A New Era in the Study of Global History is Born but it Needs to Be Nurtured

Please note: This is a preprint of a manuscript that has been accepted for publication in the *Journal of Cognitive Historiography*. The final published version may contain small differences such as formatting and copy editing. Please direct correspondence to harvey.whitehouse@anthro.ox.ac.uk.

Recommended citation: Whitehouse, H., François, P., Savage, P. E., Currie, T. E., Feeney, K. C., Cioni, E., ... Turchin, P. (In press). A new era in the study of global history is born but it needs to be nurtured. *Journal of Cognitive Historiography*. doi: 10.1558/jch.39422

Harvey Whitehouse
University of Oxford
Email: harvey.whitehouse@anthro.ox.ac.uk

Pieter François
University of Oxford
Email: pieter.francois@stb.ox.ac.uk

Patrick Savage
Keio University
Email: psavage@sfc.keio.ac.jp

Thomas E. Currie
University of Exeter
Email: T.Currie@exeter.ac.uk

Kevin C. Feeney
Data Chemist
Email: kevin@datachemist.com

Enrico Cioni
Seshat project
Email: enrico.a.cioni@gmail.com

Rosalind Purcell
Seshat project
Email: rosalindpurcell@gmail.com

Robert M. Ross

Macquarie University
Email: robross46@gmail.com

Jennifer Larson
Kent State University
Email: jl Larson@kent.edu

John Baines
University of Oxford
Email: john.baines@orinst.ox.ac.uk

Barend ter Haar
University of Hamburg
Email: barend.ter.haar@uni-hamburg.de

Alan Covey
University of Texas
Email: r.alan.covey@austin.utexas.edu

Peter Turchin
University of Connecticut
Email: peter.turchin@uconn.edu

Abstract

In this issue, Slingerland et al. criticize the quality of the data from *Seshat*: Global History Databank utilized in our *Nature* paper entitled “Complex Societies Precede Moralizing Gods throughout World History”. Their critique centres around the roles played by research assistants and experts in procuring and curating data, periodization structure, and so-called ‘data pasting’ and ‘data filling’. We show that these criticisms are based on misunderstandings or misrepresentations of the methods used by *Seshat* researchers. Overall, their critique does not call into question any of our main findings, but it does highlight various shortcomings of Slingerland et al.’s rival database project. Our collective efforts to code and quantify features of global history hold out the promise of a new era in the study of global history but only if critique can be conducted in good faith, rivalries kept in check, and both promises and pitfalls of open science fully recognized.

Keywords

Religion, evolution, social complexity, global history databank, open science.

Introduction

In this issue of the *Journal of Cognitive Historiography*, Slingerland et al. (2019) criticize the methodology in our recent paper in *Nature* (Whitehouse et al. 2019), hereafter the ‘moralizing gods paper’, based on our analysis of *Seshat*: Global History Databank (for an introduction, see Turchin et al. 2019, in this issue). Although framed as a critique of the moralizing gods paper, none of the

arguments presented by Slingerland *et al.* would significantly alter the main findings of that paper (see also Savage *et al.*, Under review). Anyone coming to their critique hoping to learn more about whether moralizing gods appeared before or after the sharpest rises in social complexity, might be disappointed.

Slingerland *et al.*'s criticisms focus on issues of data quality, arguing that almost 70% of Seshat data represent 'coding errors'. They attempt to detail issues that they claim underlie such coding errors including the role of research assistants, periodization structure, and so-called "data pasting", "data filling", and alleged failure to freeze the datasets used in publications. We have shown elsewhere (Savage *et al.* Under review) that the statistical analysis leading to their claim of 70% coding errors is invalid and instead measures limitations of their own rival project, the Database of Religious History (DRH; Slingerland and Sullivan 2017), more than it measures Seshat coding errors. Here we will show that their qualitative criticisms also reflect misunderstandings by Slingerland *et al.* and/or limitations of the DRH. More worryingly still, Slingerland *et al.* make serious allegations of research malpractice that are demonstrably false. Although we fully support open science and are pleased that the debate has been carried out using some of its principles, we argue that abuse of the affordances of open science could endanger the future of newly emerging and fragile fields, to the detriment of all.

Data that Disagree

Slingerland *et al.*'s central argument is that the quality of Seshat's data is poor. The main evidence they cite for this is in their Table 1, where they claim to find almost 70% disagreement between codings of key ritual and religion variables in Seshat and the DRH. As we argue elsewhere, this analysis is invalid because it explicitly assumes that where data coded in Seshat fails to agree with data coded in the DRH, the data in Seshat constitute 'coding errors', but in fact there is so much disagreement and missing data within the DRH itself that this method guarantees high levels of Seshat "coding errors" regardless of whether or not Seshat data is correct. For example, Slingerland *et al.*'s method compares Seshat's coding of supernatural enforcement of fairness in Wei Dynasty China against five different DRH entries, but these five different entries give three different answers (two present, two absent, and one unknown). Thus, whatever coding Seshat gives is guaranteed to be measured as at least 60% "Seshat coding errors", simply because that is the level of *internal* disagreement within the five different DRH entries for this variable. Slingerland *et al.*'s analysis is further invalidated by inappropriate comparison of different time periods, inappropriate treatment of missing DRH data as Seshat coding errors, and basic arithmetic errors (see Savage *et al.*, Under review for details).

