



HAL
open science

Laws of Divine Power in the Greek Pantheon

Laurent Gauthier

► **To cite this version:**

| Laurent Gauthier. Laws of Divine Power in the Greek Pantheon. 2021. hal-03289609v1

HAL Id: hal-03289609

<https://univ-paris8.hal.science/hal-03289609v1>

Preprint submitted on 17 Jul 2021 (v1), last revised 5 Aug 2021 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Laws of Divine Power in the Greek Pantheon*

Laurent Gauthier[†]

July 17, 2021

Abstract

In ancient Greek polytheism, worshippers could choose which gods they would address, and in doing so they expected some form of benefit in a quid pro quo relationship. We look into the optimal choice of which god to worship as a function of the presumed strategy of the gods for returning favors to the worshippers, and relate it to a form of divine efficiency measure. At the equilibrium the model also shows that the least-worshipped god receives at least a certain volume of devotion. We propose two different approaches that may account for the assumed divine efficiency measure, one based on projecting characteristics of human performance onto the gods, and the other based on a random growth model for the benefits of addressing each god. Both approaches imply that observed acts of worship would follow a type of power law. We gathered data from a large volume of epigraphic and literary sources on actual acts of worships from the ancient Greeks, and it allows us to show that the distributions of these acts at the polis level effectively follow power laws with a particularly high degree of regularity. The number of votive acts towards the least-worshipped gods also match the model's prediction. We test the extent to which the known characteristics of the poleis affect the shape of these distributions, and find little explanatory power. The shape of the distribution of votive acts across gods hence appears to have followed a general law for the Ancient Greeks.

Keywords: Ancient Greek polytheism, religious studies, theory of religious economy, game theory, random growth models, power laws, human performance modeling

*Draft working paper - comments welcome! This research has benefited from substantial input from Corinne Bonnet and Sylvain Lebreton, from the ERC Advanced Grant 741182 “Mapping Ancient Polytheisms. Cult Epithets as an Interface between Religious Systems and Human Agency”, Université Toulouse – Jean Jaurès (<https://map-polytheisms.huma-num.fr>); I thank them for their many remarks. I also thank the participants in the 2021 LED Macro Seminar at Paris 8 University.

[†]Email: laurent.o.gauthier@gmail.com. Laboratoire d'Économie Dionysien, Université Paris 8 Saint-Denis Vincennes, EA 3391 - Bâtiment D, 2, rue de la Liberté, 93526 Saint-Denis, France.

It is well known that the ancient Greeks venerated many gods, and the names of the pantheon of twelve gods are generally well known; Zeus, Apollo and Poseidon are not far from household names. This familiarity hides in fact the extent to which these ancient Greek gods are still actively researched by historians and philologists. One particular angle through which the understanding of ancient religion has been furthered is the study of the ways in which the gods were addressed, which led to the creation of a few large electronic datasets. As Brulé (1998), Brulé (2015) and Parker (2003) notably stressed, one cannot think of a god as a perfectly defined entity, because the divine notion attached to a name varies across cities¹, and the qualifications with which they are addressed vary significantly. At the first order, one may wonder how the logic in which people worshipped multiple gods could optimally operate.

In this paper, we seek an endogenously optimal explanation for the choices by the ancient Greeks of which god(s) to address. We develop a simple economic model that relates the choices of gods to these gods' assumed performance in a competitive setting. The study of ancient Greek religion saw an opposition build up in the 1970s, according to Versnel (2011), between those seeing *kosmos* (order) into it and those seeing *chaos*. The proponents of *kosmos*, following Jean-Pierre Vernant's historical psychology approach, considered that logic in the Greek gods had to be found through structuralist study, and that each one had to be understood in relation with the others, and was defined² "by the structure of relations that oppose and unite it to the other powers that constitute the divine universe". The proponents of *chaos*, following Walter Burkert, considered the gods formed³ a "scattered and heterogeneous pantheon, a mythology of bits and pieces". These two perspectives have converged in some ways over the years, as Versnel argues, but the present work will clearly insert itself in the Vernantian vein, since we will explicitly assume a strong logic at play within the Greek pantheon.

Resorting to formal modeling to account for phenomena related to religion is not new. Indeed, in religion studies, game theory has been used, albeit not extensively, to explain the existence and salient characteristics of religion. Extending the field of religious studies, Bulbulia (2010) and Bulbulia and Frean (2010) proposed an evolutionary model for religious cooperation, and showed how sacred culture enhances cooperative prediction and response among anonymous partners⁴. The evolution of costly displays is discussed in Henrich (2009). In his small book, Chwe (2001) discussed many aspects in which religious ritual can be analyzed as a rational action

¹The notion of a "city" in ancient Greece is very different from our modern perspective; keeping with usage in classics, we will make use of the specific term of *polis* which is the only one that can capture this specificity. The *polis* is an independent group of people who have their own specific political rule.

²See Versnel (2011), p. 27.

³See Versnel (2011), p. 28.

⁴In other words, a common sacred culture is shown to be one way of avoiding a lack of cohesion in the classical "stag hunt" problem, as originally defined by Jean-Jacques Rousseau.

from the standpoint of simple game theory. Using a distinct approach from game theory, Kaše, Hampejs, and Pospíšil (2018) have considered the religious rituals of early Christianity from the perspective of agents-based modeling, in contrast with system dynamics.

More specifically from the standpoint of economics research, some early work by Carr and Landa (1983) addressed religion as part of the study of identity and decision to belong to a group. Covering a wider ground, Iannaccone (1998) and Iyer (2016) present literature reviews for the economics of religion in all its facets. Economics research has focused on the structure of religious organizations, among other things. For example, Levy and Razin (2007) and Levy and Razin (2012) used symmetric Prisoner Dilemma games to represent religious organizations and showed that more demanding religious organizations attract less members, but these members are more cohesive. Iannaccone, Haight, and Rubin (2011) show how secular leaders influence the location of sacred places. The sacred places then gain widespread authority precisely because they lie beyond the centers of secular power. In McBride (2016), the focus is on showing how rituals create religious authority. Levy and Razin (2014) develop a model of social signalling of displays of religious behavior and cooperative acts in religious organizations. They show that ritual-based religions require a costly involvement, which acts as a signal and also implies a higher level of coordination in social interactions and a higher incidence of mutual cooperation⁵. In this body of literature, religions and some of their specific features are shown to optimally exist in an equilibrium. Applying economic theory to questions of beliefs has raised some debates and questions⁶, as it has been argued that religious acts may not be equated to picking products on a shelf. From the standpoint of the theory of religious economy, one can consider religion as a social good more than simply a product.

Empirical tests have confirmed many economics-based theoretical models. In Ruffle and Sosis (2007), time-consuming religious rituals are explained as promoting cooperation, and the authors offered empirical tests on the cooperative behavior of members of religious and secular Israeli kibbutzim. Also, Johnson (2005), relying on the theoretical models that suggest human cooperation may be promoted if people believe in supernatural punishment, empirically tested across a large number of cultures⁷ whether the likelihood of supernatural punishment was related to observable cooperation. In Power (2017), the author empirically showed that people are attending to the full suite of religious acts carried out by their peers, using these visible signals to discern multiple aspects of their character and intentions.

⁵This framework could have interesting implications for the study of ancient Greek rituals, which were precise and constraining; this is not our focus in the present work however.

⁶See Bankston (2002) for example. There are also some discussions on how academics from a given religious background may be able to better apply this economic framework, see Iannaccone (2005).

⁷Using the so-called Standard Cross-Cultural Sample (SCCS) of 186 human societies around the globe, spanning time and space.

While the literature mentioned above, both in religion studies and in economics, has concentrated on fundamental explanations for the existence of religion or of some of its rituals, our approach will differ. Our initial focus is not on the reason why religion may optimally exist in a certain form, but rather, given that it exists, how it may be optimally practiced from the standpoint of both the population and the assumed characteristics of the gods. In particular, the specific aspects of polytheistic religion will play an important role in our analysis. There are few empirical studies in religion science that examined polytheism. Gries, Su, and Schak (2012) proposed specific methods of study and applied them to a survey in Taiwan.

There is also little research on polytheism within the framework of the theory of religious economy. Ferrero and Tridimas (2018) developed a model where people assumed certain behavior for the gods and made their own decisions as a consequence of this assumed decision logic. They argue that the overlap between different god's prerogatives created inefficiencies, and that could be an explanation for the convergence towards monotheism that was observed at the end of Antiquity. We will not in our approach distinguish between different gods' jurisdictions, but will instead consider their differences in terms of how easily it supposedly came to them, or how efficient they were, to give benefits back to their mortal worshippers.

In the next section, we examine some aspects of ancient Greek religious thought, drawing some examples from Greek literature, and we identify important motivations in ritual behavior. In the following section, we develop a theoretical model that explains the distribution of choices of which god or gods to worship within a *polis* as a function of certain assumed qualities of the gods. We then relate such qualities to human performance in competitive sports or games, and infer that observations of which god or gods were addressed could follow a power law with exponential decay mixed with a uniform distribution. We also develop a random growth model whereby people update their perception of the gods' efficiency based on random life events, and show that in this case the distribution of votive acts would be expected to approximately follow a power law distribution as well. Then, in the third section, we combine data from two separate sources: the BDEG project on ancient Greek divine epithets and the extensive ancient Greek *poleis* survey, the POLIS data. We structure this data so that we can measure the distribution of votive acts across gods, and test the theoretical model from the first section. We find that the model's predictions are close to empirical observations, and also look for certain explanatory factors for the degree of efficiency that worshippers assumed for the gods.

Our approach fits within the methodological framework developed in Gauthier (2020b). We are directly exploiting primary sources for the most part in order to better comprehend polytheist behavior, at a large scale. By rescaling the data in a particular way, we are not attempting to explain it, but rather to make its internal logic clearer.

1 Some Aspects of Ancient Greek Religion

We will argue that from the ancient Greek’s perspective, the relationship between men and gods was fundamentally an exchange, and that we can set this exchange in economic terms. On one hand, people would carry out votive acts and rituals which were expected to please the gods⁸. Reciprocally, the gods were expected to give back thanks, *charis*, a form of *quid pro quo* through which the gods would effectively pay thanks to the worshippers through some actions⁹.

The nature of this exchange between men and gods is discussed and probed in a good amount of detail by Plato in *Euthyphro* throughout the book. In particular, the eponymous priest states:

“In summary, what is pious is to know what to say and do that is agreeable to the gods, either by praying or by doing sacrifices, and it ensures the safety of families and cities¹⁰”

In this sentence, the term for “agreeable” is literally expressed as “bringing *charis*”. With an interestingly explicit comparison, in one of his replies, Socrates remarks further:

“Hence it seems to me piety is a kind of commercial traffic between men and gods.¹¹”

Aphroditis expresses the gods’ need for being honoured in Euripides’s play *Hippolytus*, and this need actually causes the whole intrigue:

“Because this is essential to the race of the gods, they rejoice in being worshipped by men¹²”

If we look further back in time in Homer’s work for example, the same notion is also clearly expressed by Telemachus in *Odyssey*, book XVI, when he mistakes his father for a god, using the same term:

“Be good for us, so that we will give you agreeable sacrifices and golden, well-made gifts: spare us.¹³”

These mentions of the relationship between men and gods only address the “amount” of effort

⁸For example, gods were supposed to enjoy and feast on the smoke raising from various meats cooked during sacrifice, see Ekroth (2011).

⁹For example, in one inscription mentioned in Bonnet and Lebreton (2019), a mother invokes the goddess Artemis for the benefit of her children.

¹⁰*Euthyph.*, [14b] τόδε μέντοι σοι ἀπλῶς λέγω, ὅτι ἐὰν μὲν κεχαρισμένα τις ἐπίστηται τοῖς θεοῖς λέγειν τε καὶ πράττειν εὐχόμενός τε καὶ θύων, ταῦτ’ ἔστι τὰ ὅσια, καὶ σώζειε τὰ τοιαῦτα τούς τε ἰδίους οἴκους καὶ τὰ κοινὰ τῶν πόλεων.

¹¹*Euthyph.*, [14e] Ἐμπορικὴ ἄρα τις ἂν εἴη, ὧς Εὐθύφρων, τέχνη ἢ ὁσιότης θεοῖς καὶ ἀνθρώποις παρ’ ἀλλήλων.

