Economic growth on spaceship Earth

Rob Hart

rob.hart@slu.se

Contents

Preface

Economic growth on spaceship Earth

The metaphor of Earth as a spaceship—made famous by Kenneth Boulding (1966)—helps to bring home the obvious truth that we live together on a planet with finite resources. Furthermore, inputs from beyond the Earth—such as the flow of energy from the sun—are limited, and there are enormous obstacles to resource extraction from other planets, let alone their colonization. The necessity of living together on the finite surface of the Earth suggests the need for cooperation and fairness, as emphasized by George Orwell in 'The road to Wigan Pier'.[1]

> The world is a raft sailing through space with, potentially, plenty of provisions for everybody; the idea that we must all co-operate and see to it that everyone does his fair share of the work and gets his fair share of the provisions seems so blatantly obvious that one would say that no one could possibly fail to accept it unless he had some corrupt motive for clinging to the present system.

On the other hand, Kenneth Boulding focuses on the necessity of the careful use of the finite natural resources available to us, and the avoidance of fouling our own nest with pollution. In this book we follow Boulding by focusing on resource use and pollution rather than cooperation and fairness. Boulding describes the transition in the human imagination from the idea of the frontier economy in which scarcity is only ever local, to the idea of global limits. We aim to understand both the 'frontier' and 'spaceship' phases of development as part of a single long-run process; even when we have the mindset of the frontier, we are already on the spaceship.

This book thus concerns the development of the global—spaceship—economy in the very long run. How can we make sense of the changes in the human economy that we have seen of the last couple of centuries, and even the last 20 000 years? In the light of our explanation of the past, what future scenarios are possible, or indeed likely, regarding long-run global production of goods and services, given the finite nature of the Earth, its natural resources, and the inflow of energy from the sun? Furthermore, what policies are called for in order to achieve desirable long-run outcomes with regard to (sustainable) long-run production and the quality of the living environment? Over the last 100 years and more the world has witnessed economic growth which is not only uniquely rapid, but also astonishingly steady. Can this increase be maintained, and if it is maintained, will that be at the expense of the quality of the Earth's environment or other species?

Questions regarding management of the economy

There are many relevant questions that could be raised regarding economic growth on spaceship Earth. Consider the following three.

(1) (a) What is desirable?
    (b) What is feasible?
    (c) What is optimal?

The first question might be asked by a philosopher. What characterizes a good life? What characterizes a good society? What characterizes a good society taking into account not only the people (and animals, plants, etc.) of today, but also those of the future? There are many different ideas about a good society. A couple of famous ones are those of utilitarianism and Rawls. According to utilitarianism we should strive to maximize utility, which is typically considered as the sum of individual utilities. Utility may be a function of various factors; in economics we typically focus in *consumption* as the key to utility. Something different utilitarian measures generally have in common is that they are *consequentialist*, that is they measure the rightness of actions and more general moral rules according to their outcomes, with the best actions leading to the greatest sum of utility. A fascinating and deep analysis of issues regarding morality and choices affecting the future is Derek Parfit's *Reasons and persons*.[2] He subjects standard utilitarian reasoning to a

---

[1] Orwell (1958), p.203. Available at `https://archive.org/details/roadtowiganpie00orwe`.
[2] Parfit (1984).

searching examination from which it emerges severely battered. If for instance we should maximize happiness, is it better to have 100 billion people scratching out a living but enjoying a tiny bit of happiness each, or 5 billion genuinely happy people? And what responsibilities do we have to future generations since our specific choices lead to their creation? (If we had done things differently, different people would have been born.)

According to Rawls (1971), good moral rules (or indeed rules for how society is run) are rules which would be chosen by the citizens of that society if they had to make their choices from behind a 'veil of ignorance', a veil which prevented each player from knowing anything about their own position in that society (i.e. whether they would be rich or poor, or have high or low status). Rawls claims that—when faced with a choice from behind the veil—we would choose moral and social principles based on liberty, equality of opportunity, and help for the least advantaged. In an intergenerational context—where the veil prevents us from knowing into which generation we would be born—these principles translate into equality of opportunity across generations (so future generations should not be denied opportunities that we have, and vice versa), and also help for the least advantaged generations, so we should work hard to help the poorest generations rather than the richest. If we assume that global growth will carry on indefinitely this implies that we should look after ourselves, at least in a material sense, since future generations will be richer!

The second question from our list—what is feasible—might be asked by an engineer. Once a feasible solution to a problem has been found, the engineer might be satisfied. Or, if there is a limit on the budget, the problem might be to find a feasible solution within budget. This brings us to the third question, what is optimal. This is of course the classic question asked in economics, where we seek not just general rules or to rank alternative outcomes, but also specific ways of organizing society (including production and trade of goods and services) which lead towards the 'optimal' allocation of resources.

How do we know whether an allocation is optimal? Clearly we need a criterion, and the criterion typically used in economics is utilitarian: we want to maximize the sum of utility or (more generally) expected utility. But how to measure utility? Economists typically assume that it is an increasing function of consumption. Furthermore, we weight future utility lower than current utility, typically by discounting exponentially. Then we have

$$\max \sum_h U^h,$$

where
$$U^h = \sum_t u^h(c_t)\beta^t.$$

Here we have a set of discrete time periods indexed by t, U^h is the net present value of the utility of household h, $u^h(c_t)$ is the instantaneous flow of utility in period t, and β is the discount factor (which is less than 1). The reduction of utility to consumption is of course highly controversial outside economics, but most economists take it for granted. Furthermore, we also tend to take for granted the assumption that all households at a given time are identical (or can be treated as such), which is typically made in order to increase the tractability of macroeconomic models.[3]

A utility function can help us to rank alternative options, but how do we know what options are available? In order to know this, we must understand how the economy works, and how it can be controlled or managed. To build up such an understanding, we use highly simplified models which describe the agents in the economy (such as the government or regulator, firms, and households), their endowments (what they own), and the technology (what possibilities there are to produce outputs using available inputs). Given such a model we can tackle questions such as the following.

(2) (a) What would happen given laissez-faire?
    (b) What would happen given business-as-usual (b.a.u.)?
    (c) What would be the effect of regulations such as high taxes on fossil fuels, or strong support for research into renewable energy?
    (d) What would a *social planner* do?
    (e) What should a *regulator* do?

The first three questions relate to what happens in the economy under different circumstances. Question 2(a): what will happen in the absence of regulation by government? That is, what will happen if the government allows all agents to do what they like without imposing rules or taxes,

---

[3]Note that economists can (and sometimes do) use a Rawlsian criterion in which the aim is to maximize the utility of the least well off, the *maximin* criterion:

$$\max\{\min\{u_1, u_2, \ldots, u_n\}\}.$$

This states that we should maximize the smallest member of the set of utilities, i.e. we should concentrate our efforts on ensuring that the least well-off individual increases their utility.

except the rule of law (implying for instance the protection of private property or the right of individuals to own things). Question 2(b) is about the future in case the government continues to apply the policies which it currently applies, no more and no less, and the 2(c) is about the effect of specific policies.

Questions 2(d) and (e) are more subtle, since to answer them we need to be able to predict the future and then choose between alternative options. Question 2(d) presupposes the existence of an agent who has the power to fully control the allocation of resources in the economy without the need for regulatory instruments; all other agents simply do as the social planner instructs them. The planner is (fortunately) benevolent, i.e. she wants the best for everyone, she wants to maximize the sum of utility U. What would such a social planner do, how would she allocate resources? Would she, for instance, instantly cut fossil-fuel extraction? Or would she invest massively in research into alternative energy sources? If we can work out what our imaginary benevolent dictator would do, this gives us a yardstick against which to measure the results of our efforts to manage the economy through regulation (emissions taxes, technology standards, etc.), i.e. a yardstick which helps us to answer the question 2(e). How close can we get to the allocation in the planned economy?

Macroeconomic methodology

Having discussed the types of questions normally tackled in economics, we now briefly discuss neoclassical economic methodology in general, and macroeconomic methodology in particular. In neoclassical economics we build precisely defined model economies in which it is possible to calculate outcomes exactly, both given laissez faire (i.e. no regulatory intervention), and given interventions such as the application of taxes or command-and-control regulations. We then use these models to draw conclusions about how real economies work and are likely to develop or react to regulatory interventions. In macroeconomics our focus is on the economy as a whole rather than a specific set of firms or markets.

The precisely defined models of neoclassical economics consist of sets of equations. The models are generally very drastic simplifications of real economies: it is common for instance to assume that only one type of good ever gets produced in a model economy, and furthermore that there is only one type of labour. Some of the single good is consumed, while a proportion is kept back by firms in the form of capital to help in further production. It is then natural to ask whether we can ever use such models to learn anything of importance about real economies?

If a highly simplified model is to teach us about the real economy, it seems reasonable to suppose that it should be through a similar mechanism to that by which a parable teaches us about life. A good model helps us to organize our analysis of the economy, and allows us an insight into how the real, complex economy works, and how it is likely to react to changed circumstances (such as shortages of raw materials, or the introduction of a tax on fossil fuels). The point can be made more explicitly through a caricature of how *not* to do macroeconomics. Consider the aggregate data shown in Figure 4.4, reproduced here in highly stylized form as Figure 0.1. The data are consistent with a model in which long-run expenditure on metals and primary energy is a constant fraction of global product. Furthermore, the following aggregate production function is also consistent with this result:

$$Y = (A_L L)^{1-\beta}(A_R R)^{\beta}, \tag{PF}$$

where A_L is labour productivity, L is labour, A_R is resource productivity, R is the quantity of resource input (price w_R), and β is a parameter less than 1. It is straightforward to show that given perfect markets (so price = marginal revenue product) $w_R R/Y = \beta$, i.e. expenditure on the resource is a constant fraction of total product. Having noticed this property of the function, the *wrong* thing to do is to draw the conclusion that the production function (PF) is an appropriate description of production in the economy, and to use this to predict the effect of policy interventions. A prediction which would follow from this is that there is no point in boosting energy-efficiency A_R to try to reduce R, because rises in A_R cause effective energy inputs $A_R R$ to become cheaper, causing producers to use more energy, negating the effect.

Why should we not draw the conclusion that the production function (PF) is an appropriate description of production in the economy, and use this to predict the effect of policy interventions? There are many different ways to answer this question. Perhaps the simplest is that we have no evidence that (PF) is the *correct* description of the economy, or even close to being correct. All we know is that it can match one subset of aggregate data. There are likely to be many other functions or models which also match the data, but which may give completely different predictions concerning the effect of boosting A_R.

FIGURE 0.1. A stylized version of Figure 4.4, showing long-run growth in consumption and prices, compared to growth in global product, for (a) Metals, and (b) Primary energy from combustion.

In order to have any confidence in our model we need much more evidence regarding its suitability as a description of the economy. In particular as economists we want our model to include a *microeconomic mechanism* or *microfoundations*: this is a mechanism through which the behaviour of individual agents—behaviour which can be explained as a result of the incentives faced by these agents—leads to the observed aggregate quantities and trends. Furthermore, we need to see evidence that this microeconomic mechanism is in fact relevant to the case in hand. For a famous statement of the case for microfoundations—or perhaps the case *against* policy models without microfoundations—see Lucas (1976). According to Lucas, we cannot expect to predict the effects of a policy experiment based only on patterns in aggregate historical observations, because the policy experiment will change the rules of the game and therefore the old patterns may no longer apply. We must instead understand the rules of the game and how it is played—in economic terms we must know agents' preferences, the technologies available to them, what resources they are endowed with, etc.—in order to understand and predict the effect of a policy intervention. This book is largely concerned with this endeavour.

This book

This book is based on economic analysis. However, this does not mean that we accept the economist's habit of equating utility and income. Instead, we focus primarily on questions 2(b) and (c) above—what will happen under various scenarios—thus conveniently obviating the need to explore the more difficult questions 1(a), (b), and (c). It turns out that predicting the future is quite hard enough as it is.

We will also consider 2(e), policy, but based on goals which are exogenously determined (such as reducing CO_2 emissions) rather than calculated within the model. For instance, when considering climate policy we simply assume that society has decided (in its wisdom) on a goal of drastically reducing CO_2 emissions, and investigate how best that goal can be achieved. Thus, again, we avoid difficult questions about optimal overall choices, and focus on more limited questions.

Our avoidance of questions such as 1(a), (b), and (c) implies in practice that we avoid some huge and important debates within the fields of economics and sustainability. One such debate is that about the appropriate *discount rate* to use when assessing public policy, a debate made famous by the report of Sir Nicholas Stern (2006) for the UK Treasury on the economics of climate change. More generally, in the main part of the book we ignore the issue of intergenerational equity and justice completely. However, in the concluding chapters we broaden the discussion somewhat.

We almost completely ignore problems linked to the fact that there is no one global government, but rather a collection of states which are very different from one another: they have different economies, different values, and also different endowments of natural resources such as fossil fuels. This leaves individual states that wish to pursue a cleaner, greener global development path with a much more complex problem than that which would face a global government; they must either seek international agreement, or—if they act unilaterally—they must weigh up both the direct effects of their actions on the global environment, and also the indirect effects via the effect on *other* countries actions. These effects may be in the opposite direction to that desired, a form of rebound effect.

Technological progress and the human takeover of spaceship Earth

1.1. The expanding choice set

The story of modern humans (Homo sapiens) and their takeover of spaceship Earth is the story of technological progress and its consequences. Over 2 million years ago the first human species evolved, and learnt to control fire and make stone tools. For most of the intervening period technological progress was minimal. Various human species existed and spread throughout Eurasia, but they did not have a dominant position in the ecosystems in which they lived; they were neither top predators nor a major influence on plant life. And their numbers and biomass were modest compared to many other species.

Within the last 100 000 years all this has changed, as Homo sapiens (the species which emerged around 300 000 years ago) developed the ability to cooperate in large numbers and transcend the limits imposed by a strictly genetic evolution. Supplementing the—relatively glacial—genetic evolution we now have *cultural* evolution, the development of new ways of understanding the world, communicating, and behaving. Part and parcel of this cultural evolution is technological progress, where I say 'progress' rather than 'change' because I mean the ability of humans to manipulate their environment (both physical and biological) to deliver desired results, and this ability has increased over time; note that the 'desired results' could be anything from a warm and dry place to live, to a thermonuclear explosion. In the language of economic modelling, technological progress expands the choice set available to humans.

1.2. Three alternative choices

Many animal species use technologies, for instance building nests and hives. And there are also many examples of animals learning behaviours and passing them on between individuals and over generations. A famous example is blue tits learning to break through the aluminium tops of milk bottles to get at the cream, as the bottles stood on 20th century British doorsteps.[1] Early in the 20th century milk was placed on British doorsteps in open bottles, and blue tits and robins learnt to drink the cream which rose to the top. However, when thin aluminium tops were applied, only the more social tits (rather than the territorial robins) learnt to break through the tops and get at the cream, the reason being that individual robins who discovered the trick did not pass the knowledge on to other members of the species, whereas the tits did. However, the discovery by the British tits was not one of a series which led to them furthering their cooperation, outcompeting all other birds, and together ruling the roost. The discovery was instead an isolated incident, and the practice died out with the increasing popularity of skimmed milk (no cream) and the decline in doorstep deliveries.

Technological progress within a species increases the choice set available to that species. There is of course no guarantee that the choices made will increase the wellbeing of that species; for instance, as mentioned above humans may use their control of their environment to make small bits of it dramatically warmer and drier and hence more comfortable to live in, or to blow large areas of it up with thermonuclear bombs. How can we categorize or analyse the choices available? What choices have we made in the past? And what will we choose in the future? To get a handle on these questions, we return to the tits, and describe their alternatives in economic rather than biological terms. What did the British population of blue tits do with their new-found ability to take the cream from milk bottles? Let us assume that the cream provided an easy and steady source of nutrition for the birds. On discovering this source, they could hypothetically have reduced the total amount of time they spent foraging for food (work) and increased the time spent on leisure activities such as hanging out in the local oak tree chatting to their friends. Alternatively, they could have continued to spend as much time working as before, but increased their consumption of goods. They could

---

[1] See `https://www.sciencedirect.com/science/article/pii/037663579400051H` and `https://www.britishbirdlovers.co.uk/articles/blue-tits-and-milk-bottle-tops`.

have eaten more, or diversified into other goods, perhaps making larger and more elaborate nests. Finally, they could have used the extra food supply to increase their reproductive success, and hence to increase the total population of tits. In this scenario they continue to show similar patterns of behaviour, but the extra food allows them to better survive tough periods (e.g. winters), give their chicks more food, and to raise more or larger broods. So we have three (potentially overlapping) categories: increase leisure, increase consumption, and increase reproductive success.

The theory of evolution tells us unambiguously that each individual tit will use the newly discovered food source to increase its reproductive success. The tits' behaviour patterns have evolved to maximize reproductive success, and the discovery of the extra food source will not change that. If individual reproductive success translates to the level of the entire species, the effect will be—over time—an increase in the population of tits, at the expense of humans (who lose some of the cream from their milk) and quite possibly at the expense of other species who compete with tits, for instance for nesting sites. On the other hand, blue tit parasites and species that live in symbiosis with the tits will benefit.

Humans have transcended evolutionary imperatives to a far greater extent than blue tits, hence there is no simple rule or theory which tells us how we humans will use the power given to us by a new technology. However, the three categories above—leisure, consumption, and reproductive success—are useful, and we can draw some general conclusions. Up to around 1700 CE, the dominant use of our increased power was to increase reproductive success; we denote this phase *Malthusian*. In the majority of developed economies today, the power of new technology is mainly used to increase consumption, a phase we call *consumerist*. In the future, the trend towards increased leisure—which has been clear but weak in the developed economies over the last century —may become more important, and could give rise to a 'technotopia' in which technology allows us to combine leisure, consumption of material goods, and high environmental quality (including respect for other species).

1.3. The Malthusian phase

In a pure Malthusian model we can think of Homo sapiens' reproductive choices as being driven by biological rather than cultural factors. Hence when new technologies allow sapiens to extract more goods (more food, better shelter) from a given area, sapiens' reproductive success increases, and the population of that area increases. Furthermore, when new technologies allow sapiens to colonize new areas which were previously inhospitable, the total population increases as the colonized areas are populated.[2]

Over most of the Malthusian period we can assume that the rate of technological progress is slow compared to the adaptation of the population level to the new possibilities opened up by each technology. For instance, the invention of clothing—probably around 170 000 years ago, and the sewing needle—at least 50 000 years ago—would have significantly increased sapiens' ability to colonize new areas. These inventions were separated by around 120 000 years, but given the right technology sapiens colonized entire continents (such as North and South America) in a matter of a few thousand years. A consequence of this, as we see in Figure 1.1, is that during the Malthusian period humans were unable to get very far 'ahead of the curve' through technological progress, by which I mean that each new discovery gave only a short-lived period of plenty before the gains had been eaten up by an increased number of hungry mouths.[3]

Perhaps the most dramatic technological breakthrough of the Malthusian phase of human conquest of the planet was the invention of agriculture (known as the first agricultural revolution among other things), which occurred independently in many areas of the globe, but first around 12 500 years ago in what is now the Middle East. Prior to this invention, humans were exclusively hunter–gatherers. Agriculture allowed humans to extract far more food from a given area of land than could be achieved through hunting and gathering. Hence initially it would have seemed like a win–win option: combine some agriculture with hunting and gathering, and have more food while retaining the benefits of the traditional lifestyle, including a varied diet and tasks. However, where agriculture was adopted population expanded rapidly through the Malthusian logic, forcing a complete transition to agriculture and making a return to hunting and gathering impossible!

---

[2]Note that we call the phase Malthusian. The name is derived from Thomas Malthus, and in particular his book *An essay on the principles of population*, first published in 1798. He argued for the existence of precisely the mechanism captured in population model: that increasing productivity from a fixed quantity of land leads to increased population and not increased income per capita. Hence humanity is permanently on the borderline of survival, and the natural tendency for population to increase must be checked by death caused by disease, starvation, or (perhaps) self-control.

[3]The discussion of this section owes a lot to *Sapiens*, citetharari14.

FIGURE 1.1. Global product and population, historical (data from Brad DeLong). Both variables are normalized to start at zero. Population grows by a factor of approximately e^5, i.e. about 150. Average yearly production per capita is close to 100 USD throughout.

So technological progress allowed Homo sapiens to 'be fruitful and multiply' during the Malthusian phase which accounts for the vast majority of the period since the emergence of dynamic human cultures around 70 000 years ago. But what was the effect of humans (and specifically homo sapiens) on spaceship Earth during this period? Was it a period of harmony and respect for nature? Or were humans too few and feeble to inflict serious damage on other species and the global environment? The answer to both these questions is a resounding 'No!' Pre-industrial human populations wreaked havoc on spaceship Earth, in particular causing the extinctions of many of its most spectacular animals.

Homo sapiens evolved in Africa, thus co-evolving with the megafauna of that continent which thus learnt to cope with sapiens predation. We then expanded into other areas, wreaking ecological destruction on the unsuspecting species wherever we went. Sapiens arrived in Australia around 47 000 years ago, and the population of megafauna collapsed over a period from 45 000 to 43 100 years ago (van der Kaars et al., 2017). The same pattern is repeated over and over again for islands from Tasmania (41 000 years) through North and South America (13 000 years), New Zealand (700 years), and the Commander Islands just 250 years ago. Furthermore, the expansion of sapiens went hand-in-hand with the disappearance of other human species such as Neanderthals and Homo floresiensis.

Was sapiens' expansion out of Africa an 'ecological disaster'? There are of course no objective criteria on which to judge such a statement. Clearly it was a disaster for the species wiped out by sapiens, but on the other hand other species (such as those with a symbiotic relationship to sapiens, such as wheat) thrived. And there is little evidence that these extinctions were a disaster for *human* populations; having wiped out the megafauna, humans moved on to other sources of sustenance. Indeed, if humans had been dependent on the megafauna then the extinctions would probably not have occurred, since declining megafaunal (prey) populations would have led to declining populations of the predator (sapiens), and hence declining hunting pressure. We return to this in the economic models of Chapter 2.

Megafaunal extinction was a blip rather than a catastrophe for human populations because they could hunt other animals, but also because they had the alternative of gathering food, and (later on) agriculture. Agriculture also allowed the first large permanent settlements, leading to an acceleration in cultural change and technological development, development which went hand-in-hand with the development of money, markets, and the capitalist system, ultimately ushering in the consumerist phase.

1.4. The consumerist phase

In a consumerist model reproductive choices are driven purely by cultural rather than biological factors, and there is no direct link between technological progress (allowing us to increase our

FIGURE 1.2. Global product and population, modern (data from Maddison (2010) and US Census Bureau).

control over the environment and therefore extract greater flows of goods and services) and population.[4]

Figure 1.2 shows the global increases in per-capita production and population over the last 200 years. Growth in total global product is simply the sum of these two curves, hence we can see that the contribution of per capita increases is greater than the contribution of population increases. More importantly, as more and more countries complete the demographic transition over the next 50 years we expect the population curve to level off, whereas there is little evidence to suggest that average GDP per capita will stop growing any time soon.

Recall that in a pure Malthusian phase all humans consume only what they need to survive, mainly food and shelter. This was a decent approximation for most societies up to 1700, although typically there was a small elite which captured a significant proportion of the production of the hoi polloi and therefore had per capita incomes way beyond what they could spend on reproductive success. In the consumerist economy it is instead the case that the vast majority of the population have sufficient income to achieve a higher rate of reproductive success than that which they actually choose. The result is that consumers are faced with a myriad of choices to make: what should they do with their surplus income, after covering their basic needs? Clearly, they will not simply consume more of the same goods which were chosen in the Malthusian economy, i.e. staple foods and a roof over the family's head.[5] Should they consume more meat? Or clothing? Or pay for education for their children? Or travel? Furthermore, note that income itself becomes a choice variable: in the pure Malthusian economy everyone works as hard as they can in order to feed their family. However, in the consumerist economy we have the option to forego consumption for the benefit of increased leisure.

A further dilemma facing agents in the consumerist phase is how much to save, with this saving translating into investment on the production side of the economy. In the pure Malthusian economy with a more-or-less fixed set of technologies for production, saving simply has to be sufficient to ensure that there are seeds, tools, and livestock for the following year's production. However, in a consumerist economy (typically based on a capitalist system of production), agents have an incentive to forego consumption today in order to earn interest and thus consume more tomorrow. And producers have the opportunity to borrow money, allowing them to invest not just in maintaining their stocks of capital goods (replacing worn-out tools), but also in upgrading to more modern capital goods. And they may even choose to invest in developing completely new capital goods or production processes, thus contributing to technological progress.

The implications of these choices—regarding labour supply, consumption patterns, and investment—are immense, both for humans and our fellow travellers on spaceship Earth. From Figure 1.2 we can read off that global product per capita has increased by a factor of around $e^{2.5} = 12$ over the last 200 years—the value of what the average global citizen produces today is around 12 times greater (in real terms) than the value of what that citizen produced 200 years ago—

---

[4]We thus ignore research showing that, for instance, fertility choices—and hence long-run population trends—are influenced by economic factors, as argued by Barro and Becker (1989).

[5]This has been well known for over a century. Note for instance Engel's law, which is that as income rises, the proportion spent on food declines. See Engel (1857).

so the vast majority of our consumption (and production) is discretionary, we are free to choose. And the significance of our choices is immense. Consider our choice of what to eat. Agricultural production covers approximately 43 percent of global land mass (excluding areas of desert and ice), hence it is by far the biggest determinant of the fate of other terrestrial species.[6] Of this area, 83 percent is used to produce animal products—including areas used for growing crops which are then used as fodder—which yields just 18 percent of the calories. Clearly, if we all switched to vegan diets we could free up a very large proportion of global land-mass for animal and plant species other than wheat, corn, soy, and cattle. Regarding investment, consider the energy sector. Given sufficient investments in alternative sources of power we could achieve a massive fall in global CO_2 (henceforth *carbon*) emissions to the atmosphere within a decade or two, without the need for significant changes in consumption patterns. Furthermore, investments in technological progress within these alternative energy sectors would be likely to reduce the costs of such a switch significantly.

What choices have we made over the last 200 years? The first thing to note about our choices in the consumerist phase is that we have only increased our leisure time to a relatively modest extent, which is the reason why we denoted this phase the consumerist phase. Secondly, throughout the consumerist phase we have actually put a large and increasing proportion of our productive effort into investment in knowledge, rather than into making consumption goods. The vast majority of this investment is, however, intended to allow us to make more or better consumption goods in the future. This investment allows each worker to produce (on average) a greater value of goods in a day's work. Furthermore, in many cases it allows us to increase the productivity with which we use materials and primary energy sources, allowing greater production of goods from the same physical inputs. For instance, given a megajoule of primary energy in the form of coal, we can produce a lot more lumens of artificial light today than we could in 1800, and we can transport the same load a lot further. Ultimately—apart from environmental and sustainability benefits—this knowledge helps us make more consumer goods, because it implies that fewer of us are needed in the mines extracting and processing natural resources.

So our investment in knowledge (and capital) and modest increases in leisure time have allowed us to increase our per capita consumption dramatically. What have we added to basic foodstuffs and shelter in our consumption basket? Broadly we can say that productive effort has shifted from agriculture, to manufactured goods, and then to services. However, note that there is a close relationship between many services and manufactured goods; one of the biggest categories of services is transport! Looked at differently, we have shifted our consumption patterns into ever more resource- and energy-intensive goods. Within the food sector, we consume more and more meat, and particularly beef, requiring the largest land area (for feed production and grazing) per calorie produced of all the major meats. Turning to transport, we demand to be transported ever greater distances at increasing speeds, and (with regard to road transport) in heavier and more powerful vehicles.

So we have the ability to make goods and services (food, telephones, transportation) using less labour and also smaller flows of physical inputs such as metals and energy. We have continued to pump in labour, thus leading to rapid growth in GDP per capita (Figure 1.2), and delivering a large amount of discretionary consumption. Figure 1.3 shows that the net effect of changes in productivity and changes in consumption patterns is that resource and energy use have (broadly) tracked global product over the past century and beyond (we have data going further back for energy). This implies that the amount of metal and primary energy used per unit of value produced has remained unchanged; put differently, the efficiency with which metals and primary energy are used in the production process has not—on the aggregate—increased. So the combined effect of the two factors discussed above—changes in the productivity of metals and primary energy in making given products, and shifts in the pattern of products produced and consumed—has been to leave total resource and energy efficiency almost unchanged, implying steep increases in resource and energy use. Given the finite nature of the stocks of physical resources on the spaceship, and the fixed inflow of sunlight, these trends raise questions about long-run sustainability and intergenerational equity which we address below.

If metals are abundant, and energy can be obtained from the sun, the trends in physical resource consumption might not be a problem. However, the pollution linked to the large-scale use of resources is undoubtedly a problem. Over the last century we have emitted pollution to the atmosphere which has caused brain damage in our children on a staggering scale (lead), partially destroyed the upper atmosphere's ability to filter out damaging ultra-violet radiation (CFCs), acidified soils and waters over vast areas thereby severely damaging forest and aquatic ecosystems

---

[6]See for instance Poore and Nemecek (2018) for a recent analysis.

FIGURE 1.3. Long-run growth in total consumption compared to growth in total global product, for (a) Metals (tons extracted), and (b) Primary energy from combustion (joules burnt). GDP data from Maddison (2010), and metals data from Kelly and Matos (2012). For energy data sources see Hart (2018a).

(SO_2 and NO_x), and significantly altered the global climate (CO_2, CH_4, etc.).[7] However, by contrast to the aggregate resource data of Figure 1.3, the pollution data show that steep rises in pollution emissions (often steeper than the rise in GDP) are often followed by even steeper declines; in Figure 1.4 we see this pattern for sulphur dioxide emissions in the UK, and CFC emissions globally. Thus is seems that some factor which does not apply to global resource use does apply to national (and sometimes global) pollution emissions. Will this factor lead to declining flows of all pollutants in the long run? And will it (assuming it exists) lead to a decline in global carbon emissions precipitous enough to avoid catastrophic damages? We return to these questions—closely linked to the environmental Kuznets curve hypothesis—in Chapter 8.
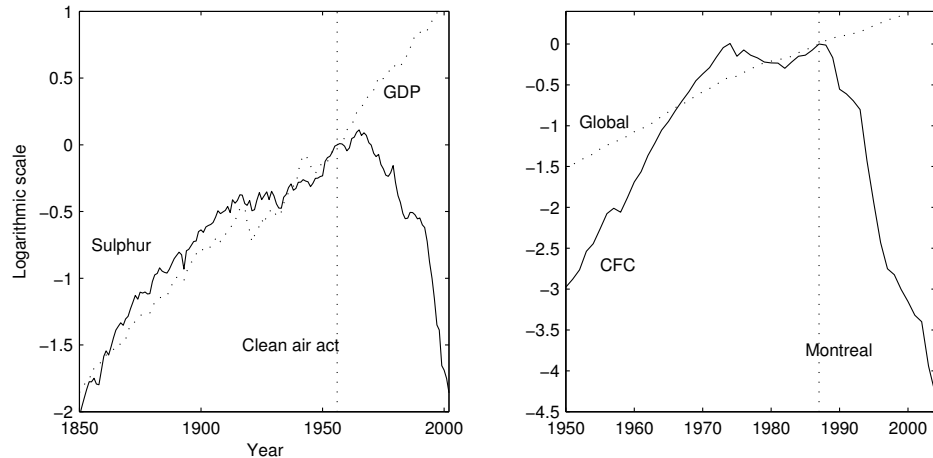
FIGURE 1.4. UK Sulphur emissions compared to total UK GDP, and global CFC production (CFC11+CFC12) compared to total global product. Sulphur: both normalized to zero in 1956, the date of introduction of the first of a long series of regulations restricting emissions. CFCs: both normalized to zero in 1987, the date of signing of the Montreal protocol. Data: Maddison (2010) (GDP), Stern (2005) (Sulphur), AFEAS (CFCs). AFEAS data downloaded from http://www.afeas.org/data.php, 9 Nov. 2014. Two anomalous points in the sulphur data have been altered.

1.5. Managing the future – A technotopian phase?

The consumerist phase has been characterized by rapidly expanding natural resource use, and ever-increasing dominance of humans over the global ecosystem. However, it has also seen the

---

[7]For more on these pollutants and their regulation see von Storch et al. (2003), Sunstein (2007), Ellerman et al. (2000), and Stern (2008) respectively.

gradual development of institutions to manage the effects of the economy on that ecosystem. Adam Smith (1776)—with his famous metaphor of the invisible hand—argued that private efforts to earn income for personal gain promote the interests of society, and this insight has evolved into a fundamental result of welfare economics, that given perfect markets with rational individuals making decisions based on self-interest, no-one can be made better off without someone else being made worse off (the First Welfare Theorem). But theory (developed by Pigou (1920) among others) and evidence regarding polluting emissions and their regulation shows that private choices must also be tempered by regulations when there are negative external effects, i.e. when choices by one agent lead to negative effects on others for which the agent does not have to pay.

The ultimate aim of economic analysis is to find better ways of managing the future. In order to better understand how to manage the future of the global economy, we first try to understand what lies behind the choices we have made during the consumerist phase described above. Based on such an understanding we can draw tentative conclusions about what policy actions are required today and in the future, and what the trajectory of the economy may be given their implementation. Given a single external effect—pollution damages—the policy problem is typically rather straightforward, although even here intractable problems arise when aggregate polluting emissions are damaging but it is costly to measure the contribution to those emissions of individual agents such as households and firms. Furthermore, in the global economy there are typically multiple external effects, including due to the existence of transboundary pollutants and knowledge externalities.

Here we mention two key underlying trends, both based on the assumption that technological progress continues. The first is that, given continued technological progress and hence an increasing ability to manipulate our environment to deliver desired results, we will increasingly prioritize higher environmental quality relative to consumption of human-made goods and services. The reason is straightforward: when technology is primitive and we have little control over our environment, environmental quality is typically high and the conditions for non-human species are good, whereas our consumption of human-made goods and services is meagre. As technology improves we therefore prioritize the former over the latter, and environmental quality declines while consumption increases. However, as consumption increases we care more and more about environmental quality, and therefore we tend to ensure that both increase over time. Since damaging effects of our actions on the environment are typically *external* (in an unregulated market we can emit pollution without having to pay for the damage), this shift is inextricably tied up with regulation of the actions of individuals and firms.[8]

The second underlying trend linked to continued technological progress regards the balance between the supply of labour (also linked to production and consumption) and leisure time. When technology is primitive then (in a Malthusian economy) we would expect humans to have little leisure time, and low consumption: we spend almost all of our available time producing sufficient goods and services to survive and reproduce.[9] If we move out of the Malthusian economy and start to deliberately control our reproduction then we can (and do) move into a state in which our production of goods and services per person is greater than what is needed for survival and reproduction, and we face a decision regarding how much labour to supply. The more labour each individual supplies, the less leisure time she has. On the other hand, when we supply more labour our income increases and we can consume more. And at the aggregate level, when total labour supply increases then aggregate production, consumption, and investment increase.

It is not obvious what the long-run trend in labour supply should be. Given consumption greater than the minimum to satisfy basic needs, we would surely expect people to choose at least some leisure time (at the expense of some of the 'excess' consumption). Furthermore, we might think that as technology improves at least some of the potential benefit should be realized in the form of increased leisure, rather than all of it being dedicated to increased consumption. The data shows very slow and uneven (both across countries and over time) increases in leisure time over the last century; the vast majority of the potential benefit of improved technology has been dedicated to increased consumption rather than increased leisure. In the last part of the book we analyse why this might be. One possible reason is that productive effort by one agent generates a negative externality for others, because the first agent's higher income (and therefore higher consumption) makes the other agents feel poorer and hence of lower status. To the extent that our choices are

---

[8]Whether or not this increasing prioritization of environmental quality will extend to non-human species is an open question.

[9]Note that it is frequently claimed that people in pre-agricultural hunter–gatherer societies had a lot of leisure time, which seems to contradict the Malthusian idea. If they did indeed have a lot of leisure time, a possible explanation is linked to the inability of hunter–gatherers to save and invest: if the population is constrained by conditions during 'tough times' when there is no leisure time, then in the remaining 'good times' there could be a lot of leisure time even though population is at its maximum sustainable level given the available technology.

driven by the pursuit of higher status rather than higher consumption *per se*, we will never be satisfied by higher aggregate consumption, and will continue to prioritize labour over leisure (*inter alia* worsening environmental problems). Furthermore, if we could coordinate and agree to chill out a bit more, every agent's utility could potentially increase, at the same time as environmental quality and conditions for other species improved: technotopia.

Malthusian growth

In this chapter we develop a model of a Malthusian economy with technological progress linked to population growth given a resilient 'spaceship' environment. We go on to consider extensions and variations in which the environment and other species are included more explicitly, allowing for the possibility of ecosystem collapse (affecting humans), or extinction of other species. The Malthusian phase of sapiens' planetary takeover is relatively straightforward to model economically, especially if we assume a resilient ecosystem which sapiens gradually learns to control more and more tightly. In a purely Malthusian model humans simply meet their basic needs, which remain the same over time. Furthermore, deliberate investment in generating new knowledge is minimal, and technological progress can be treated as being linked to the total human population, but not linked to specific incentives to perform research. When we add ecosystem dynamics then things get more interesting, and we can investigate the mechanisms behind megafaunal extinctions (as do for instance Bulte et al., 2006) and the potential of humans to destroy the conditions needed for their own survival (as do Brander and Taylor, 1998).

2.1. Technological progress and Malthusian population growth

Recall the data shown in Figure 1.1, where we saw that over a very long timescale up to 1700 AD, per capita production scarcely increased, whereas population increased at an increasing rate. In this section we build an economic model which captures two fundamental ideas which together can explain these results. The first idea is that when technological progress allows production per hectare to increase, this leads in the medium term to population increase rather than an increase in consumption per capita. And the second is that technological progress is driven by discoveries made by people, so when the population increases, the rate of progress tends to increase. Hence population growth tends to accelerate.

We start in a situation with a small number of intelligent beings with the ability to use tools and develop and apply new technologies over time. We call these beings *people*. For simplicity we can think of the 'spaceship' on which they live as consisting of a single island. Their initial technology is primitive, and they use the abundant mineral resources, the trees, and the land, in combination with sunshine and rainfall, to produce food and shelter. Their production technology is such that

$$Y = (A_L L)^{1-\alpha}(A_R R)^{\alpha},$$

where Y is total production of food and shelter, A_L is an index of labour productivity (linked to technology), L is the population (assumed the same as labour supply), A_R is an index of the productivity of the land, R is the area of land, and α is a positive parameter less that 1.[1] We assume that both A_R and R are fixed, so both the quantity and the quality of the land are fixed. In other words, we have a stable, resilient ecosystem.

We are most interest in production per person (or *per capita*), hence we divide through by L to obtain

$$Y/L = A_L^{1-\alpha}(A_R R/L)^{\alpha}. \tag{2.1}$$

So production per capita declines in population, since higher population dilutes the land available. The size of this effect depends on α: in the limit of $\alpha = 1$ then the quantity of land is the sole determinant of production, whereas for smaller α more people can extract greater total production from a given area of land. In the parameterized model we set $\alpha = 0.3$.

We model changes in population using a standard growth function from biology, the logistic:

$$L_{t+1} - L_t = \theta L_t(1 - L_t/L_t^*). \tag{2.2}$$

---

[1]Note that the people may also use capital goods in their production process, things like axes and hoes. However, we leave these goods out of the production function by incorporating them into A; that is, we assume that given their technology, the people have sufficient capital to achieve maximum productivity. So for instance for the technology of 'stone axes', there is no value in having more than one such axe per forestry worker.

Here θ is a parameter, and L_t^* is equilibrium population, often called the *carrying capacity*: when $L = L^*$, $L_{t+1} = L_t$ and population is constant, but when $L < L^*$ population tends to increase, and when $L > L^*$ it tends to decrease.

To link equations 2.1 and 2.2 we assume a level of income per capita Y/L at which the population is stable: we denote this level as $\bar{y}$. When $Y/L > \bar{y}$ then we have 'good times' and the population grows, and when $Y/L < \bar{y}$ then the population falls. It follows that when $Y/L = \bar{y}$, $L = L^*$. Insert these equations into equation 2.1 to yield

$$\bar{y} = A_L^{1-\alpha}(A_R R/L^*)^\alpha.$$

and rearrange:

$$L^* = \frac{A_R R}{\bar{y}^{1/\alpha}} A_L^{(1-\alpha)/\alpha}. \tag{2.3}$$

The equation shows that the bigger the island and the higher the productivity of the land, the higher the population that can be sustained. Furthermore, the less people need in terms of food and shelter, the higher the population. And finally, the higher is labour productivity A_L, the higher is the sustainable level of population L^*.

Now recall that the data shows us that on spaceship Earth the human population has been increasing for tens of thousands of years. How can we explain this increase? The reason must be a long-run increase in the carrying capacity L^*, driven in turn either by in increase in A_R (land quality) or an increase in A_L (labour productivity). So, either exogenous changes are making the global environment more hospitable for the species, or the species itself is developing in some way (possibly managing its environment) such that a higher population is sustainable.[2] Put this way, it should be obvious why human population has risen since 20 000 BC: over time we have learnt more and more about how to control and utilize our environment, and this has allowed us to increase in number, partly by populating our existing (20 000 BC) range more densely, and partly by extending our range to habitats that were previously inhospitable for us. Put another way, technological progress has allowed our population to grow. But how do we model technological progress?

We model technological progress using the following equation:

$$A_{Lt+1} = A_{Lt}[1 - \delta + \zeta(\Omega L_t)^\phi], \tag{2.4}$$

where δ is knowledge depreciation, Ω is the proportion of the population (or the proportion of total labour) which is devoted to the generation of new ideas, ζ is a productivity parameter equal to the progress resulting from the idea coming from 1 person's full time idea-generation, and ϕ is a parameter $\in (0,1)$; the lower is ϕ, the greater is the overlap between the ideas, and in the limit of $\phi = 0$ everyone has the same idea. We thus assume that individuals exogenously come up with new ideas about how to organize the world, once per period. Regarding these ideas, we assume that the size of the step forward through a new idea is in proportion to the stock of knowledge on which it builds. Furthermore, since these ideas 'flow from the same spring' in the sense that the individuals all have approximately the same knowledge on which to build, there is a lot of overlap between them hence the progress yielded by each idea is not additive.

We now have all the equations we need to simulate the development of the economy over time from a given starting point. Figure 2.1 shows the simulated global economy—compared to the data presented previously—using the following parameters (with 20-year periods): $\bar{y} = 1$, $R = 1$, $\alpha = 0.4$, $\theta = 0.06$, $\phi = 0.2$, $\delta = 0.02$, and $\zeta = 0.020012$, with the starting point $A = 1$, $\Omega L = 1$. (We have subsequently normalized to match the levels in the data.) Note that the model produces a reasonable match to the data. Two points must however be raised at this point: firstly, the data are very uncertain; secondly, the model is extremely simple and is only intended to show one way in which population, production, and technology can be linked in a dynamic model.

2.2. A fragile ecosystem and megafaunal extinction

An obvious objection to the above model is that the ecosystem is not resilient. A potential proof of this lack of resilience is the wave of megafaunal extinctions which occurred as modern humans (Homo sapiens) expanded their range around 'spaceship Earth' starting around 50 000 years ago, starting with Australia and nearby islands. In the terminology of Hart (2002), following Thompson et al. (1990), nature may be assumed to be benign (corresponding to our resilience but not as extreme), capricious (changing quality randomly), perverse/tolerant (benign within limits

---

[2]The model developed here is in some ways related to the models of Oded Galor, see for instance Galor (2005). However, Galor's models rely on a completely different—and more complex—mechanism to deliver the acceleration in technological progress which we ascribe to population growth.
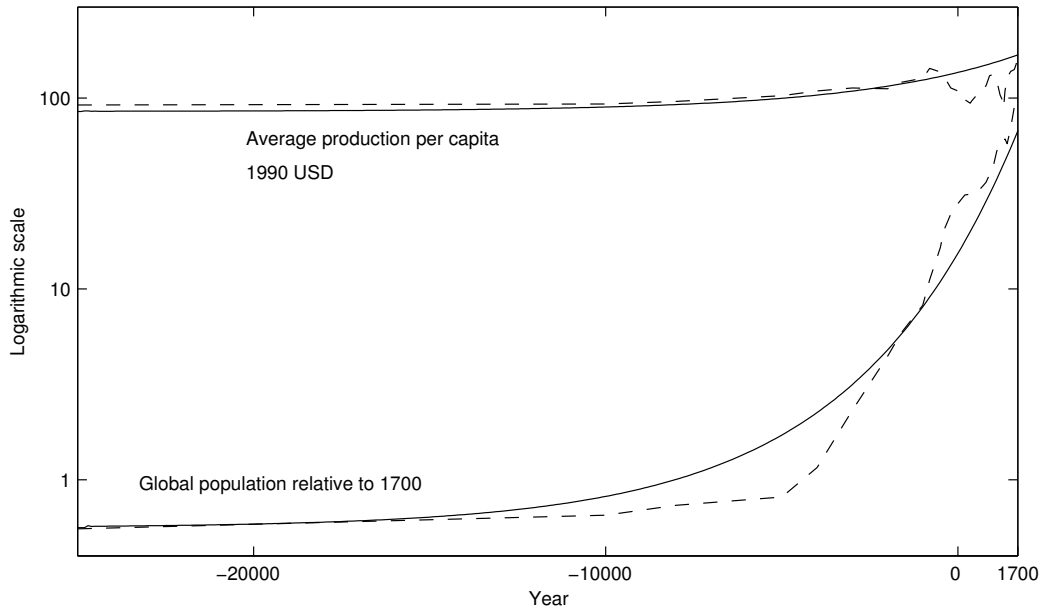
FIGURE 2.1. Global product and population. Continuous lines, model simulation; dashed lines, historical data from Brad DeLong.

but subject to collapse), or ephemeral (ready to collapse at any moment). These assumptions about nature are illustrated in Figure 2.2, where they are linked to alternative ideologies.
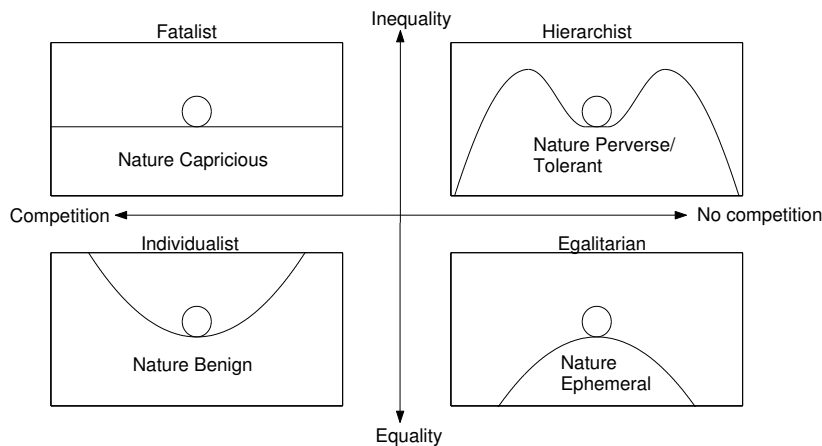
FIGURE 2.2. States of nature, adapted from Thompson et al. 1990

I illuminate the correspondence between the ideologies and the assumptions about nature— borrowing from Hart (2002)—with the following caricature of the debate about climate change. If nature is benign, then even if the climate does change nature will adapt and there will be no 'catastrophe' or even major damages. This belief may be held by an individualist, perhaps a businessman, who believes in a world in which agents are free to use their imagination and initiative in order to get the best out of the 'raw materials' that nature provides. Mistakes will be made, maybe even big mistakes, but 'luckily' nature is benign, always returning to its stable equilibrium, hence unbridled capitalism is socially optimal. The faith in 'nature benign' is thus essential to support the ideology (individualism). The hierarchist could be an environmental economist. She believes that in the good society, careful control of economic activity in general, and greenhouse-gas emissions in particular, is essential to protect public goods. Without control, disaster is certain, but with control it can be avoided—the supporting myth is thus nature perverse/tolerant. The egalitarian is an environmental activist who argues that even hierarchical control is not enough to save nature from the individualists—the only solution is to overthrow capitalism and close ranks in eco-communes. Again, the myth supports the ideology. Finally, the fatalist has little control over events, and thus adopts the myth of nature capricious; climate change is not caused by humans, but is an 'act of god' over which we have no control.

Thompson et al. (1990, p. 86–93) argue that types in opposite corners converse more readily than adjacent types. Hierarchists and individualists form an alliance which may be termed 'the establishment', keeping egalitarians and their demands at arm's length. Individualists need the stability—e.g. the rule of law—provided by the hierarchy, and hierarchists need the individualists' capacity for innovation. Egalitarians try to persuade fatalists that the collective offers an alternative to isolation and powerlessness, but are suspicious of alliances with coercive hierarchists or selfish individualists. These diagonal relationships are very clear in an economic context, the classic debate within the economics establishment being along the individualist–hierarchist diagonal. To what extent are hierarchical controls needed in order to achieve desired economic (or eco–eco) outcomes? In a world of perfect markets, the hierarchy (and its economists) are unnecessary. Neoclassical economics can then be interpreted as the study of hierarchical exceptions to the individualistic rule. Those who reject the use of 'neoclassical economic methods' (such as some ecological economists) may be placed on the fatalist–egalitarian diagonal.

Returning to the megafauna, we could claim that the extinctions were caused by capricious nature and had nothing to do with human action, following the fatalist line. Or we could follow the individualist and argue that the extinctions were entirely natural and benign, with little or no significance to humans: humans didn't cause the extinctions, and even if they did then they don't suffer from them. Or we could take the egalitarian's position and argue that the extinctions were an inevitable consequence of human greed and stupidity. However, for obvious reasons—I am an environmental economist—we choose the hierarchist's approach of trying to explain the extinction as the regrettable result of bad management of the 'resource'.[3]

Modelling the coexistence (or not) of humans and megafauna is in principle straightforward. Since the extinctions occurred over relatively short timescales compared to preindustrial rates of technological progress, we can safely ignore technological progress in our model.[4] The simplest possible approach would be to build a predator–prey model in which humans are completely dependent on 'harvesting' megafauna, the population of which decline under harvesting pressure.

We start with equation 2.1, but since the technology terms are constant we leave them out:

$$Y/L = (R/L)^\alpha.$$

Now R/L is interpreted as the harvest of meat per capita. Now we link this to the population of fauna, and hunting effort: $R = \phi E N$, where ϕ is a parameter, E is effort and N is the population, so the denser the population of animals, the easier it is to hunt. If hunting is the only activity then $\alpha = 1$ and we can assume $E = L$, so we have

$$Y/L = \phi N.$$

Food per capita is simply proportional to the animal population.

