

Chapter 25

Sustainability

John Hassler, Per Krusell, and Conny Olovsson

25.1 Introduction

With climate change surfacing as one of the major challenges of modern society, many economists have of course taken an interest in the subject. This includes macroeconomists. In a very clear sense, fighting global warming involves aggregates. Climate change itself is global in nature—there is warming all around the world—and it is fundamentally caused by aggregate CO₂ emissions, which are a natural byproduct of the use of fossil fuel as an energy source. Since fossil fuel, since long, has constituted our main source of energy, this is an aggregate issue, at least to the extent energy use is important. The services provided by energy generate a cost that roughly amounts to 5% of GDP globally. This is not a small amount and, moreover, there are many reasons to believe that sharp reductions in energy use might cause significant recessions; the oil-price hikes during the 1970s are often argued to have had such an effect. Relatedly, suppose our world leaders decide to enact policies that are meant to contain global warming. What, then, would be good such policies? Climate change is a long-run issue and direct evidence on the pros and cons of different policy choices in this arena, not only in the short run but also in the long run, is not available. For that reason, macroeconomic modeling appears useful. Thus, one of the main purposes of the present chapter is to develop some simple tools with the goal of assessing the impact of different policy paths on the climate and on our welfare. Fortunately, rather minor departures from our basic macroeconomic models are required for this endeavor, so we will be able to build significantly on earlier chapters.

A broader, global issue is that of *sustainability*: do humans, those presently alive and those of the future, have enough of what they need, for survival and for living happy lives? To what extent should we economize on natural resources of various kinds? Which, if any, broad policy tools should be considered? To be more concrete, we need to pursue policies that aim to limit world output growth or even decrease output: *de-growth*? These are questions that are obviously hard to answer definitively, but the chapter will at least begin to address them with available tools, drawing mostly on those taught in this textbook and on classic results in public/environmental economics.

The chapter begins in Section 25.2 with a general, but short, discussion of sustainability,

placing “the economy”—the subject of the field of economics—in the broader context of natural resources and how the field of economics has, or has not, made contact with natural resources. This connection will receive attention in the second part of the chapter—in Section 25.5—where we discuss basic theory as well as technological change in this context. Before that, however, in Section 25.3, in order to build toward a full analysis of climate change, we explain “how the climate works”: we offer a summary of the natural-science knowledge required for understanding how to analyze global warming from an economics perspective. This section also summarizes some basic facts about fossil fuel supplies as well as the rapidly growing literature on the economic—broadly defined—consequences of climate change. With this background, we turn to our main analysis of the climate-economic interactions in Section 25.4. In this chapter, we employ a summary of how the climate works in the form of a system of dynamic equations and add these equations to our macroeconomic model and, hence, obtain an *Integrated Assessment Model* (IAM). The word “integrated” here thus refers to the combination of the climate sciences and the economic sciences; “assessment” refers to using the model as a laboratory: to evaluate how different policy actions work their way toward climate outcomes and economic outcomes. We also simulate the model to illustrate its usefulness for policy evaluation.

25.2 The economy and the environment

The term “environment” is a standard term in microeconomic theory to represent the basic elements of an economy (preferences, technology, etc.). Here, we use it to refer to “nature”: what planet Earth (and the solar system) gives us access to for free. Today, there is clearly a movement—especially in advanced economies—to preserve the environment in a variety of ways, and this chapter will make contact with some of the issues pertinent to this movement.

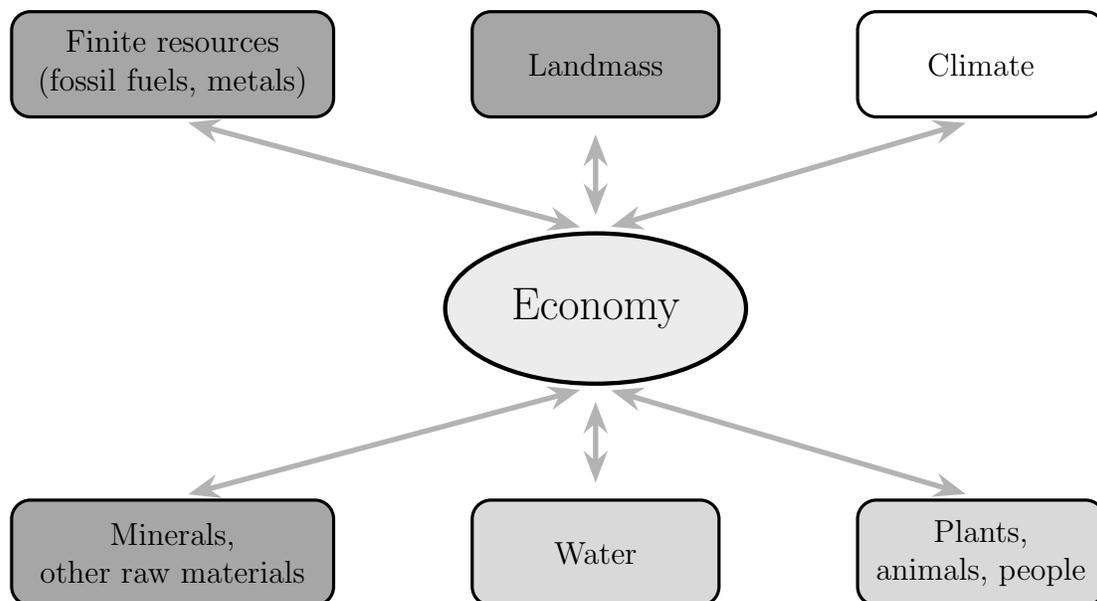


Figure 25.1: The economy and the environment. Darker shades of grey reflect a higher degree of property rights.

Figure 25.1 illustrates and depicts a moment in time, say a year. The “economy” consists of households, firms, governments, and other organizations, along with their preferences and production possibility sets (as well as information), as well as capital stocks and infrastructure built up in the past. The economy in turn interacts with Earth’s environment, which we separate schematically into different components. On the top left we have land mass, which of course is used for a number of economic activities. Land mass has the darkest shade: it is the category which we guess has the highest degree of property rights attached to it. Second, we have finite resources, such as fossils (which can be used as fuels), minerals, and so on. We approximate these to be non-renewable, even though fossils and minerals are reproduced but at a time scale that is so long that we can consider these resources to be in fixed supply. So we use up these resources over time but can also recycle. Much of these resources are owned by someone. Ownership is relevant since Coase’s insights can potentially be used: the owner could themselves manage the resources over time in a purposeful way as a function of market prices, and so it is not obvious at the outset that these resources are overconsumed. The management of natural resources in finite supply is discussed in Section 25.5 below.

Our “global climate” is the average weather (temperature, wind speeds, precipitation, etc.) across the globe and it is determined by a number of factors exogenous to humans, but also—as we will discuss in detail in Section 25.3—by the amount of greenhouse gases in the atmosphere and some of these are directly caused by human economic activity. It is not possible to attach property rights to the atmospheric greenhouse gas concentration mainly because this concentration is even around the world: any emissions spread around the Earth very fast. Thus the global climate is in the lightest in the picture. Of course local climate differences are tied to local land ownership and so local climate can be “transacted”, for example through tourism.

Human activity also crucially uses water—oceans, lakes, and rivers—in a variety of ways, and we can affect the amount of water available, too. Then there are forests, plants, etc., i.e., living organisms; like water, these are inputs into production and consumption and relevant for leisure activities. The concept *biodiversity* refers to the amount of variety within this category.

Finally, there is solar energy: it is a form of radiation and gives us energy without which we would not survive. The solar energy we have access to is a flow. It does vary somewhat over time and is an exogenous determinant of our climate.¹ Solar energy can be captured to some extent by land owners, hence giving it a relatively dark color.

The amount and quality of water, biodiversity, and land are all endogenous to human activity as well as to the climate. In Section 25.3.5 below, we discuss how climate change causes “damages” by reducing the value (to humans) of elements within these categories.

The chapter will first discuss the third category, the climate, and its interaction with the economy. To study this interaction, we employ so-called Integrated Assessment Models—IAM from now on. Even though the details may differ substantially between different IAMs, they all consist of three main building blocks: an economy, a carbon cycle, and a climate module. The term “integrated” refers to the fact that the building blocks affect and interact

¹The solar energy flow itself cannot be affected by humans but the part of it that reaches our planet can, at least in principle. One form of *geoengineering* involves sending up giant “parasols” into space with the purpose of reflecting sunlight before it reaches us.

with each other, as illustrated in Figure 25.2.

Starting with the economy, this block typically consists of a growth model with households and firms. The firms use fossil energy to produce, thus drawing down on an element of the first environmental category in the previous picture. This gives rise to CO₂ emissions as a by-product when the fossil fuel is burnt. The generated carbon dioxide then enters the next block, i.e., the carbon circulation (or carbon cycle).

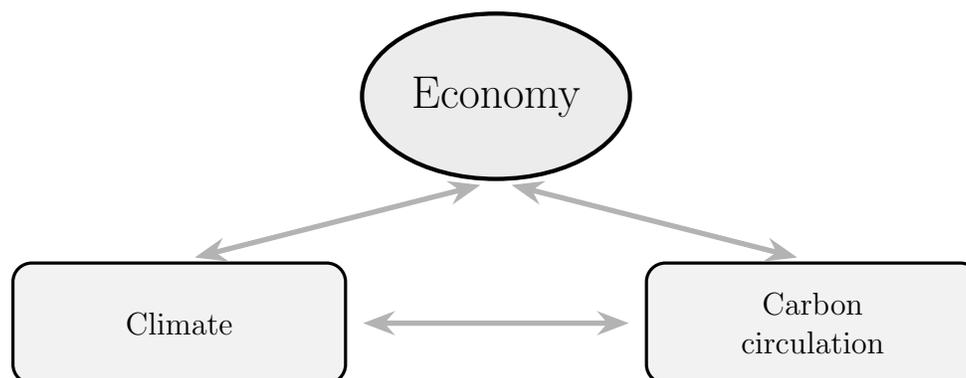


Figure 25.2: Integrated Assessment Models, IAMs, consist of three building blocks: an economy, a representation of the carbon cycle, and a climate module.

As we will discuss below, carbon (dioxide) circulates in a continuous flow between ground, air and sea, but what matters for climate change is the amount of CO₂ in the atmosphere. We will discuss this phenomenon too below. A time path for the stock of CO₂ in the atmosphere is thus the output from the carbon cycle block that enters the climate block where it affects the balance between the incoming energy from the sun and the outgoing heat radiation from Earth to space. Finally, the climate affects the economy through climate-related damages. It is this clockwise rotation that typically is modelled in IAMs, but as is indicated in Figure 25.2, in general, the arrows go both ways.²

The IAM we will build below is a neoclassical growth model—as climate change plays out on the time scale of decades and centuries, a long-run perspective is necessary—together with a climate module and a carbon cycle module, each in the form of a set of dynamic equations. To understand the nature of the equations we add to the growth model, we will go through the natural-science basics next, in Section 25.3. This is a textbook for economists, however, and it is possible to skip over this section, thus taking as given the extra equations we later use. Sections 25.3.4 and 25.3.5 may still be relevant: they discuss the fossil reserves left on the planet, along with the costs of extracting them, and how we measure damages (and how large they are estimated to be).

²For example, the climate affects the carbon circulation because warmer oceans absorb carbon less easily; increased CO₂ concentration in the oceans makes them more acid, which has negative side effects on our economy, broadly defined; and the economy's burning of fossil fuel directly affect the climate by emitting soot particles into the air, which affects warming.

25.3 Climate change: natural-science background

Let us start with a basic primer on the interaction between the economy and the climate. We will then go into detail on each topic in the subsection below. For over 100 years, it has been known that the amount of carbon dioxide (CO_2) in the atmosphere affects the Earth's *energy balance*. This balance captures the difference between the inflow of energy—in the form of solar radiation—and the outflow of energy, with 1/3 of the outflow consisting of direct reflection and 2/3 consisting of heat radiation.

The carbon dioxide does not affect the energy inflow, because solar radiation passes right through the CO_2 . The outflow of energy, however, is affected by the CO_2 in that it makes it more difficult for heat to radiate back into outer space. If the system is initially in balance, an increase in the amount of CO_2 in the atmosphere will lead to a surplus in the energy balance, and as a result, the temperature will rise.

This phenomenon is known as the greenhouse effect. The name is motivated by the fact that it acts exactly like a greenhouse, where solar radiation passes right through the glass of the house whereas the heat gets trapped inside the greenhouse. The greenhouse effect is not entirely a bad thing. Without greenhouse gases in the atmosphere, the Earth would, in fact, be covered in ice and the average surface temperature would be around minus 20°C instead of the close to 15°C that currently prevails.

The emission of greenhouse gases thus leads to climate change and global warming and these changes then affect our economies. Even though these effects from climate change—referred to as *damages*—will vary across space and take many different forms, the aggregate result is projected to be negative for human welfare. We will discuss climate-related damages in detail below, but for now it suffices to state that the damages will be the result of more intensive storms, increased frequency of forest fires, rising sea levels and floods, and so on. These changes in the climate are expected to destroy property and capital, kill and injure people, and may lead to mass migration, political instability, and economic crises.

