


## Supernatural agents and prosociality in historical China: micro-modeling the cultural evolution of gods and morality in textual corpora

Ryan Nichols, Edward Slingerland, Kristoffer Laigaard Nielbo, Peter Kirby & Carson Logan


To cite this article: Ryan Nichols, Edward Slingerland, Kristoffer Laigaard Nielbo, Peter Kirby & Carson Logan (2020): Supernatural agents and prosociality in historical China: micro-modeling the cultural evolution of gods and morality in textual corpora, Religion, Brain & Behavior, DOI: [10.1080/2153599X.2020.1742778](https://doi.org/10.1080/2153599X.2020.1742778)

To link to this article: <https://doi.org/10.1080/2153599X.2020.1742778>

 View supplementary material [↗](#)

 Published online: 30 Mar 2020.

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

 Article views: 38

 View related articles [↗](#)

 View Crossmark data [↗](#)



# Supernatural agents and prosociality in historical China: micro-modeling the cultural evolution of gods and morality in textual corpora

Ryan Nichols<sup>a,b</sup>, Edward Slingerland<sup>c,d</sup>, Kristoffer Laigaard Nielbo<sup>e,f</sup>, Peter Kirby<sup>g</sup>, and Carson Logan<sup>h</sup>

<sup>a</sup>Department of Philosophy, California State University, Fullerton, CA, USA; <sup>b</sup>Centre for Human Evolution, Cognition, & Culture, University of British Columbia, Vancouver, BC, Canada; <sup>c</sup>Department of Asian Studies, University of British Columbia, Vancouver, BC, Canada; <sup>d</sup>Center for Advanced Study in the Behavioral Sciences, Stanford University, Stanford, CA, USA; <sup>e</sup>Interacting Minds Centre, School of Culture and Society, Aarhus University, Aarhus, Denmark; <sup>f</sup>Department of History, University of Southern Denmark, Odense, Denmark; <sup>g</sup>Computer Science, Georgia Institute of Technology, Atlanta, GA, USA; <sup>h</sup>Computer Science, University of British Columbia, Vancouver, BC, Canada

## ABSTRACT

A major source of attention paid to high gods in the fields of cultural evolution and cognitive science is the social effects of belief in high gods. Belief in high gods is both hypothesized to catalyze a cognitive punishment-avoidance mechanism at the level of individual minds, and a group cultural evolutionary mechanism that amplifies in-group cooperation. Recent research into non-Western contexts not only indicates a multiplicity of supernatural influences on the individual-level and group-level mechanisms but raises questions about theoretical presuppositions about how a supernatural agent is classified as a high god or as something else. Our exploratory study operationalizes the question “Does historical China have high gods?” through the assessment of semantic associations between each of several supernatural agent categories (alleged high gods, low gods, ancestors, sage kings, and emperors) and each of several social functional content categories (punishment, reward, morality, monitoring, and religion). Analyzing collocations in a corpus of 5.7 m Chinese characters, representing all of the most influential historical Chinese-language texts, our preliminary results suggest social functions of supernatural agents in historical China were widely distributed across many species of supernatural agent thereby complicating a claim that high gods constitute a special category in relation to these social functions.

## ARTICLE HISTORY

Received 27 August 2018  
Accepted 31 January 2020


## KEYWORDS

Data mining; cognitive science of religion; cultural evolution; corpus linguistics; China; gods; ancestors

## 1. Introduction

What has come to be referred to as “moralizing high god theory” purports to explain the cultural evolutionary success and growth through time of groups that are committed to a class of gods with particular properties. Moralizing high god theory contends that big gods, along with other factors, generate big societies. This section highlights three components of moralizing high god theory (MHGT): an explanatory mechanism at the level of individuals’ minds, an explanatory mechanism at the level of the group, and a presupposition about high gods themselves.

**CONTACT** Ryan Nichols  [rnichols@fullerton.edu](mailto:rnichols@fullerton.edu)

 Supplemental data for this article can be accessed <https://doi.org/10.1080/2153599X.2020.1742778>

© 2020 Informa UK Limited, trading as Taylor & Francis Group

At the individual level, a cognitive mechanism operates by taking as inputs beliefs about traits of an agent's god and yielding as outputs an increase in the agent's punishment-avoidance behavior. Consensus in MHGT indicates that the semantic content of the beliefs input into this mechanism ranges over three key traits predicated of an agent's high god. High gods can punish wrongdoers, have an interest in human morality, and are constantly monitoring behavior (Shariff et al., 2010). When received as inputs, these semantic contents produce punishment avoidance cognition, emotion and behavior. (See Galen (2016, p. 29) on "semantic priming of religion" studies.)

According to MHGT, when a sufficient but unspecified ratio of in-group members activate this cognitive mechanism with sufficient frequency, social effects result. Specifically, transaction costs and in-group defections are lowered and public trust and group solidarity are raised. Between-group competition contributes to cultural group selection processes in which some groups survive and succeed while others do not. Group competition involves factors such as warfare, economic production, demographic expansion, cultural transmission and more (Atran & Henrich, 2010). According to MHGT, groups whose members are committed to the same high god disproportionately succeed in the process of cultural group selection because they exhibit more unity and in-group cooperation. Further, since streams of demic and cultural transmission merge, successful high god groups will transmit to future generations the religious ideas, rituals, and behaviors responsible for making them successful.

As the terms "high god" and "big god" suggest, MHGT has presupposed a distinction between the set of supernatural agents that qualifies as high gods and the set that does not. However, preliminary studies of MHGT tested for correlates between prosocial behavior and belief in *any supernatural agent at all*. For example, Shariff and Norenzayan (2007) used sentence unscramble tasks to prime supernatural concepts generally and found confirmation for a positive relationship between priming with supernatural concepts and increased prosociality in an economic game. This contrasts with study designs that include independent experimental conditions for each species of supernatural agent (ghosts, gods, angels, etc.). Here we assist in the conceptual clarification of MHGT. An efficient way to diagnose a possible confusion attending MHGT is by explicitly identifying the following presupposition:

Belief in a high god, but not belief in a set of low gods, catalyzes the individual punishment-avoidance mechanism and cultural group selection mechanism that are jointly responsible for growing big societies.

This presupposition made by MHGT's advocates, and associated conceptual issues, are the focus of this paper.

Consider that most supernatural agents are able either to monitor people, or possess the power to punish, or care about human morality, even if most supernatural agents cannot do all three. Within Roman religion, among many other deities, the goddess Poena punishes and the goddess Cura rewards. Cura, also a goddess, cares about human morality. Depending on where an agent is, a god will be able to watch a person's behavior. For example, if someone is in Rome on Palatine Hill, the goddess Palatua will monitor them. Now suppose that a group of people (in Rome) are committed to belief in all four of these lowly deities in the Roman pantheon. This state of affairs strongly appears to activate both individual and group mechanisms at work in MHGT. Despite this, the high god concept at the root of prominent formulations of MHGT (see below) does not appear to leave room for diffusion of supernatural responsibility of this sort.

MHGT has been subjected to a stream of recent criticism (e.g., Baumard et al., 2015; Galen, 2016; Ma-Kellams & Blascovich, 2013; McKay & Whitehouse, 2016; Welch et al., 2007; Whitehouse et al., 2019). Some of this work questions the causal role of prosocial supernatural punishment in general (Whitehouse et al., 2019; but see critiques in Beheim et al., 2020 and Slingerland et al., 2019), whereas other work suggests that the functions attributed to high gods might be more diffusely distributed among multiple types of supernatural agents.

Consider an analysis of Viking religion by Raffield et al. (2019), which further demonstrates how functions attributed to high gods can be operationalized in a variety of ways or even outsourced to

other supernatural agents. The highest god in the Viking pantheon is Odin, who himself is neither omnipresent nor omniscient. Yet the fact that any stranger requiring hospitality might be Odin in disguise conveys on him a degree of functional omnipresence; the fact that Odin deploys an army of ravens to observe people wherever they might be give him functional omniscience. These are all important correctives or extensions to MHGT. The research group behind MHGT has been actively exploring functional alternatives to big gods, such as the role played by impersonal cosmic forces such as karma (White et al., 2019).

This project stands aside from these types of critiques because it is interested in comparing relationships between a variety of supernatural agents and a variety of social-functional traits. The first goal is to offer a preliminary answer to the question “What are the roles of gods in early China with regard to punishment, reward, monitoring and morality?” Since Chinese civilization is the biggest, most populous, and one of the most historically extant societies in our species, it is difficult to imagine a more promising culture within which to explore the “big gods for big societies” hypothesis of MHGT. Providing considerable interest to this goal, some researchers hypothesize that historical Chinese moral psychology stands apart. A preliminary corpus-based project (Nichols & Logan, 2017) found that alleged Chinese high gods had stronger semantic associations to punishment than did Chinese low gods, offering partial support for expectations drawn from MHGT. This study, however, only compared two species of supernatural agent, which inherits inferential limitations akin to those of dichotomous cross-cultural studies that compare only two cultural samples. If historical China effectively used a wider set of supernatural agents to improve in-group cooperation, Nichols and Logan (2017) were not positioned to find out. Although its coverage of current scholarly opinion is only fragmentary, querying the variable, “A supreme high god is present” in the Database of Religious History (DRH; religiondatabase.org), from 3000 BCE to 220 CE in the region of China (live link to query: <https://religiondatabase.org/visualize/1515/>) returns 81% “yes” (13 answers) and 19% “no” (3 answers). Nonetheless, noteworthy experts say that, compared to historical Western cultures, Chinese culture appears based on non-divine sources such as “Big Gobs” (Huebner & Sarkissian, 2016; Sarkissian, 2015) and an embedded, well-transmitted kinship ethic known as “filial piety” or *xiao* 孝 (Nichols, 2013). Other sinologists say China is unique because it lacked Western-style high gods (Hall & Ames 1995). Yet other sinologists say China had high gods (Clark & Winslett, 2011). Yet a third group says China seems to have had high gods, but adds that they were not *moralizing* high gods (Dutton & Madison, 2016).

