



City Research Online

City, University of London Institutional Repository

Citation: Rigoli, F. & Lennon, J. (2024). The Gods as Latent Causes: A Statistical Inference Theory of Religion. *The International Journal for the Psychology of Religion*, pp. 1-26. doi: 10.1080/10508619.2024.2422173

This is the published version of the paper.

This version of the publication may differ from the final published version.

Permanent repository link: <https://openaccess.city.ac.uk/id/eprint/34038/>

Link to published version: <https://doi.org/10.1080/10508619.2024.2422173>

Copyright: City Research Online aims to make research outputs of City, University of London available to a wider audience. Copyright and Moral Rights remain with the author(s) and/or copyright holders. URLs from City Research Online may be freely distributed and linked to.

Reuse: Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.



The Gods as Latent Causes: A Statistical Inference Theory of Religion

Francesco Rigoli & Jack Lennon

To cite this article: Francesco Rigoli & Jack Lennon (04 Nov 2024): The Gods as Latent Causes: A Statistical Inference Theory of Religion, The International Journal for the Psychology of Religion, DOI: [10.1080/10508619.2024.2422173](https://doi.org/10.1080/10508619.2024.2422173)

To link to this article: <https://doi.org/10.1080/10508619.2024.2422173>



© 2024 The Author(s). Published with license by Taylor & Francis Group, LLC.



Published online: 04 Nov 2024.



Submit your article to this journal [↗](#)



Article views: 120



View related articles [↗](#)



View Crossmark data [↗](#)

The Gods as Latent Causes: A Statistical Inference Theory of Religion

Francesco Rigoli ^a and Jack Lennon^b

^aDepartment of Psychology, City, University of London, London, UK; ^bSchool of Archaeology and Ancient History, University of Leicester, Leicester, UK

ABSTRACT

The statistical brain hypothesis posits that the brain constructs probabilistic models of the environment. Here we examine whether this perspective can provide any insight on religion. We propose that religious ideas represent an attempt to explain away residuals, that is, to explain discrepancies between observations and predictions. The framework postulates probabilistic generative models where gods are described by latent variables whose possible states correspond to the actions available to the gods. As examined in the paper, our proposal offers a plausible interpretation of typical religious phenomena such as miracles, omens, and divination. Moreover, it captures important characteristics of religious beliefs including the notion that gods control multiple spheres of reality, are organized hierarchically, and control aspects that are salient for believers. Besides offering an intriguing new perspective on religion, the paper corroborates the possibility that the statistical brain hypothesis represents a unifying theory of the mind.

Introduction

From the perspective of an outsider, the religion professed by a different culture often appears as paradoxical. For example, a Christian might wonder how it is possible to believe, as the Lenape people of North America did (Dalton, 2004), that a giant turtle holds the world upon its shell. And yet, the same person might be surprised to learn that, in antiquity, most Jews manifested the same puzzlement when acquainted with the Christians' claim that the messiah was a humble man crucified by the Romans (Latourette & Winter, 1975). The apparently bizarre nature of many religious creeds, at least as seen be an external eye, has led some commentators to the conclusion that religion is one of the clearest manifestations of human fundamental irrationality (Freud, 1927; Shermer, 2002; Wolpert, 1994). Yet, most scholars have not abandoned the idea that, despite the apparent paradoxes, religion has its own underlying logic. This perspective fits with the functional outlook that prevails in evolutionary and social sciences: after all, humans are believed to be rational, or more precisely bounded-rational, creatures (Simon, 1997). Among the most influential theories embracing this perspective are those maintaining that, whatever its limitations and biases, religion represents an attempt to explain how reality works – we call this an *epistemic* view of religion (Barrett, 2000; Boyer, 1994; Guthrie, 1995; Iannaccone et al., 1998; McCauley, 2017; Spilka et al., 2019; Stark & Finke, 2000; Tylor, 1871). In this view, religion would not differ much from philosophy and science with regard to its purposes (i.e., the purpose of understanding reality), but rather with regard to its assumption that supernatural forces are key factors at play. The notion that epistemic motives underly religion is far from universally accepted (Rigoli, 2021). Many scholars advocate other motives such as sedating

CONTACT Francesco Rigoli  francesco.rigoli@city.ac.uk  Department of Psychology, City, University of London, Northampton Square, London EC1V 0HB, UK

© 2024 The Author(s). Published with license by Taylor & Francis Group, LLC.

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

anxiety (Atran & Norenzayan, 2004; Hume, 1757; Jong & Halberstadt, 2017; Kay, Gaucher, et al., 2010), fostering community bonds and moral values (Batson et al., 1993; Baumeister, 1991; Durkheim, 1912; Krause & Wulff, 2005; McKay & Whitehouse, 2015; Pargament et al., 1983; Saroglou, 2011), or promoting self-interest (Gramsci, 1971; Jost et al., 2014) as being central to religion. Still, epistemic theories, or at least the idea that epistemic considerations, although not unique, are crucial to religion, remain compelling according to many.

An epistemic theory of religion should ensure to be consistent with established knowledge in psychology and neuroscience. In these disciplines, the *statistical brain hypothesis* is one of the most influential views about how epistemic processes work in the brain (Clark, 2013; Friston & Kiebel, 2009; Knill & Pouget, 2004; Oaksford & Chater, 2007; Rao & Ballard, 1999). This argues that, broadly speaking, the brain works like a statistical inference machine. The proposal is that the brain functioning is akin to a process that, based on noisy and ambiguous information, aims at building statistical models of reality. Does this perspective have anything to add to our understanding of religion? Addressing this question appears to be promising for at least two reasons. First, it seems reasonable to expect that one of the most compelling cognitive theories available, that is, the statistical brain hypothesis, can contribute to understand any cultural phenomenon, including religion. Second, the argument that the statistical brain hypothesis offers a general theory of cognition requires to assess the theory in virtually all domains, including in complex phenomena like religion. Recent pioneering work has advocated the statistical brain hypothesis to explain aspects of religion (this literature will be examined in detail below) (e.g., Andersen, 2019; Schjoedt et al., 2013; Taves & Asprem, 2017; van Elk & Aleman, 2017; van Elk et al., 2016). However, despite invaluable insight, this literature has focused on somewhat specific aspects; how the statistical brain outlook can be adopted to frame the religious phenomenon at large remains an open question. The present paper aims at addressing this question. Our argument will start by considering the main tenets of the statistical brain hypothesis. Next, this framework will be applied to explain the nature of the cognitive processes underlying religious beliefs. We will then assess whether a statistical interpretation of religion can offer any new insight about important religious phenomena such as miracles, omens, and divination, and about the nature of religious beliefs as expressed by various cultures. We will conclude by evaluating our statistical interpretation of religion in the context of previous theories and, finally, we will discuss a broader set of issues arising from our argument.

The statistical brain hypothesis

Evolutionary theory posits that, for any living organism, the fundamental imperative for survival is to approach certain stimuli (e.g., food, water, shelter) and to avoid others (e.g., predators or potential sources of disease). This requires the brain to assess reality with some degree of accuracy, by addressing questions such as: is there a predator hidden behind the tree? Is this fruit edible or poisonous? Addressing these questions is far from easy: the world around us is complex, uncertain, and ambiguous. The animal behind the tree might be to a large extent concealed, meaning that its identity needs to be inferred based on partial visual features. There might be fog, rendering the scene even harder to decipher. And some predators can mimic the behavior of preys, thus misleading the observer. The statistical brain hypothesis argues that, to deal with the inherent uncertainty of the environment, the brain has evolved strategies that, to a large degree, resemble the methods employed by statisticians to analyze data (Clark, 2013; Friston & Kiebel, 2009; Knill & Pouget, 2004; Oaksford & Chater, 2007; Rao & Ballard, 1999).

According to this framework, the brain constructs *generative models* of the environment, that is, probabilistic models about the causes of sensory experience (Rigoli et al., 2017). Generative models include two different families of variables: those that, by conveying information from sensory organs, can be directly observed, and those that cannot be directly observed and thus remain hidden or latent. The latter, despite being unavailable directly, are assumed by the brain to be *latent causes* of observable variables; in other words, unobservable variables are treated as factors generating observable variables (Gershman & Niv, 2010; Gershman et al., 2015). It follows that, once observable variables are known, these can be relied upon to infer their latent causes. It is precisely the inference of latent causes which,

according to the statistical brain hypothesis, is the key process performed by the brain. To see how this works, consider a set of observable features such as sharp teeth, long claws, brown fur, and four legs. The brain would interpret these as being caused by the presence of a lion. Put another way, although the lion is a latent cause that cannot be observed directly, it manifests itself through a set of features that can be observed. The presence of such features allows the brain to infer that a lion is present.

As the example of the lion indicates, in a generative model multiple observable variables (e.g., teeth, claws, fur, and legs) typically map to one single latent variable (e.g., animal). This occurs because the function of latent variables is to simplify the picture by extracting the relevant information conveyed by observable variables (Rigoli et al., 2017). Put another way, latent variables capture the underlying *structure* that can be filtered out from observable variables (Gershman & Niv, 2010; Gershman et al., 2015). The same idea is at the basis of standard methods for dimension reduction adopted in statistics, such as factor analysis or principle component analysis.

It is important to highlight the fundamental difference between the statistical approach just outlined and classical formal logic (Oaksford & Chater, 2007). The latter requires the presence of a necessary set of features; for instance, the presence of sharp teeth, long claws, brown fur, and four legs are all required for inferring the presence of a lion. On the contrary, a statistical approach is grounded on probability. Here a set of latent causes (e.g., lion, tiger, leopard) compete and, based on the observed variables (sharp teeth, long claws, brown fur, and four legs), the probability of each latent cause is estimated. For instance, one might conclude that, based on the features observed, the lion has 0.7 probability, the tiger 0.2, and the leopard 0.1. This inference can be derived according to Bayesian rules, which enable the brain to capture the role of prior probabilities (Oaksford & Chater, 2007). For instance, before assessing any observable variable, the brain might expect the tiger to be more probable than the lion and the leopard (e.g., because presence of the tiger was more frequent in the past). After appraisal of observed variables, this prior probability can be updated by computing a posterior probability, the latter derived in such a way that both the prior probability and novel information count, each appropriately weighted.

Within a generative model, latent variables can be organized hierarchically in such a way that more abstract latent variables are treated as causes of more specific latent variables (Friston, 2005; Rigoli et al., 2017). This enables the brain to build probabilistic hierarchical taxonomies of objects in the environment. For instance, on a scale going from more abstract to more specific, the brain can represent variables such as “plants versus animals,” “carnivorous versus herbivorous,” and “lion versus tiger.” Based on the context, the brain can adopt the hierarchical model to infer latent causes at different levels of abstraction. For instance, in the wild, it is preferable to focus on a relatively abstract distinction between carnivorous and herbivores, whereas in an ethology class one can indulge on examining the subtleties distinguishing lions and tigers.

What are the mechanisms through which the brain builds generative models of the world? Besides the constraints derived from genetics, the statistical brain hypothesis attributes a key role to experience. The proposal is that generative models do not reflect experience in a neutral fashion, but in a way that depends on the interaction between the agent and the environment (Barsalou, 2008; Clark, 2013; Friston et al., 2017; Ramstead et al., 2020; Rigoli et al., 2017). This implies that the generative model is tuned to the specific characteristics of the environment that are salient for the behavior of the agent. As an example, consider an animal living in the forest who feeds on certain types of fruits available high on the trees. The animal’s generative model will be specialized in subtle discriminations between edible and poisonous fruits, and between fruits of different taste, shape, or of varying nutrient qualities. The same generative model will ignore other features of the environment (such as those concerning the terrain, or the tree roots, or the wood), because these are not relevant. The model’s focus on action-relevant variables is referred here to as *action-specificity*. For the human species, three factors can be proposed as shaping generative models according to action-specificity principles. The first of such factors comprises biological constraints: humans feed on a restricted number of objects, can live within a limited range of temperature, and are susceptible to infections of specific micro-organisms. These are all aspects that constrain the formation of generative models. The second key factor is the

environment: the generative model of someone living in the arctic will be dramatically different compared to the model of a person dwelling in the Sahara Desert. Yet, biological and environmental forces are not enough: the social structure represents the third independent factor affecting generative models. Here we define social structure broadly, as encompassing aspects such as the subsistence strategy, the organization of labor, the technology available, the presence of formal institutions, the distribution of military power, etc. To some extent, the social structure itself depends on biological and environmental forces, but it also partially transcends them. In turn, the social structure is key in shaping the generative models embraced by people participating in it.¹ For example, besides any biological or environmental factor, availability of advanced technologies such as cars and computers has arguably a huge impact upon people's generative models.

