

Memory and Belief in the Transmission of Counterintuitive Content

Aiyana K. Willard¹ · Joseph Henrich² ·
Ara Norenzayan³

© Springer Science+Business Media New York 2016

Abstract Cognitive scientists have increasingly turned to cultural transmission to explain the widespread nature of religion. One key hypothesis focuses on memory, proposing that that minimally counterintuitive (MCI) content facilitates the transmission of supernatural beliefs. We propose two caveats to this hypothesis. (1) Memory effects decrease as MCI concepts become commonly used, and (2) people do not believe counterintuitive content readily; therefore additional mechanisms are required to get from memory to belief. In experiments 1–3 ($n=283$), we examined the relationship between MCI, belief, and memory. We found that increased tendencies to anthropomorphize predicted poorer memory for anthropomorphic-MCI content. MCI content was found less believable than intuitive content, suggesting different mechanisms are required to explain belief. In experiment 4 ($n=70$), we examined the non-content-based cultural learning mechanism of credibility-enhancing displays (CREDS) and found that it increased participants' belief in MCI content, suggesting this type of learning can better explain the transmission of belief.

Keywords MCI theory · Memory · Religion · Anthropomorphism · Credibility-enhancing displays (CREDS) · Cultural learning

Electronic supplementary material The online version of this article (doi:10.1007/s12110-016-9259-6) contains supplementary material, which is available to authorized users.

✉ Aiyana K. Willard
aiyana@psych.ubc.ca

¹ Department of Psychology, University of Texas at Austin, 116 Inner Campus Dr Stop G6000, Austin, TX 78712-0187, USA

² Department of Human Evolutionary Biology, Harvard University, Cambridge, MA, USA

³ Department of Psychology, University of British Columbia, Vancouver, BC, Canada

One of the more prominent and influential hypotheses in the cognitive science of religion is minimally counterintuitive (MCI) content (see Atran & Norenzayan 2004a; Barrett 2004; Boyer 2001; Boyer 2003). This hypothesis suggests that religious representations are widespread because they are uniquely adapted to exploit the structures of our minds by violating reliably developing intuitive understandings of the world (Barrett & Nyhof 2001; Boyer & Ramble 2001). More specifically, supernatural content violates our core ontological assumptions of the world. It is therefore distinctive and more memorable than non-supernatural content. This memory bias makes religious ideas “stick” and remain in human minds across generations. In this way, MCI theory is a parsimonious and wide-reaching account of the cultural contagion of religious beliefs.

Recent developments in the cognitive science of religion have called this hypothesis into question (Gervais et al. 2011; Henrich 2009; Purzycki & Willard 2015; Russell & Gobet 2013). One of the early objections is termed the *Mickey Mouse problem* (Atran 2002; Atran & Norenzayan 2004a). Simply put, if minimal counterintuitiveness is all that is required for religious concepts to take root, then why is Mickey Mouse, an MCI concept, not treated as a god? This problem has been broadly reframed as the *Zeus problem*: why don't people readily believe in, and worship, other people's gods as much as they do their own gods? Why don't Christians believe in Zeus (and start worshipping him), if he has similar MCI properties to the Christian God (Gervais & Henrich 2010)? The exact same mental representation of a supernatural agent can ignite faith in some people but only incredulous mirth in others. The MCI hypothesis is silent about why some MCI content becomes the center of a powerful belief system, whereas others do not.

Other criticisms have come in the form of the cognitive structures required for MCI theory to create a unique and persistent memory effect (see Purzycki & Willard 2015). MCI violations are violations of core ontological categories, or innate intuitions about how the mind works. This view has been countered with the suggestion that many MCI violations are schematic rather than ontological in nature (Purzycki 2010; Upal 2010). Simply put, this would place MCI violations in the “things we find surprising” category more generally, and not as specific and persistent violations of ontological assumptions. This would mean that there is no difference between how we process information about something unusual, like a pink elephant, and something impossible, like an invisible elephant, other than the degree of surprise we feel. Further, this would suggest that once MCI ideas are accepted and commonplace, new schemas would be created that incorporate these ideas. New ideas that are consistent with these schemas will not be violating any of our expectations and should therefore not show a bias in memory. This implies that while counterintuitiveness might account for the spread of certain mental representations, it cannot explain the persistence of supernatural or other religious beliefs in particular. We will call this perspective the schema or schematic view of MCI content.

Here, we underline these concerns and provide a partial solution (see Atran & Henrich 2010; Henrich 2009). We wish to address the issue of belief and distinctiveness in three hypotheses:

1. We hypothesize that MCI content is not readily believed, and thus it demands additional mechanisms to prompt belief. Minimal counterintuitiveness may, under

- some conditions, make a concept more memorable, but it consistently makes it less believable.
2. We further hypothesize that people who more readily engage in the MCI-like thought processes of anthropomorphism should have an attenuated memory effect for this type of content. Frequent use should change schemas. This would be evidence that MCI violations are schematic in nature and that these types of beliefs can be updated with new experiences and expectations of the world.
 3. We hypothesize that non-content-based cultural learning mechanisms can play a role in overcoming the problem of belief for MCI content. We propose that exposure to credibility-enhancing displays (Henrich 2009) can instill commitment to MCI content despite its inherent unbelievability.

What Is MCI?

The MCI hypothesis has played an important role in the cognitive science of religion (Atran & Norenzayan 2004a; Barrett 2004; Barrett 2008b; Boyer 1994; Boyer 2001; Boyer 2003). The hypothesis rests on the idea that human minds possess innate or reliably developing ontological categories governed by core knowledge domains, such as folk physics, folk biology and folk psychology (see Spelke & Kinzler 2007; Wellman & Gelman 1992). A concept becomes minimally counterintuitive when an implicit feature that pertains to one of these categories is mildly but systematically violated (e.g., an object that can pass through walls), or an implicit feature gets transferred to an object or animal to which it should not apply (e.g., a statue that knows your thoughts). These violations are attention-arresting and inferentially rich and therefore encourage further cognitive processing that facilitates their cognitive stabilization and cultural transmission. However, they are not radical departures from common sense and do not rupture completely ordinary expectations that would otherwise render them incomprehensible (Atran & Norenzayan 2004a; Sperber 1996).

What differentiates MCI concepts from ones that are simply bizarre (see Lang 1995; McDaniel et al. 1995; Riefer & Lamay 1998) is that the violations are about *core ontological categories* such as physical object, living form, or person—not just our learned expectations of the world (Barrett 2008a). It has been argued that violations of these knowledge domains are surprising to us in a more fundamental way (Barrett & Nyhof 2001), and we pay more attention to these violations and remember them better than we do ideas consistent with expectations. The MCI hypothesis asserts that this increased memorability allows MCI concepts to spread at a greater rate than intuitive concepts or maximally counterintuitive concepts, leading to a world full of myths and stories about the supernatural.

Previous Research on the MCI Effect

Two foundational studies established this MCI effect (Barrett & Nyhof 2001; Boyer & Ramble 2001). Both of these studies used stories that contained sets of intuitive and MCI items, and found that people remembered more of the MCI items than the intuitive ones. Norenzayan and colleagues (2006, study 2) analyzed Grimms' fairytales and

found that stories with two or three MCI violations (but not more) were the most well-known. Another showed that the Roman prodigies with minimally counterintuitive content,¹ collected between 218 and 44 BCE, are more likely to be recorded than ones with only bizarre or intuitive content (Lisdorf 2004).

