

EPISTEMIC UNCERTAINTY, SUBJECTIVE PROBABILITY AND ANCIENT HISTORY¹

Myles Lavan

This is the author-final draft of an article forthcoming in the Journal of Interdisciplinary History

Periods from which few data survive pose a major challenge for history in the quantitative mode. There are many important historical quantities that can be estimated only on the basis of sparse and disparate information. Take the example of the population of the Roman empire. We do not have any census data for the empire as a whole. But that does not mean that we are in a position of complete ignorance. There is scattered information, both quantitative and qualitative, that allows us to reason in terms of likelihood. The methodological question is how to report our inevitably uncertain and subjective conclusions. Ancient historians tend to debate in terms of point estimates, disputing whether 54 m or 45 m is a better estimate of the population of the Roman empire in 14 CE. Estimates of this type are hard to interpret since they convey no information about the margin of error, which is often large and sometimes asymmetric. Historians sometimes resort to ranges as a concession to uncertainty, but their ranges tend to be arbitrary rather than being grounded in any measure of confidence or credibility. They signal the existence of uncertainty without offering any real guidance as to its magnitude. The problem of uncertainty becomes particularly acute when historians combine estimates for multiple uncertain quantities. It is increasingly common in ancient history for highly uncertain quantities to be estimated by expressing them as a function of other quantities about which we have better (if still limited) knowledge. For example, GDP has been modelled as a function of total population and average per capita consumption. But the total population is itself highly uncertain. The

¹ Thanks to Daniel Jew, Bart Danon and the participants in a workshop on probabilistic modelling in ancient history held in St Andrews in July 2017 for their help in refining the approach. I am also grateful to Michael Papathomas and Charles Paxton for discussions about uncertainty.

proliferating uncertainties as additional quantities are added to the equations pose a major challenge to the credibility of any point estimate.²

This paper discusses an alternative framework that allows for a more rigorous accounting of uncertainty. It formalises the probabilism that is already implicit in most historical reasoning by using probability as a measure of degree of belief. I open with a brief survey of work on the population of the Roman empire to illustrate the methodological problem. The following sections introduce the ‘subjective’ interpretation of probability as degree of belief and describe how it can be used as a tool of historical analysis. The final sections show how subjective probabilities can be used to combine uncertainties. Conceptualising uncertainty in terms of probability is a useful discipline in itself, because it forces historians to assess the uncertainty in their estimates, but its greatest value lies in the scope for aggregation.³

The probabilistic approach discussed here would be familiar to scholars and practitioners in future-oriented fields because it informs much current work in forecasting, risk assessment and decision analysis. The fact that it is unfamiliar to many historians, even those engaged in quantitative analysis, can probably be attributed to a mistaken assumption that the challenges that uncertainty poses for history are qualitatively different from those faced by other disciplines. Ancient historians may believe that they face unique challenges in the need to base quantification on subjective assessments of what is likely, rather than hard data. But the problem of reliance on subjective assessment is shared by many other fields. Consider an observation by the authors of

² For the state of the art in the estimation of Roman GDP, see Walter Scheidel and Steven J. Friesen, ‘The size of the economy and the distribution of income in the Roman empire’, *Journal of Roman Studies* 99 (2009), 61-91.

³ An earlier paper applied this approach to a long-standing problem in ancient history: Myles Lavan, ‘The spread of Roman citizenship, 14-212 CE: Quantification in the face of high uncertainty’, *Past and Present* 230 (2016), 3-46. My collaborator Daniel Jew is applying it to the problem of Athenian population in *The probable past: Agriculture and carrying capacity in ancient Greece* (Cambridge, forthcoming). This paper expands on the theoretical premises, particularly the underpinning conceptions of epistemic uncertainty and subjective probability.

a textbook on risk analysis: ‘Probabilistic risk analysis treats events with a low intrinsic rate of occurrence, and large amounts of data are seldom available. Since its inception, *expert opinion in the form of subjective probabilities* has been a dominant source of data for failure probabilities’. Surprising as it may seem, our predicament is shared by many other fields. My suggestion is that ancient historians might profitably learn from the techniques that other disciplines have developed to manage epistemic uncertainty, i.e. uncertainty that arises from the limits of our knowledge.⁴

THE POPULATION OF THE ROMAN EMPIRE

Population is relevant to a wide range of questions in the social and economic history of the Roman empire. Unfortunately the population data that the Roman state collected through regular provincial censuses are almost entirely lost to us. The most significant exception is a

⁴ Quotation from Tim Bedford and Roger Cooke, *Probabilistic risk analysis: Foundations and methods* (Cambridge, 2001), 191 (my emphasis). Others have productively applied Monte Carlo simulation to historical questions. Donald Schaefer and Thomas Weiss, ‘The Use of Simulation Techniques in Historical Analysis: Railroads versus Canals’, *Journal of Economic History* 31 (1971), 854-884 was a pioneering application of the method to the relative profitability of railways and canals in the US ca. 1840, an example followed by several subsequent papers in modern economic history. The method is widely used in historical demography, as in the CAMSIM microsimulation of kin sets: J. E. Smith and J. Oeppen, ‘Estimating numbers of kin in historical England using demographic microsimulation’, in D. S. Reher and R. S. Schofield (eds.), *Old and new methods in historical demography* (Oxford, 1993), 413-25. It is increasingly being adopted by archaeologists to manage uncertainty about chronologies (see n. 12). It has also been applied in an ad hoc way to miscellaneous problems in ancient history. Ellen Janssen, et al., ‘Fuel for debating ancient economies. Calculating wood consumption at urban scale in Roman Imperial times’, *Journal of Archaeological Science: Reports* 11 (2017), 592-9 use Monte Carlo simulation to estimate fuel consumed by pottery production and baths in Roman Sagalassos (but revert to traditional interval analysis for their estimate of total fuel consumption and its impact on local woodland, which is precisely the type of problem where epistemic uncertainties could usefully be represented as subjective probabilities). Yet none of these contributions ground the method in the Bayesian conception of uncertainty and probability. Instead, the uncertainty in question tends to be conceptualised as aleatory (i.e. related to variability or random processes). As a result, the scope for the method to be redeployed to situations of gross epistemic uncertainty has been neglected.

series of census figures for the Roman citizen population through to 14 CE. Since citizens were still concentrated in Italy, this statistic is a reasonable proxy for the population of Italy. But a major ambiguity as to the unit of measure (all persons or just adult males) and further uncertainties involved in extrapolating from the count of citizens to the population of Italy leave even the population of Italy in 14 CE contested: probably somewhere in the region of 6 m, but conceivably high as 14 m. There are a few even more problematic population data for other sub-regions, such as north-west Iberia and Egypt. But we are otherwise dependent on estimates of carrying capacity and crude judgements about regional variation in population density and about population levels relative to the late medieval and early-modern periods.⁵