¹²*Hipp.*, [7–8] ἔνεστι γὰρ δὴ κὰν θεῶν γένει τόδε: / τιμώμενοι χαίρουσιν ἀνθρώπων ὑπο.

¹³*Od.*, book XVI[184–5] ἀλλ’ ἴληθ’, ἴνα τοι κεχαρισμένα δῶμεν ἱρὰ / ἥδὲ χρύσεια δῶρα, τετυγμένα: φείδεο δ’ ἡμέων.

going in both directions, and not the fact that the gods may have specific jurisdictions or specialties. The list of the skills typically associated with each god is well known, and Ferrero and Tridimas (2018) based their approach of polytheism on the particular fields for which each god is recognized. By projecting the multidimensional characteristics of the gods onto a single measure, our work hypothesis is strong.

However, as we will see in our discussion of the data, the gods that received the most dedications within various *poleis* were not always the same. At the same time, it would seem a reasonable assumption that the things people wanted help with across *Hellas* were generally the same. Hence, the roles of the gods were to some extent fungible, and one could not argue that choosing between one or the other only depended on the gods' particular domains¹⁴.

Defining a god precisely, and in particular in opposition to other gods, is not straightforward, and simply stating that there are a certain number of gods is not innocuous. Should *Poseidon* and *Poseidon Phukios* ("of the algae") be considered the same¹⁵? In Mantinea, one counted five different ways of addressing Zeus¹⁶. The porous border lines between god denominations make the identification of each one non trivial, as well as their counting. We will nevertheless consider that some lines are drawn between all the acts of worship that one could consider, and that they can be grouped along the particular divine characteristics that our formal approach will need to use; so for our purposes, we will count a single *Poseidon* in the example above¹⁷. The complex onomastic formulae that may have been used¹⁸ are not accounted for in our approach.

Relying on these observations, we can ascribe basic motivations to worshippers, expressed in economic terms as utility maximization, which not only drive their own decisions but also the behaviors they assume the gods to follow. Since votive acts by the ancient Greeks were generally for a specific request, it seems reasonable that people would act rationally in order to maximize what they obtained in return.

The essential features that could underpin a formal model are as follows:

1. The different gods are assumed by the worshippers to offer a certain quantity of something that is beneficial to the worshippers. The quality or the nature of what is asked is also naturally important and would be carefully specified by the people participating in ritual; nevertheless we assume that this can be reduced down to a quantifiable metric, the *charis*.

¹⁴We will consider some data that supports this argument in Table 4 further down.

¹⁵Lebreton (2019) cites some examples of this kind.

¹⁶See Zaidman and Schmitt Pantel (1992), p. 212.

¹⁷There are some cases in the data where we actually separate the denominations, when the epithets are relative to a location.

¹⁸Associating several gods in a single invocation, listing particular characteristics (epicleses) for them.

2. The people who participate in worship share the total benefit. Note that in the models that generally explain religion itself, most approaches necessitate that the utility derived from religion would be increased by common participation. See Iannaccone (1998) or Iyer (2016) for descriptions of the models of religious structure relying on club goods. This is not in opposition with our approach, which takes the existence of religion for granted; we are only addressing the manner in which religion is practiced, but the participants could still withdraw large benefits from the fact that they all have a religious system in common. We consider that the total *charis* that a god can offer must be shared by all the worshippers; there is therefore no explicit advantage in everybody choosing the same god¹⁹. The sharing may not be equally received: a king praying to a god among a crowd may expect (and may be expected) to receive much more than the other participants²⁰.
3. Each god is assumed to be seeking to maximize the benefits they gain from the attendance of their worship, and that they derive a benefit equal to the number of people who worship them.
4. The gods incur some internal cost in providing the *charis* from point 1, and that cost is a function of the total amount they give to worshippers. The worshippers assume that the gods need to spend time and effort to give them benefits, and hence it is not a free resource. For simplicity, we consider that this cost is proportional to the total amount of *charis* they give back. Different gods incur different proportional costs in giving blessings to humans, as they all have different types and degrees of power²¹.

2 Worship Decision Model

In this section, we develop a theoretical model that could account for the distribution of votive acts in a given polity. In the following subsection, we address the decisions of people worshipping, and after that, the assumed decisions of the gods who determine how much they would give back to the worshippers, as a function of their internal costs for doing so. Finally, we will look into the patterns that one may expect in the assumed behavior of the gods, first based on the regularity patterns in human performance, and then from the perspective of a random growth model.

¹⁹Although there may be in choosing the same religion, in accordance with the club goods approach.

²⁰We will nevertheless first consider the case of equally shared benefits and then generalize it.

²¹It is generally accepted that the powers or domains of the gods are known to be far from fixed, they depend on the context, type of request, also on location. We will in a later section develop a model that accounts for some evolution in the assumptions about the gods. We will nevertheless assume that, locally, the common assumptions about the gods are stable enough.

2.1 People's Devotion Decisions

Focusing on people's decisions, we begin by holding the total amount of *charis* provided by each god as a constant in order to determine a first equilibrium; we will later allow it to vary when we consider the god's (assumed) decisions. We assume there are $K + 1$ gods that may be worshipped. We consider a population of n people, or groups of people. For our purposes we do not distinguish between individuals and groups of individuals. People expect to obtain a share of the *charis* from participating in the ritual pertaining to a specific god. So when people participate in ritual for god $k \in [0..K]$, there is a total volume χ_k of *charis*, for now considered a given.

When several people choose to worship the same god at the same time, we will first assume that the total benefit given back by the god is shared evenly. If we write h_k the number of people choosing god k , then each receives $\frac{\chi_k}{h_k}$. In our framework, the only metric that people consider in the practice of ritual is the amount of *charis* they expect to receive.

Each person (or group of persons) $i \in [1..n]$ is assumed to be risk neutral and seeks to maximize the share of *charis* they receive. As a result, and since we do not make any particular preference assumptions that would distinguish one person from another, the situation is represented by a symmetric game with perfect information. In particular, the problem of optimally accessing a resource shared in common is well known, a form of "El Farol Bar" or "Kolkata Paise" problem, as discussed in Arthur (1994) or Chakrabarti et al. (2009). We have the following result²².

Proposition 2.1 (Optimal Worship Strategy). *The Nash equilibrium in the game for the worshippers is a symmetric mixed strategy with a probability distribution $(\pi_k)_{k \in [0..K]}$ across gods, and $\forall k \in [0..K]$:*

$$\lim_{n \rightarrow \infty} \pi_k = \frac{\chi_k}{\sum_{i=0}^K \chi_i}.$$

Proof. We assume a symmetrical Nash equilibrium in the game in the form of a mixed strategy: all players draw their choice for the god to choose from the same distribution $(\pi_k)_{k \in [0..K]}$ where $\pi_k > 0$. The equilibrium is characterized by the fact that the expected gain for any player is independent from the specific choice they make.

Set e_k the expected share in *charis* received from choosing god k when all others follow the optimal strategy, and can hence be expressed as a function of π_k . The number of other people choosing the same god follows a binomial distribution and:

$$e_k = \sum_{i=0}^{n-1} C_i^{n-1} \pi_k^i (1 - \pi_k)^{n-1-i} \frac{\chi_k}{i+1}.$$

²²The treatment of this model is an expansion of Gauthier (2020a), p. 438-443.

This quantity is known, see Chao and Strawderman (1972) for example, and simplifies to:

$$e_k = \chi_k \frac{1 - (1 - \pi_k)^n}{n\Pi_k}.$$

Since the Nash equilibrium is characterized by $e_k = e_0$ for all k , we have:

$$\frac{1 - (1 - \pi_0)^n}{n\Pi_0} = \frac{\chi_k}{\chi_0} \frac{1 - (1 - \pi_k)^n}{n\Pi_k}.$$

As n becomes large, both $(1 - \pi_0)^n$ and $(1 - \pi_k)^n$ go to zero and the equation becomes:

$$\pi_k = \frac{\chi_k}{\chi_0} \pi_0.$$

By definition $\sum_{k=0}^K \pi_k = 1$, so $1 = \pi_0 + \sum_{k=1}^K \pi_k$ and hence $1 = \pi_0 + \sum_{k=1}^K \frac{\chi_k}{\chi_0} \pi_0$. Solving for π_0 we obtain after simplification:

$$\pi_0 = \frac{\chi_0}{\sum_{i=0}^K \chi_i}.$$

Similarly, we also obtain that $\pi_k = \frac{\chi_k}{\sum_{i=0}^K \chi_i}$. □

Note that scaling the share received up or down by the same factor has no impact on the optimal probabilities. We can see that the higher the total amount of *charis* a god can provide, the greater the optimal probability of choosing that god for worshipping, although the optimal strategy randomizes across all the possible choices, and even a god who would offer very little *charis* sees a non zero probability of seeing some devotion acts. All gods hence see some worship all the time.

It is worth stressing that if the benefit of participating in worship was not shared but was simply unaffected by the choice of other people, or even was increased by other people participating, then there would be no mixing and the same god with the highest benefit would be worshipped all the time. Then in a polytheist context, it appears that by construction there needs to be a declining benefit in participating in the same rituals if one is to explain the multiplicity and spread of god choices. One could argue that the gods offered some specialization, and therefore there could be a need for a diversity of gods; but then why would not some of them be merged together? As we will see in the framework of the random growth model for the observed benefits associated to each god developed in a later section, the multiplicity of gods may be maintained simply due to chance.

As we have stressed, we made a particularly stringent assumption in that benefits would be shared equally among participants. We know this was likely not considered to be the case by the ancient Greeks. Even the most democratic polities had strong hierarchies, and honors during

ritual were concentrated on a few participants. It makes sense, therefore, to assume that some people would collect a greater share of *charis* than others, relative to their number. In order to capture this situation, let us assume that there is a proportion r of privileged citizens²³ of the *polis*. These people do not receive their fair equal of *charis*, but a multiple $\rho > 1$ of it. We write $(\pi_k^c)_{k \in [0..K]}$ for the probability that a normal citizen would choose a god, and $(\pi_k^h)_{k \in [0..K]}$ for that probability for a privileged citizen.

With these assumptions, we obtain the following proposition.

Proposition 2.2 (Optimal Worship Strategy with Unequal Shares). *The Nash equilibria in the game for the worshippers with unequal benefits consist of all mixed strategies where the probability distributions $(\pi_k^c)_{k \in [0..K]}$ and $(\pi_k^h)_{k \in [0..K]}$ verify at the limit as n goes to ∞ , $\forall k \in [0..K]$:*

$$\frac{1-r}{1-r+r\rho}\pi_j^c + \frac{r\rho}{1-r+r\rho}\pi_j^h = \frac{\chi_k}{\sum_{i=0}^K \chi_i}.$$

In addition, for each $k \in [0..K]$, if either π_k^c or π_k^h is equal to $\frac{\chi_k}{\sum_{i=0}^K \chi_i}$, then the other is also equal to that value.

Proof. As the number of people increases, we can consider directly as in the proof of Proposition 2.1 that the number of worshippers who choose a god with probability p behaves almost surely as np . With our assumptions, the total number of people weighted by their shares of *charis* worshipping god k is $n \left((1-r)\pi_k^c + r\rho\pi_k^h \right)$.

The expected gains from worshipping god k for the members of each class hence write:

$$e_k^c = \frac{\chi_k}{n \left((1-r)\pi_k^c + r\rho\pi_k^h \right)}$$

$$e_k^h = \frac{\rho\chi_k}{n \left((1-r)\pi_k^c + r\rho\pi_k^h \right)}.$$

Since all the high-class participants are identical, and all the normal citizens are identical, the equilibrium conditions are that their gains would not be dependent on their choice of god, and therefore $e_i^c = e_j^c$ and $e_i^h = e_j^h$ for i and j in $[0..K]$. These conditions result in the same thing:

$$\frac{\chi_i}{(1-r)\pi_i^c + r\rho\pi_i^h} = \frac{\chi_j}{(1-r)\pi_j^c + r\rho\pi_j^h}.$$

From this, we obtain by summing over i , and using the fact that $\sum_{i=0}^K \pi_i^c = \sum_{i=0}^K \pi_i^h = 1$:

$$\frac{(1-r)\pi_j^c + r\rho\pi_j^h}{(1-r) + r\rho} = \frac{\chi_j}{\sum_{i=0}^K \chi_i},$$

²³Who may be kings, aristocrats, hierarchs, or any such relevant category.

which is the first statement in the proposition. Set $a = \frac{1-r}{1-r+r\rho}$ for simplicity, then if we have $\pi_j^c = \frac{\chi_j}{\sum_{i=0}^K \chi_i}$, we obtain:

$$\pi_j^h = \frac{1}{1-a} \frac{\chi_j}{\sum_{i=0}^K \chi_i} - \frac{a}{1-a} \pi_j^c = \frac{\chi_j}{\sum_{i=0}^K \chi_i}.$$

The same calculation shows that the reciprocal relationship is also true. □

We can see that with these more general assumptions on the manner in which *charis* is shared, we obtain comparable results as with the case when it is shared equally.