The qualitative criticisms presented by Slingerland *et al.* are similarly invalid. For example, their claim that "the first instance of a "moralizing" god in the Seshat dataset for Susiana does not occur until the Achaemenid period, almost two thousand years later" is wrong: Figure 1, Extended Data Table 1, and Table S2 of our original paper show that the first attestation of moralizing gods in Susiana as

being in the Akkadian Empire ca. 2200BCE, where Slingerland *et al.* agree it should be.

Another misrepresentation is found in Slingerland *et al.*'s criticism of "such absurdities as using an observation about Late Shang (1800 BCE) [*sic*] ritual practice to code the 'ritual frequency' variable for every polity from the Western Zhou (1122 BCE) to the Ming (1643 CE)". We do nothing of the sort. Ritual frequency is coded in *Seshat* as daily in the Late Shang (1250-1045BCE [not 1800BCE]), seasonally in Western Zhou, and monthly in the Ming, each supported by one or more different references (Keay 2009 and Thorp 2006 for Late Shang, Hsu and Linduff 1988 for W. Zhou, and Zhu 1991 for Ming). Even had this criticism been legitimate, it would have been entirely irrelevant to our main finding that complex societies precede moralizing gods.

The goals of *Seshat* and the DRH are not identical and failure to realize this might account for some of Slingerland *et al.*'s misunderstandings. *Seshat* seeks to gather and store vast amounts of data on world history in a way that allows people to access, review, and comment on coding decisions. This allows data to be discussed and improved prior to scientific analysis and publication, a process requiring rigorous quality checks with expert input on particular variables (such as those in Table S2 of the moralizing gods paper). By contrast, the DRH seeks to ensure that all the data it contains, whether being used for statistical analysis or not, have been subjected to exactly the same quality control checks. One possible reason for this is that the DRH is intended to serve as a resource for undergraduate teaching (Slingerland and Sullivan, 2017; Slingerland *et al.*, 2019). Much of the Slingerland *et al.* critique stems from the assumption that the DRH approach is superior when it is simply a product of different goals. The key point to emphasise is that key data relevant to findings reported in *Seshat* publications are subjected to stringent quality checks as described in our methods sections.

Seshat vs. DRH: Appearance of moralizing gods after China's Shang-Western Zhou transition

In an effort to understand how Slingerland *et al.* arrived at putative 'coding errors', it is illuminating to consider a more detailed example: the first evidence of moralizing gods in Middle Yellow River Valley China. Slingerland *et al.* claim that our coding of this in Western Zhou (ca. 1000BCE) rather than a century earlier in Late Shang China is a crucial error, and that correcting it "would seriously undermine Whitehouse *et al.*'s conclusion". This is not true: Savage *et al.* (Under review) show that reanalysis using Slingerland *et al.*'s preferred date has no qualitative effect on our conclusions. Nevertheless, this is a good example of a putative "coding error" that reflects limitations of the DRH, not of *Seshat*. The data coded in *Seshat* and analyzed in our *Nature* paper suggest that moralizing gods appear in North China during the Western Zhou period (ca.1040–771 BCE). The first truly large-scale societies in North China appeared at least a millennium earlier (Liu and Chen 2012). During the Erligang periods (1650–1250 BCE) the population of the Early Shang polity was at least 1 million, and probably

more (Liu and Chen 2012, Wang 2014). The Shang city at Zhengzhou was huge, sprawling over 2500 hectares (Bagley 1999) with a population numbering in the hundreds of thousands. In other words, the North China sequence – first large-scale societies, then moralizing gods – supports the general conclusions of the moralizing gods paper.

Because records became abundant during the Late Shang (1250–1040 BCE), we can document such a transition in the case of China. Di (帝), the High God of the Shang, was the god of rain, snow, hail, wind, thunder, and disasters (Keightley 2000: 97). Di could summon natural phenomena to, for example, ruin harvests, but also could support or ruin political and military endeavours (Eno 2009: 71). The Shang king acted as an intermediary to appease Di through the correct ritual sacrifices. There is no evidence in the oracle bone records for Di as a moralizing force (Eno 2009: 100). The Mandate of Heaven (Tian, 天) became a central concept in the Western Zhou, making a significant change in the Chinese religious landscape. The first mention of the Mandate of Heaven is in a 998 BCE Western Zhou inscription in which King Kang claims that the Shang had lost the Mandate because of poor behaviour like drunkenness and overall bad governance (Eno 2009: 101; Li 2013: 144). It seems that Tian took on the role of ethical guardian, punishing rulers according to moral standards and the quality of their stewardship of the state (Eno 2009: 101).