Because it is generally free to emit greenhouse gases into the atmosphere, even though these emissions gives rise to economic costs, the emissions constitute negative externalities. When market failures such as externalities are present, policy is needed to restore efficiency.

We now describe the basics of the climate system (Section 25.3) and the carbon cycle (Section 25.3.2). In Section 25.3.3 we then describe a relatively recent finding that characterizes the combined carbon-cycle and climate systems into an approximately constant *carbon-climate response* (CCR) to accumulated emissions. Further, in Section 25.3.4, we discuss the world supplies of different kinds of fossil fuel and, finally, in Section 25.3.5 we briefly cover the literature on damage measurements.

25.3.1 The climate system

Let us now describe how CO_2 affects the dynamics of the average atmospheric surface temperature, denoted T and measured as deviations from the pre-industrial temperature. The energy balance is also defined with respect to the pre-industrial situation. The greenhouse effect has a positive effect on the energy balance by reducing the energy outflow. Such changes to the energy balance are called *forcing*.

A second term in the energy balance is due to the fact that hotter objects radiate more

energy. Thus, as Earth gets warmer, more energy is radiated to space. We approximate this effect to be proportional to T , with a coefficient κ . Finally we approximate the change dT/dt as being proportional to the surplus in the energy balance with a coefficient σ_1 . For a given forcing f time path, the evolution T (in continuous time) is

$$\frac{dT(t)}{dt} = \sigma (f(t) - \kappa T(t)). \quad (25.1)$$

The coefficient κ incorporates *feedback effects*, such as the phenomenon that melting ice reduces the direct reflection of incoming sunlight and the effects of increased temperature on cloud formation. The sum of all the feedbacks is highly uncertain and hence the value of κ is uncertain. After an increase in forcing to a constant level f , the temperature increases and the system eventually reaches a new steady state when $T(t) = f/\kappa$.

Human activity has substantially increased the concentration of greenhouse gases in the atmosphere. The dominant greenhouse gas that stems from human activity is CO_2 . It is currently responsible for about two thirds of the forcing from greenhouse gases in the atmosphere.³ Since the beginning of the Industrial Revolution, we have increased the concentration of CO_2 in the atmosphere by around 50 per cent. Humans also emit particles and aerosols that have a direct effect as well as effects on cloud formation. In sum, these effects are likely cooling, but the uncertainty here is large.

So far, we have learned that higher concentration of CO_2 in atmosphere reduces the outgoing flow of energy in the form of infrared radiation. This phenomenon is well approximated by a logarithmic function (Arrhenius greenhouse law, 1896). If we denote the stock of carbon dioxide in the atmosphere by S , and the pre-industrial level (i.e., the level before the Industrial Revolution) by S_0 , then forcing due to carbon dioxide is well approximated by

$$f_{\text{CO}_2} = \frac{\eta}{\log 2} \log \frac{S}{S_0}. \quad (25.2)$$

An often used value of η is 3.7. Combining the steady state of (25.1) with (25.2), we obtain the steady-state temperature associated with a given amount S :

$$T(f(S)) = \frac{\eta}{\kappa \log 2} \log \left(\frac{S}{S_0} \right). \quad (25.3)$$

The ratio η/κ is labelled the *equilibrium climate sensitivity* (ECS), and it measures how much the long-run temperature will increase from a doubling of the CO_2 level. The IPCC's (Intergovernmental Panel on Climate Change) 6th report states that the ECS is "likely" between 2.5°C and 4°C , with a "best" estimate of 3°C . A 90-% interval is 2°C - 5°C . The IPCC has revised their estimate of the ECS several times. The interval reported in the 6th report is more narrow than that in the 5th report where it was given as being between 1.5°C and 4.5°C , with no best estimate given.

Equation (25.1) does not take into account the fact that the heating of the atmosphere involves interaction with the oceans. In particular, the oceans change temperature quite

³However, by far the most important cause of the greenhouse effect is water vapor, but water vapor does not derive directly from human activity.

slowly, creating a long-lasting but not permanent cooling effect on the atmosphere. Denoting the ocean temperature by T^L (again measured as a deviation from its pre-industrial level), a difference between the atmospheric and ocean temperatures will create a third term in the energy balance for the atmosphere. This term is the flow of energy from the atmosphere to the ocean and is approximately proportional to the difference between the two temperatures, with coefficient σ_2 . Finally, we need to specify the law-of-motion for T^L and we assume it to be proportional to the temperature difference $T - T^L$ (and thus to the energy flow from the atmosphere to the oceans), with coefficient σ_3 . Discretizing the resulting system yields

$$T_t = T_{t-1} + \sigma_1 \left(\frac{\eta}{\log 2} \log \left(\frac{S_{t-1}}{S_0} \right) - \kappa T_{t-1} - \sigma_2 (T_{t-1} - T_{t-1}^L) \right) \quad (25.4)$$

and

$$T_t^L = T_{t-1}^L + \sigma_3 (T_{t-1} - T_{t-1}^L). \quad (25.5)$$

It is straightforward to show that for a given S , this system has a steady state that is also given by equation (25.3). However, since an empirically reasonable parameterization requires σ_3 to be much smaller than σ_1 , convergence to the steady state is slower than for equation (25.1).

25.3.2 The carbon cycle

The strength of the greenhouse effect depends on the amount of CO_2 in the atmosphere. CO_2 is emitted into the atmosphere and then circulates between three main reservoirs—the atmosphere, the surface oceans and biosphere, and the deep oceans—and we refer to this system as the *carbon cycle*. Denoting the stock of CO_2 in the surface oceans and biosphere and in the deep oceans by S^U and S^L , respectively, the carbon cycle block can be formulated as follows.

$$S_t - S_{t-1} = \phi_{12} S_{t-1} + \phi_{21} S_{t-1}^U + E_{t-1} \quad (25.6)$$

$$S_t^U - S_{t-1}^U = \phi_{12} S_{t-1} - (\phi_{21} + \phi_{23}) S_{t-1}^U + \phi_{32} S_{t-1}^L \quad (25.7)$$

$$S_t^L - S_{t-1}^L = \phi_{23} S_{t-1}^U + \phi_{32} S_{t-1}^L, \quad (25.8)$$

where E denotes CO_2 emissions.

The above specification is used in Nordhaus's RICE model. Folini, Kubler, Malova, and Scheidegger (2021) show that the carbon-cycle model above closely replicates the mean behavior of the most advanced Earth System Models (CMIP5), if the parameters are chosen appropriately.⁴ The parameter values suggested by Folini et al. (2021) are $\phi_{12} = 0.053$, $\phi_{21} = 0.0536$, $\phi_{23} = 0.0042$, and $\phi_{32} = 0.001422$, when a period is one year. The initial stocks are set to $S_{2015} = 850$, $S_{2015}^U = 765$, and $S_{2015}^L = 1799$.

Note the important difference between measuring emissions in units of CO_2 and carbon units. One kg of carbon corresponds to approximately 3.67 kg of CO_2 .⁵ The distinction

⁴CMIP5 stands for Coupled Model Intercomparison Project, Phase 5; it is used by the IPCC and is considered to be the state-of-the-art model in climate science.

⁵A mole of carbon atoms weighs 12 grams, whereas a mole of oxygen weighs 16 grams and we obtain $(2 \times 16 + 12)/12 \approx 3.67$

sometimes causes confusion since, e.g., proposed policy tax rates are sometimes expressed per ton of CO₂ and sometimes per ton of carbon.

An alternative to a two-dimensional, linear system, is a one-dimensional reduced form for excess carbon, S :

$$S_t - S_0 = \sum_{s=0}^{\infty} (1 - d_s) E_{t-s}, \quad (25.9)$$

where the key feature again is linearity.⁶ Further simplifying approximations can be made with reference to the key characteristics of the carbon cycle as described by the IPCC and Archer (2005). Specifically, about 20–25 per cent of all carbon dioxide emitted stays in the atmosphere for up to over a thousand years. About 50 per cent disappears in a few decades. The rest takes a few centuries to circulate onwards to the oceans (where it contributes to acidification). Thus,

$$1 - d_s = \phi_L + (1 - \phi_L)\phi_0(1 - \phi)^s \quad (25.10)$$

is proposed in Golosov, Hassler, Krusell, and Tsyvinski (2014) to match these features directly, with the share of emissions that remains in the atmosphere forever represented by ϕ_L , the share that leaves the atmosphere within a period given by $1 - \phi_0$, and the remainder $(1 - \phi_L)\phi_0$ depreciating geometrically at rate ϕ . For a decadal time scale, the calibration $\phi_L = 0.2$, $\phi_0 = 0.393$, and $\phi = 0.0228$ matches the data quite well.

25.3.3 Constant carbon-climate response and the carbon budget

A key feature of the system described by equations (25.4)–(25.8) is that T_t is approximately proportional to cumulative emissions from pre-industrial times until t . This implies that temperature keeps increasing as long as emissions continue and stay constant at the level it has then reached for at long time after emissions have come to an end. The latter is due to two opposite effects that neutralize each other in equation (25.4). On the one hand, CO₂ slowly leaves the atmosphere, reducing forcing (the first term within parenthesis). On the other, the oceans slowly get warmer, reducing their cooling effect (the third term). These terms approximately balance on a centennial scale. These features also hold for the advanced climate models used by IPCC. Formally, we have

$$T_t \approx \sigma_{CCR} \sum_{s=0}^t E_s,$$

where T denotes the global mean surface temperature, E global emissions and σ_{CCR} the constant of proportionality that is referred to as the carbon-climate response (CCR), or the Transient Climate Response to Cumulative Carbon Emissions (TCRE). Note of course that this finding is one that combines the climate and carbon modules and generates a direct mapping from emissions to climate.

In the latest (6th) IPCC report, σ_{CCR} is stated as “likely” being between 1.0°C and 2.3°C per 1,000 GtC, which corresponds to 0.27°C–0.63°C/TtCO₂).⁷ The relationship between

⁶The word “excess” here refers to the level above that which prevailed before the Industrial Revolution.

⁷The term likely is explained to be interpreted as a 2/3 confidence interval.

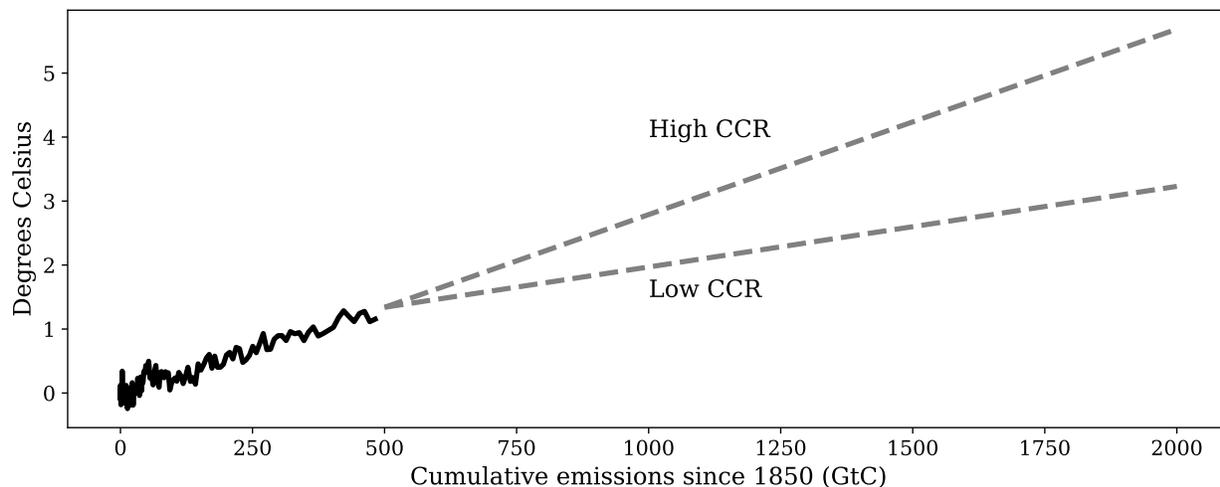


Figure 25.3: Cumulative carbon emissions and the global mean temperature. Solid line: historical data 1850-2022. Dashed lines: forecasts based on constant CCR.

cumulative emissions and the change in global mean temperature based on these results is represented in Figure 25.3. The solid line represents data, i.e., observed cumulative emissions and temperatures from 1850 to 2022. As we see, apart from substantial short-run natural variation, the historical relation is approximately proportional. The dashed lines represent forecasts of the future relationship based on the two endpoints of the uncertainty interval provided by IPCC.⁸

Several valuable insights can be obtained from the finding that the CCR is approximately constant. One is irreversibility: the relation gives temperature at point in time as a function of total past emissions. Hence, there is no depreciation, despite the fact that the carbon cycle fundamentally builds on depreciation, with carbon emitted to the atmosphere, where it acts as a greenhouse gas and causes warming, slowly moving into other sinks where it does not cause warming directly. Thus the conclusion is that if warming is to be stopped, emissions need to stop. Second, a constant CCR also suggests a notion of a *carbon budget*: given that we wish to maintain warming, by a given date, by a certain target amount, we can use the σ_{CCR} to compute how much can be emitted before that date in order to meet the target. Of course, the uncertainty about the CCR coefficient makes this computation somewhat problematic. Up until 2022, we had burnt around 650 gigaton (GtC) of carbon. If the CCR is 1, the warming caused by all historic CO₂ emissions is $0.65 \times 1 = 0.65^\circ\text{C}$. To cause a warming of 1.5°C , another 850 GtC would need to be emitted, which would take another 85 years with the current emissions rate. If, on the hand, the CCR is at the upper end of its likely range at 2.3, the 650 GtC already emitted cause a warming of $2.3 \times 0.65 = 1.5^\circ\text{C}$.⁹ Third, the constant CCR means that if a target is set for a certain future date, and a carbon budget has been correspondingly calculated, then there are many paths that could

⁸To account for the effect of other greenhouse gases, we have, somewhat arbitrarily, assumed the same proportional relation between this and the effect of CO₂ as has been previously observed. This amounts to multiplying σ_{CCR} by 1.26.

