Supplementary Materials
The Moralization Bias of Gods' Minds: A Cross-Cultural Test

Benjamin Grant Purzycki[a,*], Aiyana K. Willard[b], Eva Kundtová Klocová[c], Coren Apicella[d], Quentin Atkinson[e,f], Alexander Bolyanatz[g], Emma Cohen[h], Carla Handley[i], Joseph Henrich[j], Martin Lang[c], Carolyn Lesorogol[k], Sarah Mathew[i], Rita A. McNamara[l], Cristina Moya[m], Ara Norenzayan[n], Caitlyn Placek[o], Montserrat Soler[p], Tom Vardy[q], Jonathan Weigel[r], Dimitris Xygalatas[s,t], Cody T. Ross[u]

[a] *Aarhus University, DK*
[b] *Brunel University, UK*
[c] *LEVYNA, Masaryk University, CZ*
[d] *Department of Psychology, University of Pennsylvania, USA*
[e] *Department of Psychology, University of Auckland, NZ*
[f] *Max Planck Institute for the Science of Human History, DE*
[g] *College of DuPage, USA*
[h] *Wadham College, University of Oxford, UK*
[i] *Arizona State University, USA*
[j] *Harvard University, USA*
[k] *Washington University in St. Louis, USA*
[l] *School of Psychology, Victoria University of Wellington, NZ*
[m] *Department of Anthropology, University of California-Davis, USA*
[n] *University of British Columbia, CA*
[o] *Department of Anthropology, Ball State University, USA*
[p] *Ob/Gyn and Women's Health Institute Cleveland Clinic, USA*
[q] *Department of International Development, London School of Economics, UK*
[r] *School of Psychology, University of Auckland, NZ*
[s] *Department of Anthropology, University of Connecticut, USA*
[t] *Department of Psychological Sciences, University of Connecticut, USA*
[u] *Max Planck Institute for Evolutionary Anthropology, DE*

Contents

*Corresponding author
Email address: bgpurzycki@cas.au.dk (Benjamin Grant Purzycki)

S1. Descriptives

The target gods' moral interest questions we asked took the form of: *How important is punishing X to deity Y?* where X was either *liars*, *thieves*, or *murderers* and Y was either the moralistic or local deity selected for each field site. Answers to these questions were on a scale of 0 to 4 (0: not important at all; 1: a little important; 2: important; 3: very important; 4: the most important thing). Figures S1 and S2 present the raw frequencies of responses across our field sites.

S2. Main Analysis

S2.1. Main model structure

We fit all models using the `rstan` package (Stan Development Team, 2017) and `R` (R Core Team, 2016). We use ordered-logistic regressions to model outcomes for deities' moral concerns. In each population, j, we asked individuals, indexed by i, a set of three questions (concerning murder, theft, and lying), indexed by q, about each of two deities (moralistic and local), indexed by d. In each case, the outcome, $K_{[i,q,d]}$, is an ordered categorical response. As such, we model them using an ordered categorical likelihood model with a vector of random cut-points, C:

$$K_{[i,q,d]} \sim \text{Ordered Categorical}(\xi_{[i,q,d]}, C) \tag{1}$$

The full linear model for $\xi_{[i,q,d]}$ is given by:

$$\xi_{[i,q,d]} = \beta_{[j(i),q,d,1]} + \beta_{[j(i),q,d,2]} A_{[i]} + \beta_{[j(i),q,d,3]} M_{[i]} + \beta_{[j(i),q,d,4]} E_{[i]} + \beta_{[j(i),q,d,5]} \Phi_{[j(i),K_{[i,q,1]}]} \delta_{[i,q,d]} \tag{2}$$

where: $j(i)$ gives the site of individual i, $A_{[i]}$ is the age of individual i (centered by site), $M_{[i]}$ is a variable indicating if individual i is male, and $E_{[i]}$ is the years of formal education completed by individual i (centered by site). In the last term, we see that $K_{[i,q,1]}$ is the response that individual i gave about his or her moralistic deity's concern in question q. The symbol $\Phi_{[j]}$ denotes a site-specific, positive, ordered, parameter vector whose last element is equal to 1. We use these parameters to account for the fact that ordered categories cannot be directly interpreted as numerical values (see McElreath, 2020, pp. 391-396). In the last term, $\delta_{[i,q,d]}$, is a binary indicator for if the question was