In our analysis so far we have considered a constant number of gods $K + 1$, but this is a strong simplification of the actual dynamics of any *polis*-level pantheon. In the ancient Greek polytheist context, foreign gods were not forbidden and could be adopted and adapted. While in the Roman Republic and Empire the adoption of new gods or rituals was regulated²⁴, this was not the case in the classical Greek world. However, there is a minimal size for any cult to be carried out, so that enough people can participate in the sacrifices and libations; there is also a minimal group size for the justification for shrines and temples to be built. Hence, we could expect that gods with extremely few followers could not give rise to votive acts simply because the organization of the rituals would be problematic. As a result of this threshold, we would expect to see a finite number of gods, the maximum number being a function of the *polis*'s resources.

2.2 Optimal Divine Strategy

So far we have considered that the total amount of *charis* χ_k was a given for every god. Now, we reflect the fact that it is provided by the gods at some internal cost, and they are assumed to choose the optimal amount. We assume that each god wants to maximize the proportion of people who come and worship them. Classical Greek literature, such as epics or tragedy, is filled with examples of the gods taking punitive actions against mortals when they did not properly respect them. As we pointed out earlier, we know too that the gods were assumed to enjoy some of the output from sacrifices, which would be all the more important, the more people partake in it. We accounted earlier for the fact that certain persons would receive a greater share of *charis*; it also appears logical to consider that from the standpoint of the gods, they count for more, so that it is the properly weighted number of people that is relevant.

We assume that the the gods are held to be risk neutral, and know how the people will make their optimal choices, and seek to maximize their gains in the form of the weighted probability

²⁴The scandal of the Bacchanalia in 186 BC, as related in *Livy* 39.15, led the Roman Senate to legislate specifically to address problems caused by Dionisiac rituals, as was confirmed through epigraphic sources.

of being chosen²⁵ minus a proportional cost attached to producing the *charis*. The gain for god k , as a function of the *charis* χ_k can therefore be written $g_k(\chi_k) = \pi_k(\chi_k) - c_k\chi_k$, where c_k is the exogenously given proportional cost. All gods will seek to maximize their gains at the same time, reflecting their knowledge of each other's strategy.

The following proposition characterizes the endogenously determined *charis*, the probability of being chosen and expected gain at the equilibrium where all gods and worshippers act optimally.

Proposition 2.3 (Optimal Divine Strategy). *The optimal strategy for all gods is for the charis they offer to verify for all k in $[0..K]$ and with $M = \{i : c_i = \max_{k \in [0..K]} \{c_k\}\}$ and $\gamma_k = c_M - c_k$, such that for $k \neq M$, $\gamma_k > 0$ and $\gamma_M = 0$:*

$$\chi_k^* = \frac{\gamma_k(K+1) + \sum_{j=0}^K \gamma_j}{4(K+1) \left(\sum_{j=0}^K \gamma_j \right)^2}.$$

At this optimum, the probability of being chosen is

$$\pi_k^* = \frac{\gamma_k}{2 \sum_{j=0}^K \gamma_j} + \frac{1}{2(K+1)},$$

and the gain for the gods is

$$g_k^* = \pi_k^* \left(1 - \frac{c_k}{2 \sum_{j=0}^K \gamma_j} \right).$$

Proof. Since the gods are measuring their attractiveness in weighted terms $(1-r)\pi_k^c + r\rho\pi_k^h$, then according to Propositions 2.1 and 2.2, we know that this is equal to $\pi_k = \frac{\chi_k}{\sum_{i=0}^K \chi_i}$. There is effectively no difference with the case where benefits are equally shared.

We first note that:

$$\frac{\partial \pi_k}{\partial \chi_k}(\chi_k) = \frac{\sum_{i=0}^K \chi_i - \chi_k}{\left(\sum_{i=0}^K \chi_i \right)^2}.$$

Further, $\sum_{i=0}^K \chi_i - \chi_k$ does not depend on χ_k so that $\frac{\partial \pi_k}{\partial \chi_k}(\chi_k)$ is decreasing, and the second-order condition is verified for the existence of a maximum, as $\frac{\partial^2 \pi_k}{\partial \chi_k^2}(\chi_k) = -\frac{2 \left(\sum_{i=0}^K \chi_i - \chi_k \right)}{\left(\sum_{i=0}^K \chi_i \right)^2}$.

The first order conditions for each god k 's gain optimization writes: $\frac{\partial \pi_k}{\partial \chi_k}(\chi_k) - c_k = 0$; which based on Proposition 2.1 simplifies into:

$$\sum_{i=0}^K \chi_i - \chi_k = c_k \left(\sum_{i=0}^K \chi_i \right)^2.$$

Since the costs c_k are arbitrary, without loss of generality we choose to rank them so that c_0 is

²⁵They reason as a function of the population size: we are implicitly assuming that the citizens of a *polis* reason from the perspective of that *polis*, not accounting for the fact that other *poleis* may have different sizes and compete for the gods' attention.

the highest and c_K the lowest. Taking differences for all l and k in $[0..K]$ we hence have

$$\chi_k - \chi_l = (c_l - c_k) \left(\sum_{i=0}^K \chi_i \right)^2,$$

and in particular, at the optimum,

$$\chi_k = (c_0 - c_k) \left(\sum_{i=0}^K \chi_i \right)^2 + \chi_0.$$

Using that expression we can write:

$$\begin{aligned} \sum_{j=0}^K \chi_j &= \sum_{j=0}^K \left((c_0 - c_j) \left(\sum_{i=0}^K \chi_i \right)^2 + \chi_0 \right) \\ &= (K+1)\chi_0 + \left(\sum_{i=0}^K \chi_i \right)^2 \sum_{j=0}^K (c_0 - c_j). \end{aligned}$$

Hence, taking the square and writing $X = \left(\sum_{i=0}^K \chi_i \right)^2$ we obtain:

$$X = (K+1)^2 \chi_0^2 + 2(K+1)\chi_0 \sum_{j=0}^K (c_0 - c_j) X + \left(\sum_{i=0}^K (c_0 - c_j) \right)^2 X^2.$$

Since all the (χ_k) are determined up to a scalar, god 0 can chose χ_0 arbitrarily: there is a continuum of possible optimal equilibria depending on χ_0 . We consider that they always choose the highest χ_0 , that is the highest volume of *charis* compatible with the existence of optimal solutions. Hence χ_0 is set so that the determinant of the equation is null, that is:

$$1 - 4(K+1)\chi_0 \sum_{j=0}^K (c_0 - c_j) = 0.$$

In that case we obtain the solution

$$X = \left(\sum_{i=0}^K \chi_i \right)^2 = \frac{1}{4 \left(\sum_{j=0}^K (c_0 - c_j) \right)^2}.$$

Replacing in the equation for the optimality we can write after some simplifications:

$$\chi_k^* = \frac{(c_0 - c_k)(K+1) + \sum_{j=0}^K (c_0 - c_j)}{4(K+1) \left(\sum_{j=0}^K (c_0 - c_j) \right)^2},$$

and

$$\frac{\chi_k^*}{\chi_0^*} = \frac{(c_0 - c_k)(K+1) + \sum_{j=0}^K (c_0 - c_j)}{\sum_{j=0}^K (c_0 - c_j)}.$$

Replacing into the expression for π_k , we obtain at the optimum

$$\pi_k^* = \frac{c_0 - c_k}{2 \sum_{j=0}^K (c_0 - c_j)} + \frac{1}{2(K+1)}.$$

Finally, the optimal expected gain for gods writes:

$$\begin{aligned} g_k^* &= \pi_k^* - c_k \chi_k^* \\ &= \frac{c_0 - c_k}{2 \sum_{j=0}^K (c_0 - c_j)} + \frac{1}{2(K+1)} - c_k \frac{(c_0 - c_k)(K+1) + \sum_{j=0}^K (c_0 - c_j)}{4(K+1) \left(\sum_{j=0}^K (c_0 - c_j) \right)^2} \\ &= \left(\frac{c_0 - c_k}{2 \sum_{j=0}^K (c_0 - c_j)} + \frac{1}{2(K+1)} \right) \left(1 - \frac{c_k}{2 \sum_{j=0}^K (c_0 - c_j)} \right) \\ &= \pi_k^* \left(1 - \frac{c_k}{2 \sum_{j=0}^K (c_0 - c_j)} \right). \end{aligned}$$

We have specified the ranking of the gods as a function of costs (c_k) with the only effective requirement that c_0 should be the largest. As a result, all the expressions at the optimum can be written by replacing the index 0 with $M = \{i : c_i = \max_{k \in [0..K]} \{c_k\}\}$, so that $c_M = \max_{k \in [0..K]} \{c_k\}$. Setting $\gamma_k = c_M - c_k$ we obtain the expressions in the proposition. \square

The following consequences are worth stressing:

- The expected gain for worshippers goes to zero as n goes to infinity. Since $\frac{1-(1-\pi_k)^n}{n\pi_k} = \frac{1}{n\pi_k} + o(n)$, then the equilibrium gain expectation for the worshippers is

$$\bar{e} = \frac{1}{2n \sum_{j=0}^K \gamma_j} + o(n).$$

One could make the costs all converge to zero, arguing that they are uniformly reduced as a function of the total population for all gods, so that γ_k is instead written $\frac{\gamma_k}{n}$; in this case, the worshippers' expected gain is not null at the limit.

- The derivative of the optimal probability relative to efficiency is always positive: for all $k \neq M$,

$$\frac{\partial \pi_k^*}{\partial \gamma_k} = \frac{\sum_{j=0}^K \gamma_j - \gamma_k}{2 \left(\sum_{j=0}^K \gamma_j \right)^2} > 0.$$

Also, the derivative of the optimal probability for one god relative to the other gods' efficiency is always negative, for $l \neq k$ and both different from M :

$$\frac{\partial \pi_k^*}{\partial \gamma_l} = -\frac{\gamma_k}{2 \left(\sum_{j=0}^K \gamma_j \right)^2},$$

and one can verify that $\sum_{k=0}^K \sum_{i=0}^K \frac{\partial \pi_k^*}{\partial c_i} = 0$.

- The optimal probability π_M^* , which is the lowest among all the π_k^* , is not affected by changes in the efficiency $(\gamma_k)_{k \in [0..K]}$ and is constant with $\pi_M^* = \frac{1}{2(K+1)}$. The model therefore predicts a minimal probability of a god being worshipped that only depends on the total number of gods. The optimal probabilities can also be understood as worshippers choosing, with a probability of $\frac{1}{2}$, a god randomly with a uniform probability across all the gods, and with a probability of $\frac{1}{2}$, a god k following a probability of $\frac{\gamma_k}{\sum_{j=0}^K \gamma_j}$.
- The simple expression $\gamma_k = c_M - c_k$, and the contribution of $\frac{\gamma_k}{\sum_{j=0}^K \gamma_j}$ in the expression for the amount of worship look interestingly like the mathematical expressions of the structuralist angle that Vernant followed in looking at Greek religion. Indeed, the full quote from Vernant²⁶, with added emphasis, is that:

“A god is a power that represents a type of action, a kind of force. Within the framework of a pantheon, each of these powers is *defined not in itself* as an isolated object but by virtue of its *relative position in the aggregate of forces*, by the structure of relations that oppose and unite it to the other powers that constitute the divine universe.”

Since the gods are believed to place themselves on a unidimensional space though the choice of *charis* they give back to worshippers, we may relate this model to the approach to competition initially developed by Hotelling (1929). However, the gods’ efficiency, which we consider as a given at this stage, condition their ability to distinguish and define themselves by their effort in giving *charis*. As a result, the optimal position is driven by their characteristics rather than by strategy, in which the outcome differs from the competitive setting envisioned by Hotelling.

