PETER TEMIN

THE ROMAN MARKET ECONOMY

Chapter 10

Economic Growth in a Malthusian Empire

There is growing evidence that ordinary Romans lived well in the early Roman Empire. The existence of many cities, and particularly the large size of Rome itself, provides indirect evidence of productivity advance. More detailed evidence is emerging of improving agricultural technology, building techniques, manufacturing plants, and land use. The widespread use of African Red Slipware pottery provides evidence that even ordinary people had access to the fruits of all this technology. And the literary evidence supports the idea of prosperity by providing insights into civilized urban lives in Roman cities. As explained in chapter 9, these are all proxies for economic growth. They are not measures, but they are suggestive. How could per capita income grow in a Malthusian economy? This chapter resolves the apparent paradox of economic growth in a Malthusian world.

Even if this evidence of an improvement in general living standards is not yet convincing to everyone, there by now is enough evidence to raise the possibility that such a movement might have taken place. In the spirit of Morley's (2001) essay on the implication of a large Roman population, this chapter offers a way to understand rising per capita incomes in Rome as the accumulating evidence becomes persuasive.

It seems paradoxical that we have evidence of a rise in per capita income when we face so much uncertainty about the size of the population whose income is growing. This is only partly a matter of evidence; it also is a result of our theories. We assume that the scattered evidence of living standards can be generalized to larger groups because we implicitly or explicitly assume that people lived at roughly the same level. This does not mean that we ignore class divisions in Rome, but that we compare the lives of ordinary Romans with their predecessors. We assume this similarity of income for working families because we assume there was a labor market that brought wages into some kind of rough correspondence. I argued in chapter 6 both that we need to

FIGURE 10.1. World economic history in one picture
Source: Clark (2007a, 2)

make such an assumption to make sense of a lot of the modern literature on ancient Rome and that this assumption is an accurate one in the early Roman Empire.

The problem with theories is that there are lots of them. We use the theory of competition between workers to support the idea that living conditions improved in the later republic and early empire, but the Malthusian theory of population change argues that our observations must be false. In that theory, changes in productivity lead to changes in the size of the population (which we cannot observe) and leave the level of per capita income (which we do observe however imperfectly) the same. The implications of this theory can be seen in figure 10.1, taken from a new economic history of the world. The level of per capita income is assumed to be roughly constant before the Industrial Revolution. Saller published a more abstract version of this graph making the assumption of constant per capita income before the Industrial Revolution more apparent. He did not allow for the fluctuations around this level indicated in figure 10.1 (Clark 2007a, chapter 1; Saller 2005). This figure has just reemerged in a new book about a "unified theory" of economic growth (Galor 2011).

The view in figure 10.1 is too reductionist for the analysis of the Roman economy. From the twenty-first century, events before 1800 appear to merge together, and it is useful to summarize them in a simple way. This tendency leads historians to interpret the Malthusian model very strictly, but there are important delays in the Malthusian system. These dynamics are integral to the Malthusian model, and historical evidence from other periods indicate that the delays can be exceedingly long—on the order of centuries rather than decades. These dynamics provide a way to acknowledge growing per capita income in the basically Malthusian world of the early Roman Empire. It also provides a way to ask if the Romans could have escaped—in some alternate history—the Malthusian constraints.

I start this argument by reviewing the evidence that per capita income was high and that it fell as the Roman Empire disintegrated. I then explain how the Malthusian model works with attention to the dynamics of how shocks affect both population size and its income over time. The model reveals how an economy could avoid the Malthusian constraints within this Malthusian model for an extended period. Breaking out—industrializing—is a different matter, and I close with some speculations about the nature of the Industrial Revolution and its relevance for Roman history.

Finley (1973) argued that the ancient economy was dominated by stagnant technology. This view became the one that historians of later periods saw when they looked back at the ancient world. In a book about the history of technology in the Western world, Mokyr (1990) spent a chapter trying to explain Roman technological stagnation.

The tide has changed in the last decade. Roman archeologists have found abundant evidence of new technologies in the Roman world, and their views are now appearing in print. Greene (2000) started the debate with a paper challenging Finley on several grounds. He argued that Finley had been misled by the literary sources; only archeology could show how technology had improved. Wilson (2008) expanded the case for a progressive Roman technology in several papers, and I follow his lead here. I discuss in turn the growth of regional specialization, the expansion of land for agricultural purposes, manufacturing, and construction.

Everyone knows about the Pax Romana. Pompey finished clearing the Mediterranean of pirates in 67 BCE and made safer transport possible. The risks before then are illustrated by the kidnapping of Caesar and its subsequent horrifying effect (Casson 1991, 181–83). The usual interest in this story is Caesar's insolence toward his captors and his subsequent revenge. I emphasize here the risk of capture that is the premise of the story. As North (1968) argued in a paper cited by the Nobel Committee in awarding him the Nobel Prize in economics, reducing the cost of defending against pirates lowers the cost of shipping.

A lower cost of shipping allowed production to be located around the Mediterranean where conditions were most suitable. Instead of growing wheat in Italy, inhabitants of Rome ate porridge and bread made from wheat grown in Sicily, Egypt, Africa, Spain, and other places. Production was scattered because it was more efficient to grow wheat in these places than in Italy. Romans did not calculate what we call efficiency; instead they purchased wheat where it was cheapest. It was cheapest in places where grain agriculture was the best use of land and labor resources (Erdkamp 2005, chapters 2, 5; Rathbone 2009, 322).

The gain to both the exporting and importing regions was shown by Ricardo in his theory of comparative advantage and explained in chapter 1. Trade that allows regions to specialize in what they produce best increases the income of both sending and receiving regions. Trade functions like an extension of resources in any one region; it loosens the constraint of limited land in an agricultural economy. One effect of Roman regional specialization was the shift in Italian farms from wheat to other crops that did not travel as well (Morley 1996; Geraghty 2007).

Roman roads performed a similar service to the economy. They were less important than a peaceful sea because overland transportation was much more expensive than water-borne transport. In addition, the roads were built for military reasons; any commercial use was incidental. The roads nevertheless made it easier for goods to get around. They promoted local specialization similar to the broad patterns just described. They also allowed goods brought from far away to reach consumers living away from seaports. Shipping and roads therefore both promoted a better life for Roman citizens.

They allowed the Romans to create the urban society that was unique in the ancient world. The agricultural system—including agriculture, trade, and coordinating institutions—was efficient enough to release substantial numbers of people from the tasks of growing food. These people could gather in cities, and they could produce other goods and services. These added products improved the quality of Roman life, and they added to per capita income.

In addition to gains from regional specialization, there was technical change in each place. Mokyr (1990, 23) opened his book on technology with a chapter that criticized Finley for being too pessimistic, but he still argued that "new ideas were not altogether absent, but their diffusion and application were sporadic and slow." Recent archaeological discoveries dispute that conclusion.

Terracing was common, extending the range of land on which crops, particularly grapes and olives, could be grown. Wine and oil presses also used the screw, enabling grapes and olives grown on new land to be processed more efficiently. The Archimedean screw was used widely in cereal agriculture to drain land, extending the range of land that could be used for this crop as well. Our evidence is spotty, but recent archaeological discoveries from many different areas suggest strongly that these innovations had diffused over large ranges of the Roman Empire (Wilson forthcoming).

Everyone knows about the Roman use of arches and concrete to construct buildings, roads, and ports. Transportation helps to achieve gains from trade as well as enhancing productivity in other ways. Aqueducts are well known, but the sophistication needed to construct them seldom is noticed. In addition to the construction of the aqueducts, the level had to be adjusted for water to flow over often large distances. Only a little thought is needed to realize that the widespread evidence of aqueducts provides evidence of the diffusion of engineering knowledge around the Roman Empire.

Water wheels also were more prevalent in Roman times than previously thought: "Today, we may state with confidence that the breakthrough of the water-powered mill did not take place . . . in the early middle ages, but rather . . . in the first century A.D., or perhaps even slightly earlier" (Wikander 2008, 149). And mass-produced Roman products were prevalent: "Pottery, glassware, bricks, coins, plate, and humble metal objects such as nails were produced in enormous quantities to standard shapes and sizes, and widely traded around the Roman Mediterranean and northern Europe" (Wilson 2008, 393).

Yet the evidence of technical change is not quite the information we need. Greater production expanded the resources available to feed people, but it could have resulted in more people rather than higher per capita incomes. As noted in chapter 9, population often is used as an indication of early growth over the long run. The Roman Empire lasted for several centuries, but the time frame of the Malthusian model is even longer. The purpose of a model is to show how individual incomes could have changed in the short run.

We need evidence on the consumption of ordinary people to show that they were better off as opposed to more numerous. Any society can support an elite that lives well; richer societies can have large elites. These elites typically are too small to affect the growth of the population as a whole, although the Roman elite was quite large. We need evidence that extends beyond the literary evidence of sumptuous Roman dinner parties and feasts.

Diet is an important part of the consumption of ordinary people and therefore a good indicator of the standard of living before the modern rise of per capita incomes. The consumption of wheat was enhanced by making it into bread rather than porridge. The Roman conquest and the resulting reduction in transportation costs led to an increasing variety of diet, including a wide range of new fruits and vegetables. More important, there is evidence of greater meat consumption in both Roman Italy and the provinces. Meat, of course, is a superior good, and its consumption rises with per capita income. There were more animal bones in the early Roman Empire than in surrounding centuries, and the number of animal bones peaks again around 500 in Roman Italy. Animals were larger in this period than before or after, adding to other suggestive evidence of improved diets and higher incomes (Jongman 2007b).

Ordinary people also had consumer durables that were better than those before or after. The most prevalent was African Red Slipware, ordinary pottery that is found everywhere in Roman settlements. The pottery was wheel-thrown and highly fired, supplying a "modern" platform on which to eat. In addition to plates, ordinary people had iron knives with which to cut their food. Good Roman pottery contrasts sharply with the friable pottery found in post-Roman Britain that was neither wheel-thrown nor highly fired. The comparison shows the decline in the living standards of ordinary people after the Romans left Britain. There must have been a previous rise in per capita income for it to fall sharply thereafter (Ward-Perkins 2005).

The omnipresent oil lamp was another consumer durable of Roman times. It extended the day and enhanced the quality of life in interior spaces for many people. The assemblages of oil lamps in many museums show their spread throughout the empire and the standardization that reveals their industrial origin (Harris 1980). Like agricultural goods, industrial goods were made in centralized locations and shipped all over the Roman world.

Evidence of widespread improvements in consumption is increasing, and Roman citizens must have had increasing incomes to buy the enhanced food and consumer durables. Jongman (2007b) cited a variety of estimates showing that real wages increased in the late republic and early empire. He surveyed the occasional evidence of documented wages, subsistence annuities, and slave prices—as an index of wages of free workers with whom Roman slaves competed. The data for any one of these measures are spotty, but the pattern of all of them is the same. This common pattern suggests that the occasional observations are capturing underlying trends whose existence is attested to by the variety of evidence that fits the pattern. Real wages rose after the Antonine Plague. Labor income was the major part of total income in the agrarian society of ancient Rome, and an increase in real wages is a good index of an increase in total income.

Allen (2009a) used data from the Diocletian Price Edict of 301 to compare real wages in Rome with those in early-modern European and Asian cities. He found that the real wage calculated from the Price Edict was close to the real wage in Florence in the eighteenth century. This is impressive for an ancient society, but it also is less than real wages at the same time in London and Amsterdam and less than Florentine real wages in the century after the Black Death. Scheidel (2010) replicated Allen's estimations for Roman Egypt.

The evidence for enhanced consumption is still very sketchy, and we hope that archaeologists will broaden the evidentiary base over time. Enough evidence has been found already to indicate that ordinary Romans lived better than ordinary people before or for many centuries after. The problem is how to square these observations with the iron law of subsistence living that is part of the Malthusian model of population change.

In particular, this evidence suggests that living standards for ordinary Romans improved in the late republic, reaching a high standard for the early empire. Given the long history of the republic, this growth did not have to be rapid to result in a substantial increase in living standards. It did, however, need to be sustained over the course of a century or more. How could these innovations result in rising living standards rather than simply more people? We need to examine the Malthusian model of population dynamics to see.

Malthus argued that the size of the population was limited by the resources available to feed it. By resources, most people now mean land, understanding that the use of land and other resources may be relevant as well. This Malthusian relation is known to economists today as the declining marginal product of labor when the number of workers on a given plot of land increases. This, of course, was Ricardo's way of making the same point at approximately the same time as Malthus, and it is Ricardo's formulation that has become central to modern economics; as the number of workers rises, wages fall and rents rise (Malthus 2004).

Wages in this summary mean "real wages," that is, the purchasing power of wages as described by Allen (2009a) and Scheidel (2010). The diet of most workers near a Malthusian equilibrium consisted largely of grain in one form or another. We therefore approximate the real wage by looking at the ratio of the money wage to an index of the price of goods workers bought, weighting grain heavily in this index. If we divide money wages by the price of grain alone, we get a measure of the marginal product of labor, since farmers typically hire workers up to the point where the last (marginal) worker produces just enough grain to pay for the wages he earns (Clarke 2007b).

Ricardo's formulation shows the need for an additional relation for Malthus to find an equilibrium point on this line. Malthus did this by specifying a relation between worker's wages—taken to be their income—and their birth and death rates. Births rise with income, both as nutrition rises and as younger marriages become feasible. Death rates fall with income as infant mortality declines and plagues, wars, and pestilence become less frequent. Modern research has confirmed the first of these relations, while generally failing to find convincing evidence of the latter (Lee 1980). For most purposes, only one relation is needed, provided the other one does not operate in the reverse direction. The full Malthusian model then was taken to restrict the range of early history. Wrigley's (1988, 29) description of the world before the Industrial Revolution is clear: "An organic economy, however advanced, was subject to negative feedback in the sense that the very process of growth set in train changes that made further growth additionally difficult because of the operation of declining marginal returns in production from the land." Clark's (2007a, 27) recent description is equally clear: "Anything that reduced the death rate schedule—advances in medical technology, better personal hygiene, improved public sanitation, public provision for harvest failures, peace and order—reduced material living standards."

The preceding discussion is summarized in figure 10.2. The horizontal axis on both graphs is the same: per capita income. The top graph shows the determination of population size. Population grows if births exceed deaths; it falls if births fall short of deaths. The equilibrium is where the birth and death rates are equal, at y^*. The bottom graph shows that the resource constraint permits only a limited population size, n^*, at this income. Note that the model works well even if there is no relation between income and the death rate. If the curve marked D in the top graph is horizontal—that is, if it shows the death rate unaffected by changing income—y^* is still the equilibrium, and the analysis proceeds as before (Lee 1980; Clark 2007a, chapter 2).

Consider the effect of a plague, like the Antonine or Justinian plagues, in this model. Let us assume that the population fell by approximately one-third, without aiming for spurious precision. The effects are shown in figure 10.3. Population fell from n^* to n_1. As population fell, income rose above the previous equilibrium income, y^*, because the marginal product of labor rises as population falls. At this higher income, n_1 in figure 10.3, birth rates exceeded death rates, as shown in the upper graph. Population grew as a result. It continued to grow until income was reduced to its previous level, y^*, where births and deaths were once again equal. As can be seen in the lower graph, per capita income was unchanged at the new equilibrium, and the population returned to its former size.

How long did this process take? From the long-run point of view of figure 10.1, this may not be important, but if it took a long time to return to y^*, then per capita incomes may have been above that equilibrium level for some time. According to Solow (2007, 39), Nobel laureate in economics, "The Malthusian process works itself out slowly. The chain of causation . . . could take years or decades to complete itself." Is this accurate, or is even this casual estimate of the delays too short?

We do not have much evidence from the Roman plagues, but we know more about the aftermath of the Black Death of 1349 in England. Immediately after the plague, money wages of farm workers shot up. As I argued for the Antonine Plague in chapter 4, the immediate effect of a plague is inflation. Wages and wheat prices both rose. Yet for the Malthusian model we need to know the path of real wages, that is, the extent to which the rise in money wages exceeded the rise in the price of grain and other consumables.

Clark (2007b) provided detailed information about the progress of English real wages after the Black Death. He showed that real wages did not rise nearly as fast as money wages in the immediate aftermath of the plague. Instead they

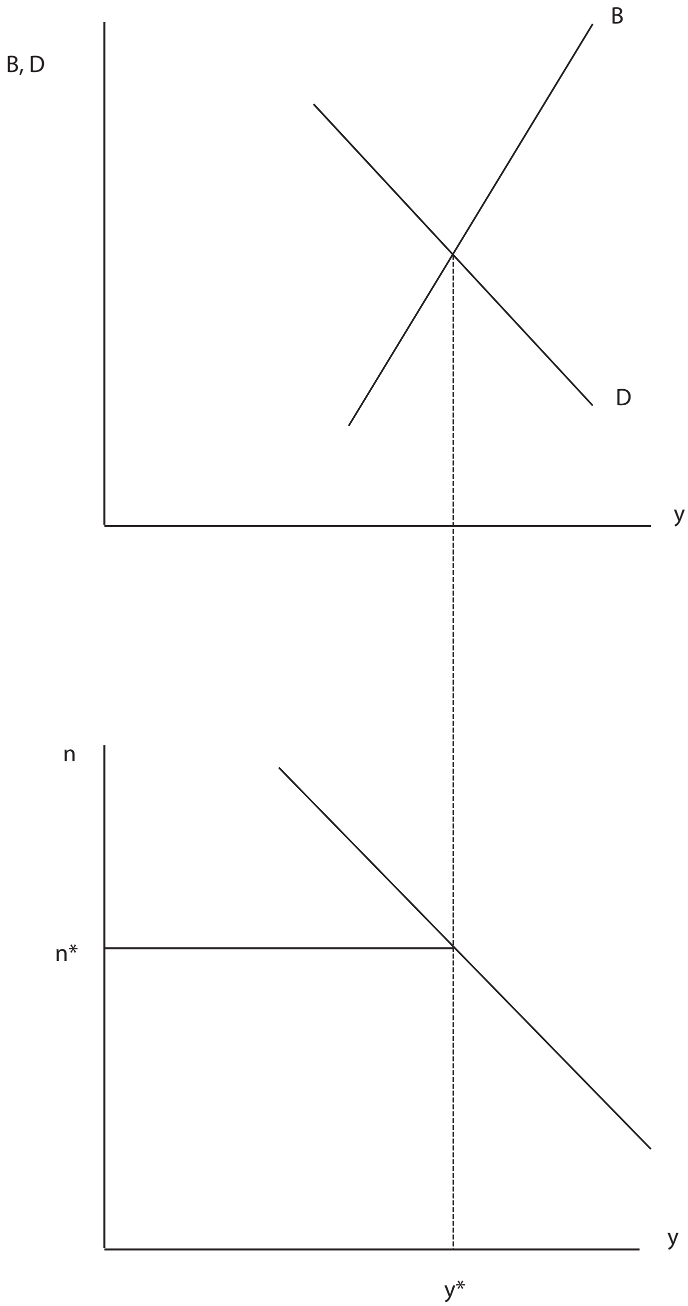

Figure 10.2. The basic Malthusian model

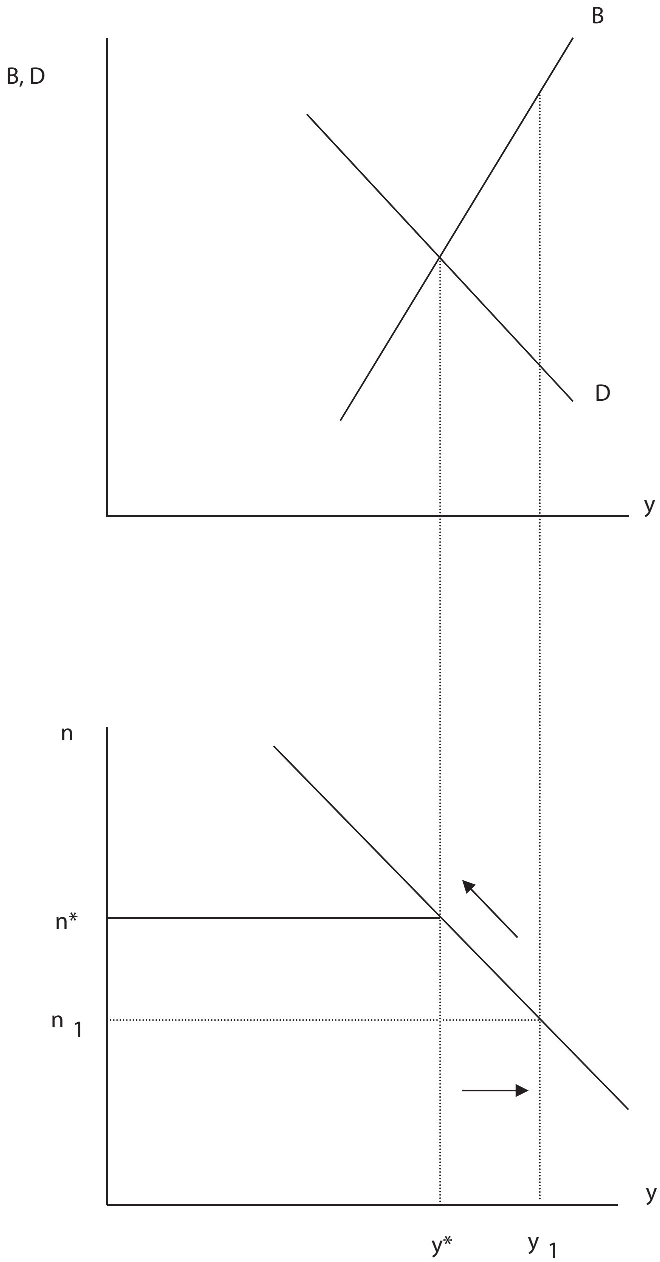

Figure 10.3. Effect of a plague

rose gradually and peaked a century after the plague, in the middle of the fifteenth century. Malanima (2007; 2009, 264) replicated this finding for Italian agricultural real wages, showing that they also peaked around 1450. Allen's estimated English real wages, reported by Malanima for comparison, peaked even later than shown in Clark's data, perhaps as late as the end of the fifteenth century. It is clear that the population did not recover to preplague levels for over a century. It took a very long time for the Malthusian system to return to its equilibrium. Clark was interested in the return to the equilibrium—as shown by figure 10.1—while I am interested here in the deviations from it.

It cannot be surprising that the return to the Malthusian equilibrium took a long time. Initially, the disruption of the plague delayed economic adjustment to the new factor proportions. As discussed in chapter 4, plagues lead also to inflation, explaining why nominal wages rose immediately after the Black Death. Only gradually did farmers take advantage of the increased land per worker and increase their incomes. Higher incomes after the Black Death may have resulted in earlier marriages, which in turn led to more children. But it took a generation or more for the effects of this change to become apparent in the agricultural labor market. If women changed their behavior slowly, it might take several generations to lead to population expansion. And when we start talking about generations, it requires only a few generations to make a century of delay. While we do not know much about family dynamics in the late fourteenth century, we do know that real wages did not start to fall until a century after the Black Death.

This is consistent with the limited evidence from the Antonine Plague. Scheidel (2002) collected fragmentary wage and price data from Roman Egypt in the second and third centuries. The ancient sources are not frequent enough to provide the detailed timing evidence found in the medieval data, but they suggest a long period after the plague when wages were high. If we regard the observations as random draws from records of wages and prices in the two centuries, we are implicitly assuming that the effects of the population decline in the Antonine Plague lasted as long as the decline after the Black Death. According to Scheidel, however, the rise in the real wage, that is, the ratio of wages to commodity prices, was smaller in the ancient world. Real wages were less than half again as large in the third century as in the second century, while real wages peaked at twice the preplague level in the fifteenth century. More ancient evidence would help us calibrate both the timing and magnitude of the ancient shock.

As a matter of logic, real wages had to rise as part of the demographic process. The question for ancient history is not whether individual incomes grew, but rather how much and for how long. Scheidel's (2010) recent estimates of real wages in Roman Egypt fail to show any rise following the Antonine

Plague. The difference appears to come from the choice of deflator, whether wheat alone is used (to maximize the number of historical observations) or a basket of consumption goods (to maximize the fit with the modern methodology described in chapter 9). The best view at the current time is that there was a significant demographic event called the Antonine Plague, although its economic effects still are only dimly seen and apparently more modest than those of the Black Death.

Consider now a different shock to the Malthusian system. Instead of assuming that the size of the population changed, assume that the Malthusian resource constraint shifted outward. This change could come from regional specialization permitted by the Pax Romana. It could come from technological change that allowed land to be used more efficiently as described by Wilson. In any case, it shifts the line in the bottom graph of figure 10.2 to the right. For any given population size, the available land now allows the marginal product of labor and income of farm workers to be higher than before.

As shown in figure 10.4, this sets up a population expansion. In the short run, the effects of this positive shock are the same as the results of a plague shown in figure 10.3, but for different reasons. The population changed dramatically during a plague, but it changes more slowly under normal conditions. Per capita income can change more rapidly, and it increases in the short run, leaving population unaffected. But in the longer run the equilibrium has changed. The excess of births over deaths causes population to rise. Equilibrium is reached when income returns to its previous level in the upper graph, y^*. Looking at the lower graph, we see that the population is larger at the new equilibrium than before, at n_2 instead of n^*. The effect of technical change has been to increase the size of the population, not per capita income, in the Malthusian equilibrium.

Note the differences between figures 10.3 and 10.4. In both of them, income rises, setting off an increase in population. This is a rightward shift in both graphs. Although population grows in both graphs, the relation of this growth to the prior level of the population is different. In figure 10.3, the population is always lower than n^*, and the growth is only to regain the losses from the plague. In figure 10.4, by contrast, population is always larger than n^*, as technology allows for a larger population. If the shift in the resource constraint is a one-time movement, then the population settles down to a new equilibrium level, n_2, larger than n^*.

As before, the economy will not move instantly to this new equilibrium. It will take a long time, perhaps more than a century. During that time, per capita income will be high and population will be growing. If the new technology diffuses slowly or perhaps continues to improve, then the resource constraint curve will continue to shift outward for a while instead of simply jumping from

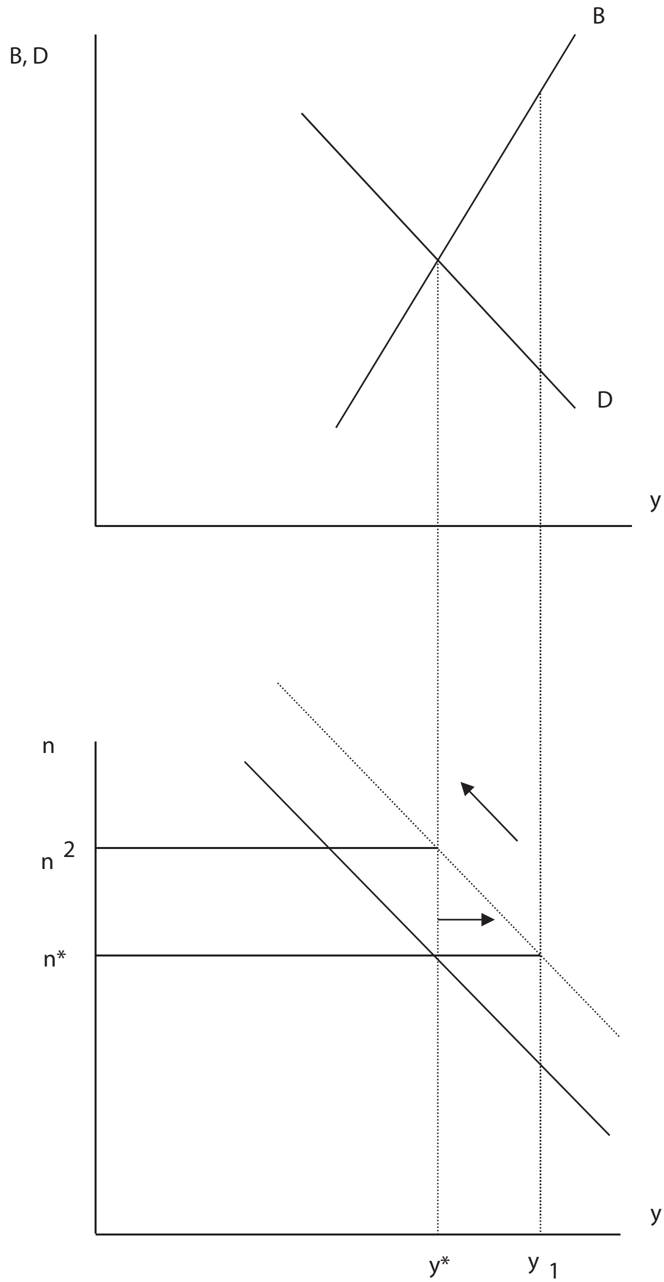

Figure 10.4. Effect of technical change

one position to another as shown in figure 10.4. In that case, both incomes and population will continue to increase for quite a while before the pull of the Malthusian equilibrium is felt. If the resource constraint continues to shift outward for a while, then income can stay above y*, Malthusian subsistence, for more than a century. If productivity continues to advance indefinitely, income can stay above y* indefinitely.

The two possibilities just mentioned can be stated as two competing hypotheses. The first hypothesis is that there was a one-time increase in productivity that had effects that gradually died out during the early Roman Empire. The second hypothesis is that there was continuing productivity growth in this time that was interrupted as the empire became less stable in the third and succeeding centuries. The long delays in the Malthusian model make differentiation between these two hypotheses difficult, but it is important to make the effort to distinguish them and to understand the nature of the Roman economy.

The two hypotheses relate to the most probable causes of increasing income. The first hypothesis of a one-time increase in productivity sees the productivity increase coming from the increase in Mediterranean trade promoted by the Romans. Making shipping safer and introducing regular sailings lowered the cost of moving even heavy and bulky goods around the Mediterranean. This allowed areas around the sea to specialize in their most productive activities and sell their products elsewhere for consumption. I showed in chapter 2 that the uniformity of wheat prices around the Mediterranean argues for a single Mediterranean market where the production of wheat could be allocated to the most productive localities.

The second hypothesis of continuing productivity growth sees this growth as coming from the improvement of technology. In this case, there is no single test of change, but rather the accumulation of evidence for technological change. Agriculture was the largest economic activity in ancient times, and improvements in agricultural productivity would have had the most impact. These changes would have required fewer agricultural workers to feed the population, allowing for the urbanization that is such a feature of the Roman world. Productivity advances in other economic processes would have had less impact, but the accumulation of productivity changes in many aspects of the economy would have increased Roman incomes.

Differentiating between hypotheses one and two is made difficult by the coincidence of shocks to the Roman economy. One might think that examination of living standards in the third century would be a way to distinguish between the two hypotheses. But in addition to the effects of technological change, there were also the effects of the Antonine Plague. In other words, the changes shown in figures 10.3 and 10.4 were superimposed on each other in the third century. The net results on population are ambiguous, but the effects on per capita income go in the same direction. This can be seen from figures 10.3 and 10.4, where both shocks increased per capita income. As a result, it will be hard to know if prosperity in the third century was the result of improving technology or declining population.

At this point in our explorations, the data are too sparse to indicate whether productivity growth was decreasing as the transition to the higher level was completed (hypothesis one) or was continuous before some kind of collapse (hypothesis two). I have tried to indicate what kind of evidence would be needed to make such a discrimination, but it may be hard to find enough data to differentiate between these two views. This lack of evidence has not kept ancient historians from debating the shape of ancient economic growth, as can be seen in two graphs from Manning and Morris (2005). The first reported Morris's (2005) data on housing sizes in Hellenistic Greece as a rough proxy for per capita income. Saller (2005) showed in a second graph an estimate of per capita income in the Roman era inferred from data indicative of trade. They both showed economic growth, but the first showed an increasing rate of growth while the second showed a decreasing rate.

Why should ancient historians be concerned about something as arcane as a different second derivative? Because these two graphs express sharply contrasting views of ancient economies. Morris's graph shows accelerating growth. Since it did not continue, it must have been interrupted by some dramatic change. Saller's graph shows decelerating growth that petered out gradually, without drama. The latter view seems more appropriate to a Malthusian process, but only in the long run; Malthus was clear that population checks could come quickly from wars or plagues.

Scheidel (2009a) made a case for hypothesis one in Roman times, a one-time increase in productivity, coming from the expansion of the Roman Republic to incorporate the whole Mediterranean. (He did not draw a speculative graph, but his argument is consistent with Saller's.) Scheidel stressed the timing of indicators of economic growth. Without explicit reference to Morris (2005), he argued that hypothesis two implied accelerating or at least continuing economic growth up to some catastrophe. He then marshaled the various indicators of economic growth described in this and the preceding chapter to argue that they peaked around the beginning of the Roman Empire. This implied that hypothesis one was correct; in his words, "a scenario of 'one-off' unsustainable growth and Malthusian pressure" (Scheidel 2009a, 69).

Scheidel presented his case forcefully, but he admitted that the underlying indexes of growth were uncertain. Jongman's (2007b) estimated consumption had different timing than the other indexes, and the shipping data are not as precise as they seem. Changes in technology could have altered the frequencies with which we find shipwrecks two thousand years later. This can be seen with a simple equation.

$$\text{Trade} = \frac{\text{Trade}}{\text{Voyages}} \times \frac{\text{Voyages}}{\text{Voyages with heavy cargoes}} \times \frac{\text{Voyages with heavy cargoes}}{\text{Shipwrecks}} \times \frac{\text{Shipwrecks}}{\text{Known shipwrecks}} \times \text{Known shipwrecks}$$

This long equation expresses an illuminating tautology. It shows the volume of trade, called simply "Trade" in the equation, as the product of a series of ratios. If you cancel all the magnitudes that appear in both the denominator and numerator of a ratio, the equation says only that the volume of trade equals the volume of trade. The shipwreck index is a good indicator of trade only if all the ratios in the equation stay constant. If we suspect they were not constant, the equation allows us to structure our investigation of changes in pursuit of a better index of trade.

For example, the ratio of voyages to those with heavy cargoes can change for many reasons. The most obvious is the nature of containers. Amphoras are heavy and durable. Ships containing amphoras will sink, and the amphoras will stay intact even as the ships themselves disappear. Late Romans learned to use barrels instead of amphoras to ship liquids. This was a gain to the porters who carried these containers, but it was a loss to archaeologists who could not recover evidence of barrel shipments. A rise in the ratio of total voyages to those containing heavy amphoras will change the ratio of known shipwrecks to the volume of trade. Put differently, a decline in the number of known shipwrecks might be an index of a decline in trade or of technological progress that reduced the dead weight being carried around in the form of amphoras. The shipwreck index may be a more accurate indicator of trade volumes during its rise in the Roman Republic than it is during its decline during the Roman Empire.

The ratio of voyages with heavy cargoes to shipwrecks is affected by the skill of ship captains. If the captains learned how to improve their navigation or weather prediction, this ratio might have changed. The ratio of shipwrecks to known shipwrecks similarly is affected by the skill of archaeologists in dating shipwrecks. McCormick (forthcoming) reported progress in this dimension, arguing that we now can date shipwrecks previous generations could not.

Wilson (2009a) argued that the indexes that Scheidel (2009a) used to date Roman economic growth were too fragile to bear this weight. He discussed the problems with the shipwreck data just noted and compared the overall index with an index of shipwrecks carrying stone cargoes. This rough device to correct for the bias in the index provided a different time path for the relevant wrecks. Other proxies were similarly flawed or ambiguous, and Wilson argued that the sum of the evidence was inconclusive. He concluded, "The apparent convergence of the proxies is very misleading" (Wilson 2009a, 71).

It is likely that both processes were in operation in the early Roman Empire. More evidence may produce a more precise accounting, but there does not seem to be evidence as yet or agreement that hypothesis one or two is correct while the other is wrong. Economists create these oppositions for intellectual clarity, but history has a way of splitting the difference or revealing a more complex story than either extreme hypothesis. At this stage, it is clear that there was a one-time gain from comparative advantage as the Romans promoted Mediterranean trade (see chapter 2). A unified Mediterranean market promoted regional specialization and associated income gains. There also were improvements in the technologies of agriculture and other economic activities. Wilson reminds of us of what we might call "hard" technologies that leave archaeological remains. To that should be added the "soft" technologies of attitudes and markets described in previous chapters that also contributed to economic prosperity and growth. We do not know the full distribution of these changes or the timing of their spread, but it is clear that much more was going on than simply unifying the Mediterranean world.

Whichever hypothesis is more correct, neither of them implies that economic growth could have continued from Roman times until today. Without industrialization, the Malthusian constraints described in the model of this chapter still held. They held loosely, allowing centuries of economic growth under favorable conditions, but they eventually would constrain this growth. The question therefore is not whether Malthusian constraints were present, but rather what changes in Roman times led to growth within these constraints and how far growth went. There were many shocks in the early centuries of this era, from the Antonine Plague to inflation, political instability, and invasions (see chapter 4). The purpose of this simple Malthusian model and the discussion it has generated is to clarify the complex interactions between these historical events.

We know that England began to industrialize in the late eighteenth century and that the English innovations spread throughout Western Europe in the nineteenth century. Why didn't a similar process happen in ancient times? The Malthusian model described here cannot answer such a complex question, and no answer will be forthcoming here. The model lets us understand what is similar and different in the two periods, sharpening the question. A more detailed model of technological change then is needed to compare ancient Rome and early modern England.

As described earlier, technological change can expand the resource constraint and allow both population and per capita incomes to rise. If technological change continues for a while, incomes can remain high while population rises. There are two reasons why population continues to rise. The technological change continues to shift outward the resource constraint, as seen in figure 10.4, which allows population and income to increase. In the short run, income increases, which then allows births to exceed deaths, which is how population increases. These two mechanisms can provide a stable situation where incomes are higher than Malthusian subsistence, y^*, and population is growing. The change from a static resource constraint to an expanding one has resulted in the growth of per capita incomes from y^* to its new level, accompanied by population growth.

A constant rate of technological progress therefore leads only to limited economic growth. Growth becomes a transitory phenomenon as the economy moves from a static Malthusian equilibrium to a dynamic one based on continuing technical change. This appears to describe the growth in real wages observed in preindustrial London and Amsterdam as the growth of trade in the seventeenth and eighteenth centuries expanded the English and Dutch resource constraint. Real wages in most European cities fell in this time as population grew, but not in these progressive cities. By the eighteenth century the difference between the high-wage cities and the low-wage cities on the European continent was about two-to-one (Allen 2001).

A small modification of the Malthusian model allows it to incorporate continuing economic growth. Instead of assuming that the rate of expansion of the resource constraint is constant, assume that it is proportional to the size of the population. In other words, the rate of technological progress rises when population rises. This is the assumption used by Kremer (1993) in his analysis of population growth for the last million years, by Diamond (1997) in his description of how the Neolithic Revolution set the stage for the Industrial Revolution, and by Galor (2011) in his presentation of his unified growth theory. I show in an appendix to this chapter that this assumption allows incomes and population to rise indefinitely.

This process describes the path of the Industrial Revolution. Technological improvement started to accelerate in the late eighteenth century and continued to increase as population also increased. Later, about a century after the start of industrialization, the rate of population growth diminished as the Demographic Transition took place. Women who had education could see the value of education for their children in the increasingly industrial world. They opted to have "better" children, that is, children with education and therefore "human capital," in place of having more children. In this context, the Malthusian model no longer provides many insights into the paths of industrial societies, as described in chapter 9.

This avenue was not open to Rome, at least not in the form that the

Industrial Revolution of the eighteenth century took place. Allen (2009b) showed that industrialization began in England as a result of two forces. The first was high wages. Both England and the Dutch Republic had high wages as a result of their profits from international trade. Roman Italy shared this prosperity two millennia earlier as a result of war booty and the profits from Mediterranean trade, but only England had the Industrial Revolution. Allen argued that high wages needed to be coupled with cheap power to generate industry. The coal industry of England shipped coal from the west and north of England to London. The price of fuel in London was about the same as in other advanced economies, but the price of fuel closer to the source was uniquely low. It was "the cheapest energy in the world" (Allen 2009b, 97).

The uniquely high ratio of wages to power costs gave rise to the Industrial Revolution, not only in England, but in only a specific part of England. Steam engines and iron technology improved in the north and west of England where coal was dirt cheap. The high ratio of wages to energy costs allowed England to produce goods that were competitive with goods produced elsewhere in Europe despite the high English wages. It explains why the first industrialization took place in England rather than in Holland or France. "The cheap energy economy was a foundation of Britain's economic success. Inexpensive coal provided the incentive to invent the steam engine and metallurgical technology of the Industrial Revolution" (Allen 2009b, 104).

The high ratio of wages to energy costs was not only absent in eighteenth-century continental Europe; it was absent as well in the Roman Empire. Despite the technical progress being made then that we are discovering more about, there was no possibility of escaping from the Malthusian constraints with the price ratios that existed then. However prosperous Rome may have been, it was not on the verge of having an Industrial Revolution. There was no analog of the British coal industry in antiquity and therefore no possibility that industrialization could have begun in the ancient world.

Even under hypothesis one, Romans in the early empire appear to have been living well. If this improvement came from the increase in trade, then the residents of Roman Italy may have been similar to the English and Dutch in the seventeenth and eighteenth centuries. Allen (2003) showed that the rise in Atlantic trade increased real wages in those countries before industrialization. Trade did the same for Roman Italy, and it may also, following Ricardo's analysis, have raised real wages elsewhere in the early Roman Empire. I pursue this further in chapter 11.

Any model is a simplification, and the one explained in this chapter is no exception. The exposition here does not aim to capture all the details of this economic expansion. Instead it provides a framework in which the details that emerge from various kinds of research can fit. It provides a mechanism to turn the odd facts gathered by archaeologists into a coherent picture. With more information, perhaps stimulated by having such a framework, we can aspire to constructing a more detailed model. A recent book stated, “Ancient economic history remains a largely undertheorized field of study” (Banaji 2001). I filled a small part of this lack by analyzing a simple Malthusian model. The model is designed to explain how per capita incomes could have grown in a predominantly Malthusian world. This is not possible in equilibrium, and this paper is about the behavior of this model out of the well-known Malthusian equilibrium.

There are several benefits of using a model like this one. Most important, it allows us to integrate recent archaeological discoveries about Roman technology into a coherent view of the Roman economy. It helps us resolve an apparent conflict between the observations we are accumulating about the good life of ordinary Roman citizens or at least structure our disagreements (Scheidel 2009a; Wilson 2009a).

The model also allows us to engage in a structured discussion of alternative histories, what economic historians of the modern world call counterfactuals. What would have happened if the western Roman Empire had not collapsed? We will never know. This model allows us to speculate in a coherent way about what might have been. The comparison with the Industrial Revolution showed an alternative history—about a far different time and place. It is clear that Rome could not have gone that way even if various other factors had been different. The Malthusian model will not help us identify which of those other factors were more important than others; it will help us understand the consequences of economic decline.

Even this simple model helps define questions for Roman archaeology. Hypothesis one above is that there was a single spurt of productivity change whose effects were gradually eroded by Malthusian pressures. Hypothesis two is that Roman productivity growth continued until some unrelated factors inhibited it, allowing living standards to stay high for a longer period. We need more detailed evidence than now exists to make this differentiation.

Finally, the use of this kind of model provides a bridge to help unify the study of ancient and more modern history. At the very least, it can integrate the analysis of ancient plagues with that of more modern ones. It can help us redraw graphs like the one shown in figure 10.1 to reveal accomplishments of people who lived long ago that have been largely forgotten by modern historians. And it raises questions about modern history as much as about ancient history. For all of the factors that doomed the Roman Empire must have been missing or at least modified two millennia later when the Industrial Revolution took place. A model that structures our discussion allows us to place ancient economic history into the general study of economic history.

Chapter 11

■ ■ ■

Per Capita GDP in the Early Roman Empire

Economic growth for a sustained time was possible in a Malthusian world (see chapter 10). There is indirect but compelling evidence that economic growth took place in the late Roman Republic and early empire. Not rapid growth of the modern era, but growth that appears to have progressed slowly over a few centuries. How far did this economic growth raise Roman incomes? Was the standard of living raised above its level before and after? If growth took place in Roman Italy—with its large cities and gains from war booty—did it extend to the provinces? These are questions for this final chapter.

Hopkins (1980) asked how large the Roman economy was in order to find out if taxes in the early Roman Empire imposed a great burden on the population. He argued that the Roman economy in the first two centuries of our era was large enough that the tax burden was light. He did not aim to perform a complete calculation, and he spoke of his result as a "metaphor" rather than an estimate. This classic paper has been widely quoted, and the estimated national expenditure has been quoted without the rest of the argument.

Goldsmith (1984) published an estimate of the gross domestic product (GDP) of the early Roman Empire as part of his exploration of the structure of historical economies. He was an economist like me, and he was seeking applications for the tools and skills he had developed for the analysis of modern economies. His estimate was approximately twice as large as Hopkins's.

I compared these two quite different estimates in Temin (2006a). I asked if the divergent treatments accorded with modern economic evaluations of national production as explained in chapter 9 and whether it is preferable to construct an estimate of total production from individual consumption of wheat or whether another approach is preferable.

My exploration gave rise to several other estimates of Roman GDP (Maddison 2007; Bang 2008; Scheidel and Friesen 2009; Lo Cascio and Malanima 2009a, 2009b). I survey this small literature in this chapter, distinguishing what

I call the first-generation estimates made by Hopkins and Goldsmith in the 1980s and the second-generation estimates of the last five years. The earlier group was concerned with comparisons within the Roman Empire: the distribution of income between the state and the people (Hopkins) and between the rich and the poor (Goldsmith). The later essays, stimulated in part by the kind of studies reported earlier in this book, were concerned with the comparison of the Roman Empire with more recent economies.

External comparisons have animated many of the previous chapters, and ancient historians disagree about them. It is hardly surprising that these disparate views of the Roman economy find their way into the abstruse calculations that support any estimate of GDP. The central question for this chapter is whether these views affect the presumptively objective calculations. In other words, do GDP calculations affect our views of the Roman economy or are they the result of our views?

All of the GDP estimates, mine included, rest on an exceedingly narrow evidentiary base. They are at best conjectural estimates based on a few observations, some about the early Roman Empire and some about modern economies. Not only are the assumptions needed to build on these observations open to question, but the observations themselves are subject to unknown errors. The whole exercise can yield only what Hopkins called a metaphor. The very shape of the metaphor is taken from other indications of the Roman economy, and GDP estimates reflect back these other indications. The devil is in the details, which in this chapter means the assumptions. They are the primary focus of the chapter.

The two early estimates are shown in tabular form in table 11.1. Hopkins dated his estimate quite broadly, from 200 BCE to 400 CE; it presumably represents an average of this long period. Goldsmith dated his estimate at death of Augustus. There is enough parallelism in the two estimates to place them in the same table, but the steps used by the two authors differ in some respects. The result is that several rows in table 11.1 have entries only in one of the two columns. Nevertheless, the table reveals clearly how two authors could start from almost the same data and end up with starkly different estimates.

Hopkins and Goldsmith worked from the same estimate of population size in the early Roman Empire. This makes comparison of the estimates easy, but it would not matter if they had differed in this dimension because the estimates were derived by blowing up per capita figures. Scheidel and Friesen (2009) assumed that the population was different than these authors assumed, and I will discuss later why this assumption is revealing.

Hopkins and Goldsmith also used the same estimated size of the modius, although this size figures only in Hopkins's estimate. This is only important because Hopkins started his estimates with kilograms and converted them to

Table 11.1.
Estimates of Roman GDP from the expenditure side

Category	*Hopkins (1980)*	*Goldsmith (1984)*
Population	54 million	55 million
Kg per modius	6.55	6.55–6.75
Price of wheat	HS 3 per modius	HS 3 per modius
Annual consumption of wheat or measured in wheat	250 kg	37.5 modii = 246–253 kg
Allowance for seed	83.3 kg	
Value of wheat production	333.3 × 3/6.55 = HS 153	37.5 × 3 = HS 112.5
Annual food expenditure		HS 112.5 × 1.8 = HS 200
Annual private expenditure		HS 200 × 1.75 = HS 350
Total per-capita expenditure	HS 153	HS 350 × 1.1 = HS 380
Total national expenditure	HS 8,244 million	HS 20,900 million
Hopkins (1995/6), from minimum to actual expenditure	1.5 × HS 8,244 million = HS 12,500 million	

Note: Figures may be rounded to agree with sources.
Source: Hopkins, (1980; 1995/6); Goldsmith (1984).

Roman units. The two authors abstracted the same price of wheat, in sesterces (HS) per modius, from the literature. This price represents an informal average of prices in Italy; Duncan-Jones (1982, 146) said they generally were between two and four sesterces per modius. We know from Cicero's Verrine orations that HS 3 was an official wheat price in Sicily, but we also know from scattered evidence that the price of wheat could vary quite widely. The importance of this price here is that it is not the source of the discrepancy between Hopkins's and Goldsmith's estimates. It is a difference among second-generation estimates.

The next row of table 11.1 shows agreement yet again on annual wheat consumption. But appearances are deceiving, and there are many problems at this juncture. It should strike any ancient historian as odd that Hopkins stated the annual consumption in a modern unit, and indeed it is. Two aspects of this estimate are of interest. First, it is not an estimate of wheat consumption. As Hopkins stated clearly, it is total annual consumption, consisting primarily of food and textiles, measured in units of wheat equivalence. The distinction between wheat consumption and total consumption measured in wheat units is critical for the derivation of any income estimate. Second, the estimate is not from ancient Rome! It is an estimate of consumption in modern subsistence agriculture, taken from a study of modern conditions by Colin Clark and associates in the 1960s. The actual source stated that "subsistence requirements, expressed in our units of grain equivalents, are somewhere between 250 and 300 kg./person/year, varying substantially, as we have seen, with climate and average body weight" (Clark and Haswell 1967, 60).

Goldsmith estimated wheat consumption from the rations given to Roman soldiers and slaves. He assumed that these were the rations for adult men, while women and children consumed less, so that his estimate of average consumption was less than the historical rations. He supported his estimate by comparing it to fifteenth-century Florence. Note that Goldsmith estimated wheat consumption, not total consumption, while Hopkins estimated total consumption in wheat equivalents. The apparent agreement in table 11.1 is not agreement at all, and only the economist's estimate is based on ancient evidence.

Both estimates implicitly assume that consumption or wages can be generalized from a few observations to the economy as a whole. This generalization only is possible if there is a labor market equalizing wages (or at least tending in that direction). As noted in chapter 6, we can speak of a labor market if wages serve to equilibrate the demand and the supply of labor, that is, to make the desired demand for labor equal the desired supply. Workers must be free to change their economic activity and their location, and they must be paid something like their labor productivity to indicate to them which kind of work to choose. That does not mean that everyone changes jobs with great frequency; it means instead that enough people are able and willing to do so to eliminate conditions where payments to labor are either excessively higher or lower than the wages of comparable work in other locations or activities. Even in the United States today, which contains the most flexible labor market in history, wages for comparable jobs are not completely equalized across the country.

Hopkins enlarged his estimated consumption by allowing one quarter of the wheat crop for seed. This has two logical problems, in addition to the problem of knowing the yield of wheat farms in antiquity. First, the estimated consumption was in wheat units, not of wheat. The seed requirement only applies to the portion of consumption that was wheat, which was not specified by Hopkins. It appears that Hopkins treated total consumption as if it consisted only of wheat. Second, seed is an intermediate good that normally is not counted in the national product, as explained in chapter 9. GDP measures the production of goods and services for final consumption, that is, for consumption by people or by business firms (in the modern world). Wheat that is produced in agriculture and then used as seed in agriculture does not appear as a part of GDP, even though it is part of farm output (Goldsmith 1984, 273n). Hopkins was right to think that his estimate of consumption might omit parts of the national product, but he was wrong to include seed in place of the missing parts.

Goldsmith took a different approach, one more in keeping with modern economics. He stepped up his estimate by appropriate multipliers to go from wheat to all food grains, from grain to total food consumption, from food consumption to total consumer spending, from consumer spending to national expenditures. These estimates came from a wide variety of ancient and modern sources; the estimated total consumer spending was checked against the allowances of the *alimenta*.

Modern national income accounting recognizes two ways of computing national and domestic products. The first way is to count all the expenditures for final goods, that is, omitting intermediate goods used in the production of final goods. This approach is spoken of as the expenditure side, and it is the approach summarized in table 11.1. The alternate approach is to sum all the earnings in the economy, which should add up to the value of final goods if there are no taxes or accounting errors. This approach is spoken of as the income side. Its result is often labeled national income, which differs from the national product in modern data due to taxes, errors and omissions, but which should be thought of as identical in ancient times (see chapter 9).

Goldsmith tested his estimate by deriving it from income data instead of expenditure data, as shown in table 11.2. The average wage was taken to be HS 3.5 from ancient sources. Days worked per year were set as 225 because that is about what modern workers work, giving annual compensation of HS 790. Assuming that the participation rate, the proportion of people in the labor force, was 0.4, then shifting from per-worker compensation to per capita, reduces it to 790 × 0.4 = 315. Goldsmith then assumed that nonwage income—rents, interest, indirect taxes, and depreciation—was 20 percent of

Table 11.2.
Estimate of Roman GDP from the income side

Category	*Goldsmith (1984)*
Average daily labor compensation	HS 3.5
Working days per year	225
Average annual labor compensation	HS 790
Participation rate	0.4
Labor income per-capita	HS 315
Step-up for non-labor income	1.2
Total income per head	HS 380

Note: Figures may be rounded to agree with sources.
Source: Goldsmith (1984).

labor compensation. Multiplying HS 315 by 1.2 gives a total income per head of HS 380, the same as estimated in table 11.1 from expenditures.

Goldsmith's estimate of per capita national expenditure of HS 380, is over twice Hopkins's estimate of HS 153. In a later essay, Hopkins (1995–96) stated that his earlier estimate was a minimum estimate of subsistence, and he multiplied it by 1.5 to get a more reasonable estimate of Roman GDP, shown in the last line of table 11.1. No source was provided for this multiplier; Hopkins identified it as a "rough guess." The resulting per capita product, HS 230, is still only two-thirds of Goldsmith's estimate. If we had to choose between them, then Goldsmith's larger estimate should be preferred. It is in accord with modern economic conventions, it was derived in two ways—from the expenditure side and the income side—and it was based, at least in part, on ancient sources. However, we can do better than simply choosing between these two existing estimates.

Let us back up a bit from the specific task at hand to inquire about assumptions underlying the whole effort. We are so familiar with the concept of the "gross product" that Hopkins set out to estimate that we do not often remind ourselves about the underlying assumptions. There are two key assumptions, both connected with prices. We assume that the goods in question, wheat in this instance, were sold on competitive markets. We do not have to assume that markets were as competitive in Rome as those monitored by the Chicago Board of Trade, but we do need to assume that there is enough competition to make prices of the same good in different sales roughly the same. Only in that case does in make sense to use the few price quotations we have from ancient times as an index of prices throughout the Roman Empire.

Some ancient historians rejected my estimated GDP because they argued that there was not enough market activity in the ancient world and not enough competition around the Mediterranean to make such an estimate meaningful. This is the position taken by Finley, who asserted that, "ancient society did not have an economic system which was an enormous conglomeration of interdependent markets" (Finley 1973, 22–23). Subsequent research has exposed the presence of substantial market activity and an international division of labor that suggests the integration of distant markets, putting the burden of proof on those who are skeptical, as presented in previous chapters.

Economists make GDP estimates for less-developed countries today and medieval economies in the past, but there comes a point where the estimate loses its meaning. If ancient markets were so rudimentary, then my estimate is inappropriate. It is worth pointing out that critics who take this line also must reject most other summary statements about ancient economies—which inevitably rest on implicit assumptions of consistent economic activity.

In addition to assuming that there were markets, we need to assume that people were free to alter their consumption in order to infer that they were better off when national income increased. The relative values of different goods then reflected the relative valuation that consumers placed on them. In other words, we need to assume that people could spend the money—or wheat—that they had in any way they chose. It is not necessary that all things were for sale to anyone; there clearly were restrictions on togas with purple stripes. But there had to be a broad range of goods that people could use their incomes to purchase and consume. Only in that case can we infer any welfare comparisons from an estimate of gross product, for only in that case will it be true that a higher product will—other things being equal—make people better off.

For example, the third line of table 11.1 lists the price of wheat as being HS 3 per modius. If the few observations of wheat prices that have survived were simply random magnitudes, or if they were administrative rates set at different times or places, there would be no sense in taking an average. Another observation, or the omission of one we have, could change the average, possibly by a large amount. Only if we think wheat was sold in markets does it make sense to see the surviving observations as indications of the usual or most frequent price. Then we can generalize from the few observations we have to the prevailing price of wheat. Even this step is problematical for the price of wheat in the early Roman Empire for the reasons given in chapter 2. Wheat prices were dependent on the distance from Rome. As Scheidel and Friesen (2009, 67) reiterated, "The only thing we can be sure of is that actual prices varied quite significantly by region, being lowest in grain-exporting Egypt and highest in the capital."

The fourth line of table 11.1 shows estimated annual consumption, which is harder to measure. There is no market in consumption, and we cannot generalize easily from a few observations. Hopkins relied on modern data, but that is unsatisfactory. He assumed that almost all Romans lived at subsistence levels, and he used modern data to estimate subsistence as if it was a biological constant that had not changed in two millennia. He assumed, in other words, that subsistence was a knife edge, below which you starved, and that this boundary is the same at all times and places. Modern research, however, has shown that subsistence is not a fixed level of income (Dasgupta 1993). In other words, people do not simply die when their consumption falls below some threshold. Their ability to work is impaired when they are malnourished, their growth is stunted, and their life expectancy is reduced. No economist today takes seriously Colin Clark's estimate of consumption; it has vanished beneath the waves of current scholarship.

There are a wide variety of countries today where most people are malnourished, and developing countries differ greatly in the extent of malnourishment. The range of poverty is great, and there does not seem to be a boundary below which a country cannot fall. Bangladesh, Haiti, and Kenya are among the poorest of countries. Their per capita income was about half that of China or Egypt in the 1990s. India's per capita income fell between those of these two groups of countries. All of these countries are poor, but some are poorer than others.

Recognizing the variety of incomes in developing countries today, the question of Roman product is not whether ordinary Romans lived at subsistence level, but rather which kind of subsistence they experienced. They could have been prosperous farmers who had adequate calorie and vitamin intake and been large and healthy people as a result. Or they could have been stunted and ill from malnutrition, even while not so poor that they died or were unable to have children. The question is not whether Rome was richer or poorer than developing countries today; it is which developing country did Rome resemble most closely, at least in respect to general nutrition.

Ancient estimates of grain consumption do not help much with this question. A survey of ancient evidence concluded that "the evidence on grain consumption is very scanty—and the situation is far worse for other food items" (Foxhall and Forbes 1982, 74). The evidence is not even about consumption; it is about grain rations for Roman soldiers and Cato's slaves. We do not know if these rations were for the workers alone or for them and their households and assistants, nor do we know what other food the soldiers and slaves consumed. The rations for soldiers varied by a factor of three, suggesting that these questions are not simply theoretical. Foxhall and Forbes (1982, 64) cast further doubt on the use of rations to estimate consumption when they speculate on the process by which rations were set: "The similarity of the rations of the army and Cato's working slaves also suggests that the Romans may have used a basic 'rule of thumb' for the estimation of projected consumption similar to that suggested for the Greeks. Possibly it too originated on farms and in households and at some stage made its way to state-level usage."

One problem with using modern averages or ancient rules of thumb is that they do not allow for the possibility of economic growth. They are static estimates that apply to no specific time. At best they could provide an estimate of the average Roman income over some period; they could never reveal to us a change in that income. It would be unfortunate if the attempt to measure Roman gross product implied that there had been no change in that product over hundreds of years, except by extension of the population (see chapter 10).

We have learned more about both the ancient economy and less-developed countries today than we knew twenty years ago. We have more price data for the ancient world and more income data for modern developing countries. I employ these new data in turn, starting with the modern data, to derive two new estimates of Roman GDP. They should, like Goldsmith's estimates, agree. But, unlike Goldsmith's data, they do not correspond to the two ways national incomes and products are calculated today. Instead, they approach the estimation from two entirely different starting points.

Many historians have noted that the Romans had an urban economy. The city of Rome is thought to have had a million residents in the early Roman Empire, and there were many smaller cities throughout the Roman Empire. A conventional view is that about 10 percent of the population lived in urban areas. Perhaps that figure is just the smallest single digit that can recognize the abundant evidence of Roman cities and towns in an agrarian society. Lo Cascio (1994) and Scheidel (2001) have warned us that all demographic estimates for Rome are open to question, and this one must be highly speculative because it is highly aggregated. Even if this number is not accurate, it captures something about the early Roman Empire.

People living in cities typically do not grow their own food. An economy with a large urban population therefore has to have a more efficient agriculture; each farmer has to feed his family and part of an urban family as well. This, of course, does not deal with the question of how food gets from countryside to cities. The implication for productivity is independent of whether this was done by taxes, tribute, or trade. This reasoning has been used for early modern economies to demonstrate the extent of economic growth. David (1967) used it to estimate the rise in United States per capita GDP in the early nineteenth century. More recently, Craig and Fisher (2000, 113–18) used data on urbanization from Bairoch, Batou, and Chèvre (1988, 259) to estimate changes in per capita GDP for several European countries between 1500 and 1750.

This index puts the early Roman Empire, at roughly 10 percent urban, at the level of Britain, France, and Germany in 1600. These countries at the time appear to have been far poorer by this index than Belgium, the Netherlands, Italy, Spain, and Portugal. The urbanization rate of the early Roman Empire cannot be known with precision. The estimate of 10 percent has to be regarded, at best, as true to only one significant digit. If we assume that the early Roman Empire was between 5 and 15 percent urban, the conclusion just stated remains valid. Roman per capita GDP was near the GDP of the three main European countries of northwestern Europe in 1600, but below that of the main trading nations both north and south of these countries.

If we consider the area that is now Italy alone, Hopkins (1978, 68) speculated that it was 30 percent urbanized in 28 CE. By this measure, the center of the early Roman Empire is estimated to have been as wealthy as Belgium and the Netherlands, the richest of the trading countries of early modern Europe above Italy and the Iberian states in 1600. Like Roman Italy, they were at the center of an extensive trading network, enjoying both gains from trade and income from the management of the trade.

I drew a sharp contrast with urbanization in Italy and the rest of the empire. While Rome was the largest city in the empire, and the largest city in history for another two millennia, it was not the only city in the early Roman Empire. There were many smaller cities scattered over Europe, Africa, and Asia. There were hundreds of cities in Roman Africa alone, suggesting that I may have overdrawn the contrast between Italy and the rest of the empire. It is important to keep the distinction between the metropolis and the empire in mind.

My estimate of Roman urbanization received support from Wilson's "very tentative estimate for the urban population in the Roman world c. A.D. 150, which suggests a population of c. 7,370 living in cities of 5,000 people or more" (Wilson 2009a, 74). Taken relative to Scheidel and Friesen's (2009) estimate of 70 million people living in the Roman Empire at this time, this implies an urbanization rate of about 10 percent. (Wilson said 12.5 percent with spurious accuracy.) Their insistence on the large size of the Roman Empire reminds us of the unusually large extent of the empire and how difficult it would have been to raise 70 million incomes.

It also helps us to distinguish between Roman Italy and the empire as a whole. Only about 10 percent of the population of the empire lived in Italy, and they could easily have had higher incomes than the provinces. As just noted, we need to compare like with like. If we are to distinguish small countries in early modern Europe, then we need to compare them with Roman Italy—or even Rome itself. If we seek to compare the GDP of the Roman Empire with early modern economies, we must compare them with Europe as a whole.

Per capita GDP grew in most European countries in the years after 1600, albeit slowly during the difficult seventeenth century. Almost all Western European countries had urbanization rates above 10 percent by 1750, the principal exceptions being Germany and Scandinavia. They had not participated in the expansion of the Atlantic trade, and they remained relatively poor and rural as a result. The early Roman Empire's per capita GDP was about equal to Germany's in both 1600 and 1750. Roman Italy was between the level of Belgium and the Netherlands in 1750, which had grown apart since 1600 by the urbanization measure.

To provide a more contemporary comparison, I collected GDP data for modern developing countries in 1995. I selected countries that had population of more than 1 million and per capita GDP of less than ten thousand dollars. This provided a sample of a hundred countries. I regressed the logarithm of per capita GDP on the proportion of the population in urban areas and several regional dummies. The result showed that per capita GDP was related to urbanization. This result does not imply that urbanization caused per capita GDP to rise—the direction of causality may well run in the opposite direction—but only that higher per capita GDP is associated with increased urbanization. The regression is as follows, with t-statistics below the coefficients.

$$\text{Log per capita GDP} = \underset{(22.88)}{7.376} + \underset{(7.30)}{.0267}\,\text{URB} + \text{Regional Dummies} \quad R^2 = .73$$

Adding .267 (.0267 × 10) to the constant gives a predicted log per capita GDP of 7.643, which implies a predicted per capita GDP of $2,085. This is comparable to India in the mid-1990s, a large country containing—like the Roman Empire—many disparate areas. This is an interesting result, which builds on Hopkins's comparison of ancient Rome and current developing countries. But the result that Roman income was like that of a current developing country has a variety of drawbacks. First, while the use of urbanization provides a way to use ancient data to refine the estimate, the effect of the ancient data is not large. The coefficient of urbanization is very small, and the estimate is close to the mean of the sample of countries used. Second, it follows that any estimate will be sensitive to the sample of countries chosen. The correlation of urbanization and per capita GDP breaks down at higher incomes, but the mean of countries clearly will be higher if richer countries are included in the sample. Third, this calculation does not provide an estimated GDP in ancient currency. It therefore can be used to compare Roman GDP with more recent conditions but not to compare GDP with other ancient magnitudes.

The previous estimate followed Hopkins's lead; the second one follows Goldsmith's, as summarized in table 11.2. I expand his calculation of the wage and explore the implications of the data for the level of the real wage, that is, the wage divided by the price of wheat. I then expand these results into an estimate of GDP and of the income distribution in the early Roman Empire.

When I started this exploration, I thought I would be able to find many more observations of Roman wages to analyze, but there are not a lot of observations, and most of them are for wages in Egypt. Given the different monetary systems used in Roman Egypt and Roman Italy, there is no thought that these wages were the same. But all is not lost. The fragmentary data that we have fit into a simple pattern that illuminates the conditions of life in ancient Rome. I rely in what follows on Egyptian wage data collected by Drexhage (1991) and Egyptian wheat price data collected by Rathbone (1997). These data have been summarized recently in Scheidel (2002). I compare these relatively abundant observations with the scatter of evidence we have about the city of Rome, mostly as collected by Duncan-Jones (1982).

The daily wage in rural Egypt from the mid-first to the mid-second centuries averaged around seven to eight obols a day. At the prevailing rate of seven obols to a drachma, the wage was approximately one drachma a day. Monthly wages were about 20 to 25 times that high, suggesting that the monthly workers were not more skilled than the daily ones. Cuvigny (1996) reported that skilled miners earned 47 drachmae a month at Mons Claudianus, approximately double the monthly rates compiled by Drexhage. The few annual wages we know of ranged from some that were 250 or 300 times the daily wage, indicating similar skills, to some that were in the range of 1,000 drachmas. The wages recorded by Cuvigny and the high annual salaries probably went to skilled workers and supervisory personnel, assuming that these wages were determined in a functioning labor market

The daily wage in Rome itself may have been about HS 3 to HS 4; Goldsmith used the average of these numbers, HS 3.5. If we convert Alexandrian drachmas to sesterces at the official rate of one-to-one, then wages in the city of Rome were three or more times as large as wages in rural Egypt. This cannot be surprising. If there was a labor market in the early Roman Empire, even an imperfect one, then this pattern is what we would expect. In less-developed countries today, wages are higher at the center of the empire than in the provinces, and they are higher in the city than in the countryside. The existence of this pattern in antiquity gives some confidence to what after all is very sparse evidence both about wages and the monetary regimes in Egypt and the rest of the Roman Empire.

Continuing this analogy, the same forces that produced higher wages in the metropolis would have generated higher prices there as well. If there had been a well-functioning labor market, the real wage, that is, the wage measured in the amount of wheat or other products it would buy, should have been similar in Egypt and Rome.

The Egyptian wheat price was about 8 drachmas for an artaba, which was the monthly ration for soldiers and other people. It then took a little over a week for a laborer in Roman Egypt to earn one artaba of wheat. We can make a similar calculation for Rome, although wages were paid in sesterces, and wheat was measured in modii. We need to know how many modii were in one artaba of wheat to make a comparison with Egypt. The artaba was a larger measure than the modius, and one artaba was the equivalent of 4.5 modii (Duncan-Jones 1982, 372).

Yet the price of wheat in Rome is difficult to use because of the prevalence of state subsidies, as discussed in chapter 2. Both Hopkins and Goldsmith stated that the price was HS 3 per modius. The cost of one artaba of wheat in Rome then would have been HS 13.5. At a daily wage of HS 3.5, that was about four and one-half days' work. This wheat price, however, was the price of wheat subsidized by the imperial authority. In other words, the cost of wheat in Rome was more than HS 3 by the amount of the subsidy. We can estimate the size of the subsidy by looking for wheat prices that do not contain a subsidy. Duncan-Jones (1982, 345) argued that "normal prices for wheat at Rome were considerably more than [HS 4]." He said that the unsubsidized price might have been as high as HS 8. If so, then it may have taken up to ten days for a Roman worker to earn one artaba of unsubsidized wheat. This is not very precise, since it is limited by the scarcity of price evidence in Rome, but it is enough to suggest that the Roman labor market might have worked exceptionally well.

Real wages, that is, wages divided by the price of wheat, in Roman Egypt and in Rome itself do not appear very different. We cannot reject the hypothesis that real wages were the same in rural Egypt and the city of Rome. This finding suggests that there was a functioning labor market that allowed workers to move to where wages were higher. It was not the case that Italian workers were impoverished relative to Egyptian, nor that Egyptian workers were exploited relative to Roman. The sparse data may even reveal that real wages were slightly higher in the city of Rome than in rural Egypt. The large difference in nominal wages translates into a smaller difference in real wages, and the best guess must be that the difference in real wages was not large. The experience of current less-developed countries suggests that there would have been a difference, as noted already, and the evidence goes in the right direction.

Scheidel (2010) disputed my estimated wage in Rome, but he confirmed my inference of similar real wages in Rome and Egypt through a comparison of his estimates for Roman Egypt with Allen's (2009a) from the Diocletian

Edict. As noted earlier, this correspondence is more troubling than comforting. Scheidel (2010) reported identical real wages in Egypt before and after the Antonine Plague in contradiction with the rise in real wages after the plague in Scheidel (2002). He used date from the earlier paper in his later paper but remained mute about the contradiction. He emphasized instead the applicability of the Diocletian Edict and implicitly the uniformity of real wages across the Mediterranean Sea.

To return to the estimation of Roman GDP, Goldsmith constructed an estimate of per capita GDP from the income side as shown in table 11.2. He used the wage in Rome. But only about one-tenth of the population of the early Roman Empire lived in Italy, and fewer than that lived in Rome. The wage data surveyed here show that wages were higher in Rome than elsewhere in the empire, implying that Goldsmith's estimate of per capita GDP was too large. It needs to be based on the average wage in the Roman Empire rather than the average wage in Rome. I retained Goldsmith's estimates of working days per year, the participation rate, and the step-up for non-labor income, changing only the estimate of the average wage.

How much lower than wages in Rome were average wages in the empire? We do not know. Wages in Egypt were only one-third as large, but we do not know if wages elsewhere were closer to those in Egypt or those in Rome. Appealing to the principle of insufficient reason, I suggested that average wages were one-half the Roman level, somewhat above the low level of Egyptian rural wages. This suggests that the GDP of the early Roman Empire was about half of that estimated by Goldsmith, or about HS 10,000 million. Surprisingly, this is almost the same as Hopkins's 1995–96 guess of HS 12,500 million. It is lower than other second-generation estimates in sesterces.

The evidence surveyed here can do more than provide another guess about the size of the Roman GDP; it can inform us about the distribution of income within the empire. We get this information by contrasting results from both ways of approaching the problem of estimation. From the modern comparison using urbanization rates, it seemed that per capita income was higher in Roman Italy than in the Roman Empire as a whole. From the comparison using ancient data on wage rates, I inferred and Scheidel (2010) agreed that real wages in Rome were equal or close to equal in Rome and at least one large province.

How can these apparently contradictory inferences be reconciled? It appears that the high income of Rome was not shared equally. Daily workers received only slightly more in real terms than daily workers outside Roman Italy. The higher incomes in Roman Italy must have been the earnings of people who were not daily workers, that is, more highly skilled workers and the owners of property. If the real incomes of richer people were higher in and around the city of Rome than in Egypt while the real incomes of the ordinary workers was the same, then the distribution of income was more unequal in Rome and Roman Italy than elsewhere in the empire. This is not to deny the presence of some very rich Romans in the provinces but rather to argue that there were only a few of them outside Roman Italy.

This suggests that the high per capita income in Roman Italy found by urbanization was the result of great fortunes of people living there. All sorts of people lived elsewhere, both rich and poor, but there was no concentration of wealth outside Italy comparable to the massing of wealth in Italy. The expansion of Roman rule had yielded great profits. It is not a new idea that these gains were captured by a few people. It may be new to learn that they were concentrated almost entirely in Rome and its surrounding countryside. They may have initiated a durable pattern of constructing luxurious villas. Two millennia later, as great wealth was being accumulated in nineteenth-century England, "millionaires made their money in the town and spent it in the country" (Crook 1999, 27). Wealthy Romans also invested in land but appear from the urbanization evidence to have retained a large presence in Rome itself.

Goldsmith (1984) argued that the distribution of income in the early Roman Empire was similar to that found in England before industrialization or in Brazil at the time he wrote. If this was true for the empire as a whole, then it follows that the distribution of income was more equal in the provinces and very highly skewed in Rome itself. Tainter (1988) argued that the concentration of wealth in the late Roman Republic led to class conflict, demagoguery, and the republic's collapse. As suggested by Tainter, the extreme income disparities in Rome itself, as opposed to conditions elsewhere under Roman rule, strengthen this argument. Phillips (2002) used the same argument to worry about the future of the American republic today.

Later estimators have all said that my estimated GDP in sesterces was too low. I used prices and wages from the provinces to be representative of the empire as a whole. If we multiply the per capita income estimate by 70 million to get aggregate income for the empire, we need to use the average price in the empire as a whole. As everyone agrees, prices were lower outside Rome and quite low in populous Egypt. But since all the estimates discussed so far were based on modern estimates of subsistence consumption, they cannot inform us of the magnitude of ancient economic growth. It is more promising to consider the alternative estimates from comparisons with early modern economies.

Maddison (2007, 43–53) started from Goldsmith after criticizing Hopkins's economics in including seeds in GDP, as described earlier. He broke down his estimates by Roman province, which added more detail, and he introduced comparison by comparing Roman GDP with early modern estimates in 1990 Geary-Khamis dollars, a hypothetical currency with the purchasing power of

1990 U.S. dollars. Purchasing power is estimated for modern countries by assuming "purchasing power parity," that is, assuming that the cost of an identical basket of internationally traded goods will cost the same around the world. In other words, the use of this unit of account implicitly assumes that the results of chapter 2 showing that the price of wheat was determined by trade across the Mediterranean were true for most items of Roman consumption. Maddison, an economist, had no trouble with this assumption; ancient historians may be suspicious. Bang (2008, 87–88) represents the opposite position. He took his cue from Hopkins, reproducing his procedures with its mistakes and accompanied by many caveats. He did not mention Geary-Khamis dollars, and he would not agree they are useful in studies of the ancient world.

Maddison's conclusions are revealing. He followed Goldsmith's lead and compared Roman incomes with those of England in 1688 when a contemporary estimate was made by valuing consumption in each period by the value of gold and then by the value of wheat. Averaging them, he concluded "that average per capita income in the Roman empire was about 40 per cent of the $1,411 in England and Wales in 1688" (Maddison 2007, 52). Since the modern estimates were in Geary-Khamis dollars, Maddison used them for Roman values.

Scheidel and Friesen drew from both ancient historians and economists. They start their estimation with their conclusion: "We merely assume that in terms of average per capital performance the Roman economy did not dramatically differ from most other pre-industrial systems and fell short of the achievements of the most advanced economies of the early modern world, those of the Netherlands and England" (Scheidel and Friesen 2009, 64). This "mere" assumption implies that Roman GDP was near subsistence and near Maddison's estimate. This becomes clear in their stated conclusions where their estimates of Roman GDP are dependent on estimated subsistence. If it was 350 kg of wheat equivalent, the Roman per capita GDP was about $610 Geary-Khamis dollars; if substance was $390—close to Maddison's $400—then Roman per capita GDP was around $700 Geary-Khamis dollars (Scheidel and Friesen 2009, 74).

They argue that their assumption of a gap between Rome and early modern Europe "does not shape our estimates . . . in an unduly arbitrary or circular way." Citing de Vries and van der Woude (1997), they elaborate as follows:

> Labeled the 'first modern economy,' the 'Golden Age' Netherlands enjoyed unusually large energy inputs, provided by fossil fuels in the form of local peat deposits, that were unavailable to other 'organic,' pre-industrial economies; attained level of formal schooling and literacy were exceptional by pre-modern standards; created a flourishing bond market; and was the first country on

record in which the share of the population engaged in farming fell below one-half. (Scheidel and Friesen 2009, 64)

There are several problems with this comparison. The main problem of this claim is that Scheidel and Friesen compare apples and oranges. As they note in the following paragraph, the Netherlands had a population of only 2 million people. How can you compare this to the Roman Empire of 70 million people, their preferred population on the eve of the plague? Rome itself had a population of 1 million, and Roman Italy had a population of around 7 million people. These entities are far closer to the early modern Netherlands than the massive Roman Empire. The Netherlands were the most advanced country in the seventeenth century, and it should be compared with the most advanced part of the Roman Empire.

Rome is smaller than the Netherlands, and Italy is larger, but they are at least the same order of magnitude. If we take Roman Italy as being comparable to the Netherlands in 1600, then the Roman Empire as a whole is comparable to Europe as a whole. According to Maddison, the population of Europe was between 60 and 70 million people, close to the 70 million maximum of the early Roman Empire. (Scheidel and Friesen were aware of this mismatch, but they compared the Roman Empire with a Europe that had the same average income after 1800 as the most advanced most advanced country in 1600. They also noted that there may have been "pockets of development such as Roman Italy or the Aegean" [Scheidel and Friesen 2009, 64]. Their reasoning is unclear.)

The putative advantages of the early modern Netherlands over Roman Italy are less clear than their advantages over the Roman Empire as a whole. Roman Italy was heavily urbanized, and it is probable that the proportion of the labor force in agriculture there was less than 50 percent. Rome did not have a bond market, but it had a more sophisticated financial system than the Netherlands (see chapter 8). It did not have peat, but it had waterwheels and needed less heat for its houses. It did not have canals, but it had roads and the Mediterranean. It had education, a functioning legal system, and an ethos of responsibility. Both places developed art and literature that still affects us today. If the previous chapters of this book are more or less accurate, then Roman Italy was comparable to the Netherlands in 1600.

The comparability is enhanced by noting that neither the Netherlands nor Rome had an industrial revolution. The Industrial Revolution took place in the north of England where the relative price of labor and power was very different than in the eighteenth-century Netherlands and ancient Rome. Dutch peat was not anywhere as cheap as English coal at the mine head, and this difference in relative prices made the early machines of the Industrial Revolution unprofitable on the European continent (Allen 2009c).

Lo Cascio and Malanima (2009a, 2009b) compared Goldsmith's and Maddison's estimates to show how similar they were. They then contested Maddison's efforts to compare Roman and early modern incomes by deflating the modern estimates by the prices of gold and wheat, arguing that gold was more expensive (and goods cheaper) than later and that Maddison had not accurately reported grain prices in early modern Europe. They followed Maddison in separating consideration of Roman Italy and the empire as a whole.

The result of these recalculations was to produce an estimate of Roman per capita GDP in 150 CE of 1,000 1990 international dollars (Lo Cascio and Malanima 2009b). This is almost twice as large as Maddison's estimates, which were the basis of Scheidel and Friesen's calculations. They follow Maddison in separating incomes in and outside Roman Italy, concluding that per capita GDP in Roman Italy was 1,400 international dollars, almost exactly equal to Maddison's estimate of per capita income in the Netherlands in 1600 of 1,381. As I argued in chapter 10, the benefits of international and interregional trade were the same in Roman Italy and the advanced countries of early modern Europe; Lo Cascio and Malanima made this same argument from a different direction.

They generalized their result to argue "that pre-modern agrarian economies underwent phases of growth and decline, but not of real long-term progress. . . . Growth was not unknown before Modern Growth, but it came about in long cycles around an overall stability of per capita income. In our view, the Roman economy was no more backward than the early modern West European economy" (Lo Cascio and Malanima 2009a, 392). In terms of the questions of this chapter, the Malthusian swing of income in the Roman Empire was as large as in early modern Europe.

Calculating Roman GDP should be a way of discovering if these two economies were equally productive. Unhappily, it appears that the answer has been prejudged before any Roman calculations are done. I therefore use a simpler and more transparent approach than that used by Lo Cascio and Malanima to confirm their results. I work backward and assume that Roman Italy was comparable to the Netherlands in 1600 and infer what the GDP of the Roman Empire might have been.

The calculations are simple, building on two observations. From Maddison (2007, 50) I take the observation that per capita income outside Roman Italy was about two-thirds of its value there. He was the only one to estimate per capita GDP by province. I couple that with the rough estimate that the population of Roman Italy (about 7 million) was about 10 percent of the Roman

Empire's population (about 70 million). Then if per capita GDP in Roman Italy was close to that of the Netherlands in 1600, per capita GDP in the Roman Empire was about 1,000 Geary-Khamis dollars. This is exactly equal to Lo Cascio and Malanima's estimate, higher than Scheidel and Friesen's high estimate of 700 for Rome and of Maddison's (2007, 382) estimate of 900 for Europe as a whole in 1600. It is almost exactly equal to Maddison's estimate of European income in 1700. Is it circular? Only if you have not been convinced by the preceding chapters about the microeconomics of Rome—or have not read them—should this short derivation seem worse than its competitors.

This easily derived estimate is the best estimate for the GDP of the Roman Empire at its peak. It can be multiplied by the size of the population to get an aggregate. As stated by Lo Cascio and Malanima (2009b): "Our discussion of the problem strongly suggests that a real difference in the level of average income between the early Roman Empire and the following European pre-modern economics did not exist at all." I look forward to more debates about how prosperous the Romans were and what their undoubted accomplishments imply for our view of their economic progress in the early Roman Empire.