Now recall (from the previous model) that we defined $\bar{y}$ such that when $Y/L = \bar{y}$ then population is stable. So there is a unique animal population consistent with a stable human population,

$$N = \bar{y}/\phi.$$

Now assume that the human population grows according to a modified logistic function

$$L_{t+1} - L_t = \theta_L L_t [1 - \bar{y}/(Y_t/L_t)]. \tag{2.5}$$

So when production per capita is very large, population grows at the maximum rate, when it is $\bar{y}$ it is stable, and when it is zero the population collapses. Insert $Y/L = \phi N$ to yield

$$L_{t+1} - L_t = \theta_L L_t [1 - \bar{y}/(\phi N_t)]. \tag{2.6}$$

This gives us the equation for the dynamics of human population given the animal population. Now we simply assume a logistic function for the animal population, adding harvest pressure:

$$N_{t+1} - N_t = \theta_N N_t [1 - N_t/N^*] - \phi L_t N_t. \tag{2.7}$$

To test the model we can do two things. Firstly, find the long-run steady state analytically. Secondly, we can simulate the dynamics of the model numerically. To find the long-run steady

---

[3]Although historically it has had its supporters, the idea that the extinctions were not caused by humans is scarcely tenable; over 50000 years and at least 13 separate locations the arrival of humans has been closely followed by the disappearance of megafauna. See for instance Burney and Flannery (2005).

[4]Extinctions in New Zealand, after arrival of the Maori around 800 years ago, occurred over a period of as little as 100 years. Technological progress over such periods—if we assume that its rate is typical for the rate shown in Figure 1.1— would be around 10 percent.

state, assume $L_{t+1} = L_t$ and $N_{t+1} = N_t$, and solve the above two equations for unique values of L and N. The result is

$$N = \bar{y}/\phi \quad \text{and} \quad L = \frac{\theta_N}{\phi}\left(1 - \frac{\bar{y}}{\phi N^*}\right).$$

Note first that if $\bar{y}/\phi \geq N^*$ then the solution makes no sense. In this case there is no stable animal population which is consistent with human survival. Assuming $\bar{y}/\phi < N^*$ then the stable human population is decreasing in $\bar{y}$ and ϕ, and increasing in θ_N and N^*. These results are straightforward with the exception of the effect of ϕ: when humans become better hunters, the stable human population falls. The reason is linked to the 'tragedy of the commons' (Hardin, 1968). In the model we effectively have a free market solution with zero cost of effort, so each agent simply exerts maximum hunting effort in all circumstances, without regard to long-run consequences for themselves or the other agents. A far higher human population could be maintained if agents (perhaps with the help of hierarchical environmental economists) could be induced to plan their hunting effort together.

If humans are to maximize their long-run population then it is easy to see that they need to manage the resource to maximize yield. To find this yield, rewrite equation 2.7 as

$$N_{t+1} - N_t = \theta_N N_t[1 - N_t/N^*] - Y_t. \tag{2.8}$$

Now we want to find the value of N_t which maximizes Y subject to the restriction that the animal population is constant, i.e. $N_{t+1} - N_t = 0$. Given this restriction we have

$$Y^* = \theta_N N_t(1 - N_t/N^*),$$

and from the first-order condition in N_t we know that the optimal animal population in this case is $N^*/2$, and the maximum sustainable yield of meat is $Y^* = \theta_N N^*/4$. At stable population we require $Y/L = \bar{y}$, hence the maximum sustainable human population is given by

$$L = \frac{\theta_N N^*}{4\bar{y}},$$

which is greater than or equal to the population in the 'market' solution.[5]

Now we turn to the simulation. Below we see a very simple Matlab program to simulate this model, and Figure 2.3 shows the output. In the figure we see that human population rises steeply initially, which drives the animal population down. The lags built in to the model lead to a certain amount of overshoot, with human population rising beyond the sustainable level, which pushes the animal population below the sustainable level, which causes a crash in the human population, recovery in the animal population, etc. These oscillations gradually attenuate, and the steady state which we derived above is approached, in which $L = 0.18$. as opposed to the maximum sustainable population of 0.5.

```
clear
clf

thetaL=0.02;
thetaN=0.02;
phi=0.1;
Nstar=100;
ybar=1;

N(1)=Nstar;
L(1)=.01;
n=2000;
```

[5]Recall the market solution:

$$\frac{\theta_N}{\phi}\left(1 - \frac{\bar{y}}{\phi N^*}\right).$$

Define $x = \phi N^*/\bar{y}$ and write the condition

$$\frac{\theta_N N^*}{4\bar{y}} \geq \frac{\theta_N}{\phi}\left(1 - \frac{\bar{y}}{\phi N^*}\right).$$

Cancel θ_N, rearrange, and substitute in x to yield

$$\frac{1}{4}x \geq \left(1 - \frac{1}{x}\right)$$

hence $\quad 4x^2 - 4x + 4 \geq 0$

and $\quad (x-2)^2 \geq 0.$

FIGURE 2.3. Development of the numerical model over time

```
for t=1:n
L(t+1)=L(t)+ thetaL*L(t)*(1-ybar/(phi*N(t)));
N(t+1)=N(t)+ thetaN*N(t)*(1-N(t)/Nstar)- phi* L(t)* N(t);
end

t=[0:1:n]
plot(t,L,'-k', t,N./Nstar,'-k')
shg
```

There are many directions in which we could now take the analysis. A typical approach would be to set up a utility function and derive the development of an optimally managed economy. However, we focus instead on the problem with which we started, the analysis of megafaunal extinctions. Since there is no extinction in our model, we clearly have a problem.

There is no extinction in our model because as the animal population declines, harvest also declines, even for a fixed human population. Furthermore, the declining harvest also drives down the human population, further reducing the total rate of harvest. This eventually allows the animal population to recover. So, more generally, resilience is built into the model in two ways: firstly, because we depend on the ecosystem in question, so ecosystem damage also puts the brakes on our destructive activity; and secondly because the worse the state of the ecosystem, the harder it is for us to inflict further damage on it (we can't find the remaining individuals of the species). Since we know that humans did actually cause megafaunal extinctions, we need to return to our model. Or, more accurately, we turn to the model developed by Bulte et al. (2006).

Following Bulte et al., we now add other sources of food to the model economy. In particular, we add small animals as well as megafauna. We continue to assume that hunting effort is exogenous, but now hunters catch both small animals (the population of which is more resilient) and large animals which risk extinction. The harvest of small animals allows the human population to remain high even when the megafauna are approaching extinction, and hence the continued high hunting pressure can push the megafaunal population to zero (in a model with discrete individuals).

We assume a harvest function

$$Y = L(\phi_N N + \phi_S S),$$

where the subscripts refer to megafauna and small animals, and S is the small-animal population. However, we simplify the growth function for the small animals by assuming that the stock S adapts instantly to hunting pressure, hence $S = S^* - \phi_S S L$, hence $S = S^*/(1 + \phi_S L)$. Then we have

$$Y/L = \phi_N N + \phi_S S^*/(1 + \phi_S L).$$

And (following a similar procedure to previously)

$$L_{t+1} - L_t = \theta_L L_t[1 - \bar{y}/(\phi_N N_t + \phi_S S^*/(1 + \phi_S L_t))]. \tag{2.9}$$

Retain

$$N_{t+1} - N_t = \theta_N N_t(1 - N_t/N^*) - \phi_N L_t N_t.$$

FIGURE 2.4. Development of the numerical model over time

To understand this model, consider first the case in which $N = 0$ so humans are completely dependent on the small animals. Then we have $Y = L\phi_S S$, $Y/L = \phi_S S^*/(1 + \phi_S L)$, and $L_{t+1} - L_t = \theta_L L_t[1 - \bar{y}/(\phi_S S^*/(1 + \phi_S L_t))]$. Solve for L_t when $L_{t+1} - L_t = 0$ to yield

$$L = S^*/\bar{y} - 1/\phi_S.$$

This shows that if $\phi_S < \bar{y}/S^*$ then the small animals are too scarce (or hard to catch) to sustain a human population on their own. In this case, their existence increases the amplitude of the fluctuations seen in Figure 2.3, since they allow the human population to keep growing for longer as the megafaunal population crashes. See Figure 2.4(a), in which $\phi_S = 0.09$ and $S^* = 10$, with other parameters unchanged. On the other hand, if $\phi_S > \bar{y}/S^*$ then the human population can be sustained on small animals alone, and now the megafaunal population approaches zero, as in Figure 2.4(b), in which $\phi_S = 0.11$.

Note that the megafauna never actually become extinct in the above models, because of the hunting success function which means that as the population approaches zero, the rate of hunting success remains a small fraction of the remaining population. However, adding realistic features such as a minimum viable population of megafauna, or exogenous shocks to population (a 'trembling hand' population), then extinction could emerge from the model. In fact, in Figure 2.4(b) extinction would be guaranteed, since in the model as it stands the megafaunal population approaches zero; on the other hand, in Figure 2.4(a) extinction would be possible depending on parameters (including the size of the minimum population or the exogenous shocks). So adding a small-animal population which could support humans would guarantee extinction of the megafaunal population. And adding a small-animal population which could *not* support humans would increase the risk of megafaunal extinction. Finally—and crucially—note that in (a) if such extinction occurred it would in turn lead to human extinction in the long run.

A model something like that illustrated in Figure 2.4(b) seems like a promising candidate to explain megafaunal extinctions. In this model megafauna provide a rich source of meat and are relatively easy to hunt. When humans arrive, the population is vast. The result is very rapid growth in the human population, and a corresponding collapse in the megafaunal population. As megafauna become scarce, the proportion of small animals in the human 'catch' increases, allowing humans to sustain their numbers and thus to continue to drive the megafauna into the evolutionary abyss. For more discussion, see Bulte et al. (2006), where a model is developed which also includes the option of dividing time between hunting and agriculture. However, Bulte et al. conclude that the addition of this option to the model has rather modest effects on its properties; the key to extinction is the fact that the both small animals and megafauna can be caught during one and the same hunt.

The model above shows how something that expands our choice set—the option of harvesting small animals as well as megafauna—can lead to a much worse long-run outcome, although it gives short-run benefits. In this specific case, the small animals may allow the human population to remain large for long enough to wipe out the megafauna on which the population depends for its long-run survival. The result follows because we have assumed that the humans simply hunt the animals without accounting for the future. This is perfectly rational at the individual level (the tragedy of the commons), but not at the societal level unless the social discount rate is zero, i.e. people only care about the present and not at all about the future.

The opposite extreme is that people only care about the future. If we retain the Malthusian set-up in which long-run income per capita is fixed, a natural assumption is that utility is maximized when population is maximized. In that case, the human population could be much better-off if they reduced the hunting effort directed at the megafauna, and sustained a megafaunal population of $N^*/2$.

It is of course possible to draw parallels with modern problems such as global warming. Here the climate system corresponds to the megafauna, which we disturb through the burning of fossil fuels (among other things). If there is some factor—corresponding to the small animals—which delays the time at which we feel the full force of changes to the climate, this may actually lead us to a catastrophe which we would otherwise have avoided, because it leads us to react too late. If we believe that these delays (and the underlying processes hidden by the delays) are essentially unknowable then this leads us towards an 'egalitarian' view of nature (Figure 2.2) in which nature is ephemeral and we need to radically reduce economic activity to reduce the risk of collapse; on the other hand, if they are in principle knowable then we are in the hierarchist's world in which nature should be wisely managed.

There are several other papers in the literature exploring related questions, of which the best known is perhaps Brander and Taylor (1998). They build a model which they argue can help us to understand what happened on Easter Island, where archeological and other evidence suggests that humans—through overexploitation—found an ecosystem which was, from a human perspective, productive, and in a relatively short period pushed it into an new unproductive state. During this process the human population first boomed, and then crashed.

2.3. The demographic transition and the post-Malthusian economy

The Malthusian models developed above—and those of Brander and Taylor (1998) and Bulte et al. (2006)—show us how technological progress can lead to population growth without long-run increases in production per capita, and how excessive extraction without regard to the long-run effects can lead to population overshoot, environmental disaster, and population collapse. The basic Malthusian model is highly relevant for understanding long-run population growth, and the extended models may have some relevance to historical events such as megafaunal extinction and the fate of Easter Island, and the pattern may even be repeated in countries with rapid population growth today. However, none of the models are directly applicable to modern societies in which population is not limited by income, but through choice. So while environmental collapse may still be a possibility in modern societies, we need a different model to capture the causes and potential solutions.

Returning to Malthus, it is ironic that his essay was written at almost precisely the time at which the link between productivity and population was broken in his homeland, i.e. the time at which increased productivity actually led to permanent increases in real income per capita rather than increases in population alone: the demographic transition. That is, the transition from a situation with high birth and death rates, to one with low birth and death rates. The transition typically occurs at the same time as a country's economy industrializes, although that does not mean the industrialization as such is the direct cause; likely direct factors driving the transition include the strengthening of women's rights and education, and a reduced probability that newborn children will die before reaching adulthood.

We do not discuss the demographic transition further here, but simply note that global population growth is not on an inherently unsustainable path; indeed, global population is likely to level off at around 10 billion people (current population 7 billion), i.e. an increase of less than 50 percent over today's level. By contrast, average global product per capita is likely to grow by up to 50 percent every 20 years for the next 100 years or more. (50 percent every twenty years corresponds to 2 percent per year.) Thus economic growth per capita is likely to be far more significant with regard to resource demand and pressure on the environment than is population growth.

Given the greater importance of growth in GDP per capita than the growth in population, in the remainder of the book we leave out population growth completely from our models. Furthermore,

for the most part we treat economic growth as exogenous, something that simply occurs in the background (as above). However, in Chapter 3 we explore reasons why firms invest in new knowledge in a market economy, and in Chapters 5 and 7 we investigate the consequences of such endogenous investment for environmental policy.

Another feature of the post-Malthusian economy is that the population has moved beyond the constraints imposed by the availability of renewable natural resources on food production; instead, most of the population is employed in industrial or service sectors dependent on energy and the extraction of non-renewable natural resources.[6] In most of the remainder of the book we therefore ignore land (in the sense of the limited surface area of the earth) as a constraint on production of final goods, since land is mainly crucial to the food sector, and the size of the food sector shrinks in the long run.[7] However, note that land is also important in many ways not directly linked to human production of goods are services; for instance, land is important for human recreation and for the survival and thriving of non-human species, or 'nature'. In turn, humans may value nature instrumentally (because of the 'ecosystem services' it provides) and for its own sake, perhaps based on moral reasoning. In Chapter 8 we argue that such values become more and more important as technology improves and our power over nature increases.

---

[6]Note that this is not to deny the importance of renewable natural resources in underpinning production. Catastrophic mismanagement of such resources could still lead to collapse of advanced economies, just as it (more obviously) can lead to collapse of agrarian economies. However, we do not study this question here.

[7]In practice we know that the share of the agricultural sector in GDP tends to shrink steadily, and in most leading economies that share is now less than 1 percent.

Post-Malthusian — industrial — growth

In the post-Malthusian era we assume that technological change and population are no longer linked; instead, we simply assume constant population for the most part. Hence technological change leads either to increases in production per capita, or increased leisure time per capita, or a combination of both. In the 'consumerist phase' defined in Chapter 1, it leads primarily to increased production rather than leisure, and this is the phase to which we now turn. In this chapter we define what we mean by economic growth, and narrow down the list of candidates with which to explain long-run growth. We show — in a neoclassical growth model without including natural resources — that in the long run increases in the productivity of labour are essential to drive economic growth; such productivity increases are inextricably linked to the concept of technological change. We therefore go on to models focusing explicitly on technological change, starting with a vintage growth model and moving on to an endogenous growth model.

3.1. Production, GDP, and growth

Economic growth is growth in GDP, gross domestic product, i.e. the value of all the goods and services produced within an economy during a given year. Consider an economy with only one final product. Then GDP is simply the quantity of that good produced multiplied by its price. *Real* GDP is quantity × price in constant dollars. Growth in real GDP (in percent per year) is then simply the growth rate of production of the good.[1]

In an economy with many products the problem of measuring GDP growth is more complex. *Nominal* GDP is simply quantity × price for each product, summed over all products, and growth in nominal GDP is the growth rate of this sum. But what about *real* GDP? If prices are all constant and the range of goods available is also constant then there is no problem: nominal and real GDP (and their growth rates) are the same. But if the prices of different goods change at different rates, and completely new goods also appear, then the problem is much more complex. We do not go deeply into this problem here: suffice it to say that growth accountants try to measure the change in consumer's (and firms') willingness-to-pay for the set of goods produced, in constant dollars. If consumers are willing to pay 3 percent more for the 'basket' of goods produced in the economy in 2014 than for the goods produced in 2013 then GDP growth must have been 3 percent.

In economic models we describe production using production functions, which describe the quantity of a good (or potentially several goods) produced as a function of the quantities of the inputs used. We now investigate the nature of production functions, focusing on one case and discussing it by comparison to descriptions from chemistry and physics. The function on which we focus is for the production of iron from iron ore. Assume we have an economy with households who supply labour L, a stock of iron (III) oxide Fe_2O_3, and a stock of coal in the form of pure carbon, C. The workers L take the iron ore and coal and — using furnaces — reduce the ore to pure iron, with a by-product of carbon dioxide. A simple chemical description of this process is

$$2\,Fe_2O_3 + 3\,C \longrightarrow 4\,Fe + 3\,CO_2.$$

Note that — as in all meaningful equations — there is a form of balance or equality between the two sides. In this case, the number of atoms of each type (iron, oxygen, and carbon) is identical on each side, and the equation shows how these atoms are rearranged through a chemical reaction.

Another way of describing the process is to track the energy released during the chemical reaction. Assume that 2 moles of iron (III) oxide and three moles of carbon react (as above) within a reaction chamber at constant temperature and pressure, where the chamber is part of an isolated system. Then we know that

$$\Delta G = \Delta H - T\Delta S_{int},$$

where ΔG is the change in Gibbs free energy, ΔH is the enthalpy change of reaction, T is the temperature, and ΔS_{int} is the entropy change within the reaction chamber. This equation is mathematical (rather than chemical) hence each side of the equation must be equal to the same

---

[1]Note that we always work with real, not nominal, quantities in this book, unless otherwise stated.

number, and furthermore these numbers must be in the same units. In this case, the units are units of energy, so we choose the standard SI unit of *joules*. Since T has units of kelvin (the temperature scale used in physics and chemistry), S must have units of *joules per kelvin*; otherwise, the units would not be equal!

The first equation above tracks changes in the arrangement of atoms when the chemical reaction occurs; the second tracks the flow of energy in the reaction. An economic production function describing the same process is an accounting relationship describing how costly inputs can be used to make a valuable output. Assume that a firm employs labour to run the process, and that vast piles of iron ore and coal are freely available; furthermore, they are so large that they are thought to be inexhaustible. Assume also that facilities such as furnaces are also abundant and freely available. So the only costly input for the firm is labour. How then do we write the firm's production function? Very generally, we can write

$$Y_i = F(l_i),$$

where Y_i is the quantity in tons of the final product (iron) made by the firm per year, F is a function, and l_i is the number of workers employed (on an annual basis). Much more specifically we could write

$$Y_i = A_L l_i.$$

Now we have assumed that there are *constant returns* in labour, so if the firm doubles its labour inputs, the quantity of final product Y_i produced will double. Note that we can assign units to A_L, since the units on each side of the equation must balance: the units of A_L must be *tons worker^{-1} year^{-1}*. Thus A_L is a measure of the firm's productivity.

The firm is also using iron ore, coal, and capital in the production process, but we do not include them in the production function because they are assumed to be free and abundant. Furthermore, the firm also produces carbon dioxide (see the chemical equation) but again since this has no value we leave it out of the production function. However, as soon as one of the other inputs or outputs becomes costly or valuable (or as soon as we realize that it is valuable) then we should include it in the production function. For instance, assume that the pile of coal starts to run out, and firms start competing on a market to buy coal inputs; or assume that coal must be dug out of the ground, requiring labour. In either case we should include coal in the production function. If we assume that the firm has no flexibility whatsoever about its production process, and that a fixed quantity of coal is required for each ton or iron produced, then we should use a Leontief production function:

$$Y_i = \min(A_L l_i, A_R r_i).$$

This reads as follows: production Y_i is equal to the smallest of the following two quantities; effective labour inputs $A_L l_i$ and effective coal inputs $A_R r_i$. The units of r_i are tons of coal per year, hence the units of A_R must be tons of iron per ton of coal. Alternatively, the firm may find that it can trade off more labour for less coal (because with more labour inputs the workers can ensure that the coal is used more effectively) in which case some more complex function F is called for where

$$Y_i = F(A_L l_i, A_R r_i).$$

In a similar way we can add capital to the production function, with each addition making the function more complex, and therefore these additions are only undertaken when they are relevant, i.e. when they add to the explanatory power of the models. In economic analysis focusing on issues other than natural resources and environment it is common to leave out natural resource inputs from the analysis, since although they are not free they do not typically make up a large proportion of firm's costs. This does not make these analyses wrong, or in contradiction with the laws of physics; it simply means that they are focused on the question at hand.

Furthermore, firms typically produce more than one output. As well as the intended output, firms often produce unintended outputs or byproducts. These may be harmless or even have some value (consider sawdust from a sawmill), but they may also be harmful pollutants such as carbon dioxide. Again, in economic analysis focusing on issues other than natural resources and environment it is common to leave out these polluting byproducts from the analysis. In this book the questions are all about resource inputs and polluting outputs, hence it is essential to include these quantities in our production functions.

3.2. What drives growth? Reasoning from first principles

We saw in Chapter 1 (Figure 1.2) that average global growth has been remarkably constant and sustained over a period of more than 200 years. What has driven this process, and if such

growth is to continue into the future, what will drive it? We now turn to these questions. First we consider an economy where there is only one scarce input—labour—then we add capital and physical resource flows.

3.2.1. Labour. We begin with a model in which production involves the application of limited labour—together with abundant machines and raw materials—to make a single final product. The final product is measured by Y, units *widgets year^{-1}*. Labour is homogeneous (all workers are the same) so we can measure its quantity in a single dimension, L, units *workers*.[2] The production function is then as follows:

$$Y = A_L L,$$

where A_L is labour productivity and has units *widgets year^{-1} worker^{-1}*. It is essential to include A_L in the function, otherwise the units on either side of the equation cannot be the same. The properties of this function are intuitively reasonable. For instance, if we double the input of labour L while holding its productivity A_L constant then production doubles. Similarly, if we double labour productivity and hold labour inputs constant then production also doubles.

Now consider growth. Assume first that L (labour) is the same as population. Then we can of course increase total production Y if population increases. However, production per capita Y/L would not then change; in the literature Y/L is frequently denoted y (as opposed to Y for total production). Another way to generate growth would be if everyone worked longer hours, or if a greater proportion of the population joined the labour force. However, the long-run possibilities using this method are clearly very limited, given the high level of labour-force participation we already see, and the limited hours in the day. The conclusion from this discussion is therefore that the only way to raise production per capita in this economy in the long run is to raise A_L, labour productivity.[3]

3.2.2. Labour and capital. Now assume that machines (or more generally *capital*) are not abundant, i.e. they are no longer free but rather they are costly and their use must be accounted for. Might an increase in the quantity of capital be the driving force for growth? Since there is still only one product, widgets, then capital K is simply the number of widgets kept back within firms to help in the production process: it therefore has units of widgets. Our general production function becomes

$$Y = F(A_L L, A_K K),$$

where K has units *widgets*, and A_K has units *widgets year^{-1} widget^{-1}*. So widgets—the final product—may either be consumed or kept back within firms, where they can be used as tools or machines. Workers need these machines in order to produce, and the more machines each worker has access to, the more she can produce. Therefore the function F must be increasing in its two arguments, effective inputs of capital and labour.

Consider now doubling both capital and labour while holding their respective productivities constant. Then we have the same number of workers per machine, hence production per worker must also be the same, and total production Y must double. This implies that there are constant returns to scale in K and L together, implying in turn that there must be decreasing returns to scale in either of K or L separately: for instance, doubling K while holding L constant will not lead to Y doubling. Finally, the more capital we have (for a fixed number of workers) the less the benefit of adding even more capital should be: formally, $F_K'' < 0$. To understand this, assume that the 'machines' are actually hammers: once workers have one hammer each there is little or no benefit to saving up additional hammers.

These arguments tell us straight away that simply accumulating capital cannot give long-run growth. Consider again an economy in which workers use hammers to make final goods. If there is initially less than one hammer each then accumulation of capital (hammers) can boost production. But when every worker has a hammer then accumulation of additional capital will scarcely boost production at all, and when there is a mountain of hammers for each worker then additional hammers will definitely not boost production.

Since neither increases in labour nor capital can drive long-run growth in this model, it must follow that technological progress—growth in A_L and A_K—lies behind long-run growth in Y. In

---

[2]One worker is then one person who works full-time throughout the year, or (for instance) two people who each work half-time throughout the year.

[3]We typically ignore the distinction between labour and population, implying that we assume that the ratio of labour supply to population is constant. However, in Chapter 10 we study labour supply explicitly.

fact we can go further, if we rewrite the production function as follows:

$$y = f(A_L, A_K K/L).$$

The function f should have the same properties as F, implying that A_L (labour productivity) must rise to deliver long-run growth, whereas A_K (capital productivity) does not necessarily need to rise, since capital per worker may rise instead.

The above conclusions are strengthened further when we note that capital is costly: if we devote effort to building up capital, we take away effort from making final goods. So a focus on making capital goods reduces current consumption, and a focus on making capital goods which have no use (due to diminishing returns) reduces consumption in all periods.

3.2.3. Natural resources. We can also add natural-resource inputs to the production function above:

$$Y = F(A_L L, A_K K, A_R R).$$

Again, even if natural resource inputs are free, using more and more of them with given technology, labour, and capital will result in diminishing returns and not long-run growth. And if natural resources are costly to extract then devoting more and more labour to their extraction will lead to falls in net production and consumption rather than increases.

3.3. Neoclassical growth models

We now leave natural resources for the remainder of the chapter, and focus on capital and labour. In this section we explain what is meant by a *neoclassical growth model*; in the following sections we focus on specific versions of such models, especially the Solow model. In neoclassical growth models capital and labour are used to produce a single good, which can then either be retained in the production sector (boosting the capital stock) or consumed. Implicitly, in neoclassical growth models we assume that resources are so cheap that they can be treated as being available to the firm for free, and therefore do not need to be included in the (economic) production function.

3.3.1. The production function. We return to the production function

$$Y = F(A_L L, A_K K),$$

and assume the properties described above. Firstly, we assume that having more capital and labour available is never a bad thing for production, in other words

$$F'_K, F'_L \geq 0.$$

Second, we assume that the more we increase one input—while holding the other constant—the smaller is the marginal increase in production, i.e.[4]

$$F''_K, F''_L \leq 0.$$

Third, we assume that each input is essential, hence there is no production when K or L are zero, and furthermore the marginal products F'_K and F'_L approach zero when the quantity of the respective input approaches infinity. Fourth, and finally, we assume *constant returns to scale* in the physical inputs K and L: that is, if we double K and double L but hold the productivities A_K and A_L constant, output doubles. In a sense this assumption is already implied by the way we have defined the productivity indices; given this definition, the new assumption is that *changing the quantities K and L alone does not affect the productivity indices*.

This production function tells us immediately that accumulation of capital will boost production, but only up to a point; as the amount of capital in the economy increases, returns to further increases in capital diminish ($F''_K < 0$). Think of the economy with workers and tools. Increasing the number of tools available may boost production if tools are short, but once there are plenty of tools available then no further benefits will be felt.

3.3.2. Technology. In neoclassical growth models we generally think of technology as knowledge, a non-rival and non-excludable good which can therefore be used simultaneously throughout the global economy. It is thus common to assume that A_L and A_K are the same in all economies. Furthermore, at the time the neoclassical growth model was developed economists had no way of modelling technological progress as an endogenous outcome of firms' optimization decisions, hence productivity was assumed to grow exogenously, i.e. for reasons not explained in the model.

---

[4]Note that we also allow for the possibility that the marginal productivity of an input may be constant in the quantity of that input, at least over some interval.

FIGURE 3.1. A Solovian economy in which hammers are the final good.

3.3.3. Savings and capital. In neoclassical growth models, since there is only one product, capital must also consist of that product. And investment is simply the choice to keep the product in the production sector rather than consuming it. Furthermore, it is common to assume that the stock of capital depreciates at a constant rate δ. That is, in discrete time, if we have 10 units of capital in period t then (in the absence of saving) there will only be $(1-\delta) \times 10$ units left in period $t+1$. Note that this depreciation is entirely independent of whether the capital goods are actually used in production or not. The equation governing the evolution of the capital stock can therefore be written

$$K_{t+1} = (1-\delta)K_t + sY_t,$$

where s is the saving rate (which may be endogenous). Recalling the production function, we can write

$$K_{t+1} = (1-\delta)K_t + sF(K_t, L_t).$$

To get an intuitive grip on the model economy, consider Figure 3.1. What are the values of K, L, Y, and s? Assume that the economy is in a steady state, i.e. nothing changes over time. What must be the value of δ?

3.3.4. Firm optimization. In neoclassical growth models we assume that there are very many small firms in perfect competition. Mathematically we can think of a unit mass of firms, the result of which is that the economy behaves *as if* there is just one firm, but that the single firm is nevertheless a price-taker. Saying that the firm is a price taker is the same as saying that the price it pays for inputs is equal to the marginal product of those inputs, and the price it takes for its product is equal to marginal cost. This one firm is known as the *representative* firm. For more on this see Appendix A.3.

The representative firm hires capital and labour on the market. The firm's problem is as follows:

$$\max_{K \geq 0, L \geq 0} F(K,L) - w_K K - w_L L. \tag{3.1}$$

The prices paid by the firm—w_K for renting capital and w_L for labour—are equal to the marginal products of the inputs, that is $w_K = F'_K$, and $w_L = F'_L$. Note that total costs are then $KF'_K + LF'_L$. If profits are to be zero these must be equal to total income which is just Y, since we normalize the price of the final good to 1.

3.4. The Solow model with constant technology

The Solow model, which we develop here, is a very simple neoclassical model in which investment in capital is exogenously fixed at a fixed proportion s of production. A more sophisticated model is the Ramsey model which we discuss briefly in Appendix 3.A.

3.4.1. Distinctive features of the Solow model. The main feature which distinguishes the Solow model from other neoclassical models is the savings rate, which is simply fixed exogenously at s:[5]

$$K_{t+1} = (1-\delta)K_t + sY_t.$$

The second feature which is generally associated with the Solow model is the choice of the Cobb–Douglas functional form for the production function. Furthermore, we assume that technological progress increases the productivity of labour alone. That is,

$$Y = (A_L L)^{1-\alpha} K^\alpha, \tag{3.2}$$

where α is between 0 and 1. Compare this to the simple economy above in which $Y = A_L L$. Now, because both labour and capital are needed for production, a doubling in labour productivity increases production, but does not give a doubling in production. You should verify that the function has all of the properties necessary for neoclassical production functions.

A special property of the Cobb–Douglas function is that the elasticity of substitution between the inputs is 1. Assume that the relative prices of the inputs change, and that we have a representative firm which is a price taker. Then the elasticity of substitution is the percentage change in relative input quantities chosen by a firm divided by the relative price change, times -1. To calculate it, note first that quantities are simply K and L. Relative prices are given by the relative marginal products, and $\partial Y/\partial L = (1-\alpha)Y/L$, whereas $\partial Y/\partial K = \alpha Y/K$. Hence

$$\frac{w_K K}{w_L L} = \frac{\alpha}{1-\alpha};$$

$$\frac{K}{L} = \frac{w_L}{w_K} \frac{\alpha}{1-\alpha};$$

$$\frac{\partial K/L}{\partial w_K/w_L} = -\frac{K/L}{w_K/w_L}.$$

This confirms the unit elasticity of substitution, and highlights another feature of the production function: the relative returns to the factors are fixed. Capital takes a proportion α of total returns, and labour takes a proportion $1-\alpha$. To get the intuition, assume that the quantity of capital available increases; one result of this increase is that the price of capital decreases relative to the wage (the price of labour). Given Cobb–Douglas, the decrease in price exactly compensates for the increase in quantity such that returns to capital remain unchanged relative to returns to labour. This is an attractive feature of the model, since in reality we observe a remarkably constant division of returns between labour (70 percent) and capital (30 percent), across different economies and different times.

Finally, in the Solow model labour productivity is typically either assumed to be constant, or to grow at a constant rate. The productivity of capital is constant (and normalized to 1). In this section we assume that labour productivity is constant.

3.4.2. Solving the model. Consider the Solowian economy of Figure 3.1 in which the production function is Cobb–Douglas and A_L is fixed. Assume that $\alpha = 1/3$ and $\delta = 0.1$. What is A_L? What is s? Is the economy in long-run equilibrium?

In Figure 3.1 there are 10 workers, each with a hammer. Thus $L = 10$ and $K = 10$. Production Y is 5 hammers per period, so $Y = 5$. So

$$Y = A_L^{1-\alpha} \cdot 10^{1-\alpha} \cdot 10^\alpha.$$

So we know that $A_L^{2/3} = 0.5$. Furthermore, we see that 20 percent of production is saved, so $s = 0.2$. Finally, the economy is in long-run equilibrium because the quantity of capital remains constant over time: savings (which are equal to investment) are 1 hammer per period, which exactly compensates for depreciation (10 percent of 10 hammers per period).

Imagine instead the same economy, but that in the initial period there is only one hammer saved as capital. Work out production and saving, and depreciation, and hence the number of hammers

---

[5]Note that we have set up this version of the model in discrete time. The choice of discrete (rather than continuous) time is essentially arbitrary; we choose discrete time because it is slightly easier to explain aspects of the intuition behind the model in such a context. For more on discrete and continuous time see Appendix A, Section A.2.

FIGURE 3.2. The transition path in the Solowian economy of Figure 3.1, starting with one hammer.

available in the next period. Describe how the economy develops over time. Your description should match Figure 3.2.

Three things are notable from the figure. Firstly, even at the desperately low initial level of capital, initial production is at almost 50 percent of its long-run level. Secondly, the transition to the long-run equilibrium is rather rapid. Each dot represents one year, so we see that in just 15 years the initially very capital-poor economy is close to 90 percent of the production of the rich economy. Finally, and most fundamentally, there is zero growth in the long run.

Mathematically, we can easily work out the long-run steady state of the model. It is when $sY = \delta K$, i.e. when investment matches depreciation. Using the production function we have

$$s(A_L L)^{1-\alpha} K^\alpha = \delta K,$$

hence
$$K^{1-\alpha} = s(A_L L)^{1-\alpha}/\delta$$

and
$$K = A_L L(s/\delta)^{1/(1-\alpha)},$$

$$Y/L = A_L(s/\delta)^{\alpha/(1-\alpha)}.$$

So long-run GDP per capita is a linear function of labour productivity, and weakly increasing in the savings rate s.

Now assume instead that capital is available for hire on a global market at price w_K per year, and that we have many separate economies. If we index the economies by i then we have (for economy i)

$$Y_i = (A_L L_i)^{1-\alpha} K_i^\alpha. \tag{3.3}$$

Economy i has exogenously given levels of labour L_i; all economies have access to the same knowledge and therefore have the same technology A_L. What is economy i's level of capital K_i and production Y_i in long-run equilibrium, i.e. when the quantity of capital K_i is optimal?

When K_i is optimal then marginal returns to capital must be equal to the price, i.e.

$$\alpha Y_i / K_i = w_K.$$

Thus we see straight away that the ratio of GDP to capital is constant across countries! Furthermore, after substituting in the expression for Y_i and rearranging we have

$$K_i = A_L L_i w_K^{-1/(1-\alpha)}.$$

This tells us the optimal quantity of capital. Substitute it back into the expression for Y to yield

$$Y_i / L_i = A_L w_K^{-\alpha/(1-\alpha)}. \tag{3.4}$$

So—again—GDP per capita rises linearly with labour productivity A_L. Now it declines slowly in the rental price of capital, rather than increasing in s, the savings rate.

3.4.3. Conclusions on Solow with constant technology. The above results show that the Solow model with constant technology is incapable of explaining any of the key empirical observations about long-run growth and differences in GDP between countries. GDP does not grow in the long run in the model, and differences in GDP between countries should be modest, even if savings rates differ drastically between the countries.

The model is nevertheless useful, because it shows us what *cannot* drive long-run growth. Capital accumulation! This is a remarkable conclusion given that the long-run aggregate production function appears to have the form[6]

$$Y/L = K/L.$$

So even though long-run production per capita is linearly *correlated* with the long-run capital per capita, simply accumulating capital will not give growth!

3.5. Solow with exogenous technological progress

Given that capital accumulation alone cannot drive growth, we turn back to labour productivity A_L. Even though we cannot explain *why* it grows, it is clear that it must grow in the long run — at least in the model economy — to explain the growth process. Note that in this section we switch to continuous time which is convenient when there is positive long-run growth.

So let's assume that A_L grows exogenously at a constant rate g_{A_L}. For good measure let L grow too, at rate n. Now assume that there exists a *balanced growth path* (b.g.p.) along which all variables grow at constant rates (note that these rates may differ from one another, and may be zero or negative). What are the characteristics of such a path, if it exists?

We have

$$Y = (A_L L)^{1-\alpha} K^\alpha$$
$$\dot{A}_L/A_L = g_{A_L}$$
$$\dot{L}/L = n$$
$$\dot{K} = sY - \delta K.$$

Now, since we know that we are on a b.g.p. we know that $\dot{K}/K$ must be constant, hence $sY/K - \delta$ must also be constant, in turn implying that Y/K must be constant. Returning to the production function we therefore know that

$$(A_L L/K)^{1-\alpha}$$

must be constant on a b.g.p., implying that

$$\dot{K}/K = \dot{A}_L/A_L + \dot{L}/L.$$

In words, the capital stock grows as the same rate as augmented labour inputs, keeping the ratio of effective labour to capital constant. But, from the production function, we know that

$$Y/L = A_L[K/(A_L L)]^\alpha,$$

implying that on such a b.g.p. Y/L grows at the same rate as A_L! That is, per-capita GDP grows at the same rate as labour productivity.

Is such a b.g.p. stable, and (if so) how quickly does the economy approach such a b.g.p. if it has been knocked off course by a shock? It is stable, for the same reasons as the economy without technological progress: when there is 'too little' capital savings outstrip depreciation and capital (per effective unit of labour) accumulates rapidly; when there is too much capital depreciation gets the upper hand and capital (per ...) declines. Furthermore, simulations show that this process is rather rapid.

Together these facts tell us very clearly that differences in the quantities of the factors of production (in particular capital) available in different countries cannot explain the very large and very persistent differences in the productivity of labour (or GDP per capita) between countries. Instead the difference must lie in the different ability to make productive use of the factors (labour, capital, resources) available.

To make intuitive sense of the model, and relate the model economy to real economies, we need a picture of how labour productivity can increase even though the single product does not change over time. The easy way out would be to assume that it is because people work harder and harder, but this is clearly nonsense: there is nothing to suggest that the average worker in the U.S. produces 34 times more (in value terms) than the average worker in India because the U.S. worker works harder. A better picture is the following. Assume that the single product is

---

[6]That is, countries' GDP tends to be in close proportion to the value of their capital.

FIGURE 3.3. Testing the Solow model. Upper panel: German real GDP, compared to an estimated constant-growth trend (growth rate 2.3 percent per year). Data: Maddison (2010). Lower panel: Model simulation assuming the same growth trend in labour productivity and parameters $\alpha = 0.3$, $\delta = 0.1$, $s = 0.2$, with a massive shock to capital (the stock is divided by a factor of 30 at the end of year 24).

a *widget*, which is a wonderful good because it can rapidly (costlessly) be transformed into any number of different goods, including both capital goods (machines) and consumption goods (food, clothes, etc.). Over time (exogenously) new designs for machines are discovered, and the newer machines augment labour better than the old ones, they make each worker more productive. So the capital stock is measured as the number of widgets, but over time these widgets are arranged into more and more productive configurations. Furthermore, when the machines are optimally allocated the newer (more productive) machines demand a greater number of widgets.[7]

This analysis is supported by empirical evidence; for instance, we know that economies recover rapidly from a sudden loss of physical capital. This is well illustrated by the case of German recovery after WW2, as shown in Figure 3.3. We see that after just 15 years the economy has almost regained its old trend line, despite the very drastic loss of capital and hence also productivity in 1945. This matches well to a model economy with standard parameters, starting on its balanced growth path.

3.6. A vintage growth model

Through the Solow model we have learned that in order to continuously increase output it is not sufficient to continuously increase the amount of capital; rather, it is essential to continuously adopt new technology. The history and growth of GDP per capita is largely the history of technological development. Is our production of goods and services higher per capita than it was 100 000 years ago because we have more stone arrowheads today? Or is it higher because we have developed technologies such as agriculture, the smelting of iron and steel, the printing press, and the computer? But how can we describe and analyse the process of adopting new technology?

3.6.1. The basic model. Technologies—such as those mentioned above—are in one sense ideas, designs, or blueprints (as famously argued by Paul Romer, see for instance Romer (1994)). However, in order for them to boost productivity in the economy they must typically be embodied in capital goods: it is not enough to have a design for a printing press with moveable type, the machine itself is also required. So the growth process should consist of a cycle of research (and invention) followed by investment (and hence application of the inventions). Then more research, etc.

Given such a cycle, a huge number of questions arise. For instance, how much research to do, and in what areas? And (given the existence of a new invention) when to scrap the old capital and invest in the new? At one extreme we could always wait for the old capital to fall apart (if capital is very expensive and new designs are not much better than the old); at the other extreme we could always invest the moment a new design is invented (if capital is cheap and advances large). Both the research and scrapping problems are complicated to model. But for our purposes it is sufficient

---

[7]Assume for instance that at time $t = 1$ the optimal arrangement of widgets is called a *hammer*, each hammer consists of 100 widgets, and in the optimum each worker has 1 hammer and has productivity $A_L = 1$. At $t = 2$ a new arrangement called a *screwdriver* has taken over. Each screwdriver consists of 50 widgets, and in the optimum each worker has 4 different screwdrivers and is twice as productive as a worker with a hammer, so $A_L = 2$. And at $t = 3$ each worker has an electric screwdriver, which takes 400 widgets to make, and makes the worker twice as productive again, i.e. $A_L = 4$.
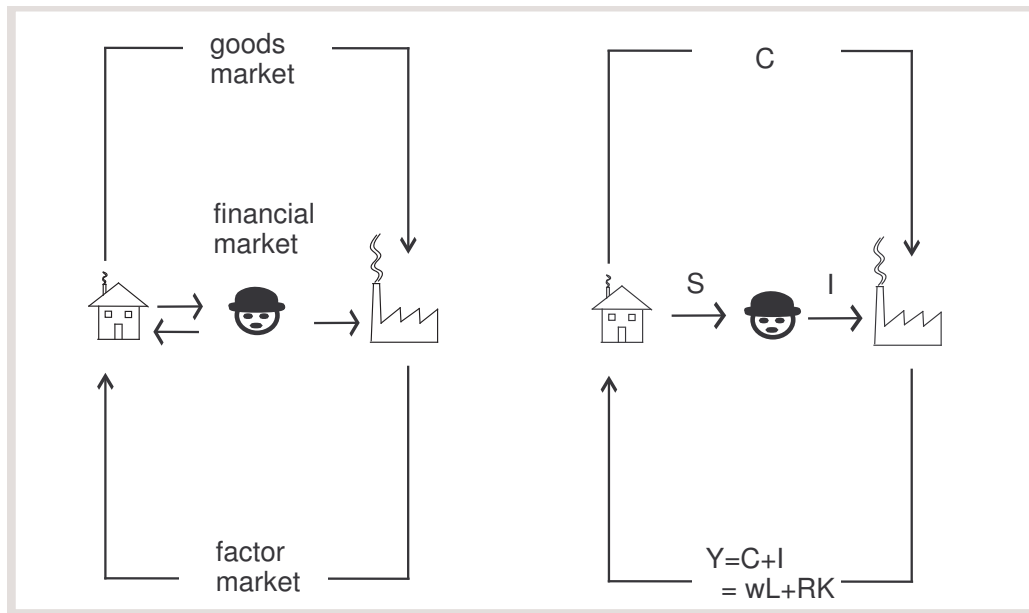
FIGURE 3.4.

for now to assume *exogenous arrival of designs* (so we follow Solow in assuming exogenous technological progress), and that it takes the same length of time to construct new machines as it takes for the existing machines to fall apart. So construction of the 'next generation' starts the moment the current generation comes into service. How might such an economy work?

Model economy 3.1.   *Assume an economy in which there are people and machines, and each machine needs one person to operate it. Machines last for ten years, at which point they fall apart irreparably and must be replaced by new machines. In the year 2000 there are 100 people and 100 (new) machines. Of these, 80 people—each with a machine—work on the production of consumer goods, while 20 people—each with a machine—work on the manufacture of the next-generation machines. 100 new machines are ready in 2010, and they are 20 percent more productive than the old. That is, they generate 20 percent higher output per period. All workers command the same wage, which is 100 crowns per period in the first period (from 1 January 2000 to 31 December 2009).*

To understand this economy, consider the fundamental picture of the circular flow of money in the economy, illustrated in Figure 3.4. (Note that money flows in the opposite direction to goods and factor inputs, which were illustrated in Figure 3.1.) Here we see the fundamental equation $Y = C + I$. We also know that Y—the flow of payments from firms to households at the bottom of the picture—is made up of wage payments to workers $w_L L$, and rental payments to capital owners, $w_K K$. So when Y increases, the sum of these payments—to capital and labour—must increase by the same amount. Furthermore, if we assume that capital and labour take *fixed shares* of the total cake, then payments to capital and labour will grow at the same rate as overall GDP. In reality it is true that the shares of capital and labour are rather constant, with capital owners typically taking around 30 percent of GDP, and workers taking the remaining 70 percent.

Now we turn to the specific economy in question, illustrated in Figure 3.5. Figure 3.5(a) shows the total flows during each period of 10 years, where Y_1 is total GDP in period 1. We know the flow of wages, but we know nothing so far about the value of the capital accumulated in the economy, nor the flow of payments to capital. Furthermore, since we do not know returns to capital we cannot work out GDP either. The fundamental problem is that we do not know the interest rate. Recall that we know that workers making machines are paid 100 crowns each per period in period 1. Who pays these workers? It must be investors who borrow money in period 1, planning to make a profit by hiring out the machines in the following period. What price do the investors set? It depends on the interest rate.

Model economy 3.1, continuation 1.   *Assume that the real interest rate per period of 10 years is 100 percent. What is GDP in period 1? What is the labour share, and what is the capital share?*[8]

---

[8]Note that an interest rate of 100 percent over 10 years corresponds to 7.2 percent per year, since $1.072^{10} = 2$.
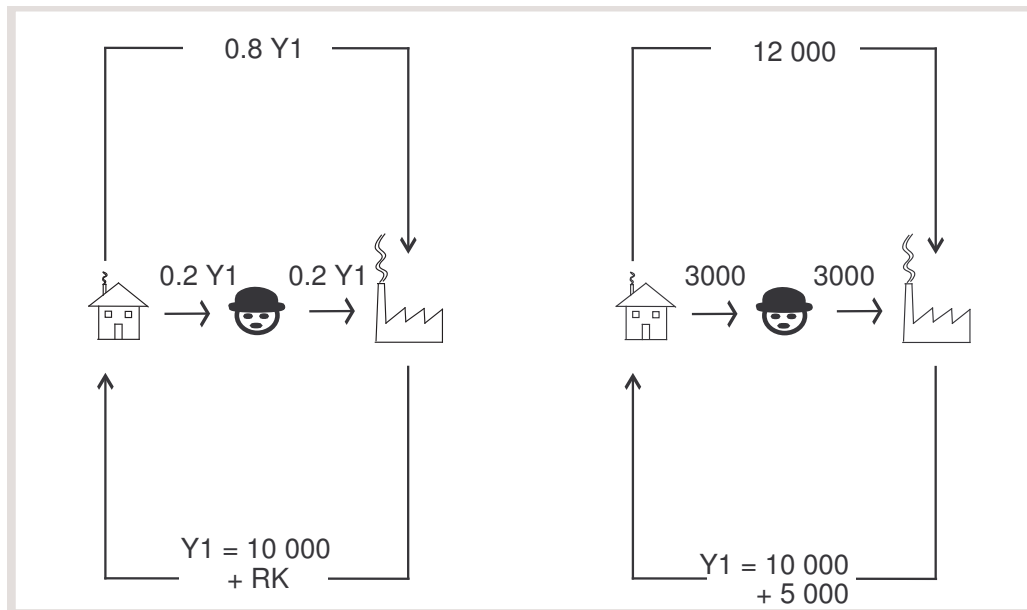
FIGURE 3.5. The circular flow in economy 3.1: (a) based on the information we are given initially; (b) after we have calculated returns to capital.

Recall that to make the 100 new machines to be used in period 2, 20 workers with 20 machines must be employed throughout period 1. Thus 20 percent of the economy's resources are directed to investment, costing 20 percent of GDP, Y_1. Therefore the total investment cost through period 1 is $0.2Y_1$ crowns. The interest rate is 100 percent per period, implying that the owners of the capital must pay back $0.4Y_1$ crowns at the end of period 2 (double the sum they borrowed in period 1). This implies in turn that they must earn $0.4Y_1$ crowns hiring out the machines during period 2 in order to break even. Therefore—assuming perfect markets—the cost of renting the machines must be $0.4Y_1$ crowns in period 2.

The final insight we need to characterize the growth path is that there is balanced growth at 20 percent per period. This means that if the cost of machine hire is $0.4Y_1$ crowns/period in period 2, it must be $(0.4/1.2)Y_1$ crowns/period in period 1, i.e. $0.33Y_1$ crowns/period. Thus payments to capital are 33 percent of GDP. Since per capita payments to labour are 100 crowns/period, per capita payments to capital must be 50 crowns/period, and GDP per capita must be 150 crowns/period. See Figure 3.5(b).

We have now characterized one growth path of this economy. However, we are far from done with our analysis, since so far we have simply assumed the investment rate (20 percent of GDP) and the interest rate (100 percent every 10 years). In real economies these numbers arise as a result of the decisions of economic agents. We now set about building a simple model to describe this process.

For now we assume that total expenditure per year (i.e. nominal GDP, PY) is fixed, and the only question is how this expenditure is allocated between consumption and investment. To analyse this allocation we consider the supply of and demand for investment funds. The supply of investment funds should be an increasing function of the interest rate, which we can think of as the price of such funds: the higher the interest rate, the more a household can earn by foregoing consumption and lending its money to those who wish to borrow. Therefore we have an upward-sloping supply curve. The demand for investment funds, on the other hand, should be a decreasing function of the interest rate: the higher is the interest rate, the fewer investment projects will be profitable. Therefore we have a downward-sloping demand curve. The result is a standard diagram, Figure 3.6, the only slightly non-standard feature being that the 'price' of investment funds, on the y axis, is the interest rate.