At this stage, with our assumptions, we have established a relationship between differences in the god’s efficiency at carrying out acts in favor of their worshippers, and the probability that they are worshipped. We will now turn to two different approaches that attempt to explain the distribution of votive acts that one should expect to observe.

2.3 Human Drivers of Divine Performance

Why would the gods be expected to possess different degrees of ability in offering *charis* to humans? It seems reasonable to consider that could be the case because the gods are assumed by people to be to some extent in competition between each other, and as a result would have acquired different levels of expertise. The ancient Greeks were familiar with competition between each other, especially in sporting events. In fact, exceptional winners at sporting events were

²⁶Versnel (2011), p. 27.

revered through poetry, in particular in Pindar's *epinicia* (odes to victors and victory). By winning and displaying rare skill or strength, they differentiated themselves from the rest, and were often qualified as divine²⁷. In the *Iliad*, the heroes, whose prowess at war was extraordinary, are also commonly qualified as divine. The literary perspective clearly establishes a link between extreme performance and divinity. Hence, it is not inconceivable that, since the attributes of extreme performance are associated to gods, then the attributes of gods be associated to extreme performance.

Human physical and cognitive performance has largely been documented to follow power law distributions. Donner and Hardy (2015) show how power laws, or combinations thereof, represent human learning curves for a variety of tasks. Record running times as a function of distance have been shown to be well captured by power laws as well, as discussed by Vandewalle (2018). Of greater importance for our purposes, the rankings of human performance have been documented to follow various types of power laws. For example, Smith (2015) documents the presence of power laws in several psycho-motor actions. More specifically for sports, Deng et al. (2012) showed that across a large set of sports (tennis, soccer, snooker, and many others), performance as a function of ranking followed power laws with exponential decays. Focusing on fluctuations of rankings in sports and games over time, Morales et al. (2016) showed that power laws with an exponential decay also accounted well for rankings in sports and competitive games such as poker. A generative model for a power law with exponential decay was proposed in Baek, Bernhardsson, and Minnhagen (2011), based on random group formation.

Note that these considerations would exclude the case where the assumed characteristics of the gods would be distributed with some variations around an average. The shape of the rankings obtained by drawing outcomes from a Gaussian distribution, for example, and sorting them, would be differ from what one gets with power laws or variations thereof.

Power laws have been applied to ranking modeling in many areas of economics and for most human activities. Gabaix (2016) offers an overview in economics: the most salient examples include the distribution of city sizes, firm sizes, stock market shocks, or individual wealth. Focusing on the ways in which the emergence of power laws may be accounted for through various theoretical arguments, Mitzenmacher (2004) lists a large section of human experience where rankings have been documented to follow these distributions, including the size of computer files, number of connections in networks, or word frequencies in texts. In particular, market shares in very diverse sectors (from cereals to rifles) show noticeable regularity and are better represented by power laws than by exponential relationships, as discussed in Kohli and Sah (2006)²⁸.

²⁷See Miller (2018).

²⁸One could directly consider the choices of gods to interact with as a similar situation to market shares for

Let us rank the gods by their efficiency $(\gamma_k)_{k \in [1..K]}$ with $\gamma_1 = \max\{\gamma_j : j \in [0..K]\}$ setting by convention $\gamma_0 = \gamma_M = 0$. We consider that people give them anthropomorphic qualities reflective of competitive activities. The terms $\frac{\gamma_k}{\sum_{j=0}^K \gamma_j}$ represent the gods' efficiency scaled to their total efficiency, which, thanks to this scaling, can be understood as a dimensionless measure, akin to a scaled number of points in a competitive sport or game. Hence, following Deng et al. (2012) or Morales et al. (2016), we would expect for all $k \in [1..K]$ that the scaled efficiency follows a power law with exponential decay²⁹, and to have an expression of the form

$$\frac{\gamma_k}{\sum_{j=0}^K \gamma_j} = ak^{-\alpha} e^{-\lambda k},$$

for some positive numbers a , α and λ . The exponential decay appears to be a necessary feature in order to account for a clear concave shape in log-log plots of competitive sports ranking data. In the analysis of sports rankings, the power exponent α has been associated with the degree of competitiveness. Stated simply, the steeper the drop in gains as a function of rank, the more the activity is deemed to be competitive.

With this definition, the probabilities with which people are expected to worship gods become, for $k \in [1..K]$:

$$\pi_k^* = ak^{-\alpha} e^{-\lambda k} + \frac{1}{2(K+1)}.$$

The differences $\pi_k^* - \pi_0^*$ would therefore be expected to follow power laws. The distribution for π_k^* itself is therefore a mixed power law with a uniform distribution.

2.4 Random Growth and Divine Performance

An alternative approach to explaining the shape of the assumed god efficiency distribution is to account for how it evolved. We can indeed formalize the distribution π_k by considering it would be based on an evolution of beliefs over time. Since people ask the gods for various things, we will model the fact that some of the wishes are granted (through sheer randomness), and that based on this observation people will update their beliefs on the gods' efficiency. In effect, we will consider that the observed benefits from worshipping each god follow a random growth model³⁰.

We have considered so far that the efficiency of the gods formed a uni dimensional continuum: their specific jurisdiction did not enter the analysis and only their efficiency at giving *charis* back was important. Now, while we keep the same broad logic, we will also consider that people

various brands. This could naturally lead us to expect some form of power law distribution for votive acts; however this would not fundamentally explain why it should be the case.

²⁹Note that a power law with exponential decay is also called a truncated power law, since the exponential term tapers off the distribution's tail.

³⁰See Gabaix (2016).

who sacrifice to a particular god will have framed their requests in the context of that god's specific skills. Hence, two people asking for the same fundamental thing of two separate gods would phrase it in different terms; as a result they would not measure whether their requests were satisfied in the same manner going forward. Random life events hence affect the people who worshipped different gods in different manners. We will also consider there are systemic life events affecting everybody the same way. The more people had started worshipping a god, the more a positive outcome of life events affecting these worshippers would be noticed. In effect, the total amount of perceived positive outcomes for a given god at some future time is the share of population initially worshipping this god compounded by all the positive outcomes through time.

Let us assume that worship is organized at discrete times $t \in \mathbb{N}$. At time $t = 0$, the gods' efficiency $(\gamma_{k,0})_{k \in [1..K+1]}$ are allocated in some unspecified manner, leading to optimal probabilities $\pi_{k,0}^*$ with a minimum $\pi_{M_0,0}^*$. We consider that the strategies that we described earlier take place without participants planning through time, their horizon is always just that of the current time. At each round, however, the collective assumptions about the pantheon are informed by some events. The events are all fundamentally random, but they are partially colored by people's perception as a function of the gods they most recently worshipped.

For the $K + 1$ gods, systemic life events are represented across the board through time by i.i.d. random variables $S_t \stackrel{(d)}{=} S$ while the perception of god-specific events are $G_{k,t} \stackrel{(d)}{=} G$, for some distributions G and S , where G is positive. These variables capture what people measure (in any unit they deem appropriate) as beneficial life events that they associate to each god. We define a measure of positive life events $p_{k,t}$ associated to each god k at each period t . The systemic events S_t represent things that can not be associated to any particular god specifically and are therefore spread out, while $G_{k,t}$ are associated to each god and therefore compounded by the current level of cumulative positive life events associated to the god. We therefore express the dynamics of that measure as a function of $G_{k,t}$ and S_t as follows:

$$p_{k,t+1} = G_{k,t} p_{k,t} + S_t.$$

At the first period, not knowing anything yet, the "goodness" of each god is simply its relative efficiency, which is equal to the share of people worshipping it (in excess of the least favored god). Hence

$$p_{k,0} = \left(\pi_{k,0}^* - \pi_{M_0,0}^* \right) = \frac{\gamma_{k,0}}{2 \sum_{j \in [1..K+1]} \gamma_{j,0}}.$$

Observing how everybody's life is affected as a function of which gods they have worshipped, people set the assumed divine efficiency accordingly, so that simply $\gamma_{k,t} = p_{k,t}$ at each period³¹.

³¹Since the distributions G and S are quite arbitrary, they can in fact reflect in any particular relationship one

We know that the probabilities $\pi_{k,t}^*$ are then set by normalizing $(\gamma_{k,t})_{k \in [1..K+1]}$ at each step. After time has advanced and the random distribution of the $p_{k,t}$ for each k converges, we consider that the perceptions of the gods' efficiency have been set, the local myths have been expressed and hence these perceptions should be held constant going forward.

With these assumptions, we have the following result.

Proposition 2.4 (Convergence to Power Law). *Assume there exists α such that $\mathbb{E}[G^\alpha] = 1$, $\mathbb{E}[G^\alpha(\ln(G) \vee 0)] < \infty$, and $0 < \mathbb{E}[|S|^\alpha] < \infty$. Also assume that $\frac{S}{1-G}$ is not degenerate and there exists no λ such that the support of $\ln(G)$ conditioned on $G \neq 0$ would be included in $\lambda\mathbb{Z}$. Then:*

- i. As t goes to infinity, for each $k \in [1..K_1]$, $\lim_{t \rightarrow \infty} p_{k,t} \stackrel{(p)}{=} \Upsilon$, and Υ asymptotically follows a power law with parameter α .*
- ii. As a result, as t goes to infinity, for each $k \in [1..K + 1]$, $(\pi_{k,t}^* - \pi_{M_t,t}^*)$ converges in probability towards a normalized random draw from an approximate power law of coefficient α ³².*

Proof. A general description of generative approaches to power laws, with applications to economics, can be found in Gabaix (2008), and we follow the outline of random growth models presented in that paper.

We use Kesten's results: Theorem B, p. 210 as well as the particular application of Theorem 4, p. 235 in Kesten (1973). The simpler presentation in Goldie (1991), Theorem 2.3, p. 129 and Theorem 4.1 p. 135, makes the theorem's conditions more explicit. The conditions of Proposition 2.4 are particular cases of the theorem's condition for positive random variables G and S . The general result is that the random growth recurrence equation has a steady-state solution and converges towards a distribution that is asymptotically a power law.

According to the theorem, as t goes to infinity, for each $k \in [1..K + 1]$, $p_{k,t}$ converges in probability to the solution Υ of $\Upsilon \stackrel{(d)}{=} G\Upsilon + S$. In addition, this random variable Υ verifies particular asymptotic conditions. There exists $a > 0$ such that $\lim_{x \rightarrow \infty} x^\alpha \mathbb{P}[\Upsilon > x] = a$. This effectively states that the tail of the distribution of Υ follows a power law of parameter α , that is corresponding to a power law of density $\mathbb{P}[\Upsilon \in dx] = \left(\frac{\alpha}{l_0}\right) \left(\frac{x}{l_0}\right)^{-(1+\alpha)} dx$, if it followed that law entirely and not just on the tail, where l_0 is the strictly positive minimum value the random variable can have.

At the limit when t goes to infinity, each $p_{k,t}$ for all k follows the same distribution with a power law tail. As a result, at the limit, the efficiency $\gamma_{k,t} = p_{k,t}$ for each k is drawn from that

would want to specify between γ and p .

³²Note that the coefficient is meant as the exponent of the counter cumulative distribution, not of the density.

distribution Υ . The terms

$$\pi_{k,t}^* - \pi_{M_0,t}^* = \frac{\gamma_{k,t}}{2 \sum_{j \in [1..K+1]} \gamma_{j,t}}$$

therefore simply represent K normalized draws from an approximate power law of parameter α . The probability distribution $(\pi_{k,t}^*)_k$ that is used by people to determine their random choice of which god to worship is therefore itself drawn from the approximate power law. This is the second statement in the Proposition. \square

Note that the proposition does not state that the distribution that people choose to determine which god to worship is a power law, but rather that some numbers are drawn from a power law, then normalized, and then the people randomly choose a god according to that normalized distribution.

Proposition 2.1 hence tells us that given any starting conditions, and given a wide range of potential life events distributions affecting people's perception of what the gods pay them as *charis*, we may empirically observe that votive acts appear to be distributed across gods following simple power laws.