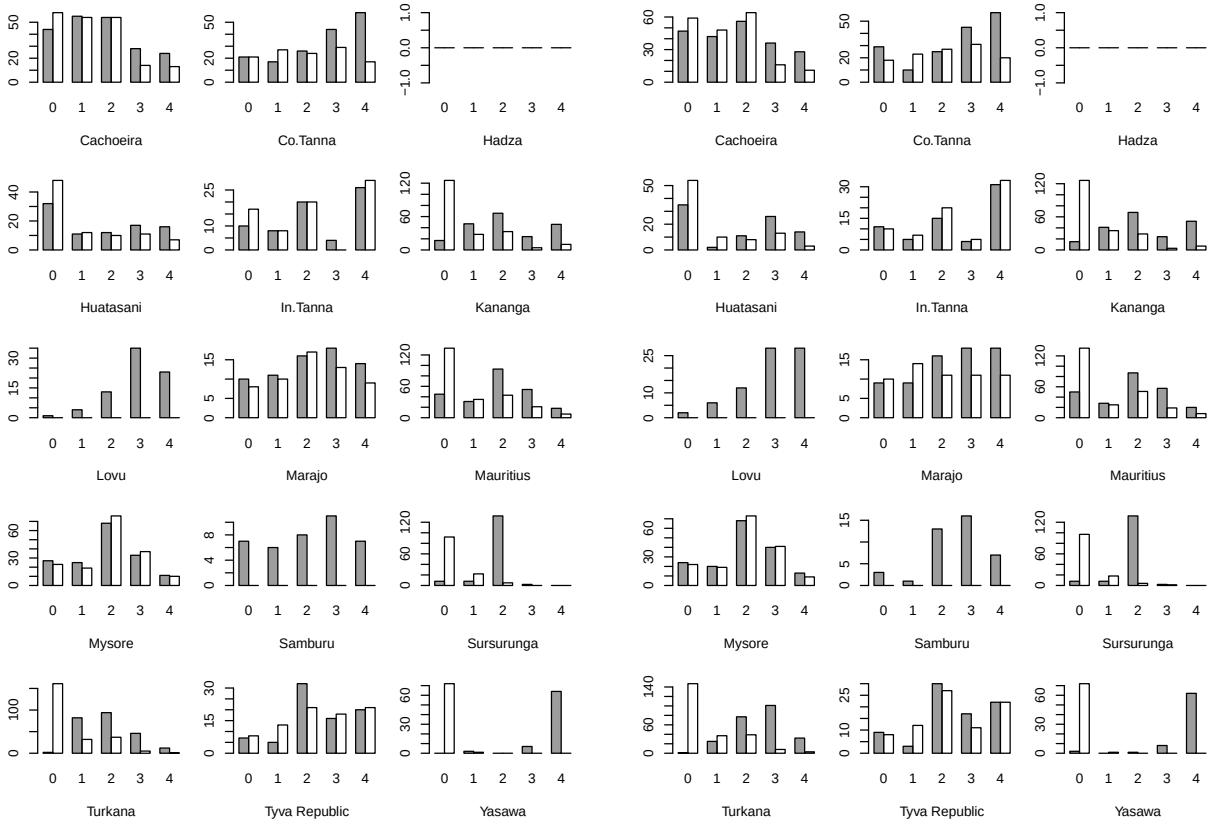

Figure S1: Barplots of liars (left) and thieves (right) questions for moralistic (grey) and local (white) deities across 15 field sites. Note that y-axes are different across field sites.

about the local deity—this ensures that the last term only enters the regression model for outcomes about the concerns of local deities.

We let $\Phi_{[j]}$ be defined as the cumulative sum of a corresponding unit simplex $\phi_{[j]}$:

$$\Phi_{[j,n]} = \sum_{m=1}^{n} \phi_{[j,m]} \tag{3}$$

The coefficient $\beta_{[j(i),q,d,5]}$ then measures the offset to the intercept if individual i responded with $K_{[i,q,1]}$ equal to the maximal category. If individual i responded with $K_{[i,q,1]}$ equal to some other category, then the intercept will be offset by some fraction of $\beta_{[j(i),q,d,5]}$, as given by the element of $\Phi_{[j(i)]}$ corresponding to that $K_{[i,q,1]}$. This model reduces to a standard regression model like $\xi_{[i,q,d]} = ... + \beta_{[j(i),q,d,5]} K_{[i,q,1]}$ in the special case that the distance between categories—i.e., the elements of $\Phi_{[j]}$—is uniform. Otherwise, it accounts for the fact that ordered categories are ordered, but with a potentially non-uniform distance between categories. Later, in Fig. S3, we plot the estimated shape of $\Phi_{[j]}$ on a site-by-site basis.

S2.2. Priors

In each population, j, we define the simplex parameter vectors, $\phi_{[j]}$, using a Dirichlet distribution:

$$\phi_{[j]} \sim \text{Dirichlet}(1, \ldots, 1) \tag{4}$$

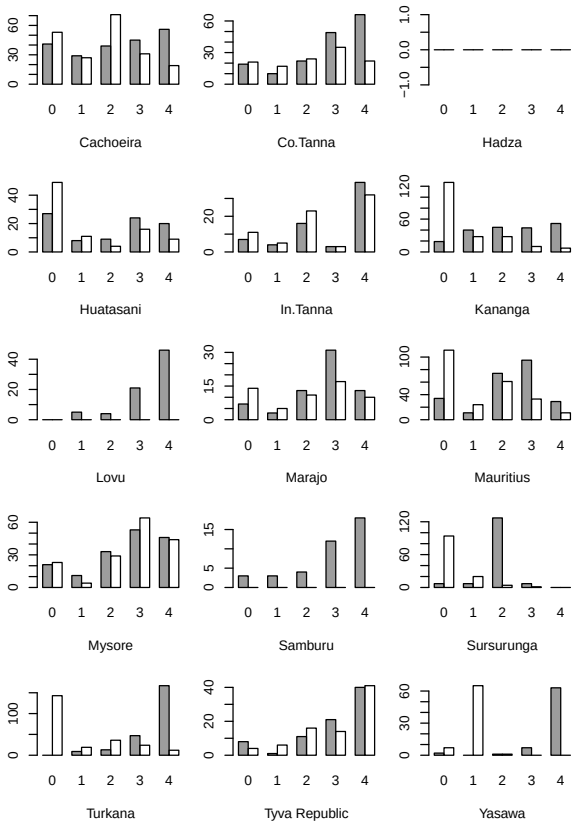

Figure S2: Barplots of murderers questions for moralistic (grey) and local (white) deities across 15 field sites. Note that y-axes are different across field sites.

The main regression parameter vectors are defined using a multi-level model for each deity type, d, that partially pools over sites, j, and questions, q:

$$\beta_{[j,q,d]} = \mu_{[d]} + \sigma_{[d]} \circ (L_{[d]} * \alpha_{[j,q,d]}) \tag{5}$$

The symbol $\circ$ denotes an element-wise product. The parameters $\mu_{[d]}$ and $\sigma_{[d]}$ are vectors, and their elements, indexed by p, have the priors:

$$\mu_{[d][p]} \sim \text{Normal}(0, 5) \tag{6}$$

$$\sigma_{[d][p]} \sim \text{Cauchy}(0, 2.5) \tag{7}$$

The parameters $L_{[d]}$ are Cholesky factors from the decomposition of correlation matrices, and they have priors:

$$L_{[d]} \sim \text{LKJ Corr Cholesky}(0, 2.5) \tag{8}$$

The parameters $\alpha_{[j,q,d]}$ are vectors, and their elements, indexed by p, have the priors:

$$\alpha_{[j,q,d][p]} \sim \text{Normal}(0, 1) \tag{9}$$

Finally, the cut-points vector, C, is given an implicit uniform prior over its support.

S2.3. Conceptual bleeding across deities

Fig. S3 shows our estimates of $\Phi_{[j]}$, the scalar to $\beta_{[j(i),q,d,5]}$.

