8 Wealth Inequality among Italian Elites

In 70 BCE, Cicero prosecuted Verres, the ex-proconsul of Sicily, on the charge of oppression and extortion. In his second speech (which was never spoken out as Verres had by then already gone into voluntary exile), Cicero complains in a more general note that a small group within the Roman elite were amassing the wealth of all nations into their own pockets.[1] While Cicero obviously aims to draw attention to the rapacity of Verres and his associates, he inadvertently points to the high degree of wealth inequality among the Roman elite. The same is implied by his sixth Stoic Paradox, in which Cicero compares the annual revenues from his own estates of HS 100,000 with that of a rich man who draws in HS 600,000.[2]

Roman elites were socially highly stratified.[3] The subdivision of the elite in senators, equestrians and decurions is a well-known example. But among these groups further hierarchies existed. Roman senators and municipal councillors were, for example, typically ranked on the basis of the highest office they had held. The *album* from Canusium is a case in point.[4] In this list, local decurions are ostentatiously subdivided over different status groups (ex-magistrates of different rank, *pedani*, *praetextati*) using explicit headings referring to their rank. Even though some links might have existed between the social and economic differentiation of Roman elites, one of the main tenets of this book is that socio-political status (rank) does not always map very precisely on economic standing (wealth or income).

The aim of this chapter is to assess the level of economic (wealth) inequality among the elites of the Italian *civitates*, without relying on these socio-political differentiations. I therefore use wealth proxy data derived from archaeological and epigraphical sources to get an idea of local elite wealth inequalities. Four different datasets are considered: the size of Pompeian

1 Cic. *Verr.* 2.5.126.
2 Cic. *Parad.* 49.
3 For aversion of social equality, see, e.g., Cic. *Rep.* 1.53 and Plin. *Ep.* 9.5.
4 *CIL* 9.338 and Salway 2000.

https://doi.org/10.1017/9781009496940.009 Published online by Cambridge University Press

houses, the pledges to two Trajanic *alimenta* schemes and the sizes of Ostian burial plots mentioned in epitaphs. The first three of these datasets are well known and well studied. The fourth has not been used in such a context before. All four datasets provide insight in the level of wealth inequality among the elite in the respective *civitates*.

These four Italian datasets imply a high level of variation in the level of elite wealth inequality in the Italian *civitates*. This is no surprise considering the high level of variation in the socio-political and economic aspects of these *civitates* as discussed in the previous chapters. It is perhaps also intuitive that the distribution of wealth looked different at a busy harbour town compared to a small inland *civitas*.

According to Walter Scheidel, this high level of variation in inequality precludes the use of local empirical datasets to infer the overall level of inequality (e.g., in the Roman Empire or even in a province).[5] He concludes that parametric modelling is the way forward. I agree that a straightforward extrapolation of the inequalities implied by a few isolated pieces of evidence to the inequality in an entire Roman province or the empire as a whole would be precipitous. However, it is also wrong to dismiss the empirical evidence out of hand. High local variation in wealth inequality was a reality and a constitutive part of the overall inequality. Empirical evidence is therefore particularly useful to get an idea of the level of inequality.

8.1 Measuring Historical Inequality

There are three interrelated challenges when trying to measure inequality in a historical society using empirical data. These are the scarcity of data, the finding of a reliable variable to represent the inequality implied by the available data and finally the complexity of giving meaning to the estimated variables. In this section, I discuss how I will approach these challenges for the estimation of the level of inequality in the Italian *civitates*.

The empirical evidence extant from Roman Italy is scant. The often-used social-table models circumvents this lack of evidence. However, this method is problematic for other reasons, first and foremost because it reflects the *social*, and not necessarily the *economic*, stratification of society.[6] I therefore use proxy data.[7] These are datasets consisting of various types of empirical evidence that is not representing wealth directly but is assumed to *proxy* it.

5 Scheidel 2020: 342–43.
6 For further discussion of this method, see Section 3.1.
7 On the use of ancient proxy data, see Verboven 2018; 2021.

In other words, I assume that there is a (rough) correlation between the studied proxy quantities and household wealth (i.e., the wealth of a person *sui iuris*).