Given that the first signs of a moralizing high god appear only during the Western Zhou period, why did the DRH experts code it differently? Consider the DRH data for “Chinese State Religion (Shang and Western Zhou)”. For China, 1750–850 BCE, the DRH asks whether there was supernatural monitoring of prosocial norm adherence. The answer provided by their expert coder is “yes.” It would be interesting to know what their expert thinks about the Shang–Zhou transition, but all we are given is a “yes”. The period encompasses four distinct eras that Seshat codes as such: (1) the last century of Erlitou, (2) Erligang (Early Shang), (3) Late Shang, and (4) Western Zhou. For the first two periods there are no written records and, therefore, it is not clear on what basis “supernatural monitoring of prosocial norm adherence” is coded as present. This surprising coding – one could use Slingerland *et al.*’s term and refer to it as a “coding error” – highlights a serious limitation of the DRH. In the end, however, it seems that what appears to be disagreement is actually agreement that moralizing gods were present in the Western Zhou dynasty, but the DRH’s 900-year time window puzzlingly extrapolates this presence back to the Shang period and beyond.

The Role of Research Assistants

Slingerland *et al.* repeatedly refer to the Seshat methodology as “Research Assistant (RA)-sourced data.” This is a misrepresentation of the data collection process on our project. For a more detailed description of the approach, see the accompanying introduction to Seshat by Turchin *et al.* Briefly, data collection in *Seshat* is accomplished as a collaboration between long-term trained RAs, expert historians/archaeologists who possess specialized knowledge about past societies, and social scientists who have deep knowledge about the scientific

theories that are to be tested with the Seshat data and analytic methods with which data will eventually be analyzed. Our experience over the past eight years has convinced us that only a truly collaborative effort can result in sufficient quantity and quality of data.

We agree with Slingerland *et al.* that expert historians and archaeologists are key to data collection. But they, by themselves, are not enough. Our initial trials showed that one cannot expect academic experts on past societies to have the time to answer hundreds of questions in detail. Coding data in Seshat is not a simple matter of answering ‘yes’ or ‘no’. Seshat strives for nuance allowing for interpretive diversity and envisaging the data as dynamic rather than fixed. A Seshat record has a complex structure, in which the code can be associated with various qualifiers that can indicate uncertainty and/or disagreement between sources. Most records include a narrative paragraph explaining the logic for selecting a particular code, accompanied by one or more references. Coding even one Seshat record is a significant amount of work and relying on university faculty to code hundreds of records is simply unrealistic. Fortunately, it is not necessary to do so, because, depending on the variable to be coded, most data can be coded by highly trained RAs using published sources (which are typically suggested by the experts). The rest of the data are coded by specific targeted questions to experts via in-person workshops (Levine *et al.* 2017), video conferencing sessions, and emails, which explain the various choices and ask for their feedback. This process takes time.

Contrary to what Slingerland *et al.* imply, Seshat RAs do not work in isolation. Their work is supervised by social scientists associated with the Seshat project, including the Directors, Editors for various regions or kinds of variables, and postdoctoral researchers (<http://seshatdatabank.info/seshat-about-us/seshat-who-we-are/>). This team-based approach has significant advantages over the DRH’s reliance on regional experts coding independently. Most professional archaeologists and historians, while deeply knowledgeable about their societies, are not steeped in the relevant social science theories that inform why (and how) we need to code a particular variable.

“Data Pasting” and “Data Filling”

Slingerland *et al.* observe that where a particular variable is judged to be present or absent at particular points in time, this is the only information that can be coded, whereas continued presence or absence of that same variable before, after, or between those time points constitutes “data pasting” and “data filling” – practices they imply are deliberately misleading. Seshat data is organized in a continuous time series, and coders seek to populate empty boxes with information about the presence, absence, or magnitude of a given variable at successive points in history for a given polity. If, for instance, we have direct evidence that iron smelting was widely practised in a given polity at certain known dates several centuries apart, we consider it reasonable to infer (in the absence of evidence to the contrary) that it was also practised in between those

time points. This is the type of entirely appropriate inference that Slingerland *et al.* describe as “data pasting”.

Closely related to this is what they call “data filling”, referring to the generalization of data from one, mostly proximate place and time source to another. This practice is essential under many circumstances. For example, when the Achaemenid Empire expands to occupy Susiana and Sogdiana, it is crucial to reflect this conquest by copying relevant data from the Achaemenid Empire to these regions during the time that it is in political control there. Most importantly, our approach is transparent, carefully justified, and its implications for data reliability and analysis open to discussion and critique.