⁹Note that this calculation does not take into account the effects of other greenhouse gases. These effects, however, are much less long-lasting than those of CO₂-emissions.

attain the target, ranging from high emissions early and gradual reductions over time to the reverse kind of time path, with radical reductions early and less and less radical reductions over time—all subject to the same total accumulated emissions. These alternative ways of reaching the target would, however, not be equivalent for human welfare, which indirectly suggests a weakness in just having a target set by a future date. To evaluate the consequences for human welfare, however, requires us to think about the climate-economy interactions, to which we turn in Section 25.4 below. A final point regards the linearity: a constant CCR also means a ruling out of non-linearities such as tipping points in the global climate.

25.3.4 The fossil energy supply

Against the background of a carbon budget, the question of how much global warming that can be expected depends not only on the CCR, but also on how large the stock of fossil fuel actually is. Here, there is an important distinction between *reserves* and *resources*. Using the definition from the BP (2021), **total proven reserves** is generally taken to be those quantities that geological and engineering information indicates with reasonable certainty can be recovered in the future from known reservoirs under existing economic and operating conditions. **Resources** includes the amount of a geologic commodity that exists in both discovered and undiscovered deposits. Hence, it is by definition a “best guess” that includes quantities that would not be profitable to extract today but that could become profitable in the future.

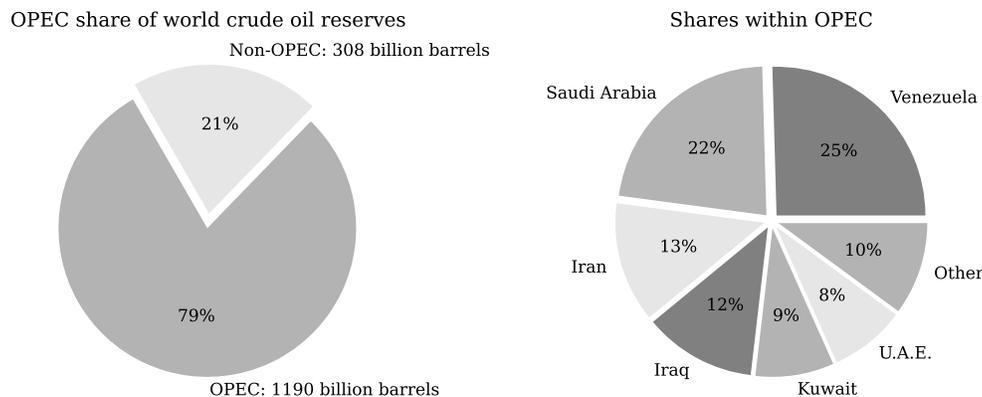


Figure 25.4: Oil Supply.

Source: OPEC (2019).

Starting with oil, world proven crude oil reserves are displayed in Figure 25.4. The figure shows that OPEC accounts for almost 80 % of world crude reserves with their 1,190 billion barrels, whereas the non-OPEC reserves consist of 308 billion barrels.

It is also useful to make a distinction between “conventional” and “unconventional” oil. The key difference is that conventional oil can be extracted with traditional drilling methods: it is liquid at atmospheric temperatures and therefore flows without additional stimulus. Unconventional oil requires more advanced extraction methods because it does not flow on its own. Marginal extraction costs are therefore (very) low for conventional oil: only a fraction of the market price. The approximation of marginal cost to 0 in the section

above on Hotelling rents precisely referred to conventional oil. For unconventional oil, the marginal extraction cost is substantially higher and much closer to (or above) the market price. For these reasons, conventional oil is very profitable, whereas unconventional oil is not. As we will see below, all the conventional oil can therefore be expected to be used up.

How much global warming would result if we were to burn all the reserves of conventional oil? It is somewhat hard to assess exactly how much of the oil that is conventional, but it seems reasonable to assume that most of the middle-east reserves are conventional. To compute the level of warming, we then make use of the approximately constant CCR: each 1,000 GtCO₂ of cumulative CO₂ emissions is likely to cause a 0.27–0.63 °C increase in global surface temperature. Converting from barrels to GtCO₂, we note that one barrel equals 1/7.33 tons, that the carbon content in oil is 85%, and that one ton of carbon equals 3.67 ton of CO₂. We thus have that 1,190 billion barrels equals $(1,190/7.33) \cdot 0.85 \times 3.67 = 500$ GtCO₂. Using constant CCR, this delivers 0.13–0.33°C of warming.

Turning to coal, total proven reserves are 1,074 GtC, according to BP (2021). With the carbon content of coal being about 70 percent of its weight. This corresponds to $1,074 \cdot 0.70 \cdot 3.67 = 2,759$ GtCO₂, which if all burnt results in an increase in the temperature of between 0.75°C and 1.73°C. It is important to point out, however, that proven coal reserves are likely underestimation: coal is not particularly profitable, giving only weak incentives to search for it.

Total proven reserves of natural gas, finally, is 190 trillion cubic meters, again according to BP (2021). Using the conversion factor 0.511 kgC/cubic meter, we see that burning all the gas results in less than a tenth of a degree of warming.

What about resources? The numbers are much more uncertain and depend on how likely it is that some resources become profitable to extract. Estimates of resources (and reserves) are provided in Table 25.1.

Table 25.1: Reserves and resources in GtC

	Reserves	Resources (Rogner)	Resources (BGR)
Oil (conv+unconv)	173	≈ 400-2,200	≈500
Coal	1,074	≈6,200	≈11,500
Gas	100	≈350	≈365

Sources: Rogner (1997), BGR (2020). BGR is the Federal Institute for Geosciences.

As can be seen the resources of coal are an order of magnitude larger than the reserves; clearly, if this whole amount of coal were burnt, the global temperature would rise by at least 10°C, again using the approximate CCR.

25.3.5 Damages

In October of 2007, the *Stern Review on the Economics of Climate Change*—a report commissioned by the British government—was released and it had great impact in the policy-making world, arguably precisely because of its discussions of how climate change would “damage” the economy. It pulled together evidence and made a quantitative assessment of

the economic value lost from warming, resulting in strikingly large numbers. This finding, the review argued, called for strong, early action by governments to halt global warming. In the present section we discuss the effects of climate change on human welfare, drawing on a growing literature, which includes touching on the findings in the Stern Review.

From an economic perspective, as the Review and the earlier work by others such as William Nordhaus had pointed out, the effects of warming on human welfare are to be regarded as classical externalities: consequences of economic behavior (emissions in the context of burning fossil fuel) for human welfare that are not priced by the market. An approach for policy, proposed in the Stern Review, by Nordhaus, as well as by many other influential economists, is therefore to aggregate all the externalities (measured in dollar terms), and tax emissions by that amount: the so-called Pigou tax.¹⁰ Moreover, in the case of climate externalities, though it was not emphasized in the previous sections, because carbon dioxide emissions spread very quickly around the globe, the total externality is the same independently of where the emission takes place. Thus, an application of a Pigou tax in this case means that the tax per unit of emission should be the same everywhere. There is no universal agreement on a Pigou tax being the best policy, and there is even less agreement that a tax, if applied, should be set at the same level in all countries. Regardless, the measurement of damages is still a very important part of the climate-economy field.

Economists have approached damage measurement in two different, and complementary, ways. One is the *bottom-up* approach, whereby all kinds of damages are identified, with a variety of econometric methods, assigned dollar values using market prices, and then added up. Clearly, there is a space/country dimension here, so the endeavor of covering all damages is extremely ambitious. Another approach is the *top-down* method, which focuses on relating aggregates—such as production, mortality, or migration—to observed changes in the climate over time or to climate differences across space. We now go through some results from this (rapidly growing) literature.

The bottom-up approach

The idea here is to measure damages of all kinds imaginable (i.e., broadly defined), assign them (dollar) values, and then add up over a given time period such as a year, thus generating a flow value. Moreover, the goal is to estimate the damage effects of each additional degree of global warming (along with the broader notions of climate change, such as changes in precipitation patterns, that warming entails), thus obtaining a *damage function*, i.e., a mapping from global temperature to a dollar value. Damages functions can be constructed separately for different countries or regions, for different sectors of production, and so on. Damages will, moreover, depend on time or economic state variables. For example, if warming affects output, then economic growth can make damages larger if they are proportional to output: there is more to lose with higher output. If people's health is negatively affected too, then the number of people factors in, so that population growth too will affect the measure of damages. Clearly, one cannot literally include everything in a global damage measure, at all

¹⁰The so-called Pigou Club, formed in 2006 by N. Gregory Mankiw with the general aim to call for governments to correct externalities with taxes and subsidies, contains a long list of very well-known economists. Arthur Pigou was a British economist who introduced the externality concept and proposed this policy; see [Pigou \(1920\)](#).

points in time and space and for all contingencies, so for the approach to be implementable in practice it is important to make attempts generalize from specific studies.

Nordhaus pursued a bottom-up approach, thus looking separately at damage effects via agriculture, sea-level rise, health, non-market amenity impacts, etc., but in adding up found limited quantitative impacts. He included the potential for catastrophes, with probability and severity assessments collected in a survey among experts rather by direct measurement and still arrived at a damage function with very modest effects: global warming of 2 degrees—the Paris agreement now aims at a maximum of 1.5 degrees—would only generate a flow damage value of less than a percent of global GDP. Nordhaus argued that the damage function was likely convex, though the degree of convexity is hard to assess.¹¹ He also formulated separate damage functions for world regions but associated great uncertainty here as well.

Later on the European Commission started the PESETA project, which has studied 11 damages categories, including heat waves, windstorms, droughts, flooding, wildfires, agriculture and energy supply, across the EU. PESETA IV reports aggregate EU damages that are nontrivial (on the same order of magnitude as those found by Nordhaus) with significant dispersion across different parts of the EU: the north is even estimated to gain from warming, but the Mediterranean economies lose significantly.

The number of studies with microeconomic damage estimates is growing rapidly, which is reassuring. We are, however, still far from a comprehensive understanding of the nature of the effects of warming on people and on our economies.

The top-down approach

The top-down approach starts from aggregate information, typically on production/GDP or other economic aggregates, such as health indicators and mortality, and then tries to relate these figures to climate variables. There are multiple econometric challenges in interpreting the typical regressions of this sort, but they have nevertheless played an important role in the literature on damages. The end results, moreover, are “reduced forms”, i.e., not pinpointing through what channels warming affects output or mortality.

One approach has been to use time-series analysis where climate change is represented by changes in weather (temperature) over time. Here one challenge is that “climate” is the “distribution of weather”, so that short-run changes in temperature (and the economic reactions to them) cannot necessarily be interpreted as climate change. [Dell, Jones, and Olken \(2014\)](#) find, based on a time series study of 125 countries over the period 1950–2005, that a one degree higher temperature leads to a one percent fall in output growth among poor countries but no effect for rich countries, and no level effects more generally. [Kahn, Mohaddes, Ng, Pesaran, Raissi, and Yang \(2019\)](#) use data from 174 countries over the period 1960–2014 and also find sizable effects on growth, e.g., that a four degree warmer global average temperature would reduce global GDP by 0.04 percent per year, accumulating to a reduction in global GDP in 2100 by 7%. A recent study on mortality ([Carleton, Jina, Delgado, Greenstone, Houser, Hsiang, Hultgren, Kopp, McCusker, Nath, Rising, Rode, Seo, Viaene, Yuan, and Zhang, 2022](#)) finds, in an extreme scenario with a 5-degree increase in

¹¹Nordhaus (1994).

global average temperature by 2100, that the total global cost of mortality (and adaptation to avoid it) amounts to 3.2% of GDP.

Another approach is to rely on cross-sectional data to glean the effects of climate on aggregates (regions or countries). This approach is related to [Hall and Jones \(1999\)](#), which accounts for differences in output by measuring the distance to the equator (which is closely related to average temperature), but it is actually preceded by [Mendelsohn, Nordhaus, and Shaw \(1994\)](#), which tries to explain agricultural output in a similar manner, with some attempts to hold constant institutions. [Dell et al. \(2014\)](#) also highlight the remarkable negative correlation between output and average temperature. Looking at regions at an even higher level of resolution, [Cruz and Rossi-Hansberg \(2022\)](#) and [Krusell and Smith \(2022\)](#) find a U-shape in the relation between average temperature and productivity, with a “sweet spot” around 11 degrees Celsius (measured as an annual average). Clearly, these cross-sectional estimates carry information about the effects of the climate on output but, equally clearly, there are confounding factors.