I use three types of proxy data: the ground-floor area of residences, the pledges to two *alimenta* schemes and the size of burial plots mentioned in epitaphs. There are many problems with the use of these proxies. Two issues are most pertinent: the strength of the correlation between the proxy quantity and wealth and the representativeness of the proxy dataset for the entire economic top layer. The following three paragraphs summarise the more elaborate discussions of these issues presented in Appendix A to C.

The first type of wealth proxy is house size. I use the ground-floor area of the intramural Pompeian residences (i.e., houses with a residential function) as a proxy for the wealth of the occupying household. I expect these residences to underestimate the local wealth inequality. The main reasons are that upper floors are ignored, the large villas outside the town walls are excluded and urban houses cannot grow commensurately with the wealth of the household (due to existing urban features). Other factors that distort the correlation between house size and household wealth include our ignorance of possible variations in land prices over town and the unknown number of uninhabited houses. The representativeness of the residences for the Pompeian economic elite is imperfect mainly for two reasons. First, more than a quarter of town remains unexcavated and it is doubtful whether the urban texture of this section was similar to that of the excavated parts. Second, as the data only include urban households (extramural villas are excluded), a specific part of the local elite is not represented.[8]

The second type of proxy data is the pledges to the imperial *alimenta* schemes. Two long inscriptions set out the details of the Trajanic *alimenta* schemes of Ligures Baebiani (Samnium) and Veleia (Aemilia).[9] Local landowners pledged estates to these schemes for a one-off loan. The interest paid on these loans (probably in perpetuity) was used to support local children. I use the total value of the estates pledged by the participants as a proxy for their wealth. There are two main caveats for the use of this proxy data. First, the participants probably did not pledge their entire (local) landholding to the scheme, which distorts the correlation between their pledge and their total wealth in an unknown manner. Second, not all local landowners participated in the scheme, which undermines the representativeness of this dataset. It is impossible to know who is missing, but the heterogeneity of

8 See Appendix A for more details.

9 Ligures Baebiani: *CIL* 9.1455 (*TALB*). Veleia: *CIL* 11.1149 (*TAV*).

https://doi.org/10.1017/9781009496940.009 Published online by Cambridge University Press

the participants (among whom there are women, minors and probably also freedmen) indicates that the scheme was at least in theory open to the entire economic elite.[10]

The third type of wealth proxy is the size of burial plots. A number of Ostian epitaphs mention the dimensions of the burial plot in which a funerary monument was set up. I use the sizes of these plots as a proxy for the wealth of the dedicator(s) of the burial. There are various factors that distort the correlation between burial plot size and the wealth of the dedicator, for instance, the different relationships between the dedicator(s) and the deceased, individual idiosyncrasies in burial practices, variations in land prices in the cemeteries (both spatially and chronologically) or the different types of tomb that were erected on the plots. Another challenge is that freedmen are grossly over-represented among the dedicators of Roman burials. It is likely that the distribution of wealth among freedmen was different compared to that in society at large, but it unfortunately remains unclear whether it would have been more or less unequal.[11]

Two further caveats for the use of all these proxy datasets need to be mentioned. First, two of these proxies have a distinct urban flavour. This is most evident for the Pompeian residences, but also applies to the Ostian burial plots (which are from urban necropoles). The inclusion of rural residences and burials might change the inequalities implied by these datasets to some extent. The *alimenta* schemes, although the inscriptions were also found in urban contexts, probably included people irrespective of their residence in town or countryside. The fact that the inscriptions were set up in town has probably more to do with the town being the *civitas*' focal point for the commemoration of benefactions than with the place of residence of the participants.[12]

Second, the chronological scope of these proxy datasets vary. While the Pompeian residences and the two alimentary inscriptions notionally describe the situation at one specific point in time, the Ostian burial plots reflect the average wealth inequality over a period of about three centuries, obfuscating any chronological developments within this period.[13]

Most economic historians use the Gini coefficient to represent the level of inequality implied by an empirical dataset.[14] The Gini coefficient is very an index which conveniently summarises the inequality implied by a dataset in