Both methods described above are explicitly designed to address the problem that the historical record has many gaps. Experts commonly disagree on quite fundamental issues. Seshat was designed to capture not only the information about past societies (“what we know”), but also uncertainty and lack of knowledge. But while we do not want the database to exaggerate the precision of what we know, it is equally an error not to reflect such evidence as we may have, even if it is imprecise.

Slingerland *et al.* accuse us of applying brief snippets across wide swathes of time, arguing that this practice “promulgates false belief in a static *Christian* religiosity and organization, which is extremely misleading for inexperienced users and not useful for academic researchers.” Nevertheless, Seshat is not an effort to code all the richness of religious ideation and organization but focuses on specific aspects of it. For example, once Christianity becomes established in a region, its internal evolution is irrelevant to the question of whether moralizing gods are present (unless, of course, this evolution takes a drastic new direction in which beliefs in moral judgement in this life and the next are abandoned).

Data Freezing

Slingerland *et al.* claim that data used in the moralizing gods paper, which should have been frozen for the purposes of replication and reanalysis purposes, was altered subsequently. This is false. Papers published as a result of the Seshat project routinely freeze the data used for replication purposes, and the *Nature* paper was no exception. In the moralizing gods paper, Table S2 summarized all of the key data used in our analyses along with over 90 references supporting them, and the entire dataset containing tens of thousands of records used for our analyses was published as a Supplementary Data file. These data were frozen at *Nature*’s online website and cannot be altered.

Nevertheless, Slingerland *et al.* mistakenly interpret the data provided on the Seshat website (<http://seshatdatabank.info/data/>) as a replication dataset and castigate us for updating the information it contains. In reality, not only did we publish the full dataset on *Nature*’s website, we also voluntarily published all analysis scripts necessary to replicate our analyses to GitHub

(<https://github.com/pesavage/moralizing-gods>), as stated in our paper. Thus, our research and the data that we published to support it meets (and indeed, exceeds) all the normal standards for openness and reproducibility in scientific research.

In contrast with the machine-readable format published at *Nature*, the Seshat web viewer was designed to allow us to share our research and the sources from which the data were drawn with colleagues and interested members of the public. We incorporated features into this site that allow us to capture feedback and spot errors in the data. As data are added or amended in the Seshat Triplestore, the web viewer automatically incorporates these changes. In addition, a data mirror on <http://Seshatdatabank.info/nature/> was provided to the reviewers of the *Nature* article so that they could see the as yet unpublished variables. Once the paper was published, we replaced it with a public link to the permanent Seshat web viewer.

Periodization Structure

We have considered above many drawbacks in relying exclusively on experts to construct a database. Another serious limitation of this approach is that it produces an unbalanced sample of past societies. There are hundreds, perhaps thousands of scholars studying Ancient Greece, many fewer experts on Ancient Persia, and one or two or even none on many regions of Melanesia. A community-based coding effort, advocated by the DRH groups, results in a plethora of data on a few well-studied past societies, and little or none on most others.

As an example, Figure 1 shows the geographic distribution of pre-1800 societies in the DRH with an answer to the DRH question as to whether there is “supernatural monitoring of prosocial norm adherence”. Note that sampling outside the “civilization belt”, stretching from the Mediterranean to China, is extremely sparse. In contrast with the DRH, Seshat’s global approach is in tune with concerns among historians for whom Western-centric bias and peripheralization of the rest of the world is no longer acceptable. The Seshat project has gone to great lengths to design a balanced sample of past societies, stratified by world region and the timing (early, intermediate, late) of the appearance of politically centralized societies (Turchin *et al.* 2015).

Figure 1. The geographic distribution of pre-1800 societies in the DRH with an answer to the DRH question as to whether there is “supernatural monitoring of prosocial norm adherence”.

The DRH sample is even less useful for analysis than the map above indicates. A key innovation of Seshat is that the databank traces how societies develop over time. For that, it is useful to have a time sequence of how particular variables changed. Thus, having data for only a single time-period in a given region (e.g., the Mongol Empire in Mongolia) tells us nothing about how religion evolved in this region.