To summarize, the statistical brain hypothesis posits that the brain constructs generative models of the world that are probabilistic in nature, reflecting beliefs about how sensory experience is produced by latent causes. These models can include latent variables that are organized hierarchically according to their level of abstraction. Action-specificity, that is, the idea that generative models reflect the interaction between an agent and its environment, is proposed to be the principle driving formation of generative models. We have now overviewed the main tenets of the statistical brain hypothesis. Is it useful to apply this outlook to interpret religion? Can any novel insight be gained by doing so? To address these questions, the next section attempts to develop a statistical inference framework applied to religion.

A statistical hypothesis of religion

Although its support among scholars is far from unanimous (Bowers & Davis, 2012), the statistical brain hypothesis represents one of the most influential perspectives in cognitive science. This framework has been proposed to explain the most diverse mental phenomena, from how color contrast affects vision (Knill & Pouget, 2004; Rao & Ballard, 1999) to how social groups form shared beliefs (Devaine et al., 2014). Given the success of this theory, it does not appear far-fetched to imagine that statistical inference processes performed by the brain are also at the root of religion. Building on recent pioneering research (Andersen, 2019; Schjoedt et al., 2013; Taves & Asprem, 2017; van Elk & Aleman, 2017; van Elk et al., 2016), we will consider this possibility in what follows.

The starting point of our argument is that, for humans, the world is only partially predictable. For example, to some degree we succeed in predicting where to find food or shelter, or where a predator is more likely to be hunting. But these predictions will not always be confirmed. Consider a neolithic village dependent on agricultural output for its survival. At the beginning of the season, peasants will attempt to predict the output based on various factors such as the units of labor employed, the amount of rainfall, and the presence of parasites. When the harvest is completed, the prediction can eventually be compared with the outcome. As modern statisticians know when dealing with similar problems, there is always some discrepancy between prediction and outcome, despite the discrepancy being sometimes small and other times large. In statistics the difference between outcome and prediction is called *residual*.² When predictions and outcomes are compared regarding something vital such as, in the example of our neolithic village, regarding the harvest, a fundamental question arises: where do the residuals come from? Broadly speaking, humanity has offered three different answers to this question. The first is that residuals come from natural factors that could be potentially known, although in fact

¹The social structure is also responsible for explaining why different groups of people within the same society often entertain substantially different generative models.

²Although this example focuses on predicting a future outcome, the same logic can be advanced in the context of inferring what happened in the past – a scenario sometimes referred as *postdiction*. Consider a historian seeking to know how many people lived in a small neolithic village that has never been excavated. As an initial guess, the historian may rely on scientific knowledge about the factors shaping population density at the time, such as the climate, the soil quality, proximity to a river, etc. Archaeologists may subsequently dig up the site and provide an accurate picture of the actual village population. This estimate may diverge greatly from the historian's initial theory. This is an example where the discrepancy between the theory and the facts pertains a past event – this discrepancy can still be interpreted using the concept of residual.

they are not. In our example, this explanation implies that if the village could measure also the temperature, the humidity, and the level of nitrate in the terrain, then a perfect match between prediction and outcome could be achieved. In other words, according to this explanation, knowing all natural factors at play (although one might concede that, in practice, knowing all such factors is impossible) would nullify any residual. This mechanistic explanation is characteristic of Newtonian physics as flourished in the 18th and 19th century in Europe, and it can also be recognized in traditions such as Buddhism (Du Sautoy, 2016).³ The second explanation of residuals maintains that randomness is intrinsic to nature: even assuming that all influential natural factors in the universe can be known, some discrepancy between predictions and outcomes will always occur. Having precursors in Greek thinkers such as Democritus and Aristotle (as well as in the notion of *fortuna* in Machiavelli's writings), this probabilistic explanation is at the center of modern quantum mechanics (Du Sautoy, 2016). A third explanation of residuals calls on the supernatural, often represented in the form of divine beings. Here the idea is that powerful supernatural agents (the gods) exist and have the power to influence reality above and beyond any natural factor. According to this view, when outcomes diverge from what can be predicted based on natural factors, it means that gods have intervened. This third explanation is at the center of religion.⁴

Support for the notion that a key difference between mechanistic and religious explanations regards how residuals are interpreted comes from a recent empirical study investigating people's beliefs about ignorance (a concept akin to the idea of residuals) concerning scientific and religious matters (Davoodi & Lombrozo, 2022). The study found that people believe that scientific ignorance occurs when something is currently unknown but can be potentially explained mechanistically with further enquiry. By contrast, the study reports that people believe that, because in the context of religious matters ignorance pertains to divine will, it cannot be mitigated by searching for possible mechanistic explanations.

An instructive anecdote about how religion is grounded on supernatural explanations of residuals comes from the anthropologist Evans-Pritchard's reports on the Azande people of Sudan (Evans-Pritchard, 1937). Evans-Pritchard describes a conversation about an episode where a roof suddenly collapsed, killing the persons dwelling in the building. Azande people commented that those killed in the accident were victims of demons. Armed with the mechanistic outlook characteristic of modern science, Evans-Pritchard remarked that, rather than being caused by demons, the roof collapse was simply caused by termites and rain that had ruined the wood supporting the roof. To this, the Azande interlocutors responded that of course termites and rain were important factors, but why did the roof collapse in the precise moment when such and such was underneath? And not one minute before or after? Only demons, they believed, could explain the coincidence. This anecdote illustrates well the idea that religious interpretations do not ignore natural factors: the Azande people acknowledged the importance of termites and rain. Yet, religion calls on supernatural forces (demons) to account for aspects that remain unexplained by natural factors (why the roof collapsed in a specific moment).

Overall, this argument implies that at the root of religion is the assumption that residuals are caused by supernatural agents. Religious beliefs arise, according to this interpretation, as an attempt to understand the supernatural causes of residuals. To describe these supernatural causes, the proposal is that religion builds a generative model where, in addition to natural factors, supernatural factors are also included. How can supernatural factors be represented within a generative model? Our proposal is that each supernatural agent or god⁵ can be described by one latent variable whose possible states

³The concept of mechanistic explanation does not apply exclusively to modern science, but to any belief system where natural factors are presupposed as causes. These factors can be of various kind (e.g., essences, astrological events, energies) as for example envisaged by traditional, holistic, or alternative belief systems.

⁴In the paper, religion is defined following a common approach in religious studies: it is broadly defined as any belief system characterized by believing in the existence of supernatural agents such as gods, spirits, and demons. Note however that alternative definitions of religion, which do not focus on the theistic aspect, can be found in the literature. On this basis, it is important to stress that the framework developed in the paper is restricted to a deity-based notion of religion.

⁵In the paper the term god refers to any supernatural agent such as spirit, ghost, demon, etc.

correspond to the actions available to the god. To understand this, let us go back to our example of the neolithic village. Imagine that people in the village believe that, in addition to natural factors (e.g., units of labor employed, amount of rainfall, presence of parasites), the agricultural output depends also on the will of the deity of agriculture (e.g., Demeter in ancient Greek). Imagine also that, according to these people, three actions are available to the deity: punish the village, reward the village, and ignore the village. In their generative model, the argument goes, the deity is represented as a variable whose possible states are: punish the village, reward the village, ignore the village. Each state is associated with an effect upon the observed variable, in this example upon the harvest outcome. For example, punishing the village leads to a diminished outcome compared to ignoring, and especially to rewarding, the village. Moreover, each action is associated with a prior probability. For example, a priori people might predict that the deity will be more likely to ignore the village rather than rewarding or punishing it. Once the harvest outcome is observed, the generative model can be employed to infer the posterior probability about the deity's actions. For example, assuming equal prior probability across the three actions, if the harvest is much worse than expected (expected based on natural factors such as units of labor employed, amount of rainfall, presence of parasites), then higher posterior probability will be attributed to punishing the village: "The deity has castigated us," the inhabitants of the village will conclude, "and this is why the harvest is so disappointing."⁶

In summary, following the statistical brain hypothesis, we propose an interpretation of religion as arising to explain away residuals by calling on supernatural factors. According to our argument, in religion a generative model includes latent variables that represent gods and whose states reflect the actions available to those gods. Based on observed variables, the generative model can be employed to infer the actions performed by the gods. This framework rises a fundamental question: why do certain people and cultures adopt a religious outlook, that is, call upon supernatural factors to explain away residuals, while other people and cultures do not? The next section addresses this question.

The origin of religion

When introducing the concept of residuals, we listed three possible attitudes toward them: a mechanistic, probabilistic, and religious attitude. Why does an individual or community embrace a religious attitude toward residuals (where residuals are thought to be caused by gods' actions) rather than a mechanistic (where residuals are thought to be caused by unknown natural factors) or probabilistic one (where residuals are thought to stem from inherent randomness)? This section explores whether the statistical brain hypothesis can offer an answer to this question.

In light of the fact that most societies throughout history have embraced religion, various authors have claimed that evolution has predisposed humans toward religion. Some have argued that a predisposition toward religion emerges because humans have an inbuilt tendency to detect agency in the environment (Atran, 2002; Bloom, 2007; Barrett, 2000; Guthrie, 1993). Yet, while empirical research has indeed documented a human bias for agency detection, evidence that this is at the root of religion is scarce (e.g., Andersen et al., 2019; Maij et al., 2019). Other scholars have claimed that humans are predisposed to religion because, in evolutionary terms, religion is particularly effective in boosting social cohesion (Dunbar, 2022; Wilson, 2002). Whatever the specific explanation proposed, if the hypothesis that humans are predisposed to religion is correct, then the implication is that, in general, humans are inclined to adopt a religious, rather than mechanistic or probabilistic, attitude toward residuals. In other words, people's innate propensity to accept religion may be accompanied by an innate tendency to attribute residuals to unobservable agents or to divine beings.

⁶Once the deity's action has been inferred, the peasants might seek an explanation for it. This typically requires considering how the peasants (or someone else) have behaved toward the deity, for example by assessing whether the right sacrifices have been performed. This sort of inference requires representing how people's actions can affect the gods' will. For the sake of simplicity, here we do not consider this aspect, but an interesting research avenue is to explore this following a statistical outlook analogous to the one adopted here – see also the Discussion.

Yet, not all scholars are persuaded by the claim that humans are predisposed to be religious. Noting that in some societies (especially modern ones) religion plays virtually no role in people's life, some authors have claimed that the context is the major determinant of whether a society becomes religious or not (Inglehart, 2020; Norris & Inglehart, 2011; M. Weber, 1921/1968). This argument raises the following question: what are the contextual conditions that facilitate a mechanistic attitude vis-à-vis a religious attitude toward residuals? We propose the following answer to this question. According to the statistical brain hypothesis, the process of *generalization* plays a fundamental role in shaping people's generative models and beliefs. If previous days were rainy, for example, generalization prescribes that rain will be likely to fall today too. Formally, the statistical brain hypothesis implies that generalizations are the product of prior beliefs. Reflecting events occurred in the past and in other domains, prior beliefs prescribe that similar events are expected to occur also in the present. The role played by generalization has implications for the issue of whether a person will embrace a religious, mechanistic, or probabilistic attitude toward residuals. Consider cultures where people's lives are largely predictable based on mechanistic explanations. Such predictability may be enabled by factors such as predictable environment and food supplies, stable social and political systems, and powerful technologies available. On this basis, a generalization process may lead these cultures to believe that, if mechanistic explanations generally work so well in many domains, then they may also be valid to explain residuals. In other words, the reasoning followed by these cultures may be that, even if a mechanistic explanation is currently unavailable for a phenomenon, this explanation might nonetheless exist and could potentially be discovered. By contrast, consider cultures where people's lives are more unpredictable. Following the same logic, these cultures may believe that, insofar as mechanistic explanations rarely work, these are unlikely to underly residuals. Thus, when confronted with an unexplained phenomenon, these cultures would be more prone to call upon religious interpretations.

