CAMBRIDGE
Cambridge University Press & Assessment
978-1-009-10065-6 — The Uncertain Past
Edited by Myles Lavan , Daniel Jew , Bart Danon
Excerpt
More Information

1 Probabilistic Modelling in Ancient History*

Daniel Jew and Myles Lavan

1.1 Introduction

Recent decades have seen ever wider recognition among ancient historians of the importance of establishing at least approximate estimates for key historical quantities such as population and gross domestic product (GDP), if we are to write adequate social, political and even cultural histories of the Greek and Roman worlds. These quantities can rarely be estimated precisely, but ancient history has always been a pragmatic discipline. We are accustomed to making the most of what we have, despite considerable uncertainties. Quantitative approaches to ancient history can be traced back at least as far as early-modern debate on Roman population.[1] The late nineteenth century brought an important advance in analytical rigour, exemplified by Julius Beloch's still fundamental work on population.[2] The late 1970s saw another step change in the prominence and sophistication of quantitative approaches, exemplified above all by the work of Keith Hopkins, in the wake of the 'cliometric' revolution in the wider discipline of history.[3] Ancient historians have borrowed methods from other fields, liberated themselves from an exclusive focus on data transmitted by ancient texts and reflected more deeply on the use of models as heuristic devices. Parametric modelling – that is, the use of models that estimate uncertain quantities as a function of other, better understood parameters – has now become mainstream. Among the most famous and influential examples are

* We wish to thank Bart Danon for implementing the model in R (see the code in the appendix), comments and preparing the figures; Peter Garnsey, Chris Hope and Robin Osborne for comments on underlying work related to Attic grain; and Alberto Corrias for an earlier joint conference paper on interdependence in historical modelling (with DJ), which explored the concepts in notes 61–66.

1 See especially David Hume's essay 'Of the Populousness of Ancient Nations' (1752), with Scheidel (2001: 3–5).

2 Beloch 1886.

3 Notably Hopkins (1978, 1983). See Ruggles and Magnuson (2020) on the explosion of quantitative approaches in the broader field of history.

Hopkins' model of Roman GDP and Walter Scheidel and Steven Friesen's model of the distribution of income in the Roman Empire.[4]

Yet approaches to uncertainty remain relatively primitive by the standard of other disciplines. Ancient historians rely almost exclusively on what more sophisticated fields would term 'deterministic' or 'best estimate' models – models in which both inputs and outputs are point estimates (i.e. single values). The approach has the merit of simplicity and produces intuitive results. It may be perfectly adequate in situations where there are only one or two unknown variables or where the goal is to produce limiting scenarios, such as a theoretical minimum or maximum. But it quickly breaks down as the number of uncertain variables increases. In such cases, the historian relying on deterministic modelling techniques is left with an unpalatable choice between producing a best estimate that cannot hope to command credibility and abandoning the problem as unquantifiable. There is much that ancient historians can learn from other fields on this count. It is easy to assume that we have a monopoly on massive uncertainty. But there are many other fields in which calculations are based on subjective assessments of what is likely. Though the scientists who forecast climate change and the demographers who forecast population growth – to cite just two examples – may have data about past trends, projecting those trends into the future always involves significant uncertainty and subjective judgement. These fields are significantly more sophisticated in their understanding of uncertainty and their methods for managing it. A key breakthrough was a shift from deterministic models, which take a single set of inputs and produce a single output, to probabilistic models, which account for uncertainties by simulating very large numbers of possible scenarios and observing how often different outcomes occur.