Now we can use the figure to analyse the effect of various shocks in the economy on the interest rate and the investment rate. First, assume (as in the economy without money) that households become concerned about the future, expecting bad times ahead. Their propensity to save therefore increases, and the supply curve for investment funds shifts to the right. From the figure we can see that the result must be a decrease in the interest rate, and an increase in the investment rate. Thus consumption does indeed decrease, and investment increases, just as it did in the economy without money. What is the effect of this shift in resource allocation, from C to I? Since $Y = C + I$, there is
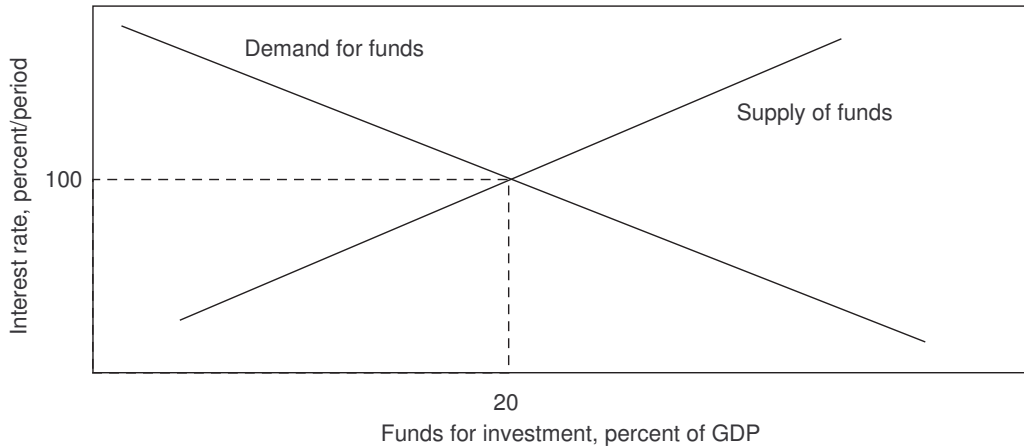
FIGURE 3.6. The supply and demand of investment funds

no immediate effect on GDP. However, over time there will be an effect. Since more resources are devoted to investment in new machines, this should allow the machine-makers to produce higher quality machines for the next period (recall that the number of machines is fixed, one per worker). Thus GDP will be higher in the next period, thanks to the higher propensity to save in the current period.

We now briefly consider two other shocks and their effects. Firstly assume that households become more optimistic about the future, causing them (ceteris paribus) to save less and consume more today. Thus the supply of money shifts to the left, the interest rate rises and investment declines. Assuming that lower investment translates into lower quality machines in the next period, the growth rate also declines. Secondly, assume that firms become much more optimistic about the returns to investment. This causes them to demand more investment funds today at a given interest rate, since their expectations about future profits are higher. Thus the demand curve shifts to the right, and both the interest rate and the investment rate rise. Assuming that higher investment translates into higher quality machines in the next period, the growth rate rises.

3.6.2. Rungs on the ladder. A wealth of evidence shows that countries far from the frontier can grow their economies very rapidly through adoption, much more rapidly than countries at the frontier can grow through R&D. Furthermore, it is not uncommon to see dramatic shifts in the trend growth rate, coinciding with policy or other changes affecting the country in question. A classic case is China: after a long period of stagnation and reversals, economic growth in China started to take off in the 1970s, and at the end of that decade it rose to the double-digit rate is has held ever since. What triggered this, and what stops other economies following a similar path? A simple adaptation of our model can help to explain the phenomenon. We assume that all technologies—including very advanced ones—are available at all times, but the more advanced technologies are more costly.

First, note that a very high rate of growth—such as that observed in China—demands a very high rate of investment. Recall the model economy 3.1. To achieve rapid growth in the model economy, firms must either renew their capital at very short intervals, or they must take very large steps forward each time they renew their capital. In either case, a high overall rate of investment is required (relative to the existing level of GDP). Returning to China, there the investment rate is around 50 percent. In a closed economy with a 50 percent investment rate, 50 percent of factors of production (workers and machines) must be dedicated to producing investment goods such as new capital, and 50 percent must be dedicated to producing consumption goods. But note that in an open economy—trading with the rest of the world—there are other alternatives. One alternative in such an economy would be to produce only consumption goods, export half of these goods, and use the proceeds to import capital goods.

Model economy 3.2. *Assume an economy similar to 3.1. However, this economy is open, and far behind the technology frontier. It produces only consumption goods, which it either consumes itself or sells abroad (exports). Using the proceeds of its exports it buys capital goods, i.e. new machines. New machines are available in a series of different qualities, where higher quality machines yield higher GDP per capita; but higher quality machines also cost more.*

Specifically, each rung of the ladder gives 20 percent higher GDP than the previous rung, and the machines to achieve it cost 20 percent more than the machines for the previous rung. The

Rung	GDP	Investment	$I/Y \times 100$	Growth
5	24883	4147	41	9.5
4	20736	3456	35	7.6
3	17280	2880	29	5.6
2	14400	2400	24	3.7
1	12000	2000	20	1.8
0	10000	1667	17	0.0

TABLE 3.1. Rungs of the quality ladder. Units of GDP and investment are USD per capita per period, and units of growth are percent per year.

current GDP per capita is 10000 *USD per ten-year period, and for* 2000 *USD per capita the country can upgrade—in the next 10-year period—to machines that yield* 12000 *USD per capita per period.*

Now assume that the country invests in new machines that are five rungs up the ladder compared to its current machines, for each member of the population. What is investment as a proportion of GDP? What is consumption? And what is the growth rate in GDP per capita, per year?

To solve this model, we calculate GDP on each rung, and the investment required to attain each rung, using the information given, and denoting the initial rung as rung 0. We then calculate the corresponding investment rates in percent, and growth rates in percent per year. This information is shown in Table 3.1.

Reading off from the table, to jump up five rungs the country must invest 4977 USD per capita in new machines, whereas GDP over the 10-year period is 10000 USD per capita. So the investment rate is 50 percent, implying that 50 percent of the country's production is exported in order to pay for the imported machinery. The country then achieves a jump in GDP from one period to the next of 149 percent, corresponding to 9.5 percent per year.

In a closed economy, high I implies high S, i.e. a high *savings rate*. But in an open economy (which trades with other economies) this is no longer the case, and the key to kick-starting growth in a stagnant economy is *not* a very high domestic savings rate: the key is instead to generate confidence that investment put into the economy will give a healthy return to the investor. Without such confidence, high domestic savings will translate into high investment in foreign countries, i.e. capital will leave the domestic economy.

What is the key to 'investor confidence'? (Note that we should not identify investors with *foreign* investors, they may equally well be domestic nationals.) After decades of searching for the key, economists have concluded that there is no one key. It is not the existence of an educated workforce, it is not the presence of infrastructure, it is not a lack of rules and regulations, it is not law and order, it is not the presence of a stable political system. Instead, the key is all of these things, and more. In short, the *institutions* of the country should be growth-friendly, investor-friendly.

Regarding the specifically economic environment, key factors include stability of the economic system, and openness. Stability of the system is crucial because of the (often long) delay between investment and returns; if investors judge that there is a high risk of economic crisis or upheaval during the lifetime of the investment this will drastically reduce the investor's willingness to invest. Openness (in particular openness for trade) is crucial because it allows new ideas to come into the economy, and it allows investors greater opportunities to get returns from their investments. Regarding openness, note that countries typically trade most with their near neighbours, and it is therefore a big advantage for domestic growth if neighbouring economies are rich, or growing rapidly.

Regarding the more general institutional environment, a lack of bureaucracy, a lack of corruption, an educated population, and a well functioning civil society would all seem to be advantageous when it comes to attracting investors (whether domestic or foreign). Again, the key is to understand that investment implies costs today and hoped-for benefits tomorrow. Factors—such as corruption—which cause investors to lose confidence in actually receiving such benefits are likely to be very negative for growth prospects. Fear of revolution or civil war would have an even more powerful negative effect.

Regarding openness, we discussed openness to trade, but another important type of openness is openness to *change*. Since growth is based on the replacement of old, less productive technologies by new ones, it is a process that is at once both creative and destructive: hence the term *creative*

destruction coined by Joseph Schumpeter.[9] When the new ideas succeed, their instigators are likely to grow in wealth and power, implying that the previously wealthy and powerful lose status. Thus it is not obvious that the powerful individuals in a given economy want economic growth at all. According to Daron Acemoglu — a highly influential macroeconomist who has written on growth and institutions among many other questions — the likelihood that leaders are anti-growth increases if they are far removed from the ordinary population.[10] One thing bringing leaders closer to the people is a functioning democracy.

3.7. Modelling endogenous growth

3.7.1. Ideas and growth. We now consider how to endogenize technological progress, or more specifically (in the context of our vintage model) how to endogenize the size of the increases in capital productivity per period. We can think of new technologies as designs or blueprints. Until they are put to use, they are not embodied in physical capital; they are intangible. From now on we simply refer to new technologies as new *ideas*. The special properties of ideas have profound consequences for the growth process. In economic jargon we say that ideas are *non-rival*. Moreover, they can be *non-excludable*.

Recall the growth model of Chapter 2, in which there was no capital, and growth was driven by increases in A_L, which were in turn linked to the size of the population. The relationship was described by equation 2.4,

$$A_{Lt+1} = A_{Lt}[1 - \delta + \zeta(\Omega L_t)^\phi],$$

in which δ is knowledge depreciation (which is presumably very slow), Ω is the proportion of the population (or the proportion of total labour) which is devoted to the generation of new ideas, ζ is a productivity parameter equal to the progress resulting from the idea coming from 1 person's full time idea-generation, and ϕ is a parameter $\in (0, 1)$; the lower is ϕ, the greater is the overlap between the ideas, and in the limit of $\phi = 0$ everyone has the same idea. We thus assumed that individuals exogenously come up with new ideas about how to organize the world, once per period.

In order to build a model in which Ω is endogenous we need to find a way to give agents an incentive to perform research or generate new knowledge or ideas. It is not obvious why agents should have such an incentive because ideas may be *non-excludable*, meaning that it may not be possible to prevent other people using one's 'own' idea, without payment. Ideas also possess a very special property — when compared to other goods such as pizza or cars — which is that they are *non-rival in consumption*, meaning that one agent's use of the good does not hinder another person from using it, either consecutively or even simultaneously. This is clearly not true of cars, and even less so of pizza. The non-rivalry of ideas makes good ideas immensely valuable to society, since a single idea can be used by any number of people, any number of times. However, the non-excludability means that an individual who comes up with an idea may have a hard time capturing any of that value.

Let us return to the context of pre-industrial human societies, starting with hunter–gatherers 50 000 years ago. Why was technological progress so slow in these societies? Firstly — as we argued in Chapter 2 — because the total global population was low, the total number of people able to invest in new ideas was also low. Hence the rate of production of ideas was low; in the vast majority of 10-year periods, there were no significant innovations whatsoever, and worn-out spears and arrow-heads were replaced by new ones of identical design. The second reason for slow technological progress in pre-industrial societies is that the incentives to innovate were very weak. If we consider a hunter–gatherer band of 100 people, what are the incentives for that band to invest in research about better ways to hunt, harvest, build, etc.? There is of course some incentive to do this: benefits of any technological advances will be felt by the group. However, the benefits to the group will be dwarfed by the potential benefits if the innovation spreads across the globe. This shows that the incentives to perform research are immeasurably greater if the researcher is able to capture a significant proportion of global benefits.[11]

Return to the Solow model and industrialized economies. Assume that someone founds a research firm and (after much work) finds a better way to make hammers such that A increases. If there are patent laws (excludability), the owner of the new knowledge will have an advantage in production, and (because of non-rivalry) can take over production in the entire economy and become fabulously rich. On the other hand, if there are no patent laws then the discoverer of the

---

[9]See for instance Schumpeter (1942), *Capitalism, Socialism and Democracy*.

[10]See for instance Acemoglu (2009), *Introduction to modern economic growth*.

[11]This suggests that in pre-market economies investment in technologies related to war should have been relatively large, since only such investments could lead the (small) group to capture benefits from outside the group.

new knowledge will gain no benefit from it whatsoever, since all firms will use the knowledge and the owner's firm will still make zero profits. In total the owner's firm will make a loss, because the firm will be unable to cover the costs of the research. Neither of these extreme results seems to make much sense. What is the way forward?

In general, we can say that a more sophisticated model is clearly required, a model in which there is a variety of products with their own production functions, and where researchers can protect their intellectual property in some way, for instance through taking out patents. Agents then have an incentive to perform research (the incentive being the expected flow of profit in the future). Furthermore, the resultant discoveries may form the basis for further advances by later researchers, thus potentially leading to a sustainable growth process.

In the literature there are two main traditions concerning how to model growth. In both research traditions an *intermediate sector* is introduced, where there are patents and market power, and the intermediate goods produced in this sector are inputs—together with labour—into the final-good sector where there is perfect competition. These traditions are the *Romer* tradition (see for instance Romer (1990)), and the *Aghion–Howitt* or *Schumpeterian* tradition (see for instance Aghion and Howitt (1992)). However, in this book we use a different basic model in which there are no intermediate goods, just a range of final goods which are imperfect substitutes for one another. This model is simpler, and in some respects it has more explanatory power as well. For the purposes of this chapter it is the best choice, partly due to its simplicity, and partly due to the fact that we later build on it in our model of directed technological change (Chapter 5).

3.7.2. The model environment. Now we turn to the model, which is illustrated in Figure 3.7. In the model there is a unit continuum of infinitely lived households, and the representative household has L_t members. The representative household maximizes its utility U, which is given by

$$U = \sum_{t=0}^{\infty} \beta^t C_t.$$

So time is discrete, and utility is linear in consumption C, discounted by a factor β per period.

The mass of products that can be made is variable, with each product made by a single firm. Denote the mass of products made by N, and assume for now that $N = 1$. Index the products by i, and focus on product i, quantity y_i, made exclusively by firm i. The production function for product i is as follows:

$$y_i = A_{Li}l_{Yi}. \tag{3.5}$$

Here l_{Yi} is the number of workers hired by the firm at price w_L. The price of labour is determined endogenously (within the economy). The firms are symmetric in the sense that each employs the same quantity of labour, and invests the same in research. We have thus dropped capital investment from our vintage capital model! This is done to simplify the equations, it actually makes very little difference to the results or the intuition.

Aggregate labour L is then the integral across the entire mass of firms of firm labour l_{Yi}:

$$L = \int_0^1 l_{Yi}\mathrm{d}i = l.$$

where l is labour demanded by the representative firm. And aggregate production Y is a function of the individual production levels y_i which are all equal to y in symmetric equilibrium:

$$Y_t = \left[\int_0^1 y_{it}^\eta \mathrm{d}i \right]^{1/\eta} \tag{3.6}$$

$$= y_t. \tag{3.7}$$

The parameter η is between 0 and 1, implying that the goods are not perfect substitutes, and each producer has a degree of market power. The price of the aggregate good is normalized to 1.[12]

In period t a firm which plans to make good i must first decide how much to invest in development of the production technology which in turn determines A_{Lit}, labour productivity. The production function for A_{Lit} is as follows:

$$A_{Lit} = A_{Lt-1}[1 - \delta + (\zeta l_{Ait})^\phi], \tag{3.8}$$

where δ is depreciation, A_{Lt-1} is general knowledge in the previous period, ζ is the productivity parameter, l_{Ait} is the quantity of research labour hired by the firm, and ϕ is a positive parameter. So the more research labour hired by the firm, the more productive it will be. The price of research

---

[12]For more on aggregation in economies with a continuum or 'mass' of firms, see Appendix A.3.

labour is simply the wage w_{Lt}.[13] Note that firm-level knowledge depreciates completely from one period to the next; firms start from scratch each period. Periods must thus be fairly long, at least a few years. Here we set the period length to 10 years.

General knowledge A_L is made up of the knowledge of each of the individual firms. We choose the simplest possible specification,

$$A_{Lt} = \int_0^1 A_{Lit}\mathrm{d}i,$$

so in symmetric equilibrium (when firm knowledge levels are all the same) we have

$$A_{Lt} = A_{Lit}.$$

We assume that the representative firm employs a constant number l_{Ai} of researchers each period, hence the aggregate number of researchers $L_A = l_{Ai}$ and is also constant. Now look at aggregate growth. From equation 3.5 and given $A_{Li} = A_L$ we have

$$Y = A_L L_Y.$$

But how does A_L change over time? Define $A_{Lt+1}/A_{Lt} = 1 + \theta$. From the equation for knowledge growth (3.8) we have

$$1 + \theta = 1 - \delta + (\zeta L_A)^\phi.$$
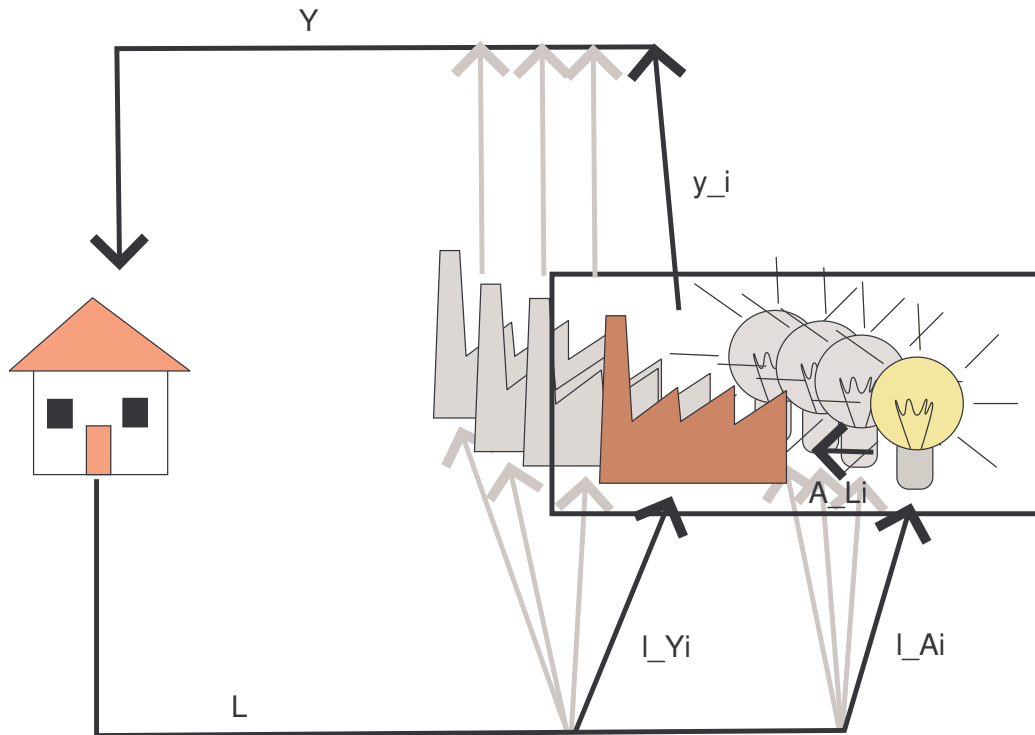
FIGURE 3.7. A schematic diagram showing the economy with endogenous growth.

3.7.3. The solution. Now we set up the optimization problem facing a firm which has decided to produce good i in period t. (Note that in Nash equilibrium no two firms will produce the same product in the same period.) The Lagrangian for the problem is as follows:

$$\mathcal{L}_{it} = p_{it}y_i - w_{Lt}l_{Yit} - w_{Lt}l_{Ait} - \lambda_{it}\left[A_{Lit} - A_{Lt-1}[1 - \delta + (\zeta l_{Ait-1})^\phi]\right].$$

So $\mathcal{L}$ is equal to revenue minus costs, minus the Lagrangian multiplier λ times the restriction on labour productivity.

To solve the problem we need to know how the price of good i, p_i, is determined. Good i is produced by a single monopolistic firm with productivity A_{Li}. To find the price of this good think

___

[13]Note that it would be more realistic—but a little more complicated—to model research occurring in period $t-1$ rather than period t.

of an entrepreneur buying goods y_i to make (and sell) the aggregate good Y at price 1 (normalized). The entrepreneur's problem is to maximize profits:

$$\pi = \left[\int_0^1 y_{it}^\eta \mathrm{d}i \right]^{1/\eta} - \int_0^1 p_i y_i \mathrm{d}i.$$

Take the first-order condition (FOC) in y_i to obtain

$$p_i = (y/y_i)^{1-\eta}$$

hence
$$\mathrm{MR} = \frac{\partial p_i}{\partial y_i} y_i + p_i = \eta p_i,$$

where MR is marginal revenue. So when firm i raises its production of good i by one unit, its revenue goes up not by p_i (the price of that unit at the start) but by ηp_i where η is less than 1. Hence firms have *market power* and can raise the price they receive for their goods by restricting production. Finally, note that in symmetric equilibrium the firms all produce equal quantities so $y_i = y$ and $p_i = 1$, hence MR $= \eta$.

Now take the FOC on the Lagrangian in labour (using the above result) to yield

$$w_L l_{Yi} = \mathrm{MR} \cdot l_{Yi} \frac{\partial y_i}{\partial l_{yi}}$$

Since MR$= \eta p_i$ and $y_i = A_{Li} l_{yi}$ we have $w_L l_{Yi} = \eta p_i y_i$. And in symmetric equilibrium $p_i = 1$ (all the goods have the same price), hence

$$w_L L_Y = \eta Y. \tag{3.9}$$

This tells us that payments to production workers use up a proportion η of total revenue (which is Y).

Now we want to know more about η. Assume that the mass of firms (1) is an endogenous outcome, as is η: when there are few firms and few goods, the firms have a lot of market power (η is low) and they make positive profits. This encourages entry, competition stiffens, and η declines. In equilibrium there is a unit mass of firms and η is such that profits are zero. So, what is the value of η?

Take the FOC in labour productivity A_{Li} to obtain (in symmetric equilibrium)

$$\lambda A_L = \eta Y.$$

And take the FOC in research labour to obtain,

$$w_{Lt} L_{Ait} = \phi \lambda_{it} [A_{Lit} - A_{Lt-1}(1-\delta)]$$

then use the assumption of symmetric equilibrium, and then the definition of θ above to show that

$$w_{Lt} L_{At} = \phi \lambda_t [A_{Lt} - A_{Lt-1}(1-\delta)]$$
$$= \phi \lambda_t A_{Lt} (\theta + \delta)/(1+\theta).$$

Finally substitute in the FOC in labour productivity to obtain

$$w_L L_A = \eta \phi Y (\theta + \delta)/(1+\theta). \tag{3.10}$$

This equation gives us the research costs of the representative firm.

The expressions for production-labour and research-labour costs (3.9 and 3.10) can be used to derive an expression for the profits made by firm i, i.e. revenue minus costs. Do this, and again switch to the aggregate level assuming symmetric equilibrium:

$$\pi_i = p_{it} y_{it} - (w_{Lt} l_{Yit} + w_{Lt} l_{Ait})$$
$$= \left[1 - \eta[1 + \phi(\theta + \delta)/(1+\theta)] \right] Y.$$

This tells us that if $1 - \eta[1 + \phi(\theta + \delta)/(1+\theta)] > 0$ then firm i will make positive profits, hence there should be an incentive for further firms to enter, raising the total mass of firms above 1. On the other hand, if $1 - \eta[1 + \phi(\theta + \delta)/(1+\theta)] < 0$ then firm profits are negative, and firms should exit thus pushing down the total mass of firms. And given that a unit mass of firms is an equilibrium, we must have zero profits and hence

$$1 = \eta[1 + \phi(\theta + \delta)/(1+\theta)].$$

Rearrange to obtain

$$\eta = \frac{1}{1 + \phi(\delta + \theta)/(1+\theta)},$$

Now we know η, we can easily find an expression for research labour L_A using equation 3.9. Since profits are zero then the sum of payments to production workers and payments to research

workers must be exactly Y. Since payments to production workers are ηY this implies that payments to research workers are $(1 - \eta)Y$, so

$$w_L L_A = (1 - \eta)Y,$$
$$= (1 - \eta)w_L L_Y / \eta,$$

and
$$L_A / L_Y = (1 - \eta)/\eta$$
$$= \phi(\delta + \theta)/(1 + \theta).$$

If we set $\phi = 0.2$, the period length to 10 years, $\delta = 0.05$, and $\theta = 0.2$ (so approximately 2 percent per year) we obtain $L_A / L_Y = 0.04$. So four percent of resources are devoted to research in this case.

3.7.4. Conclusions about endogenous growth. We have shown how we can take an 'exogenous growth model' in which workers devote a fixed proportion of their time to research for no obvious reason, and turn it into an endogenous growth model in which time spent on research is the result of competing firms solving optimization problems. For our purposes, the key use of the endogenous model will be in later chapters when we wish to analyse how firms choose between investment in alternative types of technology, such as 'clean' and 'dirty'. However, growth models such as the above raise a number of other issues relating to policy for overall growth and welfare.

The key issue is that regulatory systems giving intellectual property rights to the discoverers of ideas tend to have two drawbacks. Firstly, the rights-holders use their position to restrict production and drive up prices. This gives them profits, but it also causes deadweight losses for society. Secondly, despite this discoverers typically get lower net benefits from their discoveries than the net benefit to society, which implies that incentives to perform research are 'too weak', and hence too little research will be performed compared to what would be socially optimal.

In order to boost research further, governments typically sponsor research through research subsidies, and also through dedicated organizations supported by the government, such as universities. However, such subsidies are also fraught with difficulties. The problem here is that it's very hard to make sure that the research paid for by the government is actually performed. Furthermore, it is hard for the government to know what would have been done in the absence of a subsidy: maybe the firm just pockets the subsidy and carries on with research it would have performed anyway. These difficulties are connected to the nature of ideas (non-rivalry) and the inherent uncertainty of the research process. If the government pays a firm to build a bridge, then if the bridge is not built the government can sue. But if the government pays a firm to perform some research, and the firm reports back at the end that the research did not lead to any useful new ideas, what can the government do? To provide state aid to university research is of course a great way to get out socially useful research for a small cost while contributing to students' education. Right?

3.A. Appendix: The Ramsey model with endogenous capital investment

The Solow model is good for explaining the medium-run evolution of the capital stock, and phenomena which are tightly linked to that. For instance, if a country's capital stock is largely destroyed due to a calamity such as war or an earthquake, the Solow model predicts that the capital stock and GDP will recover rather rapidly, and the evidence bears this out (Figure 3.3).

A better model for doing the above—and many other things—is the Ramsey model, which takes the Solow model and endogenizes the saving rate s; that is, in the Ramsey model households choose the balance between consumption and saving in order to maximize their utility, rather than simply by following a rule of thumb.[14] The result is that (for instance) a catastrophic loss of capital may lead to a rise in the saving rate because of the greater effect of saving on growth, but a countervailing decline in the saving rate due to the desire to smooth consumption over time. What actually happens depends on the balance between these effects.

We do not go into the solution of the Ramsey model, since this is beyond the scope of this book. However, you should know that the key to endogenizing the savings rate is the household utility function. Recall from the Introduction that we typically write a household utility maximization problem as follows:

$$\max \sum_h U^h,$$

where
$$U^h = \sum_t u^h(c_t)\beta^t.$$

Here U^h is the net present value of the utility of household h, $u^h(c_t)$ is the instantaneous flow of utility at time t, and β is the discount factor (which is less than 1). But what is the form of the instantaneous utility function $u^h(c_t)$? The standard utility function used in the Ramsey model is the CIES, as follows:

$$u = \frac{c^{1-\sigma} - 1}{1 - \sigma}.$$

As long as $\sigma > 0$ this implies that marginal utility is declining in consumption, which means that households have a preference for consumption smoothing. Furthermore, since $\beta < 1$ they also have a preference for consuming now rather than later. In an economy in which technological progress drives consumption growth, both of these factors militate against investment and towards consumption today.

Note that in an open economy we should also allow for international capital flows, which will unambiguously favour rapid recovery from a negative shock to capital.

3.B. Appendix: A model without capital

The above analysis shows that the analysis of the capital stock and investment is principally relevant for explaining medium-run fluctuations in GDP, rather than long-run trends. Since we are interested in the long run in this book, this raises the question of whether we could simplify matters by ignoring capital completely, effectively assuming that it grows in line with effective labour.

If we write the production function as

$$Y = A_L L [K/(A_L L)]^\alpha,$$

and note that in balanced growth K grows at the same rate as $A_L L$, so $K/(A_L L)$ is constant, then it should be clear that in long-run analysis we lose little by simplifying the production function to

$$Y = A_L L.$$

The same conclusion applies when we include resources in the production function. Then we have (for instance)

$$Y = A_L L)^{1-\alpha-\beta} K^\beta R^\alpha$$
$$= (A_L L)^{1-\alpha} [K/(A_L L)]^\beta R^\alpha.$$

The above argument breaks down if (for instance) changes in resource flows R are so great that the overall growth rate changes significantly; growth is not simply driven by increases in A_L, but is affected in a major way by changes in R. In that case this will affect rates of capital investment and the interest rate in the economy, and we should switch to an analysis including capital and based on the Ramsey model with endogenous investment and an endogenous interest rate.

---

[14]Also called the Ramsey–Cass–Koopmans model. See Ramsey (1928), Cass (1965), Koopmans (1963).

The DHSS model

We now return to the neoclassical growth model of Chapter 3, and add a nonrenewable resource to the production function. We thus have the DHSS model.[1] We start with two simplified versions of the model: in the first of these the resource input is fixed (land), and in the second the flow of resources R comes from a non-renewable stock which is very large and homogeneous, in which case increasing demand—driven by the increasing productivity of labour—drives increases in resource extraction while not affecting the resource price. Next we consider the standard case in which a homogeneous resource is available in a fixed total quantity, and costs nothing to extract. In this case, if the resource is essential to production then—if production is to be sustainable—it must be extracted and consumed at a decreasing rate over time, at least in the long run. After noting the problems with all of the above cases, we extend the model to allow for an inhomogeneous stock of the resource. This greatly increases the realism of the model with regard to resource supply, and throws the spotlight onto resource demand, to which we turn in subsequent chapters.

4.1. The DHSS model: two simplified versions

4.1.1. DHSS 1: Land (and 'flow renewables').
We begin this chapter with a neoclassical growth model with zero population growth, and add the need for land in the production function. Instead of land we can also think of a renewable resource which flows at an exogenous rate and where there is no stock; consider for instance sunshine, wind, etc.: the exogenous, constant flow of the renewable resource is equivalent to the exogenous, constant quantity of land available. The quantity of land is simply fixed at R, and there is exogenous labour-augmenting technological progress:

$$Y = (A_L L)^{1-\alpha-\beta} K^\alpha R^\beta$$

$$\dot{A}_L/A_L = g_{A_L}$$

$$\dot{K} = sY - \delta K.$$

To characterize the development of the economy, assume that there exists a balanced growth path on which all variables grow at constant rates. This implies that $\dot{K}/K$ is constant, which implies (from the capital-accumulation equation) that $sY/K - \delta$ is constant, implying in turn that the ratio of production Y to capital K is constant. This implies that $\dot{Y}/Y = \dot{K}/K$, and hence (differentiating the production function with respect to time and substituting)

$$\dot{Y}/Y = (1-\alpha-\beta)g_{A_L} + \alpha\dot{Y}/Y.$$

Rearrange to obtain

$$g_Y = g_y = \frac{1-\alpha-\beta}{1-\alpha} g_{A_L}. \tag{4.1}$$

This confirms the reasonableness of the original assumption, i.e. that there exists a balanced growth path on which all variables grow at constant rates. Equation (4.1) shows us that given labour-augmenting technological progress, when $\beta > 0$ (i.e. land is a non-negligible factor of production) the growth rate of production is slowed. The underlying reason why land in the production function slows growth whereas capital does not is that land is non-reproducible (by contrast to capital).

The price of land w_R can be obtained by considering the representative producer's profit-maximization problem,

$$\max \pi = p_Y(A_L L)^{1-\alpha-\beta} K^\alpha R^\beta - (w_L L + w_K K + w_R R).$$

Normalize the price of the final good, p_Y, to one. Then take the first-order condition in R to show that $w_R = \beta Y/R$, and hence that

$$\dot{w}_R/w_R = \dot{Y}/Y.$$

---

[1]DHSS stands for Dasgupta–Heal–Solow–Stiglitz. The model was developed in the 1970s, the original papers being Dasgupta and Heal (1974), Solow (1974), and Stiglitz (1974).

So the price of land tracks growth in GDP.

Finally, if we set $\beta = 0$ (implying that land is not important) then we obtain $g_Y = g_y = g_{A_L}$, just as we did in the basic neoclassical model with the production function $Y = (A_L L)^{1-\alpha} K^\alpha$. Note that we also obtained the same result in the even simpler model in which $Y = A_L L$.

How relevant is this model to reality? Clearly, the total supply of land is fixed. The only testable prediction of the model (apart from results that more-or-less follow by assumption) is that the price of land should grow at the overall growth rate. Long-run data on land prices is hard to come by, but the prediction seems to be broadly correct.

4.1.2. DHSS 2: An abundant resource, costly to extract. Now we consider a mineral resource input instead of land. We assume that there is a very large homogeneous stock of this mineral (consider for instance iron ore). The input is costly to extract, since extraction requires the use of labour, capital, and the resource input, inputs which could otherwise have been used to produce the final good. However, in the model we assume that all labour, capital, and resources are used in the final-good sector, and that some of the goods produced in this sector are then sent to the extraction sector and used (without the need for any other inputs) in extraction.[2] There is exogenous labour-augmenting technological progress. Denoting total final-good production by Y and denoting final goods devoted to extraction as X we have

$$Y = (A_L L)^{1-\alpha-\beta} K^\alpha R^\beta;$$
$$\dot{A}_L / A_L = g_{A_L};$$
$$\dot{K} = s(Y - X) - \delta K;$$
$$R = \phi X.$$

Here ϕ is a parameter determining the relative productivity of the inputs in extraction compared to final-good production. The economy is illustrated in Figure 4.1.

To solve the model, note first that total extraction costs are simply X, since the price of the final good is normalized to 1. And since we assume perfect markets, the resource price w_R is equal to unit extraction costs X/R, hence

$$w_R = 1/\phi.$$

Now, as previously, assume a b.g.p. and then characterize it. On a b.g.p. $\dot{K}/K$ is constant, which implies (from the capital-accumulation equation) that $s(Y-X)/K - \delta$ is constant, implying in turn that Y, X, and K all grow at equal rates, and hence (from the extraction function) R also grows at this rate. Differentiate the production function with respect to time, and substitute for $\dot{K}/K$ and $\dot{R}/R$ to yield

$$\dot{Y}/Y = (1 - \alpha - \beta)g_{A_L} + \alpha \dot{Y}/Y + \beta \dot{Y}/Y.$$

Rearrange to yield

$$\dot{Y}/Y = g_{A_L}.$$

So the need for the resource imposes no brake at all on the growth rate of production, since resource use tracks production (by contrast to land, which is fixed).

Note that the predictions of this model do a pretty good job of matching the data in Figure 4.4, especially the left panel (metals): the resource price is constant, and resource extraction tracks overall growth. The broad outline of the story told by the model is simple and almost certainly correct: on the demand side, technological progress drives increasing demand for the resource; on the supply side, it implies that when a fixed proportion of labour and capital is devoted to resource extraction, the flow of extraction increases exponentially while price remains constant.

The success of this simple model in accounting for historical data is striking, but so is its unsuitability for predicting the future. The model predicts that resource extraction will continue to increase exponentially, indefinitely. As mentioned in the introduction, the current physical extraction rate of minerals is of the order of 10^{10} tonnes per year globally. Assume that this rate continues to grow at the same rate as it has done over the past 100 years, i.e. approximately 3 percent per year. Then extraction would be multiplied by a factor 20 each century, and in 700 years we would be mining and using minerals roughly equal to the entire earth's crust *every year*, based on a figure of 2×10^{19} tonnes for the Earth's crust. To predict the future we clearly need a model with limited, inhomogeneous resource stocks. The rest of this part of the book is devoted to such models.

---

[2]This way of formulating the model is simpler, and equivalent as long as there is sufficient symmetry between the final-good production function and the resource extraction function.
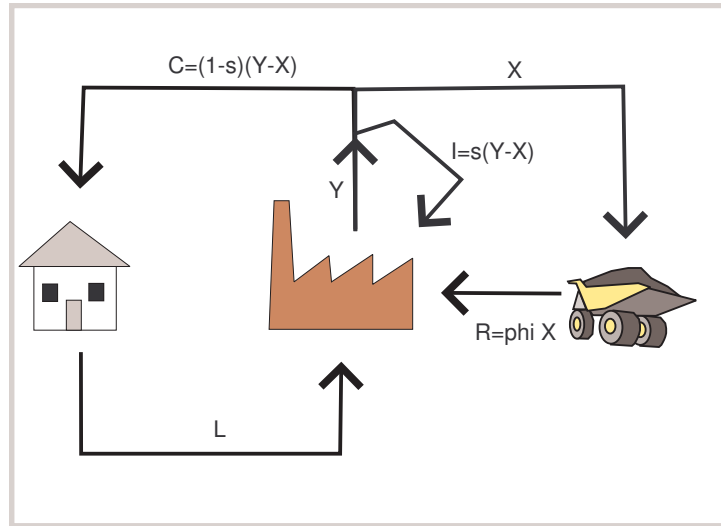
FIGURE 4.1. A schematic diagram showing the economy with an unlimited resource, costly to extract.

4.2. The standard DHSS/Hotelling model: A limited resource, costless to extract

In the previous section we had a resource that was costly to extract but available in unlimited quantity. Now we go the other way and assume a resource which is free to extract but available in a known, finite quantity. This takes us into the heart of the literature from the 20th century, with the Hotelling rule from 1931 and the DHSS model from 1974.

4.2.1. General set-up. Compared to the models above we now need to add an equation limiting total resource use to no more than S, the exogenous level of the total stock. On the other hand, there are no extraction costs. However, we add a fixed rate of population growth n, and we generalize from the above models (and from Solow) by allowing the saving rate to vary, so instead of fixing this rate at s we simply state that investment is the difference between production Y and consumption C.

$$Y = K^\alpha R^\beta (A_L L)^{1-\alpha-\beta}$$
$$\dot{A}_L/A_L = g_{A_L}$$
$$\dot{L}/L = n$$
$$\dot{K} = Y - C - \delta K$$
$$S \geq \int_0^\infty R_t \mathrm{d}t.$$

Finally, utility is some function of the consumption path over a specified time period. Typically we have $U = \int_0^\infty u(c(t))e^{-\rho t}t$ where ρ is the pure rate of time preference, but (for instance) Solow uses a maximin utility function.

Note that we can keep producing final goods Y even as R approaches zero; so if the final good is hammers, and the resource input is iron, we can make an unlimited number of hammers from a fixed stock of iron (without recycling). The model is thus set up from the start to favour feasibility of sustainable production (if $g_{A_L} = 0$) and sustainable growth (if $g_{A_L} > 0$). Note that with a different choice of production function—such as Leontief—we can rule out the above property directly; given Leontief (and no resource-augmenting technological progress), you need x grams of iron per hammer, full stop.

In the original models—see for instance Dasgupta and Heal (1974) and Solow (1974)—it is common to set depreciation $\delta = 0$, and the rate of technological progress $g_{A_L} = 0$. We can then study whether capital accumulation can—in the very long run—compensate for necessary reductions in the flow of resources to allow constant production. However, in this book we are interested in explaining and predicting phenomena in real economies, so we assume that g_{A_L} is strictly positive. Regarding δ we set it to zero initially, but only for convenience; later on we show that the results are essentially unchanged with positive δ.

4.2.2. Exogenous technological progress. What if we have the general set-up above, with $g_{A_L} > 0$, a fixed saving rate s, and $n = \delta = 0$? That is,

$$Y = K^\alpha R^\beta (A_L L)^{1-\alpha-\beta},$$

$$\dot{A}_L/A_L = g_{A_L},$$

$$\dot{K} = sY,$$

$$S \geq \int_0^\infty R_t \mathrm{d}t.$$

Under these circumstances it is easy to show that we can maintain constant consumption if we set its initial level low enough.

(1) First assume balanced growth, i.e. all growth rates are constant.
(2) Second, note that for any initial level of resource use R_0 we can always find a rate of decay in resource consumption θ such that the resource is asymptotically exhausted:

$$S = \int_0^\infty R_0 e^{-\theta t} \mathrm{d}t = [-(R_0/\theta)e^{-\theta t}]_0^\infty = R_0/\theta.$$

The lower is R_0/S, the lower is the rate of decline θ.
(3) Third, note that if Y grows at some constant rate g_Y, while a constant proportion s is invested, and $\dot{K}/K$ is constant, then Y/K must also be constant. Thus Y and K must grow at the same rate, $g_Y = g_K$. This follows since $\dot{K}/K = sY/K$.
(4) Fourth, differentiate the production function w.r.t. time and use $g_Y = g_K$ to find an expression for g_Y in balanced growth:

$$g_Y = \left(1 - \frac{\beta}{1-\alpha}\right) g_{A_L} - \frac{\beta\theta}{1-\alpha}. \tag{4.2}$$

So by choosing initial resource consumption, and hence θ, we can choose the long-run growth rate g_Y up to a maximum of $g_{A_L}[1 - \beta/(1-\alpha)]$.

Recall that when the flow of resource input was fixed (the case with land), we had

$$g_Y = \left(1 - \frac{\beta}{1-\alpha}\right) g_{A_L}.$$

This is identical to the expression above when θ approaches zero, as we would expect, since when $\theta = 0$ the resource flow is constant. Put differently, the need for resource flows to decline puts a penalty on to the growth rate.

4.2.3. Hotelling. What will the solution be in a market economy? To answer this question we need to know the resource price. The resource is free to extract, but because it is available in finite quantity it is *scarce* and hence it will have a non-zero price on the market due to scarcity rent (also known as Hotelling rent, resource rent, etc.).

Imagine you own a stock of some resource, which you keep in a large pressurized underground tank. To extract the resource, you simply turn the tap. You have no alternative uses for the tank once the resource is exhausted, and there are no environmental considerations.

In order to decide when to extract the resource you consider the price path. Is the resource price w_R increasing over time? If so, then all else equal you will keep the resource and sell it later. On the other hand, if you extract and sell the resource then you can put the money in the bank, and it will grow at rate w_m, where w_m is the interest rate (we express it this way to emphasize that the interest rate is the rental price of money). So your rule is to extract if the resource price is rising at a rate slower than w_m, and not extract if the price is rising at a rate faster than w_m. If $\dot{w}_R/w_R = w_m$ then you are indifferent.

Now imagine you are one of a vast number of such resource holders, so the market for selling the resource is competitive (as are all other markets in the economy). How do the others plan their extraction? Presumably, they think the same way as you do. Assume that the current price is high, and thus all resource owners expect slowly-rising prices in the future; then everyone wants to sell. Now assume instead that the current price is low, and thus all resource owners expect rapidly-rising prices; then no-one wants to sell. In the former case—when everyone wants to sell—the price must drop in a discrete step down until it reaches a point at which it is no longer the case that everyone wants to sell. In the latter case it must rise in a discrete step up until it reaches a point at which it is no longer the case that no-one wants to sell. The equilibrium price is the same in both cases; it is the point from which the price is expected to rise at exactly the rate of interest. At this point resource owners are indifferent between holding on to their resources and selling their resources: we have an equilibrium.

We have thus intuited the equilibrium price path for the resource:

$$\frac{\dot{w}_R}{w_R} = w_m,$$

where w_R is the price of the resource and w_m is the price of (borrowing) money, i.e. the interest rate.


FIGURE 4.2. A schematic diagram showing the economy with a limited resource, free to extract.

Now return to the overall model, which is illustrated in Figure 4.2. Given perfect markets we know that the Hotelling rule must be obeyed, i.e.

$$\dot{w}_R/w_R = w_m.$$

But the resource price is just the marginal product of the resource for the representative producer, i.e.

$$w_R = \beta Y/R.$$

Differentiate this expression and use the Hotelling result, and our result regarding the b.g.p., to show that on a b.g.p. we have

$$w_m = \dot{w}_R/w_R = (g_{A_L} + \theta)[1 - \beta/(1-\alpha)].$$

If the interest rate w_m is simply equal to the pure rate of time preference ρ—based on the simplest possible model of preferences over time, $U = \int_0^\infty c_t e^{-\rho t} dt$—then we have

$$\theta = \frac{\rho}{1 - \beta/(1-\alpha)} - g_{A_L},$$

so the optimal rate of decline of resource use is increasing in the degree of impatience of agents (ρ) and in the weight of resources in the production function (β), but decreasing in the rate of technological progress g_{A_L}. We can now put this back into the equation for g_Y to yield

$$g_Y = g_{A_L} - \rho\frac{\beta/(1-\alpha)}{1 - \beta/(1-\alpha)},$$

and hence $$\theta = \rho - g_Y.$$

So the growth penalty due to the need to gradually cut resource consumption depends on the rate of time preference ρ and the weight of the resource in the production function compared to the weight of labour. Note that for a meaningful solution we require that $\rho > g_Y$. (Among other things, the rate of decline of resource use must be strictly positive.)

As an exercise, you should redo the above assuming capital depreciation at rate δ. You should find that it makes no difference, since when capital depreciates we have $\dot{K}/K = sY/K - \delta$. On a b.g.p. (with $\dot{K}/K$ constant) this again implies that Y/K is constant, hence $\dot{Y}/Y = \dot{K}/K$, just as when δ was zero. Since capital still grows at the same rate as overall production in this economy, we obtain the same result for the growth rate. Note however that the *level* of production at a given time will be lower given depreciation, since the stock of capital will be lower. More precisely, assume an economy in which $\delta = 0$ is on its b.g.p. Now assume that, suddenly, the capital stock starts to depreciate at rate $\delta > 0$; then the economy will shift gradually towards a new b.g.p. on which (a) the level of Y at a given time is lower than on the previous b.g.p., and (b) the growth rate of Y is the same. This shift is illustrated in Figure 4.3.


FIGURE 4.3. The transition from one b.g.p. to another (lower) b.g.p., after δ increases from 0 to 0.1. Note that the figure is for a discrete-time economy with the following parameters: $\alpha = 0.33$; $\beta = 0.05$; $\theta = 0.1$ (assumed exogenous and constant); $s = 0.2$; $g_{A_L} = 0.03$. The economy starts with a capital stock such that it is on the initial b.g.p.

Note that we also require a *transversality condition* to solve the model fully. It is not enough to know the rate of change of prices, we must also know the level of the price as some point in time. There are two particularly simple possibilities: one is that the resource should only be exhausted asymptotically, the other is that the price should be equal to some backstop price at the time of exhaustion. For the case of exhaustion with a backstop, a slower rate of price increase implies (for a given transversality condition) that the initial price must be higher, and hence the initial price path will be higher and the date of exhaustion will be put back.[3]

4.2.4. Hartwick. The majority of the older literature focused on the case with no technological progress, i.e. $g_{A_L} = 0$. Then we know straight away (from the Solow model) that long-run growth is not possible, since there are diminishing returns to capital accumulation. This conclusion is strengthened when there are also resource constraints: non-negative growth is impossible in the baseline case with a fixed resource stock free to extract. Given such a stock, the extraction rate must approach zero in the long run, and the only way to compensate is if the capital stock (and hence also K/Y) approaches infinity. But this is impossible to sustain in the presence of depreciation.

The response to these results was not to reject the model with zero technological progress, but instead to do further violence to the original Solow set-up: having taken away technological progress, the next step was to take away capital depreciation as well. In this set-up long-run growth remains impossible (due to diminishing returns to capital) but sustaining production and

---

[3]A full analysis of how transversality conditions help us to solve dynamic optimization problems is beyond the scope of this book.
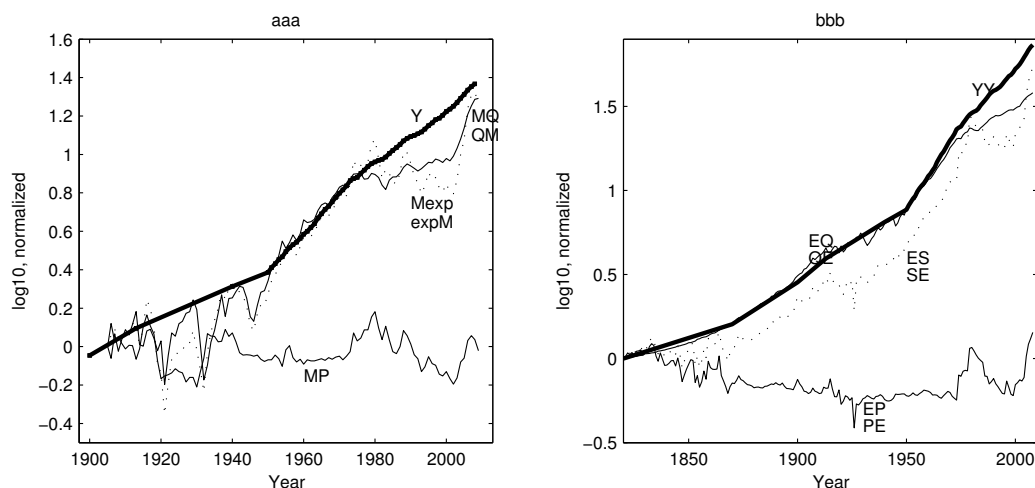
FIGURE 4.4. Long-run growth in consumption and prices, compared to growth in global product, for (a) Metals, and (b) Primary energy from combustion.

consumption may be possible: all that is required is a sufficiently high rate of capital investment. How high is sufficiently high? This question is tied up with the famous Hartwick rule. For an analysis of this rule see Appendix 4.B.

4.2.5. Historical data. Recall that the key results of the DHSS/Hotelling model are that on a balanced growth path the following relationships should hold:

$$g_Y = g_{A_L} - \rho \frac{\beta/(1-\alpha)}{1-\beta/(1-\alpha)}$$

and
$$\theta = \rho - g_Y,$$

while
$$\dot{w}_R/w_R = \rho.$$

Here θ is the rate of decline of resource use over time, ρ is the pure rate of time preference, and w_R is the resource price. Furthermore, α and β are the respective factor shares of capital and the resource in final-good production. The results show how the need for resource inputs ($\beta > 0$) puts a penalty on the growth rate that can be sustained in the economy. The size of the penalty is increasing in the rate of time preference, ρ, and the weight of the resource in the production function, β. Furthermore, we see that resource use declines at a constant rate, and that the resource price increases rapidly (recall that $\rho > g_Y$).

How do these results match up to the data for real economies? The answer to this question is to be found in Figure 4.4, which builds on Figure 1.3. Here we see that—far from declining exponentially—resource and energy use have risen exponentially, tracking global product. Furthermore, the prices of resources and energy have been remarkably constant in the long run, whereas according to the model they should grow faster than global product (the discount rate is always faster than the growth rate in optimal growth models). So the model predictions seem to be completely, hopelessly wrong.

Things are perhaps not as bad as they seem for the DHSS model. The reason is that the problems can all be traced to the same root, which is the wildly incorrect model of the resource extraction sector. Resources are not in reality available in a fixed known quantity, free to extract. Instead they are expensive to extract, and furthermore the costs of extraction vary in a complex way due to factors such as technological change and cumulative extraction (as more is extracted we must dig deeper, raising costs).

Recall our previous model in which resource stocks were unlimited and extraction costs constant. This model did a much better job of fitting the data, but at the expense of ignoring resource scarcity, making it useless for very long run analyses. In the next section we develop a model in which we allow for both extraction costs and scarcity.

4.3. An extended DHSS model: A limited resource, costly to extract

We have established that the scarcity rent of major non-renewable resources is currently low. However, we also know that in the very long run the rate of extraction of non-renewable resources cannot keep on increasing as it has over the last century. How will this slowdown occur, and what will happen to prices and the scarcity rent?

The above analysis shows that we need a model with both extraction costs and limits on resource stocks. With such a model we have a chance of both explaining historical data and predicting the future. The simplest thing to do is to assume that marginal extraction costs are constant, c. Continue to write the price as w_R, and the resource rent as λ. Then (in the case with perfect competition) the price is the sum of the resource rent and the extraction cost, $w_R = \lambda + c$, and since c is constant $\dot{w}_R = \dot{\lambda}$. The resource rent rises at the interest rate ($\dot{\lambda}/\lambda = w_m$) and the modified Hotelling rule is

$$\frac{\dot{w}_R}{w_R} = w_m \cdot \frac{w_R - c}{w_R}.$$

So when $c = 0$ this reduces to the original rule, but when c approaches w_R (implying that the resource rent is only a small proportion of the total price), then $\dot{w}_R/w_R$ approaches zero, i.e. constant price.

Unfortunately the model with constant extraction costs is still all-too naive compared to reality in which a variety of factors affect extraction costs. Consider the picture below, and with its help try to identify as many such factors as you can. Furthermore, categorize them according to whether they should make extraction costs rise or fall over time.

The three key factors are the wages paid to workers, their productivity, and the depth of the resource (where the latter should be interpreted broadly as the physical quality of resource deposits). Both wages and productivity should rise over time, and these trends may be expected to cancel each other out. (Why?) Depth should also rise over time, hence overall it seems that we should expect resource prices to rise over time. Note also that if the 'width' of the resource stock increases, then the rate of increase in depth for a given extraction rate will be lower. We now capture this—and more—in a simple model, based on the more complex model of Hart (2016).

4.3.1. The basic environment.
The basic model is illustrated in Figure 4.5. There is a constant population and utility is a linear function of consumption of the aggregate good Y:

$$U = \int_0^\infty e^{-\rho t} Y_t \, dt.$$

The interest rate is thus constant and equal to ρ. The price of the final good is normalized to 1. All markets are perfect—there are many resource owners, and many final-good producers.

As we can see in Figure 4.5, labour is allocated either to final-good production or to extraction. The extraction rate grows linearly in effective labour inputs, $A_X L_X$, where A_X (productivity) grows exogenously. However, the rate declines in A_D, which we denote 'economic depth'. Economic depth is the exponential of the physical depth of the marginal resource, D. The physical depth increases over time, with the rate of increase $\dot{D}$ equal to the rate of extraction divided by the surface area of the resource, ϕ. The resource flow X is one of two inputs into final-good production, the other being effective labour $A_Y L_Y$, where A_Y grows exogenously. It follows that the rate of increase in economic depth A_D is simply X/ϕ, so given a constant rate of extraction, the economic depth increases at a constant rate. For simplicity, we assume that labour productivity rises at the same rate in both sectors, so $\dot{A}_Y/A_Y = \dot{A}_X/A_X = \theta_A$. Finally, as hinted at in Figure 4.5, there is a limit to the depth at which the resource can be extracted.

To solve the model at a given time, the key is that the marginal revenue product of the resource in the final-good sector must be equal to the price at which the (representative) extraction firm is willing to sell the resource. This price will be equal to the sum of the marginal extraction cost and the resource rent (which the extraction firm must pay to the resource owner). Calculating the resource rent is tricky, so we 'park' the rent and simply denote it r_X. Furthermore, for convenience

FIGURE 4.5. A schematic diagram showing the economy with an inhomogeneous resource, costly to extract.

we define

$$r_X^* = r_X/w_L,$$

where w_L is the wage. So r_X^* is a (re)normalized resource rent. The other two terms are straightforward. The MRP of the resource is obtained directly from the production function, and marginal extraction costs are simply total extraction costs divided by extracted quantity (since the extraction function is linear). Put this all together and rearrange to obtain

$$X = \frac{\alpha}{1-\alpha}\left(\frac{A_D}{A_X} + r_X^*\right)^{-1}.$$

Now assume a *primitive* economy in which A_Y and A_X are both very small. What happens in the short run, and in the long run? Since labour productivity is low, both resource extraction X and final-good production Y must be low. This implies that the rate of increase of the depth of the resource must also be low. In the limit (when this rate of increase is extremely low), the resource rent must be zero since extraction has a negligible effect on the size of the stock. So then we have

$$X = \frac{\alpha}{1-\alpha}\frac{A_X}{A_D}.$$

And since A_D is effectively constant (when the extraction rate is very low), the extraction rate grows at the same rate as extraction productivity, i.e. at rate θ_A. It follows (by inspection) that the allocation of labour is constant, as is the resource price. We thus have a simplified version of the DHSS model with an abundant but costly resource, as in Section 4.1.2. (The simplification is that there is no capital.) And in the limit of very low productivity, we have a balanced growth path:

$$\dot{Y}/Y = \theta_A$$
$$\dot{X}/X = \theta_A$$
$$w_X X = \alpha Y$$
$$\dot{w}_X/w_X = 0.$$

Starting in the primitive economy, over time the level of labour productivity, and with it the rate of resource extraction, will increase. Gradually, depth D will also start to increase significantly. And when depth increases, the resource price will tend to increase, braking the increase in the extraction rate. Here we show that there exists a second b.g.p., on which the extraction rate X is constant, and production, the resource rent, and the resource price all grow at the same rate, θ_A.

Crucially, for balanced growth we have

$$X = \phi\theta_A,$$

implying that

$$\dot{A}_D/A_D = \theta_A.$$

Hence the increase in economic depth exactly balances the increase in labour productivity, holding the extraction rate constant for the given labour allocation to extraction. Since resource use is not growing, final-good production Y grows slightly more slowly than on the previous b.g.p., while the resource price tracks Y:

$$\dot{Y}/Y = (1-\alpha)\theta_A$$
$$\dot{X}/X = 0$$
$$w_X X = \alpha Y$$
$$\dot{w}_X/w_X = \dot{Y}/Y.$$

Note also that on this b.g.p. the resource rent will also grow at the overall growth rate.[4]

For the b.g.p. derived above to hold, we must be a long way from the time of exhaustion. The reason is that as exhaustion approaches, the resource rent will rise faster than θ_A, because in addition to the increasing depth, extraction will also bring the time of exhaustion closer. This is the standard reason for the existence of a resource rent, which applies when the resource stock is homogeneous. When exhaustion is very close then (if the resource is essential for production) the resource price must be very high and the extraction rate very low, and hence the depth essentially constant. In the limit we have a simple Hotelling economy in which extraction costs are zero (since depth is effectively constant, the rising extraction productivity pushes extraction costs towards zero) and the resource price rises at the interest rate ρ. We leave it to the reader to derive the equations for the rates of change of the other variables in the model on this growth path. Note however that the path differs from the others in that labour allocation changes over time: as the resource price rises at the discount rate ρ and extraction declines, extraction labour also declines, approaching zero. Resource extraction also approaches zero, with the resource being exhausted asymptotically.

4.3.2. Economic development in theory and practice. We now turn to the use of the model to understand real economies. First note that to do so we must generalize the model to allow for different types of resource stock, e.g. stocks where abundance declines with depth, or first increases and then declines. Such a model is developed in Hart (2016). Based on this model we first discuss transition paths in general, then we look at the specific cases of copper and petroleum. The parameterizations are illustrative rather than strictly predictive, the main reason being the great uncertainty concerning many of the assumptions. Nevertheless, the model succeeds in explaining observations from the last 100 years, and makes apparently reasonable predictions for the next several hundred in the cases of oil and copper.

Transition paths. The overall picture emerging from the model is straightforward. In the case of a strictly limited resource which is essential for production then we expect the economy to pass through three phases. In the initial or frontier phase resource depth is constant, extraction increases rapidly, and the resource price is roughly constant. In the mature phase resource depth increases, the rate of increase of extraction is moderated, and the resource price rises. In the exhaustion phase the depth is again constant, the extraction rate approaches zero, and the price rises at the interest rate.

The case of a strictly limited and essential resource is however not very realistic. The only truly non-renewable and essential resource is energy, but energy is not limited since we have the option of harvesting the inflow of energy from the sun. On the other hand, although minerals such as metals are of course in strictly limited supply they are not consumed in the production process, hence we have the option of recycling. Furthermore, if we focus on each metal separately then none of them are essential to the production process: if one metal runs out (including options for recycling) then we will of course do without it, substituting it with other materials.

Since there are—very generally—substitutes for non-renewable resources it is important to include this fact in the model. The simplest way to do so is to assume a *backstop resource*, i.e. a substitute available in unlimited quantity at an exogenous price. When the resource price hits the backstop price extraction stops. Clearly at this point the value of remaining resource stocks is zero, implying that the the scarcity rent λ is zero. We can solve the model in this case, and the behaviour

---

[4]The reason that there is a resource rent is that extracting resources today means that deeper resources must be extracted (at greater cost) tomorrow.

of the economy as the backstop price is approached depends very strongly on that price: if the backstop price is very high then the behaviour of the economy will be similar to the exhaustion phase as the backstop price is approached, with rapidly increasing price and declining extraction. On the other hand, if the backstop price is 'moderate' then the mature b.g.p. may evolve seamlessly into the backstop economy, and if the backstop price is very low then the mature b.g.p. may never be approached: instead, the frontier economy may evolve directly into the backstop economy.

Copper. The parameterization of the model for copper is complex, since there are two key dimensions along which the quality of copper deposits varies. The first of these is grade (the fraction of copper in the rock, by mass), and the second is depth. Combining grade and depth into one measure of 'depth–grade' r_x we obtain the relationship described in Figure 4.6. Here we see that extraction so far is only scratching the surface of total stocks, and furthermore that marginal stocks are (for now) rapidly increasing in depth. Furthermore, our calculations suggest that depth–grade r_x is only relatively weakly linked to 'economic depth' a_n: $a_n = r_x^{0.42}$. This leads to the result that the upward pressure on copper prices due to increasing depth is weak and will remain so for a long time into the future.


FIGURE 4.6. The relationship between the combined measure of depth and grade, r_x, and cross-sectional area m for copper, based on our interpretation of the literature (continuous line), and the parameterization of our economic model (dashed line). Notice the difference in scale on the two panels. The shaded area shows extraction from 1900–2011, 5.75×10^8 tons, and the dotted lines show the relationship between depth and cross-section layer-by-layer.

Feeding all this into the overall general-equilibrium model we obtain the results shown in Figure 4.7. The economy starts close to the first b.g.p. which applies for the initial stock, and price declines by around 0.4 percent per year. Once the initial stock is used up in 2054, the rate of increase in depth a_n increases, and the economy starts moving towards the second b.g.p. on which price rises by 0.8 percent per year. Finally, from around 2200 the scarcity rent starts to rise as exhaustion approaches, at least in the case with a high backstop price. With a low backstop price the scarcity rent hardly rises, and exhaustion occurs a few years earlier. Note the close agreement between the model and observed trends in prices and extraction rates.

Petroleum. If the copper simulation is an advertisement for the power of the model, the petroleum simulation highlights its weaknesses. There are two key aspects of the petroleum market which the model cannot handle as it stands: firstly, the significance of market power in the petroleum market, and secondly the inextricable links between petroleum and its substitutes, including natural gas, coal, and other energy sources such as nuclear power. Of course, market power and substitutes also exist in the market for copper, but their scale and influence is greater in the oil market. Concerning market power, consider for instance the fact that petroleum extraction occurs simultaneously from deposits for which marginal extraction costs differ by a factor of 5 or more (compare for instance the Ghawar field in Saudi Arabia to the Athabasca oil sands of Alberta). Concerning substitutes, petroleum demand is linked tightly to markets for coal and other energy

FIGURE 4.7. Observed price and extraction rate of copper, and the paths of price and extraction rate—up to the time of exhaustion—predicted by the model. Two model scenarios are shown, which differ in the assumed backstop price. Note that the extraction rate is plotted on a (natural) logarithmic scale, normalized by the rate in 2000. Prices are in 1998 USD.

sources, and strongly affected by technological change. Consider for instance the substitution from coal to oil driven by the development and refinement of the internal combustion engine. Given these problems—which are evident in Figure 4.9—the model calibration is at best illustrative, showing possible future scenarios and highlighting the effect of backstop energy sources.

The data regarding petroleum resources in the ground are uncertain. Furthermore, the data regarding the cost of extraction of these resources are even more uncertain. The most frequently cited paper on the subject is probably Rogner (1997). However, Rogner's curve relating cumulative extraction to extraction cost (see for instance his Figure 6) shows estimated extraction cost *at the time of extraction*. Its calculation must therefore involve (implicit or explicit) calculations of (i) current extraction costs, (ii) expected decline in extraction costs, and (iii) expected rate of extraction. Since we model the latter two, we need data on the first factor alone, i.e. unit extraction costs for each type of deposit making up the reserves, if full-scale extraction were to be carried out today. This is estimated by the International Energy Agency in their World Energy Outlook 2008 (p.218). The data are very approximate, but can be broadly summarized as follows: considering initial resource stocks, there was a large rather homogeneous stock of easily accessible stocks, approximately 2000 billion barrels at an economic depth of around 18 USD/barrel. Regarding the remaining stocks—about 7000 billion barrels—economic depth a_n rises approximately linearly with cumulative extraction, reaching approximately 115 USD/barrel for the deepest stocks. We capture this in the model by assuming an initial stock with low ψ ($\psi = -2.2$), so that the entire near-homogeneous stock is at a depth of 10–20, switching to the deeper stock with $\psi = 1$ from 20–115. The cross-section of the second stock is determined by its size (assumed to be 6.7×10^9 barrels), and the parameters for the first stock are then fixed by the limits on depth (10–20), the size (2.3×10^9 barrels), and the need for m to be continuous over the boundary between the stocks. The result is shown in Figure 4.8. Note that the curve shows unit extraction costs, in 2008 USD with today's technology, for all petroleum resources including the (hypothetical) current extraction cost of resources already extracted. (Note that we ignore the fact that a significant proportion of cumulative extraction has been from deeper stocks.)

Having parameterized the model, and given the assumption about the total stock of resources, the future development of prices and quantities predicted by the model depends on what we assume about the price of the backstop (i.e. the substitutes for oil that will take over when oil is exhausted or too expensive). Here we make two alternative assumptions to demonstrate the role played by the backstop resource. In the first case we assume that a backstop is available at a fixed price of 150 US dollars (2011); in the second case we assume that a backstop is available *today* at that price, and that this price will decline at the rate $\theta_{ax} - \theta_{ay}$; that is, the backstop price declines as long as
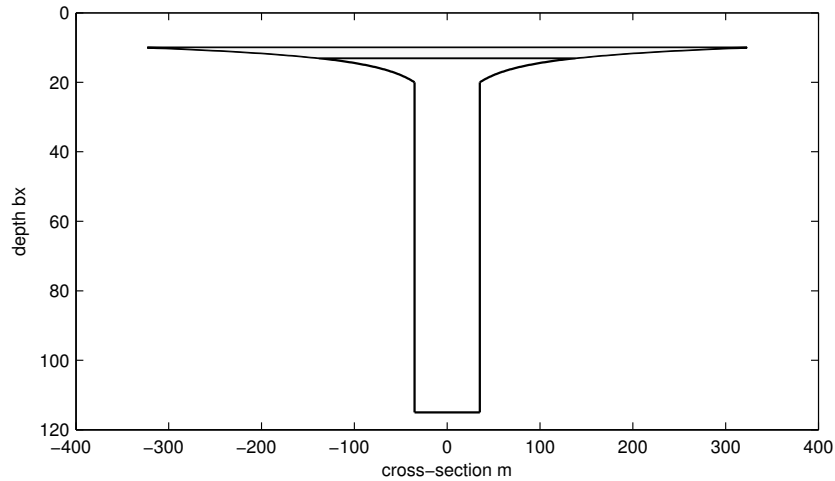
FIGURE 4.8. The relationship between economic depth, a_n, and cross-sectional area m for oil, based on our interpretation of the International Energy Agency World Energy Outlook 2008. Depth is measured in 2008 USD, and cross-sectional area in billion barrels per USD. The shaded area shows extraction from 1900–2008, 1100 billion barrels.

manufacturing productivity growth outstrips TFP growth. The result is that the backstop price is around 65 USD at the time of exhaustion, rather than 150.

FIGURE 4.9. Observed price and extraction rate of petroleum, and the paths of price and extraction rate—up to the time of exhaustion—predicted by the model. Two model scenarios are shown, which differ in the assumed backstop price. Note that the extraction rate is plotted on a logarithmic scale, normalized by the rate in 2000. Prices are in 2012 USD. Scenarios: continuous lines, backstop price 150 USD (year 2012); dashed lines, current backstop price 150 USD, declining at a rate $\theta_{ax} - \theta_{ay}$ per year. Price data from BP (2012), consumption data from Boden et al. (2012), assuming a linear relationship between CO_2 emissions and petroleum consumption.

The results are as follows. Note first that there is no market power in the model economy, hence the results are what the model predicts in an economy similar to the actual global economy but without the exercise of market power by oil producers. Turning now to the results, up to the exhaustion of the upper stock, depth is almost constant, the scarcity rent is close to zero, and price declines at a rate equal to the difference between the growth rates of extraction productivity and labour productivity in final-good production, i.e. 0.6 percent per year. However, as the upper stock nears exhaustion depth starts to rise at a significant rate, and the economy heads back towards the

b.g.p. for the stock, for which $\psi = 1$; the mature extraction phase. On this b.g.p. the growth rate of extraction is halved, the resource price rises by 0.6 percent per year, and the scarcity rent makes up 21 percent of the price.[5]

In the latter half of the 21st century the price paths of the alternative backstop scenarios diverge significantly: the upper path (high backstop price) is slightly above the b.g.p. price path, while the lower path is below it. Hence when the backstop price is fixed at 150 USD the scarcity rent rises above 21 percent of the total price as exhaustion approaches, whereas given the lower backstop price the rate of price increase slows down as exhaustion approaches, and the rent actually declines as a proportion of the price.

4.3.3. Sensitivity of the model to assumptions. The above simulations are sensitive to the assumptions made, the most uncertain of which are those regarding future demand, and future development of extraction productivity. On the one hand, our assumption about future demand is essentially at the upper bound of what is realistic, i.e. that demand per capita continues to grow indefinitely at a similar rate to the rate observed over the last 100 years. At least two factors might be expected to lead to lower future demand: firstly, if global growth slows in the long run, and secondly if there is a transition from 'early' growth based on manufacturing and hence resources, and 'post-industrial' growth based on services and hence labour.[6] The effect of lower demand would be to reduce extraction rates and hence also reduce the growth rate of prices predicted by the model. On the other hand, our assumption regarding future development of extraction productivity is also an upper bound; again, we assume that it continues to increase indefinitely. This is unlikely, not least because in reality resource extraction requires energy, and there are physical limits to the efficiency with which this energy can be used. Since these limits are already coming close in some cases, this implies that even if labour productivity continues to increase, energy productivity will not do so and hence the proportion of the energy cost in the unit cost will rise, and the rise of overall extraction productivity will slow. This assumption therefore biases the results towards lower prices and higher extraction rates than are likely to be observed.

The effect of assuming both lower future demand and lower productivity growth rates in extraction is therefore that the price path is likely to be relatively unchanged, whereas the extraction path will be lower. Furthermore, the proportion of the price accounted for by the scarcity rent will be lower. Given a finite stock, the lower extraction path will lead to later exhaustion, and hence any price spike as exhaustion approaches is also likely to be delayed.

4.4. Limitations of the Cobb–Douglas production function

So far in this chapter we have discussed alternative assumptions about the nature of natural resource stocks, arriving (in Section 4.3) at a reasonably general model which can be used to both explain historical observations and predict future trends. However, we have scarcely discussed how natural resources enter into the production function. It turns out that the choice of the Cobb–Douglas production function is an even bigger problem for the DHSS model than the assumption of finite resource stocks, free to extract.

The choice of the Cobb–Douglas was quite extensively discussed in the original papers. For instance, Dasgupta and Heal consider a more general (CES) production function, but note that (p.14) "Only the Cobb-Douglas form may be said to have properties that are reasonable at the corner", and (p.17) "The Cobb-Douglas case is particularly interesting since the analysis can relatively easily be taken further."[7] Similarly to Dasgupta and Heal, Solow (1974) moots the idea of using a production function other than Cobb–Douglas, but concludes (p.34) that "Any extra generality [gained by departing from Cobb–Douglas] hardly seems worth striving for." And Stiglitz (1974) sets up exactly the above model, again mentioning alternative production functions but stating (p.124) that Cobb–Douglas case is "special, but . . . central".

---

[5]Note that after the transition to the deeper stock with $\psi = 1$ the economy approaches the mature b.g.p. for that stock from *above*, i.e. the state variable a is above its level in the steady state. As the economy approaches the new b.g.p. x_1 falls back, which is why prices rise quite steeply throughout the 21st century.

[6]To get a feel for the sizes of demand changes in the model, we consider each simulation in turn. For copper, the extraction rate peaks at around 50 times the observed rate in year 2000. Compare this to the arbitrary assumption that the entire future global population consumes copper at the same rate as the average U.S. citizen in year 2000; this would lead to a global extraction rate approximately 6.3 times greater than that observed in 2000. If demand levels off in this way then the copper stocks will last for many centuries rather than just two or three. For petroleum, the extraction rate in the model peaks at around 3 times the year 2000 rate, which is less than the rate which would arise if all countries matched the U.S. per-capita rate from year 2000.

[7]Interestingly, the last section of the paper is devoted to an extended model in which a technological leap—occurring exogenously and at unknown time—allows the use of an alternative to the exhaustible resource which arrives in a constant inexhaustible flow. This is related to mechanism 3 defined below.

On the other hand, criticism of the DHSS model from (for instance) ecological economists has often focused on how resources enter the production function. Consider for instance Herman Daly's critique of Solow–Stiglitz [i.e. DHSS], Daly (1997) p. 263:

> In the Solow–Stiglitz variant, to make a cake we need not only the cook and his kitchen, but also some non-zero amount of flour, sugar, eggs, etc. This seems a great step forward until we realize that we could make our cake a thousand times bigger with no extra ingredients, if we simply would stir faster and use bigger bowls and ovens.

The point here is that given the Cobb–Douglas production function (especially with the parameter values typically chosen, with $\alpha = 0.3$ and β no greater than 0.05) we can greatly increase production of final goods (cake) from a given flow of resources (ingredients) by either increasing inputs of effective labour (i.e. stirring faster) or capital (bigger bowls). This clearly makes no sense at the short-run, disaggregated level of analysis.

Macroeconomic models, building on aggregate production functions, are always gross simplifications of reality, and it is frequently claimed that we need to interpret the production function of the DHSS model flexibly. Taken literally, the model shows capital substituting for resource flows in the production function, where that 'capital' is simply foregone consumption of the single final good. But Groth (2007) p.10–11 argues that capital accumulation should be interpreted as a move towards clean technology, recycling, substitution between inputs, and changes in the composition of final output. However, the problem with this approach is that the interpretation of the model is so far removed from its actual assumptions as to make the model meaningless and impossible to test. Put differently, if we really think that moves towards clean technology, recycling, substitution between inputs, and changes in the composition of final output are the key to sustainability, then we should build models in which these processes are explicitly accounted for, and find ways of testing the strengths of these different processes and hence their ability to deliver sustainable growth.

The point is effectively conceded by the authors of the DHSS model themselves. Prior to the development of the DHSS model, Solow (1973)—in an essay where he is unconstrained by the need for mathematical formalism—argued for the ability of the economy to adapt to resource scarcity. In this context he sets out a series of mechanisms as key, without ranking them in importance; ironically, the DHSS mechanism—Solow's focus in subsequent quantitative modelling—is not mentioned at all. Regarding demand for resource inputs Solow sets out three mechanisms:

(1) Increase—through technological change—resource efficiency in production of one or more product categories;

(2) Substitute on the consumption side away from product categories in which the production process is resource-intensive.

(3) Increase—through technological change—the efficiency of an alternative (substitute) resource in production of one or more product categories.

Solow pays less attention to the supply of resources, but there we have improving technology of extraction and—working in the other direction—the impoverishment of deposits. He could also have mentioned recycling.

In the same vein, Dasgupta (1993) also discusses our ability to substitute for physical resources in the production process. Dasgupta argues that there are nine ways in producers can substitute for non-renewable resource inputs, of which 7 are different forms of technological progress, number 8 is switching to lower-grade inputs, and number 9 is the mechanism of the DHSS model, i.e. the substitution of capital for resources. Dasgupta admits (p.1115) that this last mechanism is 'limited', indeed 'beyond a point fixed capital in production is complementary to resources'.

Why did literature focus for so long on the substitution of capital for resources? One possible answer is to be found in a quote from Solow (1974) (pages 10–11):

> It is clear without any technical apparatus that the seriousness of the resource-exhaustion problem must depend in an important way on two aspects of the technology: first, the likelihood of technical progress, especially natural-resource-saving technical progress, and, second, the ease with which other factors of production, especially labor and reproducible capital, can be substituted for exhaustible resources in production.
>
> My own practice, in working on this problem, has been to treat as the central case (though not the only case) the assumption of zero technological progress. This is not because I think resource-saving inventions are unlikely or that their capacity to save resources is fundamentally limited. Quite the contrary ... I think there is virtue in analyzing the zero-technical-progress case because it is easy to see how technical progress can relieve and perhaps

> eliminate the drag on economic welfare exercised by natural-resource scarcity.
> The more important task for theory is to try to understand what happens or can
> happen in the opposite case.

However, this is contradicted by Solow's own analysis of the 'three mechanisms' discussed above, and also the analysis of Dasgupta (1993). The true answer is probably more prosaic: until recently the methodological tools necessary for the analysis of Solow's three mechanisms had not been developed. We argue—following Solow—that the three mechanisms above are key to understanding how the economy adapts in the long-run to changes in resource or energy availability, or to policy measures regarding resources or energy. How important are they relative to one another, and how does their existence affect optimal policy? We now move on to an analysis of these questions.

4.A. Appendix: Hotelling with a resource supply monopoly

Imagine you have a monopoly over the resource; you are the only supplier. This makes your problem more complex, since you must account for the effect of your own extraction on the price. To see what's going on, we need to set up a *dynamic optimization problem*. The simplest way to do this is to use *discrete time*, in which case we need to set up a *Lagrangian*. We want to maximize discounted profits, subject to the resource restriction. The profit function and the restriction can be written as follows.

$$\pi = w_{rt}R_t; \qquad \sum_{t=0}^{\infty} R_t \leq S.$$

Here w_{rt} is the resource price, and R_t is the quantity extracted and sold in period t. We can use these two equations to set up the Lagrangian, which is

$$\mathcal{L} = \sum_{t=0}^{\infty} \left(\frac{1}{1+w_m}\right)^t w_{rt}R_t + \lambda(S - \sum_{t=0}^{\infty} R_t) \tag{4.3}$$

Here λ is the *resource rent* or (expressed in a more standard way) the shadow price of the resource stock.[8] Now take the f.o.c. in R_t to obtain

$$(1+w_m)^t \frac{\partial \mathcal{L}}{\partial R_t} = w_{rt} + R_t \frac{\partial w_{rt}}{\partial R_t} - \lambda(1+w_m)^t = 0. \tag{4.4}$$

Now the first two terms together represent the change in revenue given a change in quantity: marginal revenue, MR. To find the Hotelling rule, consider the first-order condition in successive periods to derive

$$\frac{MR_{t+1}}{MR_t} = 1 + w_m.$$

Now let's close the model by adding a demand function. For reasons that will become clear, we assume an inverse demand function of the form

$$w_{rt} = A + BR_t^{-\epsilon}.$$

Then

$$MR_t = w_{rt} + R_t \frac{\partial w_{rt}}{\partial R_t} = w_{rt}(1 - \epsilon) + \epsilon A,$$

and the rule is

$$\frac{w_{rt+1} + \epsilon A/(1-\epsilon)}{w_{rt} + \epsilon A/(1-\epsilon)} = 1 + w_m.$$

To get a handle on the above result assume that our monopolist is actually a representative resource owner with no market power. Then marginal revenue is simply price, and

$$\frac{w_{rt+1}}{w_{rt}} = 1 + w_m.$$

To link with our results above (in continuous time) note that the expression corresponding to $\dot{w}_R/w_R$ is $(w_{rt+1} - w_{rt})/w_{rt}$. But we can rearrange the above result to yield

$$\frac{w_{rt+1} - w_{rt}}{w_{rt}} = w_m,$$

which is thus equivalent to the continuous time result we obtained earlier (see the mathematical appendix, Chapter A.2). So when $A = 0$ the growth rate of prices is identical to the growth rate given perfect markets; when $A < 0$ prices grow more slowly given market power; and when $A > 0$ prices grow faster given market power. The intuition comes from the elasticities: when $A = 0$ we have constant elasticity demand, so as the price rises over time the elasticity of demand is unchanged. But when $A < 0$ the elasticity of demand increases as price increases, making the monopolist more reluctant to increase prices, and when $A > 0$ the elasticity of demand decreases as price increases, encouraging the monopolist to raise prices further.

---

[8]This quantity has many names in the literature; two other common ones are the *Hotelling rent* and the *scarcity rent*.

4.B. Appendix: Capital, natural resources, and the Hartwick rule

The Hartwick rule—Hartwick (1977), following Solow (1974)—has generated a huge amount of confusion, and many incorrect claims have been made.[9] Here we state some of the correct results in a simple way. The key result is that—in an economy with perfect markets, constant population and constant technology—if the value of capital stocks is kept constant then production will also be constant.

What does the above result imply about sustainability? How difficult is it to maintain a constant capital stock? If capital doesn't depreciate then all we need to do to keep the total stock constant is to invest the proceeds of using up one stock of capital into boosting some other stock of capital. On the other hand, if capital *does* depreciate over time then we have a much harder task holding the total stock of capital constant, indeed it will typically be impossible to do so.

Consider a DHSS economy such as those we have analysed above, and assume that investment in capital equipment K is determined by the Hartwick rule, implying that it is not determined by market. In such an economy there are two capital stocks, K and R, and the rule is

$$w_K \dot{K} + w_R \dot{R} = 0. \tag{4.5}$$

So we assume that a regulator lets the market determine resource extraction and price, and then invests according to the rule.

The production function is

$$Y = (A_L L)^{1-\alpha-\beta} K^\alpha R^\beta$$

and the capital accumulation equation is

$$\dot{K} = sY - \delta K,$$

where we now treat s as an unknown variable. Furthermore, w_K and w_R are the (respective) marginal revenue products of K and R, hence

$$w_K K = \alpha Y$$

and $\qquad w_R R = \beta Y.$

The first equation tells us that the marginal product of capital $w_K = \alpha Y / K$. Consider now a capital owner hiring out a single unit capital, value 1 (recall that capital and the final good are the same, and that the price of the final good is normalized to 1). The gross return (income flow per unit of capital hired out, value 1) is just w_K, but the net return is $w_K - \delta$, because the capital depreciates (disappears!) at a constant rate of δ. In equilibrium capital owners must be indifferent between hiring and selling capital, which implies that the interest rate—here we call it ρ rather than w_m—must be equal to $w_K - \delta$. So we have

$$\rho = w_K - \delta = \alpha Y / K - \delta. \tag{4.6}$$

Now return to $w_K K = \alpha Y$ and use the capital accumulation equation to write

$$w_K \dot{K} = \alpha Y / K (sY - \delta K). \tag{4.7}$$

Now turn to R. Since $w_R R = \beta Y$ and the Hotelling rule applies we know that

$$\dot{R}/R = \dot{Y}/Y - \rho,$$

where ρ (the discount rate) is determined by (4.6). Furthermore, from the production function we know that

$$\dot{Y}/Y = \alpha \dot{K}/K + \beta \dot{R}/R,$$

hence $\qquad \dot{R}/R = \alpha \dot{K}/K + \beta \dot{R}/R - \rho.$

Rearrange to yield $\qquad \dot{R}/R = \dfrac{1}{1-\beta} \left[\alpha \dot{K}/K - \rho \right],$

hence $\qquad \dot{R}/R = \dfrac{1}{1-\beta} \left[\alpha sY/K - \alpha \delta - \rho \right],$

and (since $w_R R = \beta Y$) $\qquad w_R \dot{R} = \dfrac{\beta}{1-\beta} \left[\alpha sY/K - \alpha \delta - \rho \right] Y. \tag{4.8}$

Now combine (4.7) and (4.8) with (4.5) to yield

$$sY = [\rho(\beta/\alpha) + \delta] K,$$

and (substituting for ρ) $\qquad sY = \beta Y + \delta(1 - \beta/\alpha) K.$

---

[9]For rigourous support for this statement see Asheim et al. (2003).

When $\delta = 0$ a Hartwick path is feasible: investment is a constant fraction of production, the interest rate approaches zero from above, and total resource extraction is bounded. However, when $\delta > 0$ then there is no feasible Hartwick path: on a feasible path, R must approach zero while K approaches infinity, but if $K \to \infty$ then s must be greater than 1 (which is impossible), and furthermore Y/K will approach zero implying that ρ must be negative hence R must grow, not decline.

Finally, note that even if we observe in the present (a) that resource rents are reinvested in physical capital, and (b) that physical capital does not depreciate, and (c) that all markets are perfect, this still does not guarantee sustainability, because we do not know the preferences based on which the market interest rate is determined. One thing we can say for sure is that if the investment rate is determined by the market and if the pure rate of time preference is non-zero then the rule will not be followed in the long run. The reason is that given a positive pure rate of time preference ρ, the interest rate must always be at least ρ, even when growth is zero. This implies that investment in capital will hit zero when returns to such investment hit ρ, and this must happen in finite time, when the capital stock gets 'too big' in relation Y. So in the market there is a limit to capital accumulation, even though in a planned economy there may not be.

Summing up, it remains to be demonstrated that the Hartwick rule is of anything other than purely academic interest.

4.C. Appendix: Are non-renewables 'scarce'? The elephant in Hotelling's room

In this section we show that—according to market agents—exhaustion of critical non-renewable resources is definitely not imminent. We borrow heavily from Hart and Spiro (2011): large sections of the text are taken directly from that paper.

4.C.1. Theory. Market agents value resources in the ground as assets, as shown by Hotelling. If those resources are valuable, and their price is failing to rise, then agents will realize those assets, i.e. sell them. This will cause the price to fall to a lower level, from which it will (in equilibrium) rise.

In terms of Hotelling's analysis, there are two possible reasons for the failure of prices to rise: either (i) resource markets systematically fail to value resources in the ground according to the theory; or (ii) the scarcity rent is well-behaved, but masked by other factors. In both cases, the implication is that factors other than the scarcity rent are important in shaping the resource price. Failure to value resources correctly could for instance be due to a failure to foresee (stochastic) discoveries, leading to a fall in the rent each time a discovery is made; this is illustrated in Figure 4.10(b). However, as was shown as early as 1982 by Arrow and Chang, rational actors will take account of the probability of new discoveries being made, and the effect of allowing for stochasticity will be for the rent to fluctuate around the original trend. Only in the case of a constant series of surprises, all in the same direction, could new discoveries hold back the long-run growth in the rent. Imputing such a series of surprises boils down to assuming that actors on the resource market are not rational, hence the analysis is hoist by its own petard; that is, if market actors are not rational then other fundamental elements of the market analysis also break down. Another reason for markets not to value resources according to the theory could be that politico-economic factors play an important role, as argued by many authors in the resource curse literature.[10] Resources are frequently state-owned, and in such cases it is a reasonable conjecture that there are other, non-market, mechanisms shaping extraction and price paths.

Turning to the possible masking of the scarcity rent, note first that other components of the price may be a combination of extraction costs and rents due to market power, which we denote as 'resource costs'. Resource costs could mask a rising scarcity rent in two different ways, illustrated in Figure 4.10(c) and (d). In the first, which we denote 'declining costs', the rent is a significant component of the price, but the fall in resource costs compensates for its rise. In the second, 'low scarcity', the rent is only a tiny component of the overall price, and hence its rise has an insignificant effect on the overall price. It is straightforward to demonstrate both cases in theory. For instance, a popular explanation for declining costs is that technological progress in extraction pushes down unit extraction costs; see for instance Lin and Wagner (2007). However, as Hart (2009) points out, this should only occur if technological progress in extraction outstrips progress in other sectors, since otherwise the prices of inputs (such as labour) should rise in line with increasing productivity.

Note that low scarcity rents can arise in a number of ways despite finite stocks, for example if there is a renewable substitute. For brevity, we follow Nordhaus (1973a) in the following argument by assuming a *backstop technology*, i.e. a technology capable of substituting for the resource at a

---

[10]For a survey of this literature see Van der Ploeg (2011).

price which is independent of demand. Then if extraction costs are equal to the cost of a backstop technology, it is obvious that the scarcity rent will be zero up to the point of exhaustion.[11] More subtly, this will also be the case if extraction costs are expected to be equal to the backstop cost *at the time of resource exhaustion*. Furthermore, it has long been known that given a sufficient degree of market power prices will go straight to the backstop price and stay there, even in the absence of extraction costs; see for instance Teece et al. (1993) or Dasgupta and Heal (1979).

4.C.2. Simulation. We report the results of some simple simulations for crude oil. The aim is not to prove the level of the scarcity rent, but instead to illustrate the necessary implications of a high scarcity rent. We assume that resource costs change at a constant rate, while the scarcity rent rises at a constant rate (implying that the rate of return on holding the resource is constant). Under these circumstances, if prices are known at three points in time—such as past, present, and at the time of exhaustion—then if any one of the rate of return on the asset, the rate of cost decline, and the current scarcity rent are known, then the other two are fixed by the model. Furthermore, the lower the rate of return, the higher the current scarcity rent, and the steeper the rate of cost decline. We illustrate this in Figure 4.11, in which we take the average prices for 1980–1989 and 2000–2009,[12] and assume that oil reserves will run out in 2050 at a backstop price of 150 dollars/barrel,[13] and illustrate two of the possible combinations which are consistent with the model.

Figure 4.11 shows how a higher level of the current scarcity rent with a given backstop price and time of exhaustion implies, ceteris paribus, that resource holders demand a lower rate of return for holding the resource in the ground, and that extraction costs are falling more steeply. However, we wish to focus on the relationship between R (the current percentage of the price made up by the scarcity rent), the time of exhaustion, the backstop price, and the rate of return demanded by resource holders. To do so we plot level curves for R as a function of the backstop price and rate of return, for two different exhaustion dates (Figure 4.12).

In Figure 4.12 we see that the rate of return demanded on *in situ* resources is crucial to the level of the scarcity rent obtained from the model. If resource holders are happy with a rate of return similar to the average returns on bonds, around 3 percent,[14] then in the baseline scenario (backstop price 150, oil runs out 2050), the Figure shows that the scarcity rent makes up between 50 and 90 percent of the current price. On the other hand, if resource holders demand returns similar to average returns on shares, around 7 percent, then the scarcity rent accounts for just 10 percent of the current price.

The question of what rate of return demanded by resource holders has received little attention in the literature. Note however that Stollery (1983) finds by estimation of a CAPM model that the

---

[11]See Tahvonen and Salo (2001) for a model in which a renewable substitute plays the role of backstop, putting a cap on resource prices.

[12]55.4 dollars/barrel and 53.5 dollars/barrel, in constant 2009 dollars. For data see BP (2010).

[13]The current reserve to production ratio is 46 implying 2055 to be the expected time of exhaustion if current production continues. See BP (2010) for the data.

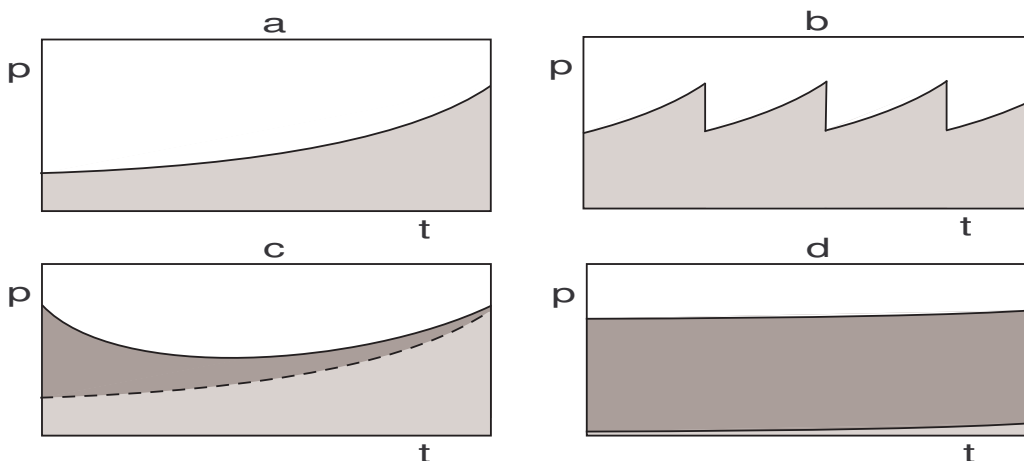[14]Average returns on 3-month UK treasury bonds are approximately 3 percent per annum.

FIGURE 4.10. Resource price as a function of time. The lightly shaded area represents the scarcity rent, the heavily shaded area represents other factors, i.e. the sum of extraction costs and rents due to market power.
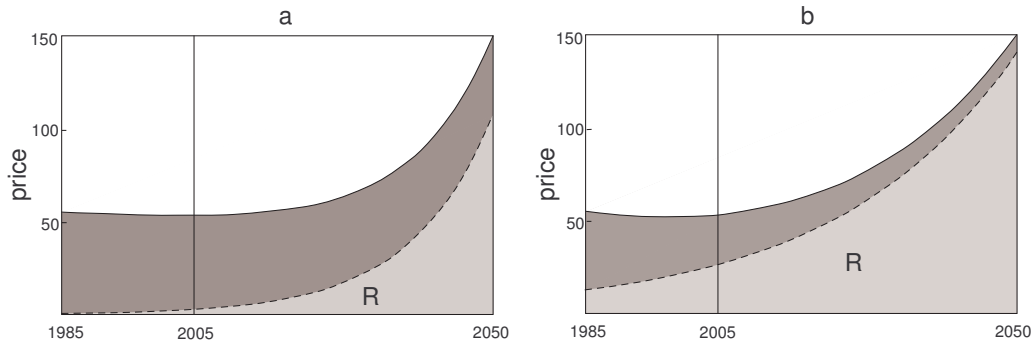
FIGURE 4.11. Graphs of price (dollars per barrel) against time for crude oil, showing alternative paths between the same observed prices for 1980–1989, 2000–2009, and in 2050 when the backstop is assumed to take over. The lightly shaded area represents the scarcity rent, which grows at rate r, and the darker area represents resource costs, which shrink at rate γ. In (a) $R_{2005} = 5$ percent, implying a high rate of return $r = 8.5$ percent/yr, and a slow decline in costs, $\gamma = 0.4$ percent/yr; in (b) $R_{2005} = 50$ percent, implying a lower rate of return $r = 3.8$ percent/yr, and a more rapid decline in costs, $\gamma = 2.4$ percent/yr.


FIGURE 4.12. Simulated level curves for the current scarcity rent as a percentage of current price of oil, where the variables are the final (backstop) oil price, and the rate of return (in percent per year) demanded by resource holders.

best fit comes from a rate of return of 14 percent for Canadian Nickel. We do not go deeply into the question here, but only note that commodity markets in general, and the oil market in particular, is considered volatile; see for instance Pindyck (2004). Thus it seems reasonable *prima facie* to assume that crude oil holders demand expected returns of at least 7 percent on *in situ* oil, probably significantly higher. This in turn suggests a scarcity rent of up to 25 percent.

From Figure 4.12 we can see that the backstop price and time of exhaustion also have significant effects. There is of course a lot of uncertainty about both numbers. However, Lindholt (2005) use a *current* backstop price of \$105/barrel (based on Manne et al., 1995), and assume a steady decline over time; it is common to assume declining backstop prices due to technological progress. Note that a reasonable first guess might be to assume that the backstop price falls at the same rate as the resource cost. If we impose this further restriction, then we can instead plot level curves for Hotelling rent as a function of *current* backstop price and rate of return. The result (not shown graphically) is that the level curves shift down significantly, because a high current scarcity rent implies rapidly falling resource costs, implying that the backstop price is also likely to fall. Specifically, a current backstop price of over \$500/barrel is required to yield a 25 percent Hotelling

rent assuming 7 percent rate of return, while if the backstop price is limited to a more reasonable $200/barrel then the current Hotelling rent is limited to 10 percent of the price.

The point of the above is that the sum of something increasing exponentially (and at a rather high rate) and something else decreasing must also start to increase within a rather short time unless the increasing part start off very very low. Since resource prices generally are not showing a clear upward trend this suggests strongly that the resource rent remains a small part of the price.

CHAPTER 5

Directed technological change and resource efficiency

In this chapter we start the process of generalizing the Cobb–Douglas production function in order to analyse the role of resources in aggregate production. Following Solow (1973), we aim to (i) capture the effect of technological progress which allows more efficient use of natural resources, (ii) include the possibility of substituting between alternative natural resources in production, and (iii) allow for the possibility of adapting consumption patterns to save on the use of natural resources. We begin with technological progress, related to Solow's first mechanism, i.e. that firms may adapt to resource scarcity by increasing—through investment in technological change—the efficiency with which they use resources in the production of their goods.

5.1. Resource efficiency in the production function

In this section we take a more detailed look at alternative ways of representing production as a function of two inputs. This is important groundwork before beginning our analysis of DTC (directed technological change). We focus on a simple case in which widgets are produced using flows of resources and labour. We abstract from (i.e. ignore) capital. In essence this implies that we assume that the capital stock grows at the same rate as the stock of effective labour $A_L L$, hence including capital would not significantly affect the analysis.

5.1.1. Familiar cases: the general production function, and Cobb–Douglas.
Assume an economy with one product, widgets, and two production inputs. For now we call them labour and resources. Denote flows of labour and resources as L and R respectively, units workers and tons/year. In general we can write

$$y = F(A_L L, A_R R). \tag{5.1}$$

Recall then that the units of L and R are widgets per worker per year and widgets per ton per year.

So far in the book we have used the Cobb–Douglas function almost exclusively. That is, we write the production function as

$$y = (A_L L)^\alpha (A_R R)^{1-\alpha} = A L^\alpha R^{1-\alpha}.$$

Thus the two separate factor-augmenting technology levels have been subsumed into a single technology index and there is no role for directed technological change. For instance, if we raise resource-efficiency A_R this serves to raise A and hence raise production, but it does not reduce the demand for resource inputs. This paradoxical result follows from the high degree of substitutability between the inputs: when resource efficiency rises the cost of resource inputs falls, causing producers to use more resources.

The Cobb–Douglas model fails when confronted by further evidence. For instance, we know that it is in fact very hard for producers to substitute for resource inputs using labour, at least in the short run. Intuitively, we can think of a producer making hammers from steel. More broadly, this is clear from the very small short-run effect of increases in the resource price on resource quantity (i.e. the inelastic short-run demand): when resource prices rise dramatically, quantity hardly falls because firms and consumers are locked in to their demand for resources by the technology and capital they possess. To go deeper we therefore need a model in which there is low short-run elasticity of substitution between labour and resources, but where there is also some mechanism through which they can be substituted in the long run. In this chapter that mechanism is DTC.

5.1.2. The Leontief production function and 'Limits to growth'.
The simplest way to specify a production function with low substitutability between the inputs is through the *Leontief* production function, in which there is no substitutability whatsoever between the inputs. The function looks like this:

$$Y = \min\{A_L L, A_R R\}. \tag{5.2}$$

This equation reads as follows: production Y is equal to the smallest of the following set of quantities: effective labour inputs $A_L L$, and effective resource inputs $A_R R$.

To interpret the equation, consider for instance production of hammers from steel. Given the design of the hammers (which determines A_R), and assuming no steel is wasted, the production rate of hammers will be limited by the flow of steel inputs R. If the design is such that each hammer requires 1 kilo of steel, and the flow of steel inputs is 10^3 tons per year, then no more than 10^6 hammers can be produced by the firm per year. However, there is no guarantee that the firm will produce that many hammers; to do so, they must have enough (productive) labour employed, where A_L is the productivity of the workers and L is their number. Note that we have—following Solow (1973)—left out capital; effectively, we assume that each worker has enough capital (machines) in order to work productively.

A profit-maximizing firm will of course make sure that it hires just enough workers, and buys just enough resources, such that $A_L L = A_R R$. The alternative is that workers stand idle waiting for resource inputs, or that resources stand idle waiting for workers to use them.

Now assume that A_L—worker productivity—rises exponentially, while A_R remains unchanged. Furthermore, assume that equation (5.2) is the aggregate production function of the economy. Then assuming that resources are available, resource flows into production will also increase exponentially, tracking the growth in labour productivity and global product. If the resource price remains constant then the share of resources in global product will also remain constant. However, such exponential growth in physical inputs cannot continue indefinitely, and the model strongly suggests that a crash is imminent when resources 'run out'.

Superficially this example matches the evidence regarding the global economy (see for instance Figure 4.4). However, it should be clear that this model is far too simple to use as a basis for drawing conclusions and designing policy. The model is obviously grossly simplified in assuming that it is impossible to increase resource efficiency. If labour efficiency can increase exponentially, why should resource efficiency be unable to increase likewise? Furthermore, as in the neoclassical model, there is no allowance for shifts in consumption patterns. Finally, there is no model of the price of the resource input, which we expect to be linked to scarcity of the input and to affect demand for the input.