Slingerland *et al.*'s critique of several aspects of the Seshat methodology shows basic misunderstandings of the Seshat sampling procedure, particularly in the Appendix which claims to provide “specific examples of errors in coding in the Seshat Databank.” Here they state, for example, that “the polity defined as ‘Latium’ is inherently problematic.” Yet Latium is not a polity in our scheme, it is a natural geographic area (NGA; cf. Turchin *et al.* 2015). In several other places, the ‘errors’ attributed to Seshat result from a failure to appreciate this key distinction between polity and NGA. Seshat’s 30 NGAs provide merely a sampling scheme, used to identify the 414 polities constituting our units of analysis. A closely related misunderstanding is apparent in the Slingerland *et al.* charge that Seshat “flattens” historical information into distinct 100–300 year periods. In fact, the Seshat periods are simply a convenience for organizing data. When there is evidence of change during these periods, such changes are reflected by attaching a specified time range to them. Of course, when a particular characteristic doesn’t change for a long period of time, the resulting trajectory is indeed a flat horizontal line. Ironically, the DRH is much more vulnerable to the charge of ‘flattening’ historical information. For example, the entry for “Religion

in Mesopotamia (3300 BCE – 539 BCE)” flattens nearly three thousand years of history in this region into a single period. Another example is for “Chinese State Religion (Shang and Western Zhou), 1750 BCE–850 BCE”, considered above.

There is no well-defined consensus among historians that we must confine ourselves to archaeological periodization. Most Roman historians do not use archaeological periodization for all periods “that pre-date the turn of the eras.” Traditional periodization (e.g., “Roman Republic”) remains widely accepted, with minor variations, in the discipline of Roman history. For the earliest part of the Roman polity, these dates are conventional and ultimately drawn from accounts which are to some degree legendary (for Romulus wholly legendary). Some historians adopt a scrupulously archaeological approach at the expense of the abundant (later) literary evidence; others prefer to use both archaeological and literary sources. For the eighth through fifth centuries, the terminology varies. Two key sources, Cornell (1995) and Forsythe (2006), draw on archaeological evidence while also engaging in extensive critical analysis of the texts which purport to give historical information about early Rome. Cornell opts for a hybrid approach (chapter titles include e.g., “The Power of Rome in the Sixth Century,” “The Beginnings of the Roman Republic”). His analysis makes clear that there was a period during which Rome was ruled by *reges*, non-hereditary kings. Similarly, Forsythe has a chapter devoted to “Archaic Italy c. 800–500 BC” and another on “Rome During the Regal Period.” While critics of Seshat may disagree with this practice on methodological grounds, it is hardly the case that Cornell and Forsythe “have failed to distinguish between emic, mythologized accounts and etic scholarly reconstructions of the past”, as claimed by Slingerland *et al.* Likewise, it does not follow that because Seshat describes a “Roman Kingdom” which virtually all historians agree existed prior to the Republic, Seshat is also suggesting that Romulus and/or Numa were historical figures. Finally, we should emphasize that while the critics deploy an attack on Seshat’s periodization in order to suggest that the information in the database is suspect, they do not point out any errors in coding that may have arisen from the periodization.

On the uses and abuses of open science

The Seshat project has gone to unprecedented lengths to achieve transparency in its research processes and outcomes. It has sought to share publicly the evidential basis on which the moralizing gods paper was based, allowing others to replicate the findings and test alternative assumptions using the same dataset. One significant cost of such openness is that it lays researchers open to hostile critique. We believe this is a price worth paying because it also helps improve the quality of the database and the work resulting from it.

The full database, which we made freely available on the web (<http://seshatdatabank.info/data/>) in tandem with the publication of the moralizing gods paper, should be considered in addition to the raw machine-readable data and analysis scripts published alongside our paper. These spreadsheets contain the precise dataset that was used in the moralizing gods

paper, presented in a format that is optimized for reproducibility. As such, our research and the data we published to support it follows the highest standards for openness and reproducibility in science. Alongside the publication of the data we also made the dataset available through a website designed for sharing our research and the sources from which the data was drawn with reviewers and interested members of the public. We incorporated features into this site that allowed us to capture feedback to help us spot errors in the data. This was in addition to, not instead of, normal standards in research data openness.

It should be noted that from a data curation perspective, Seshat presents many special challenges (Peregrine *et al.* 2018). Our computational experts have rolled out more than a dozen software updates to attempt to clarify aspects of the sourcing and checking of the data that we have published. They have updated and corrected hundreds of variables, improved references, and responded to multiple critiques of our process. We continually strive to improve our data and the clarity of the checking, provenance, and validation and to publish and present our data in a way that allows us to enhance it over time and to provide more contextual information. But the perfect should not be the enemy of the good. We have made our data publicly available inviting feedback, and we have collectively received hundreds of minor criticisms, corrections, and improvements through that process. We would like to be able to publish the entire data curation process pertaining to every variable that we use, and we continue our efforts to improve our software in this regard, but we will probably never be able to fully capture, crystallize, and communicate the breadth of expert opinion that has been applied to overseeing our data collection, modelling, and results.