Probabilistic methods have since become mainstream in the sciences and many other fields, but they remain largely unknown in pre-modern history, outside a few niches such as demographic microsimulation, carbon-14 dating or the dating of archaeological assemblages.[5] This volume aims to illustrate the untapped potential for probabilistic modelling to illuminate a much wider range of problems in ancient history. This chapter introduces key concepts and the underpinning theorisation of uncertainty and probability.[6] The remaining chapters are divided into two parts. The first (Uncertainty)

[4] Hopkins 1980 (updated by Hopkins 1995), Scheidel and Friesen 2009.
[5] See n. 24–5.
[6] The chapter reworks some material from Lavan (2019b: 95–9) and incorporates material from Jew's forthcoming study of the problem of Athenian carrying capacity. It also draws on several years of discussion of methods among the editors.

demonstrates the application of probabilistic modelling to a selection of problems in which epistemic uncertainty – that is, uncertainty caused by the limits of our knowledge – is a major obstacle to estimating important historical quantities: the rate of expropriation of property in the Greek world (Mackil), the number of families of senatorial wealth in first-century Pompeii (Danon) and the money supply in the second-century Roman Empire (Bransbourg). The second part (Variability and Missing Data) demonstrates the application of the same modelling approach to problems of a different type. Here two chapters aim to quantify variability within a population: economic outcomes for small tenants in Roman Egypt (Kelly) and the longevity of grain funds in the Greek cities of the Roman Empire (Solonakis et al.). A third chapter uses a probabilistic framework to impute missing data about the urban area of cities of the Roman Empire (Hanson). Though these are distinct exercises, they use the same mathematical tools – probability distributions and Monte Carlo simulation – as we explain in Section 1.4.6.

1.2 From Deterministic to Probabilistic Modelling

We illustrate the concept and the potential of probabilistic modelling by revisiting one of the oldest problems in the field: Classical Athens' dependence on imported grain.[7] The sustained interest in the question reflects the historical stakes, Attica's capacity to feed itself being an important determinant of Athens' geopolitical position. To simplify the presentation, we focus on one half of the problem: land carrying capacity, in the specific sense of the number of persons who could be fed by local grain.[8] Since production and consumption fluctuated from year to year, the question is one of long-term averages. We want to estimate the mean number of persons supported by grain grown in Attica over the period 500–300 BCE. This is an abstract quantity that is a simple mathematical function of two theoretically measurable, but unknown quantities:

$$\text{Land Carrying Capacity} = \frac{\text{Total Net Grain Production}}{\text{Per Capita Grain Consumption}}$$

[7] This description draws on Jew (forthcoming).

[8] Carrying capacity is a complex ecological term with a range of contested meanings in the literature (Abernethy 2001, Cheng et al. 2017). Here we use 'land carrying capacity for grain' in the simple sense of 'the population that can be supported by a region's grain output' at given levels of consumption (Feng et al. 2009: 52), within a particular system of agriculture and land usage (Allan 1965: 8–9).

Since we are interested in the average land carrying capacity for the fifth and fourth centuries, the total production and per capita consumption figures we need are long-term averages (which are easier to estimate than the figures for any particular year).

There is general agreement about the basic structure of the problem. Grain production can be disaggregated into a number of component variables that can be estimated separately: the amount of land cultivated for grain, the relative mix of barley and wheat, net yields for the two crops. Yet, uncertainty about the actual historical values of these variables leaves considerable scope for disagreement. Table 1.1 catalogues a selection of recent estimates for the carrying capacity of Attica, with brief notes on how they differ in their assumptions.

With estimates ranging from 71,000 to 150,000, it is no surprise that the problem remains contested. The estimate is highly sensitive to changes in the assumptions about the input variables, which are themselves clearly very uncertain. Ian Morris' damning criticism of the whole enterprise bears repeating:

> Models of the Athenian economy are simply less robust than those of Rome. Errors of ±15 percent in estimating population or production would have little impact on the overall shape of Hopkins' models of the Roman economy, but for Athens, they are devastating. Assuming 69,000 hectares of arable land, biennial fallow, and high population, Starr concluded that 'the fields of Attica could not have fed an urban center of even 10,000; once Athens rose toward that figure [in the seventh century?], it would have been necessary to import seaborne grain'. Assuming 96,000 hectares and less fallow, Garnsey suggests that 'a serious disequilibrium between Athens' food needs and its capacity to meet them did not develop until well into the post-Persian-War period'. Relatively minor changes to the numbers totally transform the models.[9]

The estimates by Starr, Osborne and Moreno are *best estimates* or *point estimates* based on *deterministic models* (i.e. models that use point-estimate inputs to produce a point-estimate output). They incorporate no measure of the margin of error. Garnsey and Sallares provide interval estimates, but the ranges are just an arbitrary concession to the uncertainty; they are not based on any formal assessment of the margin of error. None of the published estimates gives an adequate sense of the uncertainty about carrying capacity. It is thus no surprise that most such models have had a mixed reception in ancient history, vulnerable to being dismissed as

[9] Morris 1994: 361.