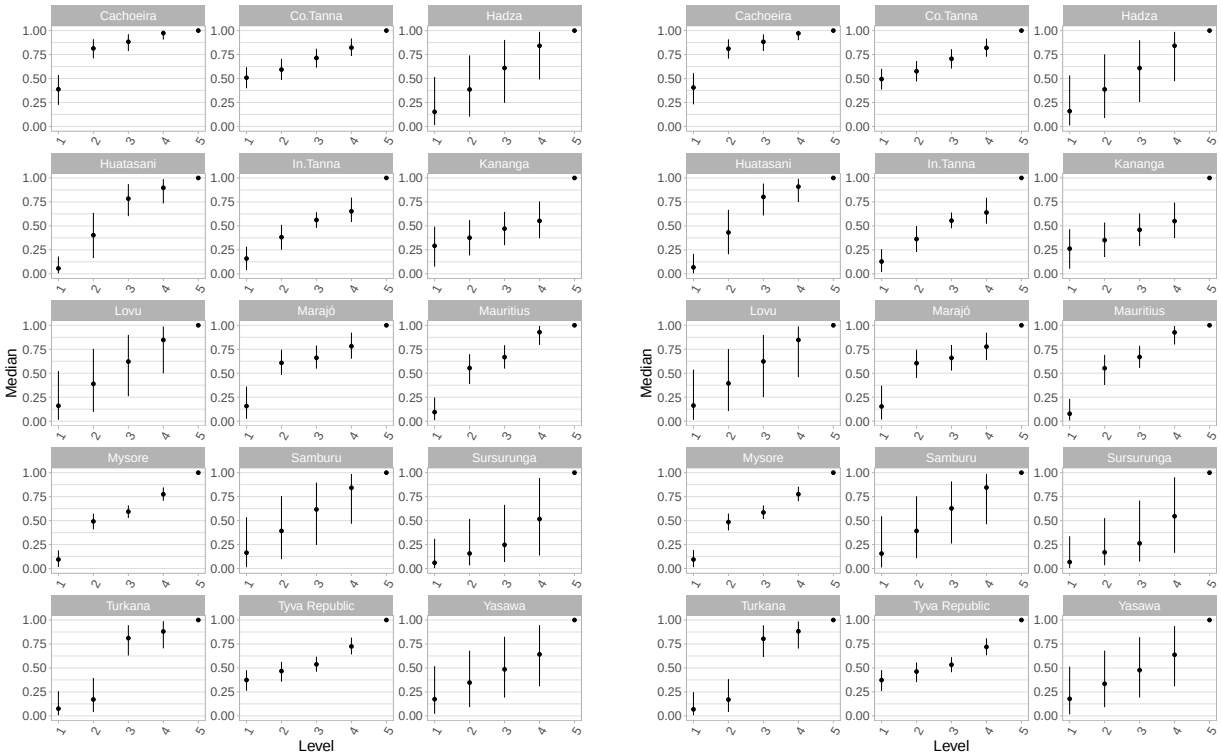

Figure S3: $\Phi_{[j]}$ from model M2 (left) and M3 (right).

S2.4. Diagnostics

Model fit was assessed with effective sample size, $\hat{r}$, and traceplots (Figs. S4-S7). Models appear to be well behaved. Details on effective sample size and $\hat{r}$ for each model can be viewed by replicating our work-flow. Traceplots are provided here for easy reference.

S2.5. Modified results from main text

In the main text, we show the "intercept" terms for the local deity model holding the MD Likert response at its median value of 2. Here, in Table S1, we set the MD Likert response to its maximum value of 4. The "intercept" parameters for the moralistic deity model still represent the predicted log odds of a female of average age and education claiming that her deities care (at least a little bit; a Likert-level of 1 or greater on a 0 to 4 scale) about moral behavior. For the local deity models, however, the "intercept" parameters represent the same quantity, but for a female of average age and education, who claims that her moralistic deity cares about moral issues at a Likert-level of 4 (rather than 2, as in the main text). In general, we find that both moralistic and local deities are reliably characterized as having at least some level of moral concern.

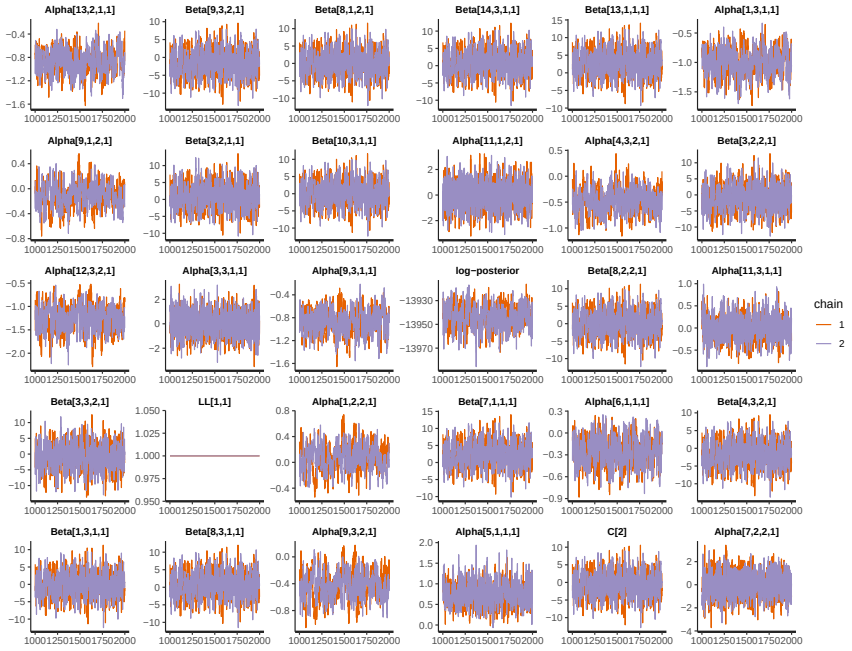

Figure S4: Traceplot for a random sample of parameters from model M0. Traceplots show good mixing and convergence of multiple chains to the same posterior region. Note that the LL parameter here reflects the Cholesky factor of a correlation matrix, whose row 1 column 1 entry is always fixed at the value of 1.0.

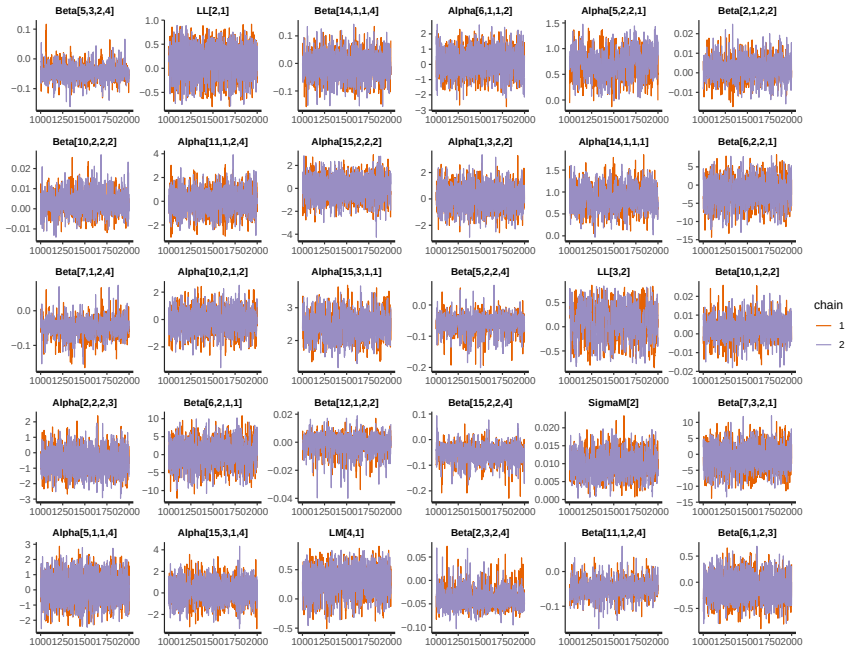

Figure S5: Traceplot for a random sample of parameters from model M1. Traceplots show good mixing and convergence of multiple chains to the same posterior region.