In 1886 Julius Beloch, the pioneer of the modern study of the Roman population, reckoned a total of 54 m inhabitants in the empire in 14 CE. He subsequently revised this up to 70 m in a later and briefer survey of the topic (assuming higher populations for Gaul, the Balkans and North Africa) and also suggested – in even more cursory fashion – that the population grew to a peak of around 100 m by the end of the second century. Where Beloch himself came to view his 1886 estimate as a minimalist figure, the next most important intervention took the opposite view. In 1978 Colin McEvedy and Richard M. Jones constructed an extraordinarily ambitious model of the evolution of the global population on a country-by-country basis from 400 BCE to the present. They had to significantly reduce Beloch's 1886 estimates for several sub-regions, particularly in Anatolia and the Levant, in order to arrive at population levels that seemed plausible in the light of the long-term history of those sub-regions, though it is worth noting that their conception of what was plausible was shaped by their assumptions about the amplitude of the second century peak in population relative to late-medieval and early modern periods and by

⁵ On the problem of the population of Italy see Walter Scheidel, 'Roman population size: The logic of the debate', in L. De Ligt and S. J. Northwood (eds.), *People, land, and politics: Demographic developments and the transformation of Roman Italy 300 BC-AD 14* (Leiden, 2008), 17-70.

their reconstructions of late-medieval and early modern populations. They put the peak population of the empire ca 200 CE at just 46 m – 10% lower than Beloch’s estimate for 14 CE and less than half his later estimate for peak population. Their work has provided the starting point for most subsequent research in the field of Roman history. Reverting closer to Beloch’s figures for Anatolia and the Levant and assuming slightly higher long-term growth, Bruce Frier proposed a peak population of 61m in 164 CE (a more plausible date for peak population, given a major epidemic in the 160s, the ‘Antonine plague’). More recently, Walter Scheidel has suggested a peak of 59-72 m in 165 CE, allowing for some uncertainty and assuming slightly higher populations in the north-western provinces.⁶

Without going into the evidence in more detail here, I merely want to observe that the debate has been conducted in a way that obscures the question of uncertainty. Most interventions have taken the form of point estimates without any real discussion of the margin of error. Scheidel’s range at least signals the problem of uncertainty, but it is far from clear how it should be interpreted. Does he mean to rule out the possibility of a population less than 59 m or higher than 72 m? I do not believe we have the evidence to categorically disprove either McEvedy and Jones’ minimalist estimate of 46 m or (late) Beloch’s maximalist 100 m – though both do now seem much less likely than a figure in the 60s.

All work on this problem – and many similar topics in ancient history – depends on an epistemology that is implicitly or explicitly probabilistic. What historians from Beloch to Scheidel have been doing is assessing the relative likelihood of different possible populations given the

⁶ Julius Beloch, *Die Bevölkerung der griechisch-römischen Welt* (Leipzig, 1886), 507; Julius Beloch, ‘Die Bevölkerung im Altertum’, *Zeitschrift für Sozialwissenschaft* 2 (1899), 618 and 620; Colin McEvedy and Richard M. Jones, *Atlas of World Population History* (London, 1978), 22; Bruce W. Frier, ‘Demography’, in A. K. Bowman, P. Garnsey and D. Rathbone (eds.), *The Cambridge Ancient History. Volume 11: The High Empire, AD 70-192* (Cambridge, 2000), 812-14; Walter Scheidel, ‘Demography’, in W. Scheidel, I. Morris and R. P. Saller (eds.), *The Cambridge economic history of the Greco-Roman world* (Cambridge, 2007), 45-9.

evidence available to them and then reporting the possibility that seems most likely. Probabilism of this sort is essential to historical reasoning, but it can lead to confusion and misunderstanding if it is left unformalised. It can be disciplined and deployed to better effect if the probabilism is made explicit and taken to its logical conclusion.

KNOWLEDGE, UNCERTAINTY & PROBABILITY

The central proposition of this paper is that our uncertainty about historical quantities like the peak population of the Roman empire can – indeed should – be represented as probabilities. This is likely to seem troubling on first acquaintance, since it contravenes some intuitions about the nature of uncertainty and probability, though those intuitions will turn out to be mistaken. 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 the result of the limits of my knowledge and could be reduced if I had access to better measurements. The first type of uncertainty is often termed *aleatory*, the second *epistemic* (or, alternatively, objective and subjective uncertainty respectively). Intuitively, probability will seem a natural way of representing aleatory uncertainty (the chance of heads is 50%) 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 association of probability with objective randomness is far less secure than it appears. There is universal agreement on the mathematics of probability, but what they represent is a profound and unresolved philosophical problem. The two most important positions are the *frequentist* and the *subjective* interpretations. Anyone with some familiarity with statistics will 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. Hence we say that the probability of heads on a coin toss is 50% because the frequency of heads would approach 50% in a very long series of tosses. On this frequentist view, it would be nonsensical to speak of the probability that that some historical quantity had some value, because it either had it or did not.⁷

A rival interpretation, variously labelled ‘subjective’, ‘personalist’ and ‘Bayesian’, has a very different understanding of probability – one that should be of significant interest to historians. Its influence is evident in the proliferation of ‘Bayesian’ approaches to inference in a wider variety of fields. On the subjective interpretation, what probability represents is my degree of belief, given all the information at my disposal. The qualification means that 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 particular sense of ‘personal’. Hence subjectivists speak of ‘my’ or ‘your’ probability rather than ‘the’ probability. Hence also De Finetti’s famous dictum that ‘probability does not exist’ – his point being that there is no such thing as ‘objective’ probability. This may seem a radical view, but it rests on the insight that there is an irreducible subjective element to all uncertainty. The apparently obvious distinction between aleatory and epistemic uncertainty is in fact deeply problematic. Phenomena that we think of as random are often the result of processes that are in fact deterministic but chaotic, in the technical sense that the outcome is highly sensitive to small changes in the initial conditions. My uncertainty about the outcome of a coin toss, for example, is really epistemic uncertainty about the initial conditions and how they determine the behaviour of the coin. Most randomness is thus a result of an observer’s lack of knowledge, not inherent in the world itself. A second qualification to our intuitions is that my assessment that the probability of heads on

⁷ For a brief overview of the interpretations of uncertainty, see M. Granger Morgan and Max Henrion, *Uncertainty: a guide to dealing with uncertainty in quantitative risk and policy analysis* (Cambridge, 1990), 48-50.

any coin toss is 50% depends on an unstated assumption that the coin in question is unbiased. Yet that assumption may not be shared by other observers, and may be mistaken. The fundamental insight of the subjectivists is that the probabilities that we conventionally think of as objective (that is as a property of the world) are always based on assumptions about the generating mechanism (the coin toss) and hence are subjective.⁸