## 2. According to moralizing high god theory, what exactly is a “high god” anyway?

The second goal of this paper is to provide a social-functional taxonomy of “species” within the supernatural agent “genus”. As of now, the question “What non-question-begging method can researchers use to determine which supernatural agent qualifies as a high god and which does not?” is without answer. In any experimental test of MHGT, researchers make presuppositions about who is and who is not a high god or, put otherwise, about *highgodship*. But the necessary conditions under which supernatural agent A possesses highgodship for population P have yet to be enumerated with care. This has led to significant indeterminacy about the scope of the high god concept at the heart of MHGT, making a data-driven taxonomy of supernatural agents of potential benefit.

A preliminary response to this objection would be to say that vagueness in the concept of a “high god” is bound to arise when a variety of non-collaborating research groups all work on the same new issue in any soft science. So, to begin to understand the importance of this issue, consider uses of terms “high god” and its synonym “big god” as they appear in a small sampling of papers by collaborating authors in the same research group. Table 1 provides four different quoted definitions for “high god” (in Column 1) authored by members affiliated with the University of British Columbia research group. Problematic features of definitions recorded in Table 1 are by no means unique to definitions provided by the UBC research group.

**Table 1.** Definitions of “high god”.

Conditions	Necessary, Sufficient, Both, or Neither?	Logical Implications
(1) High Gods are “gods who (1a) cared about cooperative – and harmony – enhancing behavior (the group’s moral norms), (1b) could and would reward and punish appropriately, and (1c) had the power to monitor all behavior all the time” (Shariff et al., 2010, 124).	<ul style="list-style-type: none"> <li>Unclear</li> <li>Probably not necessary</li> </ul>	<ul style="list-style-type: none"> <li>If (1a)-(1c) are intended as necessary conditions, then (c) implies all high gods must be omnipresent. This narrows highgodship.</li> <li>(1a) and (1b) refer to properties possessed by degree, deferring unanswered questions. E.g. how often must a god reward appropriately to qualify as a high rather than low god?</li> </ul>
(2) (2a) High gods are the gods “of the major world religions” (2012, 3273) High gods “exhibit ideal qualities for effective punishment: they are (2b) omniscient (perfect monitors), (2c) morally infallible (often the ultimate source of norms) and (2d) infinitely powerful (making punishment for the wicked seem guaranteed)” (Laurin et al., 2012, 3273)	<ul style="list-style-type: none"> <li>(2a) necessary</li> <li>(2b) unclear, perhaps sufficient?</li> </ul>	<ul style="list-style-type: none"> <li>(2a) Historical Chinese religion is not a world religion, therefore Chinese gods like Shangdi and Tian are not high gods.</li> <li>(2b) If intended as a necessary condition, then a high god must have maximalized social-functional properties. This narrows highgodship.</li> <li>(2a) and (2b)-(2d) are inconsistent with (3).</li> </ul>
(3) Big gods (a) “must exercise supernatural powers by transcending limitations that otherwise constrain human abilities”, (b) “must also share enough properties of human mind perception to make them intuitively graspable and emotionally potent,” and must (c) “have privileged access specifically to behaviors that have moral consequences.” (d) A “feature of potent Big Gods is spatial position” (Norenzayan, 2013, 26, 29)	<ul style="list-style-type: none"> <li>(3a)-(3c) are necessary</li> <li>(3d) is neither necessary nor sufficient</li> <li>Unlikely (3a)-(3d) are jointly sufficient</li> </ul>	<ul style="list-style-type: none"> <li>(3a)-(3c) refer to properties possessed by degree, deferring unanswered questions. E.g. how much access must a god have to qualify as a high rather than a low god?</li> <li>(3a)-(3d) nowhere refer to punishment</li> <li>(3d) Are high gods are omnipresent? Unknown.</li> <li>(3d) Is the set of “potent big gods” a subset of the class of “big gods”? Unknown</li> <li>(3a)-(3d) are inconsistent with (2a) and (2b)</li> </ul>
(4) (4a) MHGT refers to a “multidimensional continuum of supernatural agents in which Big Gods occupy a particular corner of the space” (Norenzayan et al., 2016, 6). (4b) The concept of “high gods” refers to “powerful, morally concerned deities who are believed to monitor human behavior. These gods are believed to deliver rewards and punishments ...” (3).	<ul style="list-style-type: none"> <li>(4a) Talk of a continuum suggests conditions are neither necessary nor sufficient</li> <li>(4b) Unclear or neither</li> </ul>	<ul style="list-style-type: none"> <li>(4) is presented explicitly in terms of the <i>beliefs of people</i> rather than <i>metaphysical traits</i>, making it inconsistent with (1), (2) and (3).</li> <li>(4b) refers to properties possessed by degree, deferring unanswered questions. E.g. how powerful must a god be to qualify for highgodship?</li> <li>(4a) and (4b) set a lower bar for highgodship than other definitions.</li> </ul>

The first question a scientist might have about definitions of terms that are explicitly used in hypotheses is whether the definition provides necessary conditions (if A then B), sufficient conditions (A only if B), both necessary and sufficient conditions, or neither type of condition. Definitions recorded in Column 1 of Table 1 generally lack the detail needed to clarify whether the conditions they state are necessary conditions, sufficient conditions, both, or neither. Clarification on this point is essential since, without it, the term’s meaning within a hypothesis is unclear. The scientific importance of this point is worth elaboration. Consider these statements:

- Neither necessary nor sufficient: “High gods care about human morality”;
- Necessary condition: “If god X cares about human morality, then X is a high god”; and
- Necessary and sufficient condition: “X is a high god if and only if X cares about human morality.”

Statement (b) supports the following inference: if during testing one finds a god that cares about human morality, it is therefore a high god. As long as researchers’ interests in hypothesis testing is

restricted to identifying who is a member of the set of high gods, then the specification in (b) is adequate. Statement (c) provides a necessary and sufficient condition for highgodship and supports the following inference: if during testing one finds a god that cares about human morality, it is therefore a high god, and also, if one finds a god that does not care about human morality, it is therefore not a high god. Statement (c) is more helpful to the scientist since (c) supports inferences about which gods are members of the set of high gods and which gods are not. In contrast, (a) inhibits scientific progress because (a) does not provide any testable condition for highgodship. Because a sufficient condition for highgodship in this context is likely composed of a set of necessary conditions, we say no more about sufficient conditions here.

Information in Table 1, Column 3 concerns the logical implications of the analyses of the term “high god”. These analyses are assumed to be offered by authors for scientific rather than other purposes making it reasonable, if not requisite, to hold them to a higher standard than dictionary definitions. To say that conditions (2b)–(2d) in analysis (2) are inconsistent with conditions (3a)–(3d) in analysis (3) is to imply that the set of members in the set of high gods according to analysis (2) is non-identical to membership in the set of high gods according to analysis (3). Put succinctly, (2) and (3) offer incompatible accounts of highgodship. An implication of this type of inconsistency is that it is possible that an empirical or experimental test using analysis (2) confirms MHGT but that a similar test using analysis (3) disconfirms MHGT. In conclusion, without clear conditions for highgodship, MHGT’s presupposition cannot be properly evaluated and appears vulnerable.

Three of the four analyses in Table 1 are noteworthy because they predicate epistemic and metaphysical traits of high gods. Only (4b) is stated in terms of *the semantic contents of beliefs about high gods*. Direct discussion of traits of gods (e.g., high gods are omnipresent) gives the discussions a theological tenor. For the sake of explaining the operation of the individual-level cognitive mechanisms promoting punishment-avoidance behavior hypothesized by MHGT, its analyses must specify semantic contents of persons’ beliefs. A charitable reading of these analyses would recognize this goal as the intent of authors of these analyses (though, for concepts at the core of one’s theory, making things explicit is surely not too much to ask). But offering theological rather than psychological definitions has led to a false assumption with which advocates of MHGT have been slow to reckon.