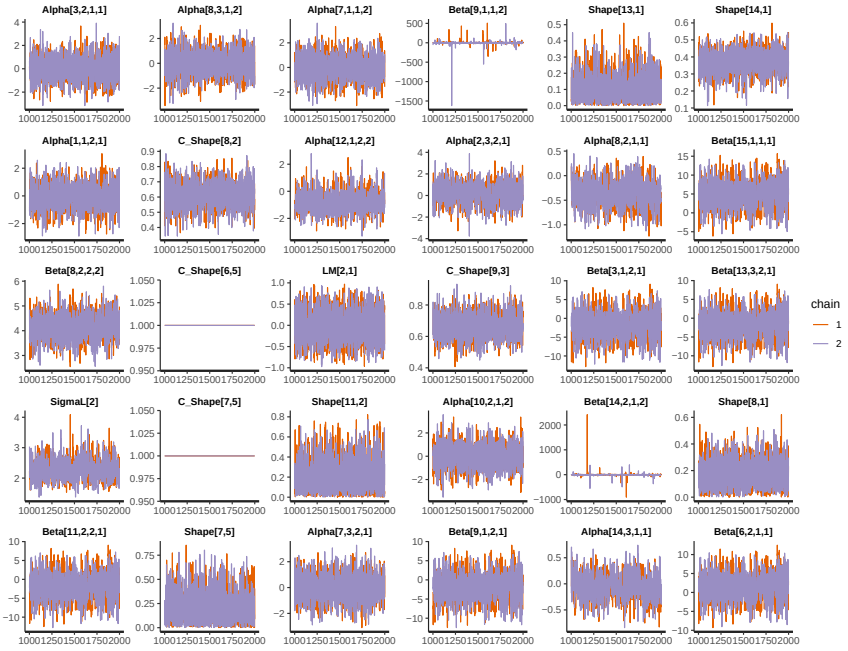

Figure S6: Traceplot for a random sample of parameters from model M2. Traceplots show good mixing and convergence of multiple chains to the same posterior region. Note that the C_Shape parameters here reflect the Φ parameters described above, whose column 5 entries are always fixed at the value of 1.0. Note also that the Beta parameters that look ill-behaved are from the conceptual bleeding term in the model for moralistic deities: since the effects of these parameters—in such a case—are removed from the model likelihood through our δ term, these traceplots simply reflect samples from a very flat, uninformative Cauchy prior distribution.

S3. Supplementary Analyses

We report here a host of supplementary analyses that are useful to assess the robustness of our main analyses. In order to estimate the chances of participants' non-zero responses (i.e., any indication of a moral association), we dummy coded all responses such that all data originally ≥ 1 were recoded as a 1. The full set of dichotomized data (i.e., the subset without missing values) includes 10,269 observations. 4,824 (47%) data rows were about the local deities, while 5,445 (53%) were about the moralistic deities. Responses about local deities were split between 44% ($n = 2{,}138$) giving a negative responses and 56% ($n = 2{,}686$) giving a positive response. Responses to moralistic deities were much starker; 88% ($n = 4{,}809$) gave a positive response, while only 12% ($n = 636$) gave a negative response. Qualitatively, the results are the same across analyses, though estimates in the main text are more conservative.

S3.1. Dichotomous Outcomes

S3.1.1. Model definition

The formal model structure of the most elaborate model (Model S3 below) is as follows:

$$y_i \sim \text{Binomial}(1, p_i)$$
$$\text{logit}(p_i) = \alpha + \gamma_{[q(i)]} + \lambda_{[s(i),d(i)]} + \beta_1 A_{[i]} + \beta_2 M_{[i]} + \beta_3 E_{[i]} + \beta_4 C_{[i]}$$

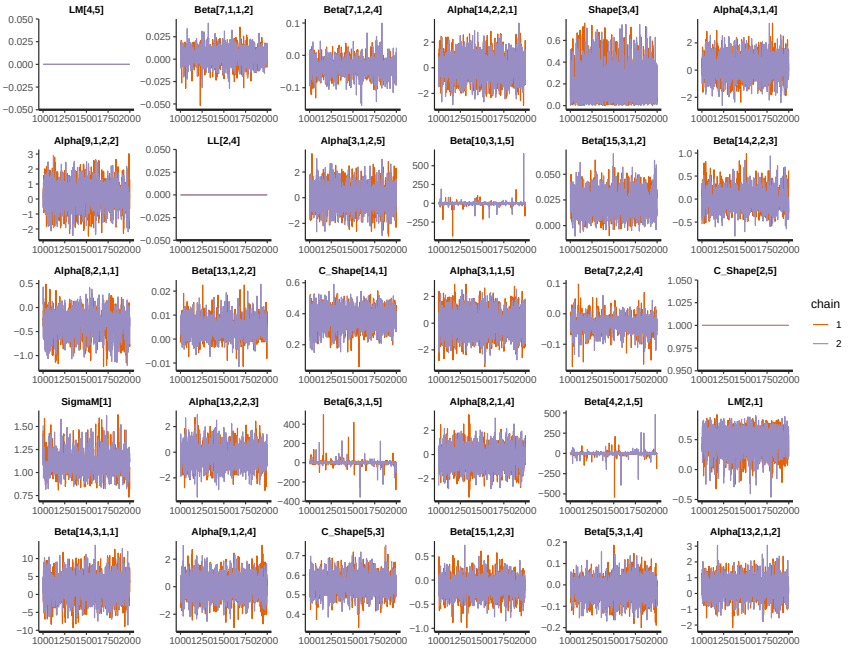

Figure S7: Traceplot for a random sample of parameters from model M3. Traceplots show good mixing and convergence of multiple chains to the same posterior region. Note that the C_Shape parameters here reflect the Φ parameters described above, whose column 5 entries are always fixed at the value of 1.0. Note also that the Beta parameters that look ill-behaved are from the conceptual bleeding term in the model for moralistic deities: since the effects of these parameters—in such a case—are removed from the model likelihood through our δ term, these traceplots simply reflect samples from a very flat, uninformative Cauchy prior distribution. The LM and LL terms are Cholesky factors of correlation matrices whose lower triangle is uniformly 0.