Summaries and modeling

We now summarize the damage estimates in the literature by looking at two meta studies and then comment on how damages are modeled in the IAMs.

Meta studies We report two meta studies relating global warming to losses in global GDP in [Figure 25.5](#). Both of these studies draw on a large number of different contributions using a variety of methods. As the figure shows, the costs of one or two degrees of warming are on average quite low but there is convexity as a function of temperature.

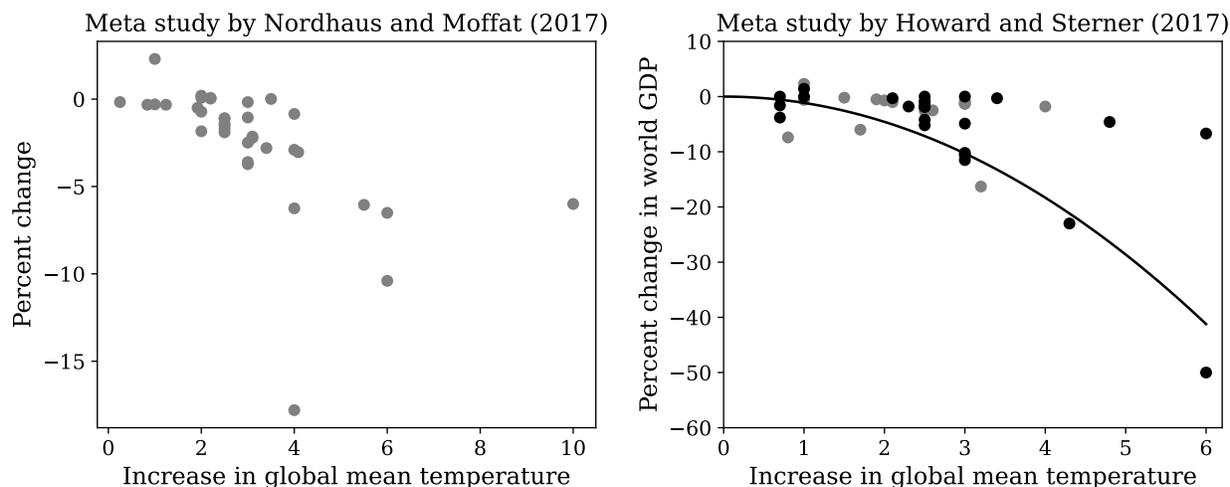


Figure 25.5: Meta studies of damages. Left: study by [Nordhaus and Moffat \(2017\)](#). Right: study by [Howard and Sterner \(2017\)](#) with black dots showing non-duplicate studies and the regression line preferred by those authors.

Damages: modeling the function(s) How can the economic damages be modeled? Nordhaus pioneered an approach where all damages are contained in a TFP factor. [Nordhaus](#)

and Sztorc (2013) specifies damages, as a share of GDP, at a temperature T degrees above the pre-industrial level as given by

$$D(T) = 1 - \frac{1}{1 + 0.00267T^2},$$

where the coefficient in front of T^2 determines the curvature. This formulation can be approximated locally around 0 by $0.024\left(\frac{T}{3}\right)^2$, a convex function indicating that at three degrees of warming, 2.4 percent of output is lost.¹² Nordhaus also considered including higher order polynomials; Weitzman argued for a polynomial of order 6.8. The higher orders, however, are not based on the use of data but rather as examples.

It is of course also possible to introduce damages in other parts of one's integrated assessment model. Endogenous life-expectancy fits naturally by introducing variable population size and a value of life, one can have endogenous depreciation by taking capital destruction (due to, say, flooding) into account, and so on. One common approach is to simply have damages additively in utility—we describe a simple model of that sort below—because in such a case there is no feedback from damages to economic activity, which makes the model simpler to solve.

An interesting case is obtained by mapping the carbon dioxide concentration in the atmosphere, S above, directly into percentage TFP losses. In particular, Golosov et al. (2014) show that the maps from S to T , by the use of the Arrhenius law described in Section 25.3, and then further from T to D according to Nordhaus's estimated function, can be very well approximated by a TFP factor $e^{-\gamma S}$. I.e., if output before damages is multiplied by this expression we obtain a surprisingly good approximation (for a specific numerical γ value) to the original formulation. This, we shall also see below, turns out to be another useful case for applied use.

Damages: adaptation In response to global warming and damages therefrom, it is possible to respond in a variety of ways to lessen the damages. This phenomenon is referred to as adaptation. For example, air conditioning can be installed in response to heat increases, walls can be built to protect against flooding, and so on. Reduced-form damage assessments thus should be interpreted as incorporating adaptation, whereas at least some microeconomic damage estimates simply do not factor in the possibility of adaptation.

There are reasons to think that damages net of adaptation are less convex than damages not taking adaptation into account. Consider the following example: total damages, in output units, are

$$\frac{(T - a)^x}{x} + a,$$

where a is adaptation. This formulation captures a possible convexity—with $x > 1$ —and the assumption that adaptation costs scale linearly: a wall twice as high costs twice as much. Minimizing this expression with respect to a delivers adaptation of $T - 1$ (so long as $T \geq 1$). The reduced-form damage is then, again for $T > 1$, $D(T) = T + \frac{1-x}{x}$, i.e., the damage function is linear, regardless of the value of x !

¹²Globally, the damage function is not convex: it asymptotes to 1 (all of output is lost) as T goes to infinity.

Damages: discounting Finally, we comment very briefly on a contentious issue: how to add up damages that occur over time. In a typical macroeconomic model based on microeconomic foundations, welfare for each agent is simply measured by the utility function of that agent, using a revealed-preference argument. So if we have measures of utility function parameters, the revealed-preference approach would be straightforward to apply. In a representative-agent model with standard discounting, we would use an estimate of β , the discount factor, to add up utils accruing at different points in time: a utility flow occurring t periods into the future is then discounted by β^t . We also explained in our early chapters that estimates of β come from observations on the average long-run (riskless) real interest rate and consumption growth together with the Euler equation and a measure for consumption curvature. Thus, a β around 0.985 is standard for an annual time period. This approach is followed by Nordhaus in his work.

In the Stern Review, an argument is put forth that a much higher β , e.g., 0.999, is appropriate. The idea is thus to “assign a higher weight to future generations”, which may not be appropriately taken into account using the revealed-preference approach above. The choice of β matters greatly on a quantitative level, since global warming is (largely) irreversible and damages long-lasting. In our view, the choice is more a matter of philosophy, not economics. It deserves to be pointed out, however, that if one disagrees with market discounting (for philosophical reasons), there may also be reason to revisit the assessments of how markets make investment decisions more generally, potentially thus implying a need for large subsidies to saving and investment.

25.4 Integrated assessment models

We are now ready to put together an integrated assessment model. The quantitative model we will use is, as advertised, a neoclassical growth model in its core, but we will actually begin with a static model, as it can provide us with significant insight, including how to think about climate policy.

25.4.1 A static one-region model

The static model is a very simple static version of the macroeconomic models earlier in the book. It might, at first thought, seem strange to think of climate change in a static setting since change by definition is non-static. However, this setting has proven useful in providing intuition for much more complex models.¹³ The model has three sectors and all sectors are perfectible competitive. The final-goods sector uses capital, labor, and energy services to produce a final good. Energy services is a new element relative to earlier chapters. We assume a very simple structure whereby these energy services are produced one-for-one from coal. Coal, finally, is produced at a constant marginal cost, labeled p , in terms of the output good.¹⁴ Coal is thus for simplicity not regarded as a resource in finite supply here. Our

¹³See, for instance, [Hassler, Krusell, and Olovsson \(2021b\)](#).

¹⁴Recall from earlier in the book that this is equivalent to regarding coal as being produced by the same basic inputs, and with the same production technology isoquants, as for final goods, only at a relative TFP of $1/p$.

dynamic model below will, however, consider fossil fuel in finite supply.

The production, Y , of the final good is given by

$$Y \equiv F(K, L, E) = AK^\alpha L^{1-\alpha-\nu} E^\nu, \quad (25.11)$$

where A is total factor productivity, K is the capital stock, L is labor, and E energy services. Unitary elasticity between inputs is a good approximation in the longer run (and the static model is meant to capture a longer run, e.g., 100 years) but not in the shorter run, where the elasticity is much lower between energy and the other inputs. In Section 25.5.3 we discuss this issue in detail. Without loss of generality, we normalize L and K to unity, which implies that output can be written as

$$Y \equiv f(1, 1, E) = AE^\nu. \quad (25.12)$$

In a decentralized allocation, the energy-service provider buys the energy input at marginal cost from the fuel producer. The profit-maximizing problem in the final-goods sector is then given by

$$\pi_e = \max_E AE^\nu - (p + \tau)E, \quad (25.13)$$

where τ is a tax on coal use. We will, below, refer to (25.13) as the private economy since this is the profit-maximization problem for a representative firm in the market economy.

This simple IAM assumes a direct link between coal use and damages. As explained above, there exists a mapping from E to S (see equation (25.10)) and then from S to T (see equations (25.4)-(25.5)). Here, these relationships are summarized in a way that implies that the damage function is quadratic in coal use:

$$D(e_1) = \gamma E^2. \quad (25.14)$$

The above function captures the typical property in climate models that was described above, i.e., a proportional relation between accumulated emissions and the temperature. Combining this property with the assumption that damages are moderately convex in temperature yields (25.14). Damages are, moreover, assumed to be separable from output, which implies that consumption is given by

$$C = AE^\nu - pE - \gamma E^2. \quad (25.15)$$

Let us now consider the equilibrium in the market economy. Because damages are additively separable, the temperature and the damages do not appear in the problem for the firm, i.e., equation (25.13). The first-order condition for this problem can then be written as

$$\nu AE^{\nu-1} = p + \tau. \quad (25.16)$$

The only unknown in the above equation is E so this variable is uniquely pinned down by (25.16); it is given by $E = ((\nu A)/(p + \tau))^{\frac{1}{1-\nu}}$. Note that there is direct negative relationship between E and τ so coal use is decreasing in the coal tax. Note also that, with mappings between E and S and T specified, it is possible to compute the temperature increases that result from different tax rates in the static period.

Pigou taxation

Turning to the social-planning problem, the objective is given by

$$\max_E AE^\nu - pE - \gamma E^2, \quad (25.17)$$

and the first-order condition can be written as

$$\nu AE^{\nu-1} = p + 2\gamma E. \quad (25.18)$$

Clearly (25.16) and (25.18) differ if $\gamma \neq 0$ and $\tau = 0$. In this case, the market economy will be inefficient and emissions will be too high relative to the efficient level. Even though we cannot solve for τ in closed form, it is clear that implementing the first best in the market economy implies setting $\tau = 2\gamma E^*$, where E^* is the solution to (25.17). This is Pigou's recipe: tax the activity that is causing a social cost and set the tax equal to the marginal externality damage that the activity is causing (Pigou, 1920). At the social optimum in this economy, the marginal damage is equal to $2\gamma E^*$.

For the private economy, i.e., ignoring the externality entirely, $\tau = 0$ constitutes a maximum. Moreover, this objective function is smoothly mountain-shaped, and concave, around a zero tax. It also implies that the costs for the private economy of modestly raising the global carbon tax above zero are negligible. The benefits for overall welfare, however, can potentially be very large if γ is large. In other words, even a modest tax can be very beneficial in alleviating the negative consequences of global warming.¹⁵

Regulation: prices vs. quantities

In practice, it seems that quantity restrictions have been preferred by politicians over affecting prices by the use of taxes. In 2005, the European Union implemented a cap-and-trade system—the EU Emission Trading System (EU ETS)—which is a quantity restriction where firms also can trade in emission rights. This raises the question of how effective quantity restrictions are relative to Pigou taxation.

To answer that question, consider a social planner that implements a quantity restriction. The regulator then decides on a quantity of coal that cannot be exceeded by the market. To ensure that the restriction is honored, anyone that wants to burn a unit of coal (and thus emit carbon dioxide into the atmosphere) also has to hand in an emission right to the regulator for doing so. A quantity restriction therefore requires the regulator to take a stance on how the emission rights are to be allocated. This can be carried out in many ways: the emission rights can, for instance, be handed out to one or several agents for free. Alternatively, the government can sell them in an auction format. Here, we simply assume that the emission rights are handed out for free to the producer of coal. This firm then sell these emission rights to the final-output firm (which, in this economy, is the only one that buys coal).

Every time that the final-good firm now wants to buy a unit of coal, it also has to buy an emission right from the coal producer. The total cost of buying a unit of coal is then $p + \lambda$, where λ is the price of the emission right. Note that the usage of a quantity restriction

¹⁵See Hassler, Krusell, and Olovsson (2023).

transforms the problem from one where, in principle, an infinite amount of coal can be supplied to one where coal effectively becomes a resource in finite supply.