10 See Appendix B for more details.

11 For more details, see Appendix C.

12 Erdkamp 2001: 337–41, Mouritsen 2005: 55.

13 What has been called 'durable' inequality; see Kohler and Smith 2018: 8–9.

14 For example, Alfani and Ammannati 2017, Kohler et al. 2017, Milanović et al. 2011. For Roman examples, see Flohr 2017, Kron 2014, Duncan-Jones 1990: 121–42.

a single number which is relatively easy to interpret. Moreover, its wide use makes it easy to contextualise new results within existing scholarship. There is however one caveat. The Gini coefficient is very sensitive to missing data. This is a big problem for historical wealth proxy datasets, which are often incomplete (or sometimes even fragmentary). In these datasets, the entire bottom of society is typically missing (the poor and the propertyless, not to mention slaves) and the representation of the rich are often also fragmented. As a result, comparing Gini coefficients based on different types of proxy data is very perilous.[15]

As this study focuses on the top of the wealth distribution, an alternative, and more reliable, index is available. This is the shaping parameter of a power-law function (*alpha*) which can be derived from the top of the wealth proxy datasets. The main advantage of using *alpha* is that a power-law function is self-similar and therefore less sensitive to missing data.[16] As a result, if a few values are missing from the top part of the distribution, the Gini coefficient drops considerably, while *alpha* remains much more stable. This is because all available datapoints carry equal weight in determining the value of *alpha*. To make this more intuitive, remember that a power-law function appears as a straight line in a Zipf plot. If a few datapoints are missing, the slope of this line (representing *alpha*) does not change dramatically.

A precondition for using *alpha* is that the top of the wealth proxy datasets resemble a power-law function. To check whether this is the case, I plot the datasets as Zipf plots (see Figure 8.1).[17] Reassuringly, the top part of all these plots follow the expected shape (a straight line with a negative slope), which implies that the highest values of these proxy datasets resemble a power-law function.

The Zipf plots additionally provide reassurance that the selected evidence is indeed proxying wealth. In particular the two largest datasets (the Pompeian residences and the Ostian burial plots) display the typical bipartite shape of a wealth distribution. At lower values the distribution is non-linear, indicating an exponential function, while at higher values it is linear, indicating a power-law function. Considering the plethora of interpretative problems and biases associated with these datasets, it is surprising how consistently these distributions resemble the expected shape of a wealth distribution.

15 For a further discussion on the use of the Gini coefficient, see Section 3.1.

16 Mitchell 2009: 245–46, Gabaix 2009: 259.

17 Following Cirillo 2013, who advises the construction of two further plots (a mean-access and Zenga plot). These plots confirm that the Italian datasets resemble power-law functions (plots not shown).

https://doi.org/10.1017/9781009496940.009 Published online by Cambridge University Press