The benefits of open science are clear and compelling, the technical limitations notwithstanding, but the approach is also open to abuse. For example, it enables attacks to be launched at a lightning pace, using data made available for replication, in ways that can be hard to regulate or rebut. Such attacks can resemble more a smear campaign than an effort at scientific debate, especially when using social media platforms, such as Twitter, that privilege sound bites over carefully weighed argument and evidence. In these circumstances, even erroneous critiques that have not been subjected to peer review such as the ones made by Beheim *et al.* (2019) can be made to appear far more devastating and authoritative than is actually the case (see Savage *et al.*, Under review for rebuttal).

In the present context, our openness with the data prompted an unexpectedly aggressive attack in pre-published papers that were widely advertised on social media, including ones tagging leading media outlets such as the *New York Times* and *Scientific American* (e.g., <https://twitter.com/JoHenrich/status/1124038389995585536>). These could have created an impression among our academic colleagues that the moralizing gods paper was of scandalously poor quality and did not deserve to be published. Most shockingly, these online statements included an allegation that Seshat

researchers lied about the involvement of one of our expert historians, Vesna Wallace, in providing professional advice on data used in the moralizing gods paper. Specifically, Slingerland *et al.* claim that Wallace played “no part in any aspect of coding, vetting or recommendation of sources”, citing “Personal communications to Brenton Sullivan and Edward Slingerland.” That allegation is disproved by documented evidence. In January 2017, Wallace participated in a two-day Seshat Axial Age workshop held in Oxford (Levine, Larson, Baines, Wallace, and ter Haar 2017). During this workshop, as part of a subgroup that also included the Seshat RA Agathe Dupeyron and the Seshat director Peter Turchin, Wallace generously shared her expertise on aspects of Buddhist religion in Mongolia, India, Japan, Indonesia, and Cambodia. Later, in a series of emails with postdoc Patrick Savage in October 2017, she confirmed that our key dates listed for the spread of Buddhism to these five regions (ultimately published in Table S2 of our original paper) were accurate, and provided additional nuance for our description of the system of karmic retribution that were also incorporated into Table S2. Note that for the purposes of the moralizing gods paper we sought and obtained her expertise on the first appearance of Buddhism, not on other characteristics of these diverse regions.

Relatedly, Slingerland *et al.* charge: “After our e-mail of March 25, 2019 to the authors, they appear to have gone through the *Seshat* data provided to *Nature* referees and systematically changed their claim about Prof. Wallace having vetted all of the religion and ritual variables to this statement: ‘verified key data regarding earliest appearance of moralizing gods/doctrinal rituals’.” But we never made any claim that Prof. Wallace or other experts “vetted *all* of the religion and ritual variables.” Our logs show that in September 2018, as a response to comments arising during the review process, we added statements indicating which experts had provided input for each polity, and included descriptions of the general topics on which expert input had focused, such as “Vesna Wallace (Religion or Normative Ideology System)”, “Barend ter Haar (Ritual variables, Religion or Normative Ideology System),” etc. Because the key variable for analyses reported in our original Table S2 was the first evidence of moralizing gods, we focused our expert checks on this question. Slingerland *et al.* claim that “Even if these scholars have verified this data concerning “earliest appearance,” this statement can logically only apply to one polity—whichever was characterized by this earliest appearance.” The logical error to which they refer is unclear. Confirming the earliest evidence of moralizing gods in a given NGA logically requires confirming the absence of evidence for earlier polities. On March 27, 2019 we clarified this statement further as follows: “Vesna Wallace (verified key data regarding earliest appearance of moralizing gods/doctrinal rituals)”. Slingerland and colleagues fail to mention that we notified them by email that we would update the website in precisely this manner to better address the concerns that we thought they had emailed us about in good faith.

Slingerland *et al.* make the following observation which, in light of the above, is chilling: “The apparent lack of appreciation for historical scholarship that... [Seshat’s] coding strategy displays runs the risk of permanently alienating the

community of academic historians". Academic historians and archaeologists are the life blood of Seshat but the multi-pronged attack launched by our critics, rather than lack of appreciation of historical scholarship, could be enough to deter some scholars from contributing to this project in the future. Indeed, the allegation that Wallace played "no part in any aspect of coding, vetting or recommendation of sources" for Seshat could hardly have been better designed to alienate experts. We take this opportunity to emphasize that Wallace's contributions to Seshat are significant and highly valued. Her authorship on our PNAS paper (Turchin *et al.* 2018) was fully deserved and her contributions to the data used in the moralizing gods paper contributed directly to our main analysis. Slingerland *et al.*'s claim above is not only unfair to Wallace, it also wrongly accuses the Seshat team of serious malpractice by allegedly lying about her involvement in the project. Seshat will continue to work to expand its growing community of expert researchers, and welcomes back any who might have been alarmed by unwarranted allegations.