Table 1.1 Survey of estimates for the land carrying capacity of Attica *c*. 600–300 BCE

	Land carrying capacity based on grain (persons)	Note on key assumptions
Starr 1977	75,000	Low estimate due to the adoption of universal biennial fallow, high per capita grain consumption and very high seed/feed losses.
Osborne 1987	150,000	High estimate due to high area under cultivation and high grain yields.
Garnsey 1988	12,000–150,000 ('likely' answer: 132,000)	Allows for a range of assumptions; high 'likely' estimate due to the rejection of universal biennial fallow.
Sallares 1991	84,000–124,000 (central estimate: 104,000)	Middling estimate due to universal biennial fallow, moderately low land under grain and low grain yields.
Whitby 1998	71,000 (notional central estimate)	Low estimate due to universal biennial fallow, low land under grain and high per capita grain consumption.
Moreno 2007	84,000	Low estimate due to universal biennial fallow and high per capita grain consumption.
Bissa 2009	82,500–98,000 (implied; average deduced: 92,500)	Allows for a range of assumptions; middling central estimate due to moderate maximum land under grain.

Notes: Starr 1977: 155 (with critique at Garnsey 1998: 185–8); Osborne (1987: 45–6); Garnsey 1988: 101–5; Sallares 1991: 79–80, 309–10, 385–6; Whitby 1998: 104–6, 117–8 and n. 29; Moreno 2007: 10, Table 1; Bissa 2009: 173–6. Whitby expressed scepticism about offering precise figures; his notional estimate can be deduced by substituting his inputs (n. 29 therein) into Garnsey's model. Bissa's implied estimate is deduced from her respective central two scenarios for barley production (36,000 ha cultivated at a yield of 16.5 and 19.6 *medimnoi* per ha) and her middle estimate for per capita consumption (6.3 *medimnoi*/person/year). Her statement that 'the average import needed would be 936,000 *medimnoi* of barley' (p. 176) implies an average carrying capacity of 92,500, calculated from her Table 11.

merely 'lend[ing] a spurious authority to what is no more than a collection of guesses'.[10]

We have not yet, however, done full justice to one of these estimates. Peter Garnsey's 1988 study of Athenian grain production stands out as revolutionary for its time. Rather than just providing a single point estimate, he also addressed the question of uncertainty by reporting a range

[10] Whitby 1998: 117.

Table 1.2 Garnsey's model and assumptions

$$\textit{Carrying Capacity} = \frac{L \times G\ \ (B \times BY \times (1 - BL) + (1 - B) \times WY \times (1 - WL))}{C}$$

	Input Variable	Minimum	Likely	Maximum
L	Surface area of Attica (ha)		240,000	
G	Attic land under grain (%)	10	17.5	30
B	Grain land under barley (%)	67	80	90
WY	Wheat yield (kg/ha)	300	600	900
BY	Barley yield (kg/ha)	500	800	1300
WL	Wheat losses (%)	16.7	33.3	50.0
BL	Barley losses (%)	16.7	25.0	50.0
C	Grain consumption (kg/person/year)	150	175	230

of possible values ('minimum', 'likely' and 'maximum') for the uncertain variables.[11] We take Garnsey's model as our point of departure for explaining the probabilistic approach, so we reproduce his mathematical model of carrying capacity and his assumptions about the input variables in Table 1.2.[12] This is *the model*. It expresses the *quantity of interest* (land carrying capacity) as a function of other quantities about which we have better information. In the context of the model, the latter are the *input variables*.