Recall that a power law random variable Υ with minimum value l_0 and parameter α has a distribution with density $\mathbb{P}[\Upsilon \in dx] = \frac{\alpha}{l_0} \left(\frac{x}{l_0}\right)^{-(1+\alpha)} dx$. The expression for the density of a normalized power law may be written in the case of two independent variables³³. However, as the number of independent variables in the normalization increases, the shape of the distribution changes drastically. Unfortunately we only know that the distribution of a sum of many power laws only shows regularity in the tail, and there are no general expressions for it³⁴. There is no simple way of expressing the distribution of the normalized $\pi_k - \pi_M$ at the limit.

A power law can nevertheless be simulated. If U is uniform over $[0, 1]$, then $l_0 U^{\frac{-1}{\alpha}}$ follows a power law of parameter α . In order to assess the nature of the distribution that may be empirically observed if the dynamics described above held, we simulated outcomes of the distribution of Υ in a number of *poleis* (one draw per god per *polis*), for a random number of gods in each *polis*, from which for each *polis* we computed the normalized $\pi_k - \pi_M$. Ranking the outcomes in each *polis*, we can directly obtain the log/log plot of probabilities relative to ranks. As Figure 1 illustrates, while there is a degree of dispersion, the curves on this scale would seem to point towards a

³³If Υ_1 and Υ_2 are independent and follow power laws of parameter α , after some calculations one gets a symmetrical distribution:


$$\mathbb{P} \left[\frac{\Upsilon_1}{\Upsilon_1 + \Upsilon_2} \in dx \right] = \frac{\alpha}{2} \mathbb{1}_{x \in [0,1]} \frac{1}{x^2} \left(\frac{1}{x} - 1 \right)^{-(1-\alpha) + 2\alpha \mathbb{1}_{x > \frac{1}{2}}}.$$

³⁴See Ramsay (2006), who identified all the particular cases where such an expression is known. Some claim that power laws are stable by summation, but that is only true asymptotically in the tail of the distribution.

power law.

We also carried out formal maximum-likelihood and goodness-of-fit tests³⁵ for each simulated *polis*, and Table 1 shows the results. We can see that the data is consistent with power law or truncated power law distributions, generally at a significant level. In addition, the parameters are fairly stable across the simulated range, and we can see that the average falls close to the parameter that was used in the simulation. Hence, the *results* from sampling and normalizing the draws from a power law would appear convincingly enough to represent power law *distributions*.

Figure 1: *Log/log Plot of Outcomes Across Simulated Poleis*



Note: Ratio is the simulated normalized draw and Rank is its rank for each set of draws. The blue line is a regression line across all the data. The parameters are set as follows: $\alpha = 0.5$, K follows a uniform distribution over $[10, 25]$, $N = 100$ *poleis*.

3 Empirical Analysis of Veneration Acts

Relying on a wealth of data that has been painstakingly gathered by historians over the years, we now turn to the empirical analysis of the model developed above. We will not only directly test the model, but also look for potential explanatory factors behind the specific values driving the differences in divine efficiency.

We discuss our sources in the first subsection, where we explain how the epigraphic and literary

³⁵Following the methodology exposed in Clauset, Shalizi, and Newman (2009). In this instance and for other applications of this methodology we adjust the parameter α to our convention of referencing the counter cumulative distribution, for consistency.

Table 1: *Summary Statistics on Distribution Fits Across Simulated Outcomes*

Statistic	Mean	St. Dev.	Min	Pctl(25)	Median	Pctl(75)	Max
Lambda Exp	0.015	0.005	0.004	0.012	0.015	0.019	0.024
Alpha Pow	0.410	0.178	0.178	0.285	0.352	0.489	0.948
Alpha Trunc	0.298	0.342	0.0000002	0.00002	0.205	0.510	1.320
Lambda Trunc	0.001	0.001	0.000	0.0001	0.001	0.002	0.004
Trunc vs Pow R	2.478	3.151	-1.075	0.427	1.705	3.086	17.079
Trunc vs Pow p	0.215	0.236	0.00004	0.007	0.135	0.429	0.780
Trunc vs Exp R	3.755	4.256	-2.939	1.456	2.790	4.641	26.997
Trunc vs Exp p	0.117	0.238	0.000	0.000004	0.005	0.056	0.999
Pow vs Exp R	2.993	4.261	-5.532	0.445	2.309	4.717	25.376
Pow vs Exp p	0.186	0.299	0.000	0.000001	0.013	0.239	0.987

Note:

The distribution names are abbreviated as follows:

Exp = exponential, *Pow* = (pure) power, *Trunc* = power law with exponential decay.

R: ratio of goodness-of-fit; a positive number means that the first law of the two is preferred.

p: significance level; the probability that the preference would be due to randomness.

The simulation parameters are: $\alpha = 0.5$, K follows a uniform distribution over $[10, 25]$, $N = 100$ *poleis*

data on votive acts can be shaped for our analysis, and merged further with geographic data. Next, we show some of the patterns directly observable in the data and in particular focus on the number of gods worshipped in each *polis*. The third subsection consists of the statistical tests of the model, and in the last subsection we will examine the empirical drivers of the worship act distribution.

As we have stressed earlier, the ancient Greek world was not a unified polity, but a constellation of independent political organizations, the *poleis*. Even though all shared polytheism, they did not necessarily worship the same gods. As we look into worship practices, it is therefore not appropriate to commingle them across *poleis*, but rather to consider each *polis* separately: comparable patterns may be apparent across various *poleis*. In our approach of the data we will therefore always keep this general framework, of not considering Greece as a whole, but rather a set of separate entities. We are effectively considering the information on gods and worship to be structured at the polity level, rather than across the entire geographic space of ancient Greece.

3.1 Epigraphic and Literary Sources Combined with *Polis*-level Data

The empirical analysis of ancient Greek religious practice is made difficult by the relative paucity of historical and material sources. There are no systematic records of religious acts, and the historical texts did not consistently keep track of the data we would need to test the models developed above. One reasonable approach is to focus on epigraphic sources; while they do not represent every single worship act, they can constitute a reasonably consistent set of observations.

Epigraphic inscriptions related to religious offerings or dedications typically contained the names of one or several gods along with some epicleses and reasons for the inscription. The centralization and analysis of such inscriptions have been carried out for over a century, starting with thousands of paper files, and recently evolved into the systematic storage of complex onomastic networks. Bonnet and Lebreton (2019) discuss the historiography of divine onomastics, with many details

on how methods have evolved over time. Since many of the sources are epigraphic, the fact that they are available or not today is driven by a conjunction of factors, including earthquakes, fires, wars, voluntary destructions, chance findings, or support nature and quality to name a few. Many of these factors may create biases, but we have no way of determining them.

For our purposes, the data from the Base de Donnée des Epiclèses Grecques (BDEG), by Lebreton et al. (2014), described in Brulé and Lebreton (2007) can be directly exploited. It contains the information from thousands of epigraphic inscriptions or literary references to votive acts³⁶, including location, the god or gods names and the nature of the inscription.

The BDEG started from the idea to allow researchers to study Greek polytheism and Greek religion in a more systematic and quantitative manner than what had been possible up to that point. More specifically, researchers gathered and made available data about Greek divine epithets from all sources, periods and region they may come from. The general notion of god name was understood as a system combining the actual name of the god, typically in the first place, with some additional qualification (commonly an adjective) into a single nominal group. The additional qualification, named cult-epithet, describes a particular feature of the god. The nominal group thus stored into the database each time refers to a precise divine entity worshipped by ancient Greeks.

The data has been entered in over 11,000 forms, with in principle one for each observation of a divinity and epithets. One significant issue affected the BDEG database: the data was initially entered as one form for all the observations of a particular epithet in a location, but later evolved so that each entry represented a separate observation. However, the forms also contain categorical data specifying a range for the number of observations of the specific divinity and epithets. In some cases, the divinities were referred to in a generic fashion, such as Θεοί, Θεός or Θεά (Gods, Goddess or God). These cases are moderately common, and we treat them as a particular God name, because they are usually associated with particular epithets, and reflect a specific choice on behalf of the worshipper.

For each inscription or group of inscriptions, unstructured data generally captures some details on the source and context:

- The source itself may be literary, epigraphic, papyrological or numismatic.
- Based on the available information, such as the archaeological context, a dating range has sometimes been made available as a period or a century.

³⁶The data also contains divine names associated to epicleses that could have appeared in the sources about a sanctuary or a celebration and hence do not match a precise votive act. We nevertheless keep these datapoints because they can account for the votive act associated to the sanctuary's dedication or to the celebration in question.

- The presence of a cult is an important qualification; some mentions in the inscription of a sacrifice, a priest, a sanctuary or just a dedication confirmed that there was a cult to the deity.

Location information is provided in some cases with coordinates, and in other cases only with the gallicized Greek name. The names were manually mapped to the corresponding anglicized Greek names. Unfortunately, in many cases neither geographic coordinates nor mappable names are provided, and we cannot properly associate these observations to a *polis* or a location.

The nature of the data is such that we cannot precisely date the observations. They need to be considered as a whole, spanning a long time period through Antiquity. In the context of the model we developed in the prior section, this effectively corresponds to many effective choices of which gods to worship cumulatively through time.

As we have alluded to before, the names of divinities are not something that is trivially defined. The BDGE tracks in its entries the epicleses associated with the gods, which implicitly defines the “name” of the god as what comes before the sequence of qualifications. In some cases, the name of the god itself is considered to include a qualification, such as *Apollo Heraklea*, for example. We keep that same logic, considering that this manner of addressing the god effectively defines it as a separate entity from simply *Apollo*. Further, gods are sometimes addressed as a list: for example “Zeus and Hera”. In these instances, we split this as two entries into each one of the separate gods. Indeed, we are implicitly assuming that in the cases where rituals were carried out for more than one god at the same time, each god would have obtained their due share of these rituals.

After the processing operations we described above, the underlying BDEG data can be represented as a series of vectors

$$\Delta_i = (\delta_i^{polis}, \delta_i^{NbObs}, \delta_i^{deity}, \dots, \delta_i^N),$$

each corresponding to one of the I entries in the BDEG, with $i \in [1..I]$ and N the number of numeric or categorical characteristics available. The set of unique god names observed across all the data is D , and we will write D_p for the set of unique god names observed for *polis* p . In the dataset, the number of gods in each *polis*, that is $|D_p|$, is called “NumGods”. The set of all *poleis* on which data is available is Π .

With each entry Δ_i , δ_i^{NbObs} represents the number of occurrences of the same votive inscription mentioned in the form. The data does not in fact precisely say how many observations were made, but specifies a range such as “2 to 5”, “>10”, etc. We convert these categories to numerical approximations. About 2,500 of the 11,000 entries in the data are given such a range.

We write the total number of acts for a given *polis* as $T_p = \sum_{\{i:\delta_i^{polis}=p\}} \delta_i^{NbObs}$. From this, we derive an empirical distribution of votive acts across deities for each *polis* $\pi_{p,d}$ for a *polis* p and a god d as:

$$\pi_{p,d}^e = \frac{1}{T_p} \sum_{\{i:\delta_i^{polis}=p \cap \delta_i^{deity}=d\}} \delta_i^{NbObs}.$$

The gods in D_p can be ranked so that for each *polis* p , they are in descending order. The difference between each god k and the smallest probability (indexed at 0) is $\pi_{p,k}^e - \pi_{p,0}^e$. We call this quantity “NumOccurRatioAdj” in the dataset, and the god d ’s rank in *polis* p , $k(d,p)$, is called “NumOccurRank”. The resulting aggregated data therefore contains $\sum_{p \in \Pi} D_p$ rows.

In addition, some inscription-level or inscription group-level data such as the presence of a cult or the source can be aggregated at the *polis* level. They could not be used in a relevant manner at the deity level, since we need to examine the data in distribution form. For example, the aggregate percentage of epigraphic observations for *polis* p can be computed as follows, where $\delta_i^{epi} \in \{0, 1\}$ is a binary flag representing whether the source is marked as epigraphic:

$$a_p^{epi} = \frac{\sum_{\{i:\delta_i^{polis}=p\}} \delta_i^{epi} \delta_i^{NbObs}}{\sum_{\{i:\delta_i^{polis}=p\}} \delta_i^{NbObs}}.$$

The same calculation can also be carried out for the evidence of a cult where the inscription was found or referenced, a_p^{cult} .

