

How many atheists are there?

Will M. Gervais

Maxine B. Najle

University of Kentucky Psychology

Contact: will.gervais@uky.edu or maxine.najle@uky.edu

Abstract

One crucible for theories of religion is their ability to predict and explain patterns of belief and disbelief. Yet, religious nonbelief is often heavily stigmatized, potentially leading many atheists to refrain from outing themselves even in anonymous polls. We used the unmatched count technique and Bayesian estimation to indirectly estimate atheist prevalence in two nationally representative samples of 2000 U.S. adults apiece. Widely-cited telephone polls (e.g., Gallup, Pew) suggest USA atheist prevalence of only 3-11%. In contrast, our most credible indirect estimate is 26% (albeit with considerable estimate and method uncertainty). Our data and model predict that atheist prevalence exceeds 11% with greater than .99 probability, and exceeds 20% with roughly .8 probability. Prevalence estimates of 11% were even less credible than estimates of 40%, and all intermediate estimates were more credible. Some popular theoretical approaches to religious cognition may require heavy revision to accommodate actual levels of religious disbelief.

Keywords: religion; atheism; social desirability; stigma; Bayesian estimation

How many atheists are there?

I am, and I wish I weren't.

Aldous Huxley, *Brave New World*

Religion is a core aspect of human nature, yet a comprehensive understanding of religion must also accommodate religious disbelief. Various models seek to explain commonalities in religious cognition (e.g., Boyer, 2008) and the possible adaptive benefits of religious beliefs and practices (e.g., Norenzayan et al., 2014). These models have advanced the naturalistic science of religion, one key challenge in interdisciplinary, consilient approaches to human nature (Wilson, 1999). The scientific success of theories of religion partially hinges on the degree to which they successfully predict and explain the distribution of belief and disbelief. The study of atheists—merely people who disbelieve or lack belief in the existence of God or gods¹—is thus a central testing ground for basic theory on the origins of religious and supernatural beliefs (Norenzayan & Gervais, 2013). Unfortunately, unbiased answers to some of the most elementary questions regarding belief and disbelief are currently unavailable. Such basic questions as: How many atheists are there?

¹ As a terminological aside, throughout this paper we use the term “atheist” to refer to people who disbelieve or lack belief in the existence of a god or gods. This definition is standard in the psychology, sociology, and philosophy of religion. Furthermore, it is the definition of “atheist” adopted by the Oxford English Dictionary. Operationally, we thus define people as atheists if they do not indicate belief in a god or gods. We focused on binary classification of atheists not because we necessarily view religious belief as a psychologically binary phenomenon, but rather for direct comparability with existing polling data utilizing binary measurements of this complex construct.

Existing global atheism estimates (Zuckerman, 2007) necessarily rely on self-report data. Yet, religious disbelief carries substantial social and reputational costs (Edgell, Gerteis, & Hartmann, 2006; Gervais, 2013). Given the centrality of religious belief to many societies (Inglehart & Norris, 2004), and the degree to which many equate religious belief with morality (Gervais, 2014a, 2014b; McKay & Whitehouse, 2014), there are profound social pressures to be—or at least appear—religious (Hadaway, Marler, & Chaves, 1993; Sedikides & Gebauer, 2009). We attempted to experimentally generate an atheist prevalence estimate in the United States that is less influenced by socially desirable responding than standard nationally representative telephone polls. This serves as an initial step towards understanding patterns of religious belief and disbelief worldwide, thus providing raw data for testing basic theories on the bases of religion.

Representative telephone polls reveal that only 3% of Americans self-identify with the term “atheist” (Pew, 2015), and only 11% deny believing in God when given binary (yes/no) response options (Gallup, 2015). Social pressures favoring religiosity, coupled with stigma against religious disbelief (Edgell et al., 2006), might cause people who privately disbelieve in God to nonetheless self-present as believers, even in anonymous questionnaires. This is especially true for telephone polls, which require individuals to verbally disclose their atheism to others. Personalized telephone polling yields inflated claims of religiosity, relative to more private internet polls (Cox, Jones, & Navarro-Rivera, 2014). However, even fully anonymous online polls may still yield underestimates of atheist prevalence. Thus, indirectly measured atheism rates—using techniques designed to obviate social desirability pressures—might be substantially

higher than telephone self-reports suggest. We tested this hypothesis in two nationally representative samples.

Study Overview

We derived two indirectly measured atheist prevalence estimates from two separate preregistered nationally representative online samples, total $N = 4000$. This two-pronged approach allowed us to 1) generate more precise pooled atheism prevalence estimates using, 2) slightly different measures across samples, while 3) assessing the validity of the indirect measure in two distinct ways. In this paper, we first describe the general indirect measure methodology, then the particulars of each individual sample. Next, we present indirect measurement results for each individual sample along with some validity assessments for the indirect measure. The primary inferences of this paper, however, stem from a single overall model pooling both samples together. This pooled model also allows us to simultaneously perform exploratory analyses in which we infer atheism prevalence rates across a number of demographic contrasts common to both samples. We report how we determined our sample sizes, all data exclusions (if any), all manipulations, and all measures in the study.

Method

Participants

Each of the two nationally representative samples consisted of 2000 American adults. We contracted these samples from YouGov, a firm specializing in omnibus nationally representative polling. They offer samples of either 1000 or 2000, and we opted for 2000 to maximize estimate precision. YouGov also provides data from a

number of standard demographic questions, full details of which are available in our freely shared datasets at <https://osf.io/4q85g/>

General Procedure

For both samples, we indirectly inferred atheism rates using the unmatched count technique (e.g., Dalton, Wimbush, & Daily, 1994; Raghavarao & Federer, 1979), a tool for inferring base rates of socially sensitive outcomes. The unmatched count technique indirectly infers underlying base rates for socially undesirable or unacceptable outcomes by randomly assigning participants to one of two versions of a count task. In one version, participants indicate how many innocuous statements from a list (e.g., I can drive a motorcycle; I exercise regularly) are true of them. In the other version, participants receive a list that is identical, save for the addition of one sensitive item (e.g., I can drive a motorcycle; I exercise regularly; I smoke crack cocaine), and they indicate how many items are true of them. Crucially, nobody indicates *which specific* items are true of them, only how many in total. The difference between the aggregate rates in these conditions can presumably be attributed to the addition of the socially sensitive item. In using this task to indirectly measure atheist prevalence, our approach mirrors recent work using the unmatched count technique to indirectly estimate the size of the LGBT community as well as antigay sentiment (Coffman, Coffman, & Ericson, 2016). Crucially, this work includes extensive validation of the task's utility in estimating the size of stigmatized groups, finding that the UCT does not appear to be driven by inattentive or random responding, and only generally diverges from self-reports of socially undesirable attributes (but not generic foil attributes). The task appears robust as well to participant inattentiveness and random responding (Coffman et al., 2016).