Conclusions

We stand at a potential turning point in the history of history. A number of highly influential scholars of the past couple of centuries, including Marx, Spencer, Tylor, Frazer, and Durkheim, dreamt of establishing generalizable theories of history but they were held back by the 'cherry picking problem'. That is, theories of history – from grandiose visions of economic and technological determinism through to the idea that the division of labour in society evolves through discernible stages – have always emphasized evidence that supported the theory, while less congenial evidence was rejected or overlooked. What is radically new about the approach adopted in our *Nature* paper is that it tests theories of history based on a serious effort to avoid bias in the selection of data by coding for features of social complexity, religion, and ritual, in exactly the same way across hundreds of societies throughout world history. All data and analysis code are publicly available so that colleagues can inspect it, replicate and criticize our efforts, and run analyses of their own.

An effort to do precisely that was pre-published online by Beheim *et al.* (2019) to coincide also with the pre-publication of Slingerland *et al.* (2019). Both papers were widely advertised by their authors on Twitter as effectively debunking our moralizing gods paper. As we have seen, Slingerland *et al.*'s critique has little or no direct relevance to our main findings. In the case of Beheim *et al.*, the claim was that reanalysis of our data "correcting" certain assumptions reverses our original analyses. But Savage *et al.* (under review) have shown that this is not the case and that their reanalyses, when appropriately performed, confirm the robustness of our original conclusions. Both Beheim *et al.* and Slingerland *et al.* took advantage of social media to advertise their hyperbolic criticisms of our work in ways that were difficult to rebut effectively at the time.

The speed of the destructive attacks on Seshat described above contrasts with the much slower, painstaking pace at which our database has been constructed.

Work to build Seshat has been underway for nearly a decade and yet we have only recently reached a point at which we believe specific subsets of that data are ready for analysis and publication. We continue to strive to improve our data and the clarity of its provenance and validation, and to publish and present our data in such a way that we can improve it further over time, as well as providing more contextual data. This goal should be tempered with the caution that it is impossible to do so for *every individual* data record in Seshat (of which there are currently ca.300,000). By using such an unprecedented degree of transparency, in the short-term we open ourselves to hostile cherry-picking of imperfections that any large databank will have. But we are willing to pay this price, because it makes the Seshat data better in the long-term.

Our goal is to use Seshat to test theories of global history in increasingly diverse ways, fully realizing that the Seshat vision will require the input of very large numbers of experts from fields as diverse as history, archaeology, classics, linguistics, cultural anthropology, and comparative religion (among others). Wariness of scientific methods in these fields is widespread, so it is no mean feat to have already attracted such large numbers to participate in the Seshat enterprise. Continuing to do so will mean not only retaining our list of current experts but attracting many more. It remains to be seen whether this is achievable. Those who are heavily invested in theories that can be refuted by Seshat data, or who aspire to building rival databases, may understandably seek to critique our work, but we hope this will be done increasingly in ways that are constructive and do not discourage humanities scholars from working with us.

If our new approach to the study of global history survives and flourishes, this will be very good news for the humanities. It will not change the fundamental methods of historical enquiry but will complement them. Existing historical research will become more thoroughly integrated with many areas of the social sciences and attract more resources. On the other hand, it will be bad news for many *theories*. Few theories will survive unscathed. But that is a desirable situation scientifically. We are confident that, despite disagreements between the various teams of scholars involved, we all share the same fundamental goal of integrating the sciences and humanities through the creation of new interdisciplinary databases, helping us to build new knowledge to answer old questions about the human condition.

Acknowledgements

We thank the following for feedback (although this does not imply that they necessarily endorse any of the views expressed here): Pascal Boyer, Daniel Hoyer, and Christopher Kavanagh. Vesna Wallace reviewed and approved our description of the nature of her work for Seshat, in the penultimate section of this paper. H.W. and P.F. were supported by an Advanced Grant from the European Research Council (ERC) under the European Union's Horizon 2020 Research and Innovation Programme ('Ritual Modes: Divergent modes of ritual, social cohesion, prosociality, and conflict', grant agreement no. 694986), H.W. by

an award from the Templeton World Charity Foundation entitled ‘Cognitive and Cultural Foundations of Religion and Morality’ (TWCF0164), P.E.S. by a Grant-in-Aid for Young Scientists from the Japan Society for the Promotion of Science, Keio Research Institute at SFC Startup Grant, and a Keio Gijuku Academic Development Fund grant, and T.C. by a Starter Grant from the European Research Council (ERC) under the European Union’s Horizon 2020 Research and Innovation Programme (‘The Cultural Evolution & Ecology of Institutions’, grant agreement no. 716212).

References

Beheim, B., Q. Atkinson, J. Bulbulia, W. Gervais, R. D. Gray, J. Henrich, M. Lang, et al. Under review. “Corrected Analyses Show That Moralizing Gods Precede Complex Societies but Serious Data Concerns Remain.” *Preprint*: <https://doi.org/10.31234/osf.io/jwa2n>.