Here, y_i is the binary response to questions about gods' moral interests modeled using a binomial probability distribution where p_i is the probability of a positive response. Across all specifications, intercepts include: a main estimate, α, with offsets for question type (lying, stealing, and murder), $\gamma_{[q(i)]}$, and the interaction of field site and deity, $\lambda_{[s(i),d(i)]}$. Here the functions $q(i)$, $s(i)$, and $d(i)$ give the question type, site, and focal deity corresponding to response i, respectively. Simple effects and their corresponding coefficients are: centered-at-site-mean age, $\beta_1 A$; sex, $\beta_2 M$ (where male is indicated by $M = 1$); centered-at-site-mean years of formal education, $\beta_3 E$; and whether or not response i occurred in a Christian site, $\beta_4 C$ (where a Christian is indicated by $C = 1$). We used the `lme4` package V1.1.23 (Bates et al., 2015) in `R` (R Core Team, 2016). See `R` script for implementation details.

On average, if people answer questions about local deities randomly, the mean estimate of answers would be around zero (i.e., OR $\approx$ 1.00), with confidence intervals that are more or less symmetrical around zero. If, however, there is a systematic bias toward replying "no," then the mean estimate will be negative and bulk of the interval will be less than zero (i.e., the upper bound of the exponentiated interval will be below or around 1.00). If there is a systematic bias toward replying "yes," then the mean estimate will be positive and the bulk of the interval mass will be greater than zero (i.e., the lower bound of the exponentiated interval will be above or around 1.00).

Table S1: **Model estimates (90% credibility intervals) for target variables and indices by site.** In an effort to make the estimates of our ordered-categorical model more intuitive and understandable, we have re-parameterized the "intercept" terms slightly, giving them a somewhat non-standard interpretation. Accordingly, we have flagged the label with an asterisk so that readers are aware that these values should be interpreted as described below. First, we note that our continuous predictor variables have been centered by site. Second, the "intercept" terms have been adjusted so that they reflect the contributions of both the standard intercept parameter and the first cut-point, C_1. Finally, the local deity "intercept" estimates that account for "conceptual bleeding" have been adjusted so that they reflect the case where the moralistic deity score is held at the maximum value of the scale (i.e., a Likert-level of 4 on a 0 to 4 scale). Jointly, these three conditions afford the following qualitative interpretation of the "intercept" terms: the "intercept" parameters for the moralistic deity model represent the predicted log odds of a female of average age and education claiming that her deities care (at least a little bit; a Likert-level of 1 or greater on a 0 to 4 scale) about moral behavior. For the local deity models, the "intercept" parameters represent the same quantity, but for a female of average age and education, who claims that her moralistic deity cares about moral issues at a Likert-level of 4. In general, we find that both moralistic and local deities are reliably characterized as having at least some level of moral concern.

Model	Intercept	Age	Male	Education	Bleeding
Moralistic Deity; M0	2.37 (2.09; 2.66)	—	—	—	—
Moralistic Deity; M1	2.37 (2.08; 2.68)	0.002 (-0.002; 0.007)	-0.02 (-0.11; 0.07)	-0.01 (-0.03; 0.01)	—
Moralistic Deity; M2	2.49 (2.21; 2.78)	—	—	—	—
Moralistic Deity; M3	2.49 (2.19; 2.80)	0.002 (-0.002; 0.007)	-0.02 (-0.11; 0.07)	-0.01 (-0.03; 0.02)	—
Local Deity; M0	0.49 (-0.07; 0.94)	—	—	—	—
Local Deity; M1	0.86 (0.37; 1.35)	0.002 (-0.002; 0.006)	-0.14 (-0.30; 0.02)	-0.05 (-0.06; -0.03)	—
Local Deity; M2	1.49 (0.79; 2.18)	—	—	—	2.98 (2.22; 3.75)
Local Deity; M3	1.69 (1.10; 2.44)	0.003 (-0.000; 0.007)	-0.14 (-0.26; -0.01)	-0.03 (-0.05; -0.01)	3.03 (2.31; 3.76)

S3.1.2. Results

Table S2 reports the results of our models. Model S0 is the null model that includes intercept offsets for item type and field site. Model S1 adds: deity, age, sex, and years of formal education. Model S2 adds Christianity as a simple effect, while Model S3 allows the effect of deity type to vary across sites (formal definition above).

	Model S0	Model S1	Model S2	Model S3
Intercept	3.98	1.73	2.51	4.74
	[2.39, 6.65]	[1.08, 2.75]	[1.31, 4.78]	[2.58, 8.72]
Deity (1 = moralistic)	—	6.58	6.58	—
		[5.92, 7.32]	[5.92, 7.32]	
Age[a]	—	0.99	0.99	0.99
		[0.99, 1.00]	[0.99, 1.00]	[0.99, 1.00]
Sex (1 = male)	—	0.86	0.86	0.84
		[0.78, 0.96]	[0.78, 0.96]	[0.75, 0.93]
Education[a]	—	0.93	0.93	0.93
		[0.91, 0.94]	[0.91, 0.94]	[0.91, 0.94]
Christian deity (1 = yes)	—	—	0.52	1.06
			[0.23, 1.18]	[0.47, 2.42]
Log likelihood	-5679.1	-4923.7	-4922.6	-4522.7
AIC	11355.72	9861.5	9861.1	9063.4
Varied Field Site?	intercept	intercept	intercept	intercept
Varied Item Type?	intercept	intercept	intercept	intercept
Varied Deity Type?	none	none	none	effect

Table S2: **Exponentiated estimates and [95% confidence intervals] of item responses to deities' moral interests.** There were 10,269 observations in each model. [a]Centered at site-specific mean. All models converged.

Across all models, the lower bound of the 95% confidence intervals for the main intercept, α, are all >1.00, showing that after holding all other factors constant, replying positively to questions about local deities is more likely than responding "no" or answering at random. Model S2 predicts that there is a 67% chance of a non-zero response to moralistic punishment questions about local deities for average-aged and educated males. If the site is also Christian, the model predicts a 47% chance of a non-zero response, a reduction of nearly 20%. Moralistic gods, however, are 92% likely to have non-zero responses. Also consistent with our expectations, age, sex, and years of formal education predict lower odds of responding positively to moral punishment index questions. As it is centered by the site-specific mean, the effect for age appears slight but stable. In summary, being male, more educated, and older predicts *less* moralization of local deities.

In Model 3, the simple effect of deity present in previous models is usurped by the fact that it also varies across sites (as does the interaction between the deity and Christian god indicators). The interaction in Model 2 shows that the Christian deity has a very strong association with positive responses. Though, as indicated by the wide confidence intervals in Model 3, once we vary the effect for moralistic deities across sites, the bulk of the interval mass is on the positive side of 1.00. As indicated by the interval breadth, however, this estimate is considerably less precise, possibly due to difficulties in fitting multi-level models in base `R`. For similar models fit in `Stan`, see the main text.