Detail of Samples

In Sample I, we wanted to compare self-reported atheism prevalence to indirect atheism prevalence estimates inferred using the unmatched count technique. Thus, participants were randomly assigned to one of three experimental conditions: 1) indicate how many of a list of nine innocuous statements were *not* true of them, 2) indicate how many from a list of ten statements (the nine previously used, plus the statement “I believe in God”) were *not* true of them, or 3) self-report whether or not they believe in God, using standard binary Gallup poll wording.

In Sample II, we sought to evaluate the robustness of our estimate to slight perturbations in the unmatched count protocol and question wording. Specifically, we omitted the self-report condition and instead had three unmatched count conditions. In the baseline condition, participants simply rated how many of the innocuous statements were true of them. In the critical condition, participants rated how many of seven statements were true of them (the six innocuous items, plus the statement “I do not believe in God”). In previous work (Coffman et al., 2016), the unmatched count technique does not diverge from self-reports of bland and socially unloaded attributes such as shirt sleeve preferences, laptop computer use, or telephone provider. Mirroring this approach, our final condition was designed to assess the sensitivity of the unmatched count technique to the addition of a bizarre additional item, rather than the atheism statement. Participants rated how many of seven statements were true of them (the six innocuous statements, plus a statement of belief in a mathematical impossibility). Tables 1 and 2 show full stimuli from both samples. Due to the availability of additional

demographic information in Sample II², we were also able to assess the validity of unmatched count estimates by comparing the indirect estimate from the first two conditions among participants who did (or did not) self-identify as atheists.

Table 1. Stimuli used in Sample I.

Self-Report	UCT Baseline	UCT Target Group
<p>Do you believe in God?</p> <ul style="list-style-type: none"> • Yes • No 	<p>Please read the following statements, and count how many of them are NOT true statements about you:</p> <ul style="list-style-type: none"> • I am a vegetarian • I own a dog • I work on a computer nearly every day • I have a dishwasher in my kitchen • I can drive a motorcycle • My job allows me to work from home more than once per week • I jog at least four times per week • I enjoy modern art • I have attended a professional soccer match <p>In the space below, please write how many of these statements are NOT true for you: _____ of these statements are NOT true for me.</p>	<p>Please read the following statements, and count how many of them are NOT true statements about you:</p> <ul style="list-style-type: none"> • I am a vegetarian • I own a dog • I work on a computer nearly every day • I have a dishwasher in my kitchen • I can drive a motorcycle • I believe in God • My job allows me to work from home more than once per week • I jog at least four times per week • I enjoy modern art • I have attended a professional soccer match <p>In the space below, please write how many of these statements are NOT true for you: _____ of these statements are NOT true for me.</p>

² YouGov added limited religious demographics between the time we fielded both samples. Thus the religious demographics were unfortunately only available for Sample II.

Table 2. Stimuli used in Sample II.

UCT Baseline	UCT Atheist	UCT Math
<p>How many of the statements below are true of you?</p> <ul style="list-style-type: none"> • I can drive a stick shift • I eat meat • I have played scrabble • I've been to the South Pole • I have visited New York City • I exercise regularly <p>In the space below, please write how many of these statements are true for you: _____ of these statements are true of me.</p>	<p>How many of the statements below are true of you?</p> <ul style="list-style-type: none"> • I can drive a stick shift • I eat meat • I have played scrabble • I've been to the South Pole • I have visited New York City • I exercise regularly • I do not believe in God <p>In the space below, please write how many of these statements are true for you: _____ of these statements are true of me.</p>	<p>How many of the statements below are true of you?</p> <ul style="list-style-type: none"> • I can drive a stick shift • I eat meat • I have played scrabble • I've been to the South Pole • I have visited New York City • I exercise regularly • I do not believe that 2 + 2 is less than 13 <p>In the space below, please write how many of these statements are true for you: _____ of these statements are true of me.</p>

General Analytic Strategy

Because our primary research goal was inferring plausible parameter values for atheism in the USA, we utilized Bayesian estimation (see, e.g., Kruschke, 2010; McElreath, 2016). Bayesian analyses offer researchers a number of pragmatic benefits (Wagenmakers, Morey, & Lee, 2016), including intuitive probability statements of inferences, the ability to compare the relative credibility of different estimates, and a full posterior distribution describing the relative plausibility that different population level atheism prevalence rates could have produced our observed data. Bayesian models simply allow researchers to allocate credibility across different parameter values, given data (Kruschke, 2010). In other words, Bayesian estimation allows us to infer the degree

to which different potential population-level atheism rates could have plausibly and credibly produced our observed data. This is in contrast to frequentist approaches (e.g., confidence intervals) that only provide endpoints of a range of parameter values for which the endpoints would be expected to contain the true parameter value in 95% of samples, were the study exactly repeated a very large number of times.

All hypotheses, methods, and planned analyses (including code) were preregistered before data collection commenced (<https://osf.io/byma4/> and <https://osf.io/st6d3/>). All indirect estimates used priors designed to allow the data to speak for themselves with minimal prior influence. Indirect priors were mildly regularizing and reflect a normal distribution, centered on Gallup's self-report estimate (.11), with a standard deviation of 1. See Online Supplement for visual depiction of these priors. For robustness, we conducted analyses with alternative priors. Estimates did not appreciably change when using alternative priors. We summarize estimates with a point estimate reflecting the most credible parameter value and represent uncertainty around this estimate using 97% highest posterior density intervals (HPDI) in brackets, which reflect the range in which the most credible 97% of parameters lie. In addition, we provide graphical summaries of all key posterior distributions so the reader can visualize the range of possible atheism prevalence rates along with their relative credibility.