Empirical studies of theological correctness investigated semantic contents of the god-beliefs of their participants. This body of research shows that *religious believers routinely fail to operationalize the metaphysical and epistemological traits that analyses found in Table 1 predicate of high gods*. For example, Barrett and Keil (1996) showed that though believers adopted theologically correct beliefs when asked formally (i.e., they reported believing God is omniscient), when participants were presented with vignettes they defaulted to anthropomorphized, theologically incorrect god concepts (i.e., they affirmed beliefs that implied that God was not omniscient; see Barrett & Keil, 1996, p. 239). Analyses and definitions of high gods contained in research papers about MHGT, including those found in Table 1, state for example that a high god has “the power to monitor all behavior all the time”, i.e., is omniscient. Studies of theological incorrectness show that believers blithely disregard the semantic content of their formal high god concepts in favor of semantic contents that far better describe low, local, limited gods (e.g., God doesn’t know everything, God can’t be in two places at once, etc.). This state of affairs raises, at least somewhat, the likelihood that high god concepts found in Table 1, Column 1, underwriting MHGT’s account of the individual-level cognitive mechanism for punishment avoidance are psychologically unrealistic.

The lack necessary or sufficient conditions, the change over time, and the inconsistency across definitions might explain a set of forceful, though not entirely clear, criticisms of MHGT. Brazil and Farias frame the problem facing moralizing high god theory as a question-begging one. They charge that Norenzayan (referring to Norenzayan et al., 2016) “seems to evade the question that is begging to be answered: Why would anyone want to believe in Big (rather than small) Gods?” (Brazil & Farias, 2016, p. 24). This rhetorical question obscures any underlying reasoning. In terms introduced in this paper, Brazil and Farias appear to be wondering why small gods do not qualify for highgodship according to some of MHGT’s own characterizations of that concept. It seems, after all, that small gods are sufficient to activate the individual and group level mechanisms posited by MHGT.

Another research group has been instrumental in crafting an alternative to moralizing high god theory (Johnson, 2015; Watts et al., 2015). Watts, Bulbulia, Gray and Atkinson were motivated in part because they believe that “clarification is needed about the causal role that different kinds of gods have played in the evolution of big societies” (Watts et al., 2016, p. 41). This comment voices skepticism about the assumption that the genus of supernatural agents is unproblematically presupposed to be described as a bimodal distribution, wherein high gods and low gods are at opposite ends of an axis. High gods and low gods, however, are not distinct natural kinds, and, to extend the metaphor, the phylogenetic branches between them are not empty. This research group subsequently developed a theory, “broad supernatural punishment theory,” that competes with MHGT.

The emergence of the broad supernatural punishment theory alternative to MHGT is due in part to the conceptual problems surrounding highgodship just discussed. “Broad supernatural punishment theory” states that belief in supernatural agents that monitor, punish and care about human morality enhances cultural evolution of groups. Unlike MHGT, broad supernatural punishment theory does not state that these functions must be carried out by a high god (Atkinson et al., 2015; Watts et al., 2016; Watts et al., 2015). Varieties of supernatural agents, in various mixes, can together have as much or more effect on the individual and group level mechanisms as might an Axial Age high god leading a monotheistic religion. Johnson (2015) has provided additional theoretical clarification on the content and hypothetical implications of the two theories.

This micro-modeling approach, and its use of techniques of corpus linguistics, affords the opportunity to turn nagging conceptual questions on their head. This is because methods enable description of *the degree to which* a given species of supernatural agent (ancestor deities, alleged high gods, etc.) is semantically associated with a given trait of social functional interest to the cognitive mechanism hypothesized by MHGT. That is, this method supports assignment to each supernatural agent (or, for simplicity, to each subspecies of supernatural agent) an individualized rate at which it is associated with punishment, with reward, with monitoring, and with morality. This informs discussion of supernatural agents and cultural evolution with more granularity about social-functional traits of supernatural agents than is possible with other methods.

### 3. Methods

#### 3.1. Data

The test corpus contains 96 texts drawn from the Chinese Text Project corpus. The Chinese Text Project is a full-text database for ancient and medieval Chinese texts ([www.ctext.org](http://www.ctext.org)) containing documents from before the Warring States period (>480 BCE), the Han Dynasty (206 BCE – 220 CE), and from the post-Han period to the Song Dynasty (960–1279 CE). See Table 2. At data retrieval, the corpus contained a total of 5,742,339 Chinese tokens distributed over 15,928 types.<sup>1</sup> See Appendix 1 for texts, genres and era-dates. Cleaning, tokenization, keyword identification (target terms) and collocation extraction (co-occurrences of pairs of target terms) was done with custom scripts in Python.<sup>2</sup>

#### 3.2. Coding and synonym groups

To avoid producing categories of terms with theoretical biases, semantic categories produced independently by University of Heidelberg’s Thesaurus Linguae Sericae project were identified then

**Table 2.** Corpus composition by era.

Era	Dates	Token Count	Percent of Corpus
Pre-Warring States	Before 480 BCE	30,447	0.53
Warring States	479–222 BCE	1,424,080	24.79
Han	221 BCE–220 CE	3,501,256	60.9
Post-Han to Song	221 CE–1044 CE	786,546	13.6
Totals		5,742,329	100

explored.<sup>3</sup> The Thesaurus Linguae Sericae (TLS) project gathered international teams of scholars together to mount a unique attempt at linguistic taxonomy of classical Chinese. Its goal was to create a conceptual map of ancient Chinese thought. These categories, falling under TLS's "synonym groups" and "lexeme groups" classifications, included sets named *deity*, *high god*, *punishment*, *reward*, and *ancestor*. (See Appendix 3 for TLS lists.)

Despite this attempt, unmodified TLS lists were unusable in the machine learning processes used here. (See Appendix 2 for the lists used in studies reported here.) This is due to the highly polysemous nature of classical Chinese characters. An example helps illustrate this problem. The TLS *deity* list includes terms such as *feng* 風 and *he* 河. These occasionally refer to minor gods, gods of the wind and gods of the river, respectively. However, far more common use of these terms is to refer simply to winds and rivers. Absent the ability to distinguish the polysemous uses of single characters, TLS lists were reduced for the machine learning process. To say a list was "reduced" means that from it were eliminated rare, marginal or excessively polysemous characters, as well as repetitious compounds, any of which would bias association measures. No new terms were added to TLS lists adapted for this study. Examples of reduced-list TLS deity or low god terms include *xian* 仙 fairy, *shen* 神 spirit, and *yao* 妖 demon. Examples of reduced-list TLS punishment terms include *xing* 刑 punishment and *zhu* 誅 execute/punish. Examples of reduced-list reward terms include *shang* 賞 reward and *ci* 賜 give/bestow favors/appoint to office.

The high god category includes three and only three proper nouns: Di 帝, Shangdi 上帝 and Tian 天. The Chinese pantheon is populated with species of supernatural agent other than high gods and deities, among which two were coded for this study, ancestors and sage kings. The TLS's "ancestor" synonym group was also reduced for use here. This group includes terms such as *zulong* 祖宗 forefathers; *xianren* 先人 ancestors, literally "first people"; and *xianzi* 曾祖 "ancestor", literally "grandfathers of the past." These and other terms in the list often refer to physically-dead human beings in one's agnatic kinship line who can still receive prayers, sacrifices, and worship on special calendrical days. Anthropologist Maurice Freedman describes Chinese ancestor worship in social functional terms pertinent to formation of in-group solidarity. "In worshipping their ancestors the Chinese are stressing harmony and unity instead of competition and individualism [because] patriliney linked the fortunes of agnates together" (1979, p. 351). Watson (1982) discusses a wealth of historical social functions of Chinese lineage systems, including the support of schools, protection, patronage, sponsorship of the clan's boys for the civil examination, and business activities.

Due to liberal coding practices in the formation of Thesaurus Linguae Sericae lists noted above, most of its synonym groups contained far too many polysemous characters that risked skewing and invalidating findings. For this reason, alternative methods for coding the remaining categories were used.

For the sage king category (there are a handful of widely recognized sage kings in Chinese history) Yao 堯 and Shun 舜 were selected because they were the two most prominent and commonly-mentioned. Yao and Shun were held in the highest regard as moral exemplars. *Analects* 6.30 reports Confucius's response to a student's question asking about the qualities of leaders who extend benevolence to the peasantry. "The Master said, 'Why stop at Good? Such a person should surely be called a sage! Even someone like Yao or Shun would find such a task daunting'" (Confucius, 2003, p. 63). Sage kings and dynastic founders were said to be sired by powerful spirits, or forces of nature that bred with human women, or dragons, and are referred to with a sense of the holy (Lewis, 1990, pp. 180–181, 307). Sage kings have temples, icons, and statues devoted to their honor even today. Although some scholars have suggested that these figures may represent euhemerized versions of previously non-human, possibly animal deities, it is nonetheless the case that, by the time of the composition of the texts in the test corpus, the sage-kings were conceived of as being extraordinary human beings with potential supernatural powers.