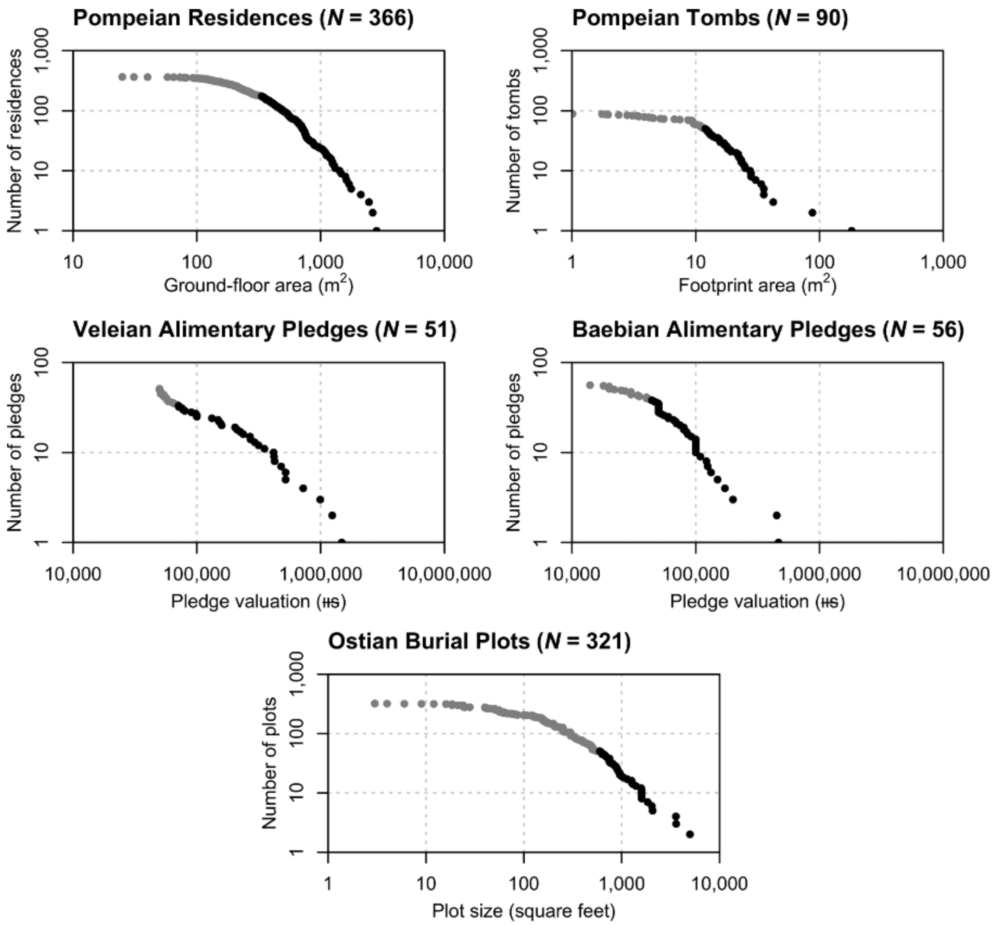

Figure 8.1 Zipf plots of the Italian wealth proxy datasets. *N* represents the number of data points.

It is however noteworthy that the data do not always form smooth curves. Most of these aberrations are due to the small size of these datasets.[18] In statistical terms, these ancient datasets are tiny. To get an idea of the error margin that this introduces, I use an analytical procedure called bootstrapping.[19] Using bootstrapping, I estimate a probability density function (PDF) instead of a point estimate for the value of *alpha*. The PDF indicates the expected level of variation in the estimated values.

[18] Noise in the top part of the Zipf plot is a common feature of this type of plot; see Cirillo 2013: 5949–50.

[19] For more details on bootstrapping in this procedure, see Clauset et al. 2009. For a similar application of this method, see Alfani 2021: 14–15.

https://doi.org/10.1017/9781009496940.009 Published online by Cambridge University Press

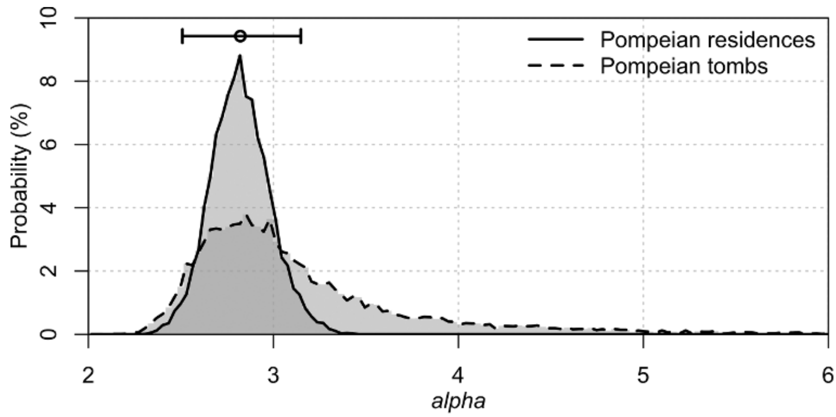

Figure 8.2 Values of *alpha* estimated on the basis of the two Pompeian wealth proxy datasets.