S3.1.3. Cross-item variation

Item types show some variation in their contribution to responses. For Model S3, the offset for murder is positive (0.18) while the offsets for lying and theft are negative (-0.09 and -0.10, respectively). The model estimates that the chances of saying "yes" to the moral punishment index questions is 82% (the logistic transform of the intercept, 1.56). If we add murder, the model estimates a 3% increase in the chances that a participant responds affirmatively. If we only consider the contribution of the theft question, the chances of saying yes only get smaller by only 1%. In other words, while "murder" is highly likely to elicit a positive response, when considering the other factors that have a negative effect, the probability remains well above 50%. While responses to gods' interest in murder might be a reflection of the extreme nature of homicide, it is curious that theft and lying were virtually indistinguishable from each other, considering that cross-culturally, theft is generally more salient than simple dishonesty in defining what it means to be immoral (Purzycki et al., 2018).

S3.1.4. Cross-cultural variation

There is also considerable cross-cultural variation (Figure S8). We tease apart this variation in Table S3. In these tables, α refers to the main intercept, while $\lambda_{s,1}$ refers to site-specific offsets for local deities. The $\lambda_{s,2}$ parameter vector refers to the varying slopes for moralistics deity type across sites. The logistic transforms of estimate summations are denoted by the columns "Prob." These can be interpreted as the probability of answering "yes" for gods' moral interests (local gods on the left side of the vertical line, moralistic gods on the right side). We include the differences between the local and moralistic deities in the "Diff." column. We graph this cross-cultural variation in Figure S8.

Culture	$\lambda_{s,1}$	$\lambda_{s,1}+\alpha$	Prob.$_{\text{LD}}$	$\lambda_{s,2}$	$\lambda_{s,2}+\lambda_{s,1}+\alpha$	Prob.$_{\text{MD}}$	Diff.
Tyva Republic	0.88	2.44	0.92	-0.20	2.24	0.90	-0.02
Mysore	0.42	1.98	0.88	-0.07	1.91	0.87	-0.01
In.Tanna	0.15	1.70	0.85	0.24	1.94	0.87	0.03
Co.Tanna	0.09	1.64	0.84	0.24	1.89	0.87	0.03
Marajo	-0.14	1.42	0.81	0.49	1.91	0.87	0.07
Samburu*	-0.39	1.17	0.76	0.93	2.09	0.89	0.13
Cachoeira	-0.58	0.97	0.73	0.49	1.46	0.81	0.09
Mauritius	-1.57	-0.01	0.50	1.69	1.67	0.84	0.35
Huatasani	-1.69	-0.13	0.47	0.87	0.74	0.68	0.21
Lovu*	-1.83	-0.27	0.43	4.34	4.06	0.98	0.55
Kananga	-2.08	-0.53	0.37	2.97	2.44	0.92	0.55
Turkana	-2.13	-0.58	0.36	5.60	5.03	0.99	0.63
Yasawa	-2.34	-0.78	0.31	4.61	3.83	0.98	0.66
Sursurunga	-2.86	-1.31	0.21	4.42	3.11	0.96	0.74

Table S3: **Culture-specific ($\lambda_{s,1}$) and main (α) intercepts and varying slopes ($\lambda_{s,2}$).** Values are from Model S3 in Table S2. Values are sorted by $\lambda_{s,1}$ in descending fashion. The Diff. column is the difference between Prob.$_{\text{LD}}$ and Prob.$_{\text{MD}}$. Dashed line indicates where Prob.$_{\text{LD}}$ is $\approx 50\%$. *Sites did not answer local deity questions.

The model estimates the lowest likelihood of replying affirmatively to local gods' moral concerns to be among the Mauritians, Huatasani, Lovu, Kanangans, Turkana, Yasawans, and Sursurungans. If we take the case of Sursurungans who have the largest estimated difference between deities, individuals effectively have a 96% chance of answering yes to the moral questions about moralistic deities, thus a difference of 74%. Note that Christianity does *not* reliably predict outcomes for local deities. What does predict local deity outcomes, however, is how local samples address commitment to local spirits (see main text). The dotted lines in Tables S3 and S5 indicate where this distinction lies.

S3.2. Teasing apart local deities

Recall from the main text that the Huatasani and Kanangan samples were split by way of questions about different local spirits. Some of the Huatasani answered questions about ancestor spirits, while others answered questions about saints. The Kanangan sample answered questions about either ancestor spirits, or *Kadima*, a deity tradition that is maligned by the local Christian doctrine. We therefore reran the regressions above, but separated the two local deities in these two groups by treating them as different sub-samples. Our results (Table S4) are qualitatively consistent with the main results.

To further assess our interpretation of the variation illustrated in Figure 2 in the main text, we examined the by-site intercepts and varying effects for deities. Table S5 reports the results. As predicted by our interpretation in the main text (we did actually conduct these focal analyses *after* the main text analyses), the *Kadima* are associated with far less moralization than the ancestors. As indicated by the logistic transformed summations of the item and site-specific intercepts, there is only a 24% chance of answering affirmatively to the questions about *Kadima* while responses about ancestors' moral concerns are 42%. The two local spirits in Peru—the *apus* or mountain spirits and saints are virtually indistinguishable in terms of their association with moralistic punishment.