Summary statistics for the BDEG after preprocessing are shown in Table 2, across all the data. One can see that many *poleis* only have a handful of observations. The same statistics are shown for the subset of *poleis* for which there were more observations in Table 3. We can see that we have some religious information on a little less than 400 different *poleis*, and filtering only for those with a minimal number of observations for very basic distributional analysis, we have about 100.

Table 2: *Summary Statistics on BDEG Data Across All Poleis*

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Median	Pctl(75)	Max
NumGods	396	5.553	6.294	1	1	3	7	45
NumOccurPolis	396	39.321	144.466	1	2	5	22.2	2,374
PctEpi	396	0.804	0.323	0	0.7	1	1	1
PctCult	396	0.833	0.261	0.000	0.752	0.955	1.000	1.000
ShareHighestPolis	396	0.625	0.290	0.135	0.357	0.562	1.000	1.000
MinOccurRatio	396	0.404	0.409	0.0004	0.045	0.200	1.000	1.000

Note: NumGods: $|D_p|$, number of separate gods observed in the location.
 NumOccurPolis: T_p , number of votive inscriptions in the location.
 PctEpi: a_p^{epi} , average percentage of epigraphic sources in the location.
 PctCult: a_p^{cult} , average percentage of cult signs in the location.
 ShareHighestPolis: $\pi_{p,|D_p|}^e$, percentage of acts towards the most favored god in the location.
 MinOccurRatio: $\pi_{p,0}^e$, percentage of acts towards the least favored god in the location.

As we noted earlier, the names of the gods that received the most attention within each *polis*

Table 3: *Summary Statistics on BDEG Data On Poleis with more than 25 Observations*

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Median	Pctl(75)	Max
NumGods	93	14.301	7.200	2	9	13	18	45
NumOccurPolis	93	148.806	271.402	26	38	60	164	2,374
PctEpi	93	0.870	0.174	0	0.8	1.0	1	1
PctCult	93	0.847	0.130	0.238	0.818	0.882	0.919	1.000
ShareHighestPolis	93	0.430	0.208	0.135	0.270	0.369	0.571	0.975
MinOccurRatio	93	0.018	0.018	0.0004	0.006	0.017	0.026	0.148

Note: NumGods, NumOccurPolis, PctEpi, PctCult, ShareHighestPolis, MinOccurRatio as defined before.

were fairly diverse. For each *polis*, we extracted the name of the most commonly mentioned god, as well as the second most commonly mentioned, and from that data we computed their frequencies across all *poleis*. Table 4 shows the twelve most common. We can see that there is substantial variety in which gods were worshipped the most in each *polis*. This further stresses that the particular powers each god was linked to were presumably not the main reason why they were selected.

Table 4: *Ranking of the First and Second Gods Most Commonly Mentioned in each Polis*

First_God	Nb_Cities_First	Nb_Obs_First	Second_God	Nb_Cities_Second	Nb_Obs_Second
Zeus	142	5,624	Zeus	52	594
Apollon	51	1,007	Apollon	34	1,032
Artemis	41	720	Artemis	30	236
Athena	29	366	Athena	23	945
Theoi	23	117	Theoi	23	167
Dionysos	13	165	Theos	15	57
Theos	13	59	Aphrodite	10	37
Aphrodite	9	30	Demeter	8	23
Hermes	9	54	Dionysos	8	44
Asclepios	7	206	Hermes	7	14
Demeter	7	15	Poseidon	7	21
Poseidon	7	37	Thea	7	27

In order to complement the data from the BDEG, we merge it with the POLIS database, Johnson and Ober (2014), the computerized and augmented version of Hansen and Nielsen (2004). The majority of *poleis* listed in the POLIS database have geographic coordinates. Hence the mapping of cult epithet observations to a particular *polis* is preferentially done using these coordinates, and when they were not available using the location names³⁷.

3.2 Analyzing Cult Data

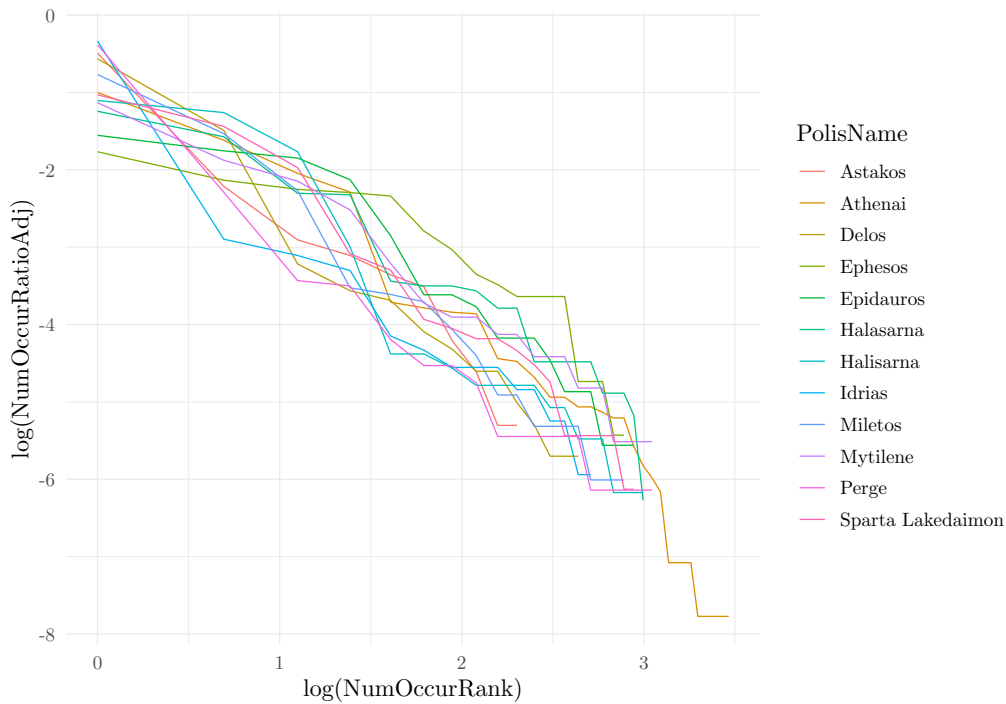
As we have discussed in the first section, we have various reasons to expect some form of power law in the data. The first reason is that people would directly assume this type of distribution for the gods' efficiency, because it is what one observes when looking at the outcomes of high level competitive activities. The second reason is that people would be inferring that some gods are more efficient than others based on some random life events taking place, which would

³⁷This required a substantial amount of manual mapping of the names, as they are not expressed in the same fashion: typically Greek or English for POLIS, and Greek or French for the BDEG.

cumulatively lead to a power law-based random draw of votive choices at each period. The first framework would call for a constant power law distribution over time, the second one for a randomly selected distribution over time, following a power law. Since we cannot distinguish time periods precisely, and consider data as a whole, through time, we could not make a difference between these two approaches.

In either case, we would expect some power law relationship, and the most direct way to visualize whether that is the case is illustrated in Figure 2, which shows the log/log plot of $\ln(\pi_{p,k}^e - \pi_{p,0}^e)$ as a function of $\ln(k)$, for *poleis* for which there were over 100 observations. All the curves do not have the same length, as the number of gods for which data has been found is not the same for all *poleis*, and the number of gods people worshipped may not be the same either. The curves appear to be fairly straight, which bodes well for our ability to represent the distributions as power laws.

Figure 2: *Log/log Plot of God Worship Distribution For Various Poleis*



Note: NumOccurAdjRatio is the term $\pi_{p,k}^e - \pi_{p,0}^e$. NumOccurRank is the rank k .

As we mentioned earlier, the number of gods worshipped by people in a given location may depend on a variety of factors, if we consider that there is some threshold under which an ill-favoured god would not receive votive acts. Table 5 shows several regressions for the number of gods in each *polis*, for which there was data in the sources. Selecting the significant factors, one can account for a reasonable fraction of this number. Interestingly, the size of the city does not

appear to play a very strong role, but the size captured in the POLIS data is an approximation of the *polis*'s surface, not its population (although the two may be related). We know there were substantial variations in the density of *poleis* in classical Greece³⁸. The factors that seem both significant and impactful are “fame” and the share of epigraphic sources for the BDEG data. A greater use of epigraphic sources in a particular place could explain that we have found more remains of votive acts than we would have otherwise. The metric for fame refers to the importance of a *polis* in literary sources, and the greater its importance, the greater number of people might have transited through the place, hence potentially leaving a trace of worship.

Table 5: *Empirical Drivers of the Number of Gods For Which References Were Found*

	Dependent variable:		
	NumGods		
	Size	OLS All Terms	Selected Terms
	(1)	(2)	(3)
SizeProxy	0.015*** (0.004)	−0.0001 (0.002)	0.003** (0.001)
Size	−0.873 (0.759)		
PctEpi		5.167*** (1.923)	3.044*** (0.932)
PctCult		2.159 (2.335)	
Fame		1.725*** (0.277)	1.523*** (0.149)
Democracy		0.701 (1.357)	
Colonies		0.051 (0.135)	
DelianLeague		2.251* (1.206)	
Greek		1.906 (2.407)	
Constant	4.756*** (1.351)	−7.360** (3.591)	−0.809 (0.864)
Observations	280	97	280
R ²	0.171	0.605	0.408
Adjusted R ²	0.165	0.569	0.401
Residual Std. Error	5.962	5.352	5.049
F Statistic	28.606***	16.829***	63.375***

Note:

*p<0.1; **p<0.05; ***p<0.01

3.3 Model Test


The theoretical model we presented earlier offers predictions for two things: the frequency of worship for the least worshipped god (which should be $\frac{1}{2}$ of a random uniform selection), and the difference between the other gods' frequency and that least worshipped god (which is expected to follow a power law).

Figure 3 plots for each *polis* p the actual relative frequency $\pi_{p,0}^e$ versus the predicted probability according to the model $\frac{1}{2D_p}$, for all *poleis* with more than 1 god observed. The actual frequencies appear somewhat lower than the predictions, and the shape of the relationship points to some heteroscedasticity.

The regressions of predicted vs. actual values shown in Table 6 confirm the visual perception. Since the smaller values for the observations and projections also correspond to the cases where there are many datapoints, it seems appropriate to use the total number of observations for each

³⁸See Ober (2015).

Figure 3: *Predicted vs. Actual Probability of Worship for the Least Worshipped God*



Note: $\text{ProjOccur} = \frac{1}{2D_p}$, $\text{NumOccurRatio} = \pi_{p,0}^e$.

polis as weights to compensate for it at the first order. Then, in order to really weight the data by the number of observations in each polis, the weights need to be squared. The coefficients, actually larger than 1, show that there definitely is mixing, whereby the least favored gods still receive a minimal amount of attention. Figure 3 shows however that in cases when there were many gods (and therefore a low projected probability), the actual values were lower than projected. This overestimation of the data could be explained by the fact that the inscriptions that have been found are not necessarily comprehensive for the gods which do not appear often, and as a result the number of separate gods derived from the data is likely underestimated.

Table 6: *Empirical Regressions for Predicted vs. Actual Probability of Worship for the Least Worshipped God*

	<i>Dependent variable:</i>		
	NumOccurRatio		
	<i>OLS</i>		
	All data	More than 5 gods	All with weights
	(1)	(2)	(3)
ProjOccur	1.922*** (0.054)	0.955*** (0.046)	1.146*** (0.051)
Observations	279	150	279
R ²	0.817	0.744	0.640
Adjusted R ²	0.817	0.743	0.639
Residual Std. Error	0.095 (df = 278)	0.031 (df = 149)	0.354 (df = 278)
F Statistic	1,244.893*** (df = 1; 278)	433.630*** (df = 1; 149)	495.202*** (df = 1; 278)

Note:

*p<0.1; **p<0.05; ***p<0.01

Now shifting our attention to the other gods who received more favors from the worshippers,

we consider several possible models to account for the relationship between $\ln(\pi_{p,k}^e - \pi_{p,0}^e)$ and k . We test a pure power model, a pure exponential, and a power law with exponential decay, since these are the kinds of models that have been used in human performance modeling, and we also know that pure power laws as well as with an exponential decay would be consistent with the random growth model of beliefs. We also include a lognormal distribution as an additional benchmark.