Other research has looked at the narrative context in which MCI content is presented (Gonce et al. 2006; Upal 2010; Upal et al. 2007). Their findings suggest that the memory bias for MCI content is only apparent when the content is not expected. When stories are prefaced with “In a dream” or presented as science fiction, the memory bias disappears. This is a potential problem for myths and fairytales, where MCI content should be expected (Gonce et al. 2006), but it also adds support to a schema view of the MCI memory effect (Upal 2010). If expecting MCI content means that the content is no longer surprising and requiring of special attention, then the possibility that MCI content is driving the spread of these types of cultural information is called into question.

Distinctiveness and MCI Theory

The first time a person hears that trees can understand and remember their conversations, clearly they will find this statement distinctive and potentially memorable. Still, what happens when that person has become well acquainted with this concept? How would they see it if they had been raised since birth with this as a belief? Humans are amazingly adaptable to novel concepts and artifacts. Magnets and helium balloons are “counterintuitive” items we see as commonplace. This problem of declining distinctiveness rests on the tension between whether these concepts exist as unchangeable core concepts or whether they are part of updatable schemas of the world. This tension has been addressed in greater detail elsewhere (see Purzycki & Willard 2015).

For the MCI hypothesis to explain the persistence of many supernatural representations, it has to go beyond a description of what happens when people come in contact with an unusual idea and account for the persistence of these beliefs over time. A key question, therefore, is whether the memory bias remains when these ideas become more schema-consistent, and therefore unlikely to violate expectations of the world.

The Problem of Belief

The fact that people can become deeply and powerfully committed to religious concepts is part of what makes them so intriguing. Though it is clear that not all supernatural ideas are believed in (consider fairytales), and not all religious traditions hold belief as the main defining component of religious devotion (Cohen et al. 2003; Purzycki & Sosis 2011; Shariff et al. 2014), belief still remains a powerful motivating factor for the way people relate to some core supernatural ideas. The very fact that these concepts are “counterintuitive” should suggest that they are less believable than our intuitive perceptions of the world. Nevertheless, people are willing to offer time and resources to, and possibly even die for, beliefs that are inherently unverifiable and contradictory

¹ Prodigies are anecdotes or microstories about transgressions in the natural order of the cosmos. Prodigies were reported to the Senate, who decided which were genuine and should be recorded.

to the way we intuitively see the material world. This puzzling finding brings us to a fundamental question in the relationship between MCI concepts and the supernatural. How do people come to believe in and commit to these unintuitive ideas?

MCI content should be, on average, less believable than intuitive content *because* MCI content violates our intuitions about the world. These intuitions exist based on our experience of the world, either in our lifetimes or as they have been encoded into our genes through generations, and they exist because of the benefit they give in accurately predicting the world.² From an evolutionary point of view, it is expected that people are skeptical of claims that defy experience or expectations (Bergstrom et al. 2006; Harris & Koenig 2006).

Based on these ideas we hypothesize here that MCI content is *less* believable than intuitive content and additional mechanisms are required to make MCI content the center of genuine belief. Still, this question of belief has yet to be empirically explored.

Cultural Learning and Credibility-Enhancing Displays (CREDS)

Thus far, we have articulated two ideas about MCI content and belief: (1) when MCI content involves ideas we call “religious,” it is not merely represented and remembered, it contains commitment or belief; (2) MCI content, as a violation of our intuitions about the world, should be particularly unbelievable—otherwise we would lose much of the benefit of having these intuitions in the first place. The second observation might seem self-evident—after all, counterintuitiveness suggests the absence of believability. Yet, to the extent that this is empirically true, we are left with an apparent paradox: if MCI content is memorable, but less believable, how do we explain both its prevalence in religions and the deep faith they can inspire? One answer, we suggest, might be found in cultural learning mechanisms that can facilitate the spread of beliefs even if they are intuitively implausible.

Several cultural learning mechanisms are likely to be involved in the transmission of religious beliefs. Henrich (Henrich 2009) has proposed a mechanism that is relevant to difficult-to-verify beliefs: credibility-enhancing displays (CRED). The CRED hypothesis suggests that we learn beliefs from others to the extent that these beliefs are backed up by credible displays. When someone behaves in a way that is credible and consistent with their beliefs, such as by giving tithes to the church they claim to believe in, or eating a mushroom they claim is not poisonous, their associated beliefs become more plausible and more likely to be acquired by observers. Their actions underwrite their words.

CREDS are actions that a person would be unlikely to engage in unless they truly believed their expressed conviction. This cultural learning bias helps learners navigate a world in which others should want to convince them of untrue things for selfish reasons (e.g., the advertising industry, lawyers, and politicians). The classical example for religion would be the pairing of verbal statements about beliefs in life after death and the actions of martyrdom. Martyrdom is powerful as a transmitter of the faith because it

² Elsewhere we have hypothesized that most religious concepts, and concepts that are labeled as MCI, are actually supported by core intuitions about the world rather than being violations of those intuitions (Purzycki & Willard 2015; Willard & Norenzayan 2013).

shows that—if nothing else—the martyr must have unreservedly held the belief he or she verbally expressed.

Overview of Experiments

We conducted four experiments to explore these ideas. Experiments 1–3 establish that counterintuitive representations are less believable than intuitive representations, even when they are more memorable. We also tested whether individual differences in trait anthropomorphism—that is, a chronic tendency to apply human characteristics widely to nonhuman things (Waytz et al. 2010a)—moderates the memory effect for similar content. We hypothesized that high anthropomorphizers should show less memory bias for anthropomorphic content than low anthropomorphizers simply because they find anthropomorphic content less distinctive. Finally, experiment 4 tests the hypothesis that MCI content can become believable through cultural learning. We assess if witnessing a CRED would enhance people’s willingness to endorse or believe in science news stories independent of MCI content.

Experiments 1–3

In the first three experiments, we address the question of believability of MCI content across three different types of stimuli. This was done by comparing ratings of believability for intuitive and counterintuitive content. We also assessed the impact of individual differences in the tendency to use one particular type of MCI content—anthropomorphic content—on the memory bias for MCI content. Anthropomorphism reliably varies across individuals and affects how people interact in the world in predictable ways (Epley et al. 2008a; Epley et al. 2008b; Waytz et al. 2010a; Waytz et al. 2010b), and this variation allows us to test the hypothesis that MCI content that is more expected is less memorable. We used a validated scale of anthropomorphism (the IDAQ; Waytz et al. 2010a) to assess the willingness to attribute human-like states to animals, nature, and technology.

Methods

Participants All participants were recruited from a university human subject pool and given course credit for their participation (Table 1).

Materials Each of the three experiments used a different set of stimuli to assess the memorability and believability of MCI versus intuitive content. They are as follows:

Exp. 1: Folktales Three different “folktales” were constructed by the experimenters (see *ESM*). Each story had three different versions pertaining to each of the categories (anthropomorphic-MCI, physical-MCI, social-intuitive), for a total of nine stories, all approximately 200 words in length. The stories were constructed to be easily divisible into individual “concepts” for coding purposes and so the content type could be

Table 1 Participants' demographics, experiments 1–4

	Exp. 1	Exp. 2	Exp. 3	Exp. 4
<i>N</i>	95	89	99	77
Female	78 (82%)	63 (71.6%)	68 (68.7%)	41 (53%)
Age (<i>sd</i>)	19.25 (1.74)	19.9 (2.6)	20.16 (3.10)	20.38 (1.88)
Ethnicity				
Asian	64 (67.4%)	55 (61.9%)	67 (67.7%)	40 (51.9%)
Caucasian	27 (28.4%)	31 (34.5%)	23 (23.2%)	34 (44.2%)
Other	4 (4.2%)	3 (3.6%)	9 (9.1%)	3 (3.9%)
Religious Affiliation				
Christian	—	41 (46.5%)	41 (41.4%)	34 (44.2%)
Other Religious	—	13 (15.1%)	16 (16.1%)	9 (11.6%)
Non-Religious	—	34 (38.4%)	42 (42.5%)	34 (44.2%)

changed without changing anything else in the story. For the most part, these concepts were single sentences (e.g., “Jane went for a walk in the woods.” Participants would be marked as remembering this “concept” if they said anything about Jane walking in the woods). If sentences contained more than one concept, they were divided up accordingly. This was done to control for anything idiosyncratic about the stories.