The model above shares crucial features with the famous 'Limits to Growth' model. The 'Limits' model dates to 1972 when the Club of Rome published a book, the Limits to Growth, which caused a sensation at the time. The analysis used a "system dynamics computer model to simulate the interactions of five global economic subsystems, namely: population, food production, industrial production, pollution, and consumption of non-renewable natural resources."[1]

The 'Limits' team programmed various scenarios, and all ended in disaster, typically by around 2050. This was either due to resource exhaustion or excessive pollution. The reasons for this are not easy to elucidate as the model is not open to examination, however, the behaviour of the model can be approximated by a very simple economic model, as follows:

$$Y = \min\{A_L L, A_R R\}; \tag{5.3}$$

$$\dot{A}_L / A_L = g; \tag{5.4}$$

$$S_0 \geq \int_0^\infty R_t \mathrm{d}t. \tag{5.5}$$

Furthermore, we have that $A_{L0} L < A_R S_0$. Thus we have a Leontief production function with exogenous growth in A_L. If A_R is constant then we get exponential growth in both Y and R, whereas if A_R is allowed to grow then the growth in R will be slower.[2] But why does A_L grow? And what determines the growth rate of A_R? And what about the resource price, does that have no effect whatsoever on the allocation? The model raises many questions, some of which we try to answer in this part of the book.

5.1.3. The CES production function. An alternative functional form—much more flexible than both Cobb–Douglas and Leontief—is the CES production function:

$$y = [\gamma(A_L L)^\epsilon + (1 - \gamma)(A_R R)^\epsilon]^{1/\epsilon}.$$

Here $\epsilon \in (-\infty, 1)$. The Cobb–Douglas emerges from the CES as a special case when $\epsilon = 0$, and Leontief when $\epsilon \to -\infty$. Finally, when $\epsilon = 1$ then the inputs are perfect substitutes, like 5 and 10-dollar bills. However, typically we assume that labour and capital, or labour and resources, are

---

[1] According to the website http://www.clubofrome.org/?p=1161+, 3 Oct. 2012.

[2] Note that we have simplified slightly here. The Limits model is actually built on an outdated growth model known as the AK model in which labour productivity A_L grows due to capital accumulation, which is possible due to a highly contrived mechanism where the greater the quantity of capital possessed by other firms, the more productive is any particular firm (irrespective of how much capital that firm has). However, the overall effect is consistent with the equations presented here.

poorly substitutable for one another, that is they are *complements*, implying that $\epsilon < 0$. When they are complements, an increase in the quantity of one of the inputs available on the market leads to a large reduction in its price such that relative returns to that input factor decline.

To get a handle on the production function, consider the following example. Assume an economy with 10 people on an island, and 10 trees/week wash up on the beach. Furthermore, the islanders have a technology called 'knives' which allows them to cut the trees into planks, which can rapidly be made into houses (final product). They manage to make 0.01 houses per week. What do they need more of to boost their rate of housebuilding? Suggest values for the elasticity of substitution between the inputs, and knowledge levels. Explain briefly.

Now assume that the islanders invent a technology called 'sawmills' (and are somehow able to obtain the necessary capital goods). What do they need more of now in order to boost their rate of housebuilding? Suggest new values for the knowledge levels. Explain briefly.

Some possible answers to the questions above follow. We have $L = 10$ and $R = 10$, and $y = 0.01$. Presumably it is rather hard to substitute workers for trees, although not impossible. For instance, if there are a lot of trees the (fixed number of) workers could select the best trees to make planks out of, rejecting less suitable ones. This would allow them to produce planks faster. So ρ should definitely be negative, but we have to guess its value. For simplicity we choose $\rho = -1$. We drop the distribution parameter γ, incorporating it into the productivity factors. So we have

$$y = \left(\frac{1}{10A_L} + \frac{1}{10A_R} \right)^{-1} .$$

Since $y = 0.01$ we can rearrange the equation to obtain $1000 = 1/A_L + 1/A_R$. Finally, and crucially, note that it is clearly workers who are in short supply in this economy, not trees; only a small fraction of the 10 trees per week can be used considering that the 10 workers only have knives with which to work. Since the supply of workers is limiting, this implies that A_L should be small compared to A_R.

More precisely, imagine that there is an abundance of trees such that $A_R R$ is very large. Then we can approximate $1/(A_R R) = 0$, and the production function becomes $y = A_L L$. Then $A_L = 0.001$ (since $y = 0.01$ and $L = 10$). If we instead imagine that labour is abundant, then A_R becomes the minimum number of houses that can be made per tree, given that infinite care is taken to avoid waste in making the planks. If this is 0.5 then the production function is

$$y = \left(\frac{1}{0.001L} + \frac{1}{0.5R} \right)^{-1} .$$

When L and R are both equal to 10, this gives 0.01 houses produced per week.[3]

When sawmills are invented, the productivity of trees in housemaking is presumably more-or-less unchanged; it still takes 2 trees to make a house. (Although we could imagine A_R changing: for instance, if the sawmills produced more waste such as sawdust then A_R would go down, whereas if they could cut thinner planks thus using the timber more efficiently then it would go up.) The productivity of labour on the other hand will go up enormously. Now we might imagine 10 people running a sawmill being able to cut up hundreds of trees each week. (Recall that the planks can 'quickly' be made into houses, implying that the time spent on this step can be ignored.) If we assume a relatively rudimentary sawmill then we can set $A_L = 10$, and the new production function is

$$y = \left(\frac{1}{10L} + \frac{1}{0.5R} \right)^{-1} .$$

Now trees are the limiting factor, and 4.8 houses can be made each week given that there are 10 trees available per week and 10 workers.

The CES production function gives us the ability to capture the effects of factor-specific (directed) technological change on production in a flexible way. This is an essential ingredient in our models of growth and sustainability. However, this is not enough on its own; we must also be able to model changes in the relative productivities of the factors as an endogenous result of other changes in the economy, such as changes in the availability of the factors. That is, we need a model of endogenous directed technological change, henceforth DTC. To build such a model, we return to our model of endogenous technological change from Chapter 3 and add the need for a resource input, with an associated level of resource-augmenting knowledge. Furthermore, we simplify the model in some other respects. Because we are primarily interested in the direction of

---

[3]More exactly, 0.00998 houses per week.

technological change rather than overall growth we simply fix the total quantity of research labour in the economy.

5.2. Foundations of an aggregate DTC model

Deriving a full DTC model for use in analysis of environmental policy questions — see for instance Hart (2018b) — is beyond the scope of this book. Here we instead derive a key general result regarding the behaviour of a small innovating firm — following Hart (2013) — and then use it in straightforward aggregate models with representative firms.

Consider a small firm i (with many competitors) buying inputs of labour l_i and a resource r_i, prices w_L and w_R. The firm is a price taker with respect to inputs. It makes a unique product y_i, and has unique firm-level productivity levels A_{li} and A_{ri}, which are achieved through investments z_{li} and z_{ri} in research (price w_Z). The elasticity of demand for the firm's product is $1 - \eta$, so

$$\frac{\partial p_i}{\partial y_i} \frac{y_i}{p_i} = -(1 - \eta).$$

The firm's production functions for knowledge and the final good are illustrated in Figure 5.1. Note that ϕ is a positive parameter, $\mathbf{A}$ is all existing knowledge in the economy, and F_r and F_l are increasing functions.[4] So firm knowledge in each sector is an increasing function of prior knowledge in the economy and its own investments. And final-good production is a CES function of the augmented inputs of labour and the resource. How much should the firm invest in the research inputs z_l and z_r?
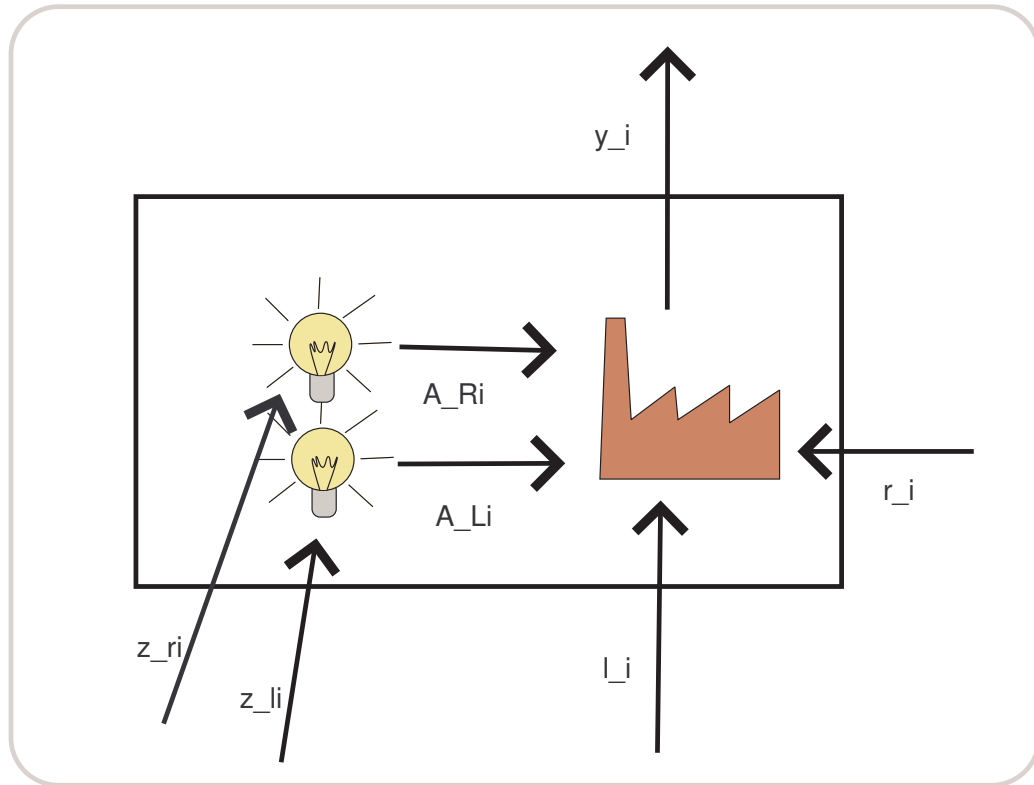

FIGURE 5.1. The single firm's production function

The firm has a straightforward, static problem, which is to maximize net revenue subject to the restrictions implied by the knowledge production functions. We can thus write down the firm's problem as a Lagrangian:

$$\mathcal{L} = p_i(y_i)[(A_{li}l_i)^\epsilon + (A_{ri}r_i)^\epsilon]^{1/\epsilon} - w_Z(z_{li} + z_{ri}) - (w_L l_i + w_R r_i)$$
$$- \lambda_{li}(A_{li} - F_l(\mathbf{A}) \cdot z_{li}^\phi) - \lambda_{ri}(A_{ri} - F_r(\mathbf{A}) \cdot z_{ri}^\phi),$$

where λ_{li} and λ_{ri} are the shadow prices of firm knowledge.

---

[4]An important feature of the production function is that firms do not build their period-t knowledge on their existing private knowledge, rather they build on all existing knowledge in the entire economy. This simplifies the dynamics of the model immensely.

To solve the problem we take first-order conditions in l_i and r_i to yield an expression for $w_L l/(w_R r)$:

$$\frac{w_L l_i}{w_R r_i} = \left(\frac{A_{li} l_i}{A_{ri} r_i}\right)^\epsilon .$$

This equation relates relative factor expenditures by the firm to the relative augmented quantities of the factors. (The augmented quantity is the physical quantity multiplied by the productivity.) Since ϵ is negative, the expression shows that the firm spends more on the factor which is relatively scarce.

Now we take first-order conditions in A_{li} and A_{ri} to yield an expression for $\lambda_{li} A_{li}/(\lambda_{ri} A_{ri})$:

$$\frac{\lambda_{li} A_{li}}{\lambda_{ri} A_{ri}} = \left(\frac{A_{li} l_i}{A_{ri} r_i}\right)^\epsilon .$$

We can actually write down this expression without the need for working through the first-order conditions, using the symmetry implicit in the Lagrangian, where the λs play the role of the ws, and the ks play the role of the qs.

Finally, we take first-order conditions in z_{li} and z_{ri} to yield a second expression for $\lambda_{li} A_{li}/(\lambda_{ri} A_{ri})$:

$$\frac{z_{li}}{z_{ri}} = \frac{\lambda_{li} A_{li}}{\lambda_{ri} A_{ri}}.$$

This simply states that firms invest in proportion to the value of the knowledge produced, which is the shadow price of that knowledge times its quantity.

Now we can eliminate the shadow prices from the two expressions above to yield an expression for relative investments as a function of prices and quantities of the inputs:

$$\frac{z_{li}}{z_{ri}} = \left(\frac{A_{li} l_i}{A_{ri} r_i}\right)^\epsilon = \frac{w_L l_i}{w_R r_i}. \tag{5.6}$$

So firms invest in proportion to the resultant relative factor expenditures. If a firm spends twice as much on labour as it does on resources, it will invest twice as much in increasing the productivity of labour compared to resources!

Note that we could also derive expressions for the quantities of labour and resources bought by the firm, the absolute quantities of investment in research, etc. However, our key focus is relative investment, so we stop at equation 5.6.

5.3. The aggregate DTC model

5.3.1. The model. We now take equation 5.6, and feed it into a simple aggregate model, illustrated in Figure 5.2. Here we see that aggregate production Y is divided between consumption C and production (or extraction) of the intermediate energy input R. There are two types of labour, L and Z, the quantities of which are exogenously given. L is production labour, and Z is research labour. Research labour may be divided between boosting labour productivity A_L and energy productivity A_R. Furthermore, the aggregate production function for the resource input R is

$$R = A_X X,$$

where A_X is the productivity of the input, and w_L is the wage. Since the price of the final good is normalized to one this gives total costs of resource production as X, and hence (assuming a perfect market) the resource price $w_R = 1/A_X$.

In the aggregate model consumers have the very simple utility function

$$U = \sum_{t=0}^\infty \beta^t C_t,$$

where β is a parameter less than 1, and C is aggregate consumption. So we have a constant, exogenous discount rate. Furthermore, we have

$$Y = \left(\int_0^1 y_i^\eta \, \mathrm{d}i\right)^{1/\eta} .$$

where η is a parameter less than 1. Thus there is monopolistic competition between the producers of the different goods. Normalize the price of the aggregate to 1. Then (by differentiating) we can obtain

$$p_{yi} = \frac{\partial Y}{\partial y_i} = \left(\frac{Y}{y_i}\right)^{1-\eta} .$$
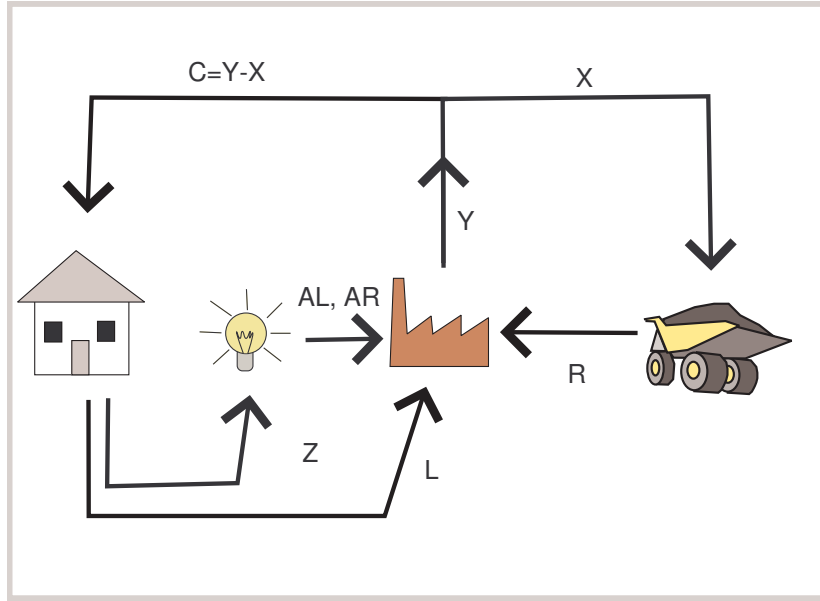
FIGURE 5.2. The aggregate flows of factors and products in the economy.

Given symmetry $Y = y_i$ so all the goods have price 1. But crucially, we see that the elasticity of demand for each good is $1 - \eta$, as in the firm model analysed above. Therefore equation 5.6 applies, and we have (in the aggregate)

$$\frac{Z_L}{Z_R} = \left(\frac{A_L L}{A_R R}\right)^{\epsilon} = \frac{w_L L}{w_R R}. \tag{5.7}$$

Finally, and crucially, we link firm knowledge A_{li} and A_{ri} to next-period economy-wide knowledge $\mathbf{A}$, and especially the functions F_l and F_r. First we must specify the relationship between the economy-wide knowledge A and the firm-specific knowledge stocks A_{li} and A_{ri}. Here we divide the economy-wide knowledge stocks into two sets, A_L and A_R, which are the sum of the individual firms' knowledge stocks:

$$A_L = \int_0^1 A_{li} \mathrm{d}i;$$

$$A_R = \int_0^1 A_{ri} \mathrm{d}i.$$

So in symmetric equilibrium with a representative firm (for which we drop the subscript is) we have $A_L = A_l$, and $A_R = A_r$. Then we must specify the functions F_l and F_r. We choose the simplest possible specification, an extreme case, as explained below.

DEFINITION 1. Independent knowledge stocks. *Knowledge stocks develop independently when*

$$F_l = A_{lt-1}/\zeta_L \quad and$$
$$F_r = A_{rt-1}/\zeta_R,$$

where ζ_L and ζ_R are positive parameters. Thus labour-augmenting knowledge builds exclusively on existing labour-augmenting knowledge, and resource-augmenting knowledge builds exclusively on existing resource-augmenting knowledge.

Thus — recalling the knowledge production functions above — we can write

$$A_{Lt} = (A_{Lt-1}/\zeta_L)Z_{Lt}^{\phi},$$

$$A_{Rt} = (A_{Rt-1}/\zeta_R)Z_{Rt}^{\phi},$$

and
$$\frac{A_{Lt}}{A_{Rt}} = \frac{A_{Lt-1}}{A_{Rt-1}} \frac{\zeta_R}{\zeta_L} \left(\frac{Z_{Lt}}{Z_{Rt}}\right)^{\phi}. \tag{5.8}$$

5.3.2. The solution. Now to finalize the solution. Take equation 5.8, rearrange, and substitute for relative investments using equation 5.7 to obtain

$$\left(\frac{A_{Lt}/A_{Lt-1}}{A_{Rt}/A_{Rt-1}} \frac{\zeta_R}{\zeta_L}\right)^{1/\phi} = \frac{w_{Lt} L_t}{w_{Rt} R_t}.$$

We can then use this equation to obtain a single equation for the development of relative knowledge stocks either as a function of the relative quantities of the inputs, or the relative prices. We assume that the relative prices of the inputs are exogenously given, i.e. w_L/w_R is given in each period, hence we want an equation in terms of relative prices. After some simple algebra we have

$$\frac{A_{Lt}/A_{Lt-1}}{A_{Rt}/A_{Rt-1}} = \left(\frac{A_{Lt-1}}{A_{Rt-1}} \cdot \frac{w_{Rt}}{w_{Lt}}\right)^{\epsilon\phi/(1-\epsilon(1+\phi))} \left(\frac{\zeta_L}{\zeta_R}\right)^{(1-\epsilon)/(1-\epsilon(1+\phi))}. \tag{5.9}$$

Equation 5.9 gives us the period-by-period development of the state of the economy. And this allows us to draw conclusions about balanced growth paths. By definition, on a b.g.p. all state variables must change at constant rates. If both A_L and A_R change at constant rates then the LHS of equation 5.9 is constant, since it is one constant divided by another. This implies in turn that the RHS is constant, so if w_R/w_L is decreasing at some fixed rate then A_L/A_R must be increasing at the same rate. So if the relative price of the resource is declining at some fixed rate, the relative productivity of the resource must decline at the same rate. This implies in turn, from equation 5.7, that relative quantities will increase at the same rate. To show this, start with

$$\frac{w_L L}{w_R R} = \left(\frac{A_L L}{A_R R}\right)^{\epsilon},$$

and rearrange to show that

$$\frac{L}{R} = \left(\frac{A_L}{A_R}\right)^{\epsilon/(1-\epsilon)} \left(\frac{w_L}{w_R}\right)^{-1/(1-\epsilon)}.$$

Hence relative investments are constant on the b.g.p.

5.3.3. Stability of the b.g.p. Whatever the values of parameters, a b.g.p. will exist. But will it be stable? In other words, will the economy approach the b.g.p. over time, or will it head off somewhere else? It turns out that the elasticity of substitution between the inputs is key.

The key is the result that relative investments are equal to relative factor shares. Assume that we are on a b.g.p. such that constant relative investments lead to constant factor shares. Then assume that there is some perturbation to the system such that the ratio of A_L to A_R moves off the balanced growth path. Does this trigger the economy to move away from the path, or does it return to the path? For concreteness, assume that A_R falls. Then we know from equation 5.6 that the factor share of the resource will rise (since ϵ is negative) causing investment in A_R to rise, thus boosting the growth of A_R. This shows that the b.g.p. of the economy is stable, as long as ϵ is negative; when $\epsilon > 0$ the analysis and results change completely, as we see in Chapter 7. The situation is illustrated in Figure 5.3

FIGURE 5.3. Illustration of how relative prices (the shape of the economic landscape) determine the relative levels of technology augmenting labour and energy.

5.3.4. The long-run aggregate production function. We have established that if w_R/w_L changes at a constant rate, the economy will approach a b.g.p. on which A_R/A_L changes at the same rate, while R/L has exactly the opposite trend. Furthermore, these results imply that both augmented inputs and the factor share are constant. So the long-run factor shares are constant despite changes in the relative quantities of the inputs. This implies that the long-run aggregate production function is, in reduced form, Cobb–Douglas:

$$Y = AL^{1-\alpha}R^{\alpha}.$$

So we are back to the production function of the DHSS model! (Although without capital.) And therefore the model can explain the data of Figure 4.4 for metals and energy, reproduced here as

Figure 5.4. Note that the result of constant prices and consumption tracking global product is that the *share* of resources in global product is also constant.[5]


FIGURE 5.4. Long-run growth in consumption and prices, compared to growth in global product, for (a) Metals, and (b) Primary energy from combustion.

The value of α will be determined by the relative values of ζ_L and ζ_R, and the relative growth rates of input quantities (or prices). If it is easier (cheaper) to develop technologies augmenting resources than it is to develop technologies augmenting labour then α will be small and the factor share of resources will also be small. Just as in the DHSS model, sustainability should be no problem as long as we manage resource stocks sensibly: thus we should gradually reduce R if there is a finite stock of the resource.

The key difference compared to the 'Limits' model is that A_R, resource-augmenting knowledge, is allowed to grow exponentially and without bound. This allows sustainable use of a finite resource stock even when the resource and other inputs (in this case labour) are highly complementary in the short run. Given such complementarity, and *without* DTC, we are very close to the limits model, and there is no way to achieve sustainability (if defined as any consumption level that can be maintained into the indefinite future).

Recall Daly's critique of the neoclassical model (page 57): can we bake more and more cake with the same quantity of ingredients? Can we, for example, produce more and more light using the same energy inputs? Can we travel further and further using the same energy inputs? Can we keep our houses warm...? Etc. It turns out that we can! Consider for instance Fouquet and Pearson (2006) on the history of light production. They conclude that the efficiency of lighting in the U.K. (measured by lumen produced per watt of energy used) increased 1000-fold from 1800 to 2000. That's a lot more cake! Regarding the production of motive power from fossil fuels, historically this concerns the efficiency of steam engines, while over the last century we must consider electric power generation and the internal combustion engine. Regarding steam engines, sources such as Hills (1993) suggest that their efficiency in generating power from coal inputs increased steadily from their invention in the early 1700s up to 1900, and by a factor of around 20 over the entire period; this growth in efficiency is at least equal to the growth in labour productivity over the same period. Subsequently, the efficiency of coal-fired power stations has continued to increase but at a declining rate; see for instance Yeh and Rubin (2007) for detailed evidence.

Does the above discussion mean Daly was wrong in using his cake analogy to criticize DHSS? In fact it does not, since in the DHSS model it is capital accumulation which allows us to bake more cake from the same ingredients, but in the above examples it is *technological change* that has allowed us to do so, not capital accumulation. Technological change does, demonstrably, allow us to get more product (value) out of given inputs, and there is no obvious limit to this process.[6] But does the model stand up to more detailed examination? We find out in the next section.

---

[5]The figures show normalized prices, quantities, etc., so they show how the factor shares of resources change over time, but nothing about the absolute levels of the factor costs compared to the value of global product. The absolute share of resources is significant but not large. For instance, the factor share of crude oil in the global economy in 2008 was 3.6 percent, whereas the factor share of the 17 major metals was just 0.7 percent.

[6]Note however that in some cases there *are* obvious limits, as with the efficiency of the production of artificial light and electricity, where there are thermodynamic limits to what can be produced from given inputs, and these limits are currently being approached.

5.4. Problems regarding the DTC model

5.4.1. Predicted growth in A_R. We have shown that our simple model can reproduce the patterns observed in the data for resource or energy demand, given the development of prices. However, the model also makes predictions about the development of energy-augmenting knowledge, and here we find that these predictions are completely at odds with the evidence.

Recall that the original production function is CES with low elasticity of substitution (labour and the resource are strongly complementary), thus we have

$$Y = [(A_L L)^\epsilon + (A_R R)^\epsilon]^{1/\epsilon},$$

where $\epsilon < 0$. Assuming perfect markets we can model the choices of a representative final-good producer whose profit-maximization problem is

$$\max_{L,R} \pi = [(A_L L)^\epsilon + (A_R R)^\epsilon]^{1/\epsilon} - w_L L - w_R R.$$

Take the first-order condition in R to show that

$$\frac{R}{Y} = \left(\frac{A_R}{w_R}\right)^{1/(1-\epsilon)} \frac{1}{A_R}.$$

Now differentiate with respect to time to show that

$$\frac{\dot{R}}{R} = \frac{\dot{Y}}{Y} - \frac{-\epsilon}{1-\epsilon}\frac{\dot{A}_R}{A_R} - \frac{1}{1-\epsilon}\frac{\dot{w}_R}{w_R}.$$

Now the historical data shows that w_R is approximately constant in the long run, while resource use tracks global product. Since ϵ is large and negative (the inputs are strongly complementary) this implies that $\dot{A}_R/A_R = 0$, i.e. there has been no resource-augmenting technological change. This is exactly what we see when we simulate the model numerically based on real data, as shown in Figure 5.5.[7] Here we see that we can parameterize the model outlined above such that when we feed in data on global energy prices the model predicts rates of energy use which match observations. The bottom panel shows what lies behind the model results: global energy use has tracked global product because the level of energy-augmenting knowledge A_R has failed to rise.

The prediction that when w_R does not rise, A_R will not rise is strongly contradicted by the evidence, especially in the case of energy, as we now show.

Estimation of productivity changes is typically performed indirectly: a productivity increase is imputed as the residual to explain changes in the value of output per unit of time from given inputs. However, in the case of energy we can use a more direct approach to measurement of factor-augmenting knowledge, since there is plenty of direct evidence about changes in our ability to extract specific physical outputs from measurable energy inputs. To illustrate we consider two products, artificial light and motive power. Light is a convenient product category for analysis since light is a consumption good which is rather homogeneous and unchanging over very long timescales, and the energy efficiency of its production is easily measured. Fouquet and Pearson (2006) study light production and consumption in the U.K. over seven centuries. They conclude that the efficiency of light production in the U.K. (measured by lumen produced per watt of energy used) increased 1000-fold from 1800 to 2000; the productivity of labour in the U.K. over the same period rose by a factor of 12–15 (estimates vary). Light production is a convenient sector within which to measure efficiency, but it is not very large. Now we turn to the production of motive power from fossil fuels, a very large sector. Motive power is typically an intermediate good rather than a final good, nevertheless increases in the efficiency with which energy inputs are used to generate motive power are very likely to be reflected in the overall efficiency with which energy is used to generate final goods, as long as the final goods are homogeneous and do not change over time. In the 19th century motive power was largely generated by steam engines, while over the last 100 years we must consider electric power generation and the internal combustion engine. Regarding steam engines, sources such as Hills (1993) suggest that their efficiency in generating power from coal inputs increased steadily from their invention in the early 1700s up to 1900, and by a factor of around 20 over the entire period; this growth in efficiency is again more rapid than the growth in labour productivity over the same period. Subsequently, the efficiency of coal-fired power stations

---

[7]When prices are exogenous the key equation is 5.7. Put into simulation-friendly form, the equation is

$$A_{lt}/A_{rt} = A_{lt-1}/A_{rt-1}\left((A_{lt-1}/A_{rt-1})(w_{Rt}/w_{Lt})\right)^{\epsilon\phi/(1-\epsilon(1+\phi))}(\zeta_L/\zeta_R)^{(1-\epsilon)/(1-\epsilon(1+\phi))}.$$

In simpler notation we can write

$$A_t = A_{t-1}(A_{t-1}/W_t)^{\epsilon\phi/(1-\epsilon(1+\phi))}\zeta^{(1-\epsilon)/(1-\epsilon(1+\phi))},$$

where $\zeta = \zeta_L/\zeta_R$.

FIGURE 5.5. The model parameterized to fit global data. The top row shows GDP
(model and data); the middle row shows energy consumption (model and data)
and energy price; the bottom row shows the growth of the knowledge stocks. All
axes are log-normalized.

has continued to increase but at a declining rate; see for instance Yeh and Rubin (2007) for detailed
evidence.

The above evidence suggests that energy productivity over this period, far from being constant,
actually increased faster than labour productivity. So the model with independent knowledge
stocks matches energy demand at the expense of wildly incorrect predictions about the growth
of energy-augmenting knowledge. On the other hand, if we lock knowledge stocks together this
comes closer to the truth in its predictions about knowledge growth, but at the expense of no longer
being able to predict energy demand correctly. And whatever we assume about knowledge stocks,
the model cannot match the data. For an analysis of what is reasonable to assume about links
between knowledge stocks, see the next chapter on technology transitions.

5.4.2. Links between knowledge stocks. A further problem with the model is the assumption
that knowledge stocks grow independently of one another. That is, a higher level of labour-
augmenting knowledge (or overall productivity in the economy) does not help at all when it comes
to performing research to raise resource-augmenting knowledge. Is this reasonable? Arguments
that it is not date at least to Nordhaus (1973b), however Norhaus's arguments appeal only to
intuition, and it would be reassuring if we had more direct microeconomic evidence.

The evidence based on intuition is nevertheless powerful. Consider the following thought
experiment. Assume there was no generation of power from wind between 1900 (when the last
windmills were decommissioned) and 1990 (when electricity generation from wind started). On
what knowledge stock would the new wind power generators build? More broadly, is the idea of
independent knowledge stocks defensible? It implies that technologies which allow us to make
better use of raw energy inputs (the steam engine, the steam turbine for the generation of electricity,
the internal combustion engine) are developed completely independently of other technological
advances in the economy. This seems to be an indefensible idea: such technologies are developed
hand-in-hand with advances in mathematics, physics, engineering etc., advances which are also
relevant to augmenting inputs of labour–capital. Thus stocks of knowledge augmenting energy
and stocks of knowledge augmenting labour–capital are intimately linked, overlapping and feeding
off one another. So, summing up, it is hardly surprising that the model fails when confronted by
evidence.

There is also direct evidence that (for instance) both resource-augmenting and labour-augmenting
knowledge draw on a common pool of general or fundamental knowledge, and that advances in this

general knowledge thus drive advances both in a_l and a_r? Such evidence can be found in studies of patent data. Very direct evidence can be found in Trajtenberg et al. (1992), who study patent citations and show that patenting firms cite patents both within their own three-digit industry, but also outside it.[8]

Popp (2002) provides more indirect—but no less convincing—evidence, in that he finds evidence for diminishing returns to investment in specific technologies over time; a rise in energy-prices induces a surge in patenting activity within energy sectors such as wind or solar, but the effect peters out rather rapidly. Within our framework we can interpret this result as follows: Discoveries of potential relevance to energy-augmentation are made frequently in other (much larger) research sectors; For instance, think of the invention of the computer. It takes time and research effort for the benefits of these discoveries to be incorporated in the energy sector. If there is a surge of research in the energy sector then initially there will be many potentially useful technologies available which have not yet been applied in that sector, but over time these 'low-hanging fruits' will be picked and the productivity of energy-augmenting research will fall.

Developing and solving a model with links between knowledge stocks is beyond the scope of this book. We simply note that it would be more reasonable to assume that knowledge stocks are linked, such that if (for instance) there is a lot of labour-augmenting knowledge this makes it easier to accumulate resource-augmenting knowledge. Given a model with links we expect the knowledge stocks to grow together (although not necessarily at exactly the same rate). What would be the results of such a change in the model?

The result of linking knowledge stocks would be to yield a better fit to the data on input productivity, at the expense of the model's ability to match the data on factor shares: if knowledge stocks grow together while the price of resources falls relative to the price of labour, then we expect the resource share to fall over time. Consider the simple case when knowledge stocks are locked together: $A_r = \gamma A_l$. Then the aggregate production function is

$$y = A_l[L^\epsilon + (\gamma R)^\epsilon]^{1/\epsilon}.$$

If the inputs are strongly complementary (ϵ is large and negative) this implies that the inputs will be bought by firms in almost fixed proportions, irrespective of their relative prices. Thus resource use will track the size of the labour force rather than overall growth. This is directly contradicted by the evidence, for instance the data shown in Figure 4.4.

To be more precise, assume perfect markets and take first-order conditions in L and R to show that

$$\frac{w_L L}{w_R R} = \left(\frac{L}{R}\right)^\epsilon$$
$$= \left(\frac{w_L}{w_R}\right)^{-\epsilon/(1-\epsilon)}.$$

Since ϵ is large and negative this shows that when resource inputs rise relative to labour inputs the resource share should decline steeply, and that when the resource price falls relative to the wage the resource share should decline steeply.

---

[8]They score patents as follows: within 3-digit scores zero, within 2-digit scores 0.33, within 1-digit scores 0.67 and outside 1-digit scores 1. The average score is 0.31. A three-digit industry is a relatively narrowly defined industrial sector, according to the Standard Industrial Classification. Two- and one-digit industries are successively more broadly defined.

Structural change

In all of our models hitherto we have either had a single product (e.g. a widget), or we have had multiple products which all have the same properties with regard to the need for natural-resource inputs. Hence — in Chapter 5 — we drew the conclusion that if resource use tracks growth it must be because resource-augmenting knowledge has not risen. In order to add alternative explanations we must include multiple final goods, the production of which differs in resource intensity. Given multiple goods resource consumption may track growth even when resource efficiency increases, if consumers switch towards resource-intensive products. This switch may be an endogenous result of increase in resource efficiency, in which case it is called *rebound*. Alternatively, the switch may be caused by other factors, such as income growth. This question is particularly relevant for energy demand, which is strongly linked to fossil-fuel demand. In this chapter we look first at some of the broad evidence which demonstrates that structural change must be central to the analysis of long-run demand for energy and resources, before turning to the key questions of what drives structural change in the energy sector and what the policy implications are. It builds to a large extent of Hart (2018a).

6.1. Introduction to structural change

FIGURE 6.1. Long-run growth in real GDP per capita: U.S. (upper line) and China.

6.1.1. Growth and structural change. The growth rate of GDP per capita in the leading industrial economy has been astonishingly constant in the industrial era. On the other hand, growth rates of other economies have varied enormously; Figure 6.1.[1] The reason is that growth in per capita production is primarily driven by new technology: leading-edge technology seems to grow at a remarkably constant rate, whereas the distance of an economy from the leading edge may vary dramatically over time. The U.S. has remained at or close to the leading edge for well over 100 years, hence its constant growth rate. China on the other hand moved further and further away from the leading edge for many decade, before this trend was turned upside down starting in the 1970s.

---

[1]Jones (2005) pointed this out. Data from Angus Maddison website, Statistics on World Population.

Note that growth—driven by new technology—is primarily about production of different, more valuable, final goods over time. Only a small proportion of our expenditure today goes on goods that were available in the same form 50 years ago. And the proportion becomes very small if we go back 100 or 150 years. Consider food. One hundred years ago food made up a large part of household budgets, everywhere. However, as we have become richer and richer our expenditure on food has changed relatively little, whereas our expenditure on other goods has risen.

Figure 6.1 shows that US GDP/capita is more than 10 times higher today than in 1870; is that growth due to U.S. citizens today consuming 10 times more of the same products today that they consumed in 1870? Consider car ownership. In 1870 there were no cars. Car ownership subsequently grew rapidly, but between 1970 and 2008 it was constant at 0.44 per capita; in the meantime, GDP per capita had more than doubled.[2] Meanwhile, ownership of home computers and mobile phones was effectively zero in 1970, whereas today it is more-or-less universal. Now, do we consume cars and smartphones today because we work longer hours, or have saved up more capital, but with the same skills and the same machines as we had in 1870? Obviously not: it is the arrival of new products, and increases in the quality of existing products, which is the major and fundamental driver of long-run growth.

Finally, note that the transformation of the economy is not simply about the addition of new products, made in new ways, i.e. increasing variety. It is better characterized as a process in which new ideas—new technologies—transform existing processes and products, as well as adding completely new ones. The most important of these new technologies are sometimes denoted *general purpose technologies*, or GPTs; see David (1990) and Bresnahan and Trajtenberg (1995) for classic discussions of such technologies. Examples include the steam engine, the electric motor, semiconductors, and the internet. Such inventions lead, over time, to fundamental changes throughout many areas of the economy, or indeed throughout the entire economy, transforming the way many existing goods or services are produced, and opening up previously unimaginable possibilities for new goods and services.

We consume a vast range of products. But the one-sector growth model lumps them all together into one aggregate, Y, and never tries to deconstruct that aggregate. (DHSS does the same; it is based on the one-sector neoclassical growth model. And the same assumption is made in the 'Limits to growth' model.) Is that OK? It would be OK if our consumption of all the different products increased at the same rate, or (less restrictively) if we could collect products into groups which were similar in a relevant sense (for instance, labour-intensive and resource-intensive products) and show that aggregate consumption of the groups of products increased at the same rate. It turns out that we cannot do so.

Return to light. We already know from Fouquet and Pearson (2006) that the efficiency of lighting in the U.K. increased 1000-fold from 1800 to 2000. In the same period, GDP per capita rose by a factor of 15. Meanwhile, consumption of artificial light per capita rose by a factor of 7000. So we have a massive substitution towards (energy-intensive) light production and consumption.

Regarding transport, the situation is complicated by the fact that the cost of personal transport is not just financial, it is also measured in time. Furthermore, transport varies in quality as well as quantity; faster is, ceteris paribus, better. The result is that rising income is correlated with more rapid forms of transport, and a greater distance travelled per person–year, but not with more time spent travelling. Schäfer (2006) shows that world travel (in terms of person-kilometres travelled per year) has grown more rapidly than global product per capita. Furthermore, rising income is correlated with a successive shift from non-motorized transport → public transport → light-duty vehicles → high-speed transport modes (such as flying).

See also Knittel (2011), who analyses technological change and consumption patterns in the U.S. automobile industry, and shows that for a vehicle of fixed characteristics in terms of weight and engine power, then fuel economy would have increased by 60 percent over the period 1980–2006 due to technological change, but that actual average fuel economy increased by just 15 percent; the difference is due to countervailing increases in the weight and power of vehicles. Rebound?

6.1.2. Why we need structural change to explain the data. Recall from the previous section that in a single-sector model the only way to explain the failure of aggregate energy efficiency to rise is through a failure of product-level energy efficiency to rise. Since we know that product-level energy efficiency has risen, it follows that we must reject the single-sector model. More specifically, since the resource efficiency of individual products has increased, the only way to explain the failure of overall resource efficiency to increase is through a shift in consumption patterns over time, from products of low intrinsic resource intensity towards products of higher

_____

[2]Sources: Popn. data from US census, and car-ownership data from the Bureau of Transportation Statistics.

intrinsic resource intensity. Such a shift is known as *structural change* in macroeconomics, and in order to understand it and predict the future we must build and test models of structural change.

6.1.3. Driving forces of structural change. As explained above—and analysed in depth by Hart (2018a)—if product-level energy efficiency has risen while aggregate energy efficiency has not, the only possible explanation is a shift in consumption patterns towards energy-intensive goods. It has been well known since Engel (1857) that economic growth goes hand-in-hand with systematic shifts in patterns of consumption, driven by income effects: as income increases, the share of necessities such as food declines while luxury goods increase their share.[3] But luxury is a relative concept, and Matsuyama (2002) argues that as productivity improves, households constantly expand the range of goods they consume, as more and more goods become affordable. He models this process using a household utility function with lexicographic preferences, i.e. households expand their consumption from one good to the next irrespective of relative prices.[4] If these goods —introduced successively—are successively more energy-intensive then this process could explain the data.

Shifts in consumption patterns towards energy-intensive goods may also be driven by substitution effects. Return to the period 1870–1970. Since the price of primary energy failed to rise during this period, while energy efficiency rose substantially, *ceteris paribus* we would expect the relative price of energy-intensive products to decline. This could induce substitution towards such products. A related process is studied by Acemoglu and Guerrieri (2008), who model substitution between labour and capital with the aim of explaining both the constant capital share and structural change. They posit two sectors with fixed—but different—capital shares, and show that if the elasticity of substitution between the sectors is less than one (in their calibration it is approximately 0.5) then capital deepening will cause relatively higher output from the capital-intensive sector, but a lower share of income to that sector. The net effect of these shifts is to leave the capital share roughly constant.

More concretely, consider transport. Since the 1930s, we have not simply travelled longer and longer distances by train. Instead we have switched to travel by car, and now increasingly from car to aeroplane. Furthermore, as car engines become more efficient, the fuel economy of the actual cars increases much more slowly, the reason being that we choose ever heavier and more powerful vehicles. Next consider lighting. The average energy efficiency of light production has increased by a factor of around 1000 over the last 200 years.[5] However, production of light has increased by a factor of 7000 over the same period, hence energy use in the lighting sector has increased by a factor of 7 despite the phenomenal increase in efficiency. What has driven this shift into energy-intensive goods such as air travel and lighting? In economic terms it could be either an income effect—rich people like energy-intensive stuff—or a substitution effect—energy-intensive stuff got cheaper compared to other stuff, and people buy more of things when they get cheaper. In reality it will of course be both, but there is some evidence to suggest that the former explanation is very important, i.e. the switch to energy-intensive goods is driven by an income effect.

6.1.4. Why is it important? If we know that structural change is happening, it is very important for policy to know what is causing it. There are several reasons for this, of which we discuss two. The first reason we need to know the causes of structural change is in order to predict future energy demand. There are many reasons why we want to able to predict future energy demand: one is that accurate prediction is important for optimal environmental policy: if demand is likely to rise steeply in the future, this will imply higher carbon emissions and this may in turn lead to higher marginal damage costs of carbon emissions today, and hence higher optimal taxes.

The second reason we need to know the causes of structural change is that it may be directly relevant to the choice of policy instruments in second best. When the only market imperfection is the failure to price carbon emissions then we know from Pigou (1920) that the optimal allocation can be achieved by applying a Pigovian tax on those emissions, i.e. a tax set at the level of marginal damages caused by the emissions.[6] But when there are multiple market imperfections, the situation is unlikely to be so simple, since some of these imperfections may be difficult or impossible to correct, and this may affect the efficacy of emissions taxes. For instance, in an international context it may be difficult for countries to agree on a uniform tax rate (or a global trading system). Under

---

[3]See Houthakker (1957) for a discussion of Engel's law.

[4]Assume rising income. Good 1 is food, and good 2 is not consumed at all until income is sufficient to satiate the desire for food. At this point, consumption of good 2 begins; when desire for that good is satiated, consumption of good 3 begins. Etc.

[5]Data from Fouquet and Pearson (2006) for the UK.

[6]The same result can of course be achieved through tradable permits as well.

these circumstances, if one country applies emissions taxes unilaterally, this may lead to leakage, i.e. the tax may cause emissions to shift out of that country and into other countries. This may be caused by energy-intensive industries relocating away from the taxing country.

If a tax is counterproductive in this way, an attractive alternative may be to subsidize investment in energy-augmenting technology, thus making energy-intensive industries more efficient and (hopefully) reducing emissions. But if the elasticity of substitution between energy-intensive and labour-intensive goods is high, the policy may not give the desired result: an increase in energy efficiency will reduce the price of energy-intensive goods, causing consumers to substitute towards such goods. This effect is known as *rebound* and will be discussed at length below. On the other hand, if the elasticity of substitution is low—implying that demand for energy-intensive goods is inelastic—then a Pigovian tax will have little direct effect on consumption, and its main effect is likely to be on technology. Under these circumstances, technology subsidies (either to energy-efficiency, or to reducing the costs of clean energy production) may be a good option.[7]

6.2. Structural change driven by substitution effects

In this section we develop a model with only substitution effects. In the following section we develop an alternative model driven by income effects. We do this for clarity and simplicity. Note that Hart (2018a) develops a model of structural change driven by both income and substitution effects.

6.2.1. Rebound and consumption patterns. Rebound is frequently ignored in theoretical literature, perhaps because of the common assumption that the economy consists of just a single sector. However, the idea has a long history, starting with Jevons (1865), who argued that future scarcity of coal would be exacerbated, not alleviated, by innovations increasing the efficiency of technologies based on coal use, the reason being that such innovations would lead to a large increase in the use of coal-based technologies. The idea has been picked up more recently by energy and ecological economists (see for instance Binswanger, 2001, and citations), where it has been named the rebound effect.

To define rebound, assume an economy in which total energy use is R, and focus on production of good i using (among other inputs) augmented energy flow $A_{ri}R_i$. Rebound is present when an increase in energy-efficiency A_{ri} by a factor x leads to a reduction of R by less than $R_i(1 - 1/x)$. Note that according to this definition rebound may occur within the production process itself: if the producer of the good has access to a more energy-efficient technology, the producer may choose to use more augmented energy and less augmented labour or capital in the production process. However, we generally think of rebound as occurring on the consumption side of the economy. Given the definition above, we can then think of the producer as having a Leontief production function with no substitutability between augmented energy and other inputs.

Both income and substitution effects may contribute to rebound: an increase in A_{ri} leads (ceteris paribus) to a fall in the price of good i, which raises the purchasing power of consumers (income effect) and induces them to substitute towards consumption of good i (substitution effect). Given the small factor share of energy, the income effect of increases in energy-augmenting technology is likely to be small; on the other hand, given the much higher energy share of some products, the substitution effect of increases of the energy-efficiency of such products may be substantial.

The evidence for rebound effects is reviewed by Sorrell (2007), who finds that they are significant but generally much less than 100 percent, implying that increases in energy efficiency of specific products do lead to large reductions in energy use associated with consumption of those products. A key reason for this is that the substitutability between energy-intensive and other products is far from perfect, just as intuition would suggest.[8] This evidence suggests that rebound alone cannot explain the shift towards consumption of energy-intensive goods, implying that income effects (driven by rising labour productivity) must also have a part to play. Although microeconomic studies of rebound abound, there have only been a few attempts to build macroeconomic models in the literature: see for instance Saunders (1992, 2000).

6.2.2. A general rebound model. To analyse rebound we must have a general equilibrium model. Since we want to focus on the substitution between products on the consumption side in the simplest possible context we assume two products y_1 and y_2, both of which are produced

---

[7]Note that if the elasticity of substitution is low, this implies that structural change must have been driven by income effects, i.e. consumers choosing more energy-intensive goods as they got richer.

[8]For the first analysis of rebound see Jevons (1865), and for another useful presentation see Binswanger (2001).

competitively by a representative firm using a Leontief technology and inputs of labour and a resource. The key equations describing the production side of the economy are as follows:

$$y_1 = \min(A_l l_1, A_{r1} r_1); \tag{6.1}$$

$$y_2 = \min(A_l l_2, A_{r2} r_2); \tag{6.2}$$

$$l_1 + l_2 = L; \tag{6.3}$$

$$r_1 + r_2 = R. \tag{6.4}$$

The first two equations are the production functions for y_1 and y_2, and the second two equations show total labour and resource inputs. We assume that the labour force L is fixed, whereas resources R are provided at a price w_r which is fixed relative to the wage w_l, i.e. $w_r/w_l = \psi$ where ψ is fixed. Both labour and the resource are traded on competitive markets.

Define p_1 as the price of good y_1, and p_2 as the price of y_2. Furthermore, let the wage be the numéraire, normalized to be equal to labour productivity A_l: $w_l = A_l$. Since markets are competitive then price must be equal to marginal cost. Since we have Leontief then marginal cost is the same as average cost, and

$$r_1 = l_1(A_l/A_{r1}); \quad r_2 = l_2(A_l/A_{r2}).$$

So total costs for good 1 are

$$c_1 = A_l l_1 + \psi A_l l_1(A_l/A_{r1}),$$

and similarly for good 2. Furthermore, production of the goods is

$$y_1 = A_l l_1 \quad \text{and} \quad y_2 = A_l l_2,$$

hence average costs, and hence prices, are given by

$$p_1 = 1 + \psi(A_l/A_{r1});$$

$$p_2 = 1 + \psi(A_l/A_{r2}).$$

Finally, use $L = l_1 + l_2$ and the expressions for r_1 and r_2 above to show that

$$R = r_1 + r_2 = A_l[l_1/A_{r1} + (L - l_1)/A_{r2}].$$

To investigate the rebound effect we assume that A_{r1} increases. What happens? From the information we have, we know that r_1/l_1 will decline, and p_1 (the price of good y_1) will also decline. In general we expect this to lead to an increase in the quantity y_1 demanded, and hence also an increase in l_1, labour employed on the production of good 1.

Mathematically we want to find the elasticity of total resource demand R to an increase in A_{r1}: we denote this elasticity η_r. Furthermore, for convenience—since we do not yet wish to specify the demand side of the model—we define the elasticity of l_1 w.r.t. the change in A_{r1} as η_l. Given these definitions, differentiate the above expression for R with respect to A_{r1} to show that

$$\eta_r = -\frac{r_1}{r_1 + r_2}\left[1 - \left(1 - \frac{A_{r1}}{A_{r2}}\right)\eta_l\right]. \tag{6.5}$$

To understand equation (6.5), assume first that $\eta_l = 0$, implying that the elasticity of substitution between the products on the consumption side is zero. Then the elasticity of total energy demand with respect to an increase in A_{r1} is simply equal to the share of product 1 in total energy demand. That is, there is zero rebound.

Now assume instead that $\eta_l > 0$, implying that the price reduction in product 1 causes some reallocation of consumption (and hence also production) towards that product. Thus the term in square brackets may differ from 1. However, as a baseline case note that when $A_{r1} = A_{r2}$ then this term is still 1, and η_r is still equal to the share of product 1 in total energy demand. The reason is that when the products (1 and 2) are of equal energy intensity then a reallocation of consumption between them does not affect total energy demand. Thus the rebound effect is zero.