While Figure 3 may entice us to simply use a linear regression to estimate the parameters and fit with a power law, this would in fact be incorrect. As was shown by Clauset, Shalizi, and Newman (2009), in order to correctly test for the presence of power laws in distributions, a more involved approach is necessary based on maximum-likelihood estimators and goodness-of-fit. Various implementations of this approach are available in R or in Python³⁹. We are facing a particular issue however in that we are not considering a single distribution, but a large number of them, with varying amounts of data. This is addressed by carrying out as many fits as there are *polis*. For each *polis* p , we will obtain some parameters, for example α_p in the case of a pure power law, fitted to the BDEG data for that *polis*.

The data consists of the empirical differences (expressed in observations counts, not as relative frequencies) $(\pi_{p,k}^e - \pi_{p,0}^e) T_p$ with the $\pi_{p,k}^e$ which we defined earlier, for each *polis* p . We select only those *polis* (a little less than a hundred of them) for which there are at least 25 observations in total⁴⁰. For each *polis*, we apply the methodology recommended in Clauset, Shalizi, and Newman (2009) in order to estimate the parameters of various potential distribution choices: exponential⁴¹, straight power law, power law with exponential decay, and also lognormal. The latter two use two parameters, versus only one for the prior two, and would hence be likely to better fit the data. We run goodness-of-fit comparison tests between these distributions, for each *polis*. For power law and truncated power law fits, the minimum value l_0 is set to 1 (it is not determined by the fit).

The tests for distribution types across all *polis* with enough data are summarized in Table 7. We can see that the power law coefficient estimates (both for the pure and truncated power laws) are very tightly concentrated. This could point towards generating mechanisms being quite similar across a large set of different polities. The research on sports and competition rankings such as Deng et al. (2012) and Morales et al. (2016) pointed to some substantial variations in the parameterization of power laws across activities. Therefore, anthropomorphizing the god's qualities drives their perceived efficiency, one may have expected some degree of variability across

³⁹See Alstott, Bullmore, and Plenz (2014) and Gillespie (2020).

⁴⁰The descriptive statistics for this subset were shown in Table 3.

⁴¹Which is the upper bound effectively defining a fat tail distribution.

space. The narrow concentration of parameter estimates could point to a more fundamental mechanism, such as the one posited by Proposition 2.1, which we know could lead empirically to observing the patterns from Table 7. The lognormal distributions fit the data well, but the resulting parameters are much less stable across *poleis*.

Table 7: *Summary Statistics on Distribution Fits across Poleis*

Statistic	Mean	St. Dev.	Min	Pctl(25)	Median	Pctl(75)	Max
Lambda Exp	0.020	0.001	0.017	0.019	0.020	0.020	0.022
Mu LogN	-0.391	0.376	-1.326	-0.652	-0.361	-0.131	0.798
Sigma LogN	3.087	0.138	2.598	3.013	3.087	3.175	3.412
Alpha Pow	0.381	0.007	0.357	0.376	0.380	0.385	0.399
Alpha Trunc	0.251	0.020	0.184	0.239	0.249	0.264	0.291
Lambda Trunc	0.001	0.0001	0.001	0.001	0.001	0.001	0.002
Trunc vs Pow R	3.106	0.102	2.989	3.048	3.072	3.132	3.752
Trunc vs Pow p	0.001	0.0004	0.00001	0.0005	0.001	0.001	0.002
Trunc vs LogN R	1.530	0.323	0.491	1.345	1.567	1.740	2.167
Trunc vs LogN p	0.146	0.105	0.030	0.082	0.117	0.179	0.624
Trunc vs Exp R	4.837	0.187	4.421	4.739	4.801	4.934	5.711
Trunc vs Exp p	0.000002	0.000001	0.000	0.000001	0.000002	0.000002	0.00001
Pow vs Exp R	4.271	0.205	3.648	4.152	4.240	4.390	5.071
Pow vs Exp p	0.00003	0.00003	0.0000004	0.00001	0.00002	0.00003	0.0003

Note: The distribution names are abbreviated as follows:
Exp = exponential, *LogN* = lognormal, *Pow* = (pure) power, *Trunc* = power law with exponential decay.
R: ratio of goodness-of-fit; a positive number means that the first law of the two is preferred.
p: significance level; the probability that the preference would be due to randomness.

As we mentioned earlier, we have had to apply an approximation to capture the number of votive acts related to entries into the BDEG that referred to multiple observations. Table 8 shows the distribution fits on the data compiled excluding these observations, that is only including the entries that referred to a unique occurrence. The slope of the power fit increases in this case, showing that entries in the data that only reflected a single observation tended to be associated with the most commonly worshipped gods. This could be a bias in the manner in which the data was entered, as more attention was potentially dedicated to the inscriptions relating to the most important gods. Nevertheless, the dispersion around this mean slope, across all *poleis*, remains extremely low.

3.4 Explaining Differences Across *Poleis*

While we can account for the overall shape of the distribution of votive acts across gods, an interesting question is to determine what in each *polis*'s characteristics may drive differences in these distributions. Since for each *polis* we possess a series of categorical or numerical variables, we may be able to find to what extent the perceived differences between the gods is impacted by certain features of the *polis*.

Figures 4, 5 and 6 show some illustrations of such a potential relationship, and plot the distributions in a log/log scale respectively as a function of the political regime, the degree of hellenicity, and the size of the *polis*, as categorized by the POLIS data. The curves for non democratic *poleis* may seem steeper than those for democratic regimes, but that is not much

Table 8: *Summary Statistics on Distribution Fits across Poleis Excluding Observations of Multiple BDEG Entries*

Statistic	Mean	St. Dev.	Min	Pctl(25)	Median	Pctl(75)	Max
Lambda Exp	0.058	0.002	0.053	0.056	0.057	0.059	0.063
Mu LogN	-4.346	0.944	-7.005	-4.833	-4.372	-3.764	-2.340
Sigma LogN	3.499	0.231	2.990	3.362	3.514	3.639	4.055
Alpha Pow	0.528	0.007	0.513	0.524	0.527	0.532	0.545
Alpha Trunc	0.515	0.015	0.467	0.502	0.515	0.524	0.554
Lambda Trunc	0.001	0.0002	0.001	0.001	0.001	0.001	0.002
Trunc vs Pow R	3.138	0.140	2.686	3.054	3.158	3.245	3.531
Trunc vs Pow p	0.015	0.004	0.002	0.013	0.015	0.017	0.023
Trunc vs LogN R	2.176	0.388	1.089	2.002	2.326	2.461	2.843
Trunc vs LogN p	0.045	0.056	0.004	0.014	0.020	0.045	0.276
Trunc vs Exp R	5.007	0.102	4.775	4.946	5.026	5.069	5.220
Trunc vs Exp p	0.000001	0.0000004	0.0000002	0.0000004	0.000001	0.000001	0.000002
Pow vs Exp R	4.801	0.117	4.540	4.730	4.819	4.872	5.050
Pow vs Exp p	0.000002	0.000001	0.0000004	0.000001	0.000001	0.000002	0.00001


Note: The distribution names are abbreviated as follows:
Exp = exponential, *LogN* = lognormal, *Pow* = (pure) power, *Trunc* = power law with exponential decay.
R: ratio of goodness-of-fit; a positive number means that the first law of the two is preferred.
p: significance level; the probability that the preference would be due to randomness.

more than an impression. *Poleis* for which the population was considered more Hellenic and less varied appear to show flatter curves than these categorized as including more people of barbarian origin. *Polis* size on the other hand does not appear to have much of an impact, on the basis of casual observation.

There are many more *polis*-level characteristics that may or may not play a role in empirically explaining the distribution of votive acts. We can use the results from the distribution fits, and will seek to explain the estimated parameters for various types of distributions across *poleis* as a function of the *poleis*' characteristics. Table 9 shows two of these regressions for the parameter α , for the power law and for the truncated power law. The parameters for these fits are so tightly distributed that the regression coefficients are very small. The R-squares are high, but essentially no variable has a strong or significant effect.


By including so many variables, we end up selecting very few datapoints that are not missing any data and the regressions in Table 9 miss many observations. Selecting only a few explanatory variables, we can obtain a regression that spans more observations, as shown in Table 10. The table also shows the regression for the pure power law coefficient on the data that excludes multiple observations. There are now more observations. With fewer variables, the coefficients are generally significant, and the R-squares remain high. In addition, the signs and magnitudes of the coefficients are consistent across the pure power fit, the truncated fit, and the pure power fit excluding multiple occurrence observations. These coefficients hence appears reasonably robust. The larger the *polis*, the steeper (less fat-tailed) the distribution; *poleis* considered more Greek have a flatter (more fat-tailed) curve, and the total number of observations tends to flatten the curve. However, their magnitude is quite small: varying the inputs, the estimated parameter α for the pure power law may change from 0.39 to 0.38 or 0.37. While this impact may be

Figure 4: *Log/log Plot of God Worship Distribution Grouped by Political Regime*




Note: NumOccurAdjRatio is the term $\pi_{p,k}^e - \pi_{p,0}^e$. NumOccurRank is the rank k . The data only includes *poleis* with more than 25 observations.

Figure 5: *Log/log Plot of God Worship Distribution Grouped by Hellenicity*



Note: NumOccurAdjRatio is the term $\pi_{p,k}^e - \pi_{p,0}^e$. NumOccurRank is the rank k . The data only includes *poleis* with more than 25 observations.

Figure 6: *Log/log Plot of God Worship Distribution Grouped by Polis Size*



Note: NumOccurAdjRatio is the term $\pi_{p,k}^e - \pi_{p,0}^e$. NumOccurRank is the rank k . The data only includes *poleis* with more than 25 observations.

statistically significant, the magnitude of the effect seems too small to consider it meaningful.

We can associate the power coefficient to the assumed degree of competitiveness between the gods, in the context of the anthropomorphic perspective that would derive the god's perceived efficiency from human performance observation. If we place ourselves in the context of a random growth model explaining the eventual god efficiency, then there would not be any particular reason for some gods to be assumed to be better than other gods, other than pure chance. In light of these considerations, the random growth hypothesis appears the most appropriate to explain the empirical distribution of epigraphic and literary references to cult.

4 Conclusion: Implications for Historical Study

We began by looking into the way the ancient Greeks considered their relationships with divine beings. Based on several salient elements, we postulated a few principles off of which we could build a simple model relating some particular characteristics of the gods and the propensity that people may have to worship them. We related the number of people worshipping each god to a measure of that god's efficiency. We proposed that this divine efficiency could have been estimated by people based on their observation of human performance. Another potential for the gods' efficiencies to evolve was based on random life occurrences in a differential growth model.

Table 9: *Regressions of Distribution Parameters with All Variables*

	<i>Dependent variable:</i>	
	‘Alpha Pow‘	‘Alpha Trunc‘
	<i>OLS</i> Power law	<i>OLS</i> Truncated Power
	(1)	(2)
PctEpi	0.007 (0.010)	-0.004 (0.026)
PctCult	-0.014 (0.014)	-0.022 (0.036)
Fame	0.001 (0.001)	0.003 (0.002)
Democracy	0.004 (0.005)	0.012 (0.013)
Size	-0.003* (0.002)	-0.012** (0.004)
Colonies	0.0001 (0.0002)	0.0003 (0.001)
DelianLeague	-0.002 (0.004)	-0.005 (0.009)
Greek	-0.003 (0.016)	-0.014 (0.041)
HasVictories	-0.0001 (0.004)	-0.006 (0.011)
GaveProxeny	0.0004 (0.004)	-0.0001 (0.009)
HasWalls	0.003 (0.005)	0.015 (0.013)
Elevation_m	0.00000 (0.00001)	0.00002 (0.00002)
NumOccurPolis	-0.00002*** (0.00000)	-0.00005*** (0.00001)
Constant	0.394*** (0.023)	0.303*** (0.057)
Observations	36	36
R ²	0.881	0.853
Adjusted R ²	0.811	0.767
Residual Std. Error	0.079	0.201
F Statistic	12.585***	9.847***

Note: *p<0.1; **p<0.05; ***p<0.01

Table 10: *Regressions of Distribution Parameters with Selected Variables*

	<i>Dependent variable:</i>		
	‘Alpha Pow‘	‘Alpha Trunc‘	‘Alpha Pow‘
	<i>OLS</i> Power law	<i>OLS</i> Truncated Power	<i>OLS</i> Power Law (Single Occ.)
	(1)	(2)	(3)
Greek	-0.013*** (0.004)	-0.035*** (0.012)	-0.014*** (0.003)
Size	0.002*** (0.001)	0.006*** (0.002)	-0.001*** (0.0005)
NumOccurPolis	-0.00001*** (0.00000)	-0.00003*** (0.00000)	-0.00001*** (0.00000)
Constant	0.388*** (0.004)	0.266*** (0.014)	0.547*** (0.004)
Observations	66	66	39
R ²	0.742	0.620	0.786
Adjusted R ²	0.730	0.602	0.768
Residual Std. Error	0.084	0.262	0.061
F Statistic	59.580***	33.776***	42.896***

Note: *p<0.1; **p<0.05; ***p<0.01

Using data from the BDEG and from POLIS, we then examined the empirical distribution of votive acts, as captured mainly by epigraphic sources.