Exp. 2: Museum Story This stimuli was taken from Boyer and Ramble (Boyer & Ramble 2001). It consists of a short introduction about Mr. Wurg and his trip to the natural history museum on the planet Zeon 3. The story lists 24 items: 12 objects and 12 people. Of each of these sets of 12 items, 6 were minimally counterintuitive and 6 were entirely intuitive. All the MCI objects use some type of animacy and/or a mental trait as their MCI violations rather than the physical violations used in the previous experiment and should be considered anthropomorphic in nature (e.g., “Objects that hide when they are scared”; “Objects that can notice people are staring at them”).

Exp. 3: Science News Stories This stimuli consisted of 12 single-paragraph-long summaries of news stories. Six of the news stories were intuitive and six had a minimally counterintuitive violation. Of the minimally counterintuitive stories, three were real news stories, and three were created by the experimenter (see [ESM](#)). Science news was selected as something that could be presented as new information that is plausible. This allowed us to assess whether believability findings from the first two experiments were due to the clear fictional nature of the stimuli.

After reading all the stimuli, participants were asked to (1) recall them and (2) rate how believable they thought each story was. This was done with a one-item measure in experiment 1 (“I believe this story could be true”) and a three-item measure for experiments 2 and 3 (“I believe this item is real”; “It’s plausible this item could exist”; “I doubt this item is true”; $\alpha=0.84$). This change was made to make the believability score more reliable and less prone to error (see Traub & Rowley 1991). Participants in

experiment 1 also rated how much they liked each story on a seven-point Likert scale. The IDAQ was used to assess each participant's tendency to anthropomorphize (Waytz et al. 2010a).

Design and Procedures For all experiments, participants came into the lab and were instructed to complete the experiment alone in a cubicle. They read stimuli displayed on a computer screen and were told that they would be asked about it again later, but they were not informed that they would be recalling it. The order of stories/items was randomized across participants. Participants were given 5 min to complete a distractor math question and then asked to recall what they had read by typing it out on the computer. For experiment 1, story recall was prompted by the main character in each story (i.e., "Recall the story about Jane") and they were asked to recall the stories as though they were writing them out for another person to read. For experiment 2 recall, participants were asked to recall as many objects and people as possible, and experiment 3 participants were asked to recall in a single sentence the main topic of as many stories as they could. After recall each story/item was displayed and participants rated how much they believed each story/item. They then completed the anthropomorphism questionnaire.

In experiments 2 and 3, participants were asked to come back for a second session, one week later. This was done to assess change through time in memory and belief. Participants were asked to recall the items from the stories they had read the previous week and then to rate belief again. This was done to see if either memory or belief changed with a short time interval (see Norenzayan et al. 2006). Finally, in experiment 3 a behavioral measure was added. After evaluating belief and memory in the second session, participants were told that some of the news stories were fake, and that we would give them \$1 US for each story they could correctly identify as real (or not real). This was done to raise the stakes of reporting belief and to move beyond simple self-reported measures of belief.

Results

Experiment 1 Stories were divided into sets of single concepts with an associated concept type (intuitive, anthropomorphic-MCI, physical-MCI), and each concept was coded as present or absent. Two independent coders coded recall ($\kappa=0.93$). The total number of recalled concepts was divided by the total number of concepts and reported as a proportion.³ We found no significant difference in the memorability of our three story types, ($F_{2, 281}=1.96, p=0.14$), or in how much participants liked each story ($F_{2, 278}=2.49, p=0.09$). There was a significant difference in the believability of each category ($F_{1.46, 205.13}=68.69, p<0.001$; using Greenhouse and Geisser (Greenhouse & Geisser

³ Since each story could be intuitive, anthropomorphic-MCI, or physical-MCI, and story type was randomized, the total number of concepts did not differ across content type even though it did differ by story.

1959) correction for sphericity; see Table 2). Using planned comparisons, we found that anthropomorphic-MCI and physical-MCI were believed less than intuitive stories ($F_{1, 282} = 136.52, p < 0.001$). No difference was found in believability between anthropomorphic and MCI stories ($F_{1, 282} = 0.09, p = 0.76$).

This analysis was confirmed using a multilevel regression with random intercepts and fixed slopes for each individual. We show that a categorical change from the combined category of anthropomorphic-MCI, physical-MCI is associated with an average change of -1.75 (95%CI: -2.10 to -1.40 ; $t_{90.34} = 10.14, p < 0.001$; $d = -2.49$) points on a seven-point belief scale when compared with intuitive items, controlling for age, gender, ethnicity, and order. A second multilevel regression using the same control variables showed no significant effect of the combined anthropomorphic-MCI, physical-MCI content on memory (-0.04 ; 95%CI: -0.08 to 0.01 ; $t_{92.77} = -1.44, p = 0.15$).

We ran an additional set of OLS regressions to examine the relationship between anthropomorphic tendencies and memory (Table 3). Our results show a significant negative effect of IDAQ on memory for anthropomorphic-MCI content and a marginal effect for physical-MCI content. We also found a negative effect of belief on memory for both anthropomorphic-MCI and physical-MCI content. No effects were found for intuitive content.

Experiment 2 Recall was recorded as number of individual items recalled in each category, using the same methods as Boyer and Ramble (Boyer & Ramble 2001). Two independent coders coded recall ($\kappa = 0.91$). MCI items were significantly more memorable after a 5-min delay ($M = 3.34, sd = 1.90$) than the intuitive items ($M = 2.34, sd = 1.83$; $t_{88} = 5.35, p < 0.001, 95\%CI: 0.63$ to 1.37 ; $d = 0.57$). MCI items were significantly less believable ($M = 3.47, sd = 1.27$) than the intuitive items ($M = 6.59, sd = 0.61$; $t_{88} = -20.51, p < 0.001, 95\%CI: -3.43$ to -2.82 ; $d = 2.17$).

After a one-week delay, we once again found that MCI items were significantly more memorable ($M = 2.20, sd = 1.81$) than the intuitive items ($M = 1.61, sd = 1.81$; $t_{82} = 3.10, p = 0.002, 95\%CI: 0.21$ to 0.97 ; $d = 0.34$) and that MCI items were significantly less believable ($M = 3.17, sd = 1.27$) than intuitive items ($M = 6.61, sd = 0.60$; $t_{81} = -20.80, p < 0.001, 95\%CI: -3.77$ to -3.11 ; $d = 2.29$).⁴ There was no change in believability from time 1 to time 2 for either MCI ($t_{78} = 0.75, p = 0.46$) or intuitive ($t_{78} = 1.41, p = 0.16$) items.

To assess the effects of belief and anthropomorphism on memory, we again conducted a series of OLS regressions (Tables 4). Religious affiliation was added as an additional control. (Information on affiliation was not collected in experiment 1.) All of the results are robust with regard to the inclusion or

⁴ Using a multilevel logistic regression with random intercepts for each individual we can show that these effects do not change when we control for age, gender, ethnicity, and religious affiliation. MCI items are significantly more memorable than intuitive items after 5 min ($z = 4.98, p < 0.001$; Odds: 1.77, 95%CI: 1.41 to 2.21) and one week ($z = 3.51, p < 0.001$; Odds: 1.66, 95%CI: 1.25 to 2.21). Belief was similarly modeled. Belief changed an average of -3.08 (95% CI -3.39 to -2.77 ; $t_{81.00} = -19.46, p < 0.001, d = -9.26$) points on a seven-point scale between intuitive items and MCI items after 5 min, and -3.46 (95%CI: -3.60 to -3.32 ; $t_{165.40} = -49.53, p < 0.001$) after one week.