Since Garnsey's approach was still deterministic, he could only simulate one scenario at a time. He presented a best estimate for carrying capacity (his 'modal answer' of 132,000 persons, i.e. 55 persons per km^2) based on his 'likely' values for each of the variables. But he included an innovative sensitivity analysis – reproduced here as Figure 1.1 – showing the sensitivity of that result to the changes in the values of the seven uncertain input variables. We return to the idea of sensitivity analysis in Section 1.9.

Garnsey did not attempt to compute a theoretical maximum or minimum by setting all the input variables to their most optimistic or most

[11] Garnsey 1988: 102. Garnsey's method of offering multiple scenarios was subsequently emulated by Sallares and Bissa.

[12] Garnsey did not include a formal presentation of his model as we do here, but the procedure is obvious and can be confirmed by reproducing his results. Note that Garnsey formulated his assumptions about yields in terms of hectolitres per hectare (as had previous scholars). This complicates the calculation (requiring production to be converted from hl to kg before it can be divided by per capita consumption) and obstructs comparison with subsequent scholarship, which is generally expressed in kg/ha. So we have converted his figures from hl/ha to kg/ha (using 64.3 and 77.2 kg/hl for barley and wheat, respectively, as implied in Garnsey's calculations; cf. Garnsey 1998: 193) and rounded to the nearest 100 to avoid a misleading impression of precision. Note that we have used the unrounded figures in our model calculations, to replicate Garnsey's results.

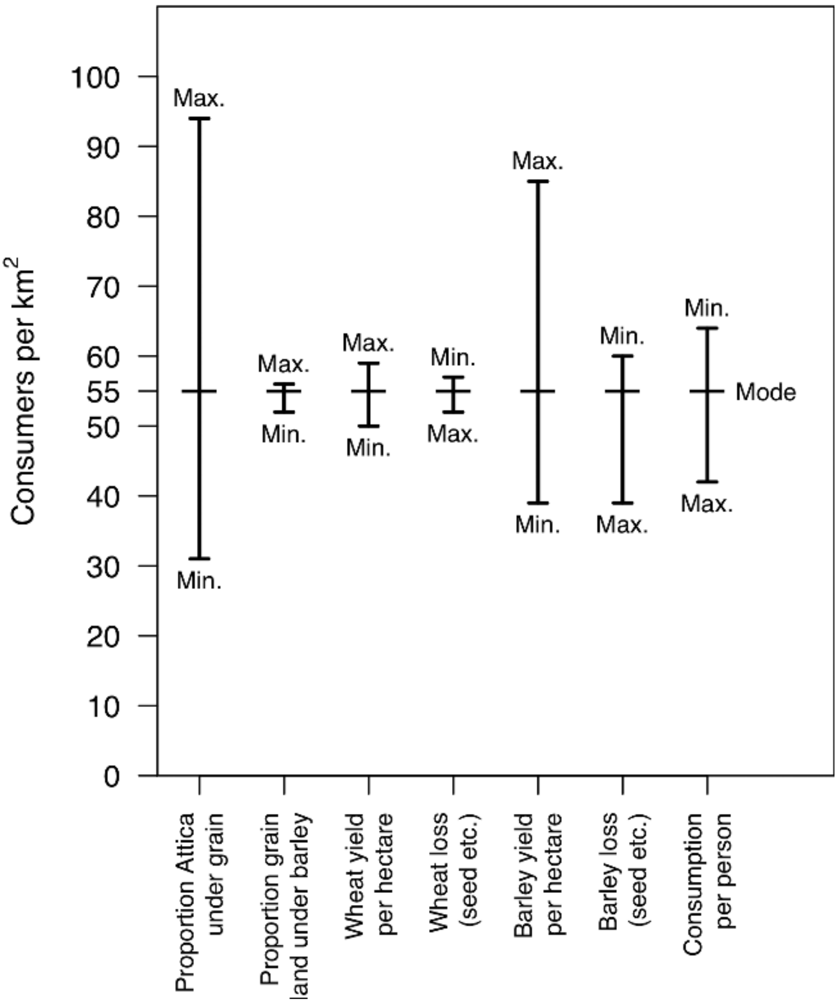

Figure 1.1 Garnsey's sensitivity analysis for Athenian land carrying capacity (redrawn from Garnsey 1988)