As we have seen, the model based on optimal god choice would have called for a somewhat greater mixing of the power law with a uniform than what we observed empirically. Nevertheless, the data did show a substantial enough amount of random mixing in the choice of deities: even the least favored gods did see quite a few offerings or dedications. This is consistent with the intuition of the model, that it is not optimal to overcrowd what is deemed as a relatively limited resource.

The distribution of choices itself across the full range of gods fitted the power law (with or without exponential decay) quite well. The fitted exponent parameters α_p were also surprisingly stable across *poleis*. Searching for explanatory factors, it appeared that not only few of them happened to be significant, but also their magnitude was extremely small. This observation seems consistent with the random growth model, for which there would be no fundamental difference from one *polis* to another. In contrast, one could expect that the particular environment the *polis* experiences could affect the collective perception of human performance, and hence the shape of the empirically observed cult occurrences.

References

- Alstott, Jeff, Ed Bullmore, and Dietmar Plenz.** 2014. “Powerlaw: A Python Package for Analysis of Heavy-Tailed Distributions.” *PLoS ONE* 9 (1): e85777.
- Arthur, W. Brian.** 1994. “Inductive Reasoning and Bounded Rationality.” *The American Economic Review* 84 (2): 406–411.
- Baek, Seung Ki, Sebastian Bernhardsson, and Petter Minnhagen.** 2011. “Zipf’s Law Unzipped.” *New Journal of Physics* 13 (4): 043004.
- Bankston, Carl L.** 2002. “Rationality, Choice and the Religious Economy: The Problem of Belief.” *Review of Religious Research* 43 (4): 311–325.
- Bonnet, Corinne, and Sylvain Lebreton.** 2019. “Mettre les polythéismes en formules ? À propos de la base de données Mapping Ancient Polytheisms.” *Kernos* 32 (December): 267–296.
- Brulé, Pierre.** 1998. “Le langage des épiclèses dans le polythéisme hellénique (l’exemple de quelques divinités féminines). Quelques pistes de recherches.” *Kernos* 11 (11): 13–34.
- . 2015. “Chapitre XV. Le polythéisme en transformation : les listes de dieux dans les serments internationaux en Grèce antique (Ve-IIe siècles).” In *La Grèce d’à côté : Réel et imaginaire en miroir en Grèce antique*, 333–363. Histoire. Rennes: Presses universitaires de Rennes.
- Brulé, Pierre, and Sylvain Lebreton.** 2007. “La Banque de données sur les épiclèses divines (BDDE) du Crescam : sa philosophie.” *Kernos* 20: 217–228.
- Bulbulia, Joseph.** 2010. “Charismatic Signalling.” *Journal for the Study of Religion, Nature and Culture* 3 (4): 518–551.
- Bulbulia, Joseph, and Marcus Frean.** 2010. “The Evolution of Charismatic Cultures.” *Method & Theory in the Study of Religion* 22 (4): 254–271.
- Carr, Jack L., and Janet T. Landa.** 1983. “The Economics of Symbols, Clan Names, and Religion.” *The Journal of Legal Studies* 12 (1): 135–156.
- Chakrabarti, Anindya Sundar, Bikas K. Chakrabarti, Arnab Chatterjee, and Manipushpak Mitra.** 2009. “The Kolkata Paise Restaurant Problem and Resource Utilization.” *Physica A: Statistical Mechanics and Its Applications* 388 (12): 2420–2426.
- Chao, M. T., and W. E. Strawderman.** 1972. “Negative Moments of Positive Random Variables.” *Journal of the American Statistical Association* 67 (338): 429–431.

- Chwe, Michael Suk-Young.** 2001. *Rational Ritual: Culture, Coordination, and Common Knowledge*. Princeton, N.J: Princeton University Press.
- Clauset, Aaron, Cosma Rohilla Shalizi, and M. E. J. Newman.** 2009. “Power-Law Distributions in Empirical Data.” *SIAM Review* 51 (4): 661–703.
- Deng, Weibing, Wei Li, Xu Cai, Alain Bulou, and Qiuping A Wang.** 2012. “Universal Scaling in Sports Ranking.” *New Journal of Physics* 14 (9): 093038.
- Donner, Yoni, and Joseph L. Hardy.** 2015. “Piecewise Power Laws in Individual Learning Curves.” *Psychonomic Bulletin & Review* 22 (5): 1308–1319.
- Ekroth, Gunnel.** 2011. “Meat for the gods.” In *Nourrir les dieux ?*, edited by Vinciane Pirenne-Delforge and Francesca Prescendi, 15–41. Presses universitaires de Liège.
- Ferrero, Mario, and George Tridimas.** 2018. “Divine Competition in Greco–Roman Polytheism.” *Homo Oeconomicus* 35 (3): 143–166.
- Gabaix, Xavier.** 2008. “Power Laws in Economics and Finance.” Working Paper 14299. Working Paper Series. National Bureau of Economic Research. <https://doi.org/10.3386/w14299>.
- . 2016. “Power Laws in Economics: An Introduction.” *Journal of Economic Perspectives* 30 (1): 185–206.
- Gauthier, Laurent.** 2020a. *Securitization Economics: Deconstructing the Economic Foundations of Asset Securitization*. Springer Texts in Business and Economics. Cham: Springer International Publishing.
- . 2020b. “New Cliometrics for Ancient History.” SSRN Scholarly Paper ID 3888407. Rochester, NY: Social Science Research Network. <https://papers.ssrn.com/abstract=3888407>.
- Gillespie, Colin.** 2020. *poweRlaw: Analysis of Heavy Tailed Distributions* (version 0.70.6).
- Goldie, Charles M.** 1991. “Implicit Renewal Theory and Tails of Solutions of Random Equations.” *Annals of Applied Probability* 1 (1): 126–166.
- Gries, Peter, Jenny Su, and David Schak.** 2012. “Toward the Scientific Study of Polytheism: Beyond Forced-Choice Measures of Religious Belief.” *Journal for the Scientific Study of Religion* 51 (4): 623.
- Hansen, Mogens Herman, and Thomas Heine Nielsen.** 2004. *An Inventory of Archaic and Classical Poleis*. Oxford ; New York: Oxford University Press.
- Henrich, Joseph.** 2009. “The Evolution of Costly Displays, Cooperation and Religion: Credibility Enhancing Displays and Their Implications for Cultural Evolution.” *Evolution and*

Human Behavior 30 (4): 244–260.

Hotelling, Harold. 1929. “Stability in Competition.” *The Economic Journal* 39 (153): 41–57.

Iannaccone, Laurence R. 1998. “Introduction to the Economics of Religion.” *Journal of Economic Literature* 36 (3): 1465–1495.

———. 2005. “Economics of Religion. Debating the Costs and Benefits of a New Field.” *Faith and Economics* 46: 1–9.

Iannaccone, Laurence R., Colleen E. Haight, and Jared Rubin. 2011. “Lessons from Delphi: Religious Markets and Spiritual Capitals.” *Journal of Economic Behavior & Organization* 77 (3): 326–338.

Iyer, Sriya. 2016. “The New Economics of Religion.” *Journal of Economic Literature* 54 (2): 395–441.

Johnson, Dominic D. P. 2005. “God’s Punishment and Public Goods: A Test of the Supernatural Punishment Hypothesis in 186 World Cultures.” *Human Nature* 16 (4): 410–446.

Johnson, Tim, and Josiah Ober. 2014. “POLIS.” Stanford University. 2014. <http://polis.stanford.edu>.

Kaše, Vojtěch, Tomáš Hampejs, and Zdeněk Pospíšil. 2018. “Modeling Cultural Transmission of Rituals in Silico: The Advantages and Pitfalls of Agent-Based Vs. System Dynamics Models.” *Journal of Cognition and Culture* 18 (5): 483–507.

Kesten, Harry. 1973. “Random Difference Equations and Renewal Theory for Products of Random Matrices.” *Acta Mathematica* 131 (0): 207–248.

Kohli, Rajeev, and Raaj Sah. 2006. “Some Empirical Regularities in Market Shares.” *Management Science* 52 (11): 1792–1798.

Lebreton, Sylvain. 2019. “Des dieux et des personnes : une approche par les réseaux (mondes grecs et ouest-sémitiques, ca. 1000 av. n. è. – 400 d. n. è.)” In *La personne en question dans les réseaux*.

Lebreton, Sylvain, Jean-Baptiste Barreau, Karine Karila-Cohen, and Pierre Brulé. 2014. “Banque de Données des Épiclèses Grecques (BDEG).” 2014. <https://epiclesesgrecques.univ-rennes1.fr/?lang=en>.

Levy, Gilat, and Ronny Razin. 2007. “A Theory of Religion: Linking Individual Beliefs.” 19. *2007 Meeting Papers*. 2007 Meeting Papers. Society for Economic Dynamics. <https://ideas.repec.org/p/red/sed007/19.html>.

- . 2012. “Religious Beliefs, Religious Participation, and Cooperation.” *American Economic Journal: Microeconomics* 4 (3): 121–151.
- . 2014. “Rituals or Good Works: Social Signaling in Religious Organizations.” *Journal of the European Economic Association* 12 (5): 1317–1360.
- McBride, Michael.** 2016. “A Rational Choice Theory of Religious Authority.” *Rationality and Society* 28 (4): 410–438.
- Miller, Peter J.** 2018. “In the Shadow of Praise: Epinician Losers and Epinician Poetics.” *Bulletin of the Institute of Classical Studies* 61 (1): 21–41.
- Mitzenmacher, Michael.** 2004. “A Brief History of Generative Models for Power Law and Lognormal Distributions.” *Internet Mathematics* 1 (2): 226–251.
- Morales, José A., Sergio Sánchez, Jorge Flores, Carlos Pineda, Carlos Gershenson, Germinal Cocho, Jerónimo Zizumbo, Rosalío F. Rodríguez, and Gerardo Iñiguez.** 2016. “Generic Temporal Features of Performance Rankings in Sports and Games.” *EPJ Data Science* 5 (1): 33.
- Ober, Josiah.** 2015. *The Rise and Fall of Classical Greece*. Princeton: Princeton University Press.
- Parker, Robert.** 2003. “The Problem of the Greek Cult Epithet.” *Opuscula Atheniensia* 28 (2003): 173–183.
- Power, Eleanor A.** 2017. “Discerning Devotion: Testing the Signaling Theory of Religion.” *Evolution and Human Behavior* 38 (1): 82–91.
- Ramsay, Colin M.** 2006. “The Distribution of Sums of Certain I.I.D. Pareto Variates.” *Communications in Statistics - Theory and Methods* 35 (3): 395–405.
- Ruffle, Bradley J., and Richard Sosis.** 2007. “Does It Pay to Pray? Costly Ritual and Cooperation.” *The B.E. Journal of Economic Analysis & Policy* 7 (1).
- Smith, Thomas Jay.** 2015. *Variability in Human Performance*. Human Factors and Ergonomics. Boca Raton: CRC Press, Taylor & Francis Group.
- Vandewalle, H.** 2018. “Modelling of Running Performances: Comparisons of Power-Law, Hyperbolic, Logarithmic, and Exponential Models in Elite Endurance Runners.” *BioMed Research International* 2018 (October): 1–23.
- Versnel, H. S.** 2011. *Coping with the Gods: Wayward Readings in Greek Theology*. Religions in the Graeco-Roman World, v. 173. Leiden ; Boston: Brill.

Zaidman, Louise Bruit, and Pauline Schmitt Pantel. 1992. *Religion in the Ancient Greek City*. Cambridge: Cambridge University Press.