Recall above that some sinologists believe that China's successful growth and cultural evolution is dependent neither upon the social functional benefits of a commitment to *high* or to *low* gods but instead is based on non-divine sources. Since unification in 221 BCE, China has had a long history



of strong central governments, or to use Sarkissian's term, a "Big Gov" (Sarkissian, 2015). To represent the role of governments in this study, the set of names of all emperors in Imperial Chinese history were collected since emperors serve as central, unequivocal authorities in civil and military matters. Emperors' names are often appended with the term *Di* 帝, god, as an affix or honorific. Qin Shi Huang 秦始皇, unifier of the warring states whose famed mausoleum is guarded by 8000 terra cotta warriors at Xian 西安, began this trend calling himself Qin Shi Huang *Di*. With this affix, he represented himself as a god. It is important to understand that the computer code used was composed so as to separate uses of "Di" that follow emperor names from all other uses of "Di". Only members of the second group, where "Di" does not occur as an affix of an emperor's proper name or title, were included in the high god category.

Categories for *high gods*, *sage kings*, and *emperors* are small and composed exclusively of proper names. Terms across categories representing *morality*, *cognition*, and *religion* are frequent. In order to create these categories, experts blind to the hypotheses read through a list of the 5,000 most frequent terms in historical Chinese texts. When each individual believed that a term on the list merited inclusion in one or another of these categories, he placed it into the appropriate category. Using common standards for coding categorical data (Fleiss, 1971; Landis & Koch, 1977), rates of interrater reliability exceeded standards for each of the three lists: moral (94% agreement), cognition (97% agreement) and religion (98% agreement).

While reward and punishment are mentioned explicitly in the definition of a high god by leading proponents of MHGT, its advocates stress that high gods also care about cooperation and morality. Use of moral terms is widespread in studies of MHGs (Roes & Raymond, 2003; Stark, 2001). Examples of morality terms include *yi* 義 rightness/morality and *ren* 仁 benevolence/goodness.

MHG's have power to *monitor* humans and their behaviors (Shariff et al., 2010, p. 124). In Norenzayan et al., 2016, Big Gods are "all-knowing," possessing maximal powers of cognition (Norenzayan et al., 2016, p. 6). Studies show, in the eponymous title of Shariff and Norenzayan (2007), that "God is watching you". Possession and use of this ability imply high gods are more likely than lesser supernatural creatures to engage in cognition, including watching, thinking, and judging. Here "cognition" refers to mentation, sensation, and perception. Examples of cognition terms include *zhi* 知 to know or to be aware of, and *si* 思 to think, reflect, long for.

Above this paper distinguished between MHGT's two major explanatory mechanisms. The first is an internal cognitive mechanism that takes in semantic content about supernatural punishment, reward, and monitoring, and yields higher marginal rates of prosocial behavior. The second is a multi-level selection mechanism in which groups composed of people committed to the same high god will be likely to increase in-group solidarity, decrease in-group cheating, and prepare their members for competition with other groups. Consideration of this latter mechanism leads to a final coding category for religion. Even if they punish, monitor, and moralize, merely theoretical gods, or old, powerful but forgotten gods, or gods introduced by minority members of the in-group, will be far less important to the group's cultural evolutionary success than will those gods who are subjects of religious devotion by the majority. Gods in the latter group are far more cognitively, emotionally, and behaviorally salient to group members. Supporters of MHGT imply that high gods are more likely to be the subjects of religious devotion (Henrich & Gervais 2010; see also Barrett, 2008). The social-functional role of religious devotion of people to their high gods is described by advocates of MHGT who say that MHGs, rather than low gods, "elicit deep devotions and extravagant rituals" and that these "ritual and devotional practices ... effectively elevate prosocial sentiments, galvanize solidarity, and transmit and signal deep faith" (Norenzayan et al., 2016, pp. 3, 6). Surrounding MHGs with traditional religious practice is said to better "exploit human psychology in a host of different ways" to "build in-group solidarity" (Norenzayan et al., 2016). Religious belief in high gods is associated with sacrifices, initiations, celibacy, fasting, and more. Culturally transmitted religious teachings and customs reinforce a group's belief that high gods are responsible for causing illnesses and harms (Barrett, 2004; Bering et al., 2005). This reasoning motivated inclusion of a *religion* category in this exploratory study. Examples of religion terms include *ming* 命 fate and *li* 禮/禮 ritual.

**Table 3.** Total character counts by category.

Agent	Types	Tokens	Content	Types	Tokens
High Gods	2	21,350	Punishment	18	31,050
Deities	5	7,770	Reward	6	8,713
Sage Kings	2	3,367	Morality	109	146,234
Ancestors	16	2,175	Cognition	70	67,073
Emperors	23	4,277	Religion	57	64,928

See Table 3 for distribution of types and tokens of terms found in the categories just discussed. See Appendix 2 for complete lists of terms in these categories.

### 3.3. Collocational descriptive data and analysis

In this exploratory study, two models of semantic associations between sets of agent categories (*high gods, deities, sage kings, ancestors, emperors*) and sets of social-functional categories (*punishment, reward, morality, cognition, religion*) are presented: a corpus-level collocation model, and a text-level regression model. This section explains the collocational data used to represent semantic association between categories of terms.

In corpus linguistics and text analytics, study of collocations is a widespread technique for modeling lexical association. Collocation analysis compares conditional probabilities between paired target terms in order to answer the question, posed of some corpus: Do paired terms *A* and *B* have greater semantic association than paired terms *A* and *C*? In the context of the social sciences and humanities the demonstrated ability to identify semantic associations between sets of terms stands out as the most promising feature of collocation analysis (Gries, 2013). Statistically tested collocations warrant inferences about patterns of thought and emotion, and, as a result, collocation analysis is practiced where literature is used as data (Manning & Schütze, 1999; Sampson & McCarthy, 2005; Teubert & Čermáková, 2007). Use of collocational analysis in religious studies is frequent and has revealed hidden features of semantic content in discursive practices related to religious identity (Baker et al., 2013).

Collocation analysis typically uses a fixed-term context window. A context window is a set of terms surrounding target terms within which collocations are assessed. While the sentence level might seem to be a natural window for lexical meaning, sentences vary considerably in length. This variance hinders statistical assessment of resulting semantic associations. Instead, most studies normalize length by opting for a fixed window size (Church & Hanks, 1990). For completeness, collocational data were calculated using two fixed window sizes, 5LR and 10LR, that is, 5 or 10 words to the left or right (LR) of a target term. For example, the 10LR window size includes the ten Chinese characters before and the ten Chinese characters after (every occurrence of) the target word. Pilot testing confirmed that the 10LR window size captures common collocations, including when target and collocate are separated by predicates, which is helpful since Chinese grammar has a subject-predicate-object structure familiar from English.

To statistically validate collocational findings the normal approximation to the binomial distribution was used to calculate observed proportions of co-occurrence along with a 95% confidence interval regarding the true parameter itself (see Formula 1). This refers to the proportion of co-occurrences in all ancient Chinese literature, of which only a portion is available to study.

Formula 1. Binomial distribution

$$\hat{p} \pm 1.96\sqrt{\frac{1}{n}\hat{p}(1 - \hat{p})}$$

where  $\hat{p}$  is the observed proportion, the constant 1.96 corresponds to the 95% level of confidence, and  $n$  is the number of observations (unordered pairs of words that could be the pair of terms of interest). The number of observations is calculated as the total number of terms in the corpus

multiplied by the word window of interest. Where the window size is 10,  $n$  represents a value that is ten times the number of terms in the texts.

Since the units are proportions, it is appropriate to compare these proportions across categories of terms. Since the goal is to explore, for example, whether it is members of the agent category of *High gods* or members of the agent category of *Deities* that have a stronger semantic relationship to members of the social-functional category of *Punishment* terms, an explicit statistical test is constructed. Use of this proportionality measure permits significance estimation by testing for the difference in proportions using the binomial approximation to the normal distribution. Given the large sample sizes involved, this is a z-test of the difference of two proportions as follows:

Formula 2. Z-test with binomial distribution

$$z = \frac{\hat{p}_1 - \hat{p}_2}{\sqrt{\hat{p}(1 - \hat{p})\left(\frac{1}{n_1} + \frac{1}{n_2}\right)}}, \text{ where } \hat{p} = \frac{\hat{p}_1 n_1 + \hat{p}_2 n_2}{n_1 + n_2}$$

where  $\alpha = .05$ . The null hypothesis is that proportions of content words within the assigned word window (10LR or 5LR) around the focal word are equal.

### 3.4. Mixed linear model for text-level associations

A z-test using descriptive collocational data may not appear to model lexical associations at the level of individual texts or control for contextual factors. So a linear mixed model was trained with proportion count as the target variable and genre with 13 levels, era of writing with 4 levels, and the interaction between agent categories (five levels) and social functional categories (five levels) as fixed factors. Because every association pair interaction (e.g., interactions between agent and social-functional categories such as *Deities* ~ *Punishment*) is extracted from each text, such that each text has 25 association pairs in the model, text was included as a random factor. The proportion counts (see previous section) were computed as the relation between two entities (an association pair), the agent category (e.g., *Deities*) and social-functional category (e.g., *Punishment*) respectively. Due to results obtained using collocational data that showed that use of the 10LR and the 5LR windows did not result in significantly different rates of semantic associations between agent and social-functional categories, to simplify the multilevel model a 10LR window was adopted. In order to test for reliable statistical effects and due to the number of observations, the  $\alpha$ -level was set to 0.005 (Benjamin et al., 2018), and the Akaike Information Criteria (AIC) were used for model selection. Association scores were computed at the 10LR level and ordered quantile (ORQ) normalization was applied because the target variable deviated from a normal assumption. ORQ was selected because it showed superior performance on the particular target variable when compared to the logarithmic, exponential, Yeo-Johnson and Lambert W x F transformations (Peterson & Cavanaugh, 2019).