To further elucidate this argument, consider an example of someone falling ill without having any mechanistic explanation of the illness. Yet, if the person trusts that a workable mechanistic explanation can be found (e.g., by going to the doctor), the prediction is that the person will not adopt a religious attitude; in other words, she will not look for a supernatural explanation of the illness. Conversely, following the same logic, if the person doubts that a workable mechanistic explanation is available, the person will adopt a religious attitude and look for a supernatural cause of the illness.⁷

The argument just proposed has implications regarding which conditions should favor a mechanistic or religious attitude toward residuals. Specifically, it implies that experiencing predictability should encourage a mechanistic attitude at the expense of a religious one. This is consistent with the observation that people's level of religiosity has widely diminished in modern societies where, thanks to progress in technology, medicine, agriculture, and industry, people's ability to interpret the world mechanistically has greatly improved (Inglehart, 2020; Norris & Inglehart, 2011; M. Weber, 1921/1968). This is also consistent with empirical evidence in social psychology showing that people's tendency to rely on religious interpretations grows when people perceive less control over the environment (Kay et al., 2008, 2009, 2010; Kay, Moscovitch, et al., 2010). Lack of control, according to the statistical brain hypothesis, implies that mechanistic interpretations are less compelling, thus discouraging the application of a mechanistic attitude toward residuals in favor of a religious one.

In conclusion, the statistical brain hypothesis offers a plausible explanation of why individuals and communities may opt for either a mechanistic or religious attitude toward residuals.⁸ While it is possible that evolution has predisposed humans to embrace a religious attitude, the process of

⁷An event is often explained with a combination of both mechanistic and religious factors. In the context of an illness, one might take a prescribed medicine while praying to a god that the medicine will work. Ultimately, explanations like these reveal a religious attitude toward residuals, insofar as they imply that supernatural forces are influential to some degree.

⁸For the sake of simplicity, in this section we have not considered the probabilistic attitude toward residuals. Notably, when looking at history, this appears to be relatively rare. A possibility is that this emerges as a further development of the mechanistic attitude, when a community develops sophisticated statistical models (e.g. quantum physics) of the world and applies these effectively to various scientific problems. At this point, the community may start interpreting residuals as being ultimately due to intrinsic randomness in the world.

generalization may implicate that a mechanistic attitude prevails when sound mechanistic explanations are widely available in other domains. We now proceed by considering another fundamental element of the theory: the role of prior beliefs.

Prior beliefs

Within a statistical framework, religious beliefs result from integrating prior beliefs and novel information. Prior beliefs, hence, play a pivotal role. In the example above, imagine that the prior probability linked with the deity's action of punishing the village is extremely low. In this case, a somewhat meager harvest may be insufficient to draw the conclusion that the deity has chosen to punish the village – an extremely bad harvest may be required to reach this conclusion. This example highlights the impact of prior beliefs on inference. Moreover, prior beliefs may be modulated by the context. To illustrate this point, compare village *x*, which is at peace, versus village *y*, which has waged war against a rival polity. People may believe that the prior probability linked with the deity's action of punishing a village is higher when a village is at war compared to when it is at peace. On this basis, a bad harvest may lead the inhabitants of village *x*, but not of village *y*, to conclude that the deity has chosen to punish them. As this example illustrates, the prior probability can change from one context to the next. The context, furthermore, can modulate the weight of prior beliefs vis-à-vis the weight of novel evidence. For example, in one context people's posterior beliefs may depend strongly on novel evidence and little on prior beliefs (e.g., the inference about the deity's action may be based primarily on observing the harvest, rather than on prior beliefs). By contrast, in another context people's posterior beliefs may depend little on novel evidence and strongly on prior beliefs (e.g., the inference about the deity's action may be based primarily on prior beliefs, rather than on observing the harvest).

Pioneering work has emphasized the key role played by prior beliefs in religion. Taves and Asprem (2017) are among the first who have interpreted religious experiences as events where culturally determined prior beliefs are integrated with novel information. A similar idea has been advanced to explain cases where agency representations are activated in a way that supports religious beliefs (Andersen, 2019). The argument is that perception of agency ensues when sensory information is unreliable, requiring the brain to rely heavily on prior religious expectations to interpret this information. Along these lines, Van Leeuwen and Van Elk (2019) have proposed that culturally-acquired general religious beliefs (e.g., “God exists”) motivate people to seek situations that trigger agency-intuition experiences. According to their argument, such experiences would in turn be integrated with prior beliefs to produce personal religious beliefs (e.g., “God appeared to me last night”).

Prior beliefs may also be crucial for explaining evidence about healing rituals in religious contexts. A recent study has examined believers' views about charismatic prayer healing rituals (Paldam & Schjoedt, 2016). The study found that, in most cases, people believe that these rituals are effective in relieving musculoskeletal pain. The authors interpret these findings as revealing a process akin to placebo analgesia in medicine, where pain diminishes simply by expecting a drug to work as painkiller. A statistical brain hypothesis has been advanced to explain placebo analgesia, asserting that pain perception results from integrating the actual pain signal with prior expectations about the effect of the drug (Büchel et al., 2014; Kiverstein et al., 2022). A similar hypothesis can be proposed to explain the finding that charismatic prayer healing rituals are effective in ameliorating pain.

In summary, a statistical hypothesis of religion is able to explain the key role played by prior beliefs in religious experience, an aspect already emphasized by pioneering literature in the field. Up to this point, we have overviewed the main tenets of a statistical brain framework to explain religion. Can this outlook offer any insight on specific aspects of religion? Let us attempt to address this question by focusing on three typical religious phenomena: miracles, omens, and divination. These will be analyzed in the following section.

Explaining religious phenomena

Miracles

Above, a theory of religion based on the statistical brain hypothesis has been introduced. This section asks whether the theory can offer any insight on key religious phenomena. Here the focus is on miracles. People who interpret an event as a miracle believe that the event violates the rules of nature and reflects the action of a god. Miracles appear to be key concepts in the religion professed by most cultures (Morris, 2006; Shanafelt, 2004). Consider for instance the three Abrahamic traditions (Schulz, 2017). In Judaism, Yahweh is believed to have repeatedly intervened upon nature, such as when parting the Red Sea or when prolonging the daylight to allow the Israelites to conquer Canaan. The Holy Ghost, reflecting God's activity upon nature in Christianity, is thought to be the source of the many miracles performed by saints. According to the Muslim hadith, Muhammad performed various supernatural deeds such as splitting the moon when his followers started to be persecuted. Miracles are not confined to the religions of Abraham, as supported by anthropologists who have recorded beliefs in miraculous events in virtually all religions worldwide (Morris, 2006; Shanafelt, 2004).

To interpret beliefs about miracles, we implement the statistical framework proposed above adopting the Bayesian network formalism (Bishop, 2006). Bayesian networks describe an actor's beliefs about a set of relevant variables. Graphically, they can be represented by circles and arrows, corresponding to the variables and to the conditional dependencies among variables, respectively. We simulate the scenario above where a village inhabitant is trying to infer whether the deity of agriculture has influenced this year's harvest or not. The network representing this scenario is displayed in Figure 1 and includes the following categorical variables: Natural factors⁹ (with possible states being *Favorable* versus *Unfavorable*), Deity (with possible states being *Intervention* versus *No intervention*) and Harvest (with possible states being *Good* versus *Bad*). The arrow projecting from Natural factors to Harvest describes the belief that a good or bad harvest is partly determined by whether natural factors are favorable or unfavorable. The arrow projecting from Deity to Harvest captures the belief that the deity's actions can influence the harvest, too. Note that the Natural Factors and Harvest variables are shaded in Figure 1. This is because these variables are treated as observed (see below). By contrast, the Deity variable is not shaded, indicating that this is a latent variable and therefore unobservable.

Once the network structure is established,¹⁰ the next step is to set up the probabilities linked with the variables. These include three probability distributions:




Figure 1. Bayesian network used in the section about miracles.

⁹The variable *Natural factors* summarizes all known natural aspects influencing the harvest, such as units of labor employed, the amount of rainfall, the presence of parasites, etc. For the sake of simplicity, these are all considered under one single variable (Natural factors)

¹⁰Formally, the structure of this specific network is described by the following joint probability: $P(\text{Natural Factors}, \text{Deity}, \text{Harvest}) = P(\text{Natural Factors}) P(\text{Deity}) P(\text{Harvest} \mid \text{Natural Factors}, \text{Deity})$.

- (1) The prior probability of Natural factors, which was set to $P(\text{Natural factors} = \text{Favorable}) = 0.8$ and $P(\text{Natural factors} = \text{Unfavorable}) = 0.2$;
- (2) The prior probability of Deity, which was set to $P(\text{Deity} = \text{Intervention}) = 0.01$ and $P(\text{Deity} = \text{No Intervention}) = 0.99$; note that here the assumption is that the deity rarely intervenes upon the harvest;
- (3) The conditional probability of Harvest, namely, $P(\text{Harvest} \mid \text{Natural factors}, \text{Deity})$. The values assigned to this are reported in [table 1](#). In short, these values indicate that (i) if the deity intervenes, then the harvest is always good, whatever the natural factors; (ii) if the deity does not intervene and the natural factors are favorable, then the harvest is good 80% of times. Regarding the probability of the harvest being good conditional on no deity's intervention and on unfavorable natural factors ($P(\text{Harvest} = \text{Good} \mid \text{Natural factors} = \text{Unfavorable}, \text{Deity} = \text{No intervention})$), this is indicated as "x" in [table 1](#) since we explored different values of it in [figure 2\(a\)](#).

Once the probability distributions are assigned, the network can be used to make inference. We assumed that the inhabitant has observed both Natural Factors and Harvest as follows: Natural Factors = Unfavorable and Harvest = Good. On this basis, Bayesian inference can be used to estimate the posterior probability for Deity ($P(\text{Deity} = \text{Intervention} \mid \text{Natural factors} = \text{Unfavorable}, \text{Harvest} = \text{Good})$), indicating the belief about whether the deity has intervened or not based on what has been

Table 1. Conditional probability table for $P(\text{Harvest} \mid \text{Natural Factors}, \text{Deity})$ relative to the section concerning miracles.

Natural Factors	Deity	Harvest	$P(\text{Harvest} \mid \text{Natural}, \text{Deity})$
Unfavorable	No intervention	Bad	$1-x$
Favorable	No intervention	Bad	0.2
Unfavorable	Intervention	Bad	0
Favorable	Intervention	Bad	0
Unfavorable	No intervention	Good	x
Favorable	No intervention	Good	0.8
Unfavorable	Intervention	Good	1
Favorable	Intervention	Good	1




Figure 2. Simulations relative to the section about miracles.

observed. We can interpret such posterior probability as reflecting the belief in a miracle. If the posterior probability is high, then the inhabitant believes that a miracle has occurred. Figure 2A plots $P(\text{Deity} = \text{Intervention} \mid \text{Natural factors} = \text{Unfavorable}, \text{Harvest} = \text{Good})$ as a function of the conditional probability $P(\text{Harvest} = \text{Good} \mid \text{Natural factors} = \text{Unfavorable}, \text{Deity} = \text{No intervention})$. The latter can be interpreted as indicating to what extent the good harvest is compatible with the presence of unfavorable natural factors – intuitively, it indicates whether the current natural conditions are sufficient to explain the harvest. As the figure illustrates, the belief in a miracle emerges only if the good harvest is highly incompatible with the presence of unfavorable natural factors (i.e., only if $P(\text{Harvest} = \text{Good} \mid \text{Natural factors} = \text{Unfavorable}, \text{Deity} = \text{No intervention})$ is very low); in other words, only if the current natural conditions cannot account for the good harvest.