Table 2 Means (and *sd*) of dependent variables for each story type

	Memory	Liking	Belief
Social-Intuitive	0.41 (0.21)	3.59 (1.51)	3.32 (1.62)
Anthropomorphic	0.38 (0.19)	3.96 (1.52)	1.57 (0.92)
MCI (physical)	0.36 (0.16)	3.82 (1.55)	1.52 (0.95)

exclusion of any of these covariates. Our analysis showed a significant negative effect of anthropomorphism on both MCI objects and people after the 5 min delay. There was also a negative effect of anthropomorphism on memory for intuitive people items. Recall at one week showed a marginal negative effect of anthropomorphism on MCI objects, but no other effects. The time 2 effects use the anthropomorphism measure taken at time 1. We did not replicate the effect of belief reducing memory found in experiment 1.

Experiment 3 Recall was coded as a yes or no for each story. Participants were coded “yes” if they remembered the main point of the news story (e.g., “A story about teleporting objects,” “A story about windows that produce solar power”). Memory was coded by two independent coders ($\kappa=0.79$). Mean differences in memory between MCI stories ($M=3.05$, $sd=1.38$) and intuitive stories ($M=2.85$, $sd=1.19$) were not significant at conventional levels after a 5-min delay ($t_{95}=1.35$, $p=0.18$, 95%CI: -0.09 to 0.49 ; $d=0.19$). Memorability was marginally different between MCI ($M=2.44$, $sd=1.48$) and intuitive stories ($M=2.12$, $sd=1.23$) after 1 week ($t_{83}=1.87$, $p=0.07$, 95%CI: -0.02 to 0.66 ; $d=0.27$). Nevertheless, a multilevel model approach with random intercepts for

Table 3 OLS regressions examining the effects of anthropomorphism and belief on story memorability, experiment 1

	<i>B</i> (<i>SE</i>)	β	95%CI	
			Lower	Upper
Anthropomorphic-MCI				
Belief	-0.05 (0.02)*	-0.22	-0.46	-0.06
IDAQ	-0.03 (0.01)*	-0.25	-0.45	-0.05
Physical-MCI				
Belief	-0.04 (0.02)*	-0.25	-0.47	-0.04
IDAQ	-0.02 (0.01)†	-0.19	-0.41	0.02
Social-Intuitive				
Belief	0.00 (0.01)	0.00	-0.22	0.22
IDAQ	-0.01 (0.02)	-0.09	-0.31	0.14

Controls: age, gender, ethnicity, and order effects

† $p < 0.10$, * $p < 0.05$

Table 4 Effects of anthropomorphism and belief on story memorability, experiment 2

	<i>B</i> (<i>SE</i>)	β	95%CI	
			Lower	Upper
5 min				
MCI objects				
Belief	0.19 (0.09)*	0.24	0.01	0.47
IDAQ	-0.23 (0.09)**	-0.31	-0.54	-0.08
MCI people				
Belief	0.03 (0.09)	0.04	-0.19	0.27
IDAQ	-0.18 (0.08)*	-0.26	-0.49	-0.03
Intuitive objects				
Belief	0.01 (0.26)	0.01	-0.27	0.28
IDAQ	-0.10 (0.08)	-0.15	-0.39	0.09
Intuitive people				
Belief	0.15 (0.21)	0.09	-0.15	0.32
IDAQ	-0.19 (0.08)*	-0.26	-0.48	-0.05
1 week				
MCI objects				
Belief	0.06 (0.09)	0.08	-0.17	0.33
IDAQ	-0.14 (0.08)†	-0.23	-0.48	0.03
MCI people				
Belief	-0.03 (0.10)	-0.04	-0.31	0.23
IDAQ	0.03 (0.08)	0.04	-0.22	0.31
Intuitive objects				
Belief	0.23 (0.23)	0.13	-0.14	0.40
IDAQ	-0.03 (0.08)	-0.13	-0.29	0.21
Intuitive people				
Belief	0.12 (0.23)	0.07	-0.22	0.36
IDAQ	-0.06 (0.07)	-0.11	-0.36	0.14

Controls: age, gender, ethnicity, and religious affiliation

† $p < 0.10$, * $p < 0.05$, ** $p < 0.01$

each individual and controlling for age, gender, ethnicity, religious affiliation, and fake stories estimates that MCI content (relative to intuitive content) increases the odds of recall by 41% ($z=2.33$, $p=0.02$; Odds: 1.41, 95%CI: 1.06 to 1.89) after 5 min. At one week the odds of recall increase by 79% ($z=3.64$, $p<0.001$; Odds: 1.79, 95%CI: 1.31 to 2.46).

Consistent with the previous experiments, MCI stories were significantly less believable ($M=3.23$, $sd=1.03$) than intuitive stories ($M=4.78$, $sd=1.18$ after 5 min; $t_{95}=-16.54$, $p<0.001$, 95%CI: -1.74 to -1.37; $d=-1.67$). Believability remained significantly lower for MCI ($M=3.25$, $sd=1.04$) than for intuitive items ($M=4.75$, $sd=1.23$) at one week ($t_{83}=14.65$, $p<0.001$, 95%CI: 1.29 to

Table 5 Effects of anthropomorphism on memorability for experiment 3

	<i>B</i> (<i>SE</i>)	β	95%CI	
			Lower	Upper
5 mi				
MCI				
Belief	0.01 (0.03)	0.03	-0.19	0.26
IDAQ	-0.02 (0.02)	-0.16	-0.42	0.26
Intuitive				
Belief	-0.03 (0.02)	-0.14	-0.34	0.07
IDAQ	-0.02 (0.02)	-0.16	-0.39	0.09
1 week				
MCI				
Belief	0.04 (0.03)	0.17	-0.05	0.39
IDAQ	-0.05 (0.02)*	-0.26	-0.48	-0.04
Intuitive				
Belief	0.02 (0.03)	0.08	-0.15	0.32
IDAQ	-0.02 (0.02)	-0.13	-0.37	0.11

* $p < 0.05$. Controls: age, gender, ethnicity, and religious affiliation

1.70; $d = -1.58$). There was no change in believability between time 1 and time 2 for either MCI ($t_{82} = 0.75$, $p = 0.46$) or intuitive items ($t_{82} = 1.41$, $p = 0.16$).⁵

The fake MCI stories were *more* believable than the real MCI stories, discrediting the hypothesis that participants were simply identifying untrue stories in their believability estimates (Fake: $M = 3.46$, $sd = 1.17$; Real: $M = 2.99$, $sd = 1.15$; $t_{98} = 4.79$, $p < 0.001$; 95%CI: 0.30 to 0.71; $d = 0.48$).

We used regression analyses to assess the effect of IDAQ and belief on memory (Table 5). Although we found a similar pattern for anthropomorphism's relationship to memory at 5 min, the effect was significant only for the second session.

The behavioral money task resulted in comparable effects to our believability scores, with nearly twice as much cash bet on intuitive stories (money per story: $M = \$0.61$, $sd = 0.25$; $t_{82} = 6.73$, $p < 0.001$) than MCI stories (money per story: $M = \$0.37$, $sd = 0.24$; 95%CI: -0.32 to -0.17; $d = 0.71$). On average, participants wagered a total of \$1.47 less (95%CI: -1.91 to -1.04) on MCI stories than on intuitive stories. Money bet on both MCI and intuitive items were positively correlated with previous belief measures, serving as a behavioral validation of the latter (MCI: $r = 0.46$, $t_{79} = 4.56$, $p < 0.001$, 95%CI: 0.26 to 0.61; INT: $r = 0.59$, $t_{79} = 6.48$, $p < 0.001$, 95%CI: 0.43 to 0.72).