I estimate the values of *alpha* using a preprogrammed function in the statistical software R.[20] I will discuss the results of the Pompeian dataset here in detail as an example. In Figure 8.2, the probability density function (PDF) for the estimated value of *alpha* is presented for the Pompeian residences (the uninterrupted line). This PDF represents the probability (on the vertical axis) that *alpha* would take a certain value (on the horizontal axis). I use two statistics to summarise this PDF; the expected value (i.e., the probability weighted mean) and the 95-per-cent HPD interval (i.e., the shortest range with at least 95 per cent of all the probability mass). In the case of the Pompeian residences, the expected value of *alpha* is 2.82, with a HPD range between 2.53 and 3.13. I present these summary statistics graphically as a circle (the expected value) with 'error bars' (i.e., the 95-per-cent HPD interval; see the example in Figure 8.2).

The example of Pompeii is interesting also because there is another dataset from this town, which can be compared with the residences. This second dataset consists of the footprint sizes of the tombs which are still standing outside the Pompeian city gates. I use the catalogue of Virginia Campbell, who comprehensively catalogued these tombs.[21] There are ninety tombs for which Campbell gives dimensions. They are all dated to the century prior to the eruption of Vesuvius in 79 CE. The largest tomb has a footprint of close to 200 m^2, while the smallest covers less than half a squared metre. The mean is 16.5 m^2 and the median is 12.5 m^2. Similarly as for the other proxy datasets, there are many problems of interpretation and

20 The R function is explained in Clauset et al. 2009.
21 Campbell 2015: 147–312.

https://doi.org/10.1017/9781009496940.009 Published online by Cambridge University Press

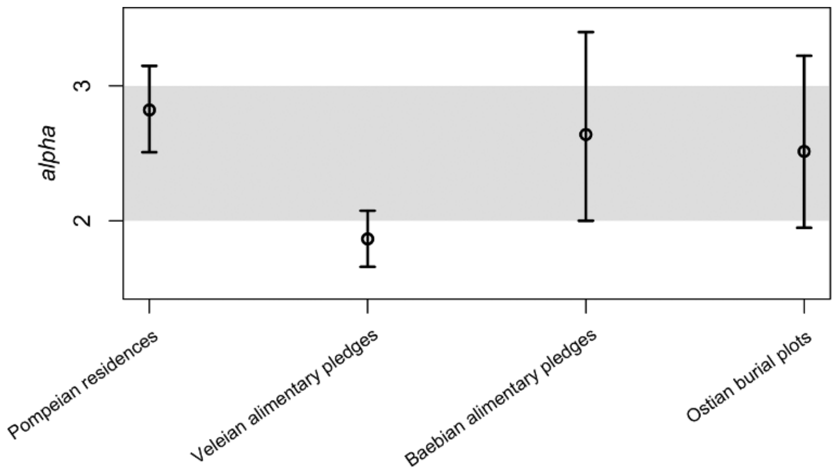

Figure 8.3 Estimated values of *alpha* for the Italian wealth proxy datasets.

representation with this dataset. Most issues are similar to those discussed for the Ostian burial plots and the reader is therefore referred to Appendix C for further details.

In Figure 8.2, the PDF of the estimated values of *alpha* based on the Pompeian tombs is represented by the dashed line. I make two observations. First, the quality of the two datasets are not the same. There are considerably fewer tombs (N = 90) than there are residences (N = 366). This is reflected in the uncertainty of the estimated value of *alpha*; the PDF based on the tombs is much wider (indicating a higher uncertainty) than that based on the residences. Second, the two datasets nonetheless suggest a very similar area of highest probability (around a value just below 3). The coincidence between the areas of highest probability is impressive, especially when one considers that these two PDFs are based on completely different data. This coincidence thus reinforces the plausibility of the values of *alpha* estimated for Pompeii (reminiscent of Keith Hopkins' wigwag argument).[22]

8.2 Italian Elite Inequalities

Figure 8.3 presents the estimated values of *alpha* (expected values and 95-per-cent HPD intervals) for all four Italian wealth proxy datasets. The reader is reminded that *alpha* is inversely correlated with inequality; a lower value of *alpha* implies a higher inequality and vice versa.