Intuitively, the results of this simulation fit with the following interpretation. Consider a condition where all natural factors (the units of labor employed, the amount of rainfall, the presence of parasites) strongly point to a very grim harvest. However, against all odds, the harvest turns out to be excellent. This leads the village inhabitants to attribute a very high posterior probability to the possibility that the deity of agriculture has rewarded them. “This is a miracle!,” the inhabitants will cheer.

The prior probability about the deity’s actions $P(\text{Deity})$ is also critical, and Figure 2(b) explores its role. In this case, we set $P(\text{Harvest} = \text{Good} \mid \text{Natural factors} = \text{Unfavorable}, \text{Deity} = \text{No intervention}) = 0.2$ and assess what happens with different values of $P(\text{Deity} = \text{Intervention})$. As Figure 2(b) illustrates, the belief that a miracle has occurred surges as $P(\text{Deity} = \text{Intervention})$ grows. This captures the idea that, other things being equal, a person is more likely to believe in miracles when miracles are viewed as more frequent a priori.

In several instances, beliefs in miracles do not arise because a person has witnessed an event, but rather because the person has heard about the event from sources such as another person or a book. This is the typical case concerning sacred texts like the Bible, where various miraculous episodes such as the parting of the Red Sea are reported. A statistical framework can be applied also to explain this sort of scenarios. From the perspective of the person hearing about the episode (e.g., the parting of the Red Sea), two forms of prior beliefs are critical. The first pertains to how probable the episode is in its particular context (e.g., given the specific deity, time, and place involved). For example, a person might believe that in ancient times God was much more prone to intervene upon human affairs than today, implying that the possibility of God parting the sea may appear plausible for the past, while appearing absurd for the present. The second relevant prior belief concerns the reliability of the source telling the episode. Sacred texts are imbued with such an authority that religious followers often accept their reports without reservations. The same story may be viewed as undisputable if reported by sacred texts while being scorned as a blatant lie when heard from less trusted sources. Thus, all in all, whether a person believes in a miracle reported by a source depends on integrating two beliefs: the belief on how likely the episode is in its context, and the belief on the source’s reliability.

In short, following a statistical framework, beliefs in a miracle emerge when an event is highly unlikely according to natural factors, and at the same time it is highly compatible with a god’s action. The next section examines the implications of the statistical brain hypothesis for another typical religious phenomenon, that is, the phenomenon in omens.

Omens

Omens can be defined as events interpreted as revealing a god’s intentions regarding an intervention (or lack thereof) upon nature in the future. As for miracles, there is evidence of beliefs in omens in virtually all cultures (Morris, 2006). As one among many examples, in the Roman empire some claimed that the flight of an eagle during the funeral of an emperor signaled that the emperor’s soul was admitted within the gods (Rüpke, 2020).¹¹ The difference between omens and miracles is subtle

¹¹It is debatable whether the belief that this claim was true was widespread. The Roman writer Cassius Dio says the eagle at Augustus’ funeral was released, and a senior magistrate swore he saw Augustus’ soul ascending to heaven (and was paid 1 million sesterces by Augustus’ widow for this!). It was certainly a pretext, but the question of whether people believed it is uncertain.

but substantial: in miracles, the event is important as such (e.g., for our neolithic village, having an excellent harvest is important as such), whereas in omens the event is not important as such, but because it signals a future important event. For example, the village inhabitants might witness a solar eclipse and interpret it as an auspicious omen that, by revealing a positive disposition of the deity of agriculture, anticipates a plentiful harvest. Here the eclipse is not important as such, but because it anticipates an important event (the harvest).

To interpret beliefs in omens according to a statistical outlook, we adopt again the Bayesian network formalism. The new generative model is displayed in [Figure 3](#) and includes also the Eclipse variable (with possible values being *Occurrence* versus *No occurrence*). As the arrows indicate, Eclipse depends both on Deity (capturing the assumption that the deity’s actions can determine whether an eclipse occurs or not) and on Natural factors (Natural factors = Favorable indicates a condition for which, based on natural considerations, occurrence of an eclipse is likely). Note that, for simplicity, the variable Harvest now depends solely on Deity.¹² Note also that now Harvest is not shaded in [Figure 3](#), meaning that it is treated as unobserved – the idea being that whether the harvest will be good or bad is unknown at present, because the time for harvesting is in the future. By contrast, the variable Eclipse is shaded, meaning that this variable has instead been observed.

Regarding the probability distributions, here we use the following: $P(\text{Natural factors} = \text{Favorable}) = 0.8$, $P(\text{Natural factors} = \text{Unfavorable}) = 0.2$; $P(\text{Deity} = \text{Intervention}) = 0.01$; $P(\text{Deity} = \text{No Intervention}) = 0.99$; $P(\text{Harvest} = \text{Good} \mid \text{Deity} = \text{Intervention}) = 1$; $P(\text{Harvest} = \text{Bad} \mid \text{Deity} = \text{Intervention}) = 0$; $P(\text{Harvest} = \text{Good} \mid \text{Deity} = \text{No intervention}) = 0.7$; $P(\text{Harvest} = \text{Bad} \mid \text{Deity} = \text{No intervention}) = 0.3$.

[Table 2](#) reports the conditional probabilities for Eclipse, that is, for $P(\text{Eclipse} \mid \text{Natural factors}, \text{Deity})$. In short, these values indicate that (i) if the deity intervenes, then the eclipse always occurs, whatever the natural factors; (ii) if the deity does not intervene and the natural factors are favorable, then the eclipse occurs 80% of times. Regarding the probability of the harvest being good conditional on no deity’s intervention and on unfavorable natural factors ($P(\text{Eclipse} = \text{Occurrence} \mid \text{Natural factors} = \text{Unfavorable}, \text{Deity} = \text{No intervention})$), this is indicated as “x” in [Table 2](#) since we explored different values of it in [Figure 4\(a\),\(b\)](#)




Figure 3. Bayesian network used in the section about omens.

Table 2. Conditional probability table for $P(\text{Eclipse} \mid \text{Natural Factors}, \text{Deity})$ relative to the section concerning omens.

Natural Factors	Deity	Eclipse	$P(\text{Eclipse} \mid \text{Natural}, \text{Deity})$
Unfavorable	No intervention	No occurrence	1-x
Favorable	No intervention	No occurrence	0.2
Unfavorable	Intervention	No occurrence	0
Favorable	Intervention	No occurrence	0
Unfavorable	No intervention	Occurrence	x
Favorable	No intervention	Occurrence	0.8
Unfavorable	Intervention	Occurrence	1
Favorable	Intervention	Occurrence	1

¹²Formally, the structure of the network is described by the following joint probability: $P(\text{Natural Factors}, \text{Deity}, \text{Eclipse}, \text{Harvest}) = P(\text{Natural Factors}) P(\text{Deity}) P(\text{Eclipse} \mid \text{Natural Factors}, \text{Deity}) P(\text{Harvest} \mid \text{Deity})$.




Figure 4. Simulations relative to the section about omens.

Once the probability distributions are assigned, the network can be used to make inference. We assumed that the inhabitant has observed both Natural Factors and Eclipse as follows: Natural Factors = Unfavorable and Eclipse = Occurrence. On this basis, Bayesian inference can be used to estimate the posterior probability for Deity ($P(\text{Deity} = \text{Intervention} \mid \text{Natural factors} = \text{Unfavorable}, \text{Eclipse} = \text{Occurrence})$), indicating the belief about whether the deity has intervened or not based on what has been observed. We can interpret such posterior probability as reflecting the belief that the eclipse is an omen. If the posterior probability is high, then the actor believes that the eclipse is an omen. Figure 4(a) plots ($P(\text{Deity} = \text{Intervention} \mid \text{Natural factors} = \text{Unfavorable}, \text{Eclipse} = \text{Occurrence})$) as a function of the conditional probability $P(\text{Eclipse} = \text{Occurrence} \mid \text{Natural factors} = \text{Unfavorable}, \text{Deity} = \text{No intervention})$. The latter can be interpreted as indicating to what extent occurrence of the eclipse is compatible with the presence of unfavorable natural factors – intuitively, it indicates whether the current natural conditions are sufficient to explain the eclipse. As the figure illustrates, the belief in an omen emerges only if occurrence of the eclipse is highly incompatible with the presence of unfavorable natural factors (i.e., only if $P(\text{Eclipse} = \text{Occurrence} \mid \text{Natural factors} = \text{Unfavorable}, \text{Deity} = \text{No intervention})$ is very low); in other words, only if the current natural conditions cannot account for the eclipse.

The network can also be used to predict the future harvest by inferring the posterior probability $P(\text{Harvest} = \text{Good} \mid \text{Natural factors} = \text{Unfavorable}, \text{Eclipse} = \text{Occurrence})$. This is displayed in Figure 4(b) as a function of the conditional probability $P(\text{Eclipse} = \text{Occurrence} \mid \text{Natural factors} = \text{Unfavorable}, \text{Deity} = \text{No intervention})$. The trend shown in Figure 4(b) tracks perfectly the trend shown in Figure 4(a), indicating that the prediction that the harvest will be good is stronger when a person believes that an omen has occurred.

Figure 4(c and d) explore the role of $P(\text{Deity})$, namely, the prior probability about the deity's actions. In this case, we set $P(\text{Harvest} = \text{Good} \mid \text{Natural factors} = \text{Unfavorable}, \text{Deity} = \text{No intervention}) = 0.2$ and assess what happens with different values of $P(\text{Deity} = \text{Intervention})$. As Figure 4(c) illustrates, the belief that an omen has occurred surges as $P(\text{Deity} = \text{Intervention})$ grows. This captures the idea that, other things being equal, a person is more likely to believe in omens when omens are viewed as more frequent a priori. Figure 4(d) indicates that, in parallel with strengthening the belief in omens, higher $P(\text{Deity} = \text{Intervention})$ also implies a more optimistic prediction about the harvest.

The discussion about omens raises an important question: what happens when the forecasted event does not occur? For example, what happens if, after the eclipse, the harvest turns out to be dismal? The statistical brain framework implies that this leads to a new inference. Depending on the specific circumstances, the new inference may reach different conclusions. In some cases, the sign previously viewed as omen (e.g., the eclipse) may be retrospectively interpreted as not being an omen after all. For example, the village inhabitants may conclude that, in this specific instance, the eclipse did not signal the deity's will to intervene upon the harvest after all. This may even end up weakening the belief that eclipses are linked with the deity's will to intervene, thus leading to a revision of the generative model. In other cases, despite the forecasted event has failed to materialize, people may nevertheless persist to consider the sign as an omen. In these cases, people often address the problem simply by postponing the forecasted event further to the future. For instance, the village inhabitants may conclude that the eclipse does not forecast this year's harvest, but next years' harvest, which has not occurred yet. Many cases like these are on record. For example, Abrahamic religions have a long history of people claiming to have identified omens revealing the date of the apocalypse (E. Weber, 1999). When their prediction failed to materialize, rather than casting doubts on the alleged omens, many revised their prediction by pushing the date of the apocalypse a few years forward.

The interpretation of omens proposed here can be extended to explain the phenomenon of divination. This will be explored in the following section.