pessimistic values simultaneously, though this is easily done. (This is not the same, it bears noting, as setting all input variables to their minimum and maximum values respectively; since carrying capacity is inversely correlated with some variables, such as per capita consumption, those variables need to be set at their minimum to maximise the result.) The resulting range of 23–500,000 persons is an *interval estimate* for land carrying capacity, bracketing it between minimum and maximum possible values (given Garnsey's assumptions). Interval estimates are easily computed, but they tend to be uninformatively wide and can legitimately be narrowed, as the rest of this chapter will show.

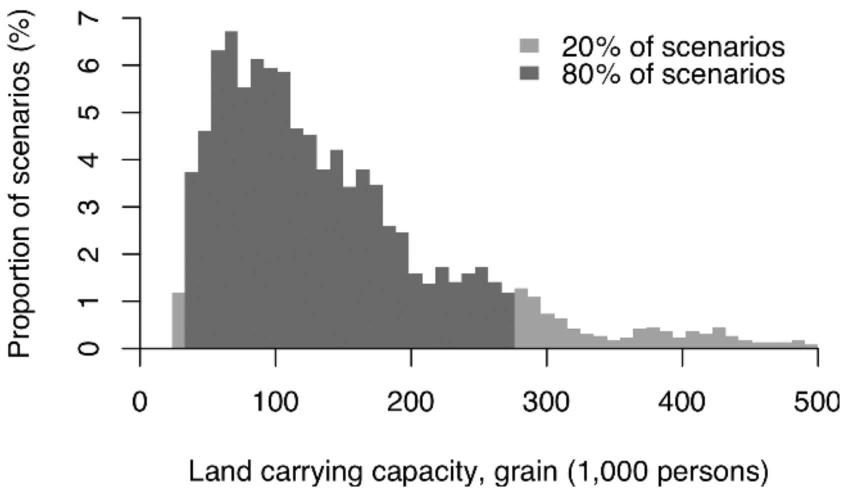

Figure 1.2 Frequency distribution of the output (land carrying capacity) across the 2,187 possible permutations of Garnsey's values

Garnsey's analysis was groundbreaking. Whereas almost all preceding scholars offered a simple point estimate or arbitrary range – answers that were easy to grasp if misleading in their confidence – Garnsey included an explicit quantification of uncertainty. His analysis identified the most important individual sources of uncertainty (the percentage of land under grain and average barley yield, whose effect dominated that of other variables in the sensitivity analysis). But it did not address the potential effect of multiple estimates being wrong simultaneously. Walter Scheidel observed that Garnsey's table of assumptions implied 2,187 (that is 3^7) different 'possible results'.[13] No human brain could intuit patterns across so many scenarios. In 1988, the reader was left to ponder the implications. Today – with cheap and ubiquitous computing power – it is easy to go further. It is a matter of moments to compute the result for all 2,187 possible permutations of Garnsey's seven variables (Figure 1.2).

This frequency distribution shows how often different results for carrying capacity occur across the 2,187 different scenarios, after the outcomes have been grouped into 'bins' of equal size (e.g. the first bin includes all outcomes in the range of 0–10,000 persons, the second includes all those in the range of 10–20,000). The vast majority of the scenarios are concentrated in the lower half of the theoretical range of 23–500,000 persons, with over 80 per cent of scenarios returning a result between 30,000 and 210,000 persons.

[13] Scheidel's addendum at Garnsey (1998: 198).