Results

We briefly report results from each sample individually (indirect estimate posteriors for both appear in Figure 1), then pool the samples for aggregate analyses, including demographic breakdowns.

Individual Sample Indirect Inferences

Sample I's unmatched count data revealed atheism rates much higher than existing self-reports suggest: the most credible indirect measure estimate from Sample I is that 32% [11%, 54%] of Americans do not believe in God, Figure 1.

Sample II included a conceptual replication effort of Sample I's indirect estimate by comparing the baseline and critical conditions. Sample II also included an additional condition assessing validity of the indirect count technique by comparing the baseline and mathematical impossibility conditions.

Sample II yielded an indirect atheism rate estimate of 20% [6%, 35%], Figure 1. This atheism estimate is lower than that in Sample I. Speculatively, this difference may reflect (among other things) a difference in how participants respond to positive versus negative framing of the unmatched count tasks. That is, Sample II primarily differed from Sample I in that it included a positive affirmation of atheism (agreeing with the statement "I do not believe in God") rather than a more passive denial of theism as in Sample I. In an attempt to assess the validity of the unmatched count technique, we tested a second model in which we included a single religious demographic contrast. Participants in Sample II indicated which of several religious identities they identified with (e.g., Protestant, Catholic, atheist, etc.). In a model comparing self-described atheists to all other religious identities, the mean difference between the critical and baseline conditions was 1.03 [.526, 1.567] for atheists and .134 [-.013, .286] for others. Put differently, the most credible indirect atheism prevalence estimate was almost exactly 100% among self-described atheists, but 13% among those not identifying as atheists. This demographic split provides some validity evidence for the technique.

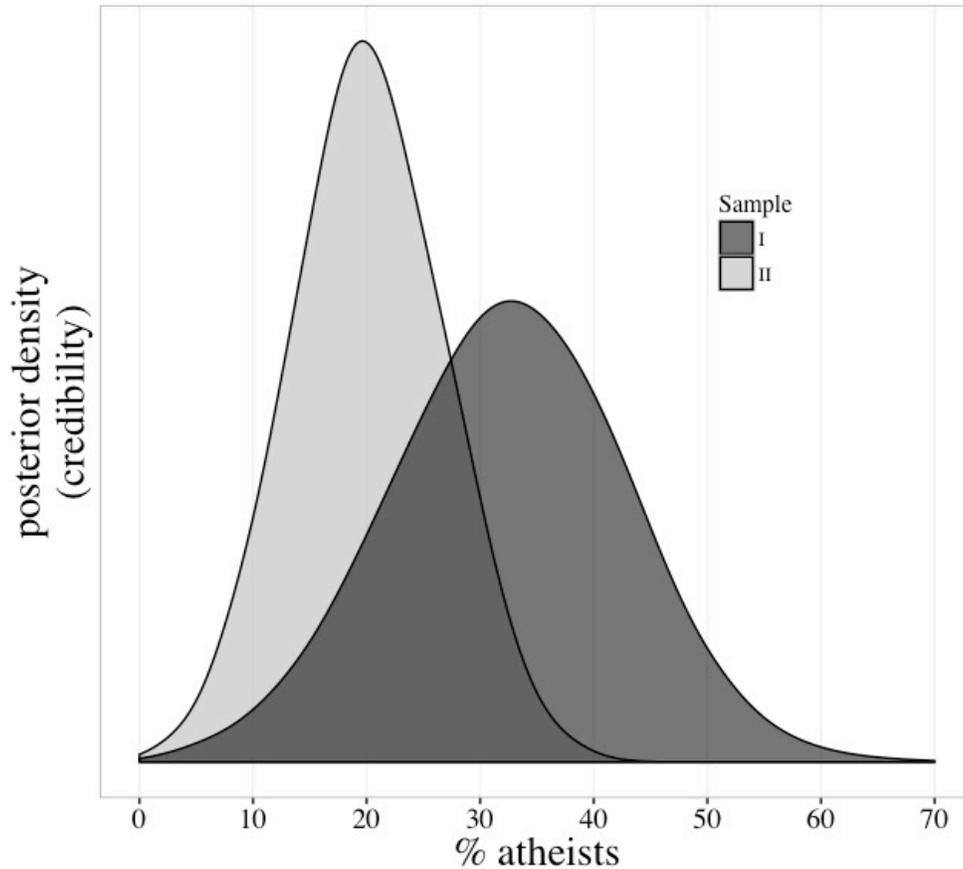


Figure 1. Posterior atheist prevalence estimates from Samples I and II. Y-axis represents the relative credibility with which different parameter values could have plausibly generated the observed data. Values higher on the y-axis represent more plausible parameter estimates; values lower on the y-axis represent less plausible parameter estimates.

Sample II included a third condition in an attempt to further gauge the validity of the unmatched count technique. Similar to previous work (Coffman et al., 2016) we included an additional condition in which the added unmatched count item was ostensibly not socially sensitive. Thus, we compared the baseline condition to a condition in which the additional “sensitive” item was endorsement of a mathematical impossibility. Rather surprisingly, and in contrast to previous unmatched count validation (Coffman et al.,

2016), we observed a reliable difference between these two conditions, with people indicating more statements true of them in the mathematical impossibility condition ($M = 3.62$, $SD = 1.16$) than in the baseline condition ($M = 3.27$, $SD = 1.15$), mean difference = .340 [.198, .477]. This result is, frankly, bizarre and we are hesitant to speculate a great deal about its causes. That said, we give it further treatment in the General Discussion.

