3 Reconstructing Historical Wealth Distributions

Inequality appears to be an intrinsic feature of life in social groups. A certain level of unequal distribution of resources can, for instance, be discerned for African great apes and early human hunter-gatherers. The advent of farming and more complex forms of socio-political organisation greatly increased disparities. The creation of empires, in which political and economic power were intimately interwoven and highly concentrated in the hands of small elites, caused inequalities to soar.[1]

Measuring this inequality is not an easy task. Even today, using highly advanced technologies of data acquisition and analysis, it is surprisingly hard to measure economic (income or wealth) inequality accurately and unequivocally. Problems arise particularly at the extremes of the distribution. At the bottom, the economic activity and assets of certain marginal groups (e.g., illegal immigrants) are typically very hard to measure. At the top, datasets tend to be imprecise due to, for example, low response rates by the rich in wealth surveys or tax evasion practices distorting fiscal records.[2]

When measuring inequality in the ancient world, the challenges are very similar. For example, the extensive and seemingly meticulous land registers of fourth-century Hermopolis only include those who owned property *and* were liable to pay taxes; tax-exempt landowners and the propertyless are both missing.[3] Polybius also notes already in the second century BCE how difficult it is to measure the property of members of the Roman elite due to their constant family squabbles.[4]

In this book, I will reconstruct the distribution of elite wealth in Early Imperial Italy. The focus is on *elite* wealth, meaning that I concentrate on the top part of the wealth distribution. This makes sense as I use the reconstructed wealth distribution to think about the political economy of Roman Italy and political officeholding was the prerequisite of the wealthy in the

1 Scheidel 2017: 25–61, Smith et al. 2018.
2 Bach et al. 2019, Vermeulen 2018.
3 Sijpesteijn and Worp 1978.
4 Polyb. 18.35.8.

Roman world. The focus is also on *wealth*, and not income. The distribution of income and wealth look fundamentally different, especially at the bottom of the distribution (a person can live without wealth but not without income). The top of the wealth and income distributions are much more similar, as most, if not all, income of the elite is derived from returns on property. The advantage of the focus on elite wealth is thus that abovementioned problems associated with measuring the bottom of a wealth distribution are less of a concern here.

This chapter is subdivided into three sections. The first discusses the approach that most economic historians employ to reconstruct wealth distributions in historical societies, so-called 'social-table' models. I will show that these models systematically underestimate the level of inequality. The second section presents an alternative model derived from the economic sciences, in which the top of a wealth distribution is assumed to follow a mathematical function (a power law). This model seems more appropriate for the present purposes. The third and final section argues that Roman empirical evidence can provide a firm foundation for this model.

3.1 Social-Table Models

Most studies which reconstruct historical wealth distributions rely on so-called social-table models. In this approach, a society is modelled as consisting of various social groups. A wealth distribution is then reconstructed by estimating the number and average wealth of the members of each social group. The main advantage of these models is that the type of historical evidence required to construct them is relatively abundant.[5]

In Table 3.1, I have summarised the results of four social-table models of the distribution of elite wealth in the Roman Empire.[6] When comparing these models, it is important to consider their social, geographical and chronological scope as well as their purpose. All these studies focus on the elite (roughly the top few percentage points of society). Scheidel and Friesen focus on the middle of the second century CE, while the other studies concentrate on the Augustan era. However, since all these scholars essentially draw on the same evidentiary material, it is probably best to see their models as applying to roughly the first two centuries CE.[7] Finally, the purpose of the

5 Milanović et al. 2011.

6 Goldsmith 1984: 276–79, Milanović et al. 2007: 64–69, Maddison 2007: 48–50, Scheidel and Friesen 2009: 75–79.

7 *Cf.* Milanović 2019: 11–12.

Table 3.1 Social-table models for the wealth distribution of the Roman elite. The percentages of elite households are based on an average household size of four persons.

	Goldsmith (1984)			**Milanović et al. (2007)**		
Region	Empire			Empire		
Period	14 CE			14 CE		
Population	55 million			55.5 million		
Elite households	400,601 (2.9%)			600,600 (4.3%)		
Elite wealth	HS 70 billion			HS 133 billion		
	Number	Wealth (HS 1,000)	(%)	Number	Wealth (HS 1,000)	(%)
Emperor	1	240,000	0.3	-	-	-
Senators	600	2,500	2	600	2,500	1
Equestrians	40,000	500	29	40,000	500	15
Decurions	360,000	133	69	360,000	133	36
'Other wealthy'	-	-	-	200,000	317	48

	Maddison (2007)			**Scheidel and Friesen (2009)**		
Region	Empire			Empire		
Period	14 CE			150 CE		
Population	44 million			70 million		
Elite households	330,601 (3.0%)			215,600–290,600 (1.2–1.7%)		
Elite wealth	HS 74 billion			HS 44–60 billion		
	Number	Wealth (HS 1,000)	(%)	Number	Wealth (HS 1,000)	(%)
Emperor	1	250,000	0.3	-	-	-
Senators	600	2,500	2	600	5,000	5–7
Equestrians	40,000	500	27	20–30,000	600	27–30
Decurions	240,000	139	45	130,000	150	33–44
'Other wealthy'	50,000	367	25	65–130,000	150	22–33

models (e.g., to calculate a minimum or best estimate) guided the scholars in their choice of assumptions, with significant consequences for the final outcome. Whereas Raymond Goldsmith expressly uses minimum estimates for his social groups, the other scholars attempt to obtain best estimates.

All four studies include the three main Roman socio-political orders (senators, equestrians and decurions), but note the differences in the estimates of the number of decurions, which ranges from 130,000 to 360,000. Only Raymond Goldsmith and Angus Maddison include the wealth of

https://doi.org/10.1017/9781009496940.004 Published online by Cambridge University Press

the emperor, which is however numerically of minor importance to the overall level of inequality. Branko Milanović and co-workers importantly note that there must have been a substantial number of wealthy households outside these three traditional socio-political orders. Accordingly, they add a group (200,000) of relatively affluent (average wealth of HS 317,000) 'other wealthy'. This group is of course of major importance for this study as they represent the households with elite wealth outside the traditional orders. In the model of Milanović et al., they are very wealthy; their aggregate wealth almost accounts for half of all elite wealth. Later studies have lowered both the number and average wealth of the 'other wealthy'. For example, the wealth of this group constitutes only about a quarter and third of all elite wealth in the models of Maddison and Scheidel and Friesen, respectively. The large differences in the estimates for their number (between 50,000 and 200,000) and their wealth (between HS 150,000 and HS 367,000) emphasise our ignorance of this group.[8]

The main problem of the social-table approach is that it invariably results in a wealth distribution that is too equal. A social-table model thus invariably underestimates the level of inequality.