⁵ Using a multilevel model with random intercepts for each individual, belief is associated with a change of -1.59 (95%CI: -1.67 to -1.39; $t_{105.29} = -15.65$; $p < 0.001$) points on a seven-point scale when comparing MCI and intuitive stories after 5 min, and an average of -1.58 (95%CI: -1.85 to -1.32; $t_{118.70} = 11.77$, $p < 0.001$) points on seven-point scale after 1 week.

Discussion

We did not find a memory effect for MCI content in the first experiment, but we did find an effect in experiments 2 and 3. The memory bias in experiment 1 may have been washed out by the other information presented in these stories (Gonce et al. 2006; Upal et al. 2007), such as social information (Mesoudi et al. 2006), or by the narrative structure itself (e.g., Brewer 1985; Ericsson & Kintsch 1995; Rubin 1995). The significance of the memory effect in experiment 3 was dependent on the statistical method used. This makes sense, given that the memory effect is small to moderate (5 min: $d=0.19$; 1 week: $d=0.27$), and our sample size lacked the power to reliably detect it based on a mean difference t test. The memory effect was the most robust in experiment 2 and was similar in magnitude to the effect found by Boyer and Ramble (effect after 5 min: $d=0.57$; Boyer and Ramble's experiment 1; $d=0.68$, estimated from F test, $n=18$). Together, these effects suggest that a memory bias is present but is sensitive to contextual changes.

As predicted, trait anthropomorphism had a negative effect on memory for anthropomorphic content across all three experiments, but the effect varied in magnitude and was less consistent across the two-session experiments. The effect was strongest in experiment 1 looking at anthropomorphic-MCI content ($\beta=-0.25$), but still marginally present for physical-MCI stories ($\beta=-0.19$). In experiment 2, an effect for both MCI objects and people, and a marginal effect on MCI object recall after 1 week, was found. Since all of these "objects" contain anthropomorphic violations, the larger effect for objects is consistent with the hypothesis. In experiment 3 the effect was only present in the second session. This may be due to the stimuli itself; the MCI content in these stories was not limited to specifically anthropomorphic information. Given the moderate size of this effect in the prior two experiments, it should be expected to be nonsignificant some of the time in samples of this size.

Interestingly, we also found a negative effect of anthropomorphism on memory for the intuitive people stimuli in experiment 2. This could be a statistical fluke, or it could represent a more general trend of the impact of anthropomorphism on memory. Though this seems unlikely and the mechanism remains unclear, there is a negative, though nonsignificant, impact of anthropomorphism on memory for intuitive content in studies 1 and 3 as well. Regardless, the effect of trait anthropomorphism on memory for anthropomorphic content is a larger and more robust effect. This effect suggests that how much we expect the world to behave according to these ontological categories can be different across individuals, and this difference can translate into individual differences in the effect of anthropomorphic content on memory.

The effects of MCI content on belief were unambiguous. Across all three experiments, MCI content was found much less believable than intuitive content (d values between -1.58 and -2.49). These effects did not decline after a one-week interval, suggesting that, at least in the short term, this effect does not change with time. It could be argued that repeated exposure, rather than just time, could lead to increased belief (e.g., Zaragoza & Mitchell 1996). Though we did not test this directly, by the time the belief measure was taken for the second time in experiments 2 and 3, participants had read the items three times and recalled them twice. This repeat exposure did not affect their willingness to believe in these items.

The monetary measure in study 3 replicated the belief findings ($d=-0.71$), indicating that people were willing to systematically bet against the reality of MCI stories. Even when rewarded for correct answers, people persistently believed that the intuitive stories were more likely to be true than MCI ones.

Experiment 4

Experiments 1–3 showed that people found MCI content less believable than intuitive content. When we consider the religious content the MCI hypothesis proposes to explain, this lack of belief becomes a problem. Religious ideas are ideas people believe in. If religious ideas include MCI content, then people must come to believe in this type of content somehow. A partial solution to this puzzle is that belief—particularly counterintuitive belief—is sustained by a variety of cultural learning mechanisms (Atran & Henrich 2010; Gervais et al. 2011; Henrich 2009; Norenzayan 2013). The theoretical problem of belief (why people come to believe in some MCI concepts and not others) has been widely claimed, but never empirically shown. In the first three experiments, we show that MCI content is found to be highly unbelievable. At the same time, we know that people do come to believe in this and other types of unbelievable content. This problem of belief is not insurmountable, and we propose it is frequently and easily overcome with non-content-based cultural learning (i.e., learning based on social transmission and not based on the content of the idea itself).

In our fourth experiment, we investigated one possible cultural learning bias, CREDs (Henrich 2009), which might explain how MCI content becomes believable through the observation of another person's behavior. Participants were given news stories from the previous experiment, as well as \$5 in dollar coins to bet on the truth of stories. Their money was doubled on correct bets and lost on incorrect bets. Any money they did not bet was theirs to keep. A confederate disguised as another participant also completed the experiment at the same time as the participant and was observable by the participant. The confederate verbally endorsed either the MCI stories or the intuitive stories and either put money down on the stories he endorsed (the CRED) or put no money down at all (no CRED). We hypothesized that participants who witness an endorsement combined with a CRED would be more likely to endorse a given story themselves by betting money on it.

Methods

For the participants, see Table 1.

Materials The stimuli consisted of 10 of the 12 stories from experiment 3. One MCI story and one intuitive story were removed from the set to make five stories of each type. This was done to make the experiment shorter so the confederate could move from one room to another without risking bumping into the next participant.

Design and Procedures Participants came into the lab and were given the news stories to read in a random order. They rated their belief in each story with the same three questions used in experiments 2 and 3. The anthropomorphism scale was excluded because memory was not theoretically relevant to this task since memory was not being assessed.

Participants were brought into a room with a confederate posing as a second participant. Both people were given five \$1 coins and told to bet the money on any story they believed was a true story. If their belief was correct, the money would be doubled, if they were incorrect, they would lose the money. They were instructed to discuss their decisions with one another. The confederate either bet all of his money, \$1 per story, on the MCI or the intuitive stories (CRED) or bet no money (no-CRED). In both the CRED and verbal (no-CRED) conditions he verbally endorsed one set of stories (MCI or intuitive) over the other. The same statements of endorsement were used in both the CRED and no-CRED conditions (see [ESM](#)). In the CRED condition, the confederate additionally stated that he would put a dollar down for each of relevant stories. In the verbal (no-CRED) condition, no comment was made about money and the confederate did not write on the paper at all. The confederate's comments about each story were scripted, but he was allowed to make unscripted small talk if it wasn't about the story. The confederate was further instructed to disagree with the participant if the participant expressed a differing option. For any disagreement that was not covered by the script, the confederate was instructed to disagree, but not to add any additional information about the content. For example, if the participant said they thought an intuitive story was correct in the MCI condition, the confederate would simply say something such as "it just doesn't sound right to me" or "it seems untrue." Though the confederate and the participant were in the same room, the confederate could not see what stories the participant marked. This interaction was cut off after 2 min. The participant and the confederate were separated and the participant was asked to rate the confederate on likability, trust, reliability, and how much they agree with his opinions. Suspicion about the confederate and the purpose of the experiment was probed with a funneled debriefing. Eight participants were removed because of their suspicion.