Divination

By allowing people to anticipate events, omens offer precious information to believers. Yet, one key problem with omens is that their occurrence is unpredictable and uncontrollable: in the example of the neolithic village, the eclipse spontaneously happens. This raises an important question: is there any way to create conditions for omens to occur? This question is what has motivated several cultures to develop divination (Morris, 2006). The term divination refers to techniques that are believed to allow omens to occur, and thus to allow gods' intentions about future intervention upon nature to be revealed.¹³ As for omens, divination practices appear in virtually all cultures (Morris, 2006). One of the first historical reports of this sort of practice comes from ancient Chinese civilization (Flad, 2008). Here, rulers of the Shang dynasty presided over a ritual where turtle shells were exposed to fire and where the gods' will was interpreted based on the shape of the ensuing cracks on the shells. Remarkably, records about this practice are the first available instance of writing in Chinese civilization.

Let us propose an interpretation of divination practices based on a statistical framework, relying on the assumption that people involved in developing these practices reasoned like proto-statisticians. Adopting this perspective, an ideal divination method needs to satisfy several conditions. The first is to identify which observable variable should be assessed during divination. This variable should be probabilistic. As an example, consider again our neolithic peasants seeking to predict whether the agricultural deity will ignore, punish, or reward them later during the harvest. Imagine that, to predict this, peasants have developed a divination ritual based on observing the outcome of a fight between a stronger and a weaker dog. Here the observable variable assessed by divination is the outcome of the fight and has three possibilities: the strong dog wins, the weak dog wins, or none wins. Note that this outcome is probabilistic because any of these events could occur. Second, how many states the observable variable can assume should correspond to the number of actions available to the god. In our example, victory of the strong dog, victory of the weak dog, and draw are associated with the

¹³In addition to cases where one attempts to forecast the future, the definition of divination can be extended also to cases where one attempts to gauge the gods' dispositions or actions in the present or even the past. For example, a divination ritual may ask whether a certain deity is angry right now, or whether the bad weather occurring in the present is due to the deity's intervention. Thus, a general definition asserts that divination encompasses any practice aiming at shedding light on the supernatural by producing certain signs in the natural world. The statistical model of divination proposed in this section can be applied to this general definition of divination.

deity's actions of rewarding, punishing, and ignoring the village, respectively. Third, an ideal divination technique will offer an appropriate parallel between the prior probability about the god's actions and the probability of each state of the observable variable. For example, imagine that the prior probability for the deity's action decreases moving from ignoring through rewarding to punishing the village (i.e., peasants believe that usually the deity ignores them rather than rewarding them, and especially rather than punishing them). This should be reflected in the probabilities of the observable variable: for instance, the probability should decrease when moving from a draw through victory of the strong dog to victory of the weak dog. Finally, an ideal divination technique should control all natural forces that affect the observable variable, in such a way that the outcome of the divination ritual should reflect solely the deity's influence. For example, the two dogs should eat an equal amount of food, should rest the same number of hours, and should be placed in two equally favorable spots at the start of the fight; in this way, the outcome of the fight can be interpreted not as depending on such natural factors, but only on the deity's will – thereby allowing peasants to guess such will.

The picture of divination just sketched is ideal: in other words, it describes how an optimal divination technique should look like according to a statistical perspective. When looking at divination practices as expressed by real people, a discrepancy with such ideal should be expected. Nonetheless, our prediction is that at least some of the elements highlighted above should be identifiable in at least some divination rituals adopted by various cultures. This would suggest that, at least to some degree, actual divination rituals are informed by the statistical considerations highlighted here.

When assessed in the context of empirical evidence, our statistical interpretation appears to be particularly appropriate to describe *cleromancy*, a divination method based on observing the outcome of a probabilistic variable (Morris, 2006). Several cultures have independently developed cleromancy practices to gain insight on the gods' disposition. In ancient Rome, an example of this is the *Sortes*, a ritual where words written on tablets were drawn as lots (Luijendijk & Klingshirn, 2018). The gods' will was said to be consulted in this way in Italian temples such as Praeneste and Caere. The Roman historian Tacitus reports similar methods among the ancient German people (Tacitus, *Germania* 10). Examples of casting lots as a means to reveal the will of God are reported in the Bible, such as when Joshua had to divide the land among the seven tribes of Israel who had not received any inheritance (Joshua 18:6). A rich cleromancy tradition can be identified in Chinese civilization too, tracing back to the turtle-shell oracle mentioned above and reaching its pinnacle with the I Ching text, one of the five Confucian Classics (Flad, 2008). Analogous techniques have been observed, among others, in India, Japan, among the Yoruba of Africa, and among the Mi'kmaq of North America. Consistent with our statistical interpretation of divination, cleromancy requires to identify a probabilistic variable to be assessed (e.g., the lottery), with a predefined set of possible outcomes (e.g., the lots). Also consistent with our proposal, cleromancy maps each possible outcome with one specific divine action: for instance, in the example taken from the Bible above (Joshua 18:6), the order of lots drawn identifies unambiguously the God's prescription concerning how to distribute land among the Israelites. Finally, as implied by our proposal, cleromancy employs highly formalized rituals where the influence of all known natural factors is controlled, in such a way that the outcome of the ritual can be unambiguously interpreted as a sign of divine will.

Altogether, our analysis suggests that, at least in the case of cleromancy, our proposal is consistent with actual divination rituals. One last aspect remains to be assessed through: our proposal postulates a parallel between the prior probability about the god's actions and the probability of each state of the observable variable. Can this parallel be identified in any actual divination practice? To answer this question, consider a form of divination performed in ancient Rome where the eating habits of sacred chickens were observed before battle (Cicero, *On the Nature of the Gods* 2.7). The assumption was that, if the chickens ate the food thrown for them, this signaled the gods' approval for the battle to take place (it did not guarantee victory, but at least signaled the gods had provided no obstacle). Here, the observed variable had two possible states: eating versus not eating. The Romans knew that the two possibilities had very different probabilities: the chicken, typically hungry, rarely refused food. Implicit to this is the assumption that, according to the Romans, the gods typically approved the battle to take

place. This is an example of a parallel between the prior probability about the god's actions (i.e., the gods' approval was deemed to be far more likely) and the probability of each state of the observable variable (i.e., eating was deemed to be far more likely too). The chickens' refusal to eat was so rare that neglecting it was considered to be foolish. The story of a commander who, ignoring the chickens' refusal to eat and fighting anyway, suffered a devastating defeat, was remembered as a cautionary tale for subsequent generations of Roman commanders (Cicero, *On the Nature of the Gods* 2.7).

Another example of a parallel between the prior probability about the god's actions and the probability of each state of the observable variable comes again from ancient Rome. Here, examination of sacrificial victims was another common method of gauging the gods' disposition. After being sacrificed, the animal victims were cut open and their organs examined for signs of abnormality or disease (Lennon, 2014; Scheid, 2003; Turfa, 2006). If everything was in order, the sacrifice had been accepted by the gods and business could proceed. This divination practice focused on observing a probabilistic variable with two possibilities: presence versus absence of abnormality. The probability of each was not equal, but much higher for the latter – i.e., abnormalities were rare occurrences. This reflects the Romans' implicit assumption that usually the gods were believed to be well disposed. Once again, we have a parallel between the prior probability about the gods' actions (i.e., the gods' approval was deemed to be far more likely than disapproval) and the probability of each state of the observable variable (i.e., absence of abnormality was deemed to be far more likely too). Presence of abnormality, reflecting the gods' bad disposition, was rare and, when it occurred, extremely concerning. In a fictional reconstruction of the outbreak of the civil war between Caesar and Pompey, the poet Lucan painted a picture of a calamitous series of events culminating in a failed sacrifice in which every part of the victim showed signs of disease and poisonous decay, interpreted as a sign that the gods of the Underworld had entered into the body of the bull and were the source of the dire warning contained within the organs (Lucan, *Pharsalia* 1.609–37).¹⁴

With this analysis of divination, we conclude our investigation of how a statistical framework can be employed to interpret important religious phenomena (including also miracles and omens). Of course, miracles, omens and divination represent only a small subset of religious phenomena. Still, the scheme proposed here may offer a useful perspective to understand other cases. For instance, ethnographic records indicate that religious interpretations are sometimes advanced to explain why a person does something wrong, breaks a taboo, or violates social norms. A statistical inference scheme akin to the one used above may clarify why these religious interpretations emerge, for instance when natural factors are unlikely to explain a person's bad behavior and therefore hint to divine intervention. While exploration of other religious phenomena is left to future research, the next section examines the statistical framework by asking whether this can offer any insight about the nature of religious beliefs.

Explaining religious beliefs

When overviewing the statistical brain hypothesis above, we highlighted three key characteristics of generative models. First, because latent variables extract the underlying structure of reality, one latent variable usually maps to multiple observable variables. Second, latent variables are often organized along a hierarchy spanning from more abstract to more specific variables. Third, generative models are formed according to the principle of action-specificity, that is, they focus on aspects that are salient for

¹⁴In Rome, the recognized experts of this kind of divination were the *haruspices*. They could be summoned by magistrates to help with the interpretation of portents and signs (typically, involving perceived distortions of the natural order such as rains of blood, earthquakes, the discovery of ill-omened animals within the city, and other occurrences judged to be especially out of the ordinary). The *haruspices* would not only confirm that the gods were indeed angry, they could also offer advice about the causes of divine displeasure and the potential remedies that should be sought. A bronze model of a cross-section of an animal's liver, found at Piacenza in Etruria, the region from which the *haruspices* were said to originate, was inscribed with the names of a number of Etruscan deities in carefully defined sections, and may have been a tool used for identifying which specific deities were voicing their displeasure or rejection of the sacrifice (Edmonds, 2019, pp. 207–9).

the behavior of an agent. What are the implications of these characteristics for our analysis of religion? How would these characteristics look like within religious generative models? And can these characteristics be observed in religions embraced by real people? Here we address these questions.

Gods control multiple domains

This section examines the notion that, in generative models, there is usually a one-to-many mapping between latent and observable variables. Applying this to religious generative models, where each latent variable represents a god, the implication is that each god will influence multiple observable variables, that is, it will influence multiple aspects of reality. This prediction appears to be confirmed empirically insofar as, in most religions, each god typically controls diverse reality domains, and not only one. This is especially evident in monotheistic faiths, where reality in its totality is subject to the power of one single divinity. Polytheistic faiths also appear to conform to this prediction. In Greek mythology, Apollo presides, to name only a few domains, over archery, music, dance, truth, prophecy, healing, disease, sun, and light (Eidinow & Kindt, 2015). No less numerous were the spheres controlled by the Egyptian deity Isis; these include, but are not limited to, motherhood, magic, death, the sky, the moon, and wisdom (Teeter, 2011). This is reflected also in the various epithets that a god may take on. So, while Romans might refer simply to Jupiter in a general sense, he could also appear as Jupiter Optimus Maximus (Best and Greatest), Jupiter Latiaris (god of the Latins), Jupiter Tonans (the thunderer), Jupiter Stator (the Stayer), Jupiter Victor, and numerous others. Considering a contemporary example, the Yoruba religion of West Africa attributes to the god Oshosi the control of hunt, forest, health, meals, the arts, and contemplation (Beier, 1980).

According to a statistical interpretation of religion, the fact that each god is linked with many domains is not accidental, but it is a necessary consequence of how the brain works. Specifically, it is the consequence of entertaining generative models where latent variables, corresponding to gods in the context of religion, capture the hidden structure of reality and thus reveal themselves in multiple spheres. An obvious question arises from this argument: why is a god linked with certain domains and not others? For instance, why is Apollo associated with the sun and with healing, and not, say, with the forest or with eros? The answer offered by a statistical interpretation is that every culture attempts to recognize the pattern of correlations underpinning observable variables (in statistics, this is akin to what is done by techniques for dimension reduction such as factor analysis). When, the argument goes, two or more aspects are believed to be correlated, a latent variable is evoked to explain them. Applying this reasoning to the god Apollo, for whatever reason the Greeks came to the conclusion that archery, music, dance, truth, prophecy, healing, disease, sun, and light were all linked together, and were thus under the control of the same divine power (i.e., of the same latent variable): the god Apollo. If he was unhappy, the Greeks thought, then usually all these domains suffered: archers missed their targets, prophecies were withheld, and diseases spread. If, on the contrary, Apollo was appeased, then archers thrived, prophecies were revealed, and health was guaranteed.¹⁵

Hierarchical pantheons

This sections investigates the second key characteristic of generative models: their hierarchical structure. In the context of religion, the ensuing implication is that religious generative models will be organized along a hierarchy where more abstract and more powerful gods stand above more specific and less powerful deities. This prediction appears to be confirmed when looking at most world religions. A revealing anecdote about this can be found in the Upanishad, one of the sacred texts of the Hindu tradition (Flood, 1996).