## 4. Results

### 4.1. Study 1: collocation results

Table 4 reports results from Study 1, which used z-tests of collocational data across agent and social-functional categories. Column 1 reports the social-functional category and column 2 delimits the social-functional category by agent category. Columns 3 and 4, Proportion, represent the raw descriptive statistic referred to above as a “proportionality” measure for the 10LR and 5LR context windows. Suppose one examines the row of data associated with the *Punishment* social-functional category and the *Deity* agent category, row 2. (Capitalized italic terms refer to the set of terms falling in the eponymous category.) Column 2 of this row reads “0.69%”. This means that 0.69% of all token terms found between 10 terms to the left and 10 terms to the right of every token represented in the

**Table 4.** Collocation proportions across category.

Content	Agent	Proportion		z-score		p-value	
		10LR	5LR	10LR	5LR	10LR	5LR
Punishment	High gods	0.89%	0.88%				
	Deities	0.69%	0.63%	5.131	4.927	<.001***	<.001***
	Sage Kings	1.28%	1.23%	-6.875	-5.155	<.999^^^	>.999^^^
	Ancestors	0.75%	0.63%	2.073	2.666	.0191*	.0038**
	Emperors	1.07%	0.95%	-3.583	-0.652	>.999^^^	.7430
Reward	High gods	0.23%	0.22%				
	Deities	0.17%	0.16%	2.859	1.875	.0021**	.0304*
	Sage Kings	0.14%	0.10%	3.08	3.582	.001**	<.001***
	Ancestors	0.22%	0.18%	0.164	0.408	.4349	.3416
	Emperors	0.31%	0.24%	-3.352	0.062	>.999^^^	.4754
Morality	High gods	4.81%	4.78%				
	Deities	4.02%	3.84%	8.837	8.152	<.001***	<.001***
	Sage Kings	4.83%	4.50%	-0.213	2.266	.5843	.017*
	Ancestors	5.54%	5.31%	-4.707	-2.882	>.999^^^	.998^^
	Emperors	6.43%	6.72%	-13.806	-14.605	>.999^^^	>.999^^^
Cognition	High gods	2.00%	1.96%				
	Deities	2.59%	2.50%	-9.646	-6.398	>.999^^^	>.999^^^
	Sage Kings	1.76%	1.57%	2.942	3.747	.0016**	<.001***
	Ancestors	2.44%	2.45%	-4.409	-5.292	>.999^^^	>.999^^^
	Emperors	1.47%	1.34%	7.243	6.332	<.001***	<.001***
Religion	High gods	3.32%	3.58%				
	Deities	4.45%	4.51%	-14.36	-6.819	>.999^^^	>.999^^^
	Sage Kings	1.80%	1.73%	14.768	13.391	<.001***	<.001***
	Ancestors	3.85%	3.91%	-4.126	-1.896	>.999^^^	.0971^
	Emperors	1.66%	1.53%	18.088	17.097	<.001***	<.001***

set of *Deity* terms in the corpus are members of the set of *Punishment* terms. Columns 5 and 6 report z-scores calculated using the raw proportionality data for the 10LR and 5LR windows, respectively. Z-scores describe the hypothesis that the category of *High gods* has a *stronger* association with the social-functional category (punishment, reward, morality, cognition, religion) than other agent categories do (*Deities*, *Sage kings*, *Ancestors*, and *Emperors*). Therefore no z-scores are reported for high gods. Row 2 represents the semantic associations found between *Punishment* ~ *High god*. There “0.89%” refers to the fact that 0.89% of all token terms found between 10 terms to the left and 10 terms to the right of every token member of the set of *High god* terms in the corpus are *Punishment* terms. The association score assigned to each social functional ~ agency pairing is calculated *relative to High god terms*. The two implications are as follows.

First, the descriptive data contained in the Proportion columns of rows 1 and 2 indicate that for both the 10LR and 5LR windows, *High gods* has a greater association with *Punishment* than does *Deities*. The z-score for the *High gods* ~ *Punishment* and *Deities* ~ *Punishment* association at 10LR window is 5.13 (row 2, column 5). This indicates that *High gods* are overrepresented by a factor of over five standard deviations than the rate at which *Deities* are represented. Columns 7 and 8 represent significance tests performed on the z-scores, where  $\alpha = .05$ . The p-values for the 10LR and 5LR z-scores confirm that the stronger semantic association between *High gods* ~ *Punishment* than between *Deities* ~ *Punishment* is statistically significant. Negative z-scores represent cases in which the collocation rate of high gods in a given social-functional category is *under* the collocation rate of the comparison agent category. For example, in row 3, column 5,  $z = -6.875$ . This indicates that *High gods* are significantly *underrepresented* in 10LR windows where punishment terms are focal by a factor of over -6 standard deviations the rate of *Sage kings*. That is, *Sage kings* have a much stronger semantic association with punishment than do *High gods*.

In this context, in order to “test” MHGT, logically derived hypotheses about the differentiation between the class of supernatural agents that qualify as high gods and the class of all other supernatural agents are needed. This in turn would require decisions about which definitions of “high god” proffered by MHGT to use to derive these hypotheses. For various reasons observed above

this is a fool's errand. So this paper finds an exploratory way forward. In [Table 4](#), asterisks indicate significant results as expected by MHGT and carats indicate significant results that were not expected by MHGT. See [Table 4](#) for additional results from Study 1.

Findings from Study 1 are now summarized. When comparing semantic associations between *High gods* terms and terms from other agent categories, a number of significant results were obtained across social-functional categories of *Punishment*, *Reward*, *Morality*, *Cognition* and *Religion*. *High gods* have significantly stronger semantic association with *Punishment* than does *Deities* and *Ancestors*, though *Sage kings* and *Emperors* have significantly stronger association with *Punishment* than does *High gods*. *High gods* have significantly stronger semantic association with *Reward* than do *Deities* and *Sage kings*, though *Emperors* have significantly stronger association with *Reward* than does *High gods*. *High gods* have significantly stronger semantic association with *Morality* than do *Deities*, though *Ancestors* and *Sage kings* have significantly stronger association with *Morality* than does *High gods*. *High gods* terms have significantly stronger semantic association with *Cognition* than do *Sage kings* and *Emperors*, though *Deities* and *Ancestors* terms have significantly stronger association with *Cognition* than does *High gods*. Finally, *High gods* have significantly stronger semantic association with *Religion* than do *Sage kings* and *Emperors*, though *Deities* and *Ancestors* terms have significantly stronger association with *Religion* than does *High gods*. Altogether, the varied pattern of collocational results indicates that at least one other agent category besides *High gods*, and usually at least two agent categories besides *High gods*, possess stronger semantic associations across every one of the five social-functional categories.

#### 4.2. Study 2: mixed linear model results

To explore collocation results at the level of individual texts, the performance of four mixed linear models with random intercepts were compared. These models were generated using data from the 10LR window. The goal of these models is to predict variation in the proportion counts on *genre*, *era*, and *association* pairs. Model 1 is an intercept-only baseline model with text as a random factor; model 2 included genre as a fixed factor; model 3 included both genre and era as fixed factors; and model 4 added an interaction predictor between agent and social-functional categories to model 3. Model 2 performed significantly better than the baseline,  $p(15) < .0001$ , but only the Confucian genre showed a reliable effect:  $\beta = 0.4$ , 95% CI [0.19, 0.62],  $t(83) = 3.42$ ,  $p < .005$ . Model 3 marginally improved the AIC,  $p(18) < .005$ , adding a significant coefficient for the Post Han era:  $\beta = -0.34$ , 95% CI [-0.53, -0.15],  $t(80) = -3.18$ ,  $p < .005$ . Finally, model 4 with association pairs added as a predictor showed a considerable improvement in predicting proportion counts  $p(42) < .0001$ . See [Table 5](#) for genres. See [Table 6](#) for the individual association pair coefficients. See [Appendix 4](#) for intercept values.

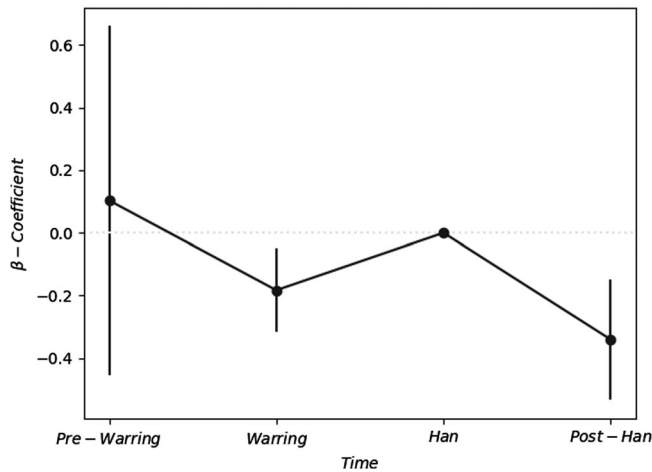
**Table 5.** Genres.