Two causes for this underestimation of the inequality can be identified. First, within-group inequalities are disregarded. As the social-table approach presupposes an average property size for the members of each social group, inequalities within these groups are ignored. It is easy to show that there were considerable within-group inequalities within Roman social groups. For example, the senator Seneca is believed to have owned around HS 300 million, while the wealth of his colleague Pliny the Younger, who lived one generation later, has been estimated at 'only' HS 20 million.[9] Such within-group disparities are thus not accounted for when using a social-table model.

Second, the social-table method does not allow for overlap between groups. It assumes perfectly sorted groups. In other words, the members of each social group cannot have wealth which is overlapping with the wealth of members of a lower or higher social group. The Roman evidence provides several explicit examples that attest to the overlap between the wealth of members of different social groups. The best-known example are the three imperial freedmen of the first century (Narcissus, Pallas, and Callistus), whose wealth probably rivalled that of Seneca.[10] These freedmen would however have belonged to the 'other wealthy' group, as they were officially barred from the traditional socio-political orders.

[8] For an elaborated discussion of this group, see Chapter 7.
[9] Seneca: Tac. *Ann.* 13.42. Pliny the Younger: Talbert 1984: 49, Duncan-Jones 1982: 17–32.
[10] Duncan-Jones 1982: 343–44.

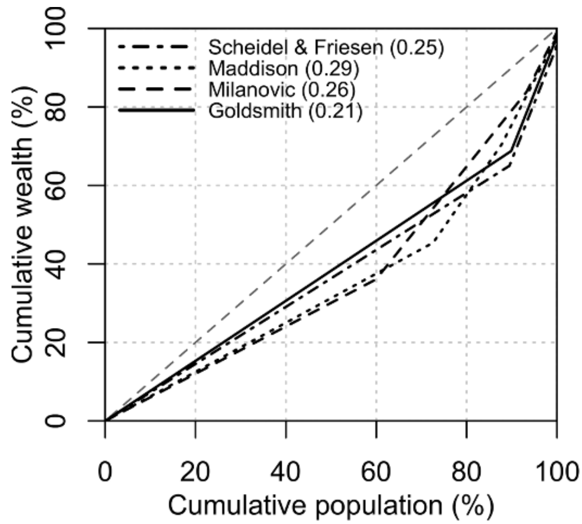

Figure 3.1 Lorenz curves of the distribution of Roman elite wealth as predicted by social-table models in previous studies (see Table 3.1). For the model of Scheidel and Friesen, I used the higher estimate for the 'other wealthy'. Gini coefficients are given in brackets in the legend.

Both of these aspects of the social-table approach thus result in an underestimation of the overall inequality. Jørgen Modalsli shows that their effect can be significant, with an underestimation of the implied inequality of up to thirty per cent.[11]

The underestimation of inequality can be further demonstrated by putting the results of the Roman social-table models in a comparative perspective. In Figure 3.1, the social-table results are plotted as Lorenz curves. These curves represent the cumulative share of wealth (represented on the vertical axis) for the respective cumulative shares of the population (represented on the horizontal axis). The straight dashed grey line is the equality line and represents perfect equality; on this line *x* per cent of the population owns precisely *x* per cent of all wealth. The farther the Lorenz curve deviates from the equality line, the higher the inequality.

The Gini coefficient is a unidimensional summary of the Lorenz curve. It is defined as the area between the Lorenz curve and the equality line, proportional to the entire area below the equality line. A Gini coefficient of zero thus implies that the area between the Lorenz curve and the equality line is non-existent (meaning that the Lorenz curve coincides with the equality line) and thus implies complete equality. Conversely, if the Gini coefficient is one, the Lorenz is the farthest away from the equality line (coinciding with

[11] Modalsli 2015: 225–29, based on Gini coefficients.

the bottom and right side of the graph), thus maximising the area between the equality line and the Lorenz curve and implying the theoretical maximum level of inequality. In other words, the higher the value of the Gini coefficient, the higher the inequality.[12]

All the Lorenz curves in Figure 3.1 are fairly close to the equality line, implying relatively low inequalities. This is confirmed by the relatively low values of the Gini coefficient (see the numbers in the legend of Figure 3.1). They range from 0.21 to 0.29. These values are significantly lower than the coefficients based on wealth proxy data from other premodern societies. For example, the Gini coefficient implied by the landholdings mentioned in a fourth-century tax document from Hermopolis in Middle Egypt is between 0.79 and 0.82.[13] The wealth mentioned in the early fifteenth-century Florence *catasto* results in a comparable coefficient of 0.79.[14] Unless wealth was distributed considerably less unequally among the Roman elite than in these comparative cases, these curves and coefficients confirm the significant underestimation of inequality by social-table models.

It is further worthwhile noting that the curves for the various social-table models are hard to distinguish. This points to the high degree of similarity between the models. On the one hand, this is not surprising as they all draw on the same sparse evidence for the number and census qualifications of the three socio-political orders. On the other hand, this is remarkable considering the relatively large differences between the estimates for the number and average wealth of some of the groups (see Table 3.1).