Fields concerned with forecasting, risk analysis and decision analysis have long recognised that predicting the future always involves epistemic as well as aleatory uncertainty and many of them have embraced the use of subjective probabilities as a measure of epistemic uncertainty. Questions have been raised about whether probabilities are adequate for assessing all forms of uncertainty, and various alternative mathematical frameworks have been suggested (including intervals, ‘imprecise probabilities’, ‘possibilities’ and ‘belief functions’). But probability remains the dominant conceptual tool for representing epistemic uncertainty in a wide range of fields engaged in describing the present and predicting the future. For its strongest proponents, indeed, it is the only rational framework for assessing uncertainty. ‘If you want to handle uncertainty, then you must use probability to do it, there is no choice.’ ‘It is very firmly our opinion that the uniquely suitable representation of uncertainty, whether aleatory or epistemic, is probability.’ On this subjective interpretation, probability is precisely the right form in which to represent uncertainty about the past.⁹

⁸ Bruno De Finetti, *Theory of Probability: A Critical Introductory Treatment* (London and New York, 1974). For a brief introduction to the subjective interpretation, see David Spiegelhalter, ‘Quantifying uncertainty’, in L. Skinnis, M. Scott and T. Cox (eds.), *Risk (Darwin College Lectures)* (Cambridge, 2011), 17-33. Dennis V. Lindley, *Understanding uncertainty* (Hoboken, 2006) provides a fuller but still very accessible overview. Morgan and Henrion, *Uncertainty*, 57-60 and 62-4, Bedford and Cooke, *Probabilistic risk analysis*, 33-4 and Terje Aven, *Foundations of risk analysis* (Chichester, 2012), 47-9 illustrate the personalist interpretation of randomness.

⁹ Quotations from Lindley, *Understanding uncertainty*, 239 and Anthony O’Hagan and Jeremy E. Oakley, ‘Probability is perfect, but we can’t elicit it perfectly’, *Reliability Engineering & System Safety* 85 (2004), 239-48, 247. Some theorists

The discussion that follows is grounded in the subjective interpretation of probability. The uncertainties that historians face are clearly epistemic. They can be represented as probabilities, but those probabilities must be understood as subjective. A probability is meaningful only in relation to a particular state of knowledge, i.e. a body of evidence. The probabilities represent the historian's degree of belief in different possibilities given that state of knowledge, not some objective randomness. Hence I will write of 'assigning' not 'estimating' probabilities. I will also use the term *beliefs* to foreground the subjectivity that is inherent in the encounter with uncertainty. By beliefs I mean a set of evidence-based probabilistic judgements about historical quantities and events. And I will always qualify this as 'my beliefs' or 'the historian's beliefs' to emphasise that the state of knowledge and hence the probabilistic beliefs based on it are always personal. Your beliefs about the past may differ from mine if you have access to more (or less) information than I have or if you interpret it differently. But describing beliefs in terms of probabilities will clarify our differences and facilitate the dialectic revision of our beliefs.¹⁰

accept the distinction between the two types of uncertainty and account for them separately, but they use probability in both cases. See e.g. Colin Howson and Peter Urbach, *Scientific reasoning: The Bayesian approach* (Chicago and La Salle, 1993), 24-5, Stan Kaplan, 'Formalisms for Handling Phenomenological Uncertainties: The Concepts of Probability, Frequency, Variability, and Probability of Frequency', *Nuclear Technology* 102.1 (1993), 137-42 (introducing the 'probability of frequency' framework in which aleatory uncertainty is represented by frequentist probabilities and epistemic uncertainty is represented by subjective probabilities) and David Vose, *Risk analysis: A quantitative guide* (Chichester, 2008), 47-9. Others have fully embraced the subjectivist perspective that all uncertainty contains an epistemic element and that all uncertainty can be represented by subjective probabilities. See e.g. Lindley, *Understanding uncertainty* and Aven, *Foundations of risk analysis*. On alternatives to subjective probability, see Franz Huber, 'Formal Representations of Belief', in Edward N. Zalt, ed., *The Stanford Encyclopedia of Philosophy* (Spring 2016 Edition), <https://plato.stanford.edu/archives/spr2016/entries/formal-belief/>.

¹⁰ On this sense of 'belief', see Lindley, *Understanding uncertainty*, 12-13. The use of subjective probability should be distinguished from formal Bayesian inference, which involves not just subjective probability but also the use of Bayes' theorem to update *a priori* probabilities given data. Caitlin E. Buck, et al., *Bayesian Approach to Interpreting*

PAST, PRESENT AND FUTURE

One might yet wonder whether an approach developed for managing uncertainty about the future is really applicable to the past. After all, the past has actually happened whereas the future is still open. But both are equally uncertain from the perspective of an observer in the present. Depending on one's philosophical position, one might hold that uncertainty about the past is purely epistemic, whereas uncertainty about the future involves an aleatory element. But all the uncertainties can be represented as probabilities.¹¹

Some archaeologists and historians have already embraced probability for the formal representation of epistemic uncertainty about dating. Absolute chronology construction (such as carbon-dating) and phylogeny (both genetic and linguistic) are two subfields in which the use of Bayesian methods of inference based on subjective probability is now familiar. But these technical fields are remote from the experience of most archaeologists and historians. A few archaeologists have also applied a probabilistic approach to the more quotidian exercise of dating artefacts based on established chronologies of artefact types. For example, the date of deposition of a sherd of African Red Slipware of Hayes form 1 – a pottery type known to have been in use

Archaeological Data (Chichester, 1996) explores the potential of Bayesian approaches in archaeology; Caitlin E. Buck and Bo Meson, 'On being a good Bayesian', *World Archaeology* 47.4 (2015), 567-84 is a recent review of progress to date. But there are many problems where we lack meaningful data with which to update our subjective beliefs and can only make the first step, quantifying informed but subjective knowledge as probabilities (analogous to the formulation of 'priors' in the Bayesian framework).