The frequency distribution – with its clustering of results and long tail of extreme values – begins to suggest a more productive and more robust way of conceptualising the problem. Rather than asking what the actual land carrying capacity was, we might ask how wide a range we need to allow in order to be confident that it includes the actual value. Rather than asking which of the results in Table 1.1 is 'right', we could ask which is most likely, and *how much* less likely the others are. In other words, it looks like probability might offer a way out of the impasse. But there are several hurdles to cross before we can pursue these ideas. First, these 2,187 discrete scenarios are far from exhausting the possibility space: we would have to accept that each of the input variables could have taken on *any* value between the minimum and maximum. Second, it ignores the fact that the input values used are not equally likely: by definition, Garnsey believed that his 'likely' values were more likely than his 'minimum' and 'maximum' values (whereas the three values are weighted equally in Figure 1.2). Last but not least, one might well recoil at the very use of probability here. Is it even legitimate to invoke the concept of probability when we are discussing a fixed quantity? We may not know the mean carrying capacity for the period, but it had a single, fixed, historical value – even if that precise value is unknown to us. This calls for some reflections on the nature of uncertainty and probability.

1.3 Knowledge, Uncertainty and Probability

The notion of applying probability to a fixed but uncertain historical quantity such as the mean carrying capacity of Classical Attica may seem problematic on first acquaintance, since it contravenes an intuition about the nature of uncertainty and probability. It will seem intuitively obvious to most that there are two fundamentally different types of uncertainty. Consider two different problems: predicting the outcome of a coin toss and estimating the distance between Cambridge and St Andrews. The uncertainty in the first case is (or rather appears to be) the result of a random process and cannot be reduced until the coin is flipped. The uncertainty in the latter case is merely a function of the limits of my knowledge and could be reduced if I had access to measurements. The first type of uncertainty is often termed *aleatory*, the second *epistemic*.[14] Intuitively, probability

[14] Alternative but equivalent terminologies distinguish between 'objective' and 'subjective' or 'irreducible' and 'reducible' uncertainty. See further Bedford and Cooke (2001: 33–4).

will seem a natural way of representing aleatory uncertainty (the chance of heads is 50 per cent) but it may seem an abuse to apply it to epistemic uncertainty. The distance is what it is; there seems no room for probability.

In fact, this intuitive distinction – and hence the association of probability with objective randomness – is far less secure than it first appears. There is universal agreement on how probabilities combine (they can be added and multiplied and must sum to unity), but what they represent remains a profound philosophical problem.[15] The two most important positions are the *frequentist* and the *subjectivist* interpretations. Readers with some familiarity with statistics may be familiar with the frequentist view, because it long dominated introductory textbooks. Frequentists see probability as an attribute of repeated events. On this interpretation, the probability of an event is the frequency with which the event would occur in a long sequence of similar trials. We say that the probability of heads on a coin toss is 50 per cent because, if the coin was flipped a very large number of times, the frequency of heads would approach 50 per cent. On the frequentist view, it would be nonsensical to speak of the probability that a historical quantity had some value, because it either had it or did not.

A rival interpretation, variously labelled 'subjectivist', 'personalist' and 'Bayesian', has a very different understanding of probability.[16] Its influence is evident in the proliferation of 'Bayesian' approaches to inference in a wide variety of fields, as can be seen from a cursory search for books, papers, projects and courses with that epithet. On this interpretation, the probability of an event such as the outcome of a coin toss is the degree of belief you have that it will occur, given all the information at your disposal.[17] Probability is a function not just of the world but also of a particular 'state of knowledge'. Since knowledge varies from observer to observer, probability is always subjective, in the sense of personal. Hence subjectivists speak of 'my' or 'your' probability rather than 'the' probability. Hence also Bruno De Finetti's famous dictum 'probability does not exist' – his

[15] 'The concept of probability is mathematically straightforward but philosophically puzzling' (Okasha 2002: 151). For a brief overview of the interpretations of probability, see Morgan and Henrion (1990: 48–50). We have opted for a brief and discursive discussion of probability here. For a more formal but still accessible overview of the mathematics of probability (aimed at archaeologists), see Buck et al. (1996: 47–65).

[16] Spiegelhalter (2011) provides a brief introduction. For a fuller but still very accessible discussion, see Lindley (2006).

[17] The probability is measured against some uncertainty standard, usually a game of chance. See Lindley (2006: 30–5) and Buck et al. (1996: 49–52).