Consider first the case where the quantity restriction requires $E \leq E^*$, i.e., we target the optimal level. The first-order condition for the final-output firm reads

$$\nu AE^{\nu-1} = p + \lambda. \quad (25.19)$$

Will E^* be an equilibrium outcome and, if so, what will the equilibrium price of emission rights be? It is easy to see that, by construction, $\lambda = \tau = 2\gamma E^*$ satisfies the first-order condition and, hence, is an equilibrium. At this quantity, λ has to be positive, since otherwise more E would be purchased than what is available. A lower E is not possible: then all emission rights would not be sold, but that can only occur if $\lambda = 0$, which clearly does not satisfy the first-order condition. The equilibrium is unique. Hence, the price and quantity regulations are equivalent.¹⁶

What is the price of an emission right if the quantity restriction instead selects another quantity? If the quantity is higher than that of the laissez-faire equilibrium, the restriction will not bind and λ will be zero. If it is lower than that, the restriction will bind and one reads off λ from the firm's first-order condition.

Finally, let us briefly consider how the efficient prices and quantities are affected by firm heterogeneity. For simplicity, consider the case with two final-output firms, labeled 1 and 2, that differ in their technologies with respect to energy intensity (ν_1 vs. ν_2). The damage function then reads $D(E_1 + E_2) = \gamma(E_1 + E_2)^2$. The planning problem reads $\max_{E_1, E_2} AE_1^{\nu_1} + AE_2^{\nu_2} - \gamma(E_1 + E_2)^2$ and the first-order condition for firm $i \in \{1, 2\}$ is then $\nu_i AE_i^{\nu_i-1} = 2\gamma(E_1 + E_2)$, which solves uniquely for $\{E_1^*, E_2^*\}$. A carbon tax satisfying $\tau = 2\gamma(E_1^* + E_2^*)$ would deliver the optimal outcome. What about a quantity regulation: would the government need to assign two values for quantities, one for each technology? That would be a possible approach—in theory, given if the government knows all the details of how firms differ—but it is not necessary: a mere restriction for total use not to exceed $E_1^* + E_2^*$ would also work (the equilibrium price of emission rights would again equal the optimal tax). In particular, it would ensure an optimal allocation of energy across firms, since they all face the same price (or tax): $\nu_1 AE_1^{\nu_1-1} = \nu_2 AE_2^{\nu_2-1}$. This is exactly how the EU system works.

25.4.2 A fully dynamic integrated assessment model

The static one-region IAM described in the previous section can straightforwardly be extended to a fully dynamic model which is also richer in other ways. The model that we will formulate here builds on that described in Golosov et al. (2014) and it incorporates 2 regions. Region 1 is inhabited by a representative oil producer that is endowed with a finite amount of conventional oil that it extracts and sells to the other region in a competitive manner. We refer to this region as the *oil producer*. The other region does not have any endowments of conventional oil but, given that they need oil, import it from the oil producer and pay for this import with a common final good. This region is referred to as the *oil consumer*. Each

¹⁶Weitzman (1974) shows that random fundamentals can, under some conditions, break the equivalence.

region is inhabited by a representative consumer with preferences given by

$$\sum_{t=0}^{\infty} \beta^t \log(C_t). \quad (25.20)$$

The law of motion for the atmospheric excess stock of carbon S_t is given by (25.9). The climate is then affected by the concentration of CO₂ in the atmosphere via the greenhouse effect. Here, we employ the direct, exponential relationship between atmospheric carbon concentration S and damages discussed in Section 25.3.5, but modify it slightly so that the production function reads

$$Y_t = \underbrace{\exp(z_t - \gamma S_t)}_{\equiv A_t} L_t^{1-\alpha-\nu} K_t^\alpha E_t^\nu, \quad (25.21)$$

where z is a potentially stochastic productivity trend and γ is a parameter that determines how climate-related damages depend on the level of the atmospheric CO₂ concentration. The climate system, i.e., the temperatures follow the energy budget model in DICE and RICE and is thus given by (25.4) and (25.5).

Energy services in the oil-consuming region are produced and supplied by a local competitive representative firm that combines different energy sources as inputs. There are three available energy sources: conventional oil, coal, and a green energy source.¹⁷ All energy sources except conventional oil are locally produced by competitive firms that are using a production technology that is linear in the final good, as in the static model.¹⁸ Specifically, $p_{\kappa,t}$ units of the final good are required to produce $e_{\kappa,t}$ units of the energy source $\kappa \in \{c, g\}$ in period t , where c and g refer to coal and green, respectively. The total amount of energy services is an aggregate of the energy inputs

$$E_t = [\lambda_o e_{o,t}^\rho + \lambda_c e_{c,t}^\rho + \lambda_g e_{g,t}^\rho]^{\frac{1}{\rho}}, \quad (25.22)$$

where ρ determines the elasticity of substitution, and $\lambda_o + \lambda_c + \lambda_g = 1$.

Final goods are either consumed, invested or used for energy production, and capital is assumed to fully depreciate between periods. This assumption is not too problematic given that a period will be calibrated to be ten years. The resource constraint for the final good is then given by

$$C_t + K_{t+1} = A_t L_t^{1-\alpha-\nu} K_t^\alpha E_t^\nu - p_{o,t} e_{o,t} - p_{c,t} e_{c,t} - p_{g,t} e_{g,t}, \quad (25.23)$$

where $p_{o,t}$ denotes the world market price for conventional oil. Total emissions in period t are given by

$$M_t = e_{o,t} + e_{c,t}, \quad (25.24)$$

where this simple summation reflects the fact that the fossil energy sources all are measured in carbon units.

¹⁷It is straightforward to include additional energy sources. Hassler et al. (2021b) also incorporate hydraulic fracturing or “fracking”.

¹⁸As before, we abstract from the finiteness of coal supply as the amount of coal is so large that the resource restriction is unlikely to bind, unless we are willing to accept global warming at the two-digit Celsius level.

The oil producer extracts oil without any extraction costs, and the total stock of oil in the ground at time t has size R_t .¹⁹ With extraction in period t given by $e_{o,t}$, the law of motion for the stock of oil is given by

$$R_{t+1} = R_t - e_{o,t}, \quad (25.25)$$

subject to

$$R_t \geq 0, \forall t.$$

The government, finally, simply sets a carbon tax and then recycle the proceeds back to the household within the period in the form of a negative income tax rate, Γ_t . The government budget constraint is then given by

$$\Gamma_t (w_t L_t + r_t K_t) = \tau_t (e_{o,t} + e_{c,t}).$$

Households supply labor inelastically and maximize (25.20) subject to the budget constraint

$$C_t + K_{t+1} = (1 + \Gamma_t) (w_t L_t + r_t K_t) = (1 + \Gamma_t) \widehat{Y}_t,$$

where $\widehat{Y}_t \equiv (1 - \nu) Y_t$ is output net of energy expenses.²⁰

Equilibrium

We now describe how to compute the market equilibrium. The problem of the oil supplier has a very simple solution. The budget constraint of this agent can be written $C_{o,t} = p_{o,t} e_{o,t}$, where C_o is the oil supplier's consumption. The maximization of $\sum_{t=0}^{\infty} \beta^t \log(C_{o,t})$ subject to this constraint and (25.25) straightforwardly delivers $e_{o,t} = (1 - \beta) R_t$ for all t . We discuss "optimal extraction" much more in Section 25.5 below.

The problem for the energy-service provider is to derive the demand functions for all fuels; it reads

$$\min_{e_{o,t}, e_{c,t}, e_{g,t}} p_{o,t} e_{o,t} + p_{c,t} e_{c,t} + p_{g,t} e_{g,t} - P_t \left([\lambda_o e_{o,t}^\rho + \lambda_c e_{c,t}^\rho + \lambda_g e_{g,t}^\rho]^{\frac{1}{\rho}} - E_t \right),$$

where P_t is both the Lagrange multiplier on the constraint and the price index of energy services. We know from Chapter 6 that cost minimization for the energy-service provider delivers that this price index is a CES function of the different input costs. The price of energy services can then be shown to be given by

$$P_t = \left(\lambda_o^{\frac{1}{1-\rho}} \widehat{p}_{o,t}^{\frac{\rho}{1-\rho}} + \lambda_c^{\frac{1}{1-\rho}} \widehat{p}_{c,t}^{\frac{\rho}{1-\rho}} + \lambda_g^{\frac{1}{1-\rho}} \widehat{p}_{g,t}^{\frac{\rho}{1-\rho}} \right)^{\frac{\rho-1}{\rho}}. \quad (25.26)$$

The first-order conditions with respect to $e_{\kappa,t}$ then yield

$$e_{\kappa,t} = E_t \left(\frac{P_t \lambda_\kappa}{p_{\kappa,t}} \right)^{\frac{1}{1-\rho}}, \text{ for } \kappa \in \{o, c, g\}. \quad (25.27)$$

¹⁹The approximation that oil is costless to produce is good for conventional oil/natural gas but not for non-conventional resources such as fracking or deep-sea oil.

²⁰The assumption of constant returns to scale implies that profits are zero in equilibrium, and that energy expenses account for the share ν of output.

Producers of the final good maximize profits taking P_t as given, implying that $P_t = \nu A_t L_t^{1-\alpha-\nu} K^\alpha E_t^{\nu-1}$, which can be rearranged to solve for energy service use:

$$E_t = \left(\nu \frac{A_t L_t^{1-\alpha-\nu} K^\alpha}{P_t} \right)^{\frac{1}{1-\nu}}. \quad (25.28)$$

Similarly, wages and the interest rate will be given by $w_t = (1-\alpha-\nu)Y_t/L_t$ and $r_t = \alpha Y_t/K_t$, respectively. The household's Euler equation reads

$$\frac{C_{t+1}}{C_t} = \beta (1 + \Gamma_t) r_{t+1}. \quad (25.29)$$

Defining the savings rate out of net output as $s_t = (\widehat{Y}_t - C_t)/\widehat{Y}_t$, we obtain $C_t = (1 - s_t)(1 + \Gamma_t)\widehat{Y}_t$ and $K_{t+1} = s_t(1 + \Gamma_t)\widehat{Y}_t$. Inserting these expressions and $r_{t+1} = \alpha Y_{t+1}/K_{t+1}$ into (25.29) yields

$$\frac{1 + s_{t+1}}{1 + s_t} = \frac{\alpha\beta}{s_t(1 - \nu)}.$$

The difference equation in s_t defined by the above expression only has one non-explosive solution: $s_t = (\alpha\beta)/(1 - \nu) \equiv s, \forall t$. The savings rate is thus constant and it defines optimal household behavior.

The equilibrium can now be defined by equation (25.21)–(25.28), $s_t = (\alpha\beta)/(1 - \nu)$, and where the state variables evolve according to $K_{t+1} = \frac{\alpha\beta}{1-\nu} (1 + \Gamma_t) \widehat{Y}_t$, $R_{t+1} = \beta R_t$, and $S_t = \sum_{v=0}^t (1 - d_{t-v}) \sum_i M_t$.

Note that the allocation is determined sequentially without any forward-looking terms. This is due to the combination of logarithmic utility, Cobb-Douglas production, full depreciation, and the way that tax revenues are rebated. Moreover, except for the world market price of oil all equilibrium conditions have closed-form solutions. Hence, when it comes to solving the model, finding the equilibrium is only a matter of finding the equilibrium oil price in each period. This is easy since the supply in each period is predetermined at $(1 - \beta) R_t$. As a result, the model typically solves in under a second also when it is expanded to include many regions.

Turning to the social planning problem, it is given by

$$\max_{\{C_t, K_{t+1}, e_{o,t}, e_{c,t}, e_{g,t}, R_{t+1}\}_{t=0}^{\infty}} \mathbb{E}_t \sum_{t=0}^{\infty} \beta^t \log(C_t)$$

subject to (25.9), (25.21), (25.23), and (25.25). In a more general version of this model, it is possible to define the marginal damage externality caused by energy source κ , Λ_t^κ , to be given by

$$\Lambda_t^\kappa = \mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \frac{U'(C_{t+j})}{U'(C_t)} \frac{\partial F_{t+j}}{\partial S_{t+j}} \frac{\partial S_{t+j}}{\partial e_{\kappa,t}}, \quad (25.30)$$

where U is a more general utility function and F is a more general production function for final goods. The intuition is clear and very robust: a unit emitted today gives $\frac{\partial S_{t+j}}{\partial e_{\kappa,t}}$ units of increased atmospheric carbon concentration j periods later, which in turn changes output

by $\frac{\partial F_{t+j}}{\partial S_{t+j}}$ per unit. Standard discounting and adding up of these future output effects then gives us the formula.