Bagley, R. 1999. “Shang archaeology.” In *The Cambridge History of Ancient China: From the Origins of Civilization to 221 BC*, ed. M. Loewe and E. L. Shaughnessy, 124–231. Cambridge: Cambridge University Press.

Cornell, T. J. 1995. *The beginnings of Rome: Italy and Rome from the Bronze Age to the Punic Wars (c. 1000-264 BC)*. London: Routledge.

Eno, R. 2009. Shang State Religion and the Pantheon of the Oracle Texts. In: *Early Chinese religion: Part one: Shang through Han (1250 BC–22 AD)*, ed. J. Lagerwey & M. Kalinowski, pp. 41–102. Brill.

Forsythe, G. 2006. *A critical history of early Rome from Prehistory to the First Punic War*. Berkeley, CA: University of California Press.

Hsu, Cho-yun and K. M. Linduff. 1988. *Early Chinese Civilization Series*. New Haven, CT: Yale University Press.

Keay, J. 2009. *China, A History*. London, HarperPress.

Keightley, D. N. 2000. *The Ancestral Landscape: Time, Space, and Community in Late Shang China (ca. 1200-1045 B.C.)*. Berkeley: University of California Press.

Levine, J., J. Larson, J. Baines, V. Wallace, and B. ter Haar. 2017. “Experts Reflect on Their Experiences Working with Seshat: Global History Databank.” *Seshat: Global History Databank Blog*. http://seshatdatabank.info/expert_reflections.

Liu, L. and X. Chen. 2012. *The Archaeology of China: From the Late Paleolithic to the Early Bronze Age*. Cambridge: Cambridge University Press.

Peregrine, P. N., R. Brennan, T. Currie, K. Feeney, P. François, P. Turchin, and H. Whitehouse. 2018. “Dacura: A new solution to data harvesting and knowledge

extraction for the historical sciences.” *Historical Methods: A Journal of Quantitative and Interdisciplinary History*. doi:10.1080/01615440.2018.1443863

Savage, P. E., H. Whitehouse, P. François, T. E. Currie, K. C. Feeney, E. Cioni, R. Purcell, et al. Under review. “Reply to Beheim et al.: Reanalyses confirm robustness of original analyses.” *SocArXiv* preprint (doi forthcoming).

Slingerland, E. and B. Sullivan, 2017. “Durkheim with Data: The Database of Religious History.” *Journal of the American Academy of Religion*, 85 (2): 312–347. <https://doi.org/10.1093/jaarel/lfw012>

Slingerland, E., W. Monroe, B. Sullivan, R. Faith Walsh, D. Veidlinger, W. Noseworthy, C. Herriott, et al. In press. “Historians Respond to ‘Complex Societies Precede Moralizing Gods throughout World History.’” *Journal of Cognitive Historiography*. Preprint: <https://doi.org/10.31234/osf.io/2amjz>.

Thorp, R. L. 2006. *Encounters with Asia: China in the Early Bronze Age: Shang Civilization*. University of Pennsylvania Press: Philadelphia

Turchin, P., R. Brennan, T. E. Currie, K. C. Feeney, P. François, D. Hoyer, J. G. Manning, A. Marciniak, D. Mullins, A. Palmisano, P. Peregrine, E. A. L. Turner, and H. Whitehouse (2015). *Seshat: The Global History Databank. Cliodynamics: The Journal of Quantitative History and Cultural Evolution* 6(1): 77-107. <http://escholarship.org/uc/item/9qx38718>

Turchin, P., T. E. Currie, H. Whitehouse, P. François, K. Feeney, D. Mullins, D. Hoyer, et al. 2018. “Quantitative Historical Analysis Uncovers a Single Dimension of Complexity That Structures Global Variation in Human Social Organization.” *Proceedings of the National Academy of Sciences* 115 (2): E144–51. <https://doi.org/10.1073/pnas.1708800115>.

Turchin, P., H. Whitehouse, P. François, et al. In press. “An Introduction to *Seshat: Global History Databank*”. *Journal of Cognitive Historiography*. doi: 10.31234/osf.io/2amjz

Wang, H. 2014. *Writing and the Ancient State: Early China in Comparative Perspective*. Cambridge: Cambridge University Press.

Whitehouse, H., P. François, P. E. Savage, T. E. Currie, K. C. Feeney, E. Cioni, R. Purcell, et al. 2019. “Complex Societies Precede Moralizing Gods throughout World History.” *Nature* 568 (7751): 226-229. <https://doi.org/10.1038/s41586-019-1043-4>.

Zhu, X. (Translated by P. B. Ebrey) 1991. *Chu Hsi's family rituals: a twelfth-century Chinese manual for the performance of cappings, weddings, funerals, and ancestral rites*. Princeton, NJ: Princeton University Press.