[22] Hopkins 1978: 19–20, see also Section 4.4.

https://doi.org/10.1017/9781009496940.009 Published online by Cambridge University Press

The interpretation of these estimated values of *alpha* is far from straightforward. First, caution is required for comparing the *alpha*'s based on different types of proxy data. Each type of evidence has its own particular problems of interpretation which will affect the estimated value of *alpha* in their own particular ways. Diligence is due when an *alpha* based on house size is compared with, for example, an *alpha* based on pledges to an *alimenta* scheme. Consequently, I can draw only broad conclusions.

What can these estimates tell us about the level of local wealth inequality among the elites of Roman Italy? First, the inequality implied for the Pompeian elite is relatively low (represented by a relatively high value of *alpha*). Different scholarly views exist on the level of Pompeian inequality. Some scholar suggest that inequality at Pompeii was higher than in earlier Greek towns.[23] Miko Flohr, for example, estimates a relatively high Gini coefficient for Pompeii.[24] His estimate of 0.62 is, for instance, much higher than the 0.44 as estimated for the empire at large by Scheidel and Friesen.[25] It is worth noting that the many small *tabernae* (which Flohr includes in his analysis) drive up the Gini coefficient considerably.[26] Conversely, a recent comparison of Gini coefficients based on the housing stock from Pompeii and a few other Roman towns points to a relatively modest inequality at Pompeii.[27]

Inequality at Ostia is significantly higher than at Pompeii. Ostia was one of the largest Italian *civitates*, which might have driven up inequality. On the other hand, the Ostian economy was heavily orientated towards commerce, which made it different from most other Italian *civitates*. Plenty of commercial opportunities, in combination with high levels of immigration, would probably have tempered local inequality. It has also to be remembered that the implied inequality at Ostia is the average for the entire Early Imperial period. Troughs in the level of inequality, for instance in the wake of the Antonine Plague, might lower the implied inequality to some extent.

The inequality implied for Veleia is the highest. One can only guess at the reasons behind this high inequality. Possibly, this is a result of the large territory of Veleia in combination with the agrarian focus of the local economy, which allowed the local elite to concentrate most wealth (i.e., land) into their own hands.

The two sets of *alimenta* pledges (in Veleia and Ligures Baebiani) imply very different inequalities. The inequality implied for Veleia is much higher

23 Wallace-Hadrill 1994: 75–77.
24 Flohr 2017: 74–75.
25 Scheidel and Friesen 2009.
26 For comparison purposes, I calculate a Gini coefficient of 0.41 based on the residences.
27 Kron 2014: 128–29.

than that implied for Ligures Baebiani.[28] Possibly, this is due to differences in the nature of the two alimentary schemes. For example, the threshold for participation in the Baebian scheme appears to have been lower than in the Veleian scheme (as inferred from the smallest pledges recorded in the two schemes). Or the Baebian landowners might have been less enthusiastic to participate than their Veleian peers.[29] It is also possible that there actually was a difference between the two towns in the distribution of elite wealth. The Veleian territory was presumably much larger than that of Ligures Baebiani, which widened the scope for differentiation within the local elite.[30] Furthermore, Ligures' possible subordination (or at least substantial loss of territory) to Beneventum might also have limited the concentration of wealth in the hands of the local elite.[31] According to Marco Maiuro, the difference is due to the share of local land that was in the hands of the emperor, senators and equestrians, who possessed much more land at Ligures compared to Veleia thus leaving less land to be owned by the local Baebian elite.[32]

A higher level of inequality in the northern town of Veleia compared to the southern town of Ligures Baebiani does not correspond with the theory of Andrik Abramenko that the Italian South (regions I to IV) was more unequal than the North (regions VIII to XI).[33] Abramenko infers this difference in inequality from distributions of burial plot sizes in the two regions. Possibly, the inclusion of *Regio* I (Latium and Campania) in Abramenko's South accounts for much of this higher inequality. In this *Regio*, many wealthy persons from Rome owned luxury estates, where they were also buried.[34] It is also possible that Veleia and/or Ligures Baebiani were exceptions or outliers in their respective regions.