With the functional form of the production function that we have considered here, we have

$$\frac{\partial F_t}{\partial S_t} \frac{1}{Y_t} = \frac{\partial (e^{-\gamma S_t} Y_t)}{\partial S_t} \frac{1}{e^{-\gamma S_t} Y_t} = -\gamma. \quad (25.31)$$

The implication of the above equation is that every additional unit of carbon in the atmosphere yields a constant percentage reduction in GDP. In other words, the “marginal damage flow” is independent of both GDP and the CO₂ concentration. Moreover, the marginal utility of consumption, $U'(C_t)$, just becomes $1/C_t$ under logarithmic utility. Finally, the derivative $\frac{\partial S_{t+j}}{\partial e_{\kappa,t}}$ is straightforwardly computed from (25.9). Combining these expressions, we can now write the marginal damage, which according to the Pigou principle also is the optimal tax on carbon, in the following way.

$$\tau_t = Y_t \left[\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \gamma (1 - d_j) \right]. \quad (25.32)$$

The expression for the optimal tax rate is surprisingly simple. The marginal externality cost of emissions as a proportion of GDP inherits the time path of GDP. The tax rate then depends on three “d:s”. The first one is *depreciation* (d), i.e., how long the marginal unit of emitted carbon dioxide is staying in the atmosphere. Since a unit will generate more damages the longer it stays in the atmosphere, the higher the tax will be. The second d is the *damage* as measured by γ . Naturally, the tax is increasing in γ since this is a direct measure of the damages that are generated by the emission. The third d , finally, is the *discount factor*, β , i.e., the weight that is put on the future. Clearly, the tax is also higher if a higher weight is put on the future. What is remarkable about the formula is that it does not contain any other parameters, for example involving how fossil fuel is produced. It is, moreover, a reasonable benchmark case since logarithmic utility and Cobb-Douglas utility are standard assumptions in the macroeconomic literature; $\delta = 1$ is a decent approximation if the time period is long enough.

Simulations

We are now ready to show some output from the model. We are particularly interested in evaluating the effect on the global temperature of different carbon taxes. We consider three levels for the tax: 0, an optimal tax computed from (25.32), and a moderate carbon tax that we take to be one third of the optimal tax. To evaluate these policies, we first need to calibrate the model; this is carried out in the appendix and follows Golosov et al. (2014) and Hassler et al. (2021b). The model has many parameters but almost all of them can be set to match standard long-run economic facts or facts from the natural-science material presented above. The observed prices and shares of different energy sources can be used to select the energy-related parameters; the elasticity of substitution between different energy sources is hard to calibrate and we set it to 2, i.e., a significantly higher substitution elasticity than for Cobb-Douglas. We use Nordhaus’s benchmark damage estimate.

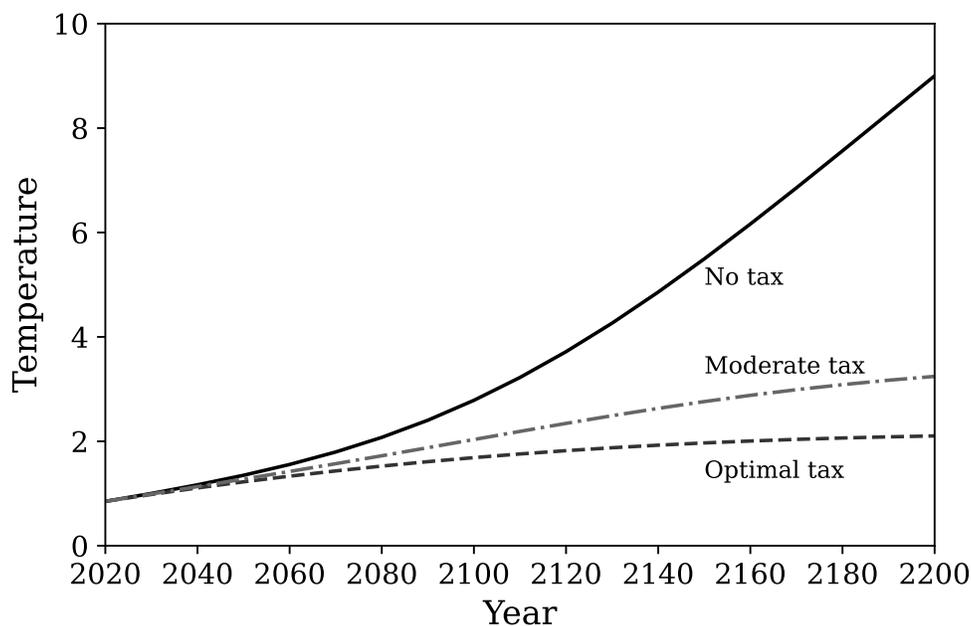


Figure 25.6: The effect on the global temperature of different carbon taxes.

Notes: The optimal tax is set to U.S. \$25 per ton CO₂ and the moderate tax is U.S. \$8 per ton CO₂.

Given the calibration, the optimal tax is about U.S. \$25 per ton CO₂ (or U.S. \$90 per ton carbon). This relatively low number is clearly sensitive to assumptions about the three “d”s in equation (25.32).²¹ The moderate tax is then only about U.S. \$8 per ton CO₂. The results associated with these different carbon taxes are presented in Figure 25.6. Note the very large difference in temperature increase between the laissez faire with no tax and the case with the optimal tax. Without any climate policy, the global temperature will have increased by as much as 9°C by the year 2200. In contrast, with a global optimal tax in place, the increase will be contained at about two degrees Celsius at the end of the considered period. This shows the strength and potency of the carbon tax as a policy for mitigating global warming. Note also that even a very low carbon tax (i.e., the moderate tax) can go a long way in making sure that the temperature increase is not too large.

25.5 Natural resources in finite supply

We now turn back to the broader issue of finite resources. When a resource is in finite supply, a basic issue is the normative one of how to manage the resource properly: at what rate should the resource be used over time? Another question is how markets, left to their own devices, will actually manage the resource. An associated question is how the resource price will evolve over time. From the perspective of climate change, it also becomes important to think about how the economy can manage without fossil fuels: can we economize on energy

²¹Again, see Golosov et al. (2014) for a detailed discussion.

use by altering the way we produce and consume? The goods and services that are very energy-intensive in production can perhaps be replaced by other products. Alternatively, we can explore other sources of energy. These issues will all be discussed in this chapter. We begin with a core theory that rests on Harald Hotelling’s seminal contribution from 1931.²²

25.5.1 Some data

Before we dive in to the theory, it is useful to look at some data. Figure 25.7 shows U.S. data on the uses of different renewable as well as non-renewable energy sources over a long time span. We see clear upward trends, while there are also (potential) peaks in use for oil and coal.

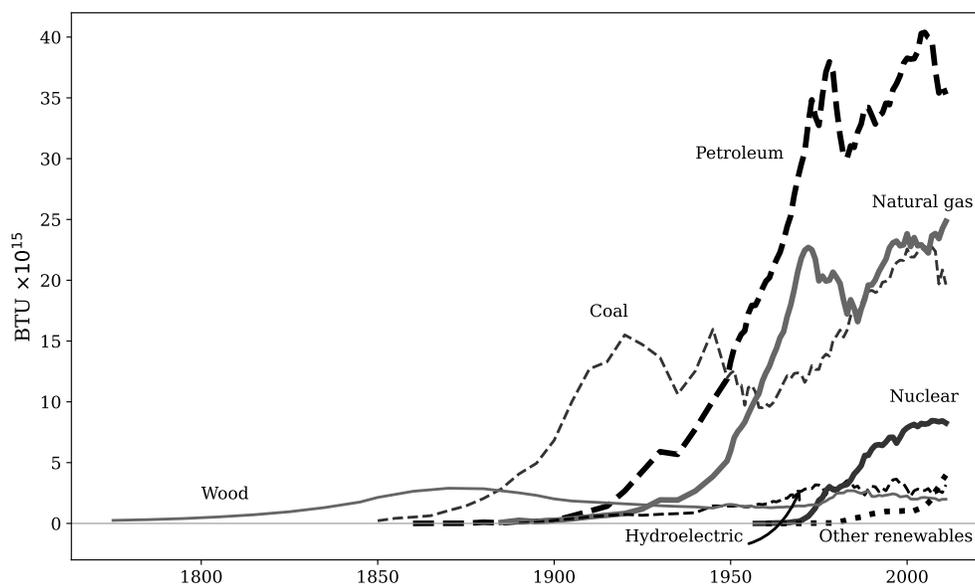


Figure 25.7: U.S. energy consumption.

Source: [Administration \(2012\)](#).

In per capita terms, the picture is different. For example, oil use peaked in the 1970s both in the U.S. and in the world as a whole and oil usage per capita has thus been decreasing since then.²³

Turning to prices, we first look at a long time series of the oil price as displayed in Figure 25.8. Maybe the most prominent feature of the time series of the price is the enormous volatility, a feature that oil shares with other natural resources. Over a long horizon, the price appear to have no trend. However, if we limit attention to the postwar period, there is potentially a positive growth rate of the oil price.

Figure 25.9 shows the prices of coal, lead, zinc, and copper. Coal’s price resembles that of oil, but the metals are somewhat different. They share the stationarity of a longer time horizon but they differ from the fossil fuels in that there is no visible upward trend in prices

²²Hotelling, born in 1895 in Fulda, Minnesota, taught Milton Friedman statistics in the early 1930s. Friedman had a number of PhD students at Chicago, one of them Neil Wallace, the PhD advisor of one of the authors of the present chapter.

²³See [Hassler, Krusell, and Olovsson \(2022\)](#).

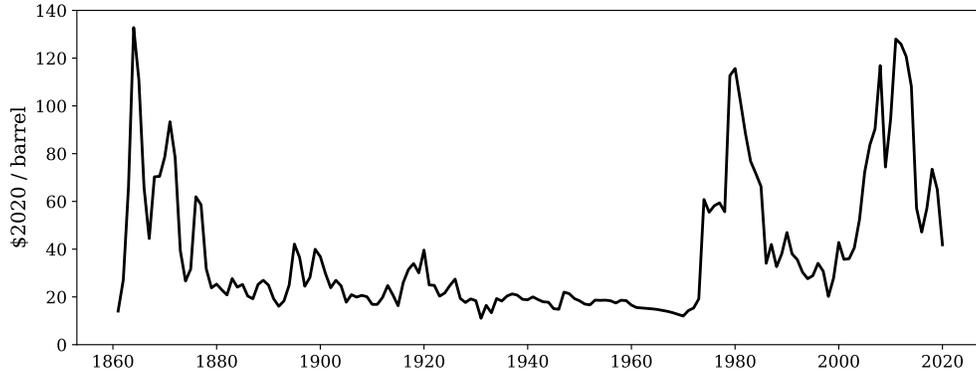


Figure 25.8: Crude oil price.

Source: [BP \(2021\)](#).

over the postwar period. A potential reason for this difference has to do with differences in extraction technologies across resources and over time. We will return to this point below.

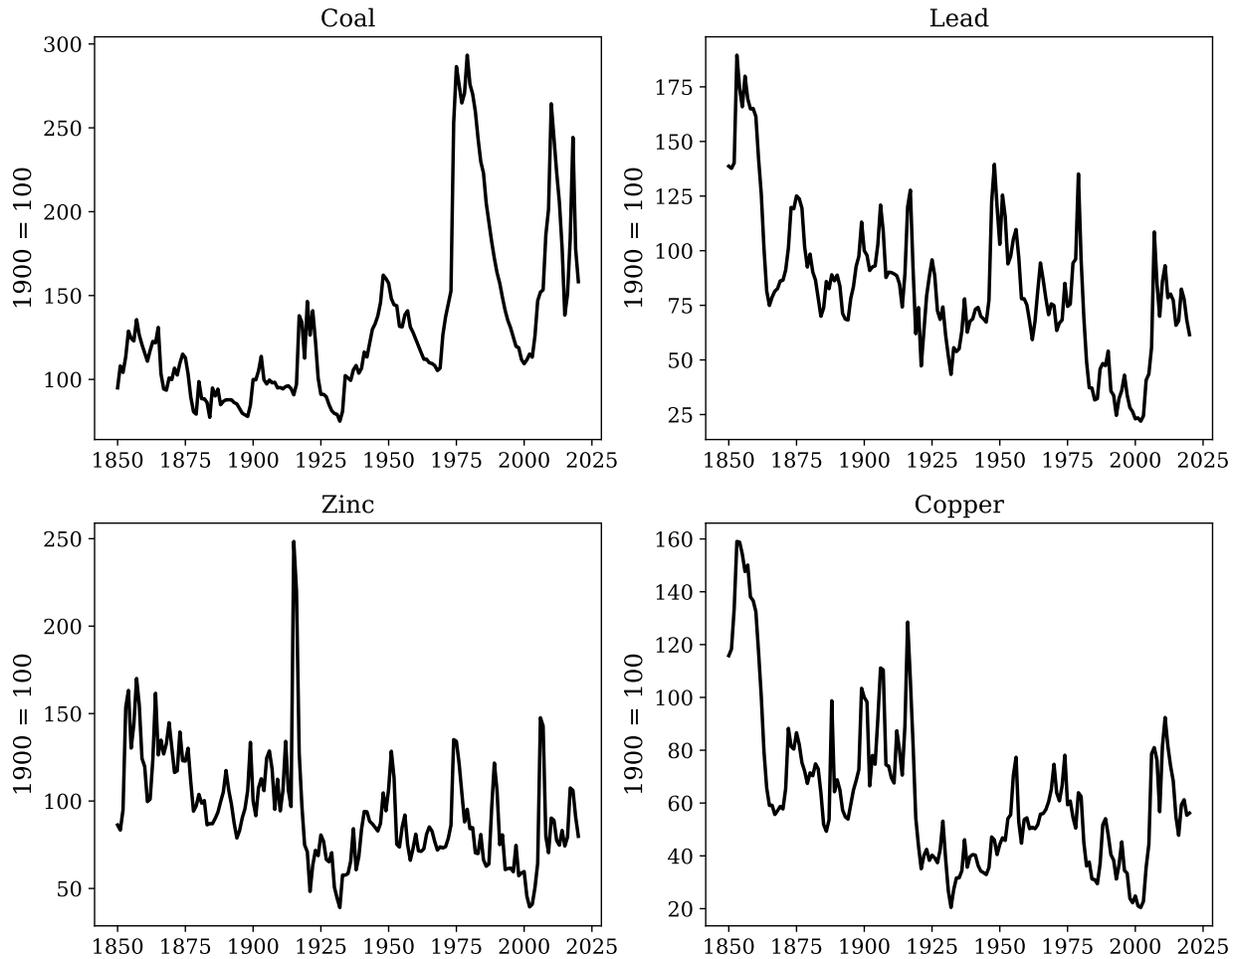


Figure 25.9: Inflation-adjusted prices for natural resources.