¹⁵According to the statistical framework introduced here, the correlation among the different life spheres does not need to be perfect. For example, the Greeks arguably believed that archery, prophecy, and health were correlated (and thus all controlled by the same deity: Apollo), but not perfectly correlated. This means that Apollo's anger, despite often targeting all three domains together, did not do this always; sometimes it might target only one or two domains, for example only health but not archery and prophecy.

Here, when asked by a disciple how many gods there are in the universe, the guru Yajnavalkya replied: “There are 3306 gods.” This was a conventional answer in Hindu tradition that left the disciple unsatisfied: “Of course,” he rebuked, “But how many gods there are truly in the universe?.” “Thirty-tree” was the next reply. “Of course,” the disciple repeated, “in one sense there are thirty-three gods. But how many gods there are truly in the universe?.” “Six” was the next guru’s answer. The disciple still remained unhappy and thus continued asking the same question again and again. Every time, the guru replied with a smaller number, until he finally declared: “There is only one true god.” The traditional interpretation of this story is that Hinduism understands the divine realm as encompassing a hierarchy where at the top stands one single supernatural being: Brahman, the ultimate divine reality underlying everything in the universe (Flood, 1996). At a lower hierarchical level, Brahman manifests itself in the form of more specific deities such as Shiva, Vishnu, and Brahma. Despite being still very abstract and powerful, these gods do not control reality in its totality, but only some aspects of it. For instance, Shiva exerts his control upon domains that include destruction, yoga, meditation, the arts, fertility, and medicine. These deities can in turn express themselves at an even lower hierarchical level in the form of more specific gods or avatars (i.e., incarnations of gods). For instance, Pashupati, the lord of animals in Hindu tradition, is believed to be an avatar of Shiva. The idea of avatar is especially prominent for cults worshipping Vishnu, who allegedly reincarnated in popular divine figures such as Rama, Buddha, and Krishna.

A hierarchical organization of gods was present also in the Greco-Roman religion (Rüpke, 2020). Here the pantheon was akin to a patriarchal hierarchy dominated by a sky god (Zeus-Jupiter) and encompassing several generations of deities. Here too, by moving down the hierarchy, one encounters deities associated with decreasing power and more specific domains. For example, the goddess Aphrodite, the deity of love broadly defined, was the mother of multiple Erotes, each linked with a specific aspect of love or sex.

We have just seen that Hinduism and ancient Greco-Roman religions provide examples supporting the idea of a hierarchy of gods. Most other polytheistic faiths appear also to be consistent with this idea. But what about monotheistic religions? How can the idea of a single god be compatible with the notion of a hierarchy? This question can be answered by looking at how monotheistic religions really work. Take Christianity. Here, although ultimately there is only one God, Christian believers accept the existence of a plethora of supernatural beings such as the virgin Mary, the devil, angels, demons, and the saints, just to name a few (Brown, 2014). These are all hierarchically subordinated to God, and yet at the same time they are independent entities that can influence specific domains of reality. For example, each saint is attributed supernatural powers in specific domains, like Saint Nicholas who is considered to be the patron of sailors, merchants, archers, and students (Brown, 2014). Similar considerations apply to other monotheistic faiths such as Islam and Judaism, where, despite one single God being at the center, a careful scrutiny reveals the presence of numerous supernatural beings hierarchically subordinated to God.

Gods are action-specific

Action-specificity is the last key characteristic of generative models highlighted in the paper. This prescribes that latent variables have a specific focus on aspects of reality that are salient for the behavior of an agent. To assess whether this characteristic can be found in religion as empirically observed, it is instructive to compare societies that rely on different subsistence strategies. The prediction is that each society will believe in gods that control spheres linked with the specific subsistence method adopted by that society (see also Campbell, 1969). This prediction appears to be confirmed in various cultures. Consider the Algonquian tribes dwelling in circumpolar North America and subsisting primarily on hunting. In the religion of these tribes, a central role is played by the masters of animals, that is, by gods controlling the abundance of games available for hunting (Martin, 2001). Compare this with Inuit peoples living further North and relying primarily on fishing. Consistently, the central figure in Inuit religion, called Sedna, is the goddess of the sea and of marine animals (Martin, 2001). A connection between methods of subsistence and religion emerges also when

considering agricultural societies. Here, a prominent role in religion and rituals is played by the god or goddess of agriculture (Randa, 2009).

In support of the principle of action-specificity is also the observation that, in the Catholic church, veneration of saints has over time adapted to the social context (Brown, 2014). An example is Saint Christopher who, traditionally venerated as the patron saint of travelers, has nowadays become the protector of motorists and train conductors. Similarly, Saint Valentin was initially honored for his martyrdom, but became the patron saint of lovers after his figure was revisited by poets who sought to make sense of the phenomenon of courtly love in the Middle Ages.

We have now completed our overview concerning how a statistical framework can be employed to interpret religion. We have examined how religious generative models should look according to a statistical outlook, and we have adopted this approach to interpret important religious phenomena such as miracles, omens, and divination. We next analyze our proposal in light of previous literature and pinpoint novel insights that can be highlighted.

Contribution to the literature

Inspired by contemporary research in cognitive science, this paper introduces a statistical inference framework to explain religion. Does this approach offer any new insight? To answer this question, it is useful to evaluate our proposal in the context of prior theories of religion. Numerous theories have been developed to explain religious beliefs and practices. A broad classification of these distinguishes between epistemic theories, claiming that religion represents an attempt to explain reality (Barrett, 2000; Boyer, 1994; Guthrie, 1993; Iannaccone et al., 1998; McCauley, 2017; Spilka et al., 2019; Stark & Finke, 2000; Tylor, 1987), and motivation theories, positing that religion arises to satisfy other motives such as sedating anxiety (Atran & Norenzayan, 2004; Hume, 1757; Jong & Halberstadt, 2017; Kay, Gaucher, et al., 2010), fostering community bonds and moral values (Batson et al., 1993; Baumeister, 1991; Durkheim, 1912; Krause & Wulff, 2005; McKay & Whitehouse, 2015; Pargament et al., 1983; Saroglou, 2011), or promoting self-interest (Gramsci, 1971; Jost et al., 2014). Available empirical evidence indicates that epistemic and non-epistemic motives are both at play in religion (Saroglou, 2011). By focusing on how people explain reality, our statistical interpretation of religion can be classified as an instance of epistemic theories. However, the role of other motives is not necessarily incompatible with our proposal; simply, in order to account for such motives, our theory is not sufficient, but it needs to be expanded further. Thus, the question of whether our proposal offers any new insight can be reformulated as the question of whether it offers any new insight specifically about epistemic aspects of religion. To answer this, let us consider the key tenets advocated by previous epistemic theories of religion.

The origin of the epistemic view of religion can be traced back at least to the writings of the French Enlightenment thinker Bernard Fontanelle (Manuel, 1959). While his contemporaries viewed the religions embraced by indigenous non-European cultures as manifestations of human irrationality, Fontanelle instead interpreted these as attempts made by intelligent creatures to understand how reality works. Indigenous people, argued Fontanelle, lacked the knowledge and technology of modern science. Thus, they had to ground their explanation of reality on something else. What else? A rational approach is to rely on what is already known and to translate this knowledge to more general domains. Applying this logic, indigenous people knew from experience that their actions could produce effects in the environment. Therefore, they probably reasoned that there might be supernatural agents whose actions elicit similar changes in the environment. This reasoning was, according to Fontanelle, at the root of religion.

The perspective pioneered by Fontanelle has since inspired countless theories of religion. Not without important distinctions, these theories all agree with Fontanelle's central idea that one of the primary motivations underlying religion is to explain how reality works. According to this view, religious beliefs reflect cognitive models describing how divine agents act upon nature. This general view is shared also by our statistical inference theory of religion. So, what does the latter add to previous proposals? We argue that the key novel insight is the idea, largely neglected so far, that religious models are *probabilistic*. Describing religious models as probabilistic might offer at

least two benefits. First, it offers an intriguing perspective to explain the key logic underlying religion. The argument presented above, interpreting religion as an attempt to explain away residuals, requires a probabilistic perspective. Furthermore, the ensuing interpretations of miracles, omens, divination, and of the nature of religious beliefs, also presuppose a probabilistic framework. The second benefit of a statistical theory of religion is that it allows to establish links with research about other psychological phenomena for which a probabilistic interpretation is already well established (Clark, 2013; Friston & Kiebel, 2009; Knill & Pouget, 2004; Oaksford & Chater, 2007; Rao & Ballard, 1999). This opens up the opportunity to translate knowledge on these phenomena to the study of religion.

The present paper is not the first one in the literature where a statistical inference framework is employed to interpret religion (see Schjoedt, 2018). Among the first who did so is a paper of Schjoedt et al. (2013). Employing a Bayesian framework, that paper argues that, because cognitive resources are typically depleted during religious rituals, people do not interpret these rituals based on their idiosyncratic experience, but, rather, based on accounts crafted by religious authorities. Along similar lines, Taves and Asprem (2017) have interpreted religious experiences as events where culturally determined prior beliefs are integrated with novel information, an idea subsequently applied to explain the role of agency detection in the formation of religious beliefs (Andersen, 2019; Van Leeuwen & Van Elk, 2019). A statistical outlook has also inspired a neuroscientific framework explaining how inferential processes implemented in specific brain regions produce typical religious and spiritual phenomena such as mystical experiences (van Elk & Aleman, 2017). Our proposal is deeply indebted to this prior research. Much of the previous work has focused on the notion that religious experiences ensue from integrating prior beliefs and novel evidence (e.g., agency detection), and on clarifying how specific circumstances (e.g., ambiguous stimuli, sensory deprivation, suggestive contexts) boost the weight of prior beliefs thereby facilitating the occurrence of religious experiences (Andersen, 2019; Taves & Asprem, 2017; Van Leeuwen & Van Elk, 2019). The present paper builds on this prior research by exploring the general theoretical foundations upon which a statistical theory of religion may be grounded. It does so by addressing the following basic questions: why are religious generative models sometimes envisaged by people? And what are the characteristics of such generative models? Moreover, we further contribute by developing a statistical interpretation of specific phenomena such as beliefs in miracles, omens, and divination.

Note that various previous works adopting a statistical framework to study religion have employed predictive coding (Schjoedt, 2018), which is a theoretical approach within the broader family of statistical theories used in cognitive sciences. Predictive coding makes some specific assumptions about the variables at play (treated as continuous) and the inference method (variational inference) (Friston & Kiebel, 2009; Rao & Ballard, 1999). Predictive coding is a very powerful framework and has provided great insight in the field of the psychology of religion. Still, to express the arguments proposed in the present paper, we rely upon basic concepts common to all statistical theories (e.g., the notion of generative models, latent variables, residuals, etc.), rather than on concepts specific to predictive coding – we believe that this more basic approach is better suited to explore the foundations of a statistical theory of religion.¹⁶

In short, a statistical inference view of religion can integrate previous epistemic theories by highlighting the importance of treating religious models as probabilistic. This might contribute to shed light on the cognitive processes underlying religion and it might bridge research on religion with research on other domains where a statistical outlook is already well developed.

¹⁶Since we do not use a predictive coding framework, we do not adopt concepts that are specific to this framework. An example is the concept precision, which is mathematically defined as the inverse of the variance and is used in predictive coding during inference to establish the relative weight of the prior mean vis-à-vis novel evidence.