The underestimation of Roman wealth inequality by the social-table models can also be demonstrated using the wealth share of the top 1 per cent (or top centile) of society. The top-centile wealth share is an alternative metric to characterise inequality in a society and is preferred over the Gini coefficient in many recent studies on inequality.[15] The downside of this metric is that it requires an estimate of not only the aggregate wealth of the top centile but also an estimate of the total capital in a society. Particularly the latter estimate is very hard to establish for historical societies. The following calculations are therefore very speculative and are meant as very crude indications of orders of magnitude only. Furthermore, I only discuss the social-table models of Goldsmith, Maddison and Scheidel and Friesen, as these scholars also present an estimate for Roman GDP (which I use to estimate total capital).[16]

12 For the calculation of the Gini coefficient, see, e.g., Milanović et al. 2011: 257–59.
13 Bowman 1985: 150.
14 Herlihy and Klapisch-Zuber 1985.
15 E.g., Scheidel 2017: 11–12, Piketty 2017: 332–36.
16 Goldsmith 1984: 263–74, Maddison 2007: 45–47, Scheidel and Friesen 2009: 63–74. Bowes 2021 critiqued the use of GDP estimates, with a reply to this critique in Scheidel 2022.

Table 3.2 Elite wealth shares based on the social-table models. The numbers printed in italics are extrapolated.

	Goldsmith (1984)	**Maddison (2007)**	**Scheidel and Friesen (2009)**
Region	Empire	Empire	Empire
Period	14 CE	14 CE	150 CE
Total population	55 million	44 million	70 million
Total households	*13.8 million*	*11.0 million*	17.5 million
Top-centile households	*138,000*	*110,000*	*175,000*
Elite households	400,601	330,601	290,600
Aggregate wealth elite (HS)	70 billion	74 billion	60 billion
Aggregate wealth top centile (HS)	*35 billion*	*43 billion*	*43 billion*
Roman GDP (HS)	21 billion	17 billion	18 billion
Total Roman capital (HS)	*147 billion*	*119 billion*	*126 billion*
Top-centile capital share	*24%*	*37%*	*34%*

I first determine aggregate elite wealth by multiplying the estimated average wealth for each group with the estimated size of these groups and adding up the results for each of the different groups. The results are overviewed in Table 3.2. It is striking that Scheidel and Friesen obtain the lowest estimate for aggregate elite wealth (HS 60 billion), even though they insist (to my mind, correctly) on an upward correction of Goldsmith's estimates for the average wealth of the elite groups (although they increase the figure for the wealth of decurions only marginally).[17] Their aggregate is nonetheless lower because they adjusted the sizes of the elite groups downward. Indeed, their total elite number ranges between a minimum of 215,600 and a maximum of 290,600, while the other studies assume totals of 330,600 or more. This is even more surprising, considering that Scheidel and Friesen notionally focus on the empire in the middle of the second century CE, according to the authors 'at the time of its putative demographic peak'.[18]

In order to derive the wealth of just the top centile from these aggregates of elite wealth, I adjusted their values in the following manner. First, I calculated the total number of households in the empire by dividing the population estimates by an estimate of the average number of household

17 Scheidel and Friesen 2009: 76–77.

18 Scheidel and Friesen 2009: 62. They assume the highest total population of 70 million inhabitants.

https://doi.org/10.1017/9781009496940.004 Published online by Cambridge University Press

members of four.[19] Next, I adjusted aggregate elite wealth by excising the households below the top 1 per cent.

Finally, the total capital in the Roman Empire needs to be estimated, which is the most precarious task. As there is not enough ancient evidence, I rely on comparative material. Thomas Piketty observes that total capital in France at the start of the eighteenth century was about seven times national income (which is close to GDP).[20] According to Walter Scheidel, early-modern France is economically fairly comparable with the Roman Empire.[21] Caution is of course in order. While the French multiplier might be too high because property rights must have been much more firmly established in France at the start of the eighteenth century than in the Roman Empire during the first two centuries CE, it might be too low as it disregards the wealth represented by Roman slaves, which could add, according to Scheidel, 10 per cent to the total wealth stock.[22] I follow Scheidel in using seven times GDP as a rough indication of total Roman capital.

The results of the calculations are presented in Table 3.2 with the extrapolated numbers in italics. The top-centile wealth shares implied by the social-table models are between 25 and 37 per cent.[23] It is worth noting that the wealth shares implied by the models of Maddison and Scheidel and Friesen are relatively high, which is mainly due to their lower estimates of Roman GDP. The wealth share implied by the work of Scheidel and Friesen would become much lower if Scheidel's more recent, higher estimate of Roman GDP were used.[24]

These wealth shares appear to be on the low side if put in a comparative perspective. For example, Thomas Piketty observes for the top centiles in France and Britain at the end of the eighteenth century wealth shares of 45 per cent and 55 per cent, respectively.[25] In other words, the results based on social tables imply a wealth concentration at the top of Roman society significantly lower than that observed in Europe at the beginning

[19] Following Scheidel and Friesen 2009: 66 and 81 note 73. For higher estimates, see, e.g., Fentress 1994: 133–35, Hanson and Ortman 2017: 308–9 note 36. The Egyptian census records imply 4.3 (Bagnall and Frier 1994: 67–69). Household size must have varied considerably according to economic and social standing (Huebner 2017: 11–12). Note that higher average household size leads to lower top-centile wealth shares.

[20] Piketty 2017: 144–49.

[21] Scheidel 2020: 315–47.

[22] *Cf.* Piketty 2017: 196–203.

[23] A similar model of Maiuro 2012: 117–32 implies much higher proportions, but this model only considers the wealth of those who own at least the equivalent of a discharge bonus for legionary veterans.

[24] Scheidel 2022: 3–8.

[25] Piketty 2017: 428–36. *Cf.* Scheidel 2020: 347.

Table 3.3 Results of the social-table models for Roman Italy.