Source: [Jacks \(2019\)](#).

To summarize, we observe large swings in prices for natural resources, in fact on the order of magnitude of stock-price fluctuations. We do not observe marked trends, although for oil there is an upward trend over the postwar period. Quantities are rising significantly over time, fossil fuels included.²⁴ For fossil fuels, however, here we do see slowdowns, at least in per-capita consumption. It will turn out to be a challenge to account for these facts quantitatively.

25.5.2 Basic theory

The so-called “cake-eating” problem that analyzes how a finite resource should be depleted over time was first studied by Hotelling (1931) and later by Gale (1967). A more recent review of the empirical performance of Hotelling’s theory can be found in Livernois (2009). The implications of natural resource scarcity then received renewed interest after the first oil-price shock in 1973 when concerns that world oil supply would run out or simply become too costly to use. The *Review of Economic Studies* featured a special issue on the economics of exhaustible resources already in 1974. The issue contains important contributions by Dasgupta and Heal, whose model we will use below, Solow, Stiglitz, among others; these papers are all concerned with the question of how economic growth is affected by the presence of an input that is depleted over time. Since a theory of endogenous technical change had not been developed at this point, however, the important concept of endogenous technical change, discussed in Chapter 13 and revisited in the present chapter, is naturally absent from the analysis in these papers.²⁵

Pure cake eating: the planner Consider first the simplest possible cake eating problem: the cake’s size is R_0 , it can be stored over time and does not depreciate, and an infinitely-lived consumer with standard preferences likes cake. There is no other good in the economy. The planning problem thus reads

Consider a planning problem under zero extraction costs.

$$\max_{\{c_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \log c_t$$

subject to

$$\sum_{t=0}^{\infty} c_t = R_0,$$

with c denoting consumption of the resource. Notice that there is no cost of production: this is the sense in which this is a cake. We assume logarithmic consumption mostly for illustrative purposes; we do, however, want to insist on balanced-growth preferences since this chapter concerns the long run. Thus, we could generalize preference somewhat, but not much.

²⁴We omit data on quantities for lead, zinc, and copper; they display significant positive trends, too.

²⁵Later contributions that are also concerned with the growth implications of natural resource scarcity but allow for endogenous technical change include Barbier (1999), Scholz and Ziemes (1999), Smulders and de Nooij (2003), Grimaud and Rouge (2003), and Hassler, Krusell, and Olovsson (2021a).

The solution to the planning problem is $c_t = (1 - \beta)R_t$, where $R_{t+1} = R_t - c_t$, so that $c_t = (1 - \beta)\beta^t R_0$.²⁶ The model is extremely stylized but allows us to make some important points. First, the problem is well defined, despite the resource use going to zero and the marginal utility of consumption going to infinity: utility is bounded. Second, consumption is higher in the beginning than later, all because of discounting.

Pure cake eating: the market outcome and Hotelling rents A third point we can learn by studying this simple economy comes from studying the market allocation. Let us imagine a perfectly competitive world where a representative consumer owns (and sells) the resource. Thus, the consumer would sell an amount of the resource e_t at time t at market price \hat{p}_t^o (“o” for oil) in terms of the numéraire, which we take to be the resource at time 0, and also buy an amount c_t at t , subject to the Arrow-Debreu budget constraint

$$\sum_{t=0}^{\infty} \hat{p}_t^o c_t = \sum_{t=0}^{\infty} \hat{p}_t^o e_t,$$

with $R_t - R_{t+1} = e_t \geq 0$ for all t . Market clearing in the simple cake-eating economy is then simply $c_t = e_t$ for all t .

Taking and combining first-order conditions with respect to c_t at two adjacent points in time, we obtain

$$\frac{\hat{p}_{t+1}^o}{\hat{p}_t^o} = \beta \frac{c_t}{c_{t+1}}.$$

The first-order conditions with respect to e_t and e_{t+1} , when combined, become $\hat{p}_t^o = \hat{p}_{t+1}^o$ for all t : it is equally costly to sell the resource at different points in time—the marginal production cost is zero—and therefore the price has to be the same; otherwise the consumer would sell all of the resource at the date with the highest price. Given this characterization of prices, we obtain, from the earlier first-order condition, that $c_{t+1} = \beta c_t$ or, imposing market clearing, $e_{t+1} = \beta e_t$. That is, markets deliver the same outcome as does the planning allocation, not surprisingly.

How should we interpret the price implications, i.e., that the price of the resource at different points in time is the same? Recall that these are Arrow-Debreu prices, in terms of the numéraire (cake at 0). If we instead think about spot prices, what would they be? There is no other consumption good than cake here, so the notion of a spot price is not a natural one. Imagine, however, that we had another consumption good whose time- t Arrow-Debreu price in terms of the numéraire would be p_t . Then p_t/p_{t+1} would be the gross real interest rate and optimal cake selling for the owner would require that the resource price satisfy the same condition as before, where p_t^o now means the spot price, in terms of the other consumption good, i.e., $\hat{p}_t^o = p_t^o p_t = p_{t+1}^o p_{t+1} = \hat{p}_{t+1}^o$, i.e.,

$$p_t^o = \frac{p_{t+1}}{p_t} p_{t+1}^o.$$

That is, the spot price of the resource today equals the *discounted* value of the spot price tomorrow. This is because consumption of the resource at all times requires indifference between selling the good at different points in time. By implication, the spot price of the resource would have to grow at the real rate of interest. This is the *Hotelling rule* governing the prices of resources in finite supply, for the case of zero extraction costs. This is a remarkably robust insight. The fact that the interest rate is zero in this simple economy is not central; if the interest rate were positive, the Hotelling rule would simply mean the price is growing.

²⁶To obtain this amounts to a calculation that is standard for this text. Simply take first-order conditions with respect to c_t and c_{t+1} to see that, when the multiplier on the constraint is eliminated, $c_{t+1}/c_t = \beta$. Then substitute all the c_t s in terms of c_0 into the constraint.

So from the market analysis, we obtain this fourth result. A fifth result, which has been implicit in the discussion so far, is that $p_t^o > 0$ for all t : even though the cost of producing the resource is zero, the resource commands a positive price. This price, called the *Hotelling rent*, is a case where price exceeds marginal cost, even though there is perfect competition. If you earn a resource in limited supply, you earn rents from it, to the extent it is valued by consumers or firms.

A production economy Moving away from cake eating and toward a setting that views the natural resource as an input into production, consider planning a problem in the context of a growth model where the natural resource is used as an input: this is the [Dasgupta and Heal \(1974\)](#) setting. We assume that production is Cobb-Douglas in capital and energy, with decreasing returns to scale (labor is assumed fixed and would command the remainder of returns) and $\delta = 1$, i.e., capital depreciates fully from period to period. The resource, however, is modeled as before.

$$\max_{\{c_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \log c_t$$

subject to

$$c_t + k_{t+1} = Az_t k_t^\alpha e_t^\nu$$

and

$$\sum_{t=0}^{\infty} e_t = R.$$

Our assumption of a Cobb-Douglas production function in k and e can be motivated if the resource costs, as a fraction of output, have been stationary. We will display the data later but a summary is that, for energy (or fossil fuel), the cost share fluctuates sharply over time, though without a strong trend. Hence, for long-run analysis, Cobb-Douglas seems like a decent starting point. For reasons we will explain shortly, we will also go beyond a Cobb-Douglas function.

The solution is straightforward to derive, following our basic chapters. For resource extraction we obtain the same solution here as for the cake-eating problem: $e_t = (1 - \beta)R_t$, where $R_{t+1} = R_t - e_t$. Hence $e_t = (1 - \beta)\beta^t R_0$.

As part of the solution, we find $k_{t+1} = \alpha\beta Az_t k_t^\alpha e_t^\nu$ and that long-run gross capital growth g is constant if z grows at a constant rate: $g = \gamma_z g^\alpha \beta^\nu = (\gamma_z \beta^\nu)^{\frac{1}{1-\alpha}}$, where $\gamma_z - 1$ is the growth rate of z . For large enough γ_z , we see that $g > 1$. I.e., technology growth, of the form assumed here, allows net positive production growth.

As for price implications, from the Euler equation, we obtain $g = \beta(1 + r)$, so that

$$1 + r = (\gamma_z \beta^\nu)^{\frac{1}{1-\alpha}} / \beta.$$

Here, the net real interest rate is positive if $g > \beta$, i.e., if γ_z is large enough. The implications for the resource price, p^o , is $p_t = \nu y_t / e_t$, which grows at $g/\beta = 1 + r$. I.e., it rises at the real rate of interest. This result holds true time period by time period, as before. What the growth rate is of course depends on the value for the real interest rate; in the pure cake-eating case, the net real interest rate was zero at all times. Here, we can choose model

parameters so that it matches observations in the data. That would mean $1 + r > 1$, for at least most of the postwar period. Hence the resource price should rise exponentially, based on the simple setting here where there are no costs from extracting the resource.

Hotelling’s rule in partial equilibrium with extraction costs Taking interest rates as given, it is straightforward to extend the Hotelling formula to a case where there are resource extraction costs, i.e., costs that resource producers have to pay in terms of goods. The idea is the same as before: for extraction to occur in two consecutive time periods, we need to have indifference between producing today and producing later. I.e., we need to have

$$p_t^o - mc_t = \frac{p_{t+1}^o - mc_{t+1}}{1 + r_t},$$

where mc_t is the “marginal extraction cost” at time t . This marginal cost should be interpreted in a broad sense: it involves the input costs of extraction at t but also potential dynamic effects of extraction: if, for example, it is more and more expensive to extract the more you extract, the marginal cost of extracting a unit now must include the cost of making extraction harder later. The specific forms for the marginal cost that would materialize depend on the detailed specification of the extraction technology, but there will always be a formula of the kind stated here (with a broad notion of mc).

This implies

$$\frac{p_{t+1}}{p_t} = 1 + r_t + \frac{1}{p_t} (mc_{t+1} - (1 + r_t)mc_t).$$

That is, if the marginal cost is rising faster than the rate of interest, the price has to rise faster to compensate. If we had data on marginal costs, we could assess whether this formula holds in the data, given a rate of interest, at least on average. As already pointed out, however, marginal costs are rarely directly observable and here they also have dynamic components, to the extent extraction at a point in time affects the extraction costs at other dates. The “given interest rate” is also a challenge, since it is not clear which interest rate to use: a low one from a comparatively safe rate of return (say, on government bonds), or a higher one that reflects risk? Many observers claim that the Hotelling rule is violated also in the postwar data, with reference to the lack of a trend whose average real increase equals real interest rate on the higher side.

Another challenge is volatility: the Hotelling-based price formula derived solely from a production perspective means a sharp focus on supply. One can imagine that various hard-to-predict shocks on the demand side cause fluctuations, but the literature has very few full quantitative general-equilibrium treatment of the world supply and demand for natural resources. For the purposes of studying fossil fuel and climate change, which is a very long-run issue, abstracting from short-run volatility seems a reasonable way forward.

25.5.3 Capital-energy complementarity and technical change

We discussed the challenges involved in accounting for the price data on natural resources. Another challenge, implicit in our discussion of the simple models above, is the strong positive trend in the use of natural resources over time, including fossil fuel. We mentioned that if the marginal production costs fell fast over time, then an upward trend would be visible. There are two factors working in opposite directions: fossil fuel is getting more and more difficult to extract, while technical change has made any given type of extraction cheaper over time. Arguably, the latter force does not dominate, or at least not so much that it can account for the strong positive trend in fossil-fuel use since the Industrial Revolution.

Another possibility arises, however, if one considers an alternative production technology, whereby energy—at least in a short-run sense—is very highly complementary with capital and labor. We now look at this case.

Strong complementarity For simplicity, consider an extreme case, a Leontief technology, whereby

$$y_t = \min \{ Ak_t^\alpha, e_t \}.$$

Note here that A is not Hicksian but merely capital-augmenting. Also, to the extent at least one of the inputs is endogenously chosen, and costly to use, a Leontief function implies that $Ak_t^\alpha e_t$ must hold at all points in time. At time 0, this means that e_0 must be given by Ak_0^α , which is exogenous, k is a state variable. However, at time 1 and on, k is endogenous and costly: consumption earlier on needs to be forsaken in order to obtain it. What is the optimal path for capital and the resource? We will argue, without fully solving the model, that the outcome, based on a plausible parameter specification, is one where e initially grows over time and then falls.