¹¹ Hence subjectivists see no qualitative difference between uncertainty about past, present and future. See e.g. Lindley, *Understanding uncertainty*, 2-7, whose twenty examples deliberately conflate the three timeframes, and the similarly deliberate conflation at Buck, et al., *Bayesian Approach*, 54: 'The view adopted in this book is that assessments of probability are subjective and made in the light of experience: there is no difference in kind between the bookmaker's estimate of odds, the architectural historian's view of a date for a medieval building, the doctor's diagnosis, the archaeologist's opinion about the provenance of a pot, or the uncertainty in a scientist's estimate of the distance of the sun from the earth.'

during the period 50-80 CE – can be represented as a probability distribution over that period (using a uniform, normal or any other distribution depending on one’s assumptions about the processes of production and deposition). Where these archaeologists use probability to represent the epistemic uncertainties about the dates of hundreds or thousands of individual sherds or other artefacts, this paper generalises this approach to epistemic uncertainty about quantities other than dates.¹²

FROM BELIEFS TO PROBABILITY DISTRIBUTIONS

My proposition is that the uncertainty about historical quantities such as the population of the Roman empire is best represented as a probability distribution. To illustrate the concept, I return to the problem of Roman population. In the probabilistic framework, the question to ask is not ‘what was the population of the empire’ (which would elicit a point estimate with spurious precision) but ‘what is my degree of belief, given the available evidence, in different possible

¹² On absolute chronologies, see Buck and Meson, ‘On being a good Bayesian’ and other articles in *World Archaeology* 47.4, a special issue on Prehistoric Bayesian Chronologies. For detailed discussions of the use of probability to date artefacts, see D. J. Carlson, ‘Computer analysis of dated ceramics: estimating dates and occupational ranges’, *Southeastern Archaeology* 2 (1983), 8-20, John M. Roberts, et al., ‘A method for chronological apportioning of ceramic assemblages’, *Journal of Archaeological Science* 39 (2012), 1513-1520, Enrico R. Crema, ‘Modelling Temporal Uncertainty in Archaeological Analysis’, *Journal of Archaeological Method and Theory* 19 (2012), 440-61, Rinse Willet, ‘Experiments with diachronic data distribution methods applied to Eastern Sigillata A in the eastern Mediterranean’, *Herom* 3 (2014), 39-69 and M.J. Baxter and H.E.M. Cool, ‘Reinventing the wheel? Modelling temporal uncertainty with applications to brooch distributions in Roman Britain’, *Journal of Archaeological Science* 66 (2016), 120-7. Practical applications include Elizabeth Fentress and P. Perkins, ‘Counting African Red Slip Ware’, in A. Mastino (ed.) *L’Africa Romana: Atti del V Convegno di studio Sassari, 11-13 dicembre 1987, Sassari, Dipartimento di Storia Università degli Studi di Sassari* 1988), 205-14, Martin Millett, ‘Pottery: population or supply patterns? The Ager Tarraconensis approach’, in G. W. W. Barker and J. Lloyd (eds.), *Roman Landscapes: Archaeological Survey in the Mediterranean Region* (London, 1991), 18-26, Andrew Wilson, ‘Approaches to quantifying Roman trade’, in A. K. Bowman and A. Wilson (eds.), *Quantifying the Roman economy: Methods and problems* (Oxford, 2009), 213-49.

values of the population' (which elicits a probability distribution). It is a matter of assessing the relative likelihood of different possibilities.¹³

Without going into detail here, I agree with Frier and Scheidel that a total population in the 60s seems more likely than a lower or higher figure. But identifying a most likely value or region is only part of the process of estimation. I also need to assess how wide a range is possible. This is a more difficult question to answer, not least because past scholarship has tended to focus almost exclusively on identifying a most likely value, with the result that much less attention has been paid to the equally important task of establishing upper and lower limits. For the purposes of this discussion, I will work with 40 m as a minimum since there is very limited scope to lower McEvedy and Jones' already minimalist figure of 46 m. Establishing a ceiling is much more difficult. Beloch's suggestion of a population of around 100 m has seemed implausibly high to most recent scholars. The combined population of the former territories of the empire probably only just exceeded that level by 1800. I will take it as an effective maximum, though the question deserves further consideration.¹⁴

Uncertainty about a continuous quantity (i.e. one that can take any value within a range) is represented by a *probability density function* (PDF). To facilitate computation, the PDF implicit in our beliefs is approximated by a known mathematical distribution. The simplest available distribution is the *uniform* distribution, which assigns an equal probability to all possible values between a minimum and maximum value. Figure 1 illustrates a uniform PDF for the peak population of the empire, assuming that any value between 40 m and 100 m is equally likely. The obvious objection to this as a representation of the uncertainty about Roman population is that it fails to take any account of my belief that a value around 65 m is much more likely than one

¹³ This paper focuses on uncertain *quantities*. There is a second type of uncertainty, in which the uncertainty concerns not the value of a quantity but the truth of a proposition. These uncertainties are termed *events* in probability theory. Much uncertainty in history concerns events in this technical sense. Although this paper limits the discussion to quantities, the framework can also accommodate uncertainty about events, by assigning a probability to the proposition that the event is true. On the concept of event, see Lindley, *Understanding uncertainty*, 12.

¹⁴ The estimate of the 1800 population of the former territories of the Roman empire is calculated from the data in McEvedy and Jones, *Atlas of world population history*.

around 40 m or 100 m. The uniform distribution has potential value in representing total ignorance or for introducing a conservative element into an analysis, but it should be clear that it would be a very crude representation of the state of knowledge in this case.¹⁵

A simple but much better alternative is the *triangle* distribution, which introduces a third parameter, the most likely value (i.e. the point of highest probability). Figure 1 illustrates a triangle distribution that represents 65 m as the most likely value. (It is perhaps worth noting that the distribution is uniquely determined by its three parameters: there is only one triangle that starts at 40, peaks at 65, ends at 100 and encloses an area of 100%.) Though still crude, it is a much better representation of the way my degree of belief in possible values decreases as I move away from the most likely value towards the minimum and maximum possible values.¹⁶ The graphic representation also reveals an aspect of the problem that is obscured by a focus on the most likely value: the uncertainty is asymmetric. The possible range extends further above 65 m than it does below. This makes the most likely value a biased estimator of the actual value. In this case the asymmetry is relatively minor, but such asymmetry is common in epistemic uncertainty and can be more pronounced. The probabilistic approach has a solution to this problem that I discuss towards the end of the paper.

¹⁵ The interpretation of a PDF deserves a brief note. With continuous variables, only intervals can be assigned a discrete probability. The probability of an interval is represented by the area under the probability density function within that interval. The total area under a PDF always sums to 100%. Strictly speaking, Roman population is a discrete rather than a continuous variable (since there are no fractional persons), but the number of possibilities (in the tens of millions) is so large that it can be represented as continuous for convenience. I have opted for a brief and discursive discussion of probability, intended only to pique the interest of historians. For a formal introduction to probability and probability distributions, see Buck, et al., *Bayesian Approach*, 47-65 or any textbook.

¹⁶ On the use of triangle distributions to represent epistemic uncertainty, see Vose, *Risk analysis*, 403 and Morgan and Henrion, *Uncertainty*, 96.