	Jongman 1988			**Maddison 2007**		
Region	Italy			Italy		
Period	50 CE			14 CE		
Population	7.5 million			7 million		
Total elite households	25,600 (1.4%)			121,600 (6.5%)		
Total elite wealth (HS)	4.6 billion			31.1 billion		
	Number	Wealth		Number	Wealth	
		(HS 1,000)	(%)		(HS 1,000)	(%)
Emperor	-	-	-	1	250,000	1
Senators	600	1,000	13	600	2,500	5
Equestrians	5,000	400	43	24,000	500	39
Decurions	20,000	100	43	80,000	139	36
'Other wealthy'	-	-	-	17,000	367	20

of the industrial revolution. If Piketty is right that the top-centile wealth shares in Antiquity were similar to those of the European nations during the long nineteenth century (of around 50–60 per cent), then these crude calculations reinforce the notion that the social-table models underestimate inequality in the Roman Empire.[26]

Before turning to considering alternatives to the social-table approach, it is worth reviewing the social-table models of Roman Italy presented by Wim Jongman and Angus Maddison (see Table 3.3).[27] The estimates of group sizes and average wealth of these two models are very different. The reason is that Jongman estimates absolute minima, while Maddison provides best-guesses. However, the two models imply a very similar level of inequality – their Lorenz curves and Gini coefficients are almost identical (see Figure 3.2). While at face value this is surprising, it can be explained by the fact that lower estimates for senators and equestrians on the one hand and decurions on the other have opposite effects on the implied inequality. It seems that the effects of Jongman's lower estimates for all groups cancel out, resulting in a similar level of inequality as Maddison's estimates.

The results for Italy further imply a higher elite wealth inequality on the peninsula than in the empire as a whole. The estimated Gini coefficients are also higher at 0.36. These results suggest that elite wealth inequality in Roman Italy was higher than that in the empire at large. This is what might

26 Piketty 2017: 436.
27 Jongman 1988: 193, Maddison 2007: 48–50.

https://doi.org/10.1017/9781009496940.004 Published online by Cambridge University Press

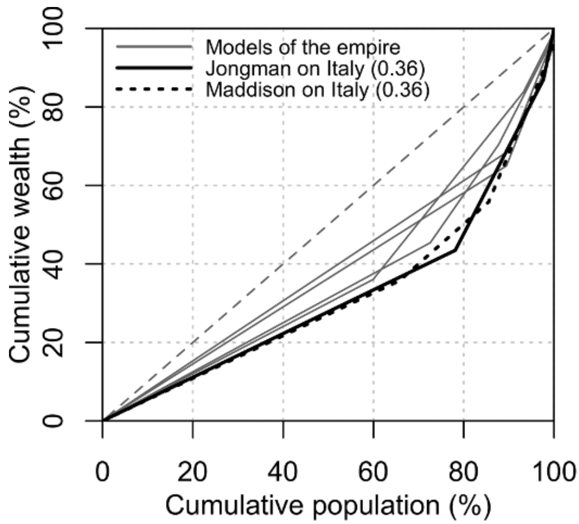

Figure 3.2 Lorenz curves of the distribution of Italian elite wealth based on the social-table models of Jongman 1988: 193 and Maddison 2007: 11–68. Gini coefficients are given in brackets in the legend.

have been expected. Taking into account that more wealth allows for more differentiation and thus higher inequality, Figure 3.2 is in accordance with a disproportionally wealthy Italy, as discussed in Chapter 1.

In conclusion, the social-table approach is not a reliable method for the reconstruction of historical wealth distributions. It invariably underestimates the wealth inequality. The fundamental problem is that social-table models are based on socio-political groups, which do not necessarily overlap with economic groups, despite the fact that the two are strongly linked. Social-table models are moreover particularly inappropriate for this study, which focuses on the (non-)overlap between socio-political and economic status. Using a wealth distribution reconstructed based on socio-political groups would thus lead to a circular argument. A different method is required, one that is independent of the socio-political structures of society. In the next section, I present an alternative model, borrowed from the economic sciences, which assumes that the top part of the Italian wealth distribution can be reasonably accurately represented by a simple mathematical function, a power law.

3.2 Power-Law Models

In the late nineteenth century, Vilfredo Pareto advanced the theory that top wealth invariably followed the same mathematical function, a power

https://doi.org/10.1017/9781009496940.004 Published online by Cambridge University Press

law.[28] He analysed the income and wealth distributions in several sixteenth- to eighteenth-century societies, observing that all of these distributions followed a power law. Later scholarship has corroborated Pareto's thesis using empirical data from both premodern and modern societies. For example, power-law distributions have been identified in the housing stock of fourteenth-century-BCE Akhetaten,[29] the number of serf families owned by Hungarian nobles in the sixteenth century CE,[30] 1996 probate data from the United Kingdom,[31] contemporary rich lists from the United States,[32] and early twenty-first-century wealth survey data from several modern Western countries.[33]

It is worthwhile to discuss the appearance of a power law in two comparative, premodern wealth proxy datasets: the landholding register from fourth-century Egyptian Hermopolis and a tax register from fifteenth-century Florence. These datasets embody two of the highest-quality wealth proxy datasets from the premodern world.

The fourth-century-CE landholding register from the Egyptian city of Hermopolis lists the inhabitants of one of the city's four quarters who own land in the local nome (territory).[34] Following Alan Bowman, I include the small number of entries of multiple owners and unspecified 'heirs', exclude civic land and ignore the distinction made between private and public land (these were probably only fiscal categories).[35] My final dataset consists of 240 landholdings.

There are of course various caveats for using these landholdings as wealth proxy data. First, the register was probably set up for tax purposes, which means that estates and/or landowners exempt from (or purposely evading) taxation are missing.[36] Conversely, the advantage of a tax document is that all landowners, regardless of their gender, legal status and so on were included. Second, the data is a wealth *proxy* in that it only includes one type of property, namely, local landholding. Many if not most of these landowners would have possessed other types of property and/or other land outside the Hermopolite nome. Third, the register only includes inhabitants of the metropolis. This is less of a concern for a study of the top of the wealth

28 Pareto 1897: 303–45.
29 Abul-Magd 2002.
30 Hegyi et al. 2007; Santos et al. 2007.
31 Drăgulescu and Yakovenko 2001.
32 Klass et al. 2006.
33 Vermeulen 2018.
34 *P. Flor.* 1.71. I use the edition of Sijpesteijn and Worp 1978: 63–103, excluding the Antinoites.
35 Bowman 1985: 141–50. For the distinction between private and public land, see also Rowlandson 1996: 63–69, Tacoma 2006: 103.
36 Tacoma 2006: 91, Vitelli 1906: 132, Bagnall 1979: 161–63.

distribution, as other Egyptian tax lists suggest that the largest landholdings were often in the hands of metropolites.[37] Fourth and finally, the register only includes the inhabitants of one of the four quarters of Hermopolis; it is unknown whether the socio-economic composition of the other three quarters was substantially different.[38]

The city of Florence instituted a new taxation system in the early decades of the fifteenth century to increase its revenues for a protracted war with Milan. For this so-called *catasto*, all households made a declaration of all their property, including land, real estate, livestock, moveable wealth (coined money, furniture etc.) and financial assets (shares in public debt). The resulting dataset is a unique source of statistical evidence for the distribution of wealth in a preindustrial society, also because the declarations has survived the centuries almost complete.[39] Moreover, the declarations from the first stage of the implementation of the *catasto*, which included the inhabitants of the city of Florence only, have been digitised and made available online.[40] This online dataset comprises 8,349 declarations of Florentine households.