Genre or School	Character Count	Percent of total
Confucianism (儒家)	1,072,675	0.1868
Daoism (道家)	198,849	0.0346
Legalism (法家)	259,225	0.0451
Mohism (墨家)	80,851	0.0141
School of Names (名家)	3,181	0.0006
School of the Military (兵家)	44,072	0.0077
Mathematics (算書)	37,994	0.0066
Miscellaneous Schools (雜家)	248,725	0.0433
Histories (史書)	3,054,843	0.5320
Ancient Classics (經典文獻)	341,970	0.0596
Etymology (字書)	163,401	0.0285
Chinese Medicine (醫學)	226,071	0.0394
Excavated texts (出土文獻)	10,682	0.0019
Totals	5,742,539	1

**Table 6.** Association pair coefficients for linear mixed model (square brackets are 95% confidence intervals), high gods and reward are reference levels.

Content	Agent	Coefficient	t-statistic
Punishment	Deities	-0.11 [-0.41, 0.2]	-0.67
	Sage Kings	-0.003 [-0.31, 0.3]	-0.02
	Ancestors	-0.44 [-0.75, -0.13]	-2.8
	Emperors	-0.39 [-0.7, -0.08]	-2.49
Morality	Deities	-0.04 [-0.34, 0.27]	-0.24
	Sage Kings	-0.41 [-0.72, -0.1]	-2.62
	Ancestors	-0.7 [-1, -0.4]	-4.42**
	Emperors	-0.82 [-1.13, -0.51]	-5.2**
Cognition	Deities	0.04 [-0.51, 0.11]	0.26
	Sage Kings	-0.37 [-0.68, -0.07]	-2.38
	Ancestors	-0.2 [-0.51, 0.11]	-1.28
	Emperors	-0.3 [-0.61, 0.01]	-1.91
Religion	Deities	-0.52 [-0.84, -0.23]	-3.32**
	Sage Kings	-1.04 [-1.35, -0.74]	-6.65**
	Ancestors	-0.54 [-0.85, -0.23]	-3.43**
	Emperors	-1.39 [-1.7, -1.08]	-8.86**

\* $p < 0.005$ ; \*\* $p < 0.001$ .

**Figure 1.** Era coefficients sorted in chronological order with 95% confidence intervals, where the Han era is the reference level.

At the level of individual texts, associations between social-functional categories and *Deities* and *High gods* respectively, behave similarly with the exception of *Religion*. The association between *High gods* and *Religion* is stronger than any other agent category. The association between *Morality* and *High gods* is also stronger than *Ancestors* and *Emperors*, but not *Deities*. The models only show very limited effects of genre and era: the Confucian genre shows a positive association with proportion counts, and the Post-Han era is negatively associated with proportion counts. The temporal development of these relationships can be explored by plotting the era coefficients in chronological order (see Figure 1). This shows that era is not a good predictor of the proportionality scores but for the fact that in the Post-Han era, a decrease in association strength between all association pairs in a 10LR window was observed.

## 5. Discussion

Here the evidential relationships between results, moralizing high god theory and broad supernatural punishment theory are tentatively assessed. Possible sources of in-group cooperation in historical

China, considering both supernatural sources and semi-secular, governmental sources, are discussed. The paper concludes by describing some of the several methodological limitations to this study and noting pathways for future research.

The testing of MHGT has included laboratory studies, experimental studies, and use of anthropological and historical databases. Landmark laboratory priming studies considered as significant evidential support for MHGT, such as Shariff & Norenzayan, 2007, were conducted primarily with WEIRD participants. Multiple recent pre-registered attempts to replicate this study have failed, including Miyatake and Higuchi (2017) and Gomes and McCullough (2015). Database work with HRAF and SCCS (Roes & Raymond, 2003) is capable of offering partial support to MHGT, but its cultures lack high social complexity. Critics show that these studies have not adequately accounted for common cultural ancestry (Atkinson et al., 2015). MHGT is said to explain religions of Mesoamerica and India's Karmic religions but, according to White, Sousa, and Prochownik (2016), it is doubtful that these religions possess high gods. The central thesis of MHGT is historical in nature and is an attempt to explain the rise of prosociality in certain cultures over time, yet laboratory studies, experimental studies, and anthropological database work are not positioned to test historical data. A recent, high-profile criticism of MHGT (Whitehouse et al., 2019) represents one of the few criticisms to address the fact that MHGT is a historical thesis about differential cultural evolution between groups, but this study has copious methodological failings (Beheim, 2019). Altogether, the current context of debate about MHGT raises the stakes of corpus analytics, especially of non-Western corpora. These considerations also raise the potential value of attempts to use historical data, specifically historical data about non-Western big societies (rather than small scale societies in SCCS or EA). This state of affairs led one researcher to comment that China is the "elephant in the [CSR] room" (Sarkissian, 2015, p. 324).

What are the roles of gods in early China with regard to punishment, reward, monitoring, and morality? This exploratory micro-modeling study finds that historical Chinese culture represented many supernatural and social sources of in-group cooperation. Consider that, in the corpus-level model from Study 1, with one exception *each one* of the supernatural agent categories bears the strongest relation to at least one of the social-functional categories. At the 10LR context window, *Sage kings* had the strongest relation to *Punishment*, *Emperors* to *Reward*, *Deities* to *Cognition*, and *Ancestors* to *Religion*. Remarkably, the exception is the *High gods* category, which never had the strongest semantic association with any social-functional category, at either 5LR or 10LR windows. This suggests both that distinct species of supernatural agent within the Chinese pantheon may have subserved distinct social functional roles, and that the social functional status of alleged high gods Shangdi, Di, and Tian is opaque.

Most if not all versions of MHGT expect that high gods are set apart from other supernatural species by their capacities to promote cooperation and prosociality in the in-group. Even the most recent, most ecumenical version of MHGT, presented in Norenzayan et al. (2016), argues that the traits of key social-functional interest to the theory are possessed to greater degrees by high gods. By contrast, broad supernatural punishment theory expects that social-functional traits useful for promoting in-groupism will be diffused across a variety of supernatural agencies, including spirits, ghosts, local deities, etc. Exploratory findings appear more in accord with expectations of broad supernatural punishment theory.

Having said this, a weighty cautionary point is in order. Though the High gods category did not have the *strongest* association with any of the social-functional categories, it did not have the *weakest* association with any of these categories. In addition, as Table 3 shows, the corpus contains far more tokens of high god terms than of deity terms. It appears Shangdi, Di, and Tian play unique and ubiquitous roles in Chinese history, though their role differentiation in promotion of in-group cooperation remains unclear, if not also dubious.

Turning to how the results address extant questions raised about the history of Chinese cooperation, a few points emerge. First, Di, Shangdi, and Tian together appear to have weaker semantic association with terms in the *Religion* category than do *Deities* and *Ancestors*. Recall Dutton

and Madison (2016) argued that Shangdi does not qualify as a high god on the grounds that Shangdi was not worshiped directly. This exploratory approach does not purport to determine the extent to which any specific supernatural agent is represented as an object of worship. In Study 1, semantic associations between species of supernatural agent and the *Religion* category may be taken as indirect support of this argument. However, in Study 2, the multilinear model suggests that under its assumptions Chinese high gods have a special relationship to religion.

Second, recall Sarkissian argued that China functions as an “elephant in the room” because China is “one of the largest and most enduring civilizations on the planet, yet one also lacking a rich tradition of belief in Big Gods or supernatural monitoring” (2015, p. 324). He adds that Shangdi and Tian (both of which are found in the *High god* category) play “a very small role in explaining how it scales up” (p. 324). One way of putting this point resembles Dutton and Madison’s observation that Chinese high gods appear to be uninvolved in activities that unify people under spiritual authority. Sarkissian’s second point is that alternatives to the MHG hypothesis, such as a “Big Gobs” hypothesis, probably better explain the fact that China is a large-scale society with extremely high rates of cooperation and prosociality (at least within the Han ethnic majority). The Big Gobs hypothesis states that ancient and medieval China had resources to increase cooperation and prosociality via the social functions of its early and vast government. Using emperor names as a proxy for government, data above suggest that the *Emperors* category has special status among supernatural agent categories. *Emperors* had significantly greater semantic association with categories of *Punishment*, *Reward*, and *Morality* than did *High gods*. As opposed to all other agent categories, *Emperors* had the least association with *Religion* terms and *Cognition* terms. These preliminary findings recommend further research to operationalize and test a “Big Gobs” hypothesis about the origins and maintenance of Chinese cultural evolution and growth.