It is probably safest to conclude that the Italian evidence points to a high level of variation in local inequality, apparently irrespective of the wider region in which they were situated. Some of this variation might be explained by the different types of wealth proxy data that have been used. But the disparities might also be due to, for example, differences in the nature of the local economy; the commercial economy of Pompeii might, for example, have led to a lower wealth inequality than the agrarian

28 *Cf.* the trend in the Gini coefficients (0.53 and 0.44 for Veleia and Ligures respectively) estimated by Duncan-Jones 1990: 129–38 (remember that a higher Gini coefficient implies higher inequality). See also Abramenko 1993: 71.

29 Veyne 1958: 219.

30 Veyne 1957: 91–112 (Ligures). Criniti 1991: 219–44 (Veleia).

31 Veyne 1957: 93–94, but see Patterson 1987: 140–42.

32 Maiuro 2012: 139–42.

33 Abramenko 1993: 68–71.

34 *Cf.* Borg 2019: 4–27.

economy of Veleia. Considering the variation in all the other aspects of the Italian *civitates* discussed in the previous chapters, this variation in economic inequality is of course not completely unexpected.

A priori, a positive correlation between the size of a *civitas* and the level of inequality might be expected.[35] The present datasets do however not substantiate such a relationship. The inequalities implied for Pompeii and Ostia (the larger *civitates*) are not notably higher than those implied for Veleia and Ligures Baebiani (the smaller *civitates*). The inequalities implied for the two larger and smaller towns are also notably different from each other. Although wealth inequality and population size must have been related, there were probably many other factors that also affected the local wealth inequality (e.g., the type of economy).

It is worth noting that the range of the implied inequalities roughly overlaps with that of inequalities implied by comparative material, which thus reinforces the plausibility of the present results.[36] For instance, the *alpha* implied by the Florentine *catasto* is around 2.45 and the Hermopolite land register around 1.70.[37] Adel Abul-Magd estimates an *alpha* of 2.59 for the houses in the fourteenth-century-BCE Egyptian city of Akhetaten.[38] Géza Hegyi et al. estimate an *alpha* based on the number of serf families owned by the sixteenth-century-CE Hungarian elite of between 1.92 and 1.95.[39] In sum, the values based on comparative evidence fall roughly within the same range (between 2 and 3) as those based on the Italian datasets. The comparative material thus broadly supports the plausibility of the values estimated for Roman Italy.

8.3 Conclusions

This chapter assessed the level of elite wealth inequality in Roman Italy on the basis of a series of wealth proxy datasets. Three different types of data from four different Italian *civitates* have been analysed. Despite the plethora of issues of interpretation and representation associated with these datasets, the top parts of all datasets appear to follow a power-law function. This not only served as a reassurance that these datasets can be used as wealth proxies, but also allowed me to estimate a value of *alpha* (the shaping parameter

35 *Cf.* Sitthiyot and Holasut 2016: 26–27, Deltas 2003. Inequality was, for example, higher at Paris than in France as a whole, Piketty 2017: 428–33, esp. figure 10.2.
36 *Cf.* Gabaix 2009: 275.
37 See Section 3.2.
38 Abul-Magd 2002. But Akhetaten might have been an exceptional city; see Montserrat 2000: 12–54.
39 Hegyi et al. 2007.

https://doi.org/10.1017/9781009496940.009 Published online by Cambridge University Press

of the power-law function) for each dataset, which I then used as an index of inequality among the local elite.

The inequalities implied by these datasets suggest first and foremost that wealth inequality among the elites of the Italian *civitates* varied considerably. In the previous three chapters, I have discussed various aspects of the Italian *civitates*: the size of their administrative centres, the number of local decurions, the local curial census qualification and the number of households with curial wealth outside the council. I concluded that all of these facets of the *civitates* varied considerably. This chapter adds wealth inequality to the list.

Based on the evidence reviewed in this chapter, there does not appear to have existed a strong correlation between the level of inequality and the size of the *civitas*. Even though *civitas* size must have affected the local level of economic inequality, it was probably just one factor among many others (including the nature and state of the local economy and the economic, social and political history of the *civitas*).

In the next chapter, I will present a new model to reconstruct the top of the wealth distribution of Roman Italy as a whole. This model is different from previous models in that it takes this multifaceted heterogeneity of the Italian *civitates* explicitly into account.

https://doi.org/10.1017/9781009496940.009 Published online by Cambridge University Press