Despite its unique quality, there are still a few problems with the dataset. First, several groups were not included, for example, foreign mercenaries and the clergy. David Herlihy however assumes that these groups constituted only a small proportion (not more than 10 per cent) of the total population.[41] A second bias is the fact that the *catasto* only includes surplus wealth. The family's primary residential house (including its furniture) and any tools used to exercise their trade were excluded from the assessments.[42] This means that the data are skewed especially at the lower end of the distribution, where such assets would constitute a relatively large part of the household's total wealth. This bias is therefore of less concern for the present study which focuses on the top part of the distribution. Third, land was assessed based on its annual agricultural yield.[43] This means that land, that was more valuable, for example, because it was closer to the city, would be assessed at the same value as less valuable remote plots, if both had the same agricultural yield. As wealthier citizens probably owned a larger share of the more valuable plots closer to the city, this bias probably results in an

37 *P.Oxy.* 44.3169 with Rowlandson 1996: 116–18. *Cf.* Bagnall 1992: 132–36.
38 Rathbone 1990: 120, Bowman 1985: 147.
39 Herlihy 1978: 131–35, Herlihy and Klapisch-Zuber 1985: 1–27.
40 Herlihy et al. 2002.
41 Herlihy and Klapisch-Zuber 1985: 24–25.
42 Herlihy 1978: 134.
43 Alfani and Ammannati 2017: 1075.

underestimation of the inequality.[44] Finally, tax evasion and corruption, as in every society or era, distorts the data to an unknown extent. Despite these problems, the *catasto* data remain an exceptionally valuable source on the distribution of wealth in a preindustrial society.

In order to see whether the top part of these wealth proxy datasets are indeed following the shape of a power-law distribution, I construct a Zipf plot. In such a plot, the number of people holding at least wealth x are plotted against this level of wealth x (both on logarithmic scale).[45] Wealth distributions typically have a bipartite shape in a Zipf plot (i.e., the shape of a hockey stick); at lower wealth levels, they are convex decreasing (implying an exponential distribution), while at higher wealth levels they are linearly decreasing (implying a power-law distribution).[46] The Zipf plots of both the Hermopolite landholdings and the Florentine wealth declarations exhibit this characteristic bipartite shape, as shown in Figure 3.3. This serves as a first reassurance that these datasets are reliably proxying wealth. More importantly, the fact that the largest values of both datasets follow a straight line implies that their upper tails (indicated by the black circles in the graphs) are shaped like a power-law distribution. The fact that the largest values of the two highest-quality premodern wealth proxy datasets are distributed as a power law underwrites Pareto's premise that top wealth generally follows a power law.

It is important to emphasise that wealth distributions with the same functional shape are not necessarily identical. The power-law function has a shaping parameter (denoted as *alpha*) which determines the 'steepness' of the distribution. The power-law tails of the wealth distributions of different societies will have different *alpha*s, reflecting the different level of elite wealth inequality in these societies. For example, the values of *alpha* estimated for the Hermopolite and Florentine wealth proxy datasets are 1.70 and 2.45, respectively.[47] As a higher value of *alpha* implies a lower level of inequality, these estimates suggest that elite wealth inequality was higher in fourth-century Hermopolis than in fifteenth-century Florence.

It might come as a surprise that the largest fortunes in ancient, medieval and modern societies are distributed following the same functional

[44] Herlihy and Klapisch-Zuber 1985: 13–15.

[45] For technical details on this plot, see Cirillo 2013. See also Newman 2005, Chakraborty and Waltl 2018: 49–50, Jenkins 2017: 273.

[46] Drăgulescu and Yakovenko 2001, Coelho et al. 2005, Patriarca et al. 2010: 149, Jenkins 2017: 273.

[47] I use a well-established computational method (Maximum Likelihood Estimation) to estimate the value of *alpha* (Clauset et al. 2009), which has been implemented as a package (poweRlaw) in the statistical software R (version 4.3.3). See also Alstott et al. 2014 and Vermeulen 2018: 379–80.

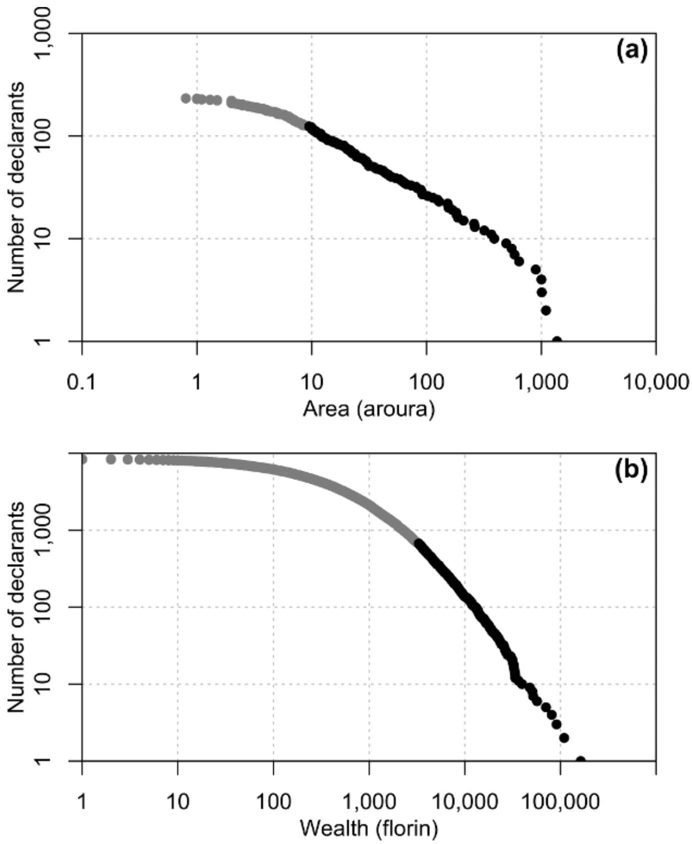

Figure 3.3 Zipf plots of (a) the landholdings registered at fourth-century Hermopolis (N=233) and (b) the declarations of the *catasto* of fifteenth-century Florence (N=8,349).