The triangle distribution remains imperfect in at least two respects: it exaggerates how quickly my degree falls off in the immediate area of the most likely value (I do not believe that a value of 65 m is so much more likely than a value of 60 m) and it assigns too high a probability to extreme values, especially in the 90-100 m range. A curve with attenuated tails would be better. There are numerous distributions that could serve that purpose, including the beta, gamma, Weibull and Burr distributions. (The more familiar normal distribution, though often used to represent aleatory uncertainty arising from variation and measurement error, is unsuitable for this context since it is strictly symmetrical – whereas the uncertainty here is clearly asymmetrical – and extends infinitely in both directions.) Fitting these distributions to one's beliefs can be computationally complex, since the parameters that define them are abstract quantities without a real world interpretation. The most intuitive to manipulate is the *PERT* distribution (a special case of the beta distribution), whose three parameters have the same interpretation as those for the triangle distribution: minimum, most likely and maximum values). Figure 1 illustrates the use of a PERT distribution to represent my beliefs. It is a marginally better approximation of my beliefs than the Triangle distribution, because of its rounded peak and more attenuated tails. But the effect on most calculations is likely to be minimal, so the computationally simpler Triangle distribution may be adequate in many cases. Practitioners in many other fields have found triangle distributions the most convenient way of representing subjective probabilities in situations of epistemic uncertainty. In any case, there is a wide range of options available and there is no obligation to use the same distribution for all uncertain quantities. Each can be modelled with the distribution whose shape best fits my beliefs.¹⁷

¹⁷ For further discussion of these and other distributions, see Vose, *Risk analysis*, 401-10 and Paul H. Garthwaite, et al., 'Statistical Methods for Eliciting Probability Distributions', *Journal of the American Statistical Association* 100(470) (2005), 688-9.

This exercise of assigning probabilities is not as outlandish as it may at first appear. The procedure of encoding subjective beliefs as probability distributions has become well-established in other disciplines. Even in fields with much better data, estimation often entails an irreducible element of subjective judgement. Forecasting and risk models regularly include quantities whose values can only be assigned on the basis of the subjective judgement of experts. The problem is so widespread that there is a whole literature devoted to the ‘elicitation’ of expert opinion in the form of probability distributions.¹⁸

It is also worth noting that the procedure of assigning a probability distribution is merely making explicit reasoning that is often implicit in historical argument. When historians like Beloch proposed a point estimate for a quantity such as the population of the empire, they were presumably reporting the value they judged most likely – i.e. the peak of their probability distribution for the quantity. Assertions that other estimates are less likely to be correct imply that they assigned a lower probabilities to those values, while claims that other suggestions are implausible or impossible should mean that their probability distributions were at or near zero at those values (otherwise their rhetoric is misleading). Making the probabilism explicit would help to clarify the claims that are being made and usefully focus attention on **the degree and direction** of uncertainty, an aspect of estimation that is too often neglected.

At this point, it is worth reiterating what the probability distribution represents. I am not trying to estimate an objective probability that exists in the world, but rather representing my degree of belief in different possible population values, given the information available. It might be useful to conceptualise the exercise not as an attempt to estimate the quantity but rather as representing

¹⁸ On widespread need to rely on subjective judgements in future-oriented fields, see e.g. Anthony O'Hagan, et al., *Uncertain judgements: Eliciting experts' probabilities* (Chichester, 2006) ch. 5, Vose, *Risk analysis* ch. 14. On the elicitation of expert opinion, see Garthwaite, et al., ‘Statistical Methods for Eliciting Probability Distributions’ for a brief overview and O'Hagan, et al., *Uncertain judgements* for a fuller synthesis of research in the field.

the uncertainty about it. I give the last word to Stan Kaplan, a pioneer in the development of probabilistic approaches to risk analysis in nuclear engineering: ‘People often think that putting forth an uncertainty curve is somehow difficult, compared with giving a single number or “point estimate”. It becomes much easier if we remind ourselves that probability curves “do not exist”, as De Finetti said. That is, they do not exist in the real world in such a way that we can try to “find” them. They exist only in our minds. *They are only a language in which we express our state of knowledge or state of certainty.* With this understanding, it is easy to put forth a curve (fat, if necessary) to express our uncertainty. What is far more difficult is to put forth a single number that people are going to believe and use for design and regulatory decisions, as if it were gospel truth.’¹⁹

COGNITIVE BIASES

Historians should be aware of a number of cognitive biases that are known to affect our probability judgements. Two pose a particular challenge to historical reasoning. It is worth emphasising at the outset that that they affect all attempts to estimate uncertain quantities, not just the formal probabilistic approach developed here.

The most significant bias is *overconfidence*. It is well established that people estimating uncertain quantities tend to underestimate the range of values that are plausible. This happens because values towards the upper and lower limits of the possible range are assigned a probability too close to zero. In other words, we are too quick to rule out extreme values as impossible when we should be assigning them low but significant probabilities. This has been demonstrated by psychological experiments with ‘almanac’ questions. Participants are asked to estimate quantities that are known or will shortly be known (such as prices or sports results) by reporting a range that is broad enough for them to be e.g. 98% confident that it contains the actual value. If their probability judgements were well calibrated, one would expect around 98% of the responses to

¹⁹ Kaplan, ‘Formalisms for Handling Phenomenological Uncertainties’, 141 (my emphasis).

contain the actual value, but repeated experiments have found that the success rate is much lower, revealing a systematic tendency to under-estimate the degree of uncertainty (hence the term overconfidence). It has also been found that the tendency to overconfidence is positively correlated with the difficulty of the problem, so that overconfidence is even more pronounced in more difficult estimation problems (the so-called ‘hard-easy’ effect). This is a major issue for historians engaged in the difficult process of estimating uncertain quantities on the basis of very limited information. It suggests that scholars who have worked on the population of the empire are likely to have been too quick to rule out relatively low or high values as implausible. The range of plausible values is probably wider than they have suggested.²⁰

A second important bias, the *anchoring* effect, arises from one of the ‘heuristics’ that Daniel Kahnemann and Amos Tversky identified in their pioneering work on the cognitive short-cuts that the human mind relies on to make decisions using limited information. They showed that our estimates of uncertain quantities are influenced by any values that are suggested to us before we consider the problem ourselves. This is a result of a heuristic that operates by first evaluating a proposed estimate as high or low and then correcting it – what Kahnemann and Tversky termed ‘judgement by anchoring and adjustment’. As a rule, the correction tends to be too small. Again this has been shown through experiments with ‘almanac’ questions. For example, subjects were asked to estimate the percentage of UN countries that were in Africa after first being asked whether the percentage was higher or lower than a random number (either 10 or 65). Those who

