Environmental Macroeconomics: the Case of Climate Change

John Hassler*and Per Krusell†

August 13, 2018

1 Introduction

Climate change is a truly macroeconomic phenomenon: it is global. Hence, to the extent it has important economic determinants and effects, it should almost by definition be covered as a part of the core macroeconomic curriculum. The analysis of climate change also overlaps in important ways with other issues traditionally covered by research in macroeconomics: growth and development, technological change, and globalization (in trade in goods and financial markets, knowledge spillovers, etc.). As a result of these connections, a number of macroeconomics-oriented researchers have recently taken an active interest in the field of climate change. Consequently, there is now a set of contributions that one might represent as the "macroeconomics and climate" literature. The aim of this chapter is to provide a compact introduction to this literature and discuss what we perceive as its main value, as well as its challenges. We pay particular attention to what we think modern macroeconomic methods can contribute in the area. The goal is thus not to survey the growing literature in the intersection of climate and macroeconomics but rather to provide background, motivation, and a methodological discussion for those potentially interested in the area.[1]

Macroeconomic research has a tradition of empirically oriented theory building that, we believe, comes in handy for the purpose of understanding climate change. We want to emphasize four important aspects here.

First, macroeconomic analysis focuses on general-equilibrium (GE) modeling, i.e., theory (i) building on microeconomic foundations and (ii) with aggregate perspective. This analysis is often mainly used as a positive tool, i.e., as a way of trying to account for

---

*Hassler: Institute for International Economic Studies (IIES) and CEPR.

†Krusell: IIES, CEPR, and NBER.

[1] Many highly relevant references will be left out in this text. This should not be seen as a reflection on our view on the relative contributions of the left-out papers but rather as an expression of the narrow aim herein.

the main aggregate facts (whether defined as short- or long-run movements). The GE approach is, in fact, a hallmark of "modern macroeconomics", to be distinguished from the traditional business-cycle modeling taught in undergraduate textbooks. With the GE structures, it is however also natural to conduct welfare analysis according to standard microeconomic principles, as in standard public-economics treatments. Therefore, these GE structures straightforwardly offer tools to evaluate policy proposals. Thus, it is natural to use these structures not just to find the optimal policy (using a planning model) but also to evaluate and compare different sub-optimal policies. In dynamic models, this can be quite involved, and climate change is an area where dynamic modeling is key.

Second, and perhaps most importantly, modern macroeconomic models are by their very nature quantitative, i.e., they do not necessarily develop new concepts or ideas but can rather be classified as applications of existing theory to address empirical issues, with parameter choices made so as to match historical data/replicate key patterns. In the context of long-run issues, one would then insist on using frameworks that can account for the main long-run facts. These include—for the United States and most rich economies—a rather balanced behavior for the macroeconomic aggregates, e.g., a stable rate of output growth per capita at about 1.5%, a stable capital-output ratio of around 3 on an annual basis, and so on. Thus, at its core, modern macroeconomics is about making statements about numbers, which is also key in the area of climate change. Certainly natural scientists engaged in the study of climate change have the same quantitative focus: they primarily apply known theoretical insights, upon which they then base measurement, causal interpretations, and forecasts. Thus, it would be to break with this tradition if the economics approach to climate change did not have a similar quantitative orientation. In their quantitative analysis, macroeconomists have thus relied on a variety of structural and other econometric tools.

Third, modern macroeconomic GE models often incorporate aspects that are thought to be relevant in the area of climate change. For one, they are typically of the dynamic, stochastic kind—they are DSGE (dynamic, stochastic general-equilibrium) models—allowing the analysis of uncertainty and risks in dynamic environments. Dynamics and uncertainty are clearly also believed to be important in the climate-change area. Many macroeconomic models nowadays also explicitly consider technical change to be endogenous, and in particular, "directed": what particular research advances are aimed for is endogenous and depends on incentives and, hence, on economic policy. Finally, many macroeconomic models are fundamentally non-linear and such features are often argued to be important—and certainly need to be examined—in the climate context.

Fourth, to solve and use quantitative DSGE models, macroeconomists have developed special computational tools. These are necessary when going beyond the simplest possible model, for a variety of robustness checks, etc., and they are not only relevant for quantitative evaluations but also often important for gaining analytical insights into mechanisms. For natural scientists, who dominate the area of climate-change research, numerical solutions are not just ubiquitous but fully accepted; they barely use stylized, theory-oriented models. Thus, whereas perhaps many theoretically minded economists are reluctant to use computational tools, the nexus area of climate change and economics is one where there does not appear to be much of an alternative.

Of these four points, all of them in principle share the methodology used in the CGE (computable general-equilibrium) literature. In many ways, the DSGE approach is simply very close to the CGE approach. The dynamic aspects of the models here, and in DSGE models in macroeconomics more generally, set them apart from CGE models and the motivation is clearly that dynamics are in focus. Dynamics can be thought of as just another dimension in a general-equilibrium model, but the typical infinite (or very long) time horizon adds computational challenges, as do time-related frictions that are often introduced, such as constraints on borrowing, adjustment costs to investment, or consumption habits (none of which are described in the benchmark models here). The presence of uncertainty is another difference, though it too is not fundamental, as uncertainty, like time, can also be thought of as just another dimension of heterogeneity in a general-equilibrium model. Typically, CGE models are much more detailed in terms of heterogeneity (goods/sectors, countries, etc.) and the DSGE models of climate are only recently moving in this direction, mainly due to the computational challenges arising from the need to deal with dynamics and uncertainty. One vision that we have is that the CGE and DSGE literature will move closer to one another and perhaps even merge into one and computational power and experience in working with numerically solved economic models continue growing.

The structure of this chapter is to describe in some detail how the macroeconomic models of climate change are formulated, motivated, parameterized, and solved. The end result is a so-called Integrated Assessment Model, IAM. The integration amounts to (i) our climate specifically being driven by human economic activity explicitly described and derived from microeconomic foundations; and (ii) our climate affecting economic activity, through various "damages" appearing in the microeconomic structure, i.e., the integration is a two-way feedback system, unlike most (but of course not all) frameworks used in this area. To this end, we need to specialize somewhat, if nothing else for the sake of efficient illustration of our main points. The core contribution upon which we build our structure is Nordhaus's DICE/RICE model framework, a setting that has many (but not all) of the ingredients we discuss.[2] Nordhaus's framework has a fundamental quantitative-theory basis. In particular, it builds on the neoclassical optimal growth model—in particular on the extension of it that incorporates a finite resource (à la Dasgupta and Heal, 1974)—that, as Solow (1956) argued and we shall also elaborate on below, offers a reasonable match to long-run economic data. It then adds two natural-science modules: a carbon-cycle model and a climate model. In order to allow nonlinearity and forward-looking, Nordhaus needed to simplify these models, without sacrificing too much of their empirical accuracy, so as to make the full integrated setup computationally tractable. Nordhaus's setting is not formulated as a general-equilibrium model, though it is consistent with one under additional assumptions. Having an explicit equilibrium setting is not only useful for communication but also because some of the most

---

[2] See, e.g., Nordhaus (1977) and Nordhaus and Boyer (2000).

pressing needs are, precisely, to compare realistic—but far from optimal—policy alternatives within a market economy. The basic DICE and RICE models also do not feature endogenous technical change or uncertainty (at least not of a very general kind), the latter because the computational methods used to solve these models are not recursive, a point we will also elaborate on below.

Since the treatment of climate change here builds on the typical macroeconomic setting used in analyses of long-run as well as short-run issues, much of the chapter will read as a primer on macroeconomic methods for environmental economists—a how-to of sorts. The end goal is to explain, motivate, and make ready for use the methods than underlie modern macroeconomic IAMs. In particular, this means going through a rather long list of stylized macroeconomic facts, since it is these facts that motivate the specific theoretical structure used. This structure is sometimes, in our experience, met with some skepticism when presented to broad audiences. The purpose here is to explain that the structure is used to account for the stylized facts and that it is hard, if not impossible, to come up with an alternative structure that can also match the facts—at least if one insists on microeconomic foundations. We mostly focus on the United States, though the focus on climate change and global macroeconomic modeling should dictate looking at global facts; the reason is data availability and quality. We add to the typical macroeconomic facts a short discussion of energy, since energy plays an important role in the climate-economy area.

As a final remark before proceeding, we need to point out that many of the estimates of the costs of climate change coming out of integrated assessment models of the sort considered here can be thought of as strikingly low, at least to the extent one feels climate change is one of the most urgent global issues facing mankind. The quantitative midpoint estimates of the the flow damage from warming under a business-as-usual scenario of the order of one percent of global GDP today, perhaps three percent in 100 years, and eleven percent in 200 years. Optimal policy according to the model, moreover, would eliminate a large part of these costs. Are these numbers small? Are they unreasonably small? First, as compared to other macroeconomic "stakes", they are large. The costs of business-cycle fluctuations are one or more orders of magnitude smaller and the business cycle is arguably much more difficult to tame (and it is not even clear it should be tamed); the same goes for the costs of inflation. As for influencing the economy's rate of growth, the stakes are obviously higher but it is also not so clear how easy it is to affect long-run growth. In the area of climate change, it is very clear what to do, however. Even though there are different views on the merits of taxes vs. quotas, these are minor disagreements; there is a consensus that either of these instruments were used (sufficiently), the climate problem would be handled optimally: climate change would then not be set to zero but it would be drastically reduced and, most importantly, the benefits of climate policy would be large and on the margin, after an optimal policy is implemented, would balance its costs. Thus, climate change is a highly relevant policy issue from a macroeconomic perspective. Of course, for any given country, it is not, because any given country cannot by itself more than marginally influence the global climate; climate change is a concern that is global in nature. Also, many might not consider climate change urgent, since the climate moves so slowly—just like concerns about long-run growth often tend to play second fiddle to discussions about an ongoing recession, and so on. It seems, however, that climate change is better able to attract attention than are concerns over long-run growth.

Second, many would argue that the benchmark cost estimates in the minds of models presented here are too low. That may be, but the point of this literature is rather a methodological one: to construct a setup identifying key inputs that all matter in the calculations of the costs of climate change are of the benefits of optimal policy. Thus, the setting is constructed so that it can generate much larger numbers, and it is up to further research to refine the basic parameters of importance. Key parameters involve the economic damages from climate change and the climate sensitivity; these are currently set to values similar to those used in, say, IPCC contexts but can easily be turned up and down for robustness checks. Another concern may be that the benchmark model used here does not have a global tipping point (nor highly non-linear dynamics). The reason for this is that there is no quantitative evidence (at least not one where there is close to a consensus) indicating a tipping point at some degree of additional warming (or at some atmospheric carbon concentration), nor is there consensus that such a global tipping point exists. But it straightforward also in this case to amend then model herein to include such a possibility.

The overall chapter is thus organized as follows. In Section 2 we first describe the core theoretical framework we use and, especially, discuss its empirical underpinnings. The first part here is the empirical background just mentioned and then the purely entirely economic model is presented, first without and then with an energy sector. This model is described and discussed in detail. The integrated climate-economy model then adds two simple and highly aggregated natural-science modules: one model describing the climate system and one representing the carbon cycle. These are presented in Section 3. The complete integrated-assessment model also requires a description of the damages from climate change; this topic is briefly discussed in Section 4. The full setup is described in Section 5, with a particular emphasis on the (suboptimal) competitive equilibrium. In this section we also briefly discuss the concept of the marginal cost of carbon. To analyze the full model, one needs to assign specific values for all the parameters and then solve the model, usually using numerical methods. We discuss parameter selection throughout the text (first for the core economic model and then whenever new elements are brought in) and computation in Section 5.3. We also discuss an *analytical IAM*, i.e., a model with all the key climate-economy components that makes some drastic simplification so as to allow analytical tractability, that we argue is a quantitative reasonable setting, and hence very useful, despite the drastic simplifications. This discussion is contained in Section 5.5. The topic endogenous technical change is discussed in Section 6.1 and multi-region models in Section 6.2. Section 7 concludes with some discussions about future challenges for the literature.

2 The neoclassical growth model: why and how?

The basic idea, as mentioned in the introduction, is that macroeconomic analysis, i.e., the analysis of our main aggregates such as output, investment, consumption, employment, and so on, strives to build a framework that is qualitatively and quantitatively capable of reproducing the main historical facts for these variables. This ambition is nowadays applied also in the context of understanding the short-run movements of output but the original contributions along these lines are due to Solow: his neoclassical growth model precisely constitutes a setting that can account for the surprising stability over a long period of time in key aggregate statistics such as the capital-output ratio, the rate of output growth, and so on. Solow's main insight was that a short list of assumptions that seem like reasonable approximations of how the market economy works will deliver *convergence* to stable values of the variables in question—a wholly nontrivial result. Solow's framework can in some dimensions be viewed as a reduced form: it hardwired two behavioral rules. One of them was a constant and exogenous rate of gross saving out of output and the other constancy of hours worked. The later literature therefore naturally looked at how these assumptions could be microeconomically founded, in particular with reference to rational choice, given utility functions in a certain class. A bit later still, another assumption Solow made—that of exogenous technical progress occurring at a constant rate—was also derived as an outcome of more basic assumptions: endogenous technical change in the market economy. We will review these elements when setting up our general model. The motivation here, however, is key: the goal of the model construction is to be consistent with a set of regularities or "stylized facts".

The quantitative macroeconomic approach to the economics of climate change also involves the role of fossil fuel, and in particular energy, in the economy. We therefore add a short discussion of energy facts to the usual macroeconomic stylized facts. We return to these questions later, as we will need them when building the extension to the theory that involves energy and energy prices. In Section 2.1, thus, we first review all the relevant facts and in Section 2.2 the task is then to build a framework that can account for these facts quantitatively: our *quantitative theory*. This theory will then later be the foundation upon our quantitative-theory approach to economics and climate will be built.

2.1 Empirical underpinnings: long-run facts

The facts we are about to review—and that in several ways echo the so-called "Kaldor facts"—involve the main macroeconomic aggregates in a way consistent with our national accounts and, furthermore, facts on input prices and interest rates, thus capturing key features of the market economy.[3] For the purpose of the climate-economy connection, we also include a short discussion of energy and the key associated facts. Our focus will be on the United States, though it has to be kept in mind—and we will return to this important point below—that the macroeconomic modeling required in the climate con-

---

[3] See Kaldor (1957).

text needs a global focus. In fact, a systematic evaluation of the Kaldor facts for a broad cross-section of countries has (to our knowledge) not yet been conducted. The U.S. focus here is motivated by (i) ease of access to statistical data whose quality is well known and considered high; and (ii) the period covered involves a relatively stable period, except for a few well-known events, such as the Great Depression and WWII.

We begin by a picture of U.S. per-capita GDP growth: Figure 2.1.

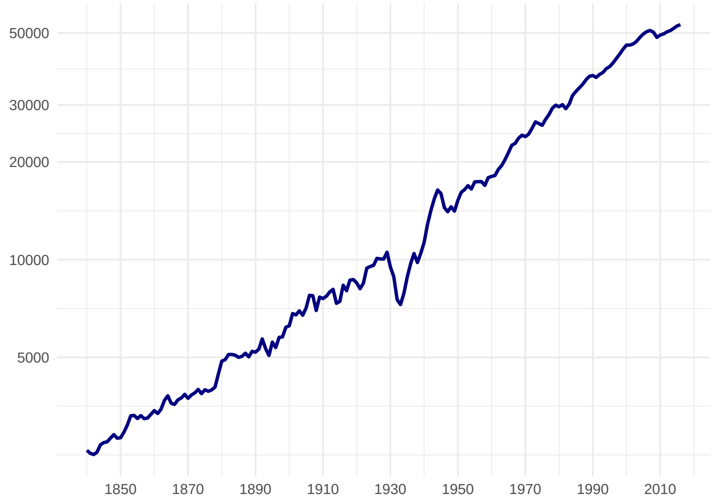

Figure 2.1: **GDP per capita in the U.S.**

U.S. GDP per capita 1840–2010, 1990 prices. Implied average annual growth rate: 1.8 percent. **Source:** the Maddison project.

What stands out is the remarkable stability of the growth path: output is well approximated by a log-linear growth path, with the main significant departures occurring the Great Depression and WWII. The recent Great Recession is barely visible.

Figure 2.2 shows the capital-output ratio, where output is measured in annual terms and capital is the sum of productive (fixed) capital and consumer durables. Clearly, the ratio hovers around 3, after sharp initial swings during the Great Depression and WWII.

Figure 2.3 goes much further back in time and uses a different, and broader, capital measure; again, the ratio hovers around a fixed value (roughly 4).

Figure 2.4 shows the consumption-output ratio. It is similarly quite stable, thought with a slight upward trend over the last fifty or so years. The figure implies—given that net exports are small in the U.S.—that the ratio of gross investment to output has been fairly stable too, but with a slight upward trend.

Hours worked in the U.S. have not displayed a trend in the postwar period, as Figure 2.5 shows.

However, looking at other countries over the same time period, we observe that the typical pattern is a downward trend (at a little below 0.5 percent per year): see Figure 2.6.

Looking over a longer period of time, we again see a clear downward trend (notice that the y axis is logarithmic, and hence a straight line means a constant rate of hours

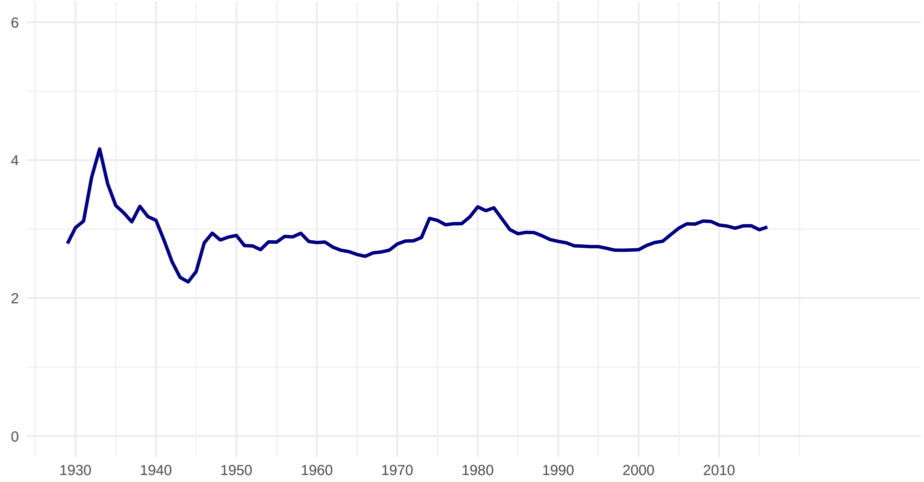

Figure 2.2: **Capital-output ratio in the U.S.**
Ratio of fixed capital and consumer durables to GDP. **Source:** NIPA table 1.1.

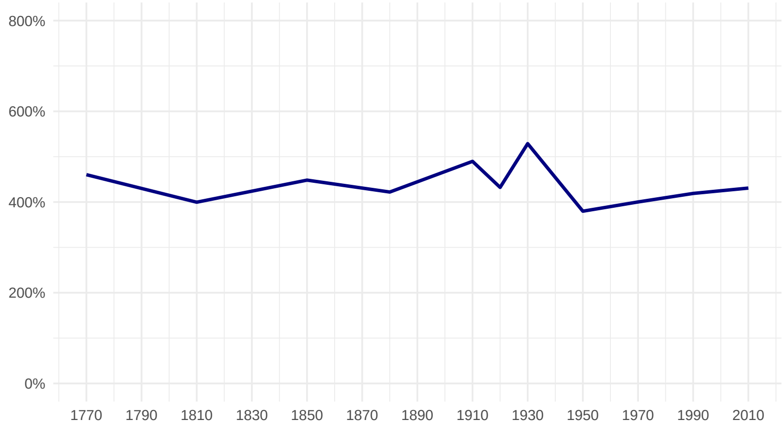

Figure 2.3: **Wealth-output ratio in the U.S., historical data**
Source: Piketty (2013).

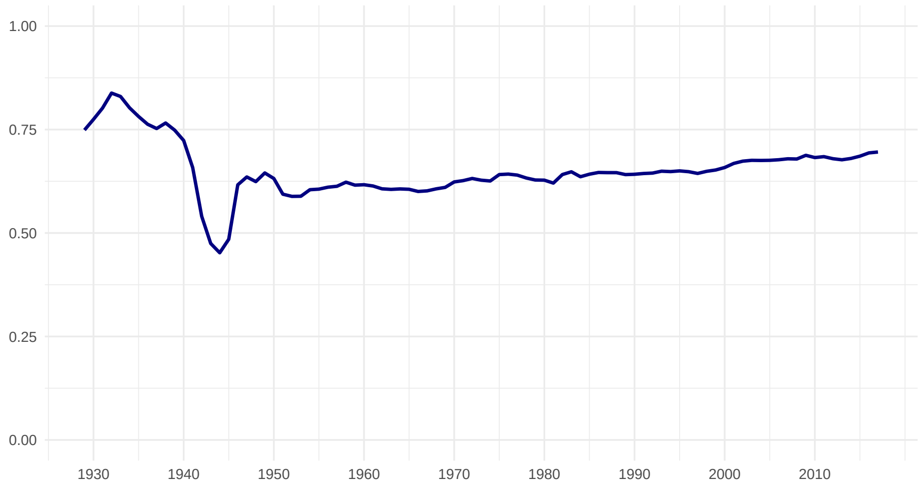

Figure 2.4: **Consumption-output ratio in the U.S.**
The ratio of private consumption to GDP 1929–2017. **Source:** NIPA table 1.1.5.

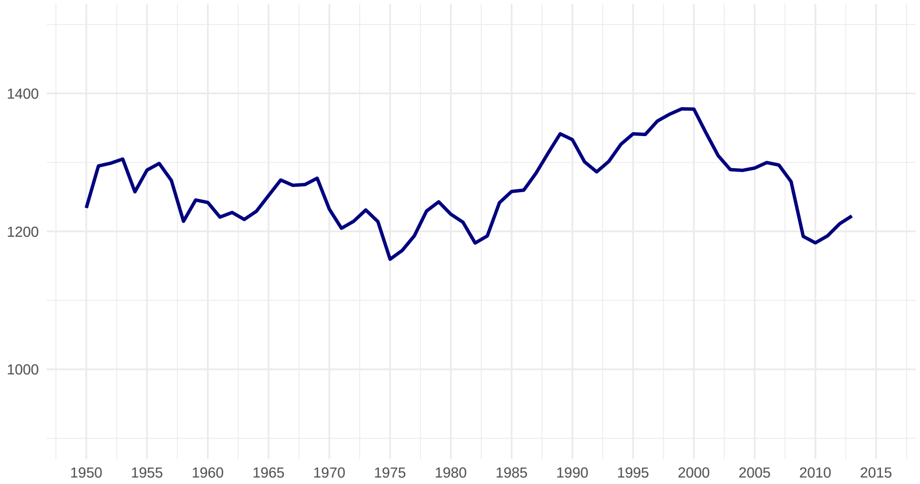

Figure 2.5: **U.S. hours worked, postwar period**
Average annual hours worked. **Source:** Boppart and Krusell (2018).

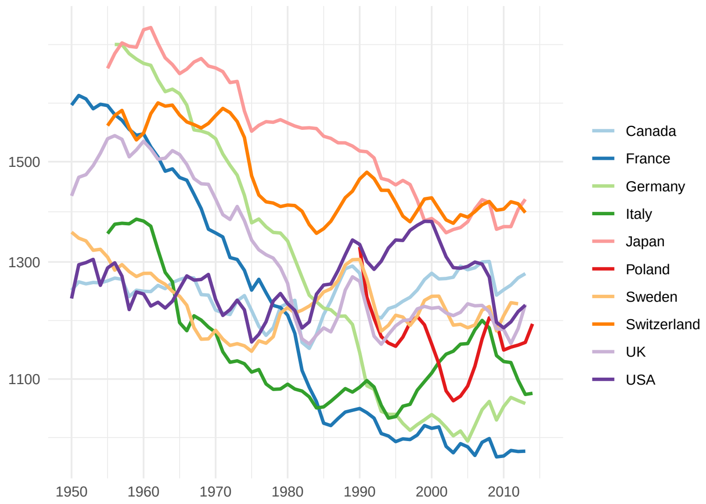

Figure 2.6: **Hours worked in the OECD, postwar period**
Average annual hours worked. **Source:** Boppart and Krusell (2018).

decline): Figure 2.7 makes this point, and here the U.S. is included, i.e., the downward trend at a constant rate is visible in the U.S. too going back further in time.

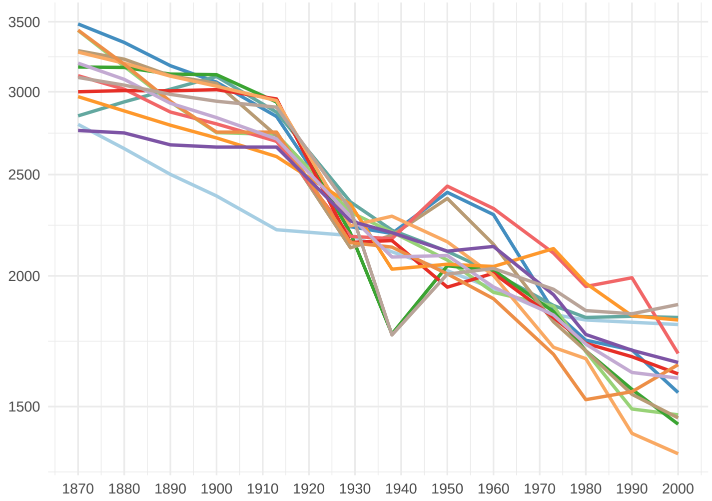

Figure 2.7: **Hours worked for a broader set of countries**
Average annual hours worked. **Source:** Boppart and Krusell (2018).

Another variable of interest is productivity. We will look at labor productivity here as well as TFP (total-factor productivity, computed from the Solow residual). Figure 2.8 shows (smoothed) growth in labor productivity in the U.S., the Euro area, Japan, and the

U.K. since the end of the 19th century. Though we see significant fluctuations, particularly for Japan and the U.K., the paths hover around 2% from a long-run perspective.

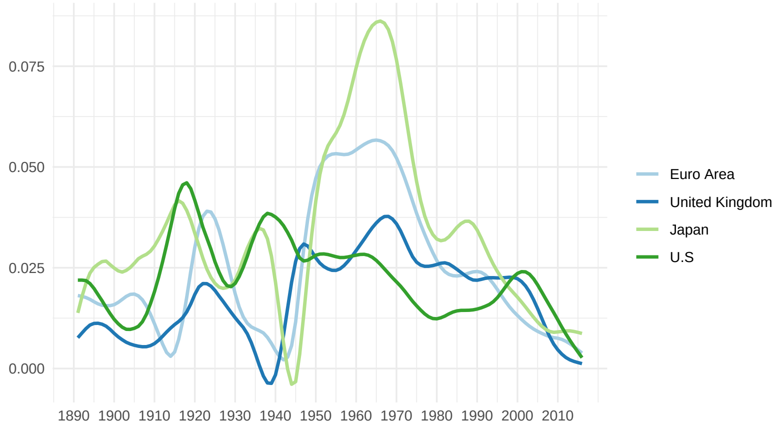

Figure 2.8: **Labor productivity for a selection of countries**
Hodrick-Prescott-filtered annual growth of labor productivity per hour U.S., the Euro Area, Japan, and the United Kingdom, 1891–2012. Source: Bergeaud, A., Cette, G. and Lecat, R. (2016).

Robert Gordon's recent calculations illustrate the slow movements around the trend in U.S. data. Figure 2.9 thus shows the slowdown periods in U.S. data: the beginning of the 1970s and fifteen years on, and the last ten or so years.

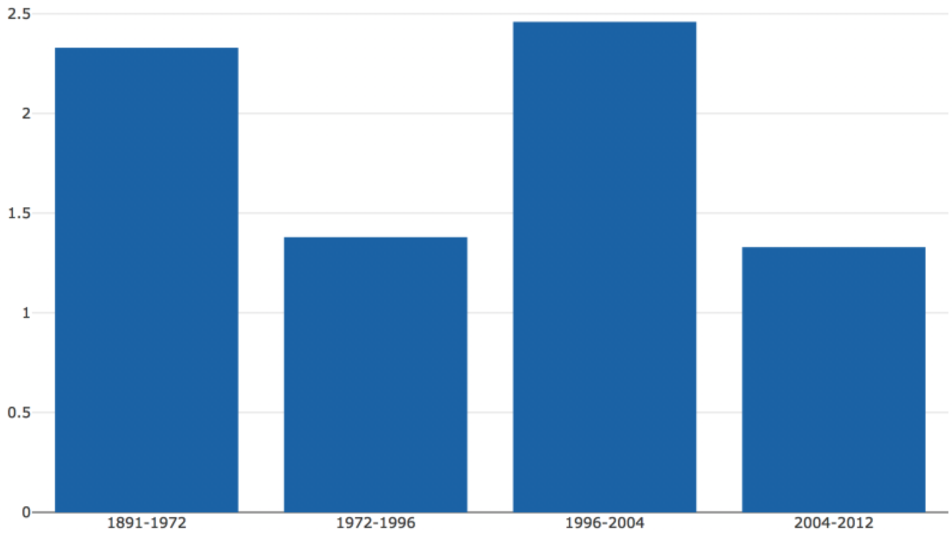

Figure 2.9: **Labor productivity in the U.S., sub-periods**
Source: Robert Gordon (2012).

Turning to TFP, we see its growth rates in the U.S. over the postwar period computed from the Penn World Tables in Figure 2.10, where we also see a more recent calculation of labor-augmenting technology growth based on BLS data available only for the last part of the period. We again see a fairly stable pattern.

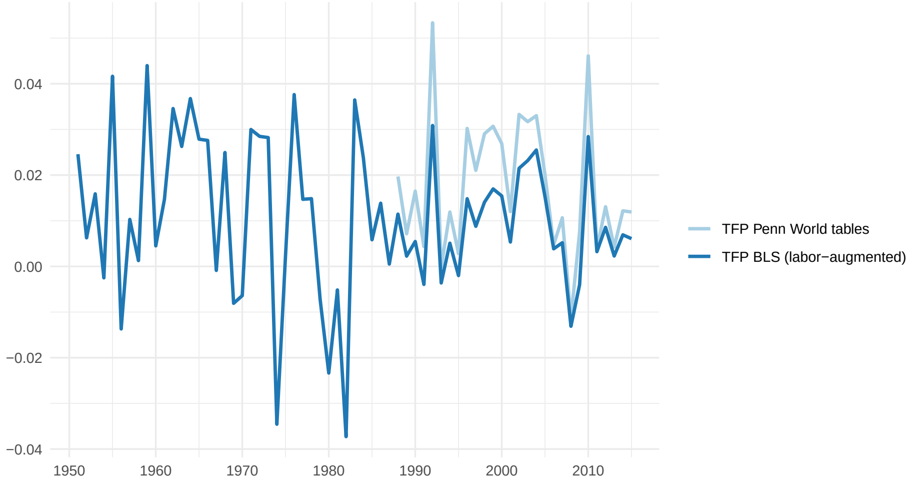

Figure 2.10: **TFP in the U.S., two measures**
Source: BLS.

Figure 2.11 shows TFP growth rates for the same regions as covered in the labor productivity figure above, with similar conclusions.

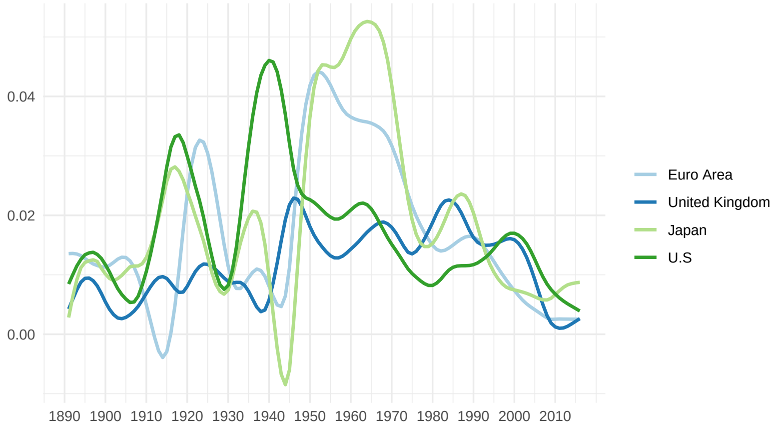

Figure 2.11: **Historical TFP for a broader set of countries**
Hodrick-Prescott-filtered annual growth of total-factor productivity in the U.S., the Euro Area, Japan, and the United Kingdom, 1891–2012. **Source:** BLS.

The above facts summarize the main data on quantities, including productivity. Turning to data involving prices, let us first look at real wages. Figure 2.12 show real wages since the early part of the 19th century. We see that real wages rise at a high rate throughout the period, with a dip toward the end of the sample.

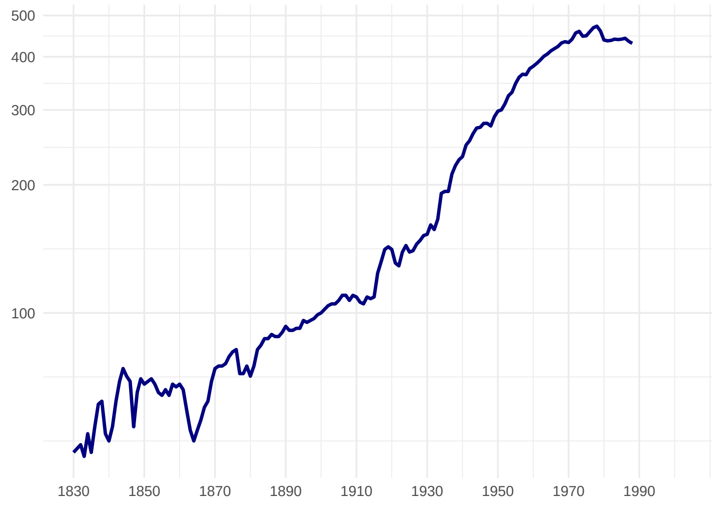

Figure 2.12: **Historical real wages in the U.S.**
Source: Williamson (1995, Table A1.1).

Turning to another, though related, topic of great recent interest, Figure 2.13 shows factor shares in the U.S. since the 1930s. We see a stable pattern, again, but closer inspection reveals a slight downward trend in the labor share over the last decades.

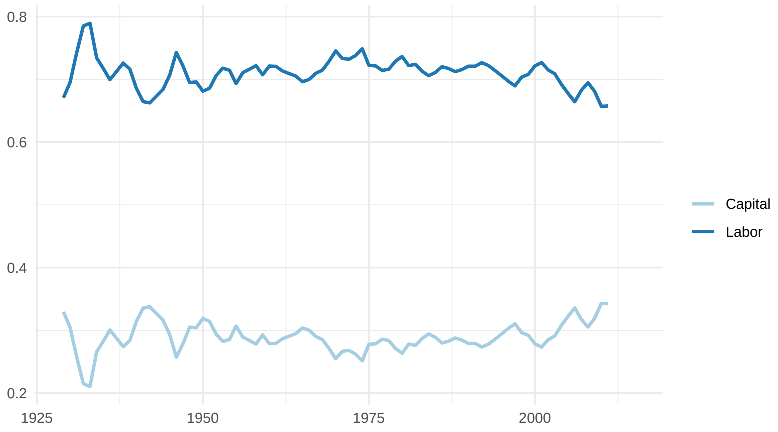

Figure 2.13: **U.S. factor shares**
Source: updated version of Piketty and Saez (2006).

The downward trend in the labor share is a world-wide phenomenon; Figure 2.14 and Figure 2.15 show the patterns for a selection of countries and as a global average.

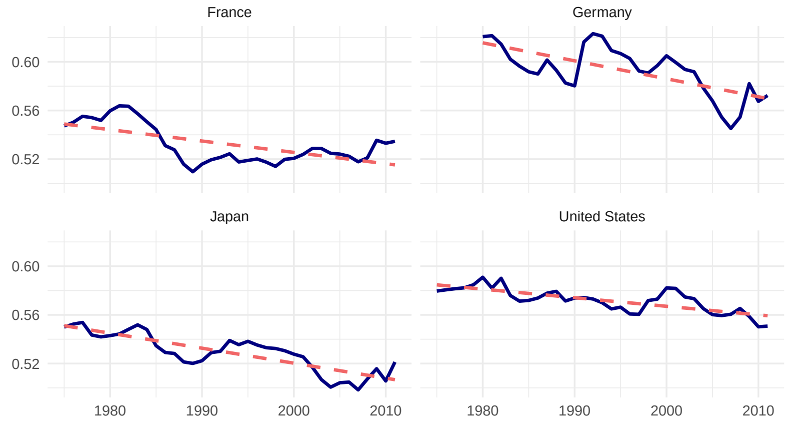

Figure 2.14: **Labor shares for a selection of countries**
Source: Karabarbounis and Neiman (2014).

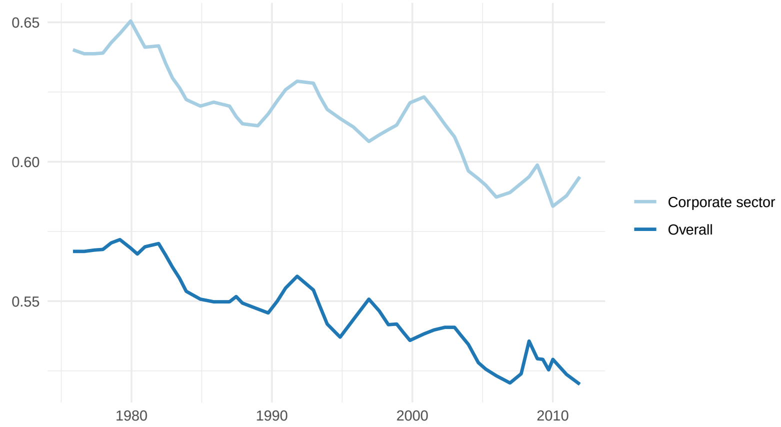

Figure 2.15: **The global labor share**
Source: Karabarbounis and Neiman (2014).

Finally, we examine interest rates. Figure 2.16 shows the return on capital, as measured by stock returns over long periods of time. The levels are rather stable at slightly below seven percent per year.

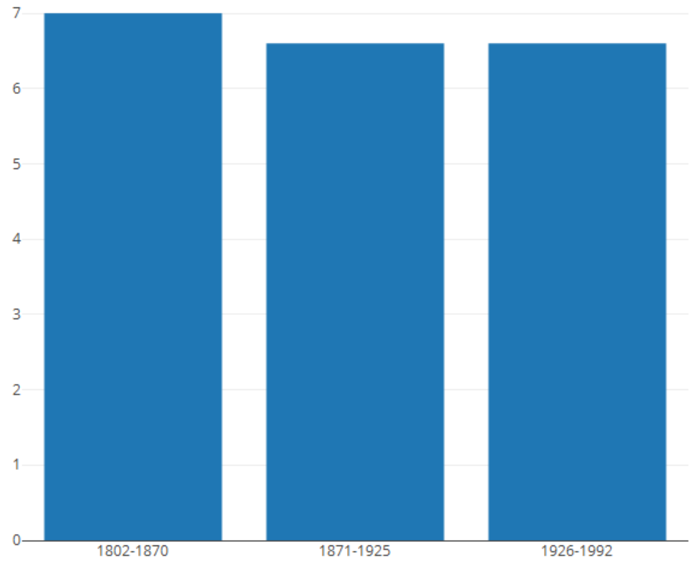

Figure 2.16: **The historical return to capital**
Source: Kongsamut, Rebelo, and Xie (2001).

The real return on riskless assets is much lower and varies quite a bit over time as well. Figure 2.17 shows the last few decades of data on the real return on government bonds—in this case averages across countries—and the much-discussed downward trend is very clear: from a rather high rate in the 1980s of about 4 percent down to essentially zero.

Figure 2.18 shows real bond returns over a longer time period. Here, the recent downward trend clearly represents a drop back toward a long-run mean toward zero, after the stark rise in bond returns in the 1970s (along with an increase in inflation during that period).

Finally, we discuss energy. Energy use is sometimes mentioned in macroeconomic analyses, though mostly then in the context of oil and especially during the period following the large price increases in the early and late 1970s. An important question for long-run modeling is how energy prices comove with economic activity. Figure 2.19, first thus shows the (real, using the GDP deflator) price per unit of energy. We see that the price movements reflect oil- and gas-price movements, since these two fossil fuels constitute an important part of the overall energy use. We note, in addition to the fluctuations, a significant upward trend in the price. Second, Figure 2.20 shows energy as a fraction of output. Here, the trend is much less apparent (there is only a slight downward movement), and it is also apparent, looking at both figures jointly, that the share follows the

Figure 2.17: **Real rates on government bonds, recent trend**
Source: Rachel and Smith (2015).

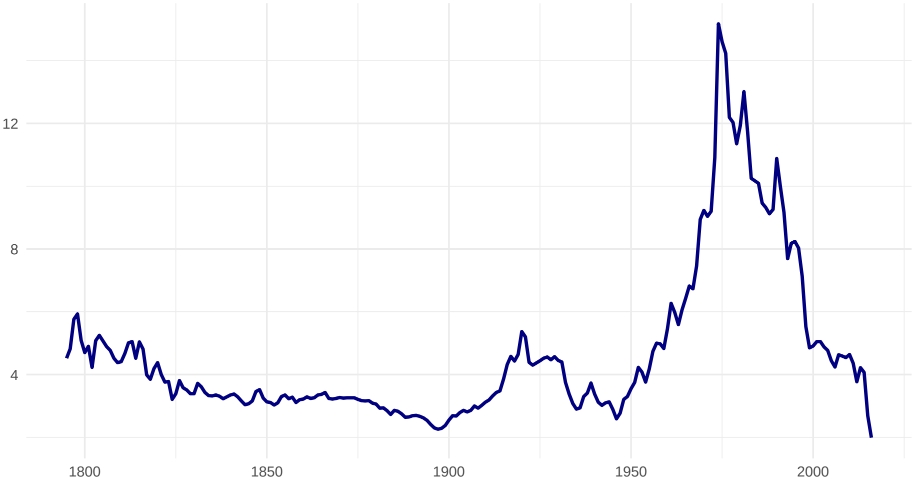

Figure 2.18: **Real rates on government bonds, long-run trends**
Source: Rachel and Smith (2015).

price closely in the short run.[4]

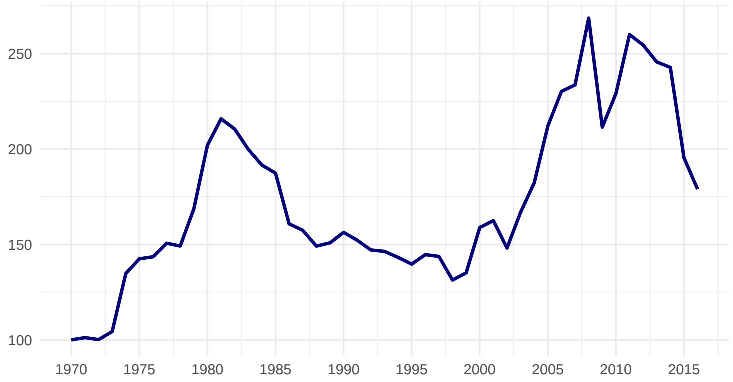

Figure 2.19: **The real price of a unit (Btu) of energy, U.S.**
Average real (using a GDP deflator) price of a Btu for the U.S., including all energy sources. **Source:** US Energy Information Administration.

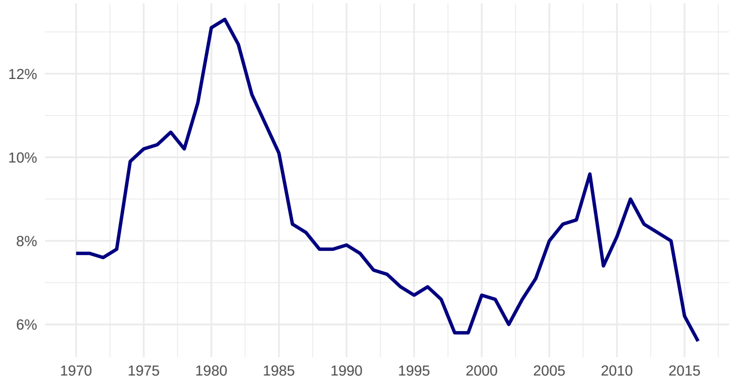

Figure 2.20: **The energy share in the U.S.**
The total nominal energy bill divided by nominal GDP. **Source:** US Energy Information Administration.

The historical time series we discussed above motivate the kinds of macroeconomic frameworks that have become the core of quantitative-theory settings used in applied work. Let us therefore summarize the "stylized growth facts" as follows, before we move on to discuss the ingredients in the theories used to account for them. In all cases, the word *approximately* is presumed as a qualifier. The facts are thus:

1. output per capita has grown at a constant rate

2. the capital-output ratio has remained constant

---

[4] A longer time series would be useful here, as well as similar facts across countries, but we have not been able to find consistently constructed series of this sort.

3. consumption as a fraction of output has been constant
4. hours worked have fallen at a constant rate
5. productivity has grown at a constant rate
6. the wage and capital shares of income have been constant
7. the real interest rate has been constant.

Some concluding comments are in order. First, the fact on hours is often stated instead as hours being constant. Here, since we take a longer-run perspective it seems appropriate to say that hours fall at a (small but significant) constant rate, in contrast to the conclusion from looking only at the postwar period of the U.S., i.e., that hours are constant. Second, there are some implications from the above facts, such as a constant rate of real wage growth (at a rate at, or slightly above, that of output growth)—given that the wage share of output is constant and hours are stable (or fall slightly)—and an investment-output ratio that is also fairly constant. Another implication is that the capital-labor ratio is increasing at a constant rate. Third, though it is clear that these facts are only approximate, the fluctuations are different for different time series; the returns series, for example, fluctuate greatly whereas the consumption-output ratio is much more stable. Fourth and finally, a variable that is often synonymous with business cycles and macroeconomics is missing in the above discussion: unemployment. We could have listed data on this variable too, with the conclusion that there does not seem to be a long-run trend in the rate of unemployment (but unemployment, clearly, is a highly volatile variable over the shorter horizon and has also experienced medium-term swings that are significant, such as the so-called hysteresis period in Europe beginning in the 1970s when unemployment rose very persistently).

In addition to the stylized facts above, we have also emphasized recent trends that indicate possible departures from the stylized facts, such as the falling labor share and the falling rate of productivity growth, but these are speculations and not large departures from the historical patterns. Therefore we will insist that a reasonable theory fit the seven facts above.

2.2 Quantitative theory

We now organize the facts below using a core theory, including a specific parameterization that not only delivers the facts qualitatively but also quantitatively. This parameterization will be outlined at the end of this section and is often called "calibration", which we take to be an informal method of estimation.

2.2.1 The setting

The core framework builds on the macroeconomic accounting identity: $C + I + G + NX = Y$. We will restrict attention here to a closed economy, in part because the application of the theory to the U.S. economy has traditionally been at center stage; the U.S. has been viewed as close to a closed economy, given its size. However, here we will ultimately have in mind a global economy, which by definition is closed. We will also lump government and private consumption together, or alternatively abstract from government consumption and investment. Thus, our economy's resource constraint will read

$$c_t + i_t = y_t$$

at time t, and we use lower-case letters to denote per-capita variables. Solow (1956) analyzed macroeconomic growth from this perspective and added two key elements: output is produced from capital and labor at any point in time, based on an *aggregate production function* and capital consists of past investments, corrected for depreciation. These elements are spelled out as follows. First,

$$y_t = F(k_t, A_t h_t),$$

where F has some key properties to be discussed below and where A_t is an exogenous labor-augmenting technology term; hours worked per capita, h, is permitted to vary here. Second,

$$k_{t+1}(1+n) = i_t + (1-\delta)k_t,$$

i.e., depreciation of capital is "geometric" at rate $\delta \in (0,1)$ and n is the rate of population growth. To save on notation, let us set $n = 0$ in what follows.

In our formulation, we use discrete time; continuous time is of course an alternative. Discrete time is useful here as we are interested in long-run issues and, thus, tracking the economy continuously over time will not be important. In particular, one type of special case is that where each discrete time period is ten years (or more) long. We will look at special cases later on with one or two periods only—these are not entirely unreasonable setting in the climate application.

The existence of an aggregate production function—a function producing GDP from the total amounts of capital and labor—is a key assumption but, so long as sectoral production functions have isoquants that are not too dissimilar, not a very restrictive one in a long-run model.[5] To elaborate slightly on this point in terms of the two uses of final output here, consumption and investment, the setting assumes that these goods are perfect substitutes with a constant relative rate of exchange but one can depart from these assumptions easily; an extension that is strictly presumed within the present one is that where consumption goods are produced using $F(k_c, A_{ct}h_c)$ and investment goods using $F(k_i, A_{it}h_i)$, with $k_c + k_i = k$ and $h_c + h_i = h$, for which it is easy to show that the relative price of investment will be pinned down by the ratio A_{ct}/A_{it}, i.e., exogenous.[6]

---

[5] Over a short time horizon, capital and labor are allocated to sectors and difficult to move, but in a longer run one can abstract from the moving costs and then by definition the economy's capacity is just a function of the total amounts of each of the inputs.

[6] For an extension to three broad sectors—agriculture, manufacturing, and services—with overall balanced growth while allowing structural, long-run change between these sectors, see Herrendorf, Rogerson, and Valentinyi (2014).

Solow assumed that F has constant returns to scale (CRS) and is "neoclassical" in the sense that there are diminishing returns to each input. Absence of overall decreasing returns can be motivated by a replication argument—it should be possible to at least produce twice as much if all inputs are available in double quantities—and the estimated degree of returns to scale in the literature lands somewhere only slighly above constant returns. F is also assumed to be (quasi-)concave and to satisfy regularity conditions along with a sufficiently high marginal product of capital at $k = 0$ and a sufficiently low marginal product at $k = \infty$ (two often-used values are ∞ and 0, respectively). Finally, Solow assumed that technology growth is labor-augmenting following Uzawa's (196xyz) analysis, proving that no other type of technical change is consistent with balanced growth; given such a long period of approximate balanced growth, this assumption seems necessary.

The setting so far leaves two variables undetermined: investment and labor input. These both reflect human choice. Solow simply assumed that investment is a constant share s (for the rate of saving) of output and set $l_t = l$, which can be normalized by setting $l = 1$. The assumption of a constant saving rate will immediately deliver the fact that the consumption-output ratio is constant. The assumption that labor input is constant is inconsistent with the longer-run data and we will amend that assumption here. Perhaps more importantly, one should view the inputs here as "utilized inputs", hence allowing for less than full factor utilization (of capital and labor) but at rates that do not vary in the long run. For labor, this assumption is justified with appeal to a long-run rate of unemployment that is not trending.

Under the assumptions stated, it is easy to show—guess and verify—that if labor-augmenting technology grows at a constant net rate g and labor input falls at a net rate g_l, then there is a (unique) exact balanced growth path for capital, investment, and output where all these variables grow at a net rate $(1+g)/(1-g_l) - 1 \approx g - g_l$; here, given an A_0 there is a unique k_0 (and associated values i_0 and y_0) leading to exact constant growth. Not only that, if the initial capital stock takes on any arbitrary value, capital (and investment and output) will converge monotonically to the balanced growth path. This result does not extend in generality to more complicated (say, multi-sector) growth models, but there are no quantitatively parameterized examples of economies where convergence does not apply, and much data analysis from a cross-section of countries support the convergence feature (see Barro, symposium paper).

Solow also assumed (in particular, see the growth accounting analysis in Solow, 1957) that prices for inputs were determined in perfectly competitive markets, hence implying that the rental rate for capital is equal to the marginal product of capital and that the real wage is equal to the marginal product of labor. With a CRS production function this implies that the capital share of output, $F_1 k/F$, and the labor share of output, $AF_2 l/F$, sum to one and are each constant along any balanced growth path (since the first derivatives of F are homogeneous of degree 0 when F is homogeneous of degree 1). Moreover, the rental rate $r = F_1$ will be constant along a balanced growth path and therefore so will the return to capital, $1 - \delta = r$, and the real wage will also grow at a constant rate since

$w = AF_2$ will grow at the rate of productivity growth, g. Hence, the setting delivers the key facts on the price-related variables: prices and shares.

The subsequent literature went on to find a more primitive basis from which the two behavioral assumptions on the consumption share of output and labor input could be derived. The optimizing growth model was such a basis, where the idea is that consumers choose their levels of consumption, along with their supply of labor, to maximize a utility function. Thus, an assumption on the population structure is needed, along with a specific utility function class, and the question then is what population and utility function assumptions are consistent with the balanced growth facts. As for the population structure, the most common practice is to assume a representative-agent dynasty. The dynasty part simply means that currently alive consumers derive utility from the consumption of their offspring in a perfectly altruistic way, i.e., their are no differences in evaluation of the consumption paths of their offspring between generations. There are alternatives, such as the overlapping-generation model, but such models lead to essentially identical restrictions on preferences. As for the representative-agent assumption—one type of agent within each cohort—there is by now a large and growing literature on heterogeneous-agent macroeconomics that explores various reasons why an aggregation theorem does not apply. In that literature, so far, the findings indicate that whereas the short-run behavior of the macroeconomic aggregates can be influenced greatly by departures from the representative-agent assumption, there is no indication that the long-run features will.[7]

As for the specific utility function, the maintained assumption is that utility is time-additive with a constant discount factor, i.e., $\dots + \beta^t u(c_t, h_t) + \beta^{t+1} u(c_{t+1}, h_{t+1}) + \dots$. For such a setting, one can then show that balanced growth obtains if and only if $u(c, h)$ is a power function of $cv(c^{\frac{\nu}{1-\nu}} h)$. [8] Here, v is a decreasing function and ν is a parameter describing the amount by which the income effect exceeds the substitution effect. Also, if productivity growth occurs at rate g, hours will grow at a gross rate $(1 + g)^{-\nu}$.

Note that the maximization of utility over time amounts to assuming rational expectations. Parts of the modern macroeconomic literature have explored departures from fully rational expectations, but few macroeconomists believe that there is no forward-looking at all—such as in regard to future taxes and policies. In the context of climate policy, it makes sense to think that firms and consumers make decision that to some extent factor in what is announced about future policy and future paths of fundamental determinants of prices (such as technology developments). Macroeconomists are fully aware that unlimitedly rational expectations is just an abstraction but short of other convenient, well tested frameworks that embody biases while maintaining a significant amount of

---

[7] See, e.g., Boppart and Krusell (2018), who explore a setting with incomplete markets and show that the restrictions placed on the representative agent's utility function to obtain all the balanced-growth facts also works well in this much more elaborate setting.

[8] A power function here is meant to include the logarithm as a special case. In addition, restrictions need to be placed on v and the power jointly so as to ensure concavity; see Boppart and Krusell (2018) for details. It is possible to consider an extensive hours choice too, in a context with heterogeneous agents, with similar conclusions.

forward-looking it remains the core assumption.[9]

In sum, we have arrived at the following structure:

- Population: representative dynasty (for illustration only: constant population size).
- Utility function: $\sum_{t=0}^{\infty} \beta^t u(c_t, h_t)$, where u is a power function of $cv(c^{\frac{\nu}{1-\nu}} h)$.
- Resource constraint: $c_t + k_{t+1} = F(k_t, A_t h_t) + (1-\delta)k_t$, where F is CRS and A_t displays geometric trend growth at net rate g.
- Rational behavior—to be specified below in the context of a market structure.

Clearly, a "social planner" here would just maximize utility subject to the constraints listed. In the next section, we will also define the dynamic market equilibrium, which will coincide with the planning outcome to the extent it is competitive and there are no relevant other distortions (such as taxes or other frictions). There we will also show how to specify this framework further in terms of specific functional-form choices and parameter values.

Solving the model above requires numerical methods, unless one makes very particular functional-form assumptions (we will consider relevant such examples below). However, the model has features that can be ascertained analytically, such as monotone convergence to a steady state. Characterization of the steady state is also straightforward; one differentiates with respect to k_{t+1} and h_t to obtain two key equations, the intertemporal Euler equation and the intratemporal effort choice condition, and then evaluates on the balanced growth path. The first step is thus

$$u_1(c_t, h_t) = \beta u_1(c_{t+1}, h_{t+1})(1 + F_1(k_{t+1}, A_{t+1}h_{t+1}) - \delta) \tag{1}$$

$$A_t F_2(k_t, A_t h_t) u_1(c_t, h_t) = u_2(c_t, h_t), \tag{2}$$

and we will then evaluate these equations on the balanced path once we have chosen functional forms for u and F. This will give us two equations in two unknowns, k_0 and h_0, i.e., the starting values from which constant growth will then occur. Notice that the

2.2.2 Market equilibrium and calibration

The market equilibrium Solow's setting presumed perfectly competitive input markets and we will adopt the same setting here. The markets for output will also be taken to be competitive—and we will define the numéraire in each period to be the consumption good. Perfect competition is often relaxed in favor of monopolistic competition, thus featuring some amount of limited market power, calibrated so as to match measures of markups; these settings are used mainly to model frictions in price and wage setting, i.e., for business-cycle analysis. Similarly, the inclusion of money and inflation is standard in these settings but as there appears to be little support in favor of monetary frictions

---

[9] For a promising path forward, see, e.g., Gennaioli and Shleifer (2010).

being important for long-run issues, we will abstract from these here, along with any monopoly elements. Common frictions considered in the business-cycle literature include borrowing constraints for consumers, limited insurance against idiosyncratic consumer risks, search and matching frictions in the labor market, and credit frictions for firms. However, none of these seems crucial for long-run modeling, especially since we can include wedges, such as taxes, as a short-hands for important distortions in different markets, such as that for investment and possible credit frictions there. Similarly, many macroeconomic frameworks for business-cycle analysis have various forms of adjustment costs (to investment, or in changing consumption habits) that operate only over shorter time horizons and hence will not be relevant here.

We assume that firms act competitively both in input and output markets. The price of the consumption good in each period is normalized to 1 and, from competitive input pricing, the relative prices of labor and capital services must then simply equal marginal products. We can therefore define a competitive competitive equilibrium mathematically as follows: (deterministic) sequences of quantities $\{c_t, h_t, k_{t+1}\}_{t=0}^{\infty}$ and prices $\{r_t, w_t, \}_{t=0}^{\infty}$ such that

1. $\{c_t, h_t, k_{t+1}\}_{t=0}^{\infty}$ maximizes
$$\sum_{t=0}^{\infty} \beta^t u(c_t, h_t)$$
subject to
$$c_t + k_{t+1} = (1 + r_t - \delta)k_t + w_t h_t$$
for all $t \geq 0$ and a no-Ponzi-game restriction (the limit present value of k is non-negative);
2. $r_t = F_1(k_t, A_t h_t)$ and $w_t = A_t F_2(k_t, A_t h_t)$ for all t; and
3. $c_t + k_{t+1} = F(k_t, A_t h_t) + (1 - \delta)k_t$ for all t.

Note that the second condition summarizes firm profit maximization and implies that profits are zero, given that F is CRS. Note also that the third condition is implied by the first two (by Walras's law).

Calibration Turning to functional-form and parameter selection, let us first discuss how to select a specific utility function. A key input here is to select the curvature of the function, i.e., the power to which $cv(h)$ is raised. This number will be related to the elasticity of intertemporal substitution of consumption: the elasticity is 1 over (1 minus the power). The macroeconomic literature contains various arguments about this value, some based on empirical microeconomic analyses and some based on macroeconomic data, and the most common assumption is to set the elasticity to 1, i.e., to use "logarithmic power"; certainly large departures from 1 are very unusual, though values like 1.5 are common as well. A unitary elasticity means that income and substitution effects cancel when one considers real interest-rate changes.

Furthermore, the v in the utility function is selected to match data on how willing consumers are to substitute labor over time—work hard one year to take advantage of higher returns from working in exchange for less hard work at other times, while being able to smooth consumption by saving. Here there is larger dispersion in the specific assumptions made but the following functional form allows us to capture the key features in the data:

$$u(c,h) = \frac{c^{1-\sigma} - 1}{1-\sigma} - \psi \frac{h^{1+\frac{1}{\theta}}}{1+\frac{1}{\theta}},$$

which is due to MaCurdy (1981).[10] Heathcote, Storesletten, and Violante (xyz) use this function and microeconomic panel data on consumers to estimate $\sigma = 1.7$ and $\theta = 0.5$, implying that hours fall at a rate consistent with the long-run data and a Frisch elasticity of labor supply—the percent increase in hours if the wage rate rises temporarily by 1%—of 0.5. A $\sigma = 1.7$ implies a elasticity of intertemporal substitution of $1/1.7$, i.e., about 0.6. Much of the macroeconomic business-cycle literature instead uses $\sigma = 1$ (logