, CA: Academic Press. http://dx.doi.org/10.1016/S0065-2601(08)60253-1

Yzerbyt, V., Muller, D., Batailler, C., & Judd, C. M. (2018). New recommendations for testing indirect effects in mediational models: The need to report and test component paths. *Journal of Personality and Social Psychology, 115,* 929–943. http://dx.doi.org/10.1037/pspa0000132

Zajonc, R. B. (1980). Feeling and thinking: Preferences need no inferences. *American Psychologist, 35,* 151–175. http://dx.doi.org/10.1037/0003-066X.35.2.151

(*Appendices follow*)

Appendix A

Belief Statements Used in Study 1

Which of the following statements best reflects your views on the use of genetic modification in food production?

I believe that genetically modified foods should be prohibited.

I believe that genetically modified foods should not be prohibited.

Which of the following statements best reflects your views on how to best combat global climate change?

I believe that global climate change is a problem that is best addressed through strong government regulation.

I believe that global climate change is a problem that is not best addressed through strong government regulation.

Which of the following statements best reflects your views on the existence of God?

I believe that God exists.

I believe that God does not exist.

Which of the following statements best reflects your views on the impact that social media has had on dating?

I believe that social media has had an overall negative effect on dating.

I believe that social media has not had an overall negative effect on dating.

Appendix B

Belief Statements Used in Study 2

Genetically modified foods should be prohibited.

Global climate change is a problem that is best addressed through strong government regulation.

God exists.

Social media has had an overall negative effect on dating.

Appendix C

Full Text of Scenarios From Posttests Associated With Study 3

Impression scale:

Extremely negative			Not negative or positive			Extremely positive
−3	−2	−1	0	1	2	3

Text From Posttest 1

Suppose you are walking down the street. Ahead of you are two people, **Jones** and **Smith**. Jones is walking about 20 paces ahead of Smith. They do not know each other.

Jones accidentally drops a $20 bill on the ground when he pulls his hand out of his pocket. He does not notice that he did this.

When Smith reaches the $20, he picks it up and puts it in his pocket and then changes direction, clearly intending to keep the money.

When you see this, you point it out to the stranger standing next to you, **Peters**. However, Peters says he doesn't think it was wrong what Smith did.

How much do these events impact your impression of each of these people?

My impression of Jones is. . . .

My impression of Smith is. . . .

My impression of Peters is. . . .

(*Appendices continue*)

Text From Posttest 2

Two people, **Jones** and **Smith**, are walking down the street. Jones is walking about 20 paces ahead of Smith. Across the street is **Peters**, who can clearly see Jones and Smith. None of them know each other.

Jones accidentally drops a $20 bill on the ground when he pulls his hand out of his pocket. He does not notice that he did this.

When Smith reaches the $20, he picks it up, puts it in his pocket, and then changes direction, clearly intending to keep the money.

Across the street, Peters sees Smith keep the $20 instead of returning it and thinks to himself "What that guy did is not wrong."

How much do these events impact your impression of each of these people?

My impression of Jones is. . . .

My impression of Smith is. . . .

My impression of Peters is. . . .

Text (and Questions) From Posttest 3

Suppose you sitting down to eat with some people and you are placed at a table with a clear view of the sidewalk.

You look down the street and see two people, Jones and Smith. Jones is walking about 20 paces ahead of Smith. They do not know each other.

Jones accidentally drops a $20 bill on the ground when he pulls his hand out of his pocket. He does not notice that he did this.

When Smith reaches the $20 bill, he picks it up and puts it in his pocket and then changes direction, clearly intending to keep the money.

You point out what you saw to the people at your table and get a few different responses.

For each person, please indicate how much their reaction to this event influences your impression of them. Assume that everyone is answering completely independently of one another—no one is giving a response based on what someone else said.

[*The four conversations below with Peters, Johnson, Smith, and Arnolds were presented in a random order.*]

You ask someone at the table, Peters, whether he could do what Smith did. Peters says, "I would not keep the money. That said, I think I could choose to keep the money if I wanted to." How would this response influence your impression of Peters?

You ask someone at the table, Johnson, whether he could do what Smith did. Johnson says, "I would not keep the money. That said, I do not think I could choose to keep the money even if I wanted to." How would this response influence your impression of Johnson?

You ask someone at the table, Arnolds, if he could think that keeping the $20 is okay. Arnolds says, "I believe what Smith did was wrong. That said, if I wanted to, I think I could choose to believe that what Smith did was not wrong." How would this response influence your impression of Arnolds?

You ask someone at the table, Brown, if he could think that keeping the $20 is okay. Brown says, "I believe what Smith did was wrong. That said, if even if I wanted to, I do not think I could choose to believe that what Smith did was not wrong." How would this response influence your impression of Brown?

Appendix D

Passages Used in the "Beliefs in General" Condition in Study 6

My current beliefs are ones that I voluntarily hold. Specifically, I could change what I believe if I wanted to even if this means I was being wrong or immoral by doing so.

My current job is one that I voluntarily hold. Specifically, I could change where I work if I wanted to even if this meant having a worse job.

My current habits are ones that I voluntarily engage in. Specifically, I could change my routines if I wanted to even if I had been doing them for years.

My current home is one that I voluntarily live in. Specifically, I could change where I live if I wanted to even if it meant changing many other parts of my life.

My behavior is voluntary. Specifically, when I make decisions to act a certain way, it is because I wanted to, and I could always act some other way if I wanted to even if this meant I was being immoral.

(*Appendix follows*)

Appendix E

Passages Used in the General Condition in Study 7

My current beliefs are ones that I voluntarily hold. Specifically, I could choose to hold different beliefs if I wanted to even if this meant being wrong or immoral.

My current job is one that I voluntarily hold. Specifically, I could change where I work if I wanted to even if this meant having a worse job.

My current home is one that I voluntarily live in. Specifically, I could change where I live if I wanted to even if it meant changing many other parts of my life.

My behavior is voluntary. Specifically, when I make decisions to act a certain way, it is because I want to, and I could act some other way even if this meant I was being immoral.