²⁰ The study of judgement under uncertainty was pioneered by Amos Tversky and Daniel Kahneman, ‘Judgment under uncertainty: Heuristics and biases’, *Science* 185 (1974), 1124-31. O’Hagan, et al., *Uncertain judgements* ch. 3 and Garthwaite, et al., ‘Statistical Methods for Eliciting Probability Distributions’, 682-4 and 685 (cf 685 on over-confidence) provide good syntheses of research in the field with a focus on the implications for the formulation of expert knowledge in probabilistic form. For more on over-confidence, see Ward Edwards, et al., *Advances in decision analysis: From foundations to applications* (Cambridge and New York, 2007), 143-4 and Garthwaite, et al., ‘Statistical Methods for Eliciting Probability Distributions’, 685

started from 10 arrived at significantly lower estimates than those that started from 65, even though they were aware that the starting value was a random number unrelated to the problem. This is another significant problem for historians, since it means that new estimates are likely to be anchored to previously published estimates, regardless of their quality. Even if we judge an existing estimate to be mistaken, we are liable to remain close to it when forming our own estimates. Each of the series of estimates of Roman population has probably been anchored to its predecessor. It is perhaps revealing that the single largest swing in the modern history of the debate was proposed by Beloch himself, when he updated his own estimate for 14 CE from 54 to 70 m. He was presumably more acutely aware of the uncertainties in his own estimate than later readers could be. Subsequent revisions have tended to be more modest. McEvedy and Jones reduced Beloch's estimate by around 14 m in their estimate for 1 CE. Frier nudged their estimate for 1 CE up by just 5 m, while Scheidel adjusted Frier's estimate for 164 by just 4 m. Given the existence of the anchoring effect, it is possible that these corrections have been too small.²¹

There has been considerable research on our capacity to suppress these biases. The most fruitful approach involves a training process where individuals are repeatedly asked to estimate a quantity and then confronted with the actual value. This has proved effective in calibrating the probability judgements of forecasters such as meteorologists. But it is obviously of little use to historians, since we rarely have the opportunity to compare our estimates to actual values. The best that we can do is be aware of the biases affecting our judgement and do our best to compensate for them. A formally probabilistic approach provides the best framework for doing so.

²¹ For the experiment, see Tversky and Kahneman, 'Judgment under uncertainty: Heuristics and biases'. See further O'Hagan, et al., *Uncertain judgements*, 47-9, Garthwaite, et al., 'Statistical Methods for Eliciting Probability Distributions', 682.

COMBINING UNCERTAINTIES

Conceptualising our beliefs about uncertain quantities as probability distributions is a useful intellectual discipline. Ancient historians are accustomed to dispute in terms of best estimates, asserting that a particular value or range is ‘most likely’, without stating how confident they are that the actual value was close to their proposed value or in their proposed range (perhaps not very confident at all). Rival estimates are rejected as ‘implausible’, without any quantification of how improbable they are (conceivably not much less probable than the new proposal). The rigour of thinking in terms of probability distributions forces the historian to think about how wide a range is plausible, rather than fixating on the most likely value alone. It is thus a valuable exercise in its own right. But its real value lies in the aggregation of uncertainties. One key advantage of probabilities is that they are easy to combine mathematically.

It is now routine in many future-oriented fields to combine uncertainties using Monte Carlo simulation. This technique involves three steps. (1) A mathematical model is constructed to represent the quantity of interest as a function of several better-understood quantities. For the purposes of the model, all uncertain quantities are represented as random variables. The quantity of interest is the output variable; the better-understood quantities are input variables. (2) The observer’s beliefs about the input variables are represented as probability distributions. This includes the assessment of epistemic interdependence between input variables (discussed below). (3) A series of random scenarios are generated by randomly selecting a set of input values from the probability distributions over the input variables. The output the model produces in each scenario can be regarded as a random sample from the probability distribution over the output variable. As the number of scenarios increases, the distribution of the output values in the sample will converge on the underlying probability distribution of the output variable.²²

²² Vose, *Risk analysis* is an excellent practical guide to Monte Carlo simulation.

The Monte Carlo method was first developed in the context of nuclear engineering during and after the second world war. It is now widely used for the purposes of forecasting, risk assessment and decision analysis. Perhaps because it is most often applied in contexts where the uncertainty may appear to be aleatory, its relevance to historical problems where the uncertainty is clearly epistemic has been neglected. Yet it offers historians a very useful tool for aggregating epistemic uncertainties. I have shown elsewhere that it offers a potential solution to a hitherto intractable problem in Roman history, namely estimating the proportion of the population that had Roman citizenship before Caracalla's universal grant of 212/213 CE. Even though the mechanisms by which Roman citizenship was disseminated were relatively well understood, quantification seemed impossible because of the many uncertainties involved. Any estimate of the proportion of persons who were citizens in 212 would require estimates for several highly uncertain quantities (including the total population of the empire). A traditional 'best estimate' based on most likely values for each of the input quantities could never hope to command credibility because of the proliferating uncertainties. A probabilistic approach using Monte Carlo simulation made it possible to account for all the component uncertainties and assess the aggregate uncertainty about the prevalence of citizenship.²³

INTERPRETING THE RESULTS

The output of a Monte Carlo simulation is most intuitively grasped through a histogram. Figure 2 shows the result of my simulation of the spread of Roman citizenship. In each random scenario, the model calculated the prevalence of citizenship in 212 CE, expressed as a percentage. The histogram shows how often different prevalences occurred in a sample of 50,000 scenarios, grouped into 1% intervals. The shape of the distribution approximates the probability density function for the prevalence of citizenship that is implied by the beliefs encoded in the model. The distribution peaks in the interval 22-23%. The most likely value is

²³ Lavan, 'Spread of Roman citizenship'.

thus around 22.5%. But the mean value is somewhat higher, at 24%. This is actually the best point estimator of the quantity (it is called the *expected value* or *expectation* of the uncertain quantity in probability theory), because it represents the probability-weighted average of all possible outcomes. Unlike the most likely value (the mode of the distribution), it takes account of any asymmetry in the uncertainty (e.g. the fact that the distribution is slightly skewed to the right in Figure 2). This is the best solution to the problem of asymmetry that I noted earlier, which can cause problems for reasoning based on most likely values alone.²⁴

In this case the improvement in point estimation is modest (though it may be more pronounced in situations of greater asymmetry). The real value of the probabilistic approach lies in the shape of the distribution. It would have been relatively easy to establish a most likely value using a traditional best-estimate approach. It would also be intuitively obvious that a value near the most likely value must be more likely than any higher or lower value. But the historian would not have had sufficient grounds to establish *how much* less likely outlying values are. The value of the Monte Carlo simulation is that it quantifies how my degree of belief falls off as I move away from the most likely value. This gives me a basis for discounting some possible values as improbable.

Estimation always entails a trade-off between confidence and precision. *Ceteris paribus*, the wider the range we allow for, the more confident we can be that the range includes the actual value. But a wider range also contains less information about the quantity. Most strategies for managing uncertainty depend on discounting some possibilities as highly unlikely. Conceptualised in terms of a probability density function representing degree of belief, this means discounting the tails of the distribution and reporting a specified *credible interval*. (A Bayesian *credible interval* – to be distinguished from a frequentist *confidence interval*, which does not admit a probabilistic interpretation – is a contiguous interval that contains a specified proportion

²⁴ On the expected value / expectation of an uncertain quantity, see Lindley, *Understanding uncertainty*, 137-9.

of the total probability mass.) Which interval to report is a matter of convention. Many disciplines operate with a 95% threshold for hypothesis testing and estimation (though the threshold is often applied in a frequentist rather than Bayesian framework). Some fields with better data hold themselves to higher standard and operate with a 99% or even higher threshold. In a field as data-poor as ancient history, a lower threshold, perhaps 80%, may well be appropriate. I suspect that most ancient historians are implicitly operating with even lower thresholds in the ranges they report. But the issue is never formally discussed. In this case, the 95% credible interval for the prevalence of citizenship is 15-33%. This estimate incorporates my uncertainty about the population of the empire and the other uncertain quantities on which it is based. The resulting range, though broad, represented an important advance in our understanding of a quantity that had hitherto resisted quantification entirely.²⁵

It is worth clarifying the reasoning process underlying the use of Monte Carlo simulation to aggregate epistemic uncertainties as I have done here. What I learn from a MC simulation is that the beliefs I hold about the input variables and the laws of probability together constrain the beliefs I can rationally hold about the quantity of interest. The underlying logic is that of *coherence* – the axiom that a set of probabilistic judgements has to be internally consistent to be valid. I have not discovered something new about the quantity of interest. I have merely represented what I know in a more insightful way. If I have discovered anything, it is that I already knew more about the quantity than I realised. The knowledge I have about other, related quantities means that I should assign a much higher probability to some possible values of the quantity than to others. My beliefs about the processes that disseminated citizenship (as encoded in the identification of variables and the mathematical model that links them to the prevalence of citizenship) and my beliefs about the historical values of those variables (as encoded in the input probability distributions) provide grounds for discounting some possibilities as highly unlikely. In

²⁵ On Bayesian credible intervals, see O'Hagan, et al., *Uncertain judgements*, 234-5.

other words, I found that I already knew enough about the mechanisms of the enfranchisement and the demography of the empire to be confident that the proportion of the population who had citizenship in 212 was in the range 15-33%.²⁶

The aggregation of probability distributions through Monte Carlo simulation is a better way of manipulating uncertain quantities than traditional approaches that collapse uncertainty by representing all variables as point estimates. In some cases, as in my quantification of the spread of Roman citizenship, Monte Carlo simulation will produce credible intervals that are usefully narrow, revealing that the historian knew more than they realised about the quantity of interest. In other cases, however, even an 80% credible interval may be very broad and hence relatively un-informative. But that too would be significant and would serve to demonstrate the vulnerability of any existing estimates for the quantity.

EPISTEMIC INTERDEPENDENCE

The core of the Monte Carlo approach is the representation of uncertain quantities as independent random variables. It is the assumption of independence (together with the shape of the input probability distributions) that causes the simulation outputs to cluster around a most likely value. But it is not always justifiable to treat uncertainties as independent. The Monte Carlo approach can cope with interdependence, but only if it is accounted for properly. It is therefore essential to consider whether there is any interdependence between the uncertainties that are being combined. The interdependence in question is specifically *epistemic*. It is a matter of the interdependence of historical problems. Two quantities are epistemically interdependent if acquiring new information about one quantity would change an observer's beliefs about the

²⁶ On the principle of coherence, see Lindley, *Understanding uncertainty*, 36-7 and 236-7 and, for the application discussed here, Aven, *Foundations of risk analysis*, 97-8.

second quantity. The question to ask is whether my probability distribution for one quantity would change if I were to learn more about the value of another quantity.²⁷

Returning to my initial example of the Roman population, it would clearly be wrong to treat the uncertainties about the population of Italy and the population of Iberia as independent, since all estimates for Iberia are explicitly or implicitly informed by an assumption that it was somewhat less densely populated than Italy. If I somehow discovered that the peak population of Italy was somewhere in the region of 15 m (rather than the 7-8 m that is commonly assumed), I would have to adjust my probability distribution for the population of Iberia accordingly. Modelling the two variables as independent would fail to account for this interdependence, with the result that the Monte Carlo simulation would not reflect my real beliefs. But there are many other uncertainties that can be regarded as independent. Take the peak population of the empire and the proportion of slaves who were freed (another variable relevant to the spread of citizenship, since slaves freed by Roman citizens became citizens themselves). In this case, discovering the exact value for the former would in no way reduce my uncertainty about the latter. The two can thus be treated as independent.

Interdependence will either reduce the variance of the output probability distribution (i.e. produce a narrower credible interval), if the effect of an extreme value for one variable is partly offset by a correspondingly extreme value for the other variable, or increase it, if the effects of extreme values compound each other (as would be the case with the populations of Italy and Spain with regard to the total population of the empire). Various strategies are available for managing epistemic interdependence once it has been identified. But unacknowledged epistemic interdependence is the potential Achilles heel of any probabilistic model. The most dangerous

²⁷ On epistemic (also termed probabilistic and subjective) interdependence, see Garthwaite, et al., 'Statistical Methods for Eliciting Probability Distributions', 686, O'Hagan, et al., *Uncertain judgements*, 107-8 and 243, Lindley, *Understanding uncertainty*, 52-3.

pitfall is ignoring strong epistemic interdependence that makes extreme outcomes more likely. This will result in credible intervals that are narrower than they should be given the historian's beliefs. The risk of such an error is increased by the incentives facing the historian, who is often hoping to produce a narrower and hence more informative estimate. Epistemic interdependence needs to be accounted for carefully.²⁸

SENSITIVITY ANALYSIS & ITERATION

Monte Carlo simulation is useful not just in arriving at an estimate but also in clarifying the structure of the problem. Various types of sensitivity analysis can be performed to measure how much each of the individual input variables contributes to the overall uncertainty about the output variable, i.e. to identify the most important components of uncertainty. Sensitivity analysis of this sort is particularly important for the ancient historian because the Monte Carlo approach works best as an iterative process. The information relevant to problems in ancient history tends to be dispersed and difficult to interpret. Even discounting the discovery of new information, an individual's assessment of the current state of knowledge can only be provisional because it will rely on the work of others (which may be ambiguous or misleading) and is likely to omit at least some relevant information (e.g. comparable data from other regions or periods). The process of assigning probability distributions to the input variables should therefore be approached as an iterative rather than one-off exercise. It is best to start with a first, rough set of probability distributions (erring on the side of exaggerating the uncertainty), run a Monte Carlo simulation to identify the variables that contribute most to the uncertainty about the quantity of interest, refine the probability distributions for those variables through a deeper review of the evidence for them, and so on, until the probability distribution for the quantity of interest begins to stabilise. It may be necessary to go through several iterations to arrive at a stable estimate

²⁸ Options for incorporating interdependence are discussed by Garthwaite, et al., 'Statistical Methods for Eliciting Probability Distributions', 687 and Vose, *Risk analysis*, 356-64.