Results

The initial pre-betting ratings of belief replicated the belief findings from the previous experiments. Participants rated the MCI stories as less believable ($M=3.18$, $sd=1.10$) than the intuitive stories ($M=4.66$, $sd=0.99$; $t_{60}=-9.95$, $p<0.001$, 95%CI: -1.76 to -1.17 ; $d=-1.22$). Across the whole sample we found a significant difference between the money bet on MCI stories ($M=\$0.98$, $sd=1.25$) and that bet on intuitive stories (money per story: $M=\$1.84$, $sd=1.71$; $t_{60}=-3.08$, $p=0.003$, 95%CI: -1.41 to -0.30 ; $d=0.32$). Critically, we found that when the betting was credibly displayed, the participants bet more money across all stories ($M=\$3.65$, $sd=1.67$) than when it was not displayed ($M=\$1.96$, $sd=2.04$; Welch's $t_{58.05}=3.58$, $p<0.001$, 95%CI: $\$0.74$ to $\$2.63$; $d=0.98$).

To assess if the credible display affected bets overall, or only items that were bet on, we regressed the four conditions as dummy codes on the amount of money placed on each story type individually. If the CRED affected betting overall, then both CREDs conditions should be significant for both story types. If the CRED affected only the stories to which the CRED was directed, we should see type-specific effects (i.e., CRED for MCI should positively affect the money placed on MCI stories, but not the money placed on intuitive stories; Table 6). When CREDs were targeted on MCI items, participants increased the amount of money they placed on MCI stories but not intuitive stories. When CREDs were targeted on intuitive items, participants increased the amount of money placed on intuitive stories but not MCI stories.

We analyzed the data a second time using a multilevel logistic regression with random intercepts and slopes for each individual. This allows us to compare the items that were endorsed with those that were not endorsed as well as whether the endorsement was verbal or involved a CRED. The random intercepts correct the standard errors for the lack of independence of our data points. We found a main effect of CRED, increasing the odds of betting by an average of 6.68, but not of verbal endorsement (Table 7). This effect is controlling for the significant effect of prior belief on betting. There was no significant interaction between CRED and story type (MCI) and a marginal interaction effect of verbal endorsement of MCI stories in our sample (Fig. 1).

Table 6 Regression using condition to predict bets above verbal endorsement

	<i>B</i> (<i>SE</i>)	β	95%CI	
			Lower	Upper
Money on MCI				
Intercept	-0.37 (1.52)	0.05	-0.56	0.65
Verbal Intuitive	-0.60 (0.36)	-0.48	-1.06	0.10
CRED Intuitive	-0.59 (0.36)	-0.47	-1.05	0.11
CRED MCI	1.30 (0.37)***	1.04	0.45	1.63
Belief MCI	0.25 (0.12)*	0.22	0.01	0.42
Money on Intuition				
Intercept	-0.75 (2.58)	-0.40	-1.07	0.27
Verbal MCI	0.04 (0.55)	0.02	-0.63	0.67
CRED Intuition	1.94 (0.55)***	1.13	0.48	1.78
CRED MCI	0.26 (0.56)	0.15	-0.51	0.82
Belief Intuition	0.42 (0.21)	0.26	-0.05	0.44

All DV ratings between \$0 and \$5. Categorical variables were not standardized, leading to non-zero intercepts for β . These betas can be interpreted as the average standard deviation change in the dependent variable for each category compared to the intercept. The unstandardized betas are the average change in dollars. Controls: age, gender, ethnicity, and religious affiliation

† $p < 0.10$, * $p < 0.05$, *** $p < 0.001$, $n = 62$

Table 7 Multilevel logistic regression using condition to predict bets

	<i>B</i> (<i>SE</i>)	Odds	95%CI	
			Lower	Upper
Place Bet (yes = 1)				
Intercept	-4.68 (0.67)***	0.01	0.00	0.03
CRED	1.90 (0.67)**	6.68	1.80	24.68
Verbal	-1.31 (0.83)	0.27	0.05	1.36
MCI	-0.45 (0.51)	0.63	0.23	1.73
Belief	0.74 (0.11)***	2.10	1.67	2.62
CRED*MCI	-0.13 (0.87)	0.88	0.51	2.33
Verbal*MCI	2.05 (1.06)†	7.77	0.97	62.35

Controls: age (centered at mean) and gender. Ethnicity and religious affiliation were not included in the model because the model failed to converge with the addition of these variables

Number Obs: 630, groups (Subject): 63

† $p < 0.10$, ** $p < 0.01$, *** $p < 0.001$

We also found a significant positive effect of the CRED manipulation on how much the participant thought the confederate was trustworthy and reliable, as well as a marginal effect on liking (Table 8). There was a negative effect of endorsing MCI stories on the agreement rating.




Fig. 1 Money bet on news stories. The None category represents a CRED or verbal endorsement of the opposite category (e.g., if MCI stories had no endorsement, it was because the intuitive stories were being endorsed). The drop from Verbal to None in MCI stories reflects the fact that endorsing intuitive stories decreased the offers on MCI stories. The opposite was not the case (see Table 6). Error bars are 95% confidence intervals of the means

Table 8 OLS regression using condition to predict participants' ratings of the confederate

	<i>B</i> (<i>SE</i>)	β	95%CI	
			Lower	Upper
Like				
CRED	0.58 (0.31)†	0.50	-0.03	1.03
MCI	-0.31 (0.32)	-0.27	-0.81	0.28
Trust				
CRED	0.87 (0.30)**	0.68	0.21	1.15
MCI	-0.26 (0.31)	-0.21	-0.69	0.27
Reliable				
CRED	0.73 (0.34)*	0.51	0.03	0.98
MCI	-0.29 (0.34)	-0.20	-0.68	0.27
Agree with				
CRED	0.50 (0.35)	0.32	-0.13	0.77
MCI	-1.40 (0.36)***	-0.89	-1.35	-0.43

† $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Discussion

The CRED performed by the confederate—that is, betting money in accordance with their opinions—increased the odds by nearly sevenfold that the participant would bet on the same stories. Participants tended to place their money on the stories that the confederate bet on rather than just picking stories at random, or picking the stories they had previously rated as believable. When participants saw the confederate bet on MCI stories, those stories were endorsed at a higher rate than intuitive stories with only verbal endorsement. Based on the pre-manipulation ratings of belief, there does not seem to be much ambiguity in which stories participants thought were false. This experiment supports the idea that witnessing a behavior that is consistent with a model's expressed beliefs increases the likelihood of endorsing those beliefs. In other words, beliefs that are backed up by CREDs can be culturally contagious.

Also consistent with the hypothesis, the CRED manipulation also affected the participants' opinions of the confederate. When the confederate acted in accordance with his opinions, the participants saw him as more trustworthy ($\beta = 0.68$), more reliable ($\beta = 0.51$), and marginally more likable ($\beta = 0.50$), suggesting that acting according to one's beliefs may lead to reputational benefits as well as spread behaviors and potentially beliefs.

Within the betting scenario, there was no significant interaction between CRED and story type. The CRED manipulation was effective for both types of information. This is consistent with the theory, which suggests that all types of information should be susceptible to social and cultural learning cues (see Henrich 2009; Henrich & McElreath 2003). Nevertheless, the CRED effect on MCI content is of particular interest because it could help resolve the apparent

paradox that MCI content is commonly found in religious traditions even though it may initially be less believable.

This is just a first step in testing a complex set of processes, with many open questions. This experiment offers some insight into the plausibility of cultural learning and social influence as a source of belief in supernatural concepts. Several other social learning mechanisms could work as additions to, or in the place of, CREDs. Still, this research suggests CREDs as one plausible candidate to aid in the transmission of religious practices and beliefs. This area of research is yet to be thoroughly explored in the cognitive science of religion.

General Discussion

Evidence from these experiments, and previous research, suggests that the memory effect of MCI content is present but fragile, and boundary conditions need to be identified (see Purzycki & Willard 2015). For example, changes in context seem to be enough to make the effect disappear (see Gonce et al. 2006; Upal et al. 2007). This and other research suggests that if an MCI memory bias exists, then further work needs to be done in exploring the boundaries of this bias. Under what conditions is this bias robust, and under what conditions is it not?

One possible boundary explored in this paper is the effect of the tendency to habitually anthropomorphize on the memory for anthropomorphic MCI content. We surmise that this decrease is due to the lack of distinctiveness of this type of MCI content to people who see this type of explanation everywhere. If one believes the ocean is conscious, and that one's computer occasionally gets mad at them, they may not be surprised when this type of anthropomorphic thinking is applied to other phenomena.

There are two potential explanations. One is a schema view of MCI content advanced by Upal (Upal 2010) and Purzycki (Purzycki 2010), rather than the original ontological category view expressed by others (Atran & Norenzayan 2004a; Boyer 2001; Boyer & Ramble 2001). If this is the case, we should expect the memory effect of MCI content to disappear and these concepts to become familiar. When we first come across these concepts, they seem surprising and therefore are memorable, but as we become familiar with them they are integrated into certain schemas we have of the world. They become commonplace, and the memory advantage disappears. A declining memory advantage such as this cannot explain religions' spread through time, but only why they might initially spread. A separate mechanism would have to explain why they persist. Second is the view that anthropomorphism is a special case, and other types of ontological violations, governed by core knowledge domains, are not as easily updated. Further work is needed to assess this possibility.

Given the general negative trend for anthropomorphism on memory for all content types, perhaps some of the effects we are seeing are due to some other cognitive capacity related to both increased anthropomorphism and decreased memory. Even if this is the case, it would not explain the type-specific

decreases we see in these experiments. If a general effect was entirely responsible for these affects, there should be no differences across conditions.

In all of these experiments, clearly people do not readily believe MCI concepts. Though the idea that counterintuitive content is hard to believe may be an uncontroversial claim, the religious phenomena that the MCI hypothesis proposes to explain entail some degree of belief or commitment—otherwise the hypothesis cannot be about religious *belief*. It could be argued that the lack of belief we found comes from our sample, which is made up entirely of educated Westerners (see Henrich et al. 2010). Perhaps people in other places and times, without Western influence and education, or those not exposed to Abrahamic religions, would not show as much skepticism toward MCI items. This conjecture runs counter to one of the foundational premises of the MCI hypothesis—namely, that these violated categories are rooted in reliably developing, core intuitions. Further, it seems unlikely that people in other places and times were less, rather than more, familiar with counterintuitive claims about the world than we are in the West today.

We also found the MCI content was consistently deemed less believable than intuitive content. This relationship between MCI content and belief illustrates three points: (1) MCI items are seen as more supernatural than intuitive items (see Atran & Norenzayan 2004b; Pyysiäinen et al. 2003); (2) MCI items are less believable than intuitive items; and (3) supernatural concepts are believed in all around the world. For all of these to be true, we need to explain how people come to believe in the unbelievable that is separate from memorability. Given the existence of complex cultures and ritual practices to support religious beliefs, multiple processes are likely to contribute to belief. One process we found evidence for is that of credible displays. These displays convey the authenticity of beliefs to others and can increase the likelihood that witnesses adopt similar practices and beliefs (Henrich 2009).

There are many other cultural learning mechanisms through which people transmit information (see Henrich & Boyd 1998; Henrich & Gil-White 2001; Henrich & McElreath 2003). We used CREDs here because it is a learning mechanism proposed to deal specifically with the type of unverifiable information we find in religious beliefs. Our findings support CREDs as a possible candidate for belief transmission. There is much more to be done to understand the role of CREDs and other cultural learning mechanisms in religious and other cultural beliefs. The study presented here is in no way meant to be a conclusive test in regard to the robustness of CREDs, but rather a preliminary demonstration that social cues can influence people's belief in the implausible.

Conclusion

MCI content is a prominent feature of religious beliefs, but more work needs to be done to explain the boundaries and conditions of this phenomenon. Showing a memory bias solves only a small part of this issue, especially if this memory effect is fragile and affected by expectations. If belief is illusive for MCI content, we require a strong mechanism to explain how this content becomes

the center of belief systems. We demonstrate CREDs to be one such mechanism. Our findings suggest that cultural learning and social influence are an important part of how we come to accept counterintuitive ideas as true. These types of influences on the belief in supernatural agents and concepts require more attention and research if we wish to understand religions and their role in human psychology and behavior.

Acknowledgments The authors would like to thank Benjamin Purzycki and Adam Baimel for their helpful comments. JH would like to thank the Canadian Institute for Advanced Research and NYU's Stern School of Business. AN acknowledges support from a Social Sciences and Humanities Research Council (SSHRC) Insight Grant (435-2014-0456). We also thank the Cultural Evolution of Religion Research Consortium, funded by a generous partnership grant (895-2011-1009) also from SSHRC.

References

- Atran, S. (2002). *In gods we trust: The evolutionary landscape of religion*. Oxford: Oxford University Press.
- Atran, S., & Henrich, J. (2010). The evolution of religion: how cognitive by-products, adaptive learning heuristics, ritual displays, and group competition generate deep commitments to prosocial religions. *Biological Theory*, 5(1), 18–30.
- Atran, S., & Norenzayan, A. (2004a). Religion's evolutionary landscape: counterintuition, commitment, compassion, communion. *Behavioral and Brain Sciences*, 27, 713–730.
- Atran, S., & Norenzayan, A. (2004b). Why minds create gods: devotion, deception, death, and arational decision making. [reply to commentators]. *Behavioral and Brain Sciences*, 27, 754–770.
- Barrett, J. L. (2004). *Why would anyone believe in God?* Walnut Creek: AltaMira Press.
- Barrett, J. L. (2008a). Coding and quantifying counterintuitiveness in religious concepts: theoretical and methodological reflections. *Method & Theory in the Study of Religion*, 20(4), 308–338.
- Barrett, J. L. (2008b). Why Santa Claus is not a god. *Journal of Cognition and Culture*, 8(1), 149–161.
- Barrett, J. L., & Nyhof, M. A. (2001). Spreading nonnatural concepts. *Journal of Cognition and Culture*, 1, 69–100.
- Bergstrom, B., Moehlmann, B., & Boyer, P. (2006). Extending the testimony problem: evaluating the truth, scope, and source of cultural information. *Child Development*, 77(3), 531–538.
- Boyer, P. (1994). *The naturalness of religious ideas: A cognitive theory of religion*. Berkeley: University of California Press.
- Boyer, P. (2001). *Religion explained: The evolutionary origins of religious thought*. New York: Basic Books.
- Boyer, P. (2003). Religious thought and behaviour as by-products of brain function. *Trends in Cognitive Sciences*, 7(3), 119–124.
- Boyer, P., & Ramble, C. (2001). Cognitive templates for religious concepts: cross-cultural evidence for recall of counter-intuitive representations. *Cognitive Science*, 25, 535–564.
- Brewer, W. F. (1985). The story schema: universal and culture-specific properties. In D. R. Olson, N. Torrance, & A. Hildyard (Eds.), *Literacy, language, and learning* (pp. 167–194). Cambridge: Cambridge University Press.
- Cohen, A. B., Siegel, J. I., & Rozin, P. (2003). Faith versus practice: different bases for religiosity judgments by Jews and Protestants. *European Journal of Social Psychology*, 33(2), 287–295.
- Epley, N., Akalis, S., Waytz, A., & Cacioppo, J. T. (2008a). Creating social connection through inferential reproduction: loneliness and perceived agency in gadgets, gods, and greyhounds. *Psychological Science*, 19, 114–120.
- Epley, N., Waytz, A., Akalis, S., & Cacioppo, J. T. (2008b). When we need a human: motivational determinants of anthropomorphism. *Social Cognition*, 26, 143–155.
- Ericsson, K. A., & Kintsch, W. (1995). Long-term working memory. *Psychological Review*, 102(2), 211–245.
- Gervais, W. M., & Henrich, J. (2010). The Zeus problem: why representational content biases cannot explain faith in gods. *Journal of Cognition and Culture*, 10, 383–389.
- Gervais, W. M., Willard, A. K., Norenzayan, A., & Henrich, J. (2011). The cultural transmission of faith: why natural intuitions and memory biases are necessary, but insufficient, to explain religious belief. *Religion*, 41(1), 389–400.