Methods employed in this study are novel and largely untested outside the area of corpus linguistics and inside cognitive science. Limitations are many. The Ctext corpus is limited, and focused on texts and documents during the Warring States period (480–221 BCE) and Han Dynasty (206 BCE–220 CE). Far more textual data is available if the timeframe is extended into the Ming (1368–1644 CE) or Qing Dynasties (1644–1911 CE). Statistical methods used in this paper are infrequently used to present conclusions about psychological states of readers and writers. As such, this paper makes a methodological assumption that robust semantic associations presented in classical texts of extreme historical authority can and regularly did influence psychological associations in the minds of their historical readers and pupils. This assumption merits continuing scrutiny since texts in this corpus are nearly the exclusive products of men. Since these texts were written by elites for elites, they should not be taken to accurately reflect the mindset of historical Chinese commoners. Further, local worship by clan members throughout historical China would have included, if not been dominated by, a different set of ancestral deities and spirits drawn from the lineage of each extended family in clan temple halls or in the presence of ancestor tablets. The present study was not positioned so as to investigate these supernatural agents or their social functions.

Statistical analyses and methods represented here for the study of corpora are uncommon and merit further scrutiny and replication. Dating of texts in this corpus was limited to eras, each of which spans hundreds of years. With more precise textual dating, future methods may support detailed conclusions about conceptual change over time. Further exploration of hypotheses from cognitive science of religion and evolutionary psychology with corpora using methods similar to or better than this one will come in due course. In addition to these strictly methodological concerns, theoretical uncertainties about the core content of moralizing high god theory means that results presented here did not test hypotheses logically derived from MHGT. Instead, this paper explored what were taken to be suggestions or expectations of this theory that contrast with suggestions or expectations from broad supernatural punishment theory.

Two types of future corpus-based machine-learning research could improve understanding of how belief in supernatural agents and their capacities influenced the individual-level and group-level mechanisms hypothesized by MHGT. First, the methods used here could be reapplied to a



much larger corpus encompassing many more texts in Chinese history. The Ctext corpus has special utility in addressing research questions raised here because it includes all of the most significant early texts written in the Chinese language, because it supported coding for era, and because it is not beyond a paywall, promoting openness and enabling easier replication. Nonetheless, in comparison with the voluminous set of writings produced in pre-Imperial and Imperial Chinese history, the Ctext corpus is undersized. Use of other corpora will provide more epistemic justification for inferences from semantic associations to social functions. The Chinese Ancient Texts Database (CHANT) housed at the D. C. Lau Research Centre for Chinese Ancient Texts, Chinese University of Hong Kong, is much larger and more authoritative than the Chinese Text Project. It is also proprietary and also behind a paywall.

The second type of corpus-based machine learning research would adopt new methods different than the micro-modeling approach used here in order to compare the social-functional roles of sets of supernatural agents as represented in vast corpora across many historical linguistic and cultural traditions. The vision of such a project would be first to understand the roles of supernatural agents *within* each tradition by exhaustive analysis of their semantic associations contained within each historical corpus (Chinese, Greek, Sanskrit, etc.). Then, second, continuing research would compare *between* social functions of supernatural agents across these culturally distinct corpora. Such a project would greatly benefit cognitive science of religion at large and, in particular, would improve understanding of the cultural evolution of high god religious groups through time. Increasing availability of large datasets and magnificent recent advances, e.g., in “word embedding,” which is a deep neural network architecture known as a multi-layer bidirectional transformer encoder (Devlin et al., 2018; Vaswani et al., 2017), put the promise of such a cross-cultural project within reach.

## Notes

1. The data repository is located here and includes the pre-processed corpus as well as associated appendices <<https://hecc.ubc.ca/quantitative-textual-analysis/data-repository/>>.
2. Downloadable at [https://github.com/petekirby/character\\_profiler](https://github.com/petekirby/character_profiler).
3. See <http://tls.uni-hd.de/>.

## Acknowledgements

We thank Donald Sturgeon for permission to use The Chinese Text Project ([www.ctext.org](http://www.ctext.org)) corpus in our analysis, and Christoph Harbsmeier for his assistance in employing the Thesaurus Linguae Sericae (<http://tls.uni-hd.de>). This work was funded by subgrants through (a) Templeton World Charities’ “The Chinese Challenge” (Barrett & Hornbeck, PIs), based at Fuller Theological Seminary, (b) the Canadian Social Sciences and Humanities Research Council (SSHRC) Partnership Grant on “The Evolution of Religion and Morality” (Edward Slingerland, PI), based at the University of British Columbia, (c) John Templeton Foundation’s “Academic Cross-Training” fellowship given to Nichols, and (d) the College of Humanities & Social Sciences at California State Fullerton University, Fullerton. Also valuable were comments from participants in sessions at the 2012 conference of the Society for the Scientific Study of Religion, the 2014 conference of the Association for Asian Studies, and the 2016 International Association for Cognitive Science of Religion conference. Special thanks to Justin Barrett, Ryan Hornbeck, Paul Bucci, Uffe Bergeton, Robert Callergård, and Scott Kleinman.

## Disclosure statement

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

## Funding

This work was supported by Academic Cross-Training Grant, John Templeton Foundation: [Grant Number 60704]; The Chinese Challenge, Templeton World Charities Foundation: [Grant Number 0020]; The Evolution of Religion and Morality, Canadian Social Sciences and Humanities Research Council: [Grant Number 895-2011-1009].

## References

- Atkinson, Q. D., Latham, A. J., & Watts, J. (2015). Are big gods a big deal in the emergence of big groups? *Religion, Brain & Behavior*, 5(4), 266–274. <https://doi.org/10.1080/2153599X.2014.928351>
- Atran, S., & Henrich, J. (2010). The evolution of religion: How cognitive by-products, adaptive learning heuristics, ritual displays, and group competition generate deep commitments to prosocial religions. *Biological Theory*, 5(1), 18–30. [https://doi.org/10.1162/BIOT\\_a\\_00018](https://doi.org/10.1162/BIOT_a_00018)
- Baker, P., Gabrielatos, C., & McEnery, T. (2013). Sketching Muslims: A corpus driven analysis of representations around the word “Muslim” in the British Press 1998–2009. *Applied Linguistics*, 34(3), 255–278. <https://doi.org/10.1093/applin/ams048>
- Barrett, J. (2008). Why Santa Claus is not a God. *Journal of Cognition and Culture*, 8(1), 149–161. <https://doi.org/10.1163/156770908X289251>
- Barrett, J. L. (2004). *Why would anyone believe in God?* AltaMira Press.
- Barrett, J. L., & Keil, F. C. (1996). Conceptualizing a nonnatural entity: Anthropomorphism in God concepts. *Cognitive Psychology*, 31(3), 219–247. <https://doi.org/10.1006/cogp.1996.0017>
- Baumard, N., Hyafil, A., Morris, I., & Boyer, P. (2015). Increased affluence explains the emergence of ascetic wisdoms and moralizing religions. *Current Biology*, 25(1), 10–15. <https://doi.org/10.1016/j.cub.2014.10.063>
- Beheim, B. A. (2019, September 19). *Missing data, moralizing gods, and the siren song of Aeiea*. <https://babeheim.com/2019/09/19/song-of-aeiea.html>
- Beheim, B., Atkinson, Q., Bulbulia, J., Gervais, W., Gray, R., Henrich, J., Lang, M., Willis Monroe, M., Muthukrishna, M., Norenzayan, A., Purzycki, B., Shariff, A., Slingerland, E., Spicer, R., & Willard, A. (2020). Corrected analyses show that moralizing gods precede complex societies but serious data concerns remain. *Nature* (under review).
- Benjamin, D. J., Berger, J. O., Johannesson, M., Nosek, B. A., Wagenmakers, E.-J., Berk, R., Bollen, K. A., Brembs, B., Brown, L., Camerer, C., Cesarini, D., Chambers, C. D., Clyde, M., Cook, T. D., De Boeck, P., Dienes, Z., Dreber, A., Easwaran, K., Efferson, C., ... Johnson, V. E. (2018). Redefine statistical significance. *Nature Human Behaviour*, 2(1), 6–10. <https://doi.org/10.1038/s41562-017-0189-z>
- Bering, J. M., McLeod, K., & Shackelford, T. K. (2005). Reasoning about dead agents reveals possible adaptive trends. *Human Nature*, 16(4), 360–381. <https://doi.org/10.1007/s12110-005-1015-2>
- Brazil, I. A., & Farias, M. (2016). Why would anyone want to believe in big Gods? *Behavioral and Brain Sciences*, 39, 24–25. <https://doi.org/10.1017/S0140525X15000370>
- Church, K. W., & Hanks, P. (1990). Word association norms, mutual information, and lexicography. *Computational Linguistics*, 16(1), 22–29. <https://doi.org/10.5555/89086.89095>
- Clark, K. J., & Winslett, J. T. (2011). The evolutionary psychology of Chinese religion: Pre-Qin high gods as punishers and rewarders. *Journal of the American Academy of Religion*, 79(4), 928–960. <https://doi.org/10.1093/jaarel/lfr018>
- Confucius. (2003). *Confucius analects: With selection from traditional commentaries*. (E. G. Slingerland, Trans.). Hackett Pub. Co.
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2018). BERT: Pre-training of deep bidirectional transformers for language understanding. *ArXiv:1810.04805 [Cs]*. <http://arxiv.org/abs/1810.04805>.
- Dutton, E., & Madison, G. (2016). Even “Bigger Gods” developed amongst the pastoralist followers of Moses and Mohammed: Consistent with uncertainty and disadvantage, but not prosociality. *Behavioral and Brain Sciences*, 39, 27–28. <https://doi.org/10.1017/S0140525X15000412>
- Fleiss, J. L. (1971). Measuring nominal scale agreement among many raters. *Psychological Bulletin*, 76(5), 378–382. <https://doi.org/10.1037/h0031619>
- Freedman, M. (1979). On the sociological study of Chinese religion. In *The study of Chinese Society* (pp. 351–369). Stanford University Press.
- Galen, L. W. (2016). Big gods: Extended prosociality or group binding? *Behavioral and Brain Sciences*, 39, e13. <https://doi.org/10.1017/S0140525X15000436>
- Gomes, C. M., & McCullough, M. E. (2015). The effects of implicit religious primes on dictator game allocations: A preregistered replication experiment. *Journal of Experimental Psychology: General*, 144(6), e94–e104. <https://doi.org/10.1037/xge0000027>
- Gries, S. (2013). 50-something years of work on collocations: What is or should be next ... *International Journal of Corpus Linguistics*, 18(1), 137–166. <https://doi.org/10.1075/ijcl.18.1.09gri>
- Hall, D. L., & Ames, R. T. (1995). *Anticipating China: Thinking through the narratives of Chinese and Western culture*. State University of New York Press.
- Henrich, J., & Gervais, W. M. (2010). The Zeus problem: Why representational content biases cannot explain faith in gods. *Journal of Cognition and Culture*, 10(3), 383–389. <https://doi.org/10.1163/156853710X531249>
- Huebner, B., & Sarkissian, H. (2016). Cultural evolution and prosociality: Widening the hypothesis space. *Behavioral and Brain Sciences*, 39, <https://doi.org/10.1017/S0140525X1500045X>
- Johnson, D. (2015). Big gods, small wonder: Supernatural punishment strikes back. *Religion, Brain & Behavior*, 5(4), 290–298. <https://doi.org/10.1080/2153599X.2014.928356>