form.[48] Some might even dismiss this assumption as a fallacious modernist approach to the ancient economy. There are however good reasons to accept the idea that top wealth is generally distributed as a power law.

First, power-law distributions are common in the distribution of both natural and man-made phenomena.[49] Outside the field of economics, famous examples include the intensity of solar flares, the magnitude of earthquakes, the diameter of moon craters, the frequency of words in a text (in most languages), the population size of cities, and the number of citations to scientific papers.[50] This is on its own an extraordinary observation. More excitingly, there is not yet a satisfactory explanation for the emergence of these mathematical regularities.[51] There seems to be a

48 Coelho et al. 2005: 516.
49 *Cf.* Scheffer et al. 2017.
50 Newman 2005, Clauset et al. 2009.
51 Mitchell 2009: 258–72.

https://doi.org/10.1017/9781009496940.004 Published online by Cambridge University Press

connection to fractal-like (self-similar), hierarchical branching network systems.[52]

For the emergence of a power-law tail in wealth distributions specifically, many different explanations have been proposed.[53] For example, Thomas Piketty and Gabriel Zucman argue that random multiplicative factors influencing the intergenerational accumulation of wealth (e.g., the age of death, the number of surviving children at death etc.) lead to a wealth distribution with its upper tail shaped as a power law.[54] Other studies emphasise factors such as the abundance of wealth exchanges, the fairness of the capital market (equal capitals having equal opportunities), preferential redistribution of wealth (larger capitals generating larger profits) and particular saving propensities of the rich (i.e., the rich putting only part of their fortune at stake in economic interactions).[55] Probably all of these factors play a role. What is important for my purpose is that all of the proffered explanations involve general characteristics of economic systems, many of which apply to the ancient economy as well.

It seems that one of the central premises of the controversy between primitivists and modernists (or substantivists and formalists for that matter), which is the degree of 'economic rationality' of the individual, does not play a major role for the shape of the wealth distribution.[56] Even if Weber, Polanyi and Finley were right that there was a big rift between the economic rationality in precapitalist and capitalist societies, a power-law tail might still be expected in precapitalist wealth distributions.

Some statisticians have questioned the applicability of a power law to top wealth. They argue that some wealth (proxy) datasets fail certain statistical tests and that other more complex mathematical functions better fit the data.[57] However, measurement errors can probably explain the statistical ambiguities these critics observe.[58] Moreover, for my purposes it matters less whether Roman wealth distributions conformed perfectly to a power-law function. It is more important that a power-law function is a reliable representation of the top part of the distribution, which stands without doubt.

[52] West 2017: 25–33, Mitchell 2009: 227–57. For the view of the ancient economy as a complex adaptive system, see Poblome 2015.

[53] Gabaix 2009 recently reviewed various theoretical explanations of the appearance of power laws in economic data.

[54] Piketty and Zucman 2015: 1351–60. *Cf.* Piketty 2017: 458–59.

[55] Exchanges: Bouchaud and Mézard 2000. Fair capital market: Solomon and Richmond 2001. Preferential redistribution: Coelho et al. 2005. Saving propensity: Modanese 2016 and Patriarca et al. 2010.

[56] For a good summary of these ideas, see Bresson 2016: 2–15.

[57] Brzezinski 2014, Chan et al. 2017, Ogwang 2011; 2013.

[58] Capehart 2014.

In sum, there are many reasons to accept that the top of ancient wealth distributions can be represented by a power-law function. Apart from the theoretical considerations and empirical examples discussed above, there are also numerous Roman wealth proxy datasets whose tails are shaped like a power-law distribution (see Chapter 8 for a discussion of several examples from Roman Italy).

If we accept that the wealth of the top layers of society can indeed be represented by a power law, the next step is to invert the logic and use a power-law function to predict the top of the wealth distribution. This is a methodology widely used by economists.[59] They employ power laws to impute the top part of modern income and wealth distributions as these parts are notoriously difficult to measure. For example, Paul Eckerstorfer et al. use the upper three deciles of an Austrian wealth survey to shape a power-law distribution, which is subsequently used to reconstruct the top 1 percentile of the survey data.[60] Similarly, Stefan Bach et al. use rich lists from Germany, France, Spain, and Greece to shape a power-law distribution to impute the top layers of the respective national wealth distributions.[61] One of the big advantages of a power-law function is that it is a relatively simple function in mathematical terms (it has only one shaping parameter).

There is however one important proviso for the application of this methodology; only the top of a society's wealth distribution can be expected to be shaped as a power law.[62] This is why wealth distributions have a bipartite shape in a Zipf plot; there is an inflexion point and the power-law assumption is valid only above this point. There is no universal threshold above which wealth is distributed as a power law; values between the top 0.05 per cent and 5 per cent are generally assumed or estimated by econo-physicists.[63] The threshold has to be determined for each dataset individually. The existence of a threshold for the application of the power-law assumption thus precludes the application of this model to wealth in society at large.[64]

A few words on the nature of a power-law distribution are in order here. A power law is a functional relationship between two quantities. This means that a relative change in one of the quantities results in a proportional relative change in the other quantity. For example, consider the relationship

59 Jenkins 2017, Vermeulen 2018, Chakraborty and Waltl 2018, Bach et al. 2019, Eckerstorfer et al. 2016.
60 Eckerstorfer et al. 2016.
61 Bach et al. 2019.
62 Newman 2005: 329–30, Jenkins 2017: 277–79.
63 Jenkins 2017: 277–79, Coelho et al. 2005: 516, Drăgulescu and Yakovenko 2001.
64 *Pace* Kay 2014: 285–97.

between the side length and area of a square; doubling the side length results in a quadrupling of the area (i.e., $2^2 = 4$ – in other words, these two quantities scale by a power of two). A power-law distribution is a so-called skewed heavy-tailed distribution; it consists of many small values and just a few large values (skewed) while these few large values nonetheless contribute significantly to the sum and mean of the distribution (heavy-tailed).[65] It resembles an exponential distribution in that it contains predominantly small values and only a few large values. The power-law distribution is however different from an exponential distribution in that these few large values (its tail) contribute significantly to the aggregate and mean of the distribution (i.e., its tail is heavy).

In their seminal article on the Roman economy, Walter Scheidel and Steven Friesen use a power-law model to predict the distribution of elite wealth in the Roman Empire.[66] It is worth discussing their model in more detail here. They start the reconstruction of their distribution by assuming that all elite households (whose number is based on the sum of the estimated sizes of the elite social groups) owned on average at least HS 125,000 (an estimate connected to the 'standard' curial census qualification of HS 100,000). They then reduce the number of households by a constant factor each time wealth doubles. For example, if there are 100,000 households with HS 125,000, then a factor of 0.5 would mean that only 50,000 households own HS 250,000, and so forth. This procedure is iteratively repeated until the number of households drops below one. At this point the distribution is completed.

The factor with which the number of households drops each time wealth doubles thus determines the shape of the power-law distribution. In order to come to a plausible value for this factor, Scheidel and Friesen construct distributions with three different factors (0.5, 0.67 and 0.75). They subsequently compare three quantities implied by the reconstructed distributions to decide which factor is most plausible: aggregate elite wealth, the income accruing from this wealth based on a 6-per-cent annual return and the number of households with equestrian wealth. Unfortunately, there is no reliable quantitative evidence for either of these three quantities to determine conclusively which of their factors performs best. They conclude that a factor of 0.67 gives the most plausible results. They further substantiate their choice by noting that this factor also results in plausible estimates for the largest wealth of an individual household and the top-centile income share.[67]

65 Alstott et al. 2014.
66 Scheidel and Friesen 2009: 79–82.
67 Scheidel and Friesen 2009: 79–81, with a critique in Bowes 2021: 10–15.

Scheidel and Friesen's shift from a social-table model to a purely economic model is ground-breaking and must be the way forward for the modelling of elite wealth in the Roman Empire. It is therefore worth exploring the results of their power-law model in more detail. They are presented in Table 3.4. These results differ in two ways from the results previously shown for the social-table models. First, the number of households with, for example, at least senatorial or equestrian wealth are model outputs, not model inputs as in the social-table models. Second, the results of the power-law model represent cumulatives; for example, there are about 30,000 households which satisfy the equestrian census, of which about 10,000 own more than 1 million sesterces.

The power-law model suggests a distinctly higher inequality than the social-table models, as indicated by the Lorenz curve, the Gini coefficient and top-centile wealth share. The Lorenz curve is shown in Figure 3.4. It is notably farther from the equality line than those of the social-table models, implying a more unequal distribution of elite wealth. The Gini coefficient at 0.43 is also higher (compare Figure 3.1). Furthermore, aggregate elite wealth predicted by the power-law model is also higher (at HS 71 billion) than that predicted by the social-table model of the same authors (HS 60 billion). Based on this aggregate elite wealth, using the same method and assumptions as for the social-table models as explained above, the predicted wealth share of the top 1 per cent reaches 41 per cent. This value comes much closer to the range observed for France and Britain at the turn of the nineteenth century. All in all, the power-law model of Scheidel and Friesen implies a higher inequality than the social-table models, which seems to paint a more plausible picture.

Finally, it is worth emphasising the remarkable number of households with senatorial wealth predicted by the power-law model of Scheidel and Friesen. According to their model, there were around 10,000 households in the empire that cleared the senatorial census qualification. This result is not noted by the authors but has major implications for our understanding of the political economy of the Roman Empire, as it provides a glimpse of the structural nature of the discordance between wealth and officeholding in the Roman world.

Assuming a power law for the top part of the Roman wealth distribution implies that this wealth distribution was not a historical anomaly. Henrik Mouritsen has very recently argued precisely the opposite.[68] In Mouritsen's view, the wealth distribution in Rome/Italy during the Late Republic was anomalous, missing a middling stratum, which was the result of

68 Mouritsen 2022: 54–56.

https://doi.org/10.1017/9781009496940.004 Published online by Cambridge University Press

Table 3.4 Results of the power-law model of Scheidel and Friesen 2009.

Region	Empire	
Period	150 CE	
Population	70 million	
Total elite households	290,600 (1.7%)	
Total elite wealth	HS 71 billion	
	Wealth (HS 1,000)	Number
Senatorial wealth	1,000	~10,000
Equestrian wealth	400	~30,000
Curial wealth	125	290,600

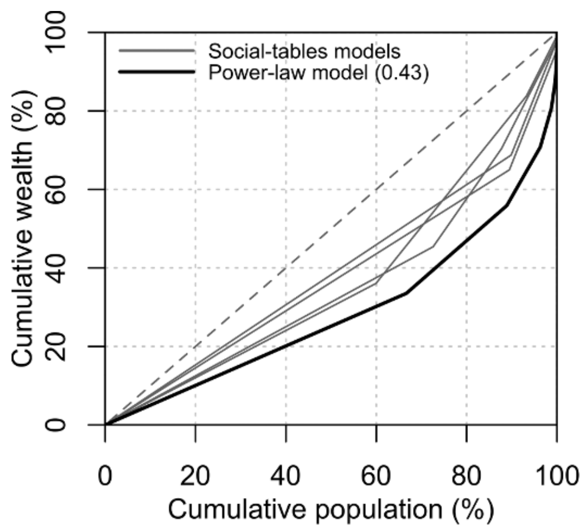

Figure 3.4 Lorenz curve of the distribution of Roman elite wealth as predicted by the power-law model of Scheidel and Friesen 2009: 79–82. The Gini coefficient is given in brackets in the legend.