- Gonce, L. O., Upal, M. A., Slone, D. J., & Tweney, D. R. (2006). Role of context in the recall of counterintuitive concepts. *Journal of Cognition and Culture*, 6, 521–547.
- Greenhouse, S. W., & Geisser, S. (1959). On methods in the analysis of profile data. *Psychometrika*, 24, 95–112.
- Harris, P. L., & Koenig, M. A. (2006). Trust in testimony: how children learn about science and religion. *Child Development*, 77(3), 505–524.
- Henrich, J. (2009). The evolution of costly displays, cooperation, and religion: credibility enhancing displays and their implications for cultural evolution. *Evolution and Human Behavior*, 30(244–260), 244–260.
- Henrich, J., & Boyd, R. (1998). The evolution of conformist transmission and the emergence of between-group differences. *Evolution and Human Behavior*, 19, 215–241.
- Henrich, J., & Gil-White, F. J. (2001). The evolution of prestige: freely conferred deference as a mechanism for enhancing the benefits of cultural transmission. *Evolution and Human Behavior*, 22(3), 165–196.
- Henrich, J., & McElreath, R. (2003). The evolution of cultural evolution. *Evolutionary Anthropology*, 12(12), 123–135.
- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, 33(2–3), 61–83.
- Lang, V. A. (1995). Relative association, interactiveness, and the bizarre imagery effect. *The American Journal of Psychology*, 108(1), 13–35.
- Lisdorf, A. (2004). The spread of non-natural concepts. *Journal of Cognition and Culture*, 4, 151–173.
- McDaniel, M. A., Einstein, G. O., DeLosh, E. L., May, C. P., & Brady, P. (1995). The bizarreness effect: it's not surprising, it's complex. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21(2), 422–435.
- Mesoudi, A., Whiten, A., & Dunbar, R. (2006). A bias for social information in human cultural transmission. *British Journal of Psychology*, 97, 405–423.
- Norenzayan, A. (2013). *Big gods: How religion transformed cooperation and conflict*. Princeton: Princeton University Press.
- Norenzayan, A., Atran, S., Faulkner, J., & Schaller, M. (2006). Memory and mystery: the cultural selection of minimally counterintuitive narratives. *Cognitive Science*, 30(3), 531–553.
- Purzycki, B. G. (2010). Cognitive architecture, humor and counterintuitiveness: retention and recall of MCIs. *Journal of Cognition and Culture*, 10, 189–204.
- Purzycki, B. G., & Sosis, R. (2011). Our gods: variation in supernatural minds. In U. J. Frey, C. Störmer, & K. P. Willführ (Eds.), *Essential building blocks of human nature* (pp. 77–93). Berlin: Springer.
- Purzycki, B. G., & Willard, A. K. (2015). MCI theory: a critical discussion. *Religion, Brain & Behavior*, 1–42.
- Pyysiäinen, I., Lindeman, M., & Honkela, T. (2003). Counterintuitiveness as the hallmark of religiosity. *Religion*, 33(4), 341–355.
- Riefer, D., & Lamay, M. (1998). Memory for common and bizarre stimuli: a storage-retrieval analysis. *Psychonomic Bulletin & Review*, 5(2), 312–317.
- Rubin, D. C. (1995). *Memory in oral traditions: The cognitive psychology of epic, ballads, and counting-out rhymes*. New York: Oxford University Press.
- Russell, Y. I., & Gobet, F. (2013). What is counterintuitive? Religious cognition and natural expectation. *Review of Philosophy and Psychology*, 4(4), 715–749.
- Shariff, A. F., Purzycki, B. G., & Sosis, R. (2014). Religions as cultural solutions to social living. In A. B. Cohen (Ed.), *Culture reexamined* (pp. 217–238). Washington, DC: American Psychological Association.
- Spelke, E. S., & Kinzler, K. D. (2007). Core knowledge. *Developmental Science*, 10(1), 89–96.
- Sperber, D. (1996). *Explaining culture: A naturalistic approach*. Oxford: Blackwell.
- Traub, R. E., & Rowley, G. L. (1991). Understanding reliability. *ITEMS*, 10, 37–45.
- Upal, M. A. (2010). An alternative account of the minimal counterintuitiveness effect. *Cognitive Systems Research*, 11(2), 194–203.
- Upal, M. A., Gonce, L. O., Tweney, R. D., & Slone, D. J. (2007). Contextualizing counterintuitiveness: how context affects comprehension and memorability of counterintuitive concepts. *Cognitive Science*, 31(3), 415–439.
- Waytz, A., Cacioppo, J., & Epley, N. (2010a). Who sees human? The stability and importance of individual differences in anthropomorphism. *Perspectives on Psychological Science*, 5(3), 219–232.
- Waytz, A., Morewedge, C. K., Epley, N., Monteleone, G., Gao, J.-H., & Cacioppo, J. T. (2010b). Making sense by making sentient: effacement motivation increases anthropomorphism. *Journal of Personality and Social Psychology*, 99(3), 410–435.

- Wellman, H. M., & Gelman, S. A. (1992). Cognitive development: foundational theories of core domains. *Annual Review of Psychology*, 43(1), 337–375.
- Willard, A. K., & Norenzayan, A. (2013). Cognitive biases explain religious belief, paranormal belief, and belief in life's purpose. *Cognition*, 129(2), 379–391.
- Zaragoza, M. S., & Mitchell, K. J. (1996). Repeated exposure to suggestion and the creation of false memories. *Psychological Science*, 7(5), 294–300.

Dr. Aiyana Willard is a postdoctoral research fellow in psychology at the University of Texas at Austin. Her research focuses on the cognitive and cultural origins of supernatural and religious belief. She is currently conducting research on “spiritual but not religious” people in North America and Europe and witchcraft beliefs around the world. She maintains an active field site in Fiji, where she conducts research with Hindu and Muslim populations.

Dr. Joseph Henrich is a professor of human evolutionary biology at Harvard University and holds the Canada Research Chair in Culture, Cognition and Coevolution at the University of British Columbia, where he's also a professor in both economics and psychology. In 2009, he won the Early Career Award for Distinguished Scientific Contributions bestowed by the Human Behavior and Evolution Society. His latest book is *The Secret of Our Success: How Culture Is Driving Human Evolution, Domesticating Our Species, and Making Us Smarter* (Princeton University Press, 2016).

Dr. Ara Norenzayan is professor of psychology and co-director of the Centre for Human Evolution, Cognition and Culture at the University of British Columbia. He received his PhD in psychology from the University of Michigan in 1999. He has published widely on the cognitive science of religious belief, the evolutionary origins of religion and religious diversity, cultural evolution, and culture and cognition. He was the recipient of the 2014–2015 James McKeen Cattell Fund Fellowship in Psychology. He is the author of *Big Gods: How Religion Transformed Cooperation and Conflict* (Princeton University Press, 2013).