Aggregate Analysis

Analysis plan. Our critical analysis pooled data from the baseline and atheism conditions of both samples. We then tested a single aggregate model to provide our most precise and comprehensive overall atheism prevalence estimate from both samples. Our final model was thus a hierarchical (i.e., multilevel) model evaluating the pooled data. We also included demographic variables common to both samples. In this model, the intercept was modeled as random across samples, but the slopes of condition, demographic splits, and their interactions were modeled as fixed across samples. This is functionally equivalent to conducting a meta-analysis of both samples (Vuorre, 2017) with demographic moderators. We generated both indirect and self-report atheism prevalent estimates for the overall sample, as well as across demographic splits of gender, education (comparing those with an education beyond high school to those with no education beyond high school), politics (Democrat, Independent/Other, Republican), and age (treated continuously, but illustrated as the difference between Millennials and Baby Boomers).

Overall estimates. Our aggregate analysis, pooling across samples, provided an indirect atheism prevalence rate of 26% [13%, 39%]. Unsurprisingly, this estimate is intermediate between both samples' individual point estimates, but with a tighter range of

plausible values than either alone. Table 3 provides a full summary of both indirect and self-reported atheism rates, both overall and across demographic splits. It also provides the posterior probability that indirect estimates are higher than self-reports for each comparison.

Our primary goal was to generate an overall indirect estimate of atheist prevalence in the USA. Our overall estimate of around 26% is substantially higher than the 11% Gallup estimates often cited in work on the psychology of religion and atheism, and reliably higher than our self-report estimate of around 17% [14%, 21%], posterior probability = .93. Though self-reported atheism rates were higher than recent Gallup and Pew estimates relying on telephone polls, this difference may be attributable to the heightened social desirability pressures inherent to telephone polling, consistent with other nationally representative religion polls using both telephone and online computer polling (Cox et al., 2014).

Figure 2 depicts the full posterior for our overall indirect atheism prevalence estimate. The y-axis of Figure 2 essentially illustrates how credible a given parameter estimate is. Inspection of this posterior enables us to make a number of statements regarding the relative credibility of a number of different estimates (Wagenmakers et al., 2016). Based on our results, we can state with around 99% certainty that more than 11% of Americans are atheists (as per Gallup's estimate), and with around 93% certainty that more than 17% of Americans are atheists (our self-report estimate). The self-report estimate of around 17% atheists is roughly as credible as there being 35% atheists, and all intermediate values are more credible. In addition, the possibility that 26% of Americans are atheists is about 3 times as credible as the possibility that only 17% of Americans are

atheists and more than 20 times as credible as the possibility that only 11% of Americans are atheists.

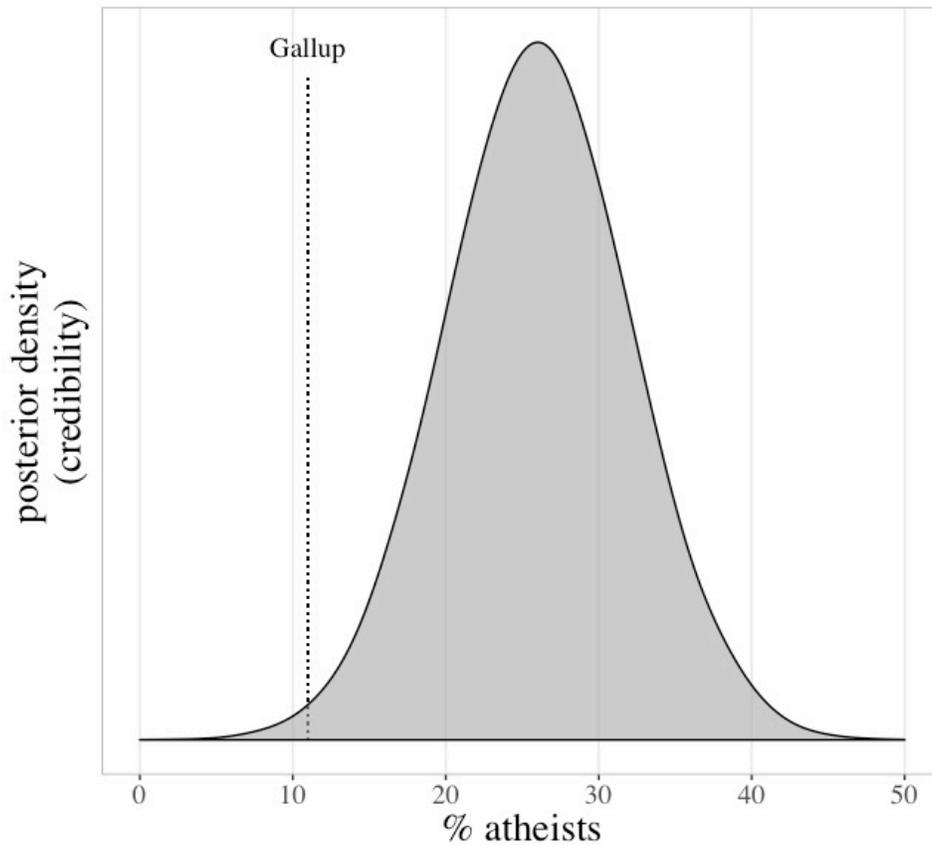


Figure 2. Posterior atheist prevalence estimate from pooled model. Y-axis represents the relative credibility with which different parameter values could have plausibly generated the observed data. Values higher on the y-axis represent more plausible parameter estimates; values lower on the y-axis represent less plausible parameter estimates.

These results have two key implications. First, that accounting for socially desirable responding, roughly 26% of American adults may actually be atheists, to the extent that the unmatched count technique captures actual population prevalence. Second, that roughly one in three atheists may not openly acknowledge their disbelief in an

anonymous online survey; the most credible indirect estimate is more than 50% higher than the most credible self-report estimate (Table 3).

Table 3. Indirect and self-report atheist prevalence estimates. Point estimates [and 97% HPDIs] are presented for the overall estimate and each demographic breakdown. $\text{Pr}(I > \text{SR})$ refers to the posterior probability that indirect measures produce higher estimates than do self-reports.

	Indirect %		Self-report %		Pr(I > SR)
Overall	26	[14, 40]	17	[14, 20]	0.93
Gender					
Female	24	[5, 41]	13	[9, 17]	0.91
Male	28	[9, 47]	22	[17, 28]	0.73
Education					
HSMax	20	[0, 40]	12	[8, 17]	0.80
HSPlus	30	[13, 47]	22	[17, 27]	0.84
Politics					
Democrat	30	[10, 53]	24	[18, 30]	0.74
Independent	39	[18, 61]	19	[13, 24]	0.98
Republican	0	[0, 26]	8	[4, 13]	0.27
Age Cohort					
Millennial (1982)	26	[7, 42]	21	[16, 25]	0.72
Baby Boomer (1952)	26	[9, 44]	14	[10, 18]	0.93

Exploratory demographic comparisons. Table 3 illustrates several potential convergences and divergences between indirect and self-report atheism prevalence estimates across demographics. Although the point estimates and HPDIs provide useful summaries, graphical inspection of posterior distributions provides for a richer set of inferences. To facilitate such inferences, we prepared an online widget that allows readers to view posterior distributions of both indirectly measured and self-reported atheist prevalence across all demographic comparisons, https://willgervais.shinyapps.io/atheist_rate/. Matching known demographic patterns, indirect atheism estimates were higher among more educated respondents and much

higher among Democrats than among Republicans. Indirect measurements and self-reports were especially discrepant among Baby Boomers, political Independents, and women. On the other hand, indirect measurements and self-reports were quite similar among Millennials, Republicans, and men.

Of potential note, two large demographic gaps in self-reports are largely absent from indirect measurements. First, men self-report atheism at a rate 77% higher than women, but are only 16% higher on indirect measurement. Second, self-reports reveal a large generational difference in atheism between Millennials and Baby Boomers—consistent with previous research—but this generational gap essentially disappears when atheism is measured indirectly. Interestingly, recent work using the unmatched count technique to indirectly assess willingness to vote for underrepresented groups (e.g., women, African Americans, Muslims) reveals a similar pattern whereby self-reports and indirect measures diverge far more among Baby Boomers than among Millennials (Brown-Iannuzzi, Najle, & Gervais, under review). We note again that the demographic comparisons were not of primary interest and were exploratory in nature. Nonetheless, these patterns may warrant additional focused investigation.

Table 4. Raw descriptive statistics for indirect measures

	<i>N</i>	<i>M</i>	<i>SD</i>
Sample I			
Atheist	696	5.81	1.88
Baseline	635	6.13	1.80
Sample II			
Atheist	619	3.27	1.15
Baseline	696	3.48	1.26
Math	685	3.62	1.16

Finally, for the sake of transparency, we include all pertinent descriptive statistics for indirect measures in Table 4, and all data are freely available at <https://osf.io/4q85g/>

Discussion

Existing nationally representative polls indicate that atheist prevalence is relatively low in the United States, perhaps only 3% (Pew, 2015) to 11% (Gallup, 2015). Given the heavy stigmatization of atheism (Edgell et al., 2006), we hypothesized that many atheists might be reluctant to disclose their disbelief to pollsters. We therefore deployed two nationally representative samples in an attempt to indirectly measure atheist prevalence using the unmatched count technique (Raghavarao & Federer, 1979). These indirect measures suggest that roughly one in four (26%) American adults may be atheists—2.4 to 8.7 times as many as telephone polls (Gallup, 2015; Pew, 2015) suggest. This implies the existence of potentially more than 80 million American atheists. The disparity between self-report and indirectly measured atheism rates underscores the potent stigma faced by atheists (Edgell et al., 2006; Gervais, 2013), as even in an anonymous online survey, about a third of American atheists may be effectively “closeted,” even in anonymous telephone polls.

Despite our observed discrepancy between directly and indirectly measured atheism rates, social pressures reinforcing religiosity in the USA may be weakening over time. Recent decades have seen increasing secularism in large parts of the world, driven in part by economic modernization and increased existential security (Inglehart & Norris, 2004). Against this backdrop, the USA stands as an apparent outlier: both highly religious and economically advanced. It is possible that general worldwide patterns in secularism are evident in the USA, albeit across generations, as younger generations

report drastically lower levels of religious belief than their predecessors, showing a cohort difference similar to that behind the secularization of modern countries (e.g., Twenge, Exline, Grubbs, Sastry, & Campbell, 2015). Our data offer an admittedly speculative complement to this work. We replicated generational differences in self-reported atheism rates between Millennials (21%) and Baby Boomers (14%), but this difference almost entirely vanished in indirect measurements (both around 26%). Put differently, our data suggest a greater than 99% probability that there are more atheists among Millennials on self-reports; when measured indirectly, this probability drops to about chance (48%). Thus, it is possible that the apparent generational atheism gap exists less at the level of disbelief in a god, but more at the level of willingness to “out” oneself as a nonbeliever in an anonymous poll. By extension, the view that the USA is a global outlier to broader secularization trends may be less tenable than assumed (see also Stark & Bainbridge, 1985).

Evaluating the Indirect Measure

Our indirect measurements relied on the unmatched count technique (Raghavarao & Federer, 1979). This task has been widely used, and its validity is typically inferred from the fact that it usually (though not universally: Coutts & Jann, 2011) reveals higher prevalence estimates of socially sensitive topics than do self-reports. Further, it appears that the unmatched count typically only diverges from self-reports of socially undesirable attributes (Coffman et al., 2016). We included one condition aimed at further assessing the validity of the technique. In Sample II, we ran an unmatched count design in which the sensitive item was endorsement of a mathematical impossibility (“I do not believe that $2 + 2$ is less than 13”). Although we hypothesized that the unmatched count would

return a prevalence estimate of essentially zero for this item, it bizarrely suggested a rate of about one third (34%). Without a doubt, this is our *most damning result* (cf. Vazire, 2016). It may reflect any combination of genuine innumeracy, incomprehension of an oddly phrased item, participant inattentiveness or jesting, sampling error, or a genuine flaw in the unmatched count technique. Fortunately, we were also able to assess validity in a second way. In Sample II, the unmatched count to generated an atheist prevalence estimate of almost exactly 100% among self-described atheists, but only 13% among all other religious identifications. It is unlikely that a genuinely invalid method would track self-reported atheism this precisely. Across two assessment attempts our validity evidence was a mixed bag. This perhaps suggests that future researchers should attempt to—as we were able in Sample II but not Sample I—include diagnostic self-reports alongside the unmatched count to assess validity. And, as the present estimates are only as strong as the method that generated them, they should be treated with some caution. In our view—given heavy social pressures to be or appear religious—the 11% atheism prevalence estimates derived solely from telephone self-reports is probably untenable. Does this imply that our most credible estimate of 26% should be uncritically accepted instead? Of course not. The present two nationally representative samples merely provide additional estimates using a different technique, and our model suggests a wide range of relatively credible estimates. We hope that future work using a variety of direct and indirect measures will provide satisfactory convergence across methods, and the present estimates are merely an initial indirect measurement data point to be considered in this ongoing scientific effort.

Limits to Generalizability and Atheism Around the Globe

Anti-atheist prejudice remains prevalent in the United States (Edgell et al., 2006), but is not exclusively an American phenomenon (Norenzayan & Gervais, 2015), suggesting that atheism may be similarly underreported elsewhere in the world where social pressures reinforce religiosity and its self-presentation. Self-reports yield an estimate of 500-700 million atheists worldwide (Zuckerman, 2007). Presumably, self-reported atheism is less biased by social desirability concerns in more highly secularized societies like those in Scandinavia (Inglehart & Norris, 2004). Thus, it would be unwise to use the present USA data as a baseline and assume that atheism rates are uniformly much higher than self-reports suggest. Instead, we predict that underreporting of atheism covaries with cultural norms promoting religion in domains like morality and cooperation (McKay & Whitehouse, 2014; Norenzayan et al., 2014). Precise global estimates are not currently possible, but we speculate that there may be actually be around two billion atheists worldwide. If true, some currently popular theoretical approaches for understanding the bases of religious cognition require substantial modification or outright abandonment.

Coda

Religion is cross-culturally universal (Boyer, 2008; Inglehart & Norris, 2004), but also highly variable across individuals (Norenzayan & Gervais, 2013) and societies (Inglehart & Norris, 2004). Theories of religion must accommodate and explain patterns of belief and disbelief. Some models view religion primarily as a reliably developing byproduct of cognitive adaptations serving other purposes, and describe atheism as rare and both cognitively unnatural and cognitively effortful (e.g., Boyer, 2008)—perhaps even psychologically superficial (Bering, 2010). Strong interpretations of these claims are

plausible primarily if atheism is a somewhat rare aberration. The prevalence of atheism—vaster than previously assumed, according to the present data, and likely deliberately concealed in large parts of the world—challenges these predictions. On the other hand, models viewing religious belief and disbelief as arising from a complex combination of factors (Geertz & Markússon, 2010; Norenzayan & Gervais, 2013)—including cognitive constraints and style, motivation, and especially cultural learning—can accommodate a much wider range of atheism rates. As basic theory on the cognitive, cultural, and evolutionary origins of religious belief advances, it is necessary to consider—and correct for—social pressures that probably distort the very data central to theory development and testing.

Finally, the present results may have considerable societal implications. Preliminary research suggests that learning about how common atheists actually are reduces distrust of atheists (Gervais, 2011). Thus, obtaining accurate atheist prevalence estimates may help promote trust and tolerance of atheists—potentially 80+ million people in the USA and well over a billion worldwide.

References

- Bering, J. M. (2010). Atheism is only skin deep: Geertz and markússon rely mistakenly on sociodemographic data as meaningful indicators of underlying cognition. *Religion, 40*(3), 166-168. doi:10.1016/j.religion.2009.11.001
- Boyer, P. (2008). Being human: Religion: Bound to believe? *Nature, 455*(7216), 1038-1039.
- Brown-Iannuzzi, J. L., Najle, M. B., & Gervais, W. M. (under review). The illusion of political tolerance: Social desirability and self-reported voting preferences. *Journal of Experimental Social Psychology*.
- Coffman, K. B., Coffman, L. C., & Ericson, K. M. M. (2016). The size of the lgbt population and the magnitude of antigay sentiment are substantially underestimated. *Management Science*.
- Coutts, E., & Jann, B. (2011). Sensitive questions in online surveys: Experimental results for the randomized response technique (rrt) and the unmatched count technique (uct). *Sociological Methods & Research, 40*(1), 169-193.
- Cox, D., Jones, R. P., & Navarro-Rivera, J. (2014). *I know what you did last sunday: Measuring social desirability bias in self-reported religious behavior, belief, and identity*. Retrieved from Washington, DC: <http://publicreligion.org/site/wp-content/uploads/2014/05/AAPOR-2014-Final.pdf>
- Dalton, D. R., Wimbush, J. C., & Daily, C. M. (1994). Using the unmatched count technique (uct) to estimate base rates for sensitive behavior. *Personnel Psychology, 47*(4), 817.

- Edgell, P., Gerteis, J., & Hartmann, D. (2006). Atheists as “other”: Moral boundaries and cultural membership in american society. *American Sociological Review*, *71*(2), 211-234.
- Gallup. (2015). *Religion*. Retrieved from <http://www.gallup.com/poll/1690/Religion.aspx>
- Geertz, A. W., & Markússon, G. I. (2010). Religion is natural, atheism is not: On why everybody is both right and wrong. *Religion*, *40*(3), 152-165.
doi:10.1016/j.religion.2009.11.003
- Gervais, W. M. (2011). Finding the faithless: Perceived atheist prevalence reduces anti-atheist prejudice. *Personality and Social Psychology Bulletin*, *37*(4), 543-556.
- Gervais, W. M. (2013). In godlessness we distrust: Using social psychology to solve the puzzle of anti - atheist prejudice. *Social and Personality Psychology Compass*, *7*(6), 366-377.
- Gervais, W. M. (2014a). Everything is permitted? People intuitively judge immorality as representative of atheists. *PLoS One*, *9*(4), e92302.
doi:10.1371/journal.pone.0092302
- Gervais, W. M. (2014b). Good for god? Religious motivation reduces perceived responsibility for and morality of good deeds. *Journal of Experimental Psychology: General*, *143*(4), 1616.
- Hadaway, C. K., Marler, P. L., & Chaves, M. (1993). What the polls don't show: A closer look at us church attendance. *American Sociological Review*, 741-752.
- Inglehart, R., & Norris, P. (2004). *Sacred and secular: Religion and politics worldwide*. Cambridge: Cambridge University Press.

- Kruschke, J. K. (2010). *Doing bayesian data analysis: A tutorial introduction with r*. Academic Press.
- McElreath, R. (2016). *Statistical rethinking: A bayesian course with examples in r and stan* (Vol. 122): CRC Press.
- McKay, R., & Whitehouse, H. (2014). Religion and morality. *Psychological Bulletin*.
- Norenzayan, A., & Gervais, W. M. (2013). The origins of religious disbelief. *Trends in Cognitive Sciences*, 17(1), 20-25.
- Norenzayan, A., & Gervais, W. M. (2015). Secular rule of law erodes believers' political intolerance of atheists. *Religion, Brain & Behavior*, 5(1), 3-14.
- Norenzayan, A., Shariff, A. F., Gervais, W. M., Willard, A. K., McNamara, R. A., Slingerland, E., & Henrich, J. (2014). The cultural evolution of prosocial religions. *Behavioral and Brain Sciences*, 1-86.
- Pew. (2015). *America's changing religious landscape*. Retrieved from
- Raghavarao, D., & Federer, W. T. (1979). Block total response as an alternative to the randomized response method in surveys. *Journal of the Royal Statistical Society. Series B (Methodological)*, 40-45.
- Sedikides, C., & Gebauer, J. E. (2009). Religiosity as self-enhancement: A meta-analysis of the relation between socially desirable responding and religiosity. *Personality and Social Psychology Review*.
- Stark, R., & Bainbridge, W. S. (1985). *The future of religion: Secularization, revival, and cult formation*: Univ of California Press.

- Twenge, J. M., Exline, J. J., Grubbs, J. B., Sastry, R., & Campbell, W. K. (2015). Generational and time period differences in american adolescents' religious orientation, 1966–2014. *PloS one*, *10*(5), e0121454.
- Vazire, S. (2016). Most damning result. Retrieved from <http://sometimesimwrong.typepad.com/wrong/2015/12/most-damning-result.html>
- Vuorre, M. (2017). Meta-analysis is a special case of bayesian multilevel modeling | matti vuorre. Retrieved from <https://mvuorre.github.io/post/2016/2016-09-29-bayesian-meta-analysis/>
- Wagenmakers, E.-J., Morey, R. D., & Lee, M. D. (2016). Bayesian benefits for the pragmatic researcher. *Current Directions in Psychological Science*, *25*(3), 169-176.
- Wilson, E. O. (1999). *Consilience: The unity of knowledge* (Vol. 31): Vintage.
- Zuckerman, P. (2007). Atheism: Contemporary numbers and patterns.

Acknowledgements

This research was supported by a grant to WMG from the John Templeton Foundation (48275). The content is solely the responsibility of the authors and does not necessarily represent the official views of JTF. We thank Drs. Jazmin Brown-Iannuzzi and Azim Shariff for feedback.

Supplementary Information
How many atheists are there?
Gervais & Najle

Core demographics	29
Sampling details.....	30
Priors: Justification and alternatives	31

Core demographics

Measure	%		Category
	<i>Sample I</i>	<i>Sample II</i>	
Gender	53.3	53.5	Male
	46.8	46.5	Female
Race	1.8	2.05	Asian
	11.1	11.55	Black
	9.5	9.25	Hispanic
	0.1	0.3	Middle Eastern
	2.0	2.2	Mixed
	1.9	1.15	Other
	0.8	0.9	Native American
	72.8	72.6	White
Education	5.5	3.65	< High school
	37.4	33.75	High school
	21.1	24.5	some college
	8.8	9.4	2 year college
	17.4	18.35	4 year college
	9.8	10.35	post-grad
Family Income	16.9	15.05	< 20k
	22.7	21.65	20-40k
	17.4	16.2	40-60k
	11.1	12.9	60-80k
	6.2	7.25	80-100k
	12.4	12.4	>100k
	13.0	14.55	no response
Political Affiliation	38.6	38.6	Democrat
	27.1	28.55	Independent
	25.0	24.9	Republican
	2.6	3.2	Other
Marital Status	6.8	4.75	Unsure
	49.2	50.1	Married
	4.7	4.55	Domestic Partnership
	11.0	10.55	Divorced
	1.9	1.8	Separated
	28.3	27.15	Single
	4.9	5.35	Widowed

Sampling details

YouGov provided the following sampling information.

Sample I:

“YouGov interviewed 2037 respondents, who were then matched down to a sample of 2000 to produce the final dataset. The respondents were matched to a sampling frame on gender, age, race, education, party identification, ideology, and political interest. The frame was constructed by stratified sampling from the full 2010 American Community Survey (ACS) sample with selection within strata by weighted sampling with replacements (using the person weights on the public use file). Data on voter registration status and turnout were matched to this frame using the November 2010 Current Population Survey. Data on interest in politics and party identification were then matched to this frame from the 2007 Pew Religious Life Survey. The matched cases were weighted to the sampling frame using propensity scores. The matched cases and the frame were combined, and a logistic regression was estimated for inclusion in the frame. The propensity score function included age, gender, race/ethnicity, years of education, and ideology. The propensity scores were grouped into deciles of the estimated propensity score in the frame and post-stratified according to these deciles.”

Sample II:

“YouGov interviewed 2217 respondents who were then matched down to a sample of 2000 to produce the final dataset. The respondents were matched to a sampling frame on gender, age, race, education, party identification, ideology, and political interest. The frame was constructed by stratified sampling from the full 2010 American Community Survey (ACS) sample with selection within strata by weighted sampling with replacements (using the person weights on the public use file). Data on voter registration status and turnout were matched to this frame using the November 2010 Current Population Survey. Data on interest in politics and party identification were then matched to this frame from the 2007 Pew Religious Life Survey. The matched cases were weighted to the sampling frame using propensity scores. The matched cases and the frame were combined and a logistic regression was estimated for inclusion in the frame. The propensity score function included age, gender, race/ethnicity, years of education, and ideology. The propensity scores were grouped into deciles of the estimated propensity score in the frame and post-stratified according to these deciles.”

Priors: Justification and alternatives

Although relatively precise data are available for self-reported belief in God in telephone polls, there is considerably less data available for online nationally representative polls, and none for nationally representative polls using the unmatched count technique. Thus, we adopted a conservative approach and used weakly regularizing priors for our analyses. The main indirect measure model is summarized below:

$$\text{belief}_i \sim \text{Normal}(\mu, \sigma)$$

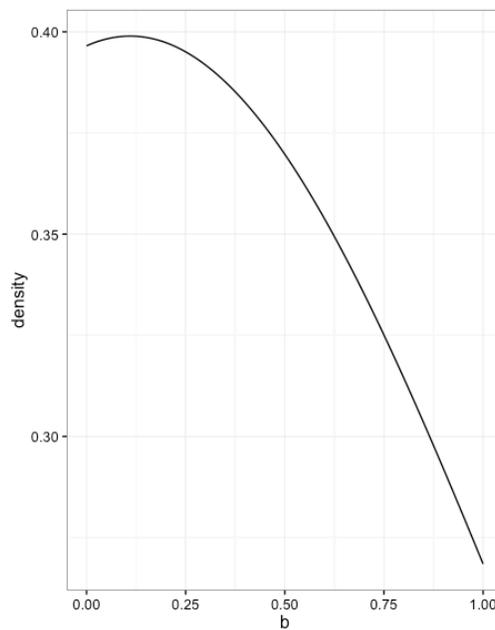
$$\mu_i = \alpha + \beta x_i$$

$$\alpha \sim \text{Uniform}(0, 10)$$

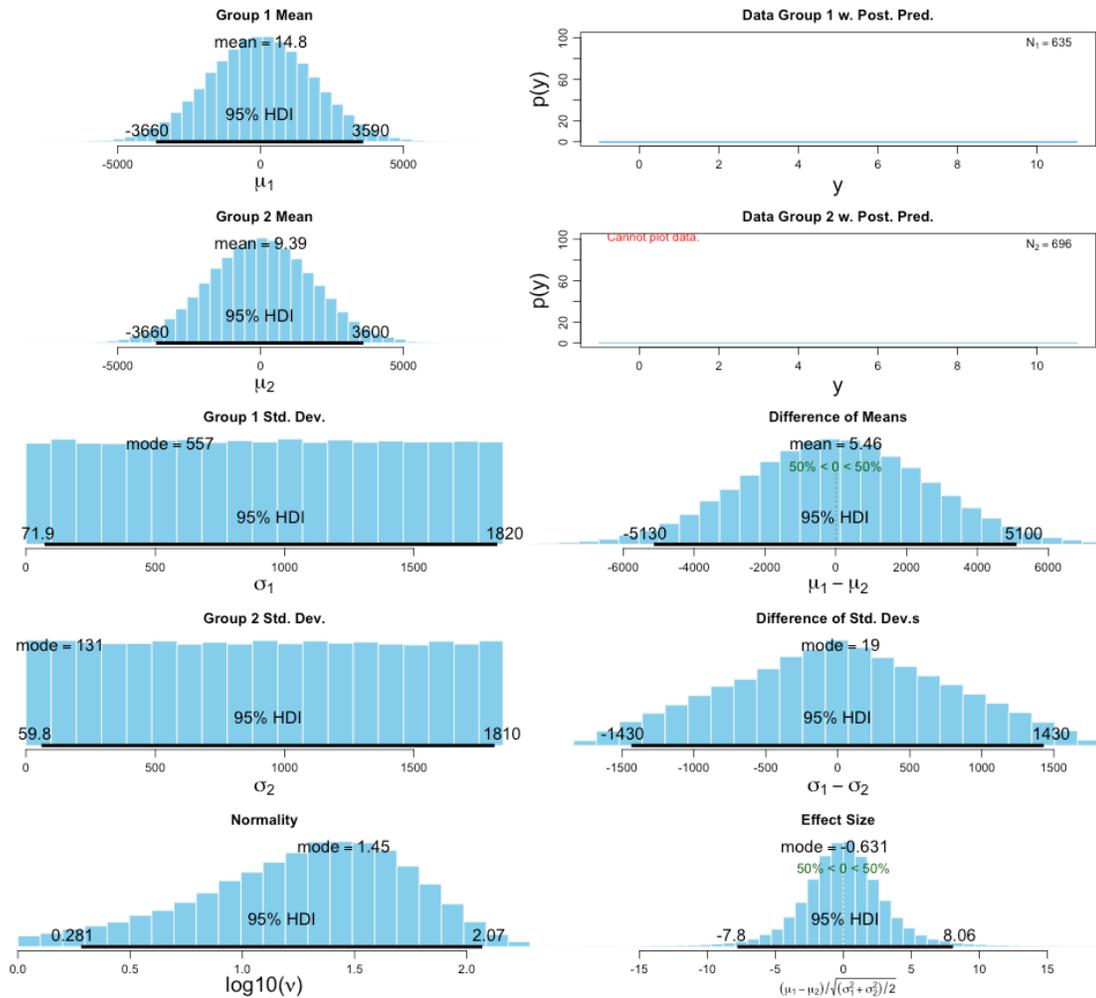
$$\beta \sim \text{Normal}(.11, 1)$$

$$\sigma \sim \text{HalfCauchy}(0, 2)$$

The prior for β , which is our indirect atheism rate estimate, was centered on Gallup's self-report atheism rate of 11%, and set to weakly regularize around this value.



Finally, we note that our initially preregistered analysis plan varied slightly from the one we eventually adopted. We preregistered analyses using the BEST package in R before the publication and reading of McElreath’s *Statistical Rethinking* book, which formed the basis for our analyses. For the sake of transparency, here are inferences from Sample I using the preregistered analyses from the BEST package. The default settings implement broad priors (Kruschke, 2013) designed to have minimal influence on posterior distributions. These include priors for the means and standard deviations of both conditions, as well as normality of data. These are the default settings for the BEST package in R. The following figure illustrates these priors:



Priors on all parameters are very broad, and are quickly washed out by data as it accumulates.

The full posterior produced by this model follows. Inferences do not appreciably change using this alternative analysis framework.