(stable in the sense that it is unlikely to change significantly based on further analysis of the available evidence).²⁹

This paper has sought to demonstrate the value of subjective probability as a tool of historical analysis. It provides a framework that can accommodate the significant epistemic uncertainty involved in estimates of historical quantities especially (but not only) for periods for which we have limited data. Thinking in terms of probability distributions is a valuable discipline for all attempts to estimate uncertain historical quantities, because it draws attention to aspects of the problem that are obscured by traditional approaches that focus exclusively on the most likely value. It is even more useful when multiple uncertain quantities have to be combined in a single analysis, a common occurrence in ancient history. Though it may appear a radical departure from current practice, it builds upon a probabilism that is already implicit in historical reasoning. Most of the estimates that circulate in ancient history are implicitly expressions of their proponents' probability distributions for the quantities in question, insofar as they have reported the value they judge most likely on the available evidence. But the traditional best-estimate approach leaves their beliefs about the likelihood of other possible values unclear or even unexamined.

It bears repeating that these probabilities have to be understood as subjective, in the technical sense. They are not estimates of some objective probabilities that exist in the world (indeed for Bayesians there are no objective probabilities), but rather representations of the epistemic uncertainty about quantities that have a fixed value, though it is unknown to us. As such, the probabilities are both conditional and personal. They are conditional because they depend on a state of knowledge. If I discover new relevant information or find that I am mistaken in my assumptions, the probabilities will change. They are personal because these are the probabilities I

²⁹ Vose, *Risk analysis*, 80-88 is a good overview of the options for sensitivity analysis.

assign based on my assessment of the available evidence. I cannot dictate your probabilities. But I can expect that my probability and your probability will converge if we agree on the model and on the probability distributions for the input variables. If we disagree, articulating our beliefs in terms of probability distributions will clarify the area of disagreement and enable us to direct further research and debate accordingly.

The avowedly subjective character of the framework may trouble some historians. But the subjectivist framework merely makes explicit a subjectivity that is inherent in historical analysis. Historians can present their evidence and their argument, but they can never coerce their colleagues' assent. Other scholars may reach different conclusions from the same evidence. The most they can hope for is that their reasoning will prove persuasive. This is no different to my expectation that other scholars will recognise in my probability distributions a careful and honest representation of the state of knowledge. The subjectivity inherent in historical analysis is too often regarded as an intellectual defect or even an outright barrier to quantification. It is thus common to find historians attempting to obfuscate this subjectivity through a misleading rhetoric of authority and certainty. One of the great merits of the subjectivist framework is that it acknowledges that subjectivity is inherent not just in history but in all empirical disciplines and shows that it is not an obstacle to quantitative analysis.

FIGURE 1
SELECTING A DISTRIBUTION

Probability density

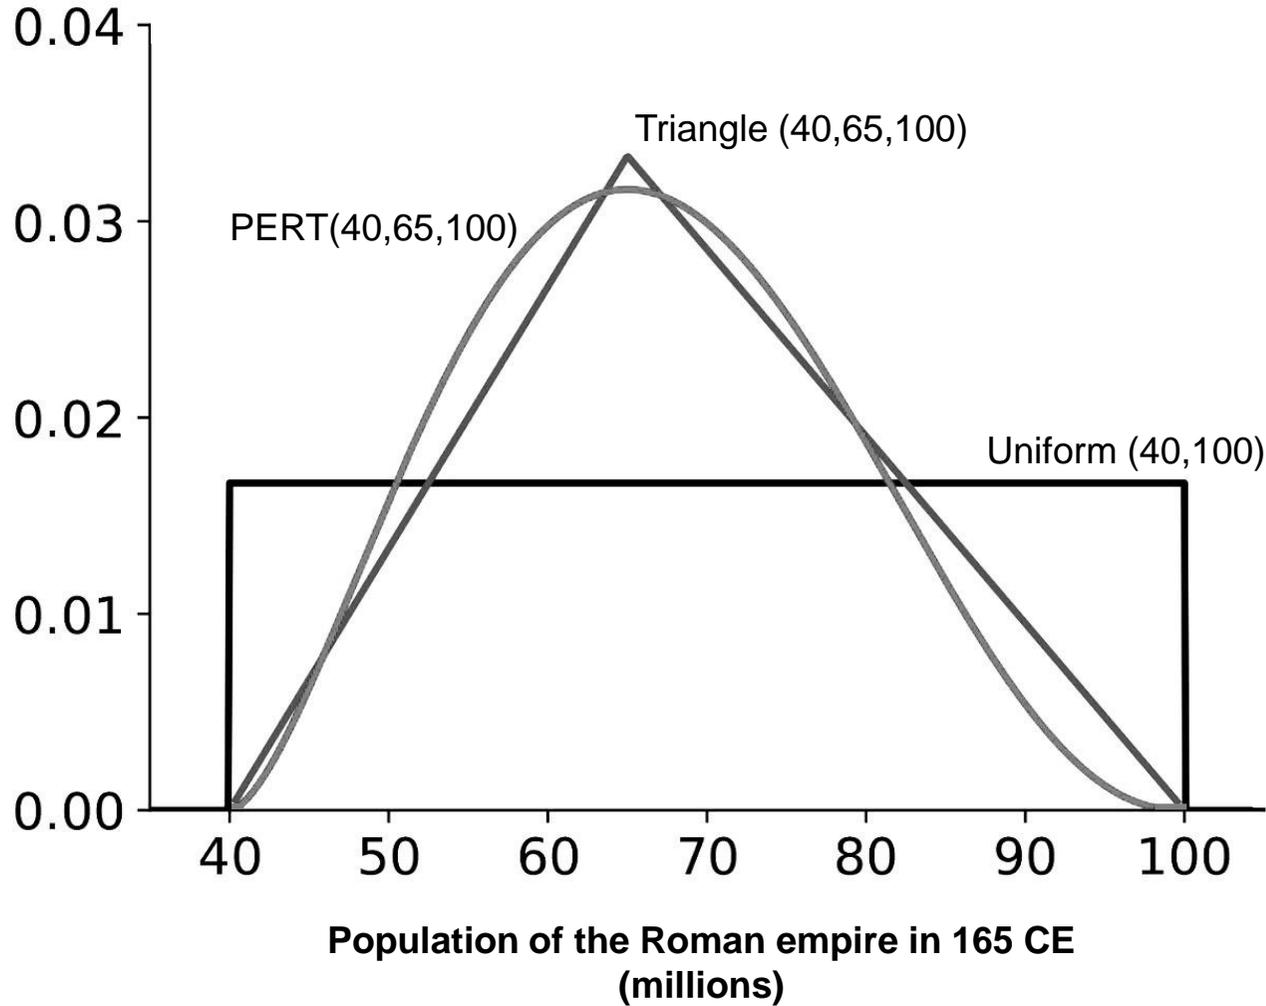


FIGURE 2

MONTE CARLO SIMULATION OF THE PREVALENCE OF CITIZENSHIP

Frequency of outcomes (%)