intensive immigration to Rome/Italy and the ubiquitous use of slaves by the Roman elite. At least for the Imperial period, his argument sits uneasy with a growing body of scholarship arguing for the existence of a substantial middling stratum in Roman society.[69]

[69] See, e.g., Flohr 2017, Zuiderhoek 2017: 106–30, Scheidel 2006a, Kehoe 2015, Haley 2003. For small-scale slaveholdings of (supposedly) middling households in Roman Egypt, see Bagnall and Frier 1994: 48–49.

https://doi.org/10.1017/9781009496940.004 Published online by Cambridge University Press

A final note on the added value of using a power-law function to assess Roman wealth distributions. While the method might *prima facie* appear overly complicated for the deficient ancient evidence, there is an urgent need for more formal analyses of the available data. This is illustrated by the divergent interpretations of the pledges to the two Trajanic *alimenta* schemes. Ramsay MacMullen plots the *alimenta* pledges of both schemes together and concludes that they attest to the extreme 'verticality' (inequality) of Roman society.[70] Walter Scheidel however argues that the distribution of these pledges suggest a 'pyramidal continuum' rather than 'an hour glass-shaped distribution'.[71] Geoffrey Kron finally observes, on the basis of the same data again, that 'extremely large farms were exceedingly rare', suggesting that Roman landholding patterns were relatively egalitarian in comparison with, for example, nineteenth-century England or Italy.[72] A more formal mathematical approach to these datasets and other similar wealth proxies allows for more objective evaluations.

3.3 The Way Forward

The article of Scheidel and Friesen is a landmark in the study of Roman inequality. The basic principle of using an *economic* model to predict *economic* aspects of Roman society has set the direction for all future studies. In this study, I will use a similar power-law model to reconstruct the top of the Italian wealth distribution in the Early Empire. There are however five points on which their model can be improved.

First, the heterogeneity of the Italian *civitates* needs to be taken into account. The bulk of our evidence on the economic and political aspects of the Italian *civitates* stems from a relatively small number of larger towns. The settlement hierarchy of Roman Italy was however very steep, meaning that most towns were small. The existing evidence can therefore not be extrapolated straightforwardly to Italy as a whole. I therefore employ a 'tessellated' approach. This entails first reconstructing the wealth distribution in each individual *civitas* (for this I use the ever-growing archaeological evidence of the inhabited areas of the Italian urban centres, see Chapter 5). Subsequently, I construct the Italian wealth distribution by aggregating all the local distributions.

Second, although the existence of the group of 'other wealthy' is endorsed by most scholars, little effort has been made to systematically investigate

70 MacMullen 1974: 95–97.
71 Scheidel 2006b: 51.
72 Kron 2008: 94.

https://doi.org/10.1017/9781009496940.004 Published online by Cambridge University Press

who belonged to this group and how numerous they were. A first attempt is made here by reviewing the evidence for the number of Italian households with curial wealth outside the councils (see Chapter 7).

Third, the power-law model of Scheidel and Friesen uses wealth groups instead of social groups. Despite the fact that this is a great improvement compared to a social-table model, it is still a discretised approach in the sense that it assumes economic groups whose members notionally possess the same (average) wealth. This means that inequalities within these groups remain unaccounted for. I will use a continuous (non-discretised) model to construct a power-law distribution, which will avoid this problem. The following expression is used to reconstruct a power-law distribution,

$$x_i = xmin\left(\frac{r_i}{N}\right)^{\left(\frac{1}{1-alpha}\right)} \tag{3.1}$$

where r_i is the rank of household i (based on its wealth x_i) and $xmin$ is the wealth of the poorest household of the total of N households in the power-law tail, while *alpha* is a parameter determining the shape of the distribution.[73]

Fourth, my model will be anchored in a broad foundation of ancient empirical evidence. The shaping parameter *alpha* determines for a large part the 'steepness' of the wealth distribution. Scheidel and Friesen accordingly elaborate to defend the choice of their factor (which is mathematically related to *alpha*). I instead use a series of local wealth proxy datasets to estimate values of *alpha*. These datasets include residences sizes, tomb sizes, burial plot sizes and pledges to the Trajanic *alimenta* schemes (see Chapter 8).

Finally, I will take account of the epistemic uncertainties involved in the modelling by using probabilistic calculations. Historical-econometric or cliometric models are notorious for the level of epistemic uncertainty involved in the estimation of the model variables. The relatively low quantity and quality of ancient evidence make these uncertainties even more pertinent for studies on ancient economies. In order to formally account for the plethora of uncertainties, I use probabilistic calculations, drawing on the pioneering work of Myles Lavan.[74]

73 For a derivation of this equation, see Eckerstorfer et al. 2016: 608–9.

74 For a good introduction of this methodology, see Lavan 2019b, Jew and Lavan 2023 and Beven 2009, esp. 49–104. For previous applications of this method to problems in ancient history, see Lavan 2016, 2019a and the chapters in Lavan et al. 2023.

3.4 Conclusions

Social-table models (based on socio-political groups) systematically underestimate the inequality in a society by disregarding within-group inequalities and the overlap of the wealth of members of different social groups. The results of four social-table models for the distribution of elite wealth in the Roman Empire indeed appear to be relatively low in a comparative perspective. The use of a purely economic model, which assumes that the top of the wealth distribution follows a mathematical (power-law) function, is preferable. This type of model has been seminally introduced to the field of Roman history by Scheidel and Friesen. I will apply their model to Roman Italy with five main improvements. These improvements relate to the extrapolation of the sparse ancient evidence, inclusion of the wealthy households outside the socio-political orders, the mathematical formulation of the model, the evidentiary basis of the model and finally the uncertainties involved in estimating the model inputs.

This chapter introduced a series of methods and analytical techniques, some quite alien to the historian. To make this methodological framework more tangible, I will in the next chapter first apply it to the well-known evidence of first-century Pompeii before applying it to Italy as a whole in Chapters 5 to 9.

https://doi.org/10.1017/9781009496940.004 Published online by Cambridge University Press