So suppose that time 0 is at the beginning of the Industrial Revolution, where presumably k_0 is near zero, and A , interpreted as capital-specific technology, is also low. However, the energy problem had been “solved”: for the purposes at the time, energy was cheap and plentiful. That is, the initial endowment R_0 was high in relation to Ak_0^α . Then the optimal allocation would be to accumulate capital over time, i.e., $k_0 < k_1 < k_2 \dots$ at least for some time, so that $e_t = Ak_t^\alpha$ is increasing over time. Unlike in the standard neoclassical growth model, however, the economy would not asymptotically move toward a steady state, because e is in finite supply (R is high but not infinity). So eventually, e will have to decline, and k along with it. Fully working out a solution to this problem is somewhat involved but not necessary for the purposes here.

A more general case The example just studied makes clear that complementarity can generate an upward-sloping path for resource use. What would, however, a more systematic quantitative approach be like? A convenient generalization of the simple case displayed is a CES formulation:

$$y_t \equiv F(A_t k_t^\alpha l_t^{1-\alpha}, A_{et} e_t) = \left[(1 - \nu) (A_t k_t^\alpha l_t^{1-\alpha})^{\frac{\varepsilon-1}{\varepsilon}} + \nu (A_{et} e_t)^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}}.$$

Here, we have reintroduced (possibly time-varying) labor, making clear that overall production has constant returns to scale; we have also allowed two technology variables, one augmenting each input, that can also change over time. The parameter ε is the elasticity of substitution between the capital-labor composite and the resource. The Leontief formulation used above is a special case: $\varepsilon = 0$.

Solow’s growth-accounting method can be used to measure TFP changes over time, i.e., technological progress affecting overall production, under relatively weak assumptions.²⁷ However, how can *factor-specific* technology movements, such as those in A , and A_e be

²⁷Recall, from Chapter 2, that the key assumptions are that inputs, outputs, and prices can be measured, that the production function has constant returns to scale, that there is perfect competition, and that firms maximize profits.

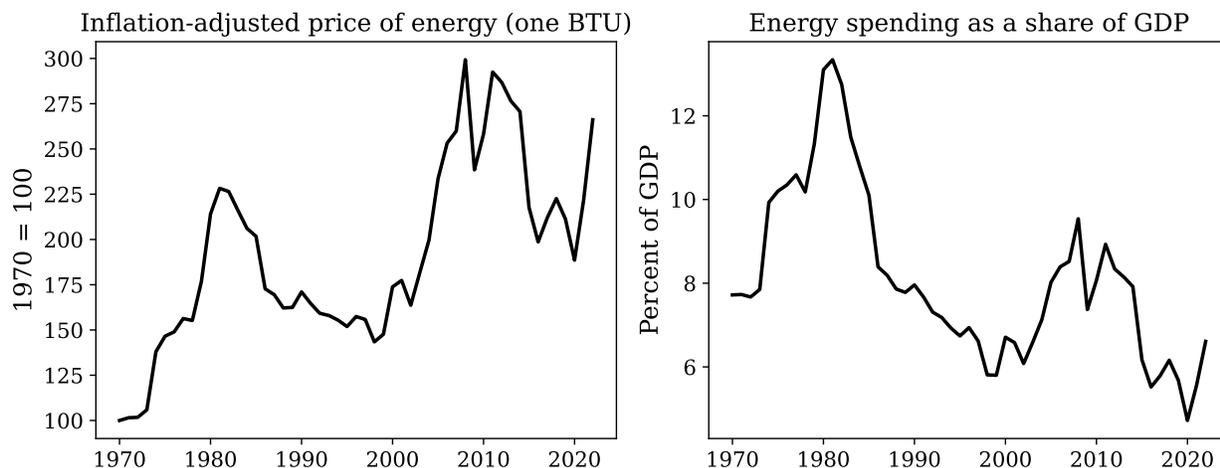


Figure 25.10: The price of fossil energy and its cost share.

measured? The short answer is that they can, under stronger assumptions on the production function. Suppose, for instance, that we knew that production was CES as just stated, and that we knew the parameters α , ν , and ε . Then the competitive firm's first-order conditions can be used to simply back out A_t and A_{et} period by period, given that all prices and quantities are observed. In this case, we obtain (after taking first-order conditions with respect to labor and the resource) and rearranging, that

$$A_t = \frac{y_t}{k_t^\alpha l_t^{1-\alpha}} \left(\frac{l_t^{share}}{1-\alpha} \right)^{\frac{\varepsilon}{\varepsilon-1}} \quad (25.33)$$

and

$$A_{et} = \frac{y_t}{e_t} (e_t^{share})^{\frac{\varepsilon}{\varepsilon-1}}, \quad (25.34)$$

where $l_t^{share} \equiv w_t l_t / y_t$ and $e_t^{share} = p_t^o e_t / y_t$ are cost shares.

An application to fossil fuel-based energy Hassler et al. (2021a) apply the ideas just discussed to the case of fossil fuel-based energy use in the United States. Relevant facts appear in Figure 25.10. Clearly, the price of fossil energy is extremely volatile over the medium run and the cost share of fossil energy appears to follow the price quite closely: when the price rises, so does the cost share. This means that a production function with ε close to zero fits the data very well; recall that a function with Cobb-Douglas or higher elasticity ($\varepsilon \geq 1$) would mean that when the price rises, the share falls (or, in the Cobb-Douglas case, does not move at all). In fact, even a Leontief function fits well. An α similar to what is used in the case without energy and a ν that matches average shares can then be added and, as a result, the movements of prices and shares match the data quite well. Given such a parameterization, it is then possible to use equations (25.33)–(25.34) to back out the factor-specific technologies. The result is depicted in Figure 25.11 together with the levels of y_t and e_t .

In Figure 25.11, e/l is (the normalized logarithm of) a fossil-fuel composite consisting of oil, coal, and natural gas per worker.²⁸ Note that energy use per worker increased up to the

²⁸The picture looks very similar if we instead plotted oil per person

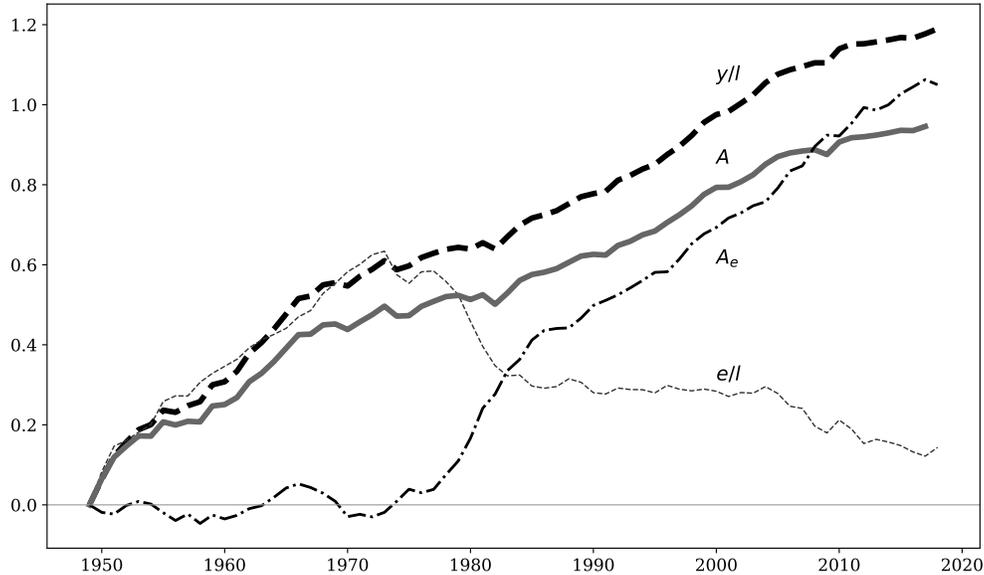


Figure 25.11: U.S. output, energy consumption and technological progress

early 1970s and it has been decreasing since. The variable y/l is (the normalized logarithm of) labor productivity. We see that e/l closely follows y/l between 1949 and 1973 and that, over this period, energy-saving technical change was effectively zero. Since 1973, however, A_e is increasing at a fast rate, while e/l has been falling. As a result, energy use in *efficiency units*, $A_e e/l$, has increased.²⁹ The reduction in physical energy use has thus been offset by faster energy-saving technical change.

The findings in Figure 25.11 strongly suggest that technology is endogenous: when prices of oil and gas shot up, energy-saving technological change started growing at a rapid rate. Moreover, technological change that augments capital and labor appears to have slowed down at the same time, suggesting that the endogeneity is *directed* in nature: away from the factors that are becoming relatively cheap toward those that are becoming relatively expensive.

Endogenous directed technology Chapter 13 developed models of endogenous technical change. An aspect that was not covered there is the multi-dimensionality of technology (such as the (A, A_e) here) and how inputs into research can be allocated based on relative profitability. In the context of Romer (1990) product variety model, A could thus be thought of as the variety range of intermediate inputs specifically used to save on capital and labor and whereas A_e would be the variety range of intermediate inputs directed toward energy-saving. Or, using Aghion and Howitt's (Aghion and Howitt, 1999) model of creative destruction, A and A_e could be the average efficiency levels of attained in (endogenous) productivity enhancements of a fixed set of intermediate inputs that are capital/labor- and energy-saving, respectively. In each of these applications, one can imagine that there is a fixed set of R&D workers that can be allocated to either of the two innovative (and patent-seeking) sectors and that market forces would dictate where the R&D workers locate each period. A richer

²⁹In fact, fossil-fuel use in efficiency units closely follows output.

model would have the total amount of R& D workers be endogenous as well.

The technology menu approach In [Hassler et al. \(2021a\)](#), a simpler setting is used where a technology *menu* G is postulated that lists different options for each firm to choose freely among. In particular, consider

$$G\left(\frac{A_{t+1}}{A_t}, \frac{A_{e,t+1}}{A_{e,t}}\right) = 0, \quad (25.35)$$

where G is strictly increasing in both arguments. Then the idea is that a given, competitive firm can choose any $(A_{t+1}, A_{e,t+1})$ at time t satisfying equation (25.35) at no cost, treating $(A_t, A_{e,t})$ as given: these are the economy-wide average chosen in the previous period. I.e., all firms choose what respective technology saving they prefer next period but do not internalize the dynamic technology impact of their choice. Under some conditions on G (jointly with assumption on the firm's production function), the firm's profit maximization problem—which now involves choosing both inputs and the input-saving technology levels—is well defined and consistent with perfect competition. A simple, static example of this procedure is as follows: suppose (i) that the production function is CES in inputs x and y , i.e., $f(A_x x, A_y y)$, and that the elasticity of substitution between x and y is less than 1 (more complementary than Cobb-Douglas); and (ii) that $G(A_x, A_y)$ is a CES minus a positive constant, so that the non-linear (CES) average of A_x and A_y equals this constant. Then it is straightforward (though tedious) to show that the technology choice will lead to a “reduced-form” production function in x and y that is a CES in x and y with a higher elasticity of substitution than in f .^a That is, when (A_x, A_y) is given and not subject to choice, x and y are harder to substitute than when they can be chosen. This captures the intuitive notion that the substitutability between inputs is higher in the long than in the short run. For example, fossil fuel may be very costly to replace as energy source in the short run but much less so in the longer run, since production techniques can then be developed that work better with other energy sources.

^aThe original paper working these results out is [Leon-Ledesma and Satchi \(2019\)](#).

Whichever way endogenous technology is added to the model, we can capture that natural notion that low short-run substitutability between energy and other input factors turns into higher substitutability in the longer run, when technology moves in such a way as to save on expensive inputs. Thus, a sharp short-run rise in the energy share in response to an increase in the price of energy will then be falling over time, even if the energy price remains high.

25.5.4 Taking stock

To conclude this section, we have presented some basic facts about natural resources in finite supply and some basic theory to begin accounting for the facts. We have focused on resources that, in contrast to the climate, have property rights associated to them, thus bracketing a broad class of environmental goods and services. The data we looked at showed increasing paths of use, presumably going hand in hand with economic growth, but also a slowdown of use toward the end of the sample. This overall pattern, we argued, seems at least qualitatively consistent with the basic theory of resource depletion that we presented, so long as we incorporate time-varying, input-specific technology levels. The data also showed

wildly fluctuating prices for natural resources, with weak time trends. The price data still remains a challenge to fully explain.

The basic theories we presented do not call for government intervention, at least not in an obvious way. The market outcome is optimal in the basic model, and although extensions involving endogenous directed technical change would suggest that some externalities to R&D be internalized with subsidies/taxes, it is not immediately obvious whether such subsidies should be directed or imply general to all kinds of R&D. Of course, the efficiency of the market outcome presumes that property owners and market participants are rational and that there are no major frictions in the relevant markets. The purpose of this section, however, was not to make the point that sustainability concerns generally are dealt with perfectly by markets; rather, it was to show how markets under ideal circumstances, given property rights, can be functioning well.