Now assume that $\eta_l > 0$ and $A_{r1} > A_{r2}$, implying that product 1 is less energy-intensive that product 2. Now the term in square brackets is greater than one, implying that the rebound effect is *negative*, i.e. the increase in A_{r1} causes a *greater* reduction in total energy demand than would be expected on the basis of a naive analysis. The reason is that the reallocation of consumption towards product 1 occurs at the expense of product 2, and product 2 is—by assumption—more energy-intensive than product 1. Therefore this reallocation leads to a reduction in energy demand over and above that which is caused directly by the efficiency increase in production of good 1.

Finally assume that $\eta_l > 0$ and $A_{r1} < A_{r2}$, implying that product 1 is more energy-intensive that product 2. Now the term in square brackets is less than one, implying that the rebound effect is positive, i.e. the increase in A_{r1} causes a smaller reduction in total energy demand than would be expected on the basis of a naive analysis. The reason is that the reallocation of consumption

towards product 1 causes a net increase in energy demand because product 2 is—by assumption —less energy-intensive than product 1. Therefore this reallocation leads to an increase in energy demand, partly or even completely cancelling out the reduction which is caused directly by the efficiency increase in production of good 1.

So we can conclude that rebound is caused by the reallocation of labour (and other inputs other than the resource) between sectors, triggered by an increase in resource efficiency in one sector. When that sector is of low initial resource-intensity then this reallocation is a further benefit, increasing the reduction in resource demand. In other words, the rebound effect is negative. When that sector is energy intensive then the reallocation diminishes—or may even reverse—the direct effect of the efficiency increase.

6.2.3. A very simple specified model. We now set out to specify our rebound model (above) and try to calibrate it to match data. One property we would like our specified model to possess is that it should yield the constant resource share we observed in Figure 4.4. We begin by testing a model that mirrors the simple assumption of a Cobb–Douglas production function of a single good with two inputs (labour and the resource), but where we now have two goods each produced with one input, and the substitutability is on the consumption side. Having checked the model against data we go on to develop a slightly more sophisticated variant.

In our simplest model the production functions for y_1 and y_2 are as follows:

$$y_1 = \gamma_l k_l l; \tag{6.6}$$

$$y_2 = \gamma_r k_r r. \tag{6.7}$$

Thus labour is the only input to y_1, and energy is the only input to y_2. In equilibrium, $l = L$ and $r = R$. To complete the overall picture we assume that aggregate consumption Y is a CES function of the two aggregate products y_1 and y_2, hence (at the level of the representative firm) we define a parameter $\epsilon \in (-1, \infty)$, and

$$y = (\alpha y_1^{-\epsilon} + (1 - \alpha) y_2^{-\epsilon})^{-1/\epsilon}. \tag{6.8}$$

Thus when ϵ is positive the two aggregate products are complements in the sense that if a product becomes increasingly scarce then its factor share rises.

As above, the energy price is w_r and the wage w_l, and $w_r/w_l = \psi$. We want to test the ability of the model economy here to reproduce the data seen in Figure 4.4. To do so we let labour L and the ratio of the input prices, w_r/w_l, evolve exogenously, and derive total energy use R from the model.

The solution is straightforward. Briefly, derive two different expressions for the ratio of the prices of the aggregate goods: firstly by comparing their marginal contribution to y, and secondly by comparing their unit production costs. Use these two expressions to eliminate the price ratio, and rearrange to show that

$$\frac{R}{L} = \left[\frac{1 - \alpha}{\alpha} \left(\frac{\gamma_r k_r}{\gamma_l k_l} \right)^{-\epsilon} \left(\frac{w_r}{w_l} \right)^{-1} \right]^{1/(1+\epsilon)}.$$

Hence the aggregate elasticity of substitution between energy and labour is $1/(1 + \epsilon)$.

Now set $\epsilon = 0$. This implies that equation 6.8 is Cobb–Douglas, and the aggregate elasticity of substitution between energy and labour is 1. Thus we have the constant-share result and 100 percent rebound (energy demand is not affected by the direction of technological change)! The result is intuitive: we have two products, one made entirely using labour, the other using only energy. When the products are combined in a Cobb–Douglas function on the consumption side the products take constant shares, and therefore labour and energy must also take constant shares.

This model tackles two of the weaknesses of the Cobb–Douglas model: the unrealistically high substitutability between the inputs on the production side, and assumption that all products have equal energy intensity and are perfect substitutes on the consumption side. However, it replaces them with an equally troubling characteristic, i.e. the assumption that final goods are produced either using pure labour or pure energy. To see how problematic this is, consider Figure 6.2, where we illustrate how energy intensity varies across sectors. In Figure 6.2(a) we see that if we divide consumption into two equal parts, one energy-intensive the other not, then the low-energy-intensity consumption accounts for just under 20 percent of energy consumption. In 6.2(b) we see the energy intensity and expenditure share of different consumption categories: different types of services—of low energy-intensity—account for more than half of expenditure, while the two major energy-intensive categories are habitation and motor transport, and the final category (with highest intensity but only a small expenditure share) is air transport. The figure shows that dividing consumption into two input-specific products does not follow naturally from the data, which suggests a continuum of products of gradually increasing intensity. Furthermore, the energy share

of the most energy-intensive product (air transport) is only around 14 percent, nothing like to 100 percent intensity assumed in the model.[9]


FIGURE 6.2. Cumulative energy use and energy intensity plotted against cumulative expenditure when consumption products are sorted in order of increasing energy intensity. All the axes are normalized. Regarding energy intensity, we only have data on relative intensities, and we normalize to give an average intensity of 4 percent.

6.2.4. A slightly more general model. Now assume two aggregate products Y_1 and Y_2, where the former is labour-intensive in production and the latter is energy-intensive. Each are produced in a constant-returns production function in which technology is exogenous. We can thus assume that they are produced by perfectly competitive firms and hence there are no aggregation problems within the two production sectors. We therefore focus on production y_1 and y_2 from the representative firm in each sector.

Because we want to build a simple model, and because labour takes more than 95 percent of returns (energy less than 5 percent), we assume that y_1 is produced using only labour, whereas y_2 uses both. The production function for y_2 is Leontief: as discussed above, the short-run elasticity of substitution between labour and energy is very low, and in the context of this model it can be ignored. Finally, we assume that technological change is completely unbiased, so we have a single knowledge stock k that grows exogenously and boosts the productivity of all inputs equally. We therefore have

$$y_1 = \gamma_{l1} k l_1; \tag{6.9}$$

$$y_2 = k \min(\gamma_{l2} l_2, \gamma_{r2} r_2); \tag{6.10}$$

$$l_1 + l_2 = L; \tag{6.11}$$

$$r_2 = R. \tag{6.12}$$

The remaining parameters are analogous to those of the DTC model: L and R are total labour and total resource use respectively. To complete the overall picture we assume that aggregate consumption Y is a CES function of the two aggregate products Y_l and Y_r, hence (at the level of the representative firm) we define a parameter $\epsilon \in (-1, \infty)$, and

$$y = (\alpha y_1^{-\epsilon} + (1 - \alpha) y_2^{-\epsilon})^{-1/\epsilon}. \tag{6.13}$$

Thus when ϵ is positive the two aggregate products are complements in the sense that if a product becomes increasingly scarce then its factor share rises.

As above, the energy price is w_r and the wage w_l. We want to test the ability of the model economy here to reproduce the data seen in Figure 4.4. To do so we let labour L and the ratio of the input prices, w_r/w_l, evolve exogenously, and derive total energy use R from the model.

To solve, first note that (from the Leontief production function) $\gamma_{l2} l_2 = \gamma_{r2} r_2$. Then derive two different expressions for the ratio of the prices of the aggregate goods: firstly by comparing their

---

[9]The data for Figure 6.2 are from Mayer and Flachmann (2011). The products—in order of increasing energy intensity—are Education services; Health services; Health services and social work; Other services; Cultural and sport services; Retail and wholesale trade; Hotel and restaurant services; Office and electrical machinery; Paper and publishing; Water transport; Auxiliary transport services; Other land transport; Furniture, jewellery, musical instruments etc.; Other products; Textiles and furs; Food and tobacco; Agricultural products; Transport via railways; Habitation; Chemical products, rubber, and plastic; Motor transport; Air transport.

marginal contribution to y, and secondly by comparing their unit production costs. Use these two expressions to show that

$$\frac{l_1}{l_2} = \left[\frac{\alpha}{1-\alpha} \left(\frac{\gamma_{l1}}{\gamma_{l2}} \right)^{-\epsilon} (1+\omega) \right]^{1/(1+\epsilon)},$$

where

$$\omega = \frac{w_r}{w_l} \frac{\gamma_{l2}}{\gamma_{r2}},$$

and represents the ratio of the energy price to the wage (in efficiency units). Use the expression for l_1/l_2 and the restriction on total labour to find l_2, and hence R using $\gamma_{l2}l_2 = \gamma_{r2}r_2$:

$$\frac{R}{L} = \frac{\gamma_{l2}}{\gamma_{r2}} \left\{ 1 + \left[\frac{\alpha}{1-\alpha} \left(\frac{\gamma_{l1}}{\gamma_{l2}} \right)^{-\epsilon} (1+\omega) \right]^{1/(1+\epsilon)} \right\}^{-1}.$$

Finally show that the elasticity of substitution between energy and labour is as follows:

$$\eta_s = \frac{W}{Q} \frac{\partial Q}{\partial W} = \frac{\omega}{1+\omega} \frac{1}{1+\epsilon} \frac{\left[\frac{\alpha}{1-\alpha} \left(\frac{\gamma_{l1}}{\gamma_{l2}} \right)^{-\epsilon} (1+\omega) \right]^{1/(1+\epsilon)}}{\left[\frac{\alpha}{1-\alpha} \left(\frac{\gamma_{l1}}{\gamma_{l2}} \right)^{-\epsilon} (1+\omega) \right]^{1/(1+\epsilon)} + 1}.$$

So when the price of energy declines relative to labour (as it has done historically), what happens to the energy share in the model economy? Start with the case of $\epsilon = 0$, in which case equation 6.13 reduces to the Cobb–Douglas. Then we have

$$\eta_s = \frac{\omega}{1+\omega} \frac{\frac{\alpha}{1-\alpha}(1+\omega)}{\frac{\alpha}{1-\alpha}(1+\omega)+1}$$

So η_s is less than 1, i.e. when the price of energy relative to labour declines by 1 percent, energy use R (relative to labour) should rise by less than 1 percent, hence the energy share declines.

However, there is a special case in which the above result does not hold, and that is when $\omega \to \infty$, implying (in the limit) that the production function for y_2 is simply

$$y_2 = \gamma_{r2}kr_2.$$

Then we have $\eta_s = 1$, i.e. the factor share of energy is constant. This is to be expected, as this case amounts to reducing the model to our previous model with input-specific products.

The special case sheds light on the general cases. When both labour and energy are used in producing the energy-intensive good then—if we hold ϵ at zero—the effect of declining energy price on the relative price of that good is not so great, hence the shift in consumption patterns caused by the price shift is not so great, hence the energy share declines as the relative price of energy declines ($\eta_s < 1$). Furthermore, as the energy share of the energy-intensive good declines the elasticity η_s declines, approaching zero in the very long run as ω approaches zero, at which point the energy share is zero.

When $\epsilon < 0$—indicating a very high degree of substitutability between the labour-intensive and the energy-intensive goods—then for ω sufficiently high the model can deliver $\eta_s = 1$ and hence a constant energy share. However, again, the decline in the relative price of energy causes ω to decline, and as ω declines then η_s will decline. That is, the constancy of the energy share is only temporary, and in the long run the energy share will—far from being constant—approach zero.

We do not parameterize the model formally, rather we look for evidence suggesting reasonable values for the parameters. Consider first Figure 6.2. Firstly, the figure shows that dividing consumption into just two products of differing energy intensity does not follow naturally from the data, which suggests a continuum of products of gradually increasing intensity. Secondly, the spread of energy intensities in the data is not very great: half of consumption expenditure goes on products with energy intensity between 25 and 50 percent of the average level, 49 percent of expenditure goes on products with energy intensity between 63 and 250 percent of the average level, and even the most energy-intensive good (air transport) at 1 percent of expenditure is just 3.5 times more energy-intensive than the average good.

The data suggests that if we must lump consumption into two goods, and assuming that average energy intensity is 5 percent, then we could choose one good with intensity 2 percent which accounts for 50 percent of consumption, and the other with intensity 8 percent accounting for the other 50 percent. Given the even more restrictive set-up of the model—with one good having zero energy intensity—then we could think of the second good as accounting for 50 percent of consumption at 10 percent intensity. If this represents the situation today, could we have got there (in the model economy) via a long-run path with constant energy share?

It is straightforward to show that the above parameterization implies that $\omega = 0.12$.[10] The only way to achieve a sufficiently high elasticity of substitution between labour and energy would then be to assume that ϵ were close to -1, i.e. labour-intensive goods such as services, and energy-intensive goods such as transport and housing should be almost perfect substitutes. The data does not support such a parameterization.

6.2.5. The failure of models without income effects. Consider for instance microeconomic studies of rebound effects. Rebound is present when an increase in the energy efficiency of some process in the economy by a factor x (where $x > 1$) leads to a reduction of total rate of energy use in the economy by less than $Q(1 - 1/x)$, where Q was the original rate of energy use in that process. There are potentially several reasons for rebound in a real economy, but in the model above there is just one: that an increase in the efficiency with which the energy-intensive product is produced leads to a reallocation of production labour towards that product, thus reducing the potential saving on energy use. The evidence for rebound effects is reviewed by Sorrell (2007), who finds that they are significant but generally much less than 100 percent: increases in energy efficiency of specific products do lead to large reductions in energy use associated with consumption of those products. In terms of the model, this indicates that the substitutability between energy-intensive and other products is relatively far from perfect, just as intuition would suggest.[11]

The rebound evidence shows that the price-elasticity of demand for specific products is insufficient to account for the historically observed aggregate elasticity of substitution between energy and labour. But the evidence over the period 1870–1970—rapidly increasing energy-efficiency in the production of specific goods (such as artificial light and motive power), combined with the rise in energy use tracking the rise in global product—nevertheless demonstrates that there must have been a shift in consumption patterns towards energy-intensive goods. Direct evidence on consumption patterns confirms that such shifts have occurred. Regarding consumption of light, for instance, Fouquet and Pearson find that per capita consumption of artificial light in the U.K. rose by a factor of 7000 between 1800 and 2000. This factor should be compared to the approximately 15-fold increase in per-capita GDP over the same period; without shifts in consumption patterns, consumption of all products should have risen by this factor over the period. Regarding transport, Knittel (2011) analyses technological change and consumption patterns in the U.S. automobile industry, and shows that—for a vehicle of fixed characteristics in terms of weight and engine power—fuel economy would have increased by 60 percent over the period 1980–2006 due to technological change; this is approximately on a par with increases in labour productivity. Furthermore, he also shows that actual average fuel economy increased by just 15 percent, the difference being due to countervailing increases in the weight and power of vehicles. Thus we have efficiency improvement leading to a fall in unit costs of energy services, combined with a countervailing increase in consumption of these services. We can thus summarize the rebound model as follows.

The above evidence supports the idea that substitution between products of differing energy intensity is important to take into account when determining energy policy, and demonstrates how such substitution has the potential to undermine efforts to reduce energy demand through increases in energy efficiency. However, the failure of our simple model above shows that we need a less restrictive model in order to capture the key mechanisms: two mechanisms which might be relevant are (i) substitution due to income effects as well as substitution effects, and (ii) substitution towards new products rather than between existing products. These two mechanisms are related to one another, as we can see by considering the historical data. Consider for instance transport: during the 20th century technological progress drove both rising incomes and the appearance of new energy-intensive consumption goods such as automobile and air transport. Such goods are luxuries to low-income households (and hence also to any household if we go back in time sufficiently) hence their income-elasticity of demand was initially high, and rising incomes caused a substitution towards such goods from less energy-intensive alternatives. In the next chapter we present a model which envisions this process of technological change and substitution towards new, more energy-intensive goods as a continuous process which will only be broken by changes in the relative price trend of energy to labour.

---

[10]If the energy share of the second good is 10 percent while its expenditure share is 50 percent then 44 percent of labour must be employed in making the second good, i.e. $l_2/l_1 = 0.8$. Furthermore, if the overall energy share is 5 percent then $w_l(l_1 + l_2) = 19w_r r_2$. Eliminate l_1 to show that $w_l l_2 = 8.4 w_r r_2$. But we know that $\gamma_{l2} l_2 = \gamma_{r2} r_2$, hence $w_l/\gamma_l = 8.4 w_r/\gamma_{r2}$, hence we have $\omega_r = 0.12$.

[11]For the first analysis of rebound see Jevons (1865), and for another useful presentation see (Binswanger, 2001).

FIGURE 6.3. The model with multiple sectors and income effects.

6.3. Structural change driven by income effects

In the previous section we argue that a model without income effects can never match the data about structural change an energy use. In this section we set up an alternative model with structural change driven only by income effects.

6.3.1. The model. We now build a model—illustrated in Figure 6.3—in which structural change is driven by income effects, in the spirit of Matsuyama (2002). A range of products can be made, and as consumers become richer they 'trade up' to the best product they can afford. The best products are also the most expensive to make, and the most energy-intensive.

There is an infinite series of products Y_i, and the production function for product is as follows:

$$Y_i = (A/2^{i-1})\min\{L_{Yi}, \alpha R_i/2^{i-1}\},$$

where A is productivity, L_Y is labour in final-good production, R is the resource input, and α is a parameter. So when productivity A doubles (holding L_{Yi} and R_i constant) Y_i doubles. However, adding 1 to i halves Y_i for given L_{Yi}, assuming that energy inputs can be doubled. The resource R is extracted using labour, and the extraction function is linear:

$$R = \phi A L_R.$$

Given the Leontief production function we have, for each good Y_i,

$$L_{Yi} = \alpha \phi A L_{Ri}/2^{i-1}.$$

Each individual supplies a unit of labour, so (using l_Y and l_R for the individuals' labour allocation) $l_Y + l_R = 1$. And if a specific individual produces good i we have (since $l_{Yi} = \alpha \phi A l_{Ri}/2^{i-1}$),

$$l_{Yi} = \frac{\alpha \phi A/2^{i-1}}{1 + \alpha \phi A/2^{i-1}}; \tag{6.14}$$

$$y_i = \frac{A}{2^{i-1}}\frac{\alpha \phi A/2^{i-1}}{1 + \alpha \phi A/2^{i-1}}; \tag{6.15}$$

$$l_{Ri} = \frac{1}{1 + \alpha \phi A/2^{i-1}}; \tag{6.16}$$

and $$r_i = \frac{\phi A}{1 + \alpha \phi A/2^{i-1}}. \tag{6.17}$$

So for given i (i.e. assuming no structural change), if A increases at a constant rate then labour in the resource sector declines, approaching zero when $A \to \infty$, resource extraction grows slowly and

approaches a constant rate, and production increases at a rate that approaches the growth rate of A in the limit.

Labour is paid the same wage w_L in both sectors, and we normalize $w_L = A$.[12] It follows that nominal GDP is AL, where L is total labour. We also define this to be real GDP.[13] Total costs for production of quantity Y_i are $w_L(L_{Yi} + L_{Ri}) = Y_i 2^{i-1}[1 + 2^{i-1}/(\alpha\phi A)]$. And unit costs (and hence p_i) are as follows:

$$p_i = 2^{i-1}\left[1 + \frac{2^{i-1}}{\alpha\phi A}\right]. \tag{6.18}$$

So we have a range of goods indexed by i, which are increasingly costly to make as i increases, both because of greater labour costs in the final-good sector, and greater resource costs. Consider for instance car travel ($i = n$) and air travel ($i = n+1$). According to the function, if we switch from car to air travel—holding A constant—the same amount of travel Y will require double the labour and four times the energy.

Prices increase steeply in i, but consumers choose goods with high i because they are more attractive, or judged to be of higher quality. More specifically, we assume that consumers have lexicographic preferences such that they always prefer to consume the good with the highest i available, subject to a restriction that $c_i > \bar{c}$. We can interpret this as consumers demanding a minimum quantity of consumption $\bar{c}$, and that given that this quantity restriction is satisfied they choose the highest affordable quality i. This is most easily understood in terms of food. Assume that food is the only consumption good. Then the utility function implies that consumers only ever demand a certain quantity of food (e.g. the quantity they need to satisfy their hunger), but subject to this restriction they choose the highest possible quality. However, we can also think of other consumption categories, such as transport. We can imagine an economy in which each individual needs a car, and for given income the individuals choose the most expensive car they can afford (rather than, for instance, two crummy cars).

To close the model we need to link production and consumption. For simplicity we assume that there are L identical individuals in the economy, supplying total labour L which is divided between L_R and L_Y. Total household income is $w_L L = AL$, and the income of each individual is A. Since the price of good i is p_i (equation 6.18), the condition for the representative consumer being able to afford good i is

$$A \geq 2^{i-1}\bar{c}\left[1 + \frac{2^{i-1}}{\alpha\phi A}\right].$$

For given i, define the minimum value of A which yields affordability as $\bar{A}_i$. Insert $\bar{A}_i$ instead of A, write the equation as an equality, and solve for $\bar{A}_i$ to yield

$$\bar{A}_i = 2^{i-2}\left[1 + \left(1 + \frac{4}{\alpha\phi\bar{c}}\right)^{0.5}\right]\bar{c}.$$

It follows by inspection that each time A doubles, the index i of the best affordable product increases by 1. And by inspection of equations 6.14–6.17, production (i.e. the number of items y_i produced) remains the same, labour allocation between the resource and final-good sectors remains the same, but resource use tracks the growth rate of A (and hence the growth rate of GDP).

6.3.2. Policy implications. Assume that the resource R is coal, and that coal causes polluting emissions with associated damages. There is no alternative technology, and technological change is exogenous. Agents are symmetrical. What is the effect of a Pigovian tax on emissions from coal burning, with the revenues recycled lump-sum? Consider a single agent. This agent has to pay a higher price for each good i, because of the tax. But the agent's income is higher, thanks to the transfer. If all agents carry on buying the same good they chose before the tax was imposed, nothing will change: the tax payments and the transfer will exactly match each other. But could a single agent raise her utility by switching to a good with lower i, hence paying less tax and being able to buy more of the good than she could in laissez-faire? The answer to this question is clearly 'No', because of the lexicographic preference function: any quantity of a good of lower i will give lower utility than the minimum quantity of the original good. So, the emissions tax has no effect!

The above example illustrates how little effect an emissions tax has on consumption-based emissions when consumers are not willing to substitute between different goods. Note that there would be some effect if the tax had a redistributive effect: if we had consumers on different

---

[12]Note that it makes no sense to normalize the price of the final good to 1 because there are many final goods which vary in quality.

[13]How real GDP should be measured is a complex question in an economy with many goods of varying quality. Given our definition, the prices of the goods decline over time, but only slowly ??

incomes, where those with highest income chose the most emissions-intensive consumption, then an emissions tax would redistribute some income to those on lower incomes, potentially saving some emissions. However, if the redistribution caused lower-income consumers to 'trade up' to higher emissions consumption the effect might be the opposite.

The above analysis does not show that emissions taxes are toothless in an economy with very powerful income effects. However, the effect (if any) must occur on the production side. That is, the tax must induce firms to shift their production technology to low-emissions alternatives, including the option of first developing such alternatives (linked to Solow's third mechanism). For instance, if producers can choose between coal, gas and renewables for electricity generation, an emissions tax may shift the choice from coal to gas, and a higher tax may shift the choice again to renewables. So the key point is that in an economy in which consumers are not prepared to substitute between alternative products, the key to reducing emissions is cleaner technology.

If an emissions tax induces a switch to a perfectly clean technology, the effect on emissions is obvious. But what if the tax causes a more expensive but cleaner resource to be used in production? That is, what if the tax results in an increase in α, but a decrease in ϕ? If the resource share is small (so $L_Y \gg L_R$) then we can approximate 6.17 as

$$r_i = \frac{2^{i-1}}{\alpha},$$

and an increase in α drives emissions down linearly, as long as i stays the same. And since the resource share is small, the effect of resource efficiencies on the chosen i will also be small, hence the effect of resource efficiency on the choice of product will typically be small. That is, in this economy the 'rebound effect' of exogenous technological change is small. And since choice of technology is key, regulations banning polluting technologies may be as efficient—and simpler to apply—than pollution taxes.

Substitution between alternative resource inputs

In the previous two chapters we saw—Chapter 5—why models based on DTC are not successful at explaining why aggregate demand for (for instance) energy from fossil fuels has tracked GDP, implying that aggregate energy efficiency has failed to rise, and then—Chapter 6 —how structural change driven by a combination of substitution and income effects can explain the aggregate data. Furthermore, we saw that where income effects drive increases in resource demand, the only way to stem resource demand—save preventing rises in income—may be to find alternative technologies using different inputs. In this chapter we return to the basic model— Cobb–Douglas in labour and the resource, no structural change—and extend it in a third direction, by including the possibility of substituting between alternative natural resources in production, using a nested production function.

We begin by building and testing a very simple model of resource substitution, in which two alternative resources are available which are substitutable for one another, and technological change in the resource sector is unbiased. We show that the model can do a reasonable job of accounting for aggregate data in two cases. We go on to consider Solow's third adaptation mechanism to resource scarcity (page 57), which was to increase—through technological change—the efficiency of an alternative (substitute) resource in production of one or more product categories. The idea here is that when there is a need to switch to an alternative resource, directed investments lead to an increase in the efficiency with which we can use the alternative resource. This idea is related to the concepts of path dependence and lock-in discussed by Arthur (1989) among others, and the idea that we are 'locked in' to fossil-fuel use by history dates at least to Unruh (2000). More recently, Acemoglu et al. (2012) have proposed a model in which a regulator can transform lock-in to break-out: through a massive but short-run effort the regulator can set in train a technology transition from dirty to clean energy, a transition which will then continue without the need for further regulation. We analyse—and question—the relevance of this model in the following chapter, on pollution.

7.1. A simple model with alternative resource inputs

7.1.1. The basic model. Recall from Chapter 4 that the simple Cobb–Douglas production function with labour, capital, and resources does a decent job of matching the aggregate long-run data, with constant factor shares for labour, capital, and resources.[1] Furthermore, in a long-run context with perfect information and relatively constant growth we can abstract from capital without any major loss of relevance for the model. We then have the following aggregate production function:

$$Y = (A_L L)^{1-\alpha}(A_R R)^{\alpha}.$$

The units of Y are *widgets year*$^{-1}$, the units of L are simply *workers*, hence the units of A_L are *widgets worker*$^{-1}$ *year*$^{-1}$. The units of R, the resource input, are now *res year*$^{-1}$, where res is a measure of resource services, the meaning of which will become clearer below.

The production function implies (as we saw above) that returns to the resource relative to labour are constant, irrespective of relative prices, directed technological change, etc. To see this set up the representative producer's problem—

$$\max \pi = p_y(A_L L)^{1-\alpha}(A_R R)^{\alpha} - w_l L - w_r R$$

—and differentiate w.r.t. R to find w_r:

$$w_r = \alpha Y/R$$

and
$$w_r R/Y = \alpha.$$

---

[1]Note that we know from the previous chapter that the diversity of products made in real economies is important for understanding resource demand. However, here we make the assumption that it can be ignored—at least for the time being —when modelling demand for alternative resources.

Now we return to R, resource services. We have two physical resource inputs (units tons year^{-1}), X_c and X_d, and the production function for resource services is

$$R_t = [(\gamma_c A_{ct} X_{ct})^\epsilon + (\gamma_d A_{dt} X_{dt})^\epsilon]^{1/\epsilon}, \qquad (7.1)$$

where the elasticity of substitution between the two is $1/(1-\epsilon)$, and $\epsilon \in (0,1)$. So the two inputs are highly substitutable in the production of resource services R. Consider for example iron and aluminium, or coal and oil. The units are shown in Table 7.1, where (for instance) ideas$_C$ means ideas augmenting resource input C.

TABLE 7.1. Units in the production function for resource services

Quantity	Unit	Quantity	Unit
γ_c	res ton$_C^{-1}$ idea$_C^{-1}$	γ_d	res ton$_D^{-1}$ idea$_D^{-1}$
A_C	ideas$_D$	A_D	ideas$_D$
C	tons of C year^{-1}	R	tons of D year^{-1}

Note that the production function for resource services implies that if we have only one resource input (perhaps only iron and no aluminium) then the function is very simple. For instance, if we have only input C then

$$R_t = \gamma_c A_{ct} X_{ct},$$

and likewise for input D. Now assume that $\gamma_c A_{ct} X_{ct}$ and $\gamma_d A_{dt} X_{dt}$ are both equal to $\bar{R}/2$. Then when they are combined in the production function we have

$$R = 2^{(1-\epsilon)/\epsilon} \bar{R},$$

which is greater than $\bar{R}$. Because the inputs are imperfect substitutes they therefore complement each other to some extent, implying that the whole is more than the sum of the parts.

Now we turn to technological change. In this section we simply assume that technological change in the resource production function is unbiased and exogenous, implying that A_C and A_D grow at equal rates. Furthermore, we assume that A_R is constant, since resource services are already augmented by technologies A_C and A_D. Finally, A_L grows exogenously at a constant rate g, which is also the growth rate of A_C and A_D, and population grows at rate n. Without loss of generality we normalize $A_R = 1$, and $A_C = A_D = A_L = A$. The two resources are extracted using final goods, and unit extraction costs are w_{ct} and w_{dt} respectively. So we have

$$Y_t = (A_t L_t)^{1-\alpha} R_t^\alpha,$$

$$R_t = A_t [(\gamma_c X_{ct})^\epsilon + (\gamma_d X_{dt})^\epsilon]^{1/\epsilon},$$

and $\qquad C_t = Y_t - (w_{ct} X_{ct} + w_{dt} X_{dt}),$

where C is aggregate consumption. If the resources are scarce then resource price will be extraction cost plus scarcity rent; however, for now we simplify by assuming that price equals extraction cost, and account for scarcity informally by allowing price to rise. When testing the model empirically we only have data on price (and hence no data on extraction cost and scarcity rent).
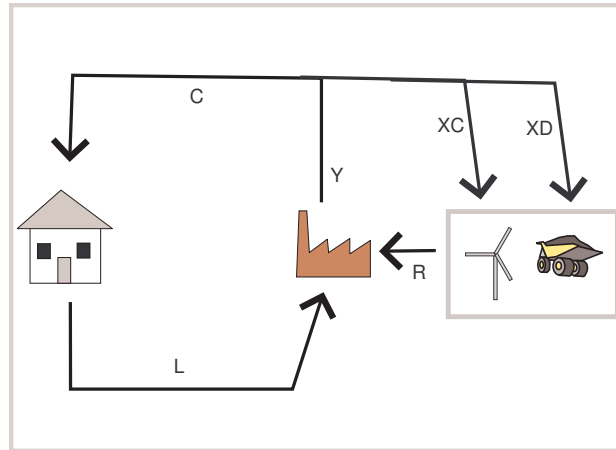

FIGURE 7.1. The economy with substitute resource inputs

7.1.2. The solution. The solution to the model is straightforward. We already have that

$$w_r R = \alpha Y.$$

Consider now production of R. Set up the producer's profit-maximization problem as follows,

$$\pi = w_{rt} A_t \left[(\gamma_c X_{ct})^\epsilon + (\gamma_d X_{dt})^\epsilon \right]^{1/\epsilon} - w_{ct} X_{ct} - w_{dt} X_{dt},$$

and take first-order conditions to show that

$$w_c X_c = w_r (R/A)^{1-\epsilon} (\gamma_c X_c)^\epsilon$$

and
$$w_d X_d = w_r (R/A)^{1-\epsilon} (\gamma_d X_d)^\epsilon,$$

and hence

$$w_c X_c = w_r^{1/(1-\epsilon)} (R/A)(\gamma_c/w_c)^{\epsilon/(1-\epsilon)}$$

and
$$w_d X_d = w_r^{1/(1-\epsilon)} (R/A)(\gamma_d/w_d)^{\epsilon/(1-\epsilon)}.$$

Because we have perfect markets, price equals unit cost so

$$w_r = (w_c X_c + w_d X_d)/R$$

$$= \left\{ A / \left[(\gamma_c/w_c)^{\epsilon/(1-\epsilon)} + (\gamma_d/w_d)^{\epsilon/(1-\epsilon)} \right] \right\}^{\epsilon/(1-\epsilon)}.$$

And since $w_r R = \alpha Y$ we have

$$w_r = \alpha (AL/R)^{1-\alpha},$$

and we can eliminate w_r to yield

$$R = AL \left\{ \alpha \left[(\gamma_c/w_c)^{\epsilon/(1-\epsilon)} + (\gamma_d/w_d)^{\epsilon/(1-\epsilon)} \right] \right\}^{1/(1-\alpha)}.$$

So if w_c and w_d are both constant then R grows at the same rate as Y, i.e. $g + n$, the sum of the growth rates of labour productivity and population. And the relative factor shares of the resources are

$$\frac{w_c X_c}{w_d X_d} = \left(\frac{\gamma_c/w_c}{\gamma_d/w_d} \right)^{\epsilon/(1-\epsilon)}.$$

This implies that the resource that is cheaper per efficiency unit takes the larger factor share, and the advantage is bigger the higher is the substitutability between the resources (i.e. when $\epsilon \to 1$).

7.1.3. Empirical tests. To test the explanatory power of this very simple model we take two pairs of resources, oil–coal and iron–aluminium. In each case we first present data on prices and factor expenditure, and compare the expenditure to GGP (gross global product). We then plot total expenditure on the two factors together against GGP, and check whether the ratio of the two is approximately constant. (Recall from the model that this ratio should be constant, α.) Finally we plot the ratio of expenditures on the two factors, and compare this to the same ratio obtained using the price data and the parameterized model. In the model economy (immediately above) the ratio is

$$\frac{w_c X_c}{w_d X_d} = \left(\frac{\gamma_c/w_c}{\gamma_d/w_d} \right)^{\epsilon/(1-\epsilon)}.$$

We take the price data (i.e. w_c and w_d), then find parameters ϵ and γ_c/γ_d to fit the factor-share data as well as possible.[2]

In Figure 7.2 we see that the Cobb–Douglas does a reasonable job of approximating the aggregate global production function in this case, although the factor share of oil and coal combined has (according to the data) actually risen during the 140 years for which we have data, rather than being constant. This can be seen most clearly in the middle panel. In the left panel we see that the price trend of oil relative to coal was a very slow decline up to 1973, after which prices became much more volatile, and the gap decreased. The model (smoothed) can easily match the increase in the oil share as the oil price declines (right-hand panel), but it doesn't do quite such a good job of accounting for what happens after the oil crisis, where the oil share continues to grow despite the increasing price of oil relative to coal. Nevertheless, considering the simplicity of the model it does a remarkably good job of matching the data.

In Figure 7.3 we see that the Cobb–Douglas does an excellent job of approximating the aggregate global production function when we assume that the inputs are labour (or labour–capital), iron, and aluminium (middle panel). In the left panel we see that the price trend was for a steady decline in the price of aluminium relative to iron, while the trend in factor share was a corresponding increase in the share of aluminium. The model (smoothed) also does a remarkably good job of

---

[2]Note that at this stage we simply eyeball the graphs, Figures 7.2 and 7.3 to determine the goodness of fit.

FIGURE 7.2. Long-run growth in prices and factor expenditure, compared to growth in global product, for crude oil and coal, and a test of the model. In the left-hand figure we see observed prices and expenditures, with expenditures compared to global product. In the middle figure we see observed total expenditure on coal and oil, compared to the model prediction (based on the prices). And in the right-hand figure we see the observed relative factor shares of coal and oil, compared to the model prediction. In the calibrated model we have $\alpha = 0.02$, $\gamma_c/\gamma_d = 0.55$, and $\epsilon = 0.76$.


FIGURE 7.3. Long-run growth in prices and factor expenditure, compared to growth in global product, for iron and aluminium, and a test of the model. In the left-hand figure we see observed prices and expenditures, with expenditures compared to global product. In the middle figure we see observed total expenditure on iron and aluminium, compared to the model prediction (based on the prices). And in the right-hand figure we see the observed relative factor shares of iron and aluminium, compared to the model prediction. In the calibrated model we have $\alpha = 0.002$, $\gamma_c/\gamma_d = 50$, and $\epsilon = 0.55$.

matching the increase in factor share triggered by the decline in relative price of aluminium, although it could be argued that this is due to a lack of variability in the data (if the long-run trend of relative prices were more complex this would be a more discriminating test of the model).

7.2. Technological change

In the above model we assumed unbiased technological change in the resource sector. We thus rule out by construction the mechanism discussed by Solow (1973) and modelled by Acemoglu et al. (2012) whereby a resource which increases in importance (factor share) attracts more investment, and therefore increases in efficiency also. The success of our model without DTC suggests that this mechanism may be of limited importance, but it is intuitively appealing and potentially important for prediction and policy, hence we investigate further here, and again in the next part of the book.

7.2.1. The basic DTC model. Our DTC model with alternative resource inputs is pictured, in the aggregate, in Figure 7.4. Here we see that DTC occurs in a sector in which the flows C and

D from two alternative resource sectors are combined to make an intermediate good R, which we can think of as electricity. R and L are then combined in a Cobb–Douglas production function to make the final good Y. At the firm level, investments z_C and z_D are determined by the relative factor shares of C and D in the electricity sector. This is reflected in a new version of equation 5.7:

$$\frac{Z_c}{Z_d} = \frac{w_C C}{w_D D} = \left(\frac{A_C C}{A_D D}\right)^{\epsilon}. \tag{7.2}$$
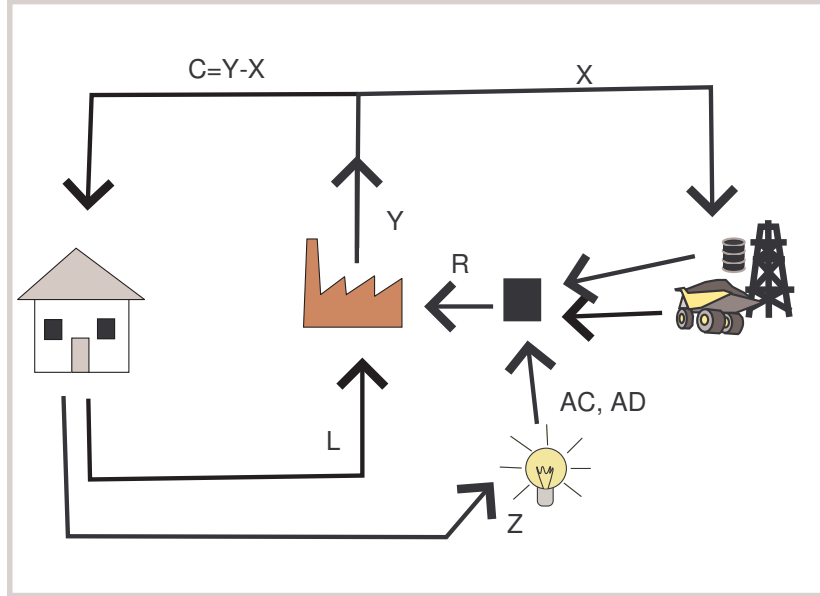

FIGURE 7.4. The aggregate flows of factors and products in the economy.

The big difference from the model of Chapter 5 is that in Chapter 5 the resource and labour inputs are complements, implying that $\epsilon < 0$ and hence an abundant factor earns a low share. However, the alternative resource inputs in equation 7.2 are substitutes, hence ϵ is positive and an abundant factor earns a big share. Think of spruce and pine trees in the latter case; if there is a lot of spruce and little pine the price of pine might be a little bit higher than otherwise, but the share of expenditure on spruce will be high compared to pine (spruce won't be *that* cheap).

If we add the assumption that knowledge stocks grow independently then we have (corresponding to equation 5.8)

$$\frac{A_{ct}/A_{ct-1}}{A_{dt}/A_{dt-1}} = \left(\frac{z_{ct}}{z_{dt}}\right)^{\phi} \left(\frac{\zeta_d}{\zeta_c}\right).$$

Putting it all together we have

$$\frac{A_{ct}/A_{ct-1}}{A_{dt}/A_{dt-1}} = \left(\frac{w_{ct} C_t}{w_{dt} D_t}\right)^{\phi} \left(\frac{\zeta_d}{\zeta_c}\right). \tag{7.3}$$

In order to say more about the development of the economy we must specify how the inputs C and D are supplied. For instance, we could assume that they are supplied in fixed (exogenous) quantities, and that the price is then determined by the market. Alternatively, we could assume that they are supplied at fixed prices, with the quantity then determined by the market. Or we could assume supply functions (linking price and quantity).

If we assume that quantities are exogenous, then all we need to do is to find an expression for relative factor costs as a function of quantities and the state of knowledge. But we already have such an expression, equation 7.2. Substitute this into 7.3 to obtain

$$\frac{A_{ct}/A_{ct-1}}{A_{dt}/A_{dt-1}} = \left(\frac{A_{ct} C_t}{A_{dt} D_t}\right)^{\epsilon\phi} \left(\frac{\zeta_d}{\zeta_c}\right).$$

Finally, multiply both sides by $\left(\frac{A_{ct}/A_{ct-1}}{A_{dt}/A_{dt-1}}\right)^{-\epsilon\phi}$ to obtain

$$\frac{A_{ct}/A_{ct-1}}{A_{dt}/A_{dt-1}} = \left[\left(\frac{A_{ct-1} C_t}{A_{dt-1} D_t}\right)^{\epsilon\phi} \left(\frac{\zeta_d}{\zeta_c}\right)\right]^{1/(1-\epsilon\phi)}.$$

Since $\epsilon > 0$ this implies that the factor which starts off more abundant earns a greater share, hence its abundance (after allowing for factor-augmenting knowledge) tends to increase! This process accelerates over time, so the economy heads for a corner in which the initially abundant factor dominates completely. (Note that we have ignored the role of the relative productivities of research. There is no obvious reason to suppose that these should differ, and if they do not differ then the term ζ_d/ζ_c disappears.)

Now assume instead that prices are exogenous. Thinking about non-renewable resources, this assumption makes more sense than the assumption of exogenous quantities; recall that it implies that the alternative resources are available to the final-good sector (in any quantity) at some exogenous price level. Of course, in the short-run we know that a sudden increase in demand will lead to a steep rise in price, but in the long run it is reasonable to suppose that prices are close to unit extraction costs in most cases, and that unit extraction costs do not vary greatly in total quantity.[3]

When prices are exogenous we have, from 7.2, that

$$\frac{w_{ct}C_t}{w_{dt}D_t} = \left(\frac{A_{ct}}{A_{dt}}\right)^{\epsilon/(1-\epsilon)}\left(\frac{w_{ct}}{w_{dt}}\right)^{-\epsilon/(1-\epsilon)}.$$

Substitute this into 7.3—and assume that the research productivities are equal—to obtain

$$\frac{A_{ct}/A_{ct-1}}{A_{dt}/A_{dt-1}} = \left(\frac{A_{ct}/w_{ct}}{A_{dt}/w_{dt}}\right)^{\epsilon\phi/(1-\epsilon)}.$$

Finally, multiply both sides by $\left(\frac{A_{ct}/A_{ct-1}}{A_{dt}/A_{dt-1}}\right)^{-\epsilon\phi/(1-\epsilon)}$ to obtain

$$\frac{A_{ct}/A_{ct-1}}{A_{dt}/A_{dt-1}} = \left(\frac{A_{ct-1}/w_{ct}}{A_{dt-1}/w_{dt}}\right)^{\epsilon\phi/(1-\epsilon(1+\phi))}.$$

Now the input with the higher ratio of productivity to price (effectively, the cheaper input) takes the higher factor share, and thus attracts more investment, and thus dominates more and more over time, and the economy heads to a corner in which only one of the inputs is used.

In terms of the balanced growth path determined in the previous chapter (for $\epsilon < 0$), such a path also exists in the case of $\epsilon > 0$, but it is not stable. On a b.g.p. we have

$$\left(\frac{A_{ct}C_t}{A_{dt}D_t}\right)^{\epsilon\phi}\frac{\zeta_d}{\zeta_c} = 1,$$

implying that the more abundant input is harder to augment. However, the situation is as illustrated in the picture below when the ball is balanced at the top of the hill; the slightest disturbance and it will start rolling one way or the other. See Figure 7.5


FIGURE 7.5. Illustration of how relative prices (the shape of the economic landscape) determine the relative levels of technology augmenting clean and dirty inputs in the model, and the role of a regulator.

7.2.2. Implications of the DTC model. The basic DTC model with independent knowledge stocks and fixed exogenous prices implies that the economy should head to a corner in which one or the other input dominates completely. Furthermore, even if prices vary (exogenously) the economy is still likely to head for a corner. Once the economy is close to a corner—such that one input is very dominant technologically—then even a radical fall in price of the other input will not shift the economy towards the other corner, as long as the technological advantage of the first input is larger.

Now assume that the first input—which has dominated for 100 years—is found to be running out. This causes its price to rise steeply, and potentially without bound. The second input must be used. However, it will take 100 years of investment to bring the second input's productivity up to

---

[3]For instance, the cost of extracting a ton of coal from the Blackwater mine in Queensland, Australia is not affected by the global extraction rate of coal.

the level of the first input, if we assume that the total quantity of research investment is constant over time. So the switch from one input to the other will be enormously costly in terms of lost production. Finally, if a completely new resource appears on the market, there will be no market for it, since there will be no technology complementing that resource initially, implying that its price in efficiency units (the price of a unit of $A_c C$ for instance) will be infinite, there will be zero demand for the resource, and hence also zero investment in technology augmenting that resource.

Fortunately, and clearly, the simple model with independent knowledge stocks is not applicable to empirical cases. This is demonstrated by the data presented in Section 7.1, where we see that substitute resources coexist rather than outcompeting one another. Furthermore, when a new resource appears (oil, aluminium) it rapidly takes market share rather than being 'locked out' by the incumbent resource. Thus we must refine the DTC model if we are still convinced that the DTC approach is potentially relevant.

7.3. An alternative to lock-in: The Fundamentalist economy

The above aggregate evidence is consistent with a model in which knowledge stocks are linked: new knowledge boosting the productivity of alternative energy inputs is produced not simply using existing knowledge regarding that input, but also using overall general knowledge, and also knowledge which is specific to the use of other energy inputs. Recall the discussion of the previous chapter (Section 5.4.2) regarding links between knowledge stocks and wind power. When the interest in wind power rose during the late 20th century, researchers did not turn to windmill designs from the 19th century. Neither did they perform their research into new designs using 19th-century techniques. They harnessed the power of the technological progress made between 1900 and 1990 in order to make very rapid progress regarding the productivity of power generation from wind. Much of the knowledge they used was general to the whole economy—for instance computers—whereas some would have been rather specific to other alternative power sectors, such as electric turbines.

A simple alternative to the lock-in/break-out economy is an economy in which relative knowledge stocks grow at equal rates, based on growth in overall general knowledge. This alternative has an obvious drawback, however, which is that is seems to rule out the technology transitions which we do actually observe, such as towards the use of oil and aluminium in the 20th century. The following model—the Fundamentalist economy—captures such transitions in a simple way.

In the Fundamentalist economy we introduce a new distinction between knowledge stocks k_c and k_d and productivities A_c and A_d, and the dynamics are driven by the difference between resource-specific parameters $\bar{k}_c$ and $\bar{k}_d$, which represent the degree of technological sophistication required to make use of each resource. Production of y_c is a function of input productivity A_c,

$$y_c = \gamma_c A_c q_c,$$

and input productivity is a function of input-related knowledge k_c and $\bar{k}_c$,

$$A_c = k_c (1 - \bar{k}_c / k_c)^{1/\omega_c} \text{ for } k_c > \bar{k}_c, \text{ otherwise } k_c = 0.$$

The parameter $\omega_c > 0$. Symmetric expressions apply for input D. For simplicity we completely short-circuit the process of DTC by assuming that there is only one type of investment z_r, and it boosts both types of knowledge equally. Since k_c and k_d are equal we define $k_{ct} = k_{dt} = k_{rt}$, and in equilibrium

$$k_{rt+1} = k_{rt} z_{rt+1}^{\phi} / \zeta_r. \tag{7.4}$$

In this economy there will therefore be no path-dependence or lock-in.

The dynamics of resource productivity are as follows. If $k_c < \bar{k}_c$ the productivity of input C is zero; technology is too primitive to make any use of the input. However, since k_c rises at a constant rate θ then at some point we have $k_c = \bar{k}_c$, and the productivity of the input rises above zero beyond this point. The initial rate of increase will be very large, approaching θ asymptotically from above. The rate of approach will depend on ω_c; in the limit as ω_c approaches infinity, A_c jumps straight to k_c as soon as $k_c > \bar{k}_c$, whereas when $\omega_c \to 0$ the rate of approach becomes slow. The productivity of resource D will follow a similar pattern, but the timing will be different if $\bar{k}_c \neq \bar{k}_d$.

The model based on technological fundamentals does a much better job of explaining the data than the lock-in / break-out model. According to this model, oil and aluminium demand a higher level of technology in order to be used productively, and once the overall economy has reached this level then they rapidly take their place alongside the other inputs (including coal and iron). The role for directed investments is limited.

If the relative productivities are governed by a process such as that in the Fundamentalist economy then policies to encourage directed research efforts are likely to be a waste of time. Assume for instance that solar PV is a technology with a high technology threshold $\bar{k}_{pv}$, but that this threshold has now been passed and hence that a_{pv} is approaching k_{pv}, and productivity growth in the sector is high. Now, if k_{pv} is high enough then solar power will soon take over from fossil power; on the other hand, if k_{pv} is not high enough then solar power will remain small relative to fossil power as long as the price of fossil power is not raised relative to solar, through for instance taxation of CO_2 emissions. So in the fundamentalist economy emissions taxes are the key instrument to yield a technology transition, not research subsidies.

Finally, note that we have in no way proved the suitability of the fundamentalist model as a description of the economy and basis for policy. We have simply presented evidence suggesting that it is a much more promising alternative than the DTC model with independent knowledge stocks.

7.4. Concluding discussion on resource stocks in the very long run

Substitutability between resource inputs has major implications for resource scarcity in the very long run, which we analysed in Chapter 4. Goeller and Weinberg (1976) argue that reserves of the majority of minerals—including staples such as iron and aluminium—are so vast that they can never realistically be consumed. With the clear exception of phosphorus they argue that society can exist on these superabundant resources indefinitely. Of course, fossil fuels are certain to run out in the not-too-distant future—unless we restrict their use for the sake of the climate—but here there are obvious substitutes including the abundant inflow of energy from the sun. Goeller and Weinberg are well aware that resource stocks are inhomogeneous. However, they show that extraction costs are relatively insensitive to grade and depth of the resource deposits, since the physical extraction and sorting of the mineral-rich material is only part of the process, and generally not the most expensive. Instead, the energy-intensive reduction of the metal ores to the pure form is typically a major part of the costs. This shows that we should—ideally—include energy in the extraction function for minerals, and hence that future energy prices will be important determinants of mineral prices. This applies particularly since the energy-efficiency of the reduction process is already close to the physical limit of what is possible.

When studying the very long run, the limiting effect of the surface area of the globe is crucial. This limit obviously puts a brake on indefinite (exponential) population growth, and also resource extraction. Thus it is clear that the current trend of constant growth in global resource extraction—tracking global product—cannot continue indefinitely. What then will make extraction flatten out or even turn downwards, and when?

Before considering the above questions, we think about the likely nature of a (very) long-run growth path. Given the constant flow of energy inputs (from the sun) and the fixed resources on Earth, it seems reasonable to suppose that use of energy will be approximately constant in the very long run. This implies that some fixed proportion of land will be devoted to energy harvesting. The remainder of the land will presumably be devoted—in fixed proportions—to other uses such as agriculture, living space, production, recreation, and (hopefully) nature. Given a fixed energy flow, what about minerals? Since we are already close to the limits of the (energy) efficiency of mineral extraction—and the energy costs of this extraction are large—it is clear that a long-run growth path must involve non-increasing mineral extraction.

So in the very long run we should expect an economy with roughly constant use of minerals, energy, and land. Furthermore, the productivity of energy in making goods such as motive power and light will be constant, as will the productivity of land in making (or harvesting) energy. On the other hand, the productivity of labour may still be able to increase even in the very long run, as it is hard to envisage a limit on human ingenuity and hence our ability to generate more value from given inputs. The key resource for energy production will be land, indeed land will be the ultimate scarce resource, strictly limited and needed for harvesting energy, production (including food), recreation, and (hopefully) nature.

Pollution

We now return to the simplest model of substitution between inputs from the previous chapter, and apply it to the analysis of polluting flows, which are linked to use of natural-resource inputs. Growth in pollution flows is driven by increasing demand for such inputs (the prices of which have no long-run trend), while dramatic falls in pollution occur when firms switch to cleaner inputs or production processes. These switches are triggered by increasing marginal pollution damages, linked to increasing income. We investigate the relevance of the model and explore possible extensions using evidence regarding emissions of SO_2 and CO_2.

8.1. Empirical observations and literature

Before we turn to the model, we return to the empirical observations of Chapter 1, especially Figure 1.4, reproduced here (Figure 8.1).


FIGURE 8.1. UK Sulphur emissions compared to total UK GDP, and global CFC production (CFC11+CFC12) compared to total global product. Sulphur: both normalized to zero in 1956, the date of introduction of the first of a long series of regulations restricting emissions. CFCs: both normalized to zero in 1987, the date of signing of the Montreal protocol. Data: Maddison (2010) (GDP), Stern (2005) (Sulphur), AFEAS (CFCs). AFEAS data downloaded from http://www.afeas.org/data.php, 9 Nov. 2014. Two anomalous points in the sulphur data have been altered.

The pattern we see in these data is repeated over and over again in countries across the world, for many different pollutants. Panayotou (1993) described this phenomenon as the environmental Kuznets curve (henceforth EKC), and Grossman and Krueger (1995) is the seminal work. In the empirically oriented EKC literature there is strong support for the idea that when the flow of a single pollutant in a single country is plotted against time, that flow will in most cases first tend to rise, and later (if enough time has passed) decline. See for instance Grossman and Krueger (1995) and Selden et al. (1999). However, if we compare paths for the same pollutant across different countries, it is hard to find clear patterns; the turning point is neither at a given time, nor at a given level of per-capita GDP. For instance, Stern (2004) concludes [p1435] that '[t]here is little evidence for a common inverted U-shaped pathway that countries follow as their income rises'.

Despite more than 20 years of research, there is still no widely accepted theoretical explanation for the phenomenon. The reason for this is that researchers building theoretical models have fallen into the trap of treating pollution as an input to production, rather than as a by-product of the use of

natural resources; see for instance Stokey (1998), Andreoni and Levinson (2001), Brock and Taylor (2010), Smulders et al. (2011), and Figueroa and Pastén (2015). In doing so they are following a tradition going back at least as far as Baumol and Oates (1975). If we think of pollution as an input in a Cobb–Douglas production function, then (using our earlier work) we know that the factor share of pollution must be constant. And if we let the marginal damage caused by pollution track income (a natural assumption) then the flow of pollution should be constant as the economy grows. This is like our DHSS-style model with a fixed quantity of land, in which the price of land tracks the growth rate. Except that here it is the price which is tied to the growth rate, and this leads (endogenously) to a constant flow of pollution.

When we treat pollution as a by-product of natural-resource use (following Murty et al. (2012)), the analysis changes completely, as we see in the next section. The social costs of natural resource use are then the sum of extraction costs and the damage costs of the concomitant pollution. At low income the pollution damages are small and the (constant) extraction cost dominates. And because the natural resource is an input in a Cobb–Douglas production function, natural resource consumption increases with growth, as do polluting emissions. As income increases, so does the WTP to avoid pollution. The social cost of natural-resource use starts to rise, and resource use levels off. However, more importantly, if there is a cleaner (but more expensive) alternative resource, there will come a point at which this resource is preferred, and pollution falls dramatically.

8.2. The specified model

We now develop a specified model economy to demonstrate the mechanism.

8.2.1. The environment. There is a unit mass of competitive firms which produce a single aggregate final good the price of which is normalized to 1. Both the firms and the population L are spread uniformly over a unit area of land. The production function of the representative firm in symmetric equilibrium hiring labour L (productivity A_L) and buying a resource-intensive intermediate input R is

$$Y(t) = [A_L(t)L(t)]^{1-\alpha} R(t)^{\alpha} e^{-P(t)^{\phi}}, \tag{8.1}$$

where α is the share of the intermediate input, which is small, P is the aggregate flow of pollution—which is uniformly mixed—and ϕ is a parameter greater than 1. Both A_L and L are exogenously given, and $A_L L$, effective labour, grows at a constant rate g:

$$\dot{A}_L(t)/A_L(t) + \dot{L}(t)/L(t) = g.$$

From now on we omit the time index whenever possible.

Intermediate production R—which we can think of as, for instance, electricity—is the sum of inputs from n different resource-based technologies, which are all perfect substitutes in production. The quantity of input from technology j is denoted D_j, so

$$R = \sum_{j=1}^{n} D_j.$$

The use of input quantity D_j leads to emission of pollution $\psi_j D_j$, where $\psi_j \geq 0$, hence aggregate pollution

$$P = \sum_{j=1}^{n} \psi_j D_j.$$

The cost of a unit of input j is w_j.

We can interpret alternative technologies j and k simply as alternative resource inputs, for instance low- and high-sulfur coal for electricity generation. However, a third technology l could be high-sulfur coal combined with flue-gas desulfurization (FGD). If the input is simply a natural resource then we can think of it as being extracted competitively from a large homogeneous stock, with each unit extracted requiring w_j units of final good as input. But for technology l the price w_l would be w_k plus the unit cost of FGD, and unit emissions ψ_l would be $\psi_k \times$ the fraction remaining after FGD.

We denote aggregate production net of extraction costs as Z, so

$$Z = (A_L L)^{1-\alpha} \left(\sum_{j=1}^{n} D_j \right)^{\alpha} e^{-(\sum_{j=1}^{n} \psi_j D_j)^{\phi}} - \sum_{j=1}^{n} w_j D_j. \tag{8.2}$$

8.2.2. The solution. In solving the model we focus throughout on the social planner's solution; given this solution the regulatory problem is straightforward. Furthermore, we focus mainly on a model with a choice between just two technologies, because this gives the clearest intuition. Given two technologies, the planner chooses the set of values (D_1, D_2) to maximize Z (equation 8.2). Take the first-order conditions on equation 8.2 in D_1 and D_2 respectively to derive the following necessary conditions for an internal optimum.

$$\text{FOC } D_1: \qquad \alpha Y/(D_1 + D_2) = w_1 + \phi(\psi_1 D_1 + \psi_2 D_2)^{\phi-1}\psi_1 Y. \qquad (8.3)$$

$$\text{And FOC } D_2: \qquad \alpha Y/(D_1 + D_2) = w_2 + \phi(\psi_1 D_1 + \psi_2 D_2)^{\phi-1}\psi_2 Y. \qquad (8.4)$$

In these equations, the marginal societal benefits of making an extra unit of intermediate good R ("electricity") using technology j are on the left-hand side, and the marginal costs are on the right-hand side. The marginal benefits are identical whether we use input 1 or 2 to make R, but the marginal costs differ. The costs are the sum of the natural-resource input costs w_j and the pollution damage costs $\phi P^{\phi-1}\psi_j Y$.

To build intuition we start with the case in which $w_1 < w_2$ and $\psi_1 < \psi_2$, so D_1 is both cheaper and cleaner, and D_2 will never be used. Then we have the following proposition.

PROPOSITION 1. *When only input D_1 is used, from any given initial state (defined by $A_L(0)L(0)$), P increases monotonically and approaches a limit of $\bar{P} = (\alpha/\phi)^{1/\phi}$. If we let $A_L(0)L(0)$ approach zero then the initial growth rate of P approaches g from below.*

PROOF. See 8.C.1. □

The interpretation is as follows. The shadow price of the polluting input to the social planner is the sum of extraction cost and marginal damages. The extraction cost is constant, whereas marginal damages increase linearly in Y. So when Y is small the shadow price is approximately equal to the constant extraction cost, and both resource use and polluting emissions track growth. As Y increases, marginal damages increase and hence the shadow price of using the polluting input increases, braking the growth in its use. When Y is large marginal damages dominate the extraction cost, the shadow price of using the input grows at the overall growth rate, and emissions (and input use) are constant. So we have a transition from emissions tracking growth towards (in the limit) constant emissions.

Now we take the more interesting case when technology 2 is more expensive but cleaner, i.e. $\psi_1 > \psi_2$.[1] In this case, as Y increases, the increasing importance of pollution damages does not just lead to pollution abatement within technology 1—i.e. the substitution of labour–capital for D_1 in production—it also narrows the gap between the social costs of D_1 (cheap and dirty) and D_2 (expensive but cleaner). At some point the social costs are equal, and a transition to the cleaner technology begins.

PROPOSITION 2. *In a two-technology economy, there exist times T_{1a} and T_{1b} (where $T_{1b} > T_{1a}$) such that up to T_{1a}, D_1 increases monotonically while $D_2 = 0$. Between T_{1a} and T_{1b}, D_1 decreases monotonically while D_2 increases monotonically. And for $t \geq T_{1b}$, $D_1 = 0$ and D_2 increases monotonically. Furthermore, T_{1a} and T_{1b} can be expressed in closed form. In the special case of $\psi_2 = 0$ (the cleaner resource is perfectly clean) then T_{1b} is not defined; instead, as $t \to \infty$, $D_1 \to 0$, and hence $P \to 0$.*

PROOF. See 8.C.2. □

It is straightforward to extend Proposition 2 to the case of multiple technologies which differ in cost and polluting emissions: Proposition 3.

PROPOSITION 3. *In an n-technology economy there is a series of m transitions (where $m \leq n-1$), starting with the cheapest input and ending with the cleanest. Each of these transitions proceeds analogously to the transition from 1 to 2 described in Proposition 2.*

PROOF. See 8.C.3, where an n-technology economy is also precisely defined. □

In Figure 8.2 we illustrate the development of the economy in a specific case with three technologies, the third of which is perfectly clean. Since $\psi_3 = 0$ the second transition is completed asymptotically, and as $t \to \infty$, $P \to 0$. In Figure 8.2 we show the paths of effective labour $A_L L$, and pollution P, and the pollution limit $\bar{P}$. We also show—using dotted lines—the paths of P which would be followed if (respectively) only technologies 1 and 2 were available.

---

[1]Furthermore, to ensure unambiguous results we require that $(\phi-1)(1+\alpha)/\alpha > \psi_1^2(w_2 - w_1)/(w_2\psi_1 - w_1\psi_2)$.
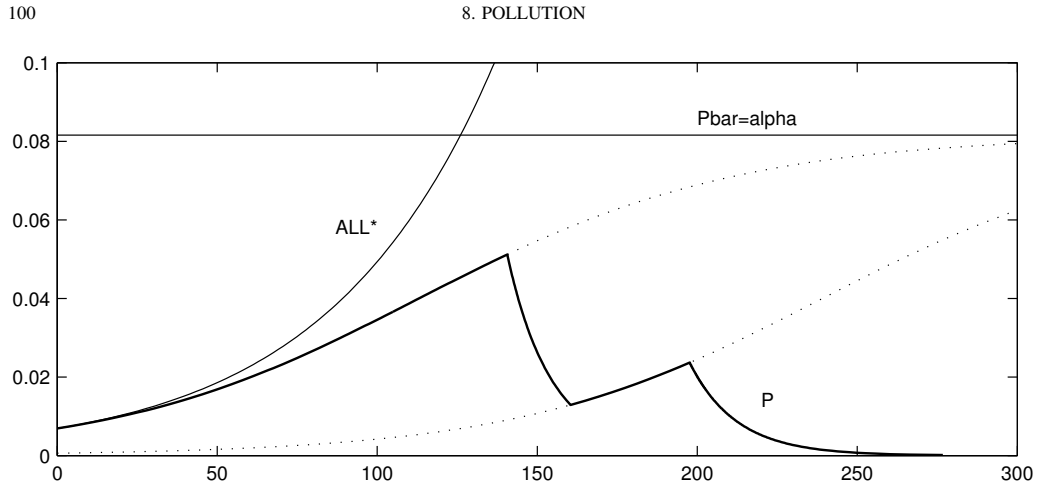
FIGURE 8.2. Pollution flow P compared to the limit, $\overline{P} = \alpha$, and normalized effective labour $(A_L L)^*$, where $(A_L L)^* = A_L L / [A_L(0)L(0)] \cdot P(0)$. Parameters: $g = 0.02$; $A_L(0)L(0) = 1$; $\phi = 1.3$; $\psi_1 = 0.0072$, $\psi_2 = \psi_1/6$, $\psi_3 = 0$; $\alpha = 0.05$; $w_1 = \alpha$, $w_2 = 2\alpha$, $w_3 = 2.5\alpha$. The dotted lines show pollution paths in case only one of the inputs is available.

8.3. A graphical treatment

We now turn to a graphical treatment of the problem, consider how the optimal choice of polluting emissions P and net production X changes over time, as technology improves. To tackle the problem graphically we need to define two sets of curves in (P, X) space, the first of which is the set of production possibility frontiers (PPFs), and the second of which is the set of indifference curves.

Each PPF shows the maximum amount of net production which is possible for each quantity of pollution emitted, at a given level of productivity. If polluting emissions were an input then we would expect the PPFs to be upward sloping: the more pollution, the more production is possible. See Figure 8.3(a).

However, since pollution is actually a by-product of natural resource use, the PPFs are hump-shaped. For given natural-resource prices and technology, there is some amount of natural resource input which will yield maximal production (and some level of pollution flow). Using more natural-resource inputs than this amount will be wasteful and lead to less net production (because firms are spending too much time extracting costly natural resources and not enough making valuable goods) and higher pollution. And using less natural resources will lead to less net production and less pollution. See Figure 8.3(b).

Finally, their are many PPFs, for different levels of productivity. As productivity increases, the PPF moves outwards: the capacity of the economy to both make final goods and to extract natural resources (leading to pollution) increases. See Figure 8.3(c).

FIGURE 8.3. PPFs. The PPF in (a) is not allowed because there is no turning point; the PPF in (b) is allowed; in (c) we see a set of PPFs for different productivity levels.

Now we turn to the indifference curves. The indifference curves are derived from the utility function. Utility U is a CES function of consumption X and environmental quality Q:

$$U = \left\{ (1-\alpha)X^{(\eta-1)/\eta} + \alpha Q^{(\eta-1)/\eta} \right\}^{\eta/(\eta-1)}, \tag{8.5}$$

where $\eta > 0$. Furthermore, $Q = 1/[d(P)]$, where d is the damage function, which is differentiable and strictly increasing, and $d(0) > 0$. Since $U'_X > 0$, we can also define the equation for the

indifference curves in (X, P) space:

$$X = V(U, P) = \left[U^{(\eta-1)/\eta}/(1-\alpha) - \alpha/(1-\alpha) \cdot d(P)^{(1-\eta)/\eta} \right]^{\eta/(\eta-1)}. \tag{8.6}$$

Furthermore, we impose one further restriction on the utility function, which is that $V''_P > 0$, implying that the indifference curves are strictly convex and guaranteeing a unique solution to the problem of maximizing utility at given A. The indifference curves then have the following three properties:

(1) When $X \to 0$, $V'_P \to 0$ for all P, so the price of pollution is zero when consumption is zero;

(2) V'_P increases monotonically in X for any $P > 0$, so the price of pollution increases with consumption;

(3) When $X \to \infty$, $V'_P \to \infty$ (as long as $P > 0$), so the price of pollution approaches infinity when consumption approaches infinity.

Returning to the graphical representation, we show an example of an allowed set of curves, and two that are ruled out, in Figure 8.4.


FIGURE 8.4. Three sets of indifference curves. The second is ruled out because X/P does not increase in X, implying that the WTP to remove a unit of pollution does not increase in income, and the third is ruled out because the curves are not strictly convex.

Putting the set of PPFs and the indifference curves together it is clear by inspection that pollution must first rise and then fall as long as (i) initial productivity is low enough (so the lowest PPF is sufficiently close to the origin) and (ii) strictly positive production is possible with zero polluting emissions. See Figure 8.5.
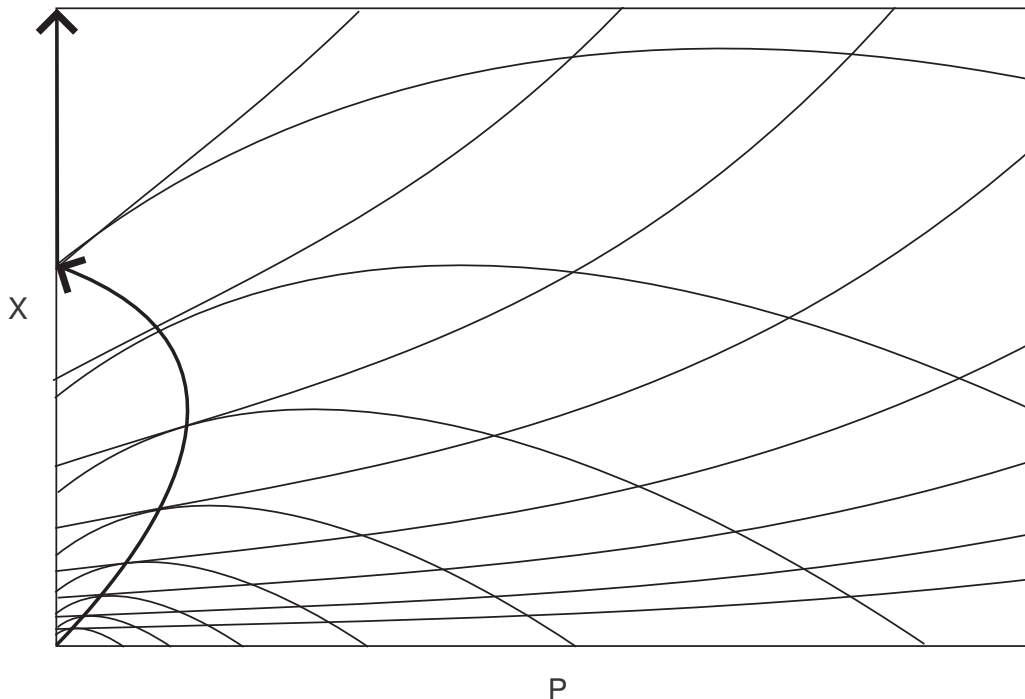

FIGURE 8.5. An illustration of the rise and fall of P as productivity increases and the PPF moves outwards.

8.4. Discussion

Here we discuss the generality of the model and possible other causes of the EKC.

8.4.1. The utility function. We postulated the CES utility function with very little discussion. Crucially, it implies that WTP for higher environmental quality Q approaches zero when income approaches zero, and approaches infinity when Q is bounded above and income approaches infinity. (These properties are all that are needed to generate the key results, the assumption of CES is made to rule out confounding mechanisms, similarly to the assumption of constant returns in the ppf.) Here we argue that these assumptions are very mild. It is hard to see how WTP for lower pollution flows P could fail to approach zero as long as $Q > 0$ and income approaches zero, and similarly it is hard to see how WTP for lower P could fail to approach infinity as long as Q is bounded above and income approaches infinity. However, there seems to be remarkably little research which systematically studies the WTP to reduce pollution or increase environmental quality as a function of income; for one example see Jacobsen and Hanley (2009). Regarding the exact specification of the utility function, the most common assumption in the non-EKC literature is that marginal damages from a given change in Q are proportional to GDP; see for instance climate models such as Nordhaus (2008) and Golosov et al. (2014), and the study of SO_2 policy of Finus and Tjotta (2003).

8.4.2. Is production possible with zero pollution? Is production possible with zero pollution? We analyse this question in two steps. In the first step we assume a situation in which pollutants are entirely independent of one another in the sense that there are no trade-offs between them: cutting one pollutant (such as SO_2) never leads to increased emissions of others (particulates, CO_2, etc.). If this condition holds then the limiting substitutability between a given pollutant and labour–capital should be high, because in most cases of polluting technologies there exists an alternative technology which is (a) a very good substitute for the polluting technology (albeit more expensive), and (b) clean (i.e. emissions of the pollutant in question are zero). If coal for electricity generation is emitting SO_2, we can use gas instead. If lead in gasoline is finding its way into our lungs and subsequently damaging our brains, we can use lead-free. If CFCs are destroying the ozone layer, we can use HFCs instead. More specifically, in many cases we expect the ppf the meet the $P = 0$ axis at positive X, and indeed at a level of X not much below $\bar{X}$.[2] On the other hand, there are of course cases where abatement is incremental and does not involve a single radical switch of technology and the ppf will be curved in the relevant segment, indicating a non-infinite elasticity of substitution between A and P: a good example are emissions of nitrogen and phosphorous to water, which primarily come from agriculture and sewage treatment, and where abatement consists of many incremental changes in technology; another example is CO_2, which has many different sources, such as road transport, air transport, electricity generation, cement production, etc., each with different abatement costs. However, even in these cases it is clear that emissions could be reduced to zero (or in the case of nitrogen and phosphorus, to natural levels) while retaining positive consumption X.

In the second step we assume instead that polluting emissions are linked to one another. There may be options which reduce a whole range of emissions (such as switching from coal to gas or renewables in energy production), but ultimately there will be trade-offs: for instance, switching to renewable energy may lead to greater noise and visual pollution. Effectively then we are defining pollution as any human-induced loss of environmental quality relative to the natural state. If we define the natural state as pristine, with perfect environmental quality, then of course it will never be achieved in the future. However, if we accept the idea the humans may actually *improve* the natural environment then the very-long-run outcome depends on whether this improvement can be achieved without sacrificing all consumption of non-environmental goods X.

8.4.3. A common pathway? The ideal next step would be to specify and calibrate a model—more general than that of Section 8.2—which could account for patterns of polluting emissions

---

[2]Consider for instance the case of lead emissions to air. Assume that they come exclusively from burning petrol in automobiles (since lead in petrol was banned, emissions have fallen by close to 100 percent, as described in the US EPA website for information on lead emissions to air, `https://www.epa.gov/lead-air-pollution/basic-information-about-lead-air-pollution`. Lead in petrol is a cheap way to achieve a high octane rating, desirable to allow a high compression ratio in the motor and hence more power. Assume that the same effect can be achieved at the expense of a 5 percent increase in fuel price. Now assume that the cost of petrol accounts for 1 percent of GDP. Then the difference between $\bar{X}$ (when lead is used without regard to emissions) and $\underline{X}$ (when lead emissions are zero) is approximately 0.05 percent, so lead emissions to the atmosphere (and consequent brain damage suffered especially by children) can be avoided at a cost of 0.05 percent of GDP, and the segment of the ppf in the segment between the $P = 0$ axis and the turning-point at $(\bar{P}, \bar{X})$ is a straight and almost horizontal line.

from different countries given panel data on countries' GDP, resource use, resource prices, etc. However, it is an almost hopeless endeavour to specify a model of national emissions for any given pollutant which can be calibrated and applied to explain patterns of aggregate polluting emissions in heterogeneous countries, because of the idiosyncratic nature of the forces driving resource use, technology choice, and consequent emissions in different countries and over time. For a specific example of the kind of idiosyncrasies that may be relevant, consider sulfur emissions to the atmosphere in the U.K. and the U.S. In the U.K. there has been a rapid decline in SO_2 emissions since 1960 (Figure 8.6(a)), driven mainly by the replacement of coal by oil and gas in the overall energy mix (Figure 8.6(b)). This shift was partly driven by the increase in road transport, but also by the 'dash for gas' in electricity generation, driven in turn by a steep decline in the price of gas relative to coal (Figure 8.6(c)). In the U.S., sulfur emissions started to decline in the mid-1970s (see for instance Stern (2005)), at least partly due to the introduction of the clean air act in 1970. However, Ellerman and Montero (1998) demonstrate that the steep decline in sulfur emissions was facilitated by the significant fall in transport costs of coal which occurred subsequent to the deregulation of the railroads in the 1980s, which reduced the cost of shipping coal from the Powder River Basin; this coal is both the cheapest and cleanest in the U.S.
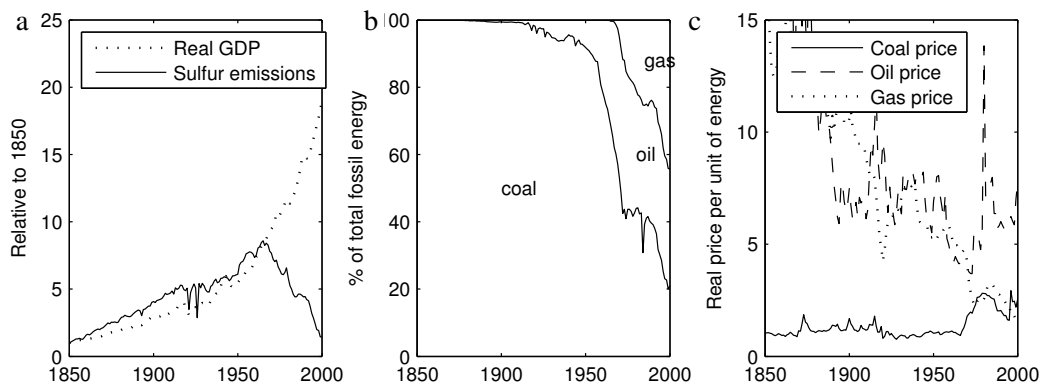

FIGURE 8.6. Fossil fuels and SO_2 in the UK since 1850: (a) Emissions track GDP, then decline abruptly after 1960; (b) Oil and then gas eat into the share of coal in total fossil energy supply, with the major shift starting in 1960; (c) The gas price approaches the coal price over time, whereas the price path for oil is more complex (prices are normalized relative to the coal price in 1850). Data sources: fossil-fuel consumption from Warde (2007); prices from Fouquet (2011); sulfur emissions from Stern (2005); real GDP from Maddison (2010).

So although the overall mechanism of the model is relevant in driving the rise and fall of polluting emissions identified by Grossman and Krueger (1995), heterogeneity between countries and over time implies that no 'common pathway' followed by different countries can be identified in the data (cf. Stern, 2004), and a much richer model than that developed above is required if it is to be calibrated to data from multiple countries. Such a model would have to explain and predict the development of aggregate demand for natural resources at country level, and the technologies applied when those natural resources are used as inputs in the economy. Such a task is far beyond the scope of this paper.

8.4.4. DTC, resource scarcity, structural change, etc. As discussed above, many factors may affect polluting emissions over time, factors which are not included in our model. If there are factors which are (a) important, and (b) consistently affect emissions in the same way (across countries and over time), then they should be included. However, we argue that the other factors do not affect emissions consistently across different cases, and are typically not of crucial importance. Here we list a few factors and discuss them very briefly.

DTC. DTC occurs when technology advances faster in one sector than another due to endogenous investments by firms. It is modelled by Acemoglu et al. (2012), whose focus is the transition from dirty to clean technology, and Smulders et al. (2011) argue persuasively that it is relevant to the EKC. If technological progress is more rapid in abatement technologies (or in the use of clean inputs) than in other sectors then it will tend to flatten the ppf when plotted in (p, x) space, creating downward pressure on optimal pollution. In a single-country context, adding DTC to our model would not change much: compared to a baseline in which all technologies are mature from the start, the need for 'catch-up' investment in abatement technology should delay the transition but make it more abrupt when it happens. In a multicountry context, if the other countries can free-ride

on the leading country's investment then their transition will occur earlier (at lower GDP, ceteris paribus). However, note that in the case of FGD the calibrated model above suggests a limited role for DTC, and Hart (2013, 2018b) argues that the power of the DTC mechanism developed in models such as Acemoglu et al. (2012) is exaggerated when compared to reality.

Imperfect information. There may often be imperfect information about the damages caused by polluting emissions. This lack of information may cause a delay in dealing with emissions when compared to the optimal pathway. Similar to the DTC case, when emissions are finally tackled the transition to clean technology is likely to be more abrupt if there has been a delay. Furthermore, transitions across countries are likely to be closer together in time, since the 'follower' countries can presumably learn from the leader.

Scarcity and structural change. Scarcity of a natural resource can drive up its price and cause a switch to other resources. This may cause pollution to decline or rise depending on which resource is cleaner. Generalized scarcity of natural resources will tend to brake the rate of increase in their use, and slow down pollution growth. Similar arguments apply to structural change. Generalized structural change away from resource-intensive goods will encourage a steeper decline in pollution. However, in reality both effects are likely to be more complicated: there is little evidence for generalized resource scarcity driving up prices any time soon (see Hart and Spiro (2011) and Hart (2016)), and scarcity of specific resources may push pollution either way. For instance, natural gas (low sulphur) will become increasingly scarce long before coal (high sulphur), and this will push the ppf to the right, tending to increase pollution. And Hart (2018a) shows that aggregate structural change over the last century or more has actually been *towards* energy-intensive goods, not away from them (consider train/bus → automobile → airplane), although this may change in the future.

Our overall model of resource supply and demand—with constant prices and resource demand driven by a Cobb–Douglas production function—is broadly consistent with the aggregate evidence, as discussed above. However, at country level the picture is more complex (Figure 8.7): in the U.S. there have been declines in the consumption of primary energy, metals, and cement relative to GDP since 1950. However, these declines are not large enough to account for the falls in polluting emissions, as shown by the analysis of Selden et al. (1999).


FIGURE 8.7. Long-run growth in US consumption compared to growth in US GDP, for (a) primary energy from combustion, (b) metals, and (c) cement.

8.5. Conclusions

We have presented a simple explanation for the rise and fall of polluting emissions based on a model in which such emissions are linked to use of natural-resource inputs, and alternative inputs differ in pollution intensity. The mechanism through which polluting emissions rise and fall may be applicable to a large number of relevant empirical cases. We have discussed a few—including SO_2 and lead emissions to air—but we could equally well have chosen many other cases in which the key criteria for our mechanism are fulfilled: marginal damages increase in income, emissions are linked to the use of natural resources in production, the prices of the relevant natural resources display little or no long-run trend, and alternative, cleaner, inputs or production processes are available. Phenomena not included in the model—such as directed technological change—may add important dimensions to the analysis and require specific policy instruments if an optimal allocation is to be achieved.

Our theoretical analysis is congruent with the empirical analysis of Stern (2004) and many others: we should not expect the exact pattern for one pollutant and location to be repeated for

another pollutant or location; key factors in the model which may be expected to vary—both between pollutants, and over time and space for the same pollutant—include the relative prices of alternative resource inputs or production processes, and the valuations of environmental quality at given income and pollution levels. For instance, both natural gas and coal are relatively expensive to transport, hence in some locations coal may have a price advantage over gas, whereas in other locations the reverse may be true.

Finally, the analysis has important implications for future emissions, not least carbon emissions. Carbon dioxide emissions are of course a by-product of the use of certain technologies, especially those involving the burning of fossil fuels, and there exist very good zero-carbon substitutes for these technologies. Currently many countries are pushing for a global agreement on major reductions in global emissions, whereas others are resisting such efforts for geopolitical or other reasons. Our model shows that as incomes rise in a given country, the price of switching to zero-carbon technology is likely to seem—to that country's citizens and decision-makers—increasingly like a price worth paying. This effect may be boosted by directed technological change, making the substitutes cheaper. Fossil scarcity, paradoxically, may work in the opposite direction if gas runs out before shale oil and coal, since the latter are more carbon-intensive.

8.A. Appendix: A calibrated model

The central hypothesis of this chapter is that rising income drives the imposition of environmental regulations which—in the long run—drive switches to cleaner technologies and hence falling emissions. Here we provide empirical support for this idea by showing that the timing of adoption of flue-gas desulfurization across six countries can be understood based on a model in which underlying preferences for clean air, and the unit cost of installing FGD, are constant across the countries and over time, and the timing of the imposition of the regulation is determined by income per capita, population, and the size of the territory.

We argued above that the shape of the PPF of pollution and production varies between countries, even those on the same income level, as does the shape of the indifference curves. Furthermore, biased technological change and new information may change PPFs and indifference curves over time.[3] It is therefore not possible to test the empirical relevance of the models above by looking for simple patterns such as turning points in pollution flows at given income levels. Instead of looking for patterns in emissions, we look for patterns in the application of environmental regulation, specifically the timing of adoption of FGD in Japan, the US, West Germany (as it was at the time of adoption), the UK, China, and India. FGD is a set of technologies used to remove sulfur dioxide from exhaust gases of coal-fired power plants (see US EPA (2003)). We choose it because of the readily available data about the timing of the implementation of FGD. We investigate the following hypothesis.

Hypothesis 1. *The unit costs of sulfur abatement through FGD are constant over time and across countries, and the time of introduction in a given country is determined by the marginal damage cost of sulfur emissions, which is a linear function of income per capita, and an increasing function of the size and population density of the country.*

We have data on the time of adoption (which we define as the first year when at least 5 percent of coal capacity has FGD installed), GDP per capita (from Maddison (2010)), population, and land area. The time of adoption ranges from 1970 (Japan) to 2016 (India).[45]

Ideally we would perform an econometric test of a structural model, but since we have only six observations we limit ourselves to a calibration exercise. We base the equation to be calibrated on equation 8.2, $W = X/\exp(P^\phi)$. That is, we assume multiplicative utility following the climate literature. Since we do not have data on measured pollution concentrations, we assume that $\phi = 1$, making marginal damages approximately independent of P as long as total damages are small in relation to total utility. This assumption is also in line with the literature on damages from SO_2 where log-linear damages are typically assumed.[6] We then approximate X by real GDP, which we denote Y, and convert to per capita terms (so w is per capita utility, and y per capita GDP):

$$w = ye^{-P}.$$

The next step is to think carefully about the implications of modelling different countries, which differ in surface area and population as well as GDP and polluting emissions. The concentration of pollution will (if the pollution is uniformly mixing and remains exclusively over the territory in question) be linearly related to emissions per unit of area, and damages (if they affect humans directly) should be a function of concentration. Denoting the area as H (recall that we previously

---

[3]For a specific example of the kind of idiosyncracies that may be relevant, consider sulfur emissions to the atmosphere in the UK and the US. In the UK there has been a rapid decline in SO_2 emissions since 1960, driven mainly by the replacement of coal by oil and gas in the overall energy mix. This shift was partly driven by the increase in road transport, but also by the 'dash for gas' in electricity generation, driven in turn by a steep decline in the price of gas relative to coal. In the US, sulfur emissions started to decline in the mid-1970s (see for instance Stern (2005)), at least partly due to the introduction of the clean air act in 1970. However, Ellerman and Montero (1998) demonstrate that the steep decline in sulfur emissions was facilitated by the significant fall in transport costs of coal which occurred subsequent to the deregulation of the railroads in the 1980s, which reduced the cost of shipping coal from the Powder River Basin; this coal is both the cheapest and cleanest in the US.

[4], where the values for Germany are adjusted upwards by 14 percent to reflect the difference between average German GDP and West German GDP

[5]The year of FGD introduction is taken as the first year when at least 5 percent of coal capacity has FGD installed. The sources are as follows: Maxwell et al. (1978), Figure 2; US EPA (1995), Figure 4; Taylor et al. (2005) Figure 4; Markusson (2012) Table 1 (we assume that the 5 percent threshold was reached in 1993); Wang and Hao (2012), where the text implies that implementation of FGD took off around 2005; and lastly for India, Black and Veatch (2016), one of many available documents showing that India announced a stringent FGD program to start in 2016. GDP data is taken from Maddison (2010), extrapolated for India using equivalent data from the World Bank.

[6]See for instance Muller and Mendelsohn (2007), especially equation 12 in the additional materials, and the dose–response function of Barreca et al. (2017).

normalized it to 1) we have

$$w = y \exp\left(-\frac{P/L}{H/L}\right).$$

This equation puts issues of *scale* into focus: it implies that if we replicate the economy (doubling P, L, and H but holding w and y constant) then the proportion of gross product y lost to pollution damages will remain the same. However, when we consider pollution transport it is clear that this will not in reality be the case: for an airborne pollutant, given a larger territory, a bigger proportion of emissions will land within the territory and thus cause damage there.

To account for pollution transport, we introduce a transport coefficient δ, where δ is the proportion of emissions transported out of the territory, and

$$\delta = \exp(-\theta H^{1/2}),$$

where θ is a positive parameter.[7] As $H \to 0$, $\delta \to 1$, and as $H \to \infty$, $\delta \to 0$, so for a very small territory almost all the pollution emitted leaves the territory without causing damage 'at home', whereas for a very large territory the reverse applies. So given δ we now have

$$w = y \exp\left(-(1-\delta)\frac{P/L}{H/L}\right).$$

Finally, and also related to scale, the above equation shows that when land area H increases, pollution damages decrease because the concentration of pollutant decreases. This effect should be straightforward if population and emissions are spread homogeneously over the territory. However, in reality they are spread inhomogeneously, and furthermore if the degree of inhomogeneity is an increasing function of the sparseness of population (because people concentrate in cities even in sparsely populated countries) then the effect of increasing H/L will be weakened. To allow for this possibility we introduce a parameter ω as follows:

$$w = y \exp\left(-(1-\delta)\frac{P/L}{(H/L)^\omega}\right).$$

So when $\omega = 1$ population is uniformly distributed, whereas when $\omega = 0$ overall population density has no effect because the population and electricity production are always confined to a sub-area in proportion to the size of the population. It remains to find marginal abatement benefits by differentiating wL w.r.t. P to obtain (after approximating $y = w$)

$$MAB = \phi(1-\delta)(L/H)^\omega y. \tag{8.7}$$

To calibrate the model we must find values for θ and ω. We choose θ to match the observation of Smith and Jeffrey (1975) that around 75 percent of UK emissions leave the territory, yielding $\theta = 0.826$, and implying that in the largest countries (the US and China) around 83 percent of emissions cause damage within the territory. This leaves us with ω, which we choose in order to fit the data as well as possible, i.e. we find the value of ω which yields the set of six estimates for MAB with the lowest variance. This yields $\omega = 0.524$, implying that a doubling in population density leads to an increase in marginal abatement benefits by a factor of approximately $\sqrt{2}$.

The results are illustrated in Figure 8.8, in which we show estimated marginal abatement benefits over time for each country, with a circle showing the time of FGD adoption. The cross-country variation in estimated MAB at the time of adoption gives an idea of the variation which is unexplained by the model. Note that—with the minor exception of Japan, which adopts 'early'—the countries adopt in the expected sequence and at expected times; small shifts in timing (between 0 and 3 years) would have all the other 5 countries adopting at the same level of estimated MAB. According to the estimates, both Japan and China adopt at somewhat lower benefit levels than the other four countries. These are also the two countries with the steepest rises in benefits of adoption, linked to their very high rates of economic growth at the time of adoption. In Japan this rapid growth—in both GDP and pollution flows—led to a dramatic increase in pressure for environmental improvements from the population, and the so-called 'pollution diet' of 1970; see Avenell (2012).

Figure 8.8 shows that we can rationalize most of the large differences in the time of adoption of FGD based on the model. Furthermore, inspection of the data shows that some of the simpler explanations that might be proposed are decisively rejected. For instance, there is no single level of GDP at which countries adopt FGD and thus reduce sulfur emissions. Furthermore, there is little

---

[7]The power of 1/2 follows because if the area of the territory doubles, the average distance to the border is multiplied by $\sqrt{2}$.

FIGURE 8.8. Estimated marginal abatement benefits for the six countries plotted over time, and at the time of FGD adoption.

evidence from the model that the unit costs of FGD have declined over time, thus encouraging lower-income countries to adopt at lower marginal benefit levels than the early-adopting higher-income countries.[8]

8.B. Appendix: Stokey (1998)

Our model is closely related to that of Stokey (1998). He we explain the similarities, and the differences. In Stokey the abatement function is defined in such a way that final-good production X is a Cobb–Douglas function of P and A: $X = P^\alpha A^{1-\alpha}$.[9] But since abatement is restricted to be non-negative, X reaches a corner when $X = P = A$. The ppf is thus constant returns, hence when plotted in (p, x) space it is invariant to changes in A. However, in (P, X) space we see that as A increases—which causes the WTP to reduce pollution to increase—we move from a corner solution in which P and X are both maximized, to an internal solution in which consumption is sacrificed in order to abate pollution. In Figure 8.9(a) we have Cobb–Douglas production (so demand for the polluting 'input' grows with total production) and the damage function is specified such that damages are proportional to total income. The result is that the pollution flow is constant for all internal solutions, as we see in the figure. (Stokey assumes that damages grow faster with income, yielding an EKC.) The big question raised by the analysis is why the production function has the form assumed. Why does the marginal product of pollution suddenly drop from a strictly positive level to zero? This question has not been satisfactorily answered in the literature.

The big difference between our model and that of Stokey is that in our model the restrictions on the ppf are derived as a necessary consequence of the nature of the pollution-producing process (cf. Murty et al. (2012)), whereas Stokey's ppf is simply assumed. Among other things this means that our solution is internal, and furthermore that we can easily specify empirically grounded models which are special cases of our general model. In Figure 8.9(b) we see that pollution first rises, and then falls very abruptly, due to an abrupt switch to the clean technology. The reason for the abruptness is that the technologies (clean and dirty) are perfect substitutes, while the indifference curves are almost straight lines since marginal damages are independent of the quantity of pollution.

8.C. Appendix: Proofs

---

[8]Note that, *ceteris paribus*, technological progress is not expected to drive down FGD costs. Technological progress implies that more goods can be produced using given inputs, however if it is neutral or unbiased then it will not change the relative prices of these goods. So in an economy with just two goods—an aggregate consumption good and sulfur capture through FGD—neutral technological progress implies that given inputs of labour–capital can produce more of both, but should not change the price of one relative to the other.

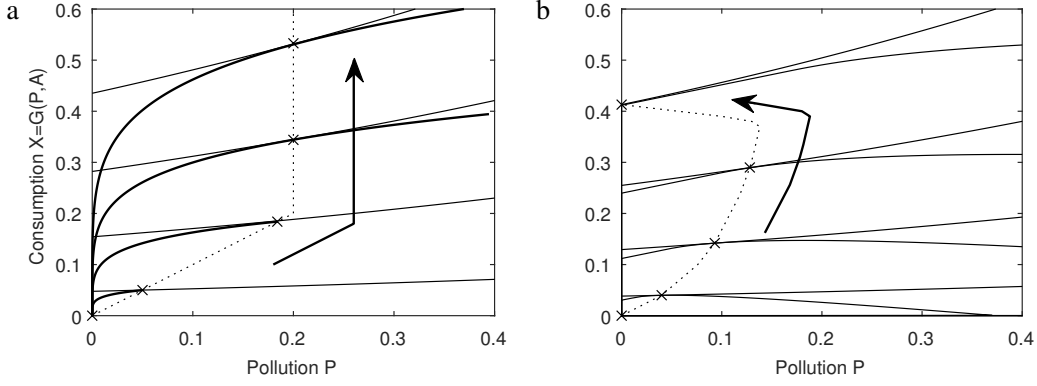[9]Although Stokey has β which is identically equal to $1/\alpha$.

FIGURE 8.9. The path of polluting emissions (a) in the Stokey model compared to (b) our model.

The utility function is the same in each case ($U = Xe^{-P}$, so $\eta = 1$). In Stokey the equation for the ppf is $X = A^{1-\alpha}P^\alpha$ up to the limit of $X = P$, whereas in our model we have the convex combination of two alternative technologies: the first is dirty, with the production function $X = A^{1-\alpha}D_1^\alpha - w_1 D_1$ where $P = D_1$; and the second is clean, with the production function $X = A^{1-\alpha}D_2^\alpha - w_2 D_2$ where $P = 0$. Parameters: $\alpha = 0.2$, $w_1 = 0.2$, $w_2 = 0.6$.

8.C.1. Proof of Proposition 1. Use equation (8.3) and the definition of Y (equation 8.1) to show that when $D_2 = 0$,

$$\left(\frac{\psi_1 A_L L}{P}\right)^{1-\alpha}\left(\alpha - \phi P^\phi\right) = w_1 e^{P^\phi}.$$

Now let $A_L L \to 0$ and show that this implies that $P \to (\alpha/w_1)^{1/(1-\alpha)}\psi_1 A_L L$, so in the limit $\dot{P}/P = g$. Then let $A_L L \to \infty$ and show that this implies that $P \to (\alpha/\phi)^{1/\phi}$.

8.C.2. Proof of Proposition 2. First we state the closed-form solutions for T_{1a} and T_{1b}, and then derive them.

$$T_{1a} = \frac{1}{g}\log\left[\frac{D_1(T_{1a})}{A_L(0)L(0)}\left(\frac{w_1 e^{(\psi_1 D_1(T_{1a}))^\phi}}{\alpha - \phi(\psi_1 D_1(T_{1a}))^\phi}\right)^{1/(1-\alpha)}\right], \tag{8.8}$$

and $\qquad T_{1b} = \frac{1}{g}\log\left[\frac{D_2(T_{1b})}{A_L(0)L(0)}\left(\frac{w_2 e^{(\psi_2 D_2(T_{1b}))^\phi}}{\alpha - \phi(\psi_2 D_2(T_{1b}))^\phi}\right)^{1/(1-\alpha)}\right], \tag{8.9}$

where $\quad D_1(T_{1a}) = \frac{1}{\psi_1}\left(\frac{\alpha\psi_1}{\phi}\frac{w_2-w_1}{w_2\psi_1-w_1\psi_2}\right)^{1/\phi}$ and $D_2(T_{1b}) = \frac{1}{\psi_2}\left(\frac{\alpha\psi_2}{\phi}\frac{w_2-w_1}{w_2\psi_1-w_1\psi_2}\right)^{1/\phi}.$

$$\tag{8.10}$$

Up to some time T_{1a}, input 1 is used exclusively, and the quantity D_1 is the unique solution to

$$A_L L = D_1\left(w_1\frac{e^{(\psi_1 D_1)^\phi}}{\alpha - \phi(\psi_1 D_1)^\phi}\right)^{1/(1-\alpha)}. \tag{8.11}$$

Given that D_1 is the solution to 8.11, we can find marginal damages from D_2 when $D_2 = 0$, using equations 8.1 and 8.3:

$$\phi(\psi_1 D_1)^{\phi-1}\psi_2(A_L L)^{1-\alpha}D_1^\alpha e^{-(\psi_1 D_1)^\phi}.$$

And then we can find the condition for starting to use input 2, which is that the marginal social costs of each input are equal:

$$w_1 + \phi\psi_1(\psi_1 D_1)^{\phi-1}(A_L L)^{1-\alpha}D_1^\alpha e^{-(\psi_1 D_1)^\phi} = w_2 + \phi\psi_2(\psi_1 D_1)^{\phi-1}(A_L L)^{1-\alpha}D_1^\alpha e^{-(\psi_1 D_1)^\phi},$$

hence $\quad \phi\psi_1^{\phi-1}D_1^\phi(A_L L/D_1)^{1-\alpha}e^{-(\psi_1 D_1)^\phi}(\psi_1 - \psi_2) = w_2 - w_1 \qquad\qquad .$

Substitute for $A_L L$ using 8.11 to yield (after some algebra)

$$(\psi_1 D_1)^\phi = \frac{\alpha}{\phi}\cdot\frac{w_2-w_1}{w_2\psi_1-w_1\psi_2}\psi_1.$$

Knowing $D_1(T_{1a})$ we can substitute in to equation 8.11 to find $A_L(T_{1a})L(T_{1a})$, and hence T_{1a}.

By a similar argument we can find the final time at which both inputs are used, which we denote T_{2a}. At T_{2a}, $D_1 = 0$, and the marginal social costs of each input are equal. Then we have, by symmetry,

$$(\psi_2 D_2)^\phi = \frac{\alpha}{\phi} \cdot \frac{w_2 - w_1}{w_2 \psi_1 - w_1 \psi_2} \psi_2,$$

and

$$A_L L = D_2 \left(w_2 \frac{e^{(\psi_2 D_2)^\phi}}{\alpha - \phi(\psi_2 D_2)^\phi} \right)^{1/(1-\alpha)}. \tag{8.12}$$

What happens between T_{1a} and T_{1b}? To answer this question, ideally we would solve for D_1 and D_2 as functions of $A_L L$ (which we now denote as A for clarity). However, we cannot solve in closed form, and must instead solve for the slopes $\partial D_1/\partial A$ and $\partial D_2/\partial A$ in this interval, using the theory of implicit functions.

We want to prove that D_2 increases while D_1 decreases across the interval when both are strictly positive. First take 8.3 and write

$$G_1(A, D_1, D_2) = w_1 + \left[\phi P^{\phi-1} \psi_1 - \alpha/R \right] Y = 0 \tag{8.13}$$

and

$$G_2(A, D_1, D_2) = w_2 + \left[\phi P^{\phi-1} \psi_2 - \alpha/R \right] Y = 0, \tag{8.14}$$

where

$$P = \psi_1 D_1 + \psi_2 D_2, \quad R = D_1 + D_2, \quad \text{and} \quad Y = A^{1-\alpha} R^\alpha e^{-P^\phi}. \tag{8.15}$$

Then the implicit function theorem tells us that

$$\begin{pmatrix} \partial D_1/\partial A \\ \partial D_2/\partial A \end{pmatrix} = - \begin{pmatrix} \partial G_1/\partial D_1 & \partial G_1/\partial D_2 \\ \partial G_2/\partial D_1 & \partial G_2/\partial D_2 \end{pmatrix}^{-1} \cdot \begin{pmatrix} \partial G_1/\partial A \\ \partial G_2/\partial A \end{pmatrix}.$$

This is relatively straightforward to solve. Firstly we have

$$\partial G_1/\partial D_1 = -\alpha w_1/R + \psi_1 w_1 \phi P^{\phi-1} + [\psi_1^2 \phi(\phi-1) P^{\phi-2} + \alpha/R^2] Y,$$

$$\partial G_1/\partial D_2 = -\alpha w_1/R + \psi_2 w_1 \phi P^{\phi-1} + [\psi_1 \psi_2 \phi(\phi-1) P^{\phi-2} + \alpha/R^2] Y,$$

$$\partial G_2/\partial D_1 = -\alpha w_2/R + \psi_1 w_2 \phi P^{\phi-1} + [\psi_1 \psi_2 \phi(\phi-1) P^{\phi-2} + \alpha/R^2] Y,$$

and

$$\partial G_2/\partial D_2 = -\alpha w_2/R + \psi_2 w_2 \phi P^{\phi-1} + [\psi_2^2 \phi(\phi-1) P^{\phi-2} + \alpha/R^2] Y,$$

and secondly

$$\partial G_1/\partial A = -(1-\alpha)w_1/A$$

and

$$\partial G_2/\partial A = -(1-\alpha)w_2/A.$$

Now use some tedious algebra, or a program such as Mathematica, to show that

$$\begin{pmatrix} \partial D_1/\partial A \\ \partial D_2/\partial A \end{pmatrix} = \frac{1-\alpha}{A} \left[\phi P^{\phi-2} \frac{\alpha}{R} (\psi_1 - \psi_2) \right]^{-1}$$

$$\left\{ (\phi-1) \left[(w_2 \psi_1 - w_1 \psi_2) + \frac{Y}{R}(\psi_1 - \psi_2) \right] - \psi_1 \frac{P}{R}(w_2 - w_1) \right\}^{-1}$$

$$\begin{pmatrix} -(w_2 - w_1)\frac{\alpha}{R^2} - (w_2 \psi_1 - w_1 \psi_2)\phi(\phi-1)P^{\phi-2}\psi_2 \\ (w_2 - w_1)\frac{\alpha}{R^2} + (w_2 \psi_1 - w_1 \psi_2)\phi(\phi-1)P^{\phi-2}\psi_1 \end{pmatrix}.$$

The signs of all the terms in this expression are unambiguous, except for the term in curly brackets. Denote $\{\cdot\}^{-1} = \Omega^{-1}$. To sign Ω, first use the first-order condition (8.3) to show that

$$R = \alpha Y \frac{\psi_1 - \psi_2}{w_2 \psi_1 - w_1 \psi_2} \qquad \text{and} \qquad P = \phi P^\phi Y \frac{\psi_1 - \psi_2}{w_2 - w_1},$$

and insert these results to yield

$$\Omega = (w_2 \psi_1 - w_1 \psi_2)/\alpha \left[(\phi-1)(1+\alpha) - \psi_1 \phi P^\phi \right].$$

We know that $(\phi-1)(1+\alpha) > 0$, but how large is $\psi_1 \phi P^\phi$ in comparison? From equation 8.10 we know that at the start of the transition

$$D_1 = \frac{1}{\psi_1} \left(\frac{\alpha}{\phi} \cdot \frac{w_2 - w_1}{w_2 \psi_1 - w_1 \psi_2} \psi_1 \right)^{1/\phi},$$

while $D_2 = 0$. So at the start of the transition,

$$\phi P^\phi = \alpha \psi_1 \cdot \frac{w_2 - w_1}{w_2 \psi_1 - w_1 \psi_2},$$

and if

$$(\phi - 1)(1 + \alpha)/\alpha > \psi_1^2 \cdot \frac{w_2 - w_1}{w_2\psi_1 - w_1\psi_2}$$

then D_1 is declining at this point while D_2 is rising. Furthermore, returning to the original expression for Ω we know that during the transition P/R declines while Y/R increases, hence the inequality continues to hold.

8.C.3. Proof of Proposition 3. First we must define an n-technology economy precisely.

DEFINITION 2. *An n-technology economy is an economy with alternative inputs $j = 1,\dots n$, with associated parameters w_j and ψ_j. Of these n inputs, input 1 is the cheapest and input m is cleanest. And input $k + 1$—where $k \in (1,\dots m - 1)$—is the input such that*

$$\frac{w_{k+1} - w_k}{w_{k+1}\psi_k - w_k\psi_{k+1}} < \frac{w_{k+m} - w_k}{w_{k+m}\psi_k - w_k\psi_{k+m}} \tag{8.16}$$

for all $m \in (2,\dots,n - k)$. Furthermore, $(\phi - 1)(1 + \alpha)/\alpha > \psi_j^2(w_{j+1} - w_j)/(w_{j+1}\psi_j - w_j\psi_{j+1})$ for $j = 1,\dots,m - 1$. Finally, the initial state $A_L(0)L(0)$ is such that only input 1 is used.

Now to the proof. During a transition from k to $k + 1$, equation 8.3 shows that

$$w_k + \phi\psi_k P^{\phi-1}Y = w_{k+1} + \phi\psi_{k+1}P^{\phi-1}Y,$$

and hence $$P^{\phi-1}Y = \frac{1}{\phi}\frac{w_{k+1} - w_k}{\psi_k - \psi_{k+1}}.$$

So during the transition, $P^{\phi-1}Y$ is constant. Now assume that the transition to $k + 2$ starts *during* the transition from k to $k + 1$. Then (analogously to equation 8.3) we have

$$w_k + \phi\psi_k P^{\phi-1}Y = w_{k+1} + \phi\psi_{k+1}P^{\phi-1}Y = w_{k+2} + \phi\psi_{k+2}P^{\phi-1}Y$$

But since $P^{\phi-1}Y$ is constant during transitions, the transition to inputs $k + 1$ and $k + 2$ should have started simultaneously, which is ruled out since equation 8.16 is never satisfied with equality. So the transition to $k + 2$ can only start after the transition to $k + 1$ is complete, which completes the proof.

CHAPTER 9

Technotopia?

The power of Sapiens over the Earth has been increasing for the last 70 000 years or more, first gradually, then rapidly, and recently explosively. This process shows no signs of abating: consider automation and robotics, which allows each human hour of work to achieve extraordinary changes to the configuration of the physical materials on Earth, both at large scales (moving mountains) and small (manufacturing microchips and editing genes). In the book so far we have analysed the Malthusian phase of rapid population growth, and the consumerist phase during which we use our increased power to increase consumption per capita. In this chapter we ask whether the consumerist phase can continue or whether it is doomed due to its internal contradictions. We then discuss the possibility of a new 'technotopian' phase during which use improved technology to give ourselves radically increased leisure time (and less work) and devote significantly greater resources to the protection and development of public goods including the global environment and other species.

9.1. Is unsustainability sustainable?

We begin by discussing whether the consumerist phase (global capitalism, if you like) is doomed due to its internal contradictions, or whether the system can carry on muddling through, leaving accidents and environmental disasters in its wake. We are unable to find convincing evidence that the system is doomed: unsustainability may be sustainable. If true this has profound implications for those campaigning for (or simply wishing for) a global economy which takes greater care of the natural world.

9.1.1. Humans (sapiens) are trashing the planet. In this section we argue for the following hypothesis.

Hypothesis 2. *We are getting richer and healthier, but trashing the planet.*

So far we have learnt that improvements in technology give humans increasing power over the global environment. Over tens of thousands of years, up to around 1700 AD, humans used this increasing power to extract a greater total quantity of food (and create better shelter) from the environment, and hence increase their population. Along the way, humans caused massive waves of extinction wherever they went, especially in areas where they had not co-evolved with the existing flora and fauna. On the other hand, some species—such as wheat—were favoured.

Since 1700 technological progress has accelerated dramatically, and with it our power over the global environment. We dig up ever increasing quantities of natural resources, and there is little or no sign that the supply is about to dwindle. We also tend to emit ever increasing quantities of pollution, but here we see a clear tendency for pollution flows to rise and then fall, due to the introduction of various regulations, typically mandating the use of alternative, cleaner technologies. We can easily understand the choice to introduce such regulations as a response to an increasing willingness to pay for environmental quality, driven in turn by increasing income. (Richer people are willing to pay more for a good environment, since the environment and consumption goods are imperfect substitutes.)

The long-run trend is that we use our increased power over the global environment to make ourselves richer and healthier. The benefits are of course unevenly spread, but nevertheless the long-run trend is clear, even for the poorest, as we see in Figure 9.1. In Figure 9.1(a) we see the rapid and remarkably steady growth in global product. The global financial crisis of 2008 shows up only as a blip, as the drastic effects of the crisis in Europe and North America are balanced by the quick recovery and subsequent rapid growth in the rest of the world, especially Asia. In 9.1(b) we see how the absolute number of people in extreme poverty has fallen steeply since the early 1990s. Note however that there has been a huge shift—over the period shown in Figure 9.1—in the distribution of those in extreme poverty (under 1.90 USD/day). Up to 1990 most were in East Asia, followed by South Asia, and finally sub-Saharan Africa. Over the following 25 years, extreme poverty in East Asia (dominated by China) has almost disappeared, and in South Asia (dominated by India) it has fallen substantially. But in sub-Saharan Africa it has increased significantly and a large majority of extremely poor people now live there.
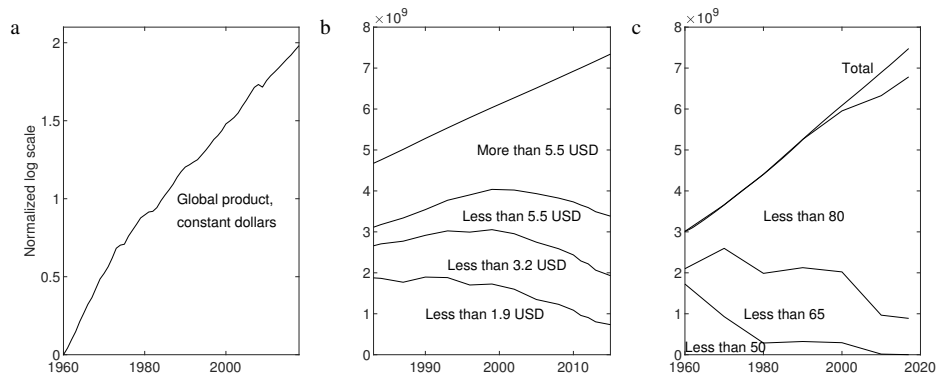
FIGURE 9.1. Trends in global production, population and poverty: (a) Global product, constant dollars, log scale; (b) The distribution of global population by income category (2011 PPP dollars per day); (c) The distribution of global population by life expectancy at birth. World Bank data.

Despite our increasing WTP for environmental quality, we continue to trash the planet. We emit vast volumes of carbon dioxide. These emissions are already causing significant changes in the climate which are certain to continue and strengthen, as well as acidification of the oceans. We are causing loss of species on a massive scale; according to Thomas et al. (2004), 18–35 percent of species will be committed to extinction by 2050. The point stands, even if **?** claim that the Thomas et al. estimate is high-end possibility which may be overestimate by a factor of around 6.

Figure 9.1(a) also shows that there is no sign that global growth is about to slow down dramatically. See also Figure 1.2. This is hardly surprising when we consider that growth is driven by technological progress, and technological progress is continuing at a rapid pace in the leading economies, at the same time as existing technologies are diffusing rapidly to economies behind the technology frontier. So if growth is to stop it will either be because of deliberate decisions we have made, for instance to work less and take more leisure, or it will be because of some form of collapse of the social or ecological systems on which the economy depends. How likely is such a collapse?

9.1.2. Ecosystem services and nature. Analysis of the likelihood of social upheaval or collapse is beyond the scope of this book. We instead focus on ecological collapse. Consider the following alternative hypotheses.

Hypothesis 3.
- *We rely on services provided by the planet, and by trashing the planet we are destroying the planet's ability to provide these services in the future.*
- *Therefore, if we carry on trashing the planet the loss of these services will lead to us getting poorer and sicker.*

Hypothesis 4.
- *We are adaptable and ingenious.*
- *Therefore we can carry on both trashing the planet and getting richer and healthier, indefinitely.*

Consider for instance climate change. What do mainstream climate models predict about the effects of climate change on the global economy? If we start with Nordhaus (nobel prize, together with Paul Romer, 2018) we see that the predictions are that damages will be modest when measured as a proportion of global product. For instance, Nordhaus (2018) estimates that damages from 3°C warming would be just 2.1 percent of global income, and even an extreme 6°C warming would cost a relatively manageable 8.5 percent of income. As long as there are only level effects (and no effect on the long-run growth rate) this is pretty modest. For instance, Brexit is likely to cost the UK far more, because it will damage the growth rate of the economy (by isolating the UK from Europe and thus reducing competition and slowing the spread of ideas). Of course, Nordhaus does not then draw the conclusion that we should not act on climate change; on the contrary, since the costs of cutting emissions are far lower than the benefits in terms of avoided damage, we should definitely act far more strongly than we are doing so far. Even those who are far less sanguine about damages, such as the late Marty Weitzman (2011), argue that the point is that although the risk of very bad outcomes is small, it is not small enough, and hence these outcomes should weigh very heavily in the policy calculus.

Climate damages strike unequally, and are likely to hit the poor hardest. There are therefore also distributional aspects to the problem. However, given how little is being done to redress global inequality, this also seems to be something that we are prepared to put up with.

Now assume that you care about the planet, and would like to help redress the balance between the pursuit of material wealth and care of the planet. What to do? Consider the following two strategies.

(1) Find evidence for hypothesis 3, or try to convince others of its veracity.

(2) Persuade others to care too: either more about the planet, or less about material wealth!

There are of course other strategies. For instance:

(3) Demonstrate that the system (e.g. 'global capitalism') is going to crash anyway, irrespective of the state of the planet. And argue that since it's going to crash we might as well slow it down gently and save the planet at the same time.

The first strategy is fine as long as hypothesis 3 holds. But what if it doesn't? If you've pinned all your arguments onto it, and it turns out not to be true, you're in trouble. Maybe following the second strategy would be a better idea? In this chapter I will suggest that hypothesis 4 is probably true, and therefore that strategy 2 is much preferable to strategies 1 and 3, which are both very risky.

If hypothesis 2 is true, then strategy 2 is almost certainly preferable to strategies 1 and 3, which are both very risky. Think about this when studying the literature. What are the approaches of the following authors?

- Jackson (2009)
- Rockström et al. (2009).
- Meadows et al. (1972).

9.2. Lessons from historical adaptation

The regulated market economy has shown a remarkable ability to adapt and react to crises when they arise, including environmental crises. When major problems are discovered with a direct effect on the welfare of the rich, decisive action is taken. CFCs, DDT, etc. Climate?

On the other hand, we know that environmental crises may often have far-reaching consequences for nature, and sometimes for human welfare. And when the consequences are *only* for nature, not a lot tends to get done. Consider for instance the Baltic Sea, or bird populations in Europe.

Finally, there are examples of civilizations that have collapsed, apparently due to environmental collapse. E.g. Easter Island. What lessons are there here? E.g. Brander and Taylor (1998).

9.3. Financial and other crises

Growth—building on technological progress—in the industrial era has been remarkably resilient, and there is no evidence that this is likely to change.

We know that financial crises—especially large-scale ones across many countries—typically have severe and long-lasting effects. However, such a crisis does not signal the death-throes of capitalism, the collapse of the system under the weight of contradictions. We know why the recent global financial crisis occurred, and we know why recovery from it is so slow. The reason is the

lack of confidence in the future which is widespread among agents, a lack of confidence which is rational for each individual in the knowledge that everyone else lacks confidence. It is a gigantic coordination problem, the solution to which is either some massive shock (such as WW2 in 1939) or gradual, inch-by-inch progress.

9.4. Uncertainty and future crises

So the global economy will very likely be able to keep going as it has up to now, for decades or even centuries to come. Growing, triggering environmental problems and even catastrophes, and then solving them. All the while, the space for the non-human or 'natural' world is likely to be circumscribed ever-more by our thirst for consumption, consumption of everything from food to wilderness experiences.

Unsustainability—defined as the trashing of the planet by global capitalism—may be sustainable for centuries to come. (Assuming that we accept this description of what is going on today.)

To turn this around we need policies driven by politicians elected by voters who value so-called global public goods more highly relative to material consumption, compared to today.

Alternatively, we need a technological breakthrough which drastically reduces or even eliminates the cost advantage of fossil fuels over clean energy sources.

An advantage of the second alternative is that it would bring on board countries which care little about the planet.

But even this wouldn't ease the pressure on habitats, both on land and in water.

A 'chilling out' phase?

10.1. Consumerism

Firms want to sell more stuff, employ more people, etc. The market economy has its own dynamic. Or something.

10.2. 'Green' consumerism

Green consumerism is a tricky business. The key problem is rebound. If you don't consume one thing, but your income is unchanged, you will consume something else instead, or invest in capital which may be just as bad. It is an impossible task for individual consumers to weigh up the environmental effects of their actions. We need consumers to elect politicians who enact laws which (a) lead to external effects being internalized in the prices of goods (in borderline cases), and (b) lead to highly damaging or unnecessary practices being banned (in black-and-white cases). A recent example of the latter is the ban on incandescent light bulbs in both the US and the EU.

10.3. Conspicuous consumption, labour, and leisure

10.3.1. Background chat, very preliminary. Are we prioritizing consumption of goods too highly, and preservation of nature too low? But we elect governments in a democratic process, and they choose policies which determine these priorities.

But what about the Easterlin paradox, the claim that increasing average income in a country is only weakly correlated with increasing happiness and well-being? Does this not suggest that the all-out effort to work harder and produce more is a rat-race with no winner?

This idea links with the ideas of Thorstein Veblen (see for instance Veblen, 1899), who coined the term 'conspicuous consumption' to capture the idea that we consume in order to be seen to consume, and by being seen to consume we raise our status which is gives us utility. However, in order to consume we must earn income, and in order to earn income we must work, and in order to work we must sacrifice leisure, and the sacrifice of leisure reduces our utility.

Putting Veblen's ideas into a modern economic context, if we consume to gain status, then there is a consumption externality: one person's higher status is her neighbours' lower status, driving down their utility. Therefore each individual's choice to work long hours to earn more money and consume more has a negative external effect on that individual's neighbours. And, in the global village, that might mean everyone else in the global population.

Consider climate negotiations. Is it possible that a major stumbling block in these negotiations is not the desire of China and India to get richer, but rather the desire of China and India to *catch up* with the OECD countries in terms of wealth and income per capita? And, by the same token, the desire of the richest countries that China and India should *not* catch up with them in terms of wealth and income per capita, implying that China and India (by virtue of their higher populations) would wield much more economic power than USA and the EU?

Why do politicians constantly exhort their citizens to adopt growth-friendly policies so that they do not lose out in the 'global race'? There is no global race: countries which adopt new technology less aggressively—and countries whose populations work less and take more leisure time—have lower GDP per capita than other countries, but they have neither higher unemployment nor necessarily lower welfare. Why then the political obsession with growth and productivity? Could it be that citizens compare their consumption rates across borders, and furthermore that politicians gain utility from observing that 'their' economies are larger and more powerful than those of their neighbours?

If the above is true (or partly true) then it is highly likely that we work *too hard* and take too little leisure time for our own good. More precisely, if everyone worked less and took more leisure time, everyone would be better off! And this applies irrespective of spin-off benefits for the environment and nature. In the next section we build and solve a simple model to demonstrate the mechanism.

10.3.2. A very simple model. Assume a population of identical households indexed by i where the utility of a given household is described by the following function:

$$u_i = c_i^{\alpha_1} r_i^{1-\alpha_1-\alpha_2} (c_i/\bar{c})^{\alpha_2}.$$

Here c_i is consumption, r_i is leisure, and $\bar{c}$ is average consumption across all households. Furthermore, α_1 and α_2 are parameters the sum of which is less than 1. For convenience define

$$\alpha = \alpha_1 + \alpha_2.$$

Production is a linear function of labour, but consumption is reduced by an income tax at a flat rate τ, and boosted by a lump-sum transfer of public goods which is equal to average production $\bar{l}$ multiplied by the tax rate. Thus we have

$$c_i = l_i(1-\tau) + \bar{l}\tau.$$

Leisure r (for *recreation*) is equal to total time R minus labour time, i.e.

$$r_i = R - l_i.$$

Now take the tax as exogenous and work out how much each household chooses to work. To do so, substitute into the utility function to yield

$$u = \left[l_i(1-\tau) + \bar{l}\tau \right]^{\alpha} (R - l_i)^{1-\alpha} / \bar{c}^{\alpha_2}.$$

Take the first-order condition in l_i and solve to show that

$$l_i = \alpha R - (1-\alpha)\bar{l}\tau/(1-\tau).$$

So when income tax is zero $l_i = \alpha R$: as labour income dominates the utility function (α high) households devote more of their time to labour and less to leisure.

Now assume a symmetric equilibrium such that average labour $\bar{l}$ is equal to the labour supplied by household i, l_i. Inserting this into the above result we have

$$l_i = \bar{l} = \frac{\alpha R}{1 + (1-\alpha)\tau/(1-\tau)}.$$

Now the question for a regulator is, what level of tax τ maximizes utility for households? Economic theory tells us that if markets are perfect then the optimal tax should be zero, hence $l_i = \alpha R$. But if there is a consumption externality—i.e. if $\alpha_2 > 0$—then this no longer holds.

To solve the problem, we insert the expression for l_i as a function of τ into the utility function —noting that in symmetric equilibrium $c_i = \bar{c}$ and (as already stated) $l_i = \bar{l}$—to obtain

$$u = \left[\frac{\alpha R}{1 + (1-\alpha)\tau/(1-\tau)} \right]^{\alpha_1} \left[R - \frac{\alpha R}{1 + (1-\alpha)\tau/(1-\tau)} \right]^{1-\alpha}.$$

Simplify to obtain

$$u = R^{1-\alpha_2} \alpha^{\alpha_1} \omega^{-(1-\alpha_2)} (\omega - \alpha)^{1-\alpha},$$

where $\qquad \omega = 1 + (1-\alpha)\tau/(1-\tau).$

Take the first-order condition in ω to solve for the optimal ω, and then use the definition of ω to solve for the optimal tax:

$$\tau = \frac{\alpha_2}{\alpha_1 + \alpha_2}.$$

So, the stronger the weight of 'conspicuous consumption' in utility, the more labour income should be taxed.

How big is the effect? Assuming conspicuous consumption has equal weight to consumption in utility then labour income should be taxed at 50 percent. The effect of the tax is to reduce labour supply by a factor ω. And if leisure has 50 percent weight in utility (implying that $\alpha = 0.5$ so in laissez-faire the individuals would work 8 hours and have 8 hours of leisure time, assuming that 8 hours are needed for sleep) then $\omega = 1.5$, so labour supply is reduced by one third.

10.3.3. Conclusions on conspicuous consumption. The above model should not be taken too seriously: it is based on assumptions which are plucked more-or-less out of thin air rather than backed by careful argument and empirical evidence. Nevertheless, it demonstrates that there may exist sound economic arguments for governments to discourage labour and encourage leisure even in the absence of environmental damage from production. However, if *international* consumption externalities are important then it is only rational for national governments to impose such policies if their neighbours do the same. Perhaps this is why European economies have (collectively) been able to hold down or even reduce working hours over recent decades, whereas the US has lurched dramatically in the opposite direction.[1]

---

[1]For a reference to give an introduction to the field, see Aronsson and Johansson-Stenman (2008). On international comparison of working hours see Prescott (2004).

Appendices

Mathematical appendix

A.1. Growth rates

What is meant by a growth rate? What is a constant rate of growth? And how is growth best represented graphically? Assume we are interested in GDP, denoted Y. A constant rate of growth in Y implies that Y grows exponentially, such that (for instance) the time it takes for global product to double is constant. Mathematically we have $Y = Y_0 e^{gt}$, where Y_0 is global product at time zero, g is the growth rate, and t indicates time.

Consider for instance an economy growing by 3 percent per year, and with initial GDP of 1×10^9 USD/year. We then have $y = y_0 e^{gt}$ where t is time measured in years, $g = 0.03$, and y_0 is initial GDP. If we plot y against time we get the familiar exponential form; Figure A.1(a). However, this is a poor way to organize data as it is hard for the eye to interpret: given a curve of increasing slope it is not generally possible to determine by eye whether it represents a constant, increasing, or decreasing growth rate. For instance, compare the curve to Figure A.1(b), which plots the function $y = y_0(1 + t^{2.5}/5300)$. This is not exponential growth, but this is far from obvious from the figure.
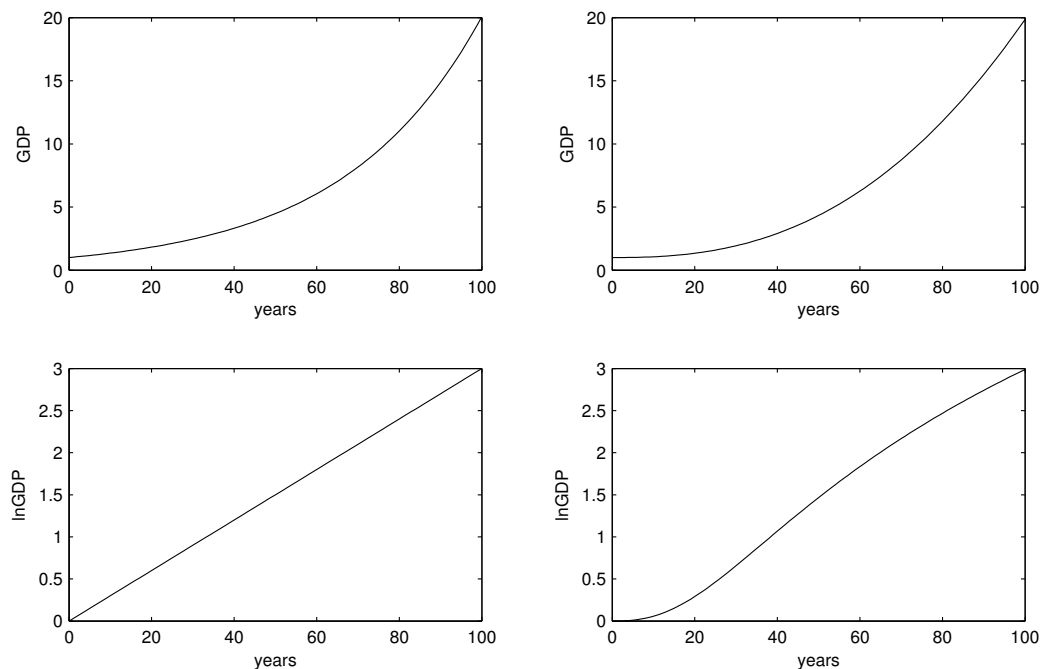

FIGURE A.1. Two different growth patterns, illustrated using a linear scale ((a) and (b)) and a logarithmic scale ((c) and (d)).

In order to see the true growth trend, take the (natural) logarithm of the equations. In the case of exponential growth we have

$$\ln(y/y_0) = gt. \tag{A.1}$$

So we have a straight line through the origin, the slope of which is g; Figure A.1(c). The second equation results in the curve shown in Figure A.1(d), where we can see that the growth rate is initially zero, subsequently increases, and then declines.

Another advantage to plotting the logarithm of a growing variable is that the relative sizes of fluctuations are also shown in proportion. Consider for instance the data showing USA's GDP from 1870 to 2008, Figure A.2. Consider the size of the fluctuations in GDP before and after the great depression and WW2. Have the fluctuations increased over time, decreased, or remained the same? From A.2(a) it is very hard to tell; a naive interpretation of the curve would be that the fluctuations

have increased. However, when we plot the logarithm (A.2(b)) we see that the fluctuations in GDP in the second period were smaller than in the first, as a proportion of GDP at the time.


FIGURE A.2. Two different representations of U.S. growth. The logarithmic scale helps us to see both that the trend growth rate is very constant, and that the fluctuations are smaller in the latter period (after WW2).

Finally, note that another way to illustrate growth data is through plotting the growth rate $\dot{Y}/Y$ against time. Given a constant growth rate g we have $\dot{Y}/Y = g$ hence this yields a horizontal line at height g.

A.2. Continuous and discrete time

Dynamic economic models may be set up either in continuous or discrete time. In continuous time we can think of time flowing, and also physical goods (such as natural resources or machines) should flow in the sense that their quantities change gradually rather than suddenly jumping from one level to another. In discrete time we divide time up into discrete periods, and the state of the system jumps between one period and the next. Models in continuous time are often more mathematically elegant, whereas models in discrete time may be more practical in some circumstances. Empirical (economic) data is typically collected at discrete intervals, so we measure for instance total production and natural resource use in a quarter (three months), rather than the flow of production and resource use at each instant. In some economic models the assumption of discrete time is crucial because it is assumed, for instance, that firms all invest simultaneously and periodically.

Mathematically the appearance of models in discrete and continuous time is typically very different, but the results are typically essentially the same (as long as it is practically possible to model the system using both). Assume a process of exponential growth in continuous time. Following the examples of the previous section we assume

$$\dot{Y}/Y = g.$$

Now consider measuring Y at discrete intervals of one year, indexed by t, starting at $t = 0$ when we assume that $Y = 1$. What is Y at $t = 1$?

$$\frac{\mathrm{d}Y}{\mathrm{d}t} = gY,$$

$$\int_{Y_0}^{Y_t} (1/Y)\mathrm{d}Y = \int_0^1 g\mathrm{d}t,$$

$$\log(Y_t/Y_0) = gt,$$

$$Y_t/Y_0 = e^{gt}.$$

So at $t = 1$ we have $Y = e^g$, at $t = 2$ we have $Y = e^{2g}$, and so on. Furthermore, we can write

$$\frac{Y_{t+1} - Y_t}{Y_t} = e^g$$

and
$$\frac{Y_{t+1}}{Y_t} = 1 + e^g.$$

So if we take measurements in discrete time we find a *growth factor* $1 + e^g$, or $1 + \theta$, whereas in continuous time we have a *growth rate g*. Note that when g is small then $\theta \approx g$, whereas when g becomes large θ becomes significantly greater than g because compound growth during each period becomes significant. Finally, if we let the period length approach zero then g also approaches zero, and θ approaches g.

A.3. A continuum of firms

In macroeconomic modelling it is very common to assume a *continuum* of firms (and indeed households). Why do we do this, and what does it mean?

To understand this, assume instead that we have just one firm. This is nice and simple in one sense, because aggregate production and aggregate demand for inputs is the same as the production and demand of the single firm. However, there is a big problem in that this single firm must then have market power, both on the market for inputs (e.g. the labour market) and on the market for the final good. Of course, we know that market power does exist in the real economy, but to start off with a model of only one firm is unsatisfactory: typically we want to start our models assuming the simplest possible case, i.e. competitive markets, and then introduce market power later when it is relevant; furthermore, even when market power is relevant we are unlikely to have pure monopsony or monopoly (only one buyer or seller).

In order to generate competitive markets in our models we need to add more firms. But how many? In a competitive market the actions of any one firm have no effect on prices, firms are price takers. But this implies that each firm must be infinitesimally small, its production must be negligible compare to the total. Thus we need an infinite number of very small firms. So firms are no longer countable, but we can measure their mass; and in some circumstances we may want to claim that the mass of firms has grown (e.g. doubled), even though the *number* of firms is not a meaningful concept. However, in most circumstances we simply assume that there is a *unit mass* of firms, or a continuum of firms measure 1. This has the advantage that, in symmetric equilibrium, production per firm is the same as aggregate production from all the firms (and the same holds for input demand).

To see this mathematically, index firms by i and assume that $i \in (0, 1)$. Production from firm i is denoted Y_i, and total production from all firms is Y. All firms produce the same good. Then total production is simply the integral across all the firms,

$$Y = \int_0^1 y_i \mathrm{d}i.$$

Now, if the equilibrium is symmetric so all firms produce the same amount then y_i is simply constant (it does not vary as i runs from 0 to 1), and

$$Y = y,$$

where y is production by any one *representative firm*.

Note that in more complex models the goods produced by the firms may differ, and the mathematics is slightly more complicated. These cases are analysed in detail in the main text when they arise.

Bibliography

Acemoglu, D., 2009. Introduction to modern economic growth. Princeton, New Jersey.

Acemoglu, D., Aghion, P., Bursztyn, L., Hemous, D., 2012. The environment and directed technical change. American Economic Review 102, 131–166.

Acemoglu, D., Guerrieri, V., Jun. 2008. Capital deepening and nonbalanced economic growth. Journal of Political Economy 116 (3), 467–498.

Aghion, P., Howitt, P., 1992. A model of growth through creative destruction. Econometrica 60, 323–51.

Andreoni, J., Levinson, A., 2001. The simple analytics of the environmental Kuznets curve. Journal of Public Economics 80, 269–286.

Aronsson, T., Johansson-Stenman, O., Jun. 2008. When the Joneses' consumption hurts: Optimal public good provision and nonlinear income taxation. Journal of Public Economic 92 (5-6), 986–997.

Arrow, K. J., Chang, S., Mar. 1982. Optimal pricing, use, and exploration of uncertain natural resource stocks. Journal of Environmental Economics and Management 9 (1), 1–10.

Arthur, W., 1989. Competing technologies, increasing returns, and lock-in by historical events. Economic Journal 99, 116–131.

Asheim, G. B., Buchholz, W., Withagen, C., Jun. 2003. The Hartwick rule: Myths and facts. Environmental and Resource Economics 25 (2), 129–150.

Avenell, S., 2012. From fearsome pollution to Fukushima: Environmental activism and the nuclear blind spot in contemporary Japan. Environmental History 17 (2), 244–276.

Barreca, A. I., Neidell, M., Sanders, N. J., June 2017. Long-run pollution exposure and adult mortality: Evidence from the Acid Rain Program. Working Paper 23524, National Bureau of Economic Research.

Barro, R. J., Becker, G. S., 1989. Fertility choice in a model of economic growth. Econometrica 57 (2), 481–501.
   URL http://www.jstor.org/stable/1912563

Baumol, W. J., Oates, W. E., 1975. The theory of environmental policy. Prentice-Hall, New Jersey.

Binswanger, M., 2001. Technological progress and sustainable development: What about the rebound effect? Ecological Economics 36 (1), 119–132.

Black and Veatch, 2016. Air quality control solutions for India: Innovative approaches and seamless execution. Tech. rep.

Boden, T. A., Marland, G., Andres, R. J., 2012. Global, regional, and national fossil-fuel CO_2 emissions. Tech. rep., CDIAC, Oak Ridge National Laboratory, U.S. Department of Energy, Oak Ridge, Tenn., U.S.A.

Boulding, K. E., 1966. The economics of the coming spaceship earth. In: Jarrett, H. (Ed.), Environmental Quality in a Growing Economy. Resources for the Future/Johns Hopkins University Press, pp. 3–14.

BP, 2012. Statistical review of world energy. Annually published data, British petroleum.

Brander, J. A., Taylor, M. S., Mar. 1998. The simple economics of Easter Island: A Ricardo–Malthus model of renewable resource use. American Economic Review 88 (1), 119–138.

Bresnahan, T., Trajtenberg, M., 1995. General purpose technologies: Engines of growth? Journal of Econometrics, Annals of Econometrics 65, 83–108.

Brock, W. A., Taylor, M. S., 2010. The green Solow model. Journal of Economic Growth 15 (2), 127–153.

Bulte, E., Horan, R. D., Shogren, J. F., 2006. Megafauna extinction: A paleoeconomic theory of human overkill in the pleistocene. Journal of Economic Behavior & Organization 59 (3), 297–323.

Burney, D., Flannery, T., 2005. Fifty millennia of catastrophic extinctions after human contact 20, 395–401.

Cass, D., 1965. Optimum growth in an aggregative model of capital accumulation. Review of Economic Studies 32 (3), 233–240.

Daly, H. E., 1997. Georgescu-Roegen versus Solow/Stiglitz. Ecological Economics 22, 261–266.

Dasgupta, P., 1993. Natural resources in an age of substitutability. In: Kneese, A. V., Sweeney, J. L. (Eds.), Handbook of natural resource and energy economics. Vol. 3. Elsevier, Ch. 23.

Dasgupta, P., Heal, G., 1974. The optimal depletion of exhaustible resources. The Review of Economic Studies 41, 3–28.

Dasgupta, P., Heal, G., 1979. Economic Theory and Exhaustible Resources. Cambridge University Press, Cambridge.

David, P. A., May 1990. The dynamo and the computer: An historical perspective on the modern productivity paradox. American Economic Review 80 (2), 355–61.

Ellerman, A., Joskow, P., Schmalensee, R., Bailey, E., Montero, J., 2000. Markets for Clean Air: The U.S. Acid Rain Program. Cambridge University Press.

Ellerman, A. D., Montero, J.-P., 1998. The declining trend in sulfur dioxide emissions: Implications for allowance prices. Journal of Environmental Economics and Management 36 (1), 26–45.

Engel, E., 1857. Die productions- und consumptionsverhaeltnisse des koenigsreichs sachsen. Zeitschrift des Statistischen Bureaus des Koniglich Sachsischen Ministeriums des Inneren 8–9.

Figueroa, E., Pastén, R., 2015. Beyond additive preferences: Economic behavior and the income pollution path. Resource and Energy Economics 41, 91–102.

Finus, M., Tjotta, S., Sep. 2003. The Oslo Protocol on sulfur reduction: the great leap forward? Journal of Public Economics 87 (9-10), 2031–2048.

Fouquet, R., 2011. Divergences in long-run trends in the prices of energy and energy services. Review of Environmental Economics and Policy 5 (2), 196–218.

Fouquet, R., Pearson, P. J. G., 2006. Seven centuries of energy services: The price and use of light in the United Kingdom (1300–2000). Energy Journal 27, 139–177.

Galor, O., 2005. From stagnation to growth: Unified growth theory. In: Aghion, P., Durlauf, S. (Eds.), Handbook of Economic Growth. Elsevier.

Goeller, H. E., Weinberg, A. M., 1976. The age of substitutability. Science 191 (4228), 683–689.

Golosov, M., Hassler, J., Krusell, P., Tsyvinski, A., 2014. Optimal taxes on fossil fuel in general equilibrium. Econometrica 82 (1), 41–88.

Grossman, G. M., Krueger, A. B., May 1995. Economic growth and the environment. Quarterly Journal of Economics 110 (2), 353–377.

Groth, C., 2007. A new-growth perspective on non-renewable resources. In: Bretschger, L., Smulders, S. (Eds.), Sustainable Resource Use and Economic Dynamics. Springer, pp. 127–163.

Hardin, G., 1968. The tragedy of the commons. Science 162 (3859), 1243–1248.
    URL http://science.sciencemag.org/content/162/3859/1243

Hart, R., 2002. Dynamic modelling of environment and economy. Ph.D. thesis, SLU, Uppsala.

Hart, R., 2009. The natural-resource see-saw: Resource extraction and consumption with directed technological change. Presented at EAERE 2009.

Hart, R., 2013. Directed technological change and factor shares. Economics Letters 119, 77–80.

Hart, R., 2016. Non-renewable resources in the long run. Journal of Economic Dynamics and Control 71, 1–20.

Hart, R., 2018a. Rebound, directed technological change, and aggregate demand for energy. Journal of Environmental Economics and Management 89, 218–234.

Hart, R., 2018b. To everything there is a season: Carbon pricing, research subsidies, and the transition to fossil-free energy. Journal of the Association of Environmental and Resource Economists. Forthcoming.

Hart, R., Spiro, D., 2011. The elephant in Hotelling's room. Energy Policy 39 (12), 7834–7838.

Hartwick, J. M., 1977. Intergenerational equity and the investing of rents from exhaustible resources. American Economic Review 67, 972–974.

Hills, R. L., 1993. Power from Steam: A history of the stationary steam engine. Cambridge University Press.

Houthakker, H. S., 1957. An international comparison of household expenditure patterns, commemorating the centenary of Engel's law. Econometrica 25 (4), 532–551.

Jackson, T., 2009. Prosperity Without Growth: Economics for a Finite Planet. Earthscan.

Jacobsen, J. B., Hanley, N., Jun. 2009. Are there income effects on global willingness to pay for biodiversity conservation? Environmental & Resource Economics 43 (2), 137–160.

Jevons, W. S., 1865. The coal question; an inquiry concerning the progress of the Nation, and the probable exhaustion of our coal-mines. Macmillan.

Jones, C., 2005. Growth and ideas. In: Aghion, P., Durlauf, S. (Eds.), Handbook of Economic Growth. Elsevier.

Kelly, T. D., Matos, G. R., 2012. Historical statistics for mineral and material commodities in the United States. Tech. rep., U.S. Geological Survey, Data Series 140.
URL http://pubs.usgs.gov/ds/2005/140/

Knittel, C. R., Dec. 2011. Automobiles on steroids: Product attribute trade-offs and technological progress in the automobile sector. American Economic Review 101 (7), 3368–3399.

Koopmans, T. C., 1963. On the concept of optimal economic growth. Cowles foundation discussion paper.

Lin, C.-Y. C., Wagner, G., Jul. 2007. Steady-state growth in a Hotelling model of resource extraction. Journal of Environmental Economics and Management 54 (1), 68–83.

Lindholt, L., Sep. 2005. Beyond kyoto: backstop technologies and endogenous prices on co2 permits and fossil fuels. Applied Economics 37 (17), 2019–2036.

Lucas, R. E., 1976. Econometric policy evaluation: A critique. Carnegie-Rochester Conference Series on Public Policy 1 (0), 19–46.

Maddison, A., 2010. Historical statistics of the world economy: 1–2008 ad. Tech. rep., Groningen growth and development centre.

Malthus, T., 1798. An essay on the principle of population. Oxford World Classics, reprinted 2008. Oxford University Press.

Manne, A., Mendelsohn, R., Richels, R., Jan. 1995. A model for evaluating regional and global effects of GHG reduction policies. Energy Policy 23 (1), 17–34.

Markusson, N., 2012. The politics of FGD deployment in the UK (1980s–2009). Tech. rep., UK ERC.

Matsuyama, K., 2002. The rise of mass consumption societies. Journal of Political Economy 110, 1035–1070.

Maxwell, M., Elder, H., Morasky, T., 1978. Sulfur oxides control technology in Japan. Interagency task force report.

Mayer, H., Flachmann, C., 2011. Extended input-output model for energy and greenhouse gases. Tech. rep., Federal Statistical Office Germany.

Meadows, D. H., Meadows, D. L., Randers, J., Behrens III, W. W., 1972. The Limits to Growth. A Report for The Club of Rome's Project on the Predicament of Mankind. Universe Books.

Muller, N. Z., Mendelsohn, R., 2007. Measuring the damages of air pollution in the United States. Journal of Environmental Economics and Management 54 (1), 1–14.

Murty, S., Russell, R. R., Levkoff, S. B., 2012. On modeling pollution-generating technologies. Journal of Environmental Economics and Management 64 (1), 117–135.

Nordhaus, W., 2008. A Question of Balance: Weighing the Options on Global Warming Policies. Yale University Press.

Nordhaus, W., 2018. Projections and uncertainties about climate change in an era of minimal climate policies. American Economic Journal: Economic Policy 10 (3), 333–60.

Nordhaus, W. D., 1973a. The allocation of energy resources. Brookings Papers on Economic Activity 4 (1973-3), 529–576.

Nordhaus, W. D., 1973b. Some skeptical thoughts on theory of induced innovation. Quarterly Journal of Economics 87 (2), 208–219.

Orwell, G., 1958. The Road to Wigan Pier. Harcourt, Brace, and company, New York.

Panayotou, T., 1993. Empirical tests and policy analysis of environmental degradation at different stages of economic development. ILO Working Papers 292778, International Labour Organization.

Parfit, D., 1984. Reasons and Persons. OUP Oxford.

Pigou, A., 1920. The Economics of of Welfare, 4th Edition. Macmillian.

Pindyck, R. S., 2004. Volatility and commodity price dynamics. Journal of Futures Markets 24 (11), 1029–1047.

Poore, J., Nemecek, T., 2018. Reducing food's environmental impacts through producers and consumers. Science 360 (6392), 987–992.

Popp, D., 2002. Induced innovation and energy prices. American Economic Review 92 (1), 160–180.

Prescott, E. C., July 2004. Why do Americans work so much more than Europeans? Federal Reserve Bank of Minneapolis Quarterly Review, 2–13.

Ramsey, F. P., 1928. A mathematical theory of saving. Economic Journal 38, 543–559.

Rawls, J., 1971. A theory of justice. Harvard University Press, Cambridge.

Rockström, J., Steffen, W., Noone, K., Persson, A., Chapin, F. S., Lambin, E. F., Lenton, T. M., Scheffer, M., Folke, C., Schellnhuber, H. J., Nykvist, B., de Wit, C. A., Hughes, T., van der Leeuw, S., Rodhe, H., Sorlin, S., Snyder, P. K., Costanza, R., Svedin, U., Falkenmark, M.,

Karlberg, L., Corell, R. W., Fabry, V. J., Hansen, J., Walker, B., Liverman, D., Richardson, K., Crutzen, P., Foley, J. A., Sep. 2009. A safe operating space for humanity. Nature 461 (7263), 472–475.

Rogner, H. H., 1997. An assessment of world hydrocarbon resources. Annual Review of Energy and the Environment 22, 217–262.

Romer, P. M., 1990. Endogenous technological change. Journal of Political Economy 98 (5), S71–102.

Romer, P. M., 1994. The origins of endogenous growth. Journal of Economic Perspectives 8 (1), 3–22.

Saunders, H. D., 1992. The Khazzoom–Brookes postulate and neoclassical growth. Energy Journal 13 (4), pp. 131–148.

Saunders, H. D., Jun. 2000. A view from the macro side: Rebound, backfire, and Khazzoom-Brookes. Energy Policy 28 (6-7), 439–449.

Schäfer, A., 2006. Long-term trends in global passenger mobility. The Bridge, 24–32.

Schumpeter, J., 1942. Capitalism, socialism and democracy. Harper.

Selden, T. M., Forrest, A. S., Lockhart, J. E., 1999. Analyzing the reductions in U.S. air pollution emissions: 1970 to 1990. Land Economics 75 (1), 1–21.

Smith, A., 1776. An inquiry into the nature and causes of the wealth of nations. William Strahan.

Smith, F., Jeffrey, G., 1975. Airborne transport of sulphur dioxide from the U.K. Atmospheric Environment 9, 643–659.

Smulders, S., Bretschger, L., Egli, H., 2011. Economic growth and the diffusion of clean technologies: Explaining environmental Kuznets curves. Environmental & Resource Economics 49 (1), 79–99.

Solow, R. M., 1973. Is the end of the world at hand? Challenge, 39–46.

Solow, R. M., 1974. Intergenerational equity and exhaustible resources. Review of Economic Studies, Symposium on the economics of exhaustible resources 41, 29–45.

Sorrell, S., 2007. The rebound effect: an assessment of the evidence for economy-wide energy savings from improved energy efficiency. UKERC.

Stern, D. I., 2004. The rise and fall of the environmental Kuznets curve. World Development 32 (8), 1419 – 1439.

Stern, D. I., 2005. Global sulfur emissions from 1850 to 2000. Chemosphere 58, 163–175.

Stern, N., 2006. Stern Review on the Economics of Climate Change. UK Treasury.

Stern, N., 2008. The economics of climate change. American Economic Review 98 (2), 1–37.

Stiglitz, J. E., 1974. Growth with exhaustible natural resources: Efficient and optimal growth paths. Review of Economic Studies, Symposium on the Economics of Exhaustible Resources 41, 123–37.

Stokey, N. L., 1998. Are there limits to growth? International Economic Review 39, 1–31.

Stollery, K. R., 1983. Mineral depletion with cost as the extraction limit – a model applied to the behavior of prices in the nickel industry. Journal of Environmental Economics and Management 10 (2), 151–165.

Sunstein, C. R., 2007. Of Montreal and Kyoto: A tale of two protocols. Harvard Environmental Law Review 31, 1–65.

Tahvonen, O., Salo, S., Aug. 2001. Economic growth and transitions between renewable and nonrenewable energy resources. European Economic Review 45 (8), 1379–1398.

Taylor, M. R., Rubin, E. S., Hounshell, D. A., 2005. Regulation as the mother of innovation: The case of SO_2 control. Law and Policy 27, 348–378.

Teece, D. J., Sunding, D., Mosakowski, E., 1993. Natural resource cartels. In: Kneese, A., Sweeney, J. (Eds.), Handbook of Natural Resource and Energy Economics. Vol. III. Elsevier, Ch. 24, pp. 1131–1165.

Thomas, C. D., Cameron, A., Green, R. E., Bakkenes, M., Beaumont, L. J., Collingham, Y. C., Erasmus, B. F. N., de Siqueira, M. F., Grainger, A., Hannah, L., Hughes, L., Huntley, B., van Jaarsveld, A. S., Midgley, G. F., Miles, L., Ortega-Huerta, M. A., Townsend Peterson, A., Phillips, O. L., Williams, S. E., Jan. 2004. Extinction risk from climate change. Nature 427 (6970), 145–148.

Thompson, M., Ellis, R., Wildavsky, A., 1990. Cultural Theory. Westview Press, Boulder.

Trajtenberg, M., Henserson, R., Jaffe, A., 1992. Ivory tower versus corporate lab: An empirical study of basic research and appropriability. NBER working paper 4146.

Unruh, G., 2000. Understanding carbon lock-in. Energy Policy 28 (12), 817–830.

US EPA, 1995. Flue gas desulfurization technologies for control of sulfur oxides: Research, development, and demonstration. Tech. rep.

US EPA, 2003. Air pollution control technology fact sheet. EPA-452/F-03-034.

van der Kaars, S., Miller, G. H., Turney, C. S. M., Cook, E. J., Nürnberg, D., Schönfeld, J., Kershaw, A. P., Lehman, S. J., Jan. 2017. Humans rather than climate the primary cause of pleistocene megafaunal extinction in australia. Nature Communications 8, 14142.
URL http://dx.doi.org/10.1038/ncomms14142

Van der Ploeg, F., 2011. Natural resources: Curse or blessing? Journal of Economic Literature Forthcoming.

Veblen, T., 1899. The theory of the leisure class: An economic study in the evolution of institutions. Macmillan.

von Storch, H., Costa-Cabral, M., Hagner, C., Feser, F., Pacyna, J., Pacyna, E., Kolb, S., 2003. Four decades of gasoline lead emissions and control policies in Europe: a retrospective assessment. Science of the Total Environment 311 (1), 151–176.

Wang, S., Hao, J., 2012. Air quality management in China: Issues, challenges, and options. Journal of Environmental Sciences 24 (1), 2–13.

Warde, P., 2007. Energy consumption in England and Wales, 1560-2000. CNR–ISSM.

Weitzman, M. L., 07 2011. Fat-tailed uncertainty in the economics of catastrophic climate change. Review of Environmental Economics and Policy 5 (2), 275–292.

Yeh, S., Rubin, E. S., 2007. A centurial history of technological change and learning curves for pulverized coal-fired utility boilers. Energy 32, 1996–2005.