Discussion

Following the logic advocated by the statistical brain hypothesis, this paper asks whether any insight can be gained by interpreting religious beliefs as expression of statistical inference processes. The proposal is that religious ideas arise to explain residuals, that is, to explain discrepancies between observations and predictions (the latter based on natural factors). The framework postulates probabilistic generative models where gods are described by latent variables whose possible states correspond to the actions available to the gods. Thus, when an event occurs, generative models can be employed to infer the gods' actions responsible for the event. As examined in the paper, our proposal offers a plausible interpretation of typical religious phenomena such as miracles, omens, and divination, as well as of important characteristics of religious beliefs (specifically, the notion that gods control multiple spheres of reality, are organized hierarchically, and conform to the principle of action-specificity). In this final section, we will consider corollaries of the theory that are important to discuss.

Acquiring religious generative models

The paper has focused on the following questions: why do people sometimes rely on religious generative models? What are the characteristics of religious generative models? And how do they produce phenomena such as beliefs in miracles, omens, and divination? Our analysis has so far neglected a fundamental question: what are the processes whereby religious generative models are developed by people? In other words, how do some people come to believe that, say, the god Apollo can behave in such and such manner and control such and such reality domains? Though a comprehensive examination of this question is beyond the scope of the manuscript, a tentative answer can be outlined. Cognitive science research suggests that generative models can be acquired in two ways: via direct interaction with the environment and via socialization (Schwartz, 1989; Zentall & Galef, 1988). On this basis, broadly speaking, religious generative models might arise from how a community of people interact with the environment and from how members of the community interact among themselves (Banerjee & Bloom, 2013; Sherkat, 2003). When children are raised within a community, they might progressively acquire religious beliefs from interacting with adult members and, by interacting with the environment in a creative way, they might partially transform their beliefs and in turn influence the beliefs of other members. Of course, this explanation is only sketched; a promising research endeavor is to employ a statistical inference framework to examine precisely how people acquire religious generative models.

Epistemic motives are not enough

As highlighted above, a statistical framework appears to be well suited to explain epistemic processes underlying religion. However, as also stressed above, epistemic facets appear to be only part of the picture. Research has identified motives such as sedating anxiety (Atran & Norenzayan, 2004; Hume, 1757; Jong & Halberstadt, 2017; Kay, Gaucher, et al., 2010), fostering community bonds and moral values (Batson et al., 1993; Baumeister, 1991; Durkheim, 1912; Krause & Wulff, 2005; McKay & Whitehouse, 2015; Pargament et al., 1983; Saroglou, 2011), and promoting self-interest (Gramsci, 1971; Jost et al., 2014) as important too. Thus, a question left open is how precisely a statistical approach to religion can be extended to encompass processes that are not epistemic. A recent paper has attempted to address this question (Rigoli, 2021). Instead of relying on the notion of inference, this previous proposal suggests that a Bayesian decision process is what drives religious beliefs (Rigoli, 2021). In this framework, probabilistic inference processes remain important, but at the same time motives such as sedating anxiety, fostering bonds and values, and fulfilling self-interest, play a role too. Reconciling this Bayesian decision approach with the analysis of religious inference processes developed in this paper represents a promising research endeavor. Another intriguing possibility is to develop an active inference model of religion where epistemic processes are reconciled with other

motives. Active inference is an influential theoretical framework employed in cognitive science to integrate Bayesian inference processes with motivational and behavioral ones (Friston et al., 2015, 2017), and therefore may help achieving a similar integration in the context of religion.

The role of religious practices

The paper focuses on religious beliefs about the gods and their actions. However, religious beliefs also encompass knowledge about how humans can interact with the gods. In the example of our neolithic village, peasants might believe that the deity of agriculture can be appeased by sacrificing animals during rituals. An intriguing research avenue is to employ a statistical framework to explore beliefs about how humans can influence the gods' will.

Relatedly, the paper focuses only on beliefs. These are of course important, but they are not the whole story: religion also encompasses a complex system of practices. A detailed examination of the relationship between beliefs and practices is beyond the scope of the paper. As a first approximation though, there is probably a bidirectional influence whereby not only beliefs guide practices, but practices in turn shape beliefs (Argyle, 2006). For example, children often find themselves involved in religious practices before understanding the deep meaning of these practices. Here participating in the religious ceremonies is itself a key factor that shapes the formation of religious beliefs (Argyle, 2006). An intriguing question for future research is whether a statistical framework can shed any light on how religious generative models are shaped by religious practices and in turn guide people's behavior in these practices.

Automatic and controlled processes

An influential distinction in the social sciences pertains to the opposition between controlled processes, which are conscious, effortful, and flexible, versus automatic processes, which, by contrast, are unconscious, effortless, and rigid (Kahneman, 2011). Do the mental processes analyzed in the paper belong to the first or to the second category? Generally speaking, the statistical brain hypothesis has been advanced to explain mental phenomena that encompass both categories (Clark, 2013; Friston & Kiebel, 2009; Knill & Pouget, 2004; Oaksford & Chater, 2007; Rao & Ballard, 1999). Indeed, from a statistical brain perspective, it is possible to identify instances of mental inference that are effortless, unconscious, and rigid, such as during perception, as much as it is possible to identify instances that are effortful, conscious, and flexible, such as during complex forms of reasoning and decision making. On this ground, the theory we propose in the paper may apply to automatic religious inferences and to controlled religious inferences alike. Examples of the former are mystical experiences, in which certain sensory signals originate the perception of the divine (Rigoli, 2023; Taves & Asprem, 2017). Examples of the latter are theological discourses, which are built upon sophisticated lines of argumentation and reasoning.

Conclusion

Scholars have warned the cognitive science community about the danger of fragmentation, that is, the danger of relying on theories about micro-phenomena that remain disjointed from other equally narrow theories (Goertzen, 2008). In recent times, the statistical brain hypothesis has offered a promising platform to avoid this danger (Chater et al., 2006; Oaksford & Chater, 2007). Armed with the idea that, in multiple domains, dealing effectively with uncertainty is a fundamental imperative for the brain, this framework has strived to identify the unifying principles that underly diverse cognitive phenomena, spanning from perception (Knill & Pouget, 2004; Rao & Ballard, 1999) to social interaction (Devaine et al., 2014). Inspired by prior work (Andersen, 2019; Schjoedt et al., 2013; Taves & Asprem, 2017; van Elk & Aleman, 2017; van Elk et al., 2016), our paper contributes to this endeavor by arguing that a statistical framework offers a plausible interpretation of religion, or at least of important aspects thereof.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

ORCID

Francesco Rigoli  <http://orcid.org/0000-0003-2233-934X>

Data availability statement

Not applicable since the paper does not analyze any empirical data.

References

- Andersen, M. (2019). Predictive coding in agency detection. *Religion, Brain & Behavior*, 9(1), 65–84. <https://doi.org/10.1080/2153599X.2017.1387170>
- Andersen, M., Pfeiffer, T., Müller, S., & Schjoedt, U. (2019). Agency detection in predictive minds: A virtual reality study. *Religion, Brain & Behavior*, 9(1), 52–64. <https://doi.org/10.1080/2153599X.2017.1378709>
- Argyle, M. (2006). *Religious behaviour*. Routledge.
- Atran, S. (2002). *In gods we trust: The evolutionary landscape of religion*. Oxford University Press.
- Atran, S., & Norenzayan, A. (2004). Religion's evolutionary landscape: Counterintuition, commitment, compassion, communion. *Behavioral and Brain Sciences*, 27(6), 713–730. <https://doi.org/10.1017/S0140525X04000172>
- Banerjee, K., & Bloom, P. (2013). Would tarzan believe in god? Conditions for the emergence of religious belief. *Trends in Cognitive Sciences*, 17(1), 7–8. <https://doi.org/10.1016/j.tics.2012.11.005>
- Barrett, J. L. (2000). Exploring the natural foundations of religion. *Trends in Cognitive Sciences*, 4(1), 29–34. [https://doi.org/10.1016/S1364-6613\(99\)01419-9](https://doi.org/10.1016/S1364-6613(99)01419-9)
- Barsalou, L. W. (2008). Grounded cognition. *Annual Review of Psychology*, 59(1), 617–645. <https://doi.org/10.1146/annurev.psych.59.103006.093639>
- Batson, C. D., Schoenrade, P., & Ventis, W. L. (1993). *Religion and the individual: A social-psychological perspective*. Oxford University Press.
- Baumeister, R. F. (1991). *Meanings of life*. Guilford press.
- Beier, U. (Ed.). (1980). *Yoruba myths*. CUP Archive.
- Bishop, C. M. (2006). *Pattern recognition and machine learning*. springer.
- Bloom, P. (2007). Religion is natural. *Developmental Science*, 10(1), 147–151. <https://doi.org/10.1111/j.1467-7687.2007.00577.x>
- Bowers, J. S., & Davis, C. J. (2012). Bayesian just-so stories in psychology and neuroscience. *Psychological Bulletin*, 138(3), 389. <https://doi.org/10.1037/a0026450>
- Boyer, P. (1994). *The naturalness of religious ideas: A cognitive theory of religion*. Univ of California Press.
- Brown, P. (2014). *The cult of the saints: Its rise and function in Latin Christianity*. University of Chicago Press.
- Büchel, C., Geuter, S., Sprenger, C., & Eippert, F. (2014). Placebo analgesia: A predictive coding perspective. *Neuron*, 81(6), 1223–1239. <https://doi.org/10.1016/j.neuron.2014.02.042>
- Campbell, J. (1969). *The masks of god: Primitive mythology*. Viking Press.
- Chater, N., Tenenbaum, J. B., & Yuille, A. (2006). Probabilistic models of cognition: Conceptual foundations. *Trends in Cognitive Sciences*, 10(7), 287–291. <https://doi.org/10.1016/j.tics.2006.05.007>
- Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, 36(3), 181–204. <https://doi.org/10.1017/S0140525X12000477>
- Dalton, A. (2004). *The Lenape of Pennsylvania*. and Ontario. The Rosen Publishing Group, Inc.
- Davoodi, T., & Lombrozo, T. (2022). Varieties of ignorance: Mystery and the unknown in science and religion. *Cognitive Science*, 46(4), e13129. <https://doi.org/10.1111/cogs.13129>
- Devaine, M., Hollard, G., Daunizeau, J., & Beck, J. (2014). The social Bayesian brain: Does mentalizing make a difference when we learn? *PLOS Computational Biology*, 10(12), e1003992. <https://doi.org/10.1371/journal.pcbi.1003992>
- Dunbar, R. (2022). *How religion evolved: And why it endures*. Oxford University Press.
- Durkheim, E. (1912). *The elementary forms of the religious life*. London: George Allen & Unwin Ltd.

- Du Sautoy, M. D. (2016). *What we cannot know: Explorations at the edge of knowledge*. Fourth Estate.
- Edmonds, R. G. (2019). *Drawing down the moon: Magic in the ancient Greco-Roman world*. Princeton University Press.
- Eidinow, E., & Kindt, J. (Eds.). (2015). *The oxford handbook of ancient Greek religion*. OUP Oxford.
- Evans-Pritchard, E. E. (1937). *Witchcraft, oracles and magic among the azande* (Vol. 12). Oxford.
- Flad, R. K. (2008). Divination and power: A multiregional view of the development of oracle bone divination in early China. *Current Anthropology*, 49(3), 403–437. <https://doi.org/10.1086/588495>
- Flood, G. D. (1996). *An introduction to Hinduism*. Cambridge University Press.
- Freud, S. (1927). *The future of an illusion*. Broadview Press.
- Friston, K. (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1456), 815–836. <https://doi.org/10.1098/rstb.2005.1622>
- Friston, K., FitzGerald, T., Rigoli, F., Schwartenbeck, P., & Pezzulo, G. (2017). Active inference: A process theory. *Neural Computation*, 29(1), 1–49. https://doi.org/10.1162/NECO_a_00912
- Friston, K., & Kiebel, S. (2009). Predictive coding under the free-energy principle. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364(1521), 1211–1221. <https://doi.org/10.1098/rstb.2008.0300>
- Friston, K., Rigoli, F., Ognibene, D., Mathys, C., Fitzgerald, T., & Pezzulo, G. (2015). Active inference and epistemic value. *Cognitive Neuroscience*, 6(4), 187–214. <https://doi.org/10.1080/17588928.2015.1020053>
- Gershman, S. J., & Niv, Y. (2010). Learning latent structure: Carving nature at its joints. *Current Opinion in Neurobiology*, 20(2), 251–256. <https://doi.org/10.1016/j.conb.2010.02.008>
- Gershman, S. J., Norman, K. A., & Niv, Y. (2015). Discovering latent causes in reinforcement learning. *Current Opinion in Behavioral Sciences*, 5, 43–50. <https://doi.org/10.1016/j.cobeha.2015.07.007>
- Goertzen, J. R. (2008). On the possibility of unification: The reality and nature of the crisis in psychology. *Theory & Psychology*, 18(6), 829–852. <https://doi.org/10.1177/0959354308097260>
- Gramsci, A. (1971). *Selections from the prison notebooks*. Lawrence and Wishart.
- Guthrie, S. E. (1995). *Faces in the clouds: A new theory of religion*. Oxford University Press.
- Hume, D. (1757). *The natural history of religion*. A. and H. Bradlaugh Bonner.
- Iannaccone, L. R., Stark, R., & Finke, R. (1998). Rationality and the “Religious Mind”. *Journal of Economic Literature*, 36(3), 1465–1495. <https://doi.org/10.1111/j.1465-7295.1998.tb01721.x>
- Inglehart, R. F. (2020). *Religion's sudden decline: What's causing it, and what comes next?* Oxford University Press.
- Jong, J., & Halberstadt, J. (2017). What is the causal relationship between death anxiety and religious belief? *Religion, Brain & Behavior*, 7(4), 296–298. <https://doi.org/10.1080/2153599X.2016.1249916>
- Jost, J. T., Hawkins, C. B., Nosek, B. A., Hennes, E. P., Stern, C., Gosling, S. D., & Graham, J. (2014). Belief in a just god (and a just society): A system justification perspective on religious ideology. *Journal of Theoretical and Philosophical Psychology*, 34(1), 56. <https://doi.org/10.1037/a0033220>
- Kahneman, D. (2011). *Thinking, fast and slow*. macmillan.
- Kay, A. C., Gaucher, D., McGregor, I., & Nash, K. (2010). Religious belief as compensatory control. *Personality and Social Psychology Review*, 14(1), 37–48. <https://doi.org/10.1177/1088868309353750>
- Kay, A. C., Gaucher, D., Napier, J. L., Callan, M. J., & Laurin, K. (2008). God and the government: Testing a compensatory control mechanism for the support of external systems. *Journal of Personality & Social Psychology*, 95(1), 18. <https://doi.org/10.1037/0022-3514.95.1.18>
- Kay, A. C., Moscovitch, D. A., & Laurin, K. (2010). Randomness, attributions of arousal, and belief in god. *Psychological Science*, 21(2), 216–218. <https://doi.org/10.1177/0956797609357750>
- Kay, A. C., Whitson, J. A., Gaucher, D., & Galinsky, A. D. (2009). Compensatory control: Achieving order through the mind, our institutions, and the heavens. *Current Directions in Psychological Science*, 18(5), 264–268. <https://doi.org/10.1111/j.1467-8721.2009.01649.x>
- Kiverstein, J., Kirchhoff, M. D., & Thacker, M. (2022). An embodied predictive processing theory of pain experience. *Review of Philosophy and Psychology*, 13(4), 973–998. <https://doi.org/10.1007/s13164-022-00616-2>
- Knill, D. C., & Pouget, A. (2004). The Bayesian brain: The role of uncertainty in neural coding and computation. *Trends in Neurosciences*, 27(12), 712–719. <https://doi.org/10.1016/j.tins.2004.10.007>
- Krause, N., & Wulff, K. M. (2005). Church-based social ties, a sense of belonging in a congregation, and physical health status. *The International Journal for the Psychology of Religion*, 15(1), 73–93. https://doi.org/10.1207/s15327582ijpr1501_6
- Latourette, K. S., & Winter, R. D. (1975). *A history of Christianity* (Vol. 2). Harper & Row.
- Lennon, J. J. (2014). *Pollution and religion in ancient Rome*. Cambridge University Press.
- Luijendijk, A., & Klingshirn, W. E. (Eds.). (2018). *My lots are in thy hands: Sortilege and its practitioners in late antiquity*. Brill.
- Maij, D. L., Van Schie, H. T., & Van Elk, M. (2019). The boundary conditions of the hypersensitive agency detection device: An empirical investigation of agency detection in threatening situations. *Religion, Brain & Behavior*, 9(1), 23–51. <https://doi.org/10.1080/2153599X.2017.1362662>
- Manuel, F. E. (1959). *The eighteenth century confronts the gods*. Harvard University Press.
- Martin, J. W. (2001). *The land looks after us: A history of native American religion*. Oxford University Press.

- McCauley, R. N. (2017). Twenty-five years in: Landmark empirical findings in the cognitive science of religion. *Religion Explained? The Cognitive Science of Religion After Twenty-Five Years*, 19(3), 17–55. <https://doi.org/10.4013/fsu.2018.193.07>
- McKay, R., & Whitehouse, H. (2015). Religion and morality. *Psychological Bulletin*, 141(2), 447. <https://doi.org/10.1037/a0038455>
- Morris, B. (2006). *Religion and anthropology: A critical introduction*. Cambridge University Press.
- Norris, P., & Inglehart, R. (2011). *Sacred and secular: Religion and politics worldwide*. Cambridge University Press.
- Oaksford, M., & Chater, N. (2007). *Bayesian rationality: The probabilistic approach to human reasoning*. Oxford University Press.
- Paldam, E., & Schjoedt, U. (2016). Miracles and pain relief: Experienced health effects of charismatic prayer healing in a large collection of Christian testimonies. *Archive for the Psychology of Religion*, 38(2), 210–231. <https://doi.org/10.1163/15736121-12341322>
- Pargament, K. I., Silverman, W., Johnson, S., Echemendia, R., & Snyder, S. (1983). The psychosocial climate of religious congregations. *American Journal of Community Psychology*, 11(4), 351–381. <https://doi.org/10.1007/BF00894054>
- Ramstead, M. J., Kirchhoff, M. D., & Friston, K. J. (2020). A tale of two densities: Active inference is enactive inference. *Adaptive Behavior*, 28(4), 225–239. <https://doi.org/10.1177/1059712319862774>
- Randa, D. G. (2009). *Fertility gods and goddesses of ancient agricultural societies*. California State University.
- Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1), 79–87. <https://doi.org/10.1038/4580>
- Rigoli, F. (2021). A computational perspective on faith: Religious reasoning and Bayesian decision. *Religion, Brain & Behavior*, 11(2), 147–164. <https://doi.org/10.1080/2153599X.2020.1812704>
- Rigoli, F. (2023). The computations underlying religious conversion: A bayesian decision Model. *Journal of Cognition and Culture*, 23(1–2), 241–257. <https://doi.org/10.1163/15685373-12340161>
- Rigoli, F., Pezzulo, G., Dolan, R., & Friston, K. (2017). A goal-directed Bayesian framework for categorization. *Frontiers in Psychology*, 8, 408. <https://doi.org/10.3389/fpsyg.2017.00408>
- Rüpke, J. (2020). *Pantheon: A new history of Roman religion*. Princeton University Press.
- Saroglou, V. (2011). Believing, bonding, behaving, and belonging: The big four religious dimensions and cultural variation. *Journal of Cross-Cultural Psychology*, 42(8), 1320–1340. <https://doi.org/10.1177/0022022111412267>
- Scheid, J. (2003). *An introduction to Roman religion* (tr. J. Lloyd). Edinburgh University Press.
- Schjoedt, U. (2018). Predictive coding in the study of religion: A believer's testimony. In A. K. Petersen, G. Ingvild Sælid, L. H. Martin, J. S. Jensen, & J. Sørensen (Eds.), *Evolution, cognition, and the history of religion: A new synthesis* (pp. 364–379). Brill.
- Schjoedt, U., Sørensen, J., Nielbo, K. L., Xygalatas, D., Mitkidis, P., & Bulbulia, J. (2013). Cognitive resource depletion in religious interactions. *Religion, Brain & Behavior*, 3(1), 39–55. <https://doi.org/10.1080/2153599X.2012.736714>
- Schulz, H. (2017). The problem of miracles: Jewish, Christian, and Islamic perspectives. *Toronto Journal of Theology*, 33(S1), 1–7. https://doi.org/10.3138/tjt.33.suppl_1.1
- Schwartz, B. (1989). *Psychology of learning and behavior*. WW Norton & Co.
- Shanafelt, R. (2004). Magic, miracle, and marvels in anthropology. *Ethnos*, 69(3), 317–340. <https://doi.org/10.1080/0014184042000260017>
- Sherkat, D. E. (2003). Religious socialization: Sources of influence and influences of agency. In M. Dillon (Ed.), *Handbook of the Sociology of Religion* (pp. 151–163). Cambridge: Cambridge University Press.
- Shermer, M. (2002). *Why people believe weird things: Pseudoscience, superstition, and other confusions of our time*. Macmillan.
- Simon, H. A. (1997). *Models of bounded rationality: Empirically grounded economic reason* (Vol. 3). MIT press.
- Spilka, B., Shaver, P. R., & Kirkpatrick, L. A. (2019). A general attribution theory for the psychology of religion. In B. Spilka (Ed.), *The psychology of religion* (pp. 153–170). Routledge.
- Stark, R., & Finke, R. (2000). *Acts of faith: Explaining the human side of religion*. Univ of California Press.
- Taves, A., & Asprem, E. (2017). Experience as event: Event cognition and the study of (religious) experiences. *Religion, Brain & Behavior*, 7(1), 43–62. <https://doi.org/10.1080/2153599X.2016.1150327>
- Teeter, E. (2011). *Religion and ritual in ancient Egypt*. Cambridge University Press.
- Turfa, J. M. (2006). Etruscan religion at the watershed. In C. E. Schultz & P. B. Harvey (Eds.), *Religion in republican Italy* (pp. 62–89). Cambridge University Press.
- Tylor, E. B. (1871). *Primitive culture: Researches into the development of mythology, philosophy, religion, art, and custom* (Vol. 2). J. Murray.
- van Elk, M., & Aleman, A. (2017). Brain mechanisms in religion and spirituality: An integrative predictive processing framework. *Neuroscience & Biobehavioral Reviews*, 73, 359–378. <https://doi.org/10.1016/j.neubiorev.2016.12.031>
- van Elk, M., Friston, K., & Bekkering, H. (2016). The experience of coincidence: An integrated psychological and neurocognitive perspective. In K. Landsman, & E. van Wolde (Eds.), *The challenge of chance*. The Frontiers Collection (pp. 171–185). Cham: Springer.
- Van Leeuwen, N., & Van Elk, M. (2019). Seeking the supernatural: The interactive religious experience model. *Religion, Brain & Behavior*, 9(3), 221–251. <https://doi.org/10.1080/2153599X.2018.1453529>

- Weber, E. (1999). *Apocalypses: Prophecies, cults, and millennial beliefs through the ages*. Harvard University Press.
- Weber, M. (1921/1968). *Economy and society*. Bedminster Press.
- Wilson, D. (2002). *Darwin's cathedral: Evolution, religion, and the nature of society*. University of Chicago press.
- Wolpert, L. (1994). *The unnatural nature of science*. Harvard University Press.
- Zentall, T. R., & Galef, B. G. (Eds.). (1988). *Social learning: Psychological and biological perspectives*. Psychology Press.