- Landis, J. R., & Koch, G. G. (1977). The Measurement of observer agreement for categorical data. *Biometrics*, 33(1), 159. <https://doi.org/10.2307/2529310>
- Laurin, K., Shariff, A. F., Henrich, J., & Kay, A. C. (2012). Outsourcing punishment to God: Beliefs in divine control reduce earthly punishment. *Proceedings of the Royal Society B: Biological Sciences*, 279(1741), 3272–3281. <https://doi.org/10.1098/rspb.2012.0615>
- Lewis, M. E. (1990). *Sanctioned violence in early China*. State University of New York Press.
- Ma-Kellams, C., & Blascovich, J. (2013). Does “science” make you moral? The effects of priming science on moral judgments and behavior. *PLoS ONE*, 8(3), Article 57989. <https://doi.org/10.1371/journal.pone.0057989>
- Manning, C. D., & Schütze, H. (1999). *Foundations of statistical natural language processing*. MIT Press.
- McKay, R., & Whitehouse, H. (2016). Religion promotes a love for thy neighbour: But how big is the neighbourhood? *Behavioral and Brain Sciences*, 39, <https://doi.org/10.1017/S0140525X15000503>
- Miyatake, S., & Higuchi, M. (2017). Does religious priming increase the prosocial behaviour of a Japanese sample in an anonymous economic game? Religion and prosociality in Japan. *Asian Journal of Social Psychology*, 20(1), 54–59. <https://doi.org/10.1111/ajsp.12164>
- Nichols, R. (2013). The origins and effects of filial piety (xiao 孝): How culture solves an evolutionary problem for parents. *Journal of Cognition and Culture*, 13(3–4), 201–230. <https://doi.org/10.1163/15685373-12342092>
- Nichols, R., & Logan, C. (2017). High gods, low gods, and morality in ancient China: Developing new methods, answering old questions. In R. G. Hornbeck, & J. Barrett (Eds.), *Religious cognition in China: Homo religious and the dragon* (pp. 161–177). Springer.
- Norenzayan, A. (2013). *Big gods: How religion transformed cooperation and conflict*. Princeton University Press.
- Norenzayan, A., Shariff, A. F., Gervais, W. M., Willard, A. K., McNamara, R. A., Slingerland, E., & Henrich, J. (2016). The cultural evolution of prosocial religions. *Behavioral and Brain Sciences*, 1–86. <https://doi.org/10.1017/S0140525X14001356>
- Peterson, R. A., & Cavanaugh, J. E. (2019). Ordered quantile normalization: A semiparametric transformation built for the cross-validation era. *Journal of Applied Statistics*, 1–16. <https://doi.org/10.1080/02664763.2019.1630372>
- Raffield, B., Price, N., & Collard, M. (2019). Religious belief and cooperation: A view from Viking-Age Scandinavia. *Religion, Brain & Behavior*, 9(1), 2–22. <https://doi.org/10.1080/2153599X.2017.1395764>
- Roes, F., & Raymond, M. (2003). Belief in moralizing gods. *Evolution and Human Behavior*, 24(2), 126–135. [https://doi.org/10.1016/S1090-5138\(02\)00134-4](https://doi.org/10.1016/S1090-5138(02)00134-4)
- Sampson, G., & McCarthy, D. (Eds.). (2005). *Corpus linguistics: Readings in a widening discipline* (Pbk. ed). Continuum.
- Sarkissian, H. (2015). Supernatural, social, and self-monitoring in the scaling-up of Chinese civilization. *Religion, Brain & Behavior*, 5(4), 323–327. <https://doi.org/10.1080/2153599X.2014.928358>
- Shariff, A. F., & Norenzayan, A. (2007). God is watching you: Priming God concepts increases prosocial behavior in an anonymous economic game. *Psychological Science*, 18(9), 803–809. <https://doi.org/10.1111/j.1467-9280.2007.01983.x>
- Shariff, A., Norenzayan, A., & Henrich, J. (2010). The Birth of high gods: How the cultural evolution of supernatural policing agents influenced the emergence of complex, cooperative human societies, paving the way for civilization. In M. Schaller, A. Norenzayan, S. J. Heine, T. Yamagishi, & T. Kameda (Eds.), *Evolution, culture, and the human mind* (pp. 119–136). Psychology Press.
- Slingerland, E., Willis Monroe, M., Sullivan, B., Walsh, R. F., Veidlinger, D., Noseworthy, W., Herriott, C., Raffield, B., Peterson, J. L., Rodriguez, G., Sonik, K., Green, W., Tappenden, F. S., Ashtari, A., Spicer, R., & Muthukrishna, M. (in press). Historians Respond to Whitehouse et al. (2019). Complex societies precede moralizing gods throughout world history. *Journal of Cognitive Historiography*.
- Stark, R. (2001). Gods, rituals, and the moral order. *Journal for the Scientific Study of Religion*, 40(4), 619–636. <https://doi.org/10.1111/0021-8294.00081>
- Teubert, W., & Čermáková, A. (2007). *Corpus linguistics: A short introduction*. Continuum.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł, & Polosukhin, I. (2017). Attention is all you need. In *Advances in neural information processing systems Vol. 30* (pp. 5998–6008). Curran Associates, Inc. <https://papers.nips.cc/paper/7181-attention-is-all-you-need.pdf>.
- Watson, J. L. (1982). Chinese kinship reconsidered: Anthropological perspectives on historical research. *The China Quarterly*, 92, 589. <https://doi.org/10.1017/S0305741000000965>
- Watts, J., Bulbulia, J., Gray, R. D., & Atkinson, Q. D. (2016). Clarity and causality needed in claims about Big Gods. *Behavioral and Brain Sciences*, 39. <https://doi.org/10.1017/S0140525X15000576>
- Watts, J., Greenhill, S. J., Atkinson, Q. D., Currie, T. E., Bulbulia, J., & Gray, R. D. (2015). Broad supernatural punishment but not moralizing high gods precede the evolution of political complexity in Austronesia. *Proceedings of the Royal Society B: Biological Sciences*, 282(1804), Article 20142556. <https://doi.org/10.1098/rspb.2014.2556>
- Welch, M. R., Sikkink, D., & Loveland, M. T. (2007). The radius of trust: Religion, social embeddedness and trust in strangers. *Social Forces*, 86(1), 23–46. <https://doi.org/10.1353/sof.2007.0116>
- White, C. J. M., Norenzayan, A., & Schaller, M. (2019). The content and correlates of belief in karma across cultures. *Personality and Social Psychology Bulletin*, 45(8), 1184–1201. <https://doi.org/10.1177/0146167218808502>

- White, C., Sousa, P., & Prochownik, K. (2016). Explaining the success of karmic religions. *Behavioral and Brain Sciences*, 39, 42–43. <https://doi.org/10.1017/S0140525X15000588>
- Whitehouse, H., François, P., Savage, P. E., Currie, T. E., Feeney, K. C., Cioni, E., Purcell, R., Ross, R. M., Larson, J., Baines, J., Ter Haar, B., Covey, A., & Turchin, P. (2019). Complex societies precede moralizing gods throughout world history. *Nature*, 568(7751), 226–229. <https://doi.org/10.1038/s41586-019-1043-4>