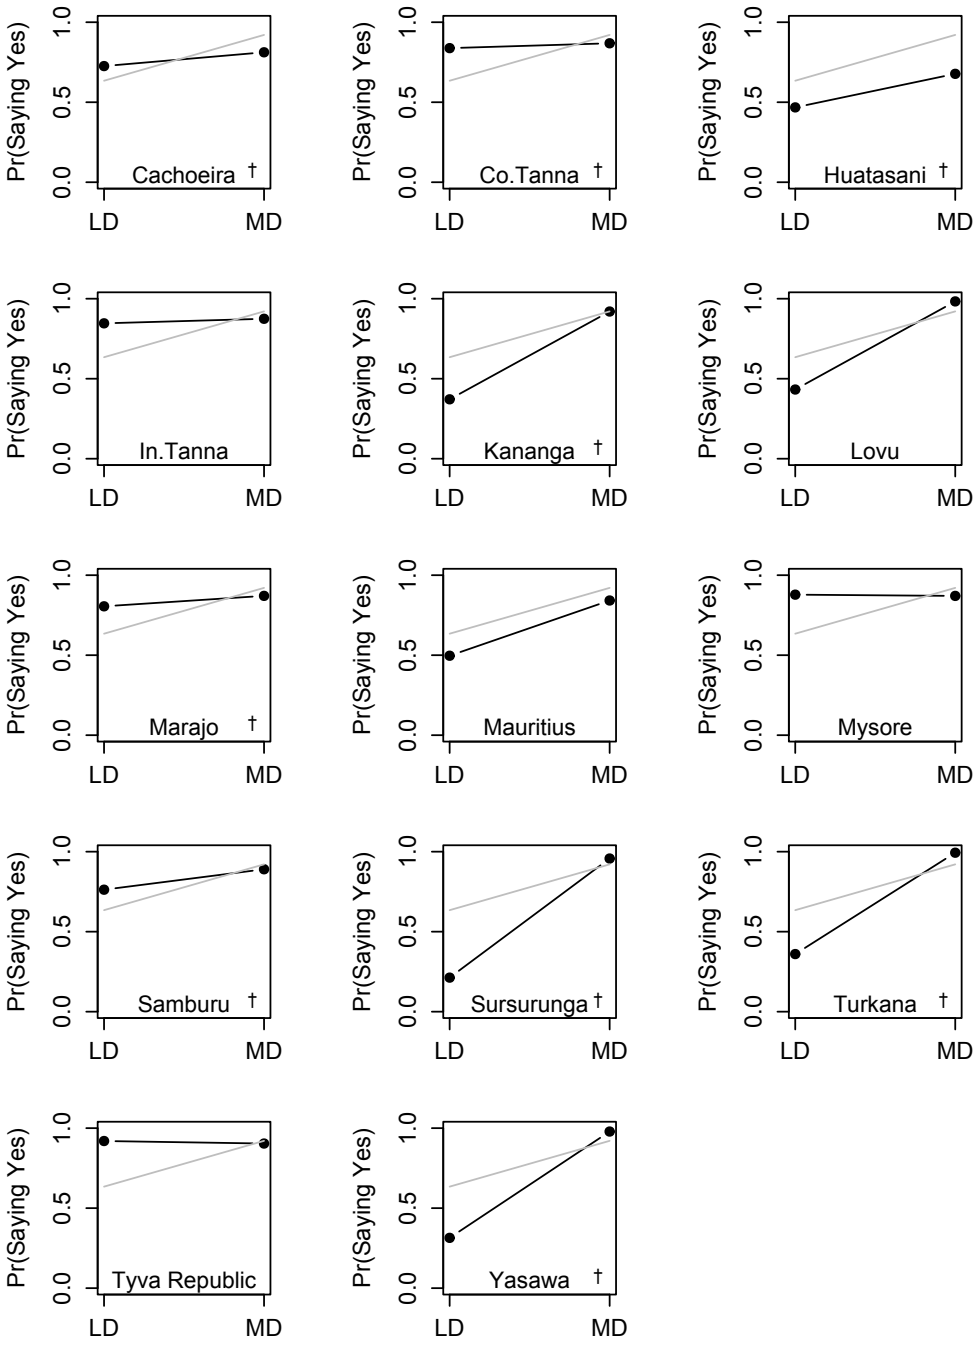

Figure S8: **Probability plot of "yes" answers to moral questions about the local (LD) and moralistic (MD) deities.** Regression lines (black) are logistic transformed summations of regression equations when deity values are 0 (local deity) and 1 (moralistic deity). Values are from Model S3 in Table S2 and values in Table S3. Baseline regression lines (gray) are the same for each site and are calculated by using the mean site-intercept and mean varying effect for deity type. †Christian sites.

	Model B0	Model B1	Model B2	Model B3
Intercept	3.51	1.53	2.00	3.98
	[2.17, 5.70]	[0.98, 2.38]	[1.08, 3.78]	[2.20, 7.18]
Deity (1 = moralistic)	—	6.62	6.61	—
		[5.95, 7.36]	[5.95, 7.36]	
Age[a]	—	0.99	0.99	0.99
		[0.99,1.00]	[0.99, 1.00]	[0.99, 1.00]
Sex (1 = male)	—	0.87	0.87	0.84
		[0.79, 0.96]	[0.79, 0.96]	[0.76, 0.94]
Education[a]	—	0.93	0.93	0.93
		[0.91, 0.94]	[0.91, 0.94]	[0.91, 0.94]
Christian deity (1 = yes)	—	—	0.61	1.21
			[0.28, 1.36]	[0.54, 2.70]
Log likelihood	-5672.34	-4917.66	-4916.95	-4512.73
AIC	11350.68	9849.32	9849.90	9043.47
Varied Field Site?	intercept	intercept	intercept	intercept
Varied Item Type?	intercept	intercept	intercept	intercept
Varied Deity Type?	none	none	none	effect

Table S4: **Exponentiated mean estimates and [95% confidence intervals] of item responses to deities' moral punishment concern when Huatasani and Kanangan sites' local deities are teased apart.** [a]Centered at site-specific mean. There were 10,269 observations in each model. All models converged.

Culture	$\lambda_{s,1}$	$\lambda_{s,1}+\alpha$	Prob.$_{\text{LD}}$	$\lambda_{s,2}$	$\lambda_{s,2}+\lambda_{s,1}+\alpha$	Prob.$_{\text{MD}}$	Diff.
Tyva Republic	1.05	2.43	0.92	-0.20	2.23	0.90	-0.02
Mysore	0.59	1.97	0.88	-0.07	1.90	0.87	-0.01
In.Tanna	0.31	1.70	0.85	0.23	1.93	0.87	0.03
Co.Tanna	0.13	1.51	0.82	0.24	1.75	0.85	0.03
Marajo	-0.10	1.29	0.78	0.49	1.77	0.85	0.07
Samburu*	-0.41	0.97	0.72	0.99	1.96	0.88	0.15
Cachoeira	-0.54	0.84	0.70	0.49	1.33	0.79	0.09
Mauritius	-1.40	-0.02	0.50	1.68	1.67	0.84	0.35
Huatasani (*apus*)	-1.47	-0.08	0.48	1.25	1.17	0.76	0.28
Huatasani (saints)	-1.53	-0.15	0.46	0.59	0.45	0.61	0.15
Kananga (ancestors)	-1.70	-0.32	0.42	2.46	2.14	0.89	0.47
Lovu*	-1.91	-0.53	0.37	4.56	4.04	0.98	0.61
Turkana	-2.09	-0.71	0.33	5.59	4.89	0.99	0.66
Yasawa	-2.30	-0.92	0.29	4.59	3.68	0.98	0.69
Kananga (*Kadima*)	-2.55	-1.17	0.24	3.73	2.56	0.93	0.69
Sursurunga	-2.83	-1.44	0.19	4.41	2.96	0.95	0.76

Table S5: **Culture-specific ($\lambda_{s,1}$) and main (α) intercepts and varying slope ($\lambda_{s,2}$) table for split subsamples (Huatasani and Kananga).** Values are from Model B3 in Table S4. Values are sorted by $\lambda_{s,1}$ in descending fashion. The Diff. column is the difference between Prob.$_{\text{LD}}$ and Prob.$_{\text{MD}}$. Dashed line indicates where Prob.$_{\text{LD}}$ is $\approx 50\%$. *Sites did not answer local deity questions.

S4. Hadza Free-List Data

S4.1. Coding

We used the following rubric (see Purzycki and McNamara, 2016, for further discussion) to code the free-list data. Codes are in bold. As these are general categories, individual could list multiple instances of a single code (e.g., if someone listed "theft" and "murder", both would get coded as **Morality**). We therefore used the option to consider only individuals' earliest-listed items with multiple iterations (i.e., we used the `MAX` argument in the `dealwithDoubles` function in the `AnthroTools` package; Jamieson-Lane and Purzycki 2016).

1. **Morality**: generalized behaviors that have a benefit or cost to other people (e.g., hurting, being generous, sharing, etc.)
2. **Virtue**: individual qualities that may or may not have social ramifications (e.g., hard-working, kind, bad conscience, etc.)
3. **People**: in reference to the quality, and/or the state of people (e.g., people, people stay in good health, be happy, etc.)
4. **Etiquette**: conventional social behaviors that have no immediate cost or benefit to others (e.g., shaking hands, wearing the proper clothes, etc.)
5. **Substance Use/Abuse**: Items that involve the use of illicit substances
6. **Religion**: any non-ritual or non-behavioral item concerned with the supernatural (e.g., faith, devotion, loving god, etc.)
7. **Ritual**: any behavior or object used in ritual devoted to the supernatural (e.g., praying, meditation, offerings, sacrifices, not participating in ritual, etc.)
8. **Ecology**: any behavior or object affecting non-human relationships (e.g., pollution, keeping sacred places clean, gardening, etc.)
9. **Food**: any item composed of food items (e.g. yam, milk, etc.)
10. **Miscellaneous**: miscellaneous items (items that cross-cut categories, etc.)
11. **D/K**: I don't know, not sure, etc.
12. **Specific**: Items that are specific to a culture (e.g., bel' leaf, artysh, etc.)

S4.2. Free-list analysis

Table S6 reports all Smith's S scores for each category type across free-list tasks. Across sub-domains (i.e., deity type and concern valence), the most salient items—item codes that are higher within and more prevalent across lists—are People, Morality, and D/K.

S4.3. Regression analysis

To estimate the relationship between contact with missionaries and associating *Haine* or *Ishoko* with moral concern, we regressed the presence of moral content in free-lists on participant sex and self-reported contact with missionaries. Recall that only participants who claimed to believe in *Haine* or *Ishoko* were asked to do the free-list task. Because these are two different deities across two sub-domains, we analyzed these data in four different models. The general model structure is as follows: the probability (p_i) of moral content in free-list responses, y_i, is distributed following a binomial distribution. Using a logit link function, our linear model predicting p_i is defined as: $\text{logit}(p_i) = \alpha + \beta_1 * \text{Sex}_i + \beta_2 * \text{MissionaryContact}_i$. We set priors uninformatively, $\alpha, \beta_1, \beta_2 \sim \text{Normal}(0, 10)$. Here, y_i is the presence of an item coded as "moral" in the four different

Table S6: **Smith's S scores of Hadza free-list data.** Bold items are two most salient catgories in each sub-domain.

Code	Pleases *Haine*	Angers *Haine*	Pleases *Ishoko*	Angers *Ishoko*
Morality	0.09	**0.49**	0.06	**0.46**
Virtue	0.23	0.05	0.15	0.06
People	**0.31**	–	**0.29**	–
Etiquette	0.02	0.00	0.02	–
Drug Use/Abuse	–	0.12	–	0.06
Religion	–	0.02	–	0.02
Ritual	0.19	0.08	0.15	0.07
Food	0.03	–	0.01	–
Miscellaneous	0.03	0.02	0.09	0.03
D/K	**0.28**	**0.35**	**0.34**	**0.38**
Specific	0.02	–	–	–

free-list tasks discussed above. Participant sex (male = 1) and missionary contact (yes = 1) are indicator variables. Because of missing data and incompatibility across data sets, we only have 40 individuals in this analysis. Table S7 consists of the raw estimates of these models, based on analysis using the `brms` package for `R` (Bürkner, 2018).

Table S7: **Estimates from a logistic regression predicting moral content in individuals' models of *Haine's* and *Ishoko's* concerns.** α is the intercept and β represents the slopes. Priors were uninformative: $\alpha, \beta \sim$ Normal(0, 10). β_2 refers to the slope coefficient for exposure to Christian missionaries (yes = 1), and β_1 refers to the slope coefficient for participant sex (male = 1).

	Haine			***Ishoko***		
	α	$\beta_2(Missionary)$	$\beta_1(Male)$	α	$\beta_2(Missionary)$	$\beta_1(Male)$
Angers	-0.58	0.32	0.36	-1.24	-0.30	0.83
	[-1.67, 0.48]	[-0.74, 1.43]	[-0.72, 1.45]	[-2.46, -0.10]	[-1.45, 0.86]	[-0.34, 2.00]
Pleases	-2.42	-0.32	-0.41	-4.33	-0.42	0.37
	[-4.23, -0.92]	[-1.90, 1.16]	[-1.98, 1.13]	[-7.62, -2.04]	[-2.23, 1.37]	[-1.38, 2.14]

References

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

Bürkner, P.C., 2018. Advanced Bayesian multilevel modeling with the R package brms. The R Journal 10, 395–411. doi:10.32614/RJ-2018-017.

Jamieson-Lane, A., Purzycki, B.G., 2016. AnthroTools: Some custom tools for anthropology. R package version 0.8.

McElreath, R., 2020. Statistical Rethinking: A Bayesian Course with Examples in R and Stan. 2nd ed., CRC Press, New York.

Purzycki, B.G., McNamara, R.A., 2016. An ecological theory of gods' minds, in: De Cruz, H., Nichols, R. (Eds.), Cognitive Science of Religion and Its Philosophical Implications. Continuum, New York, pp. 143–167.

Purzycki, B.G., Pisor, A., Apicella, C., Atkinson, Q.D., Cohen, E., Henrich, J., McNamara, R.A., Norenzayan, A., Willard, A.K., Xygalatas, D., 2018. The cognitive and cultural foundations of moral behavior. Evolution and Human Behavior 39, 490–501.

R Core Team, 2016. R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing. Vienna, Austria.

Stan Development Team, 2017. RStan: the R interface to Stan. URL: http://mc-stan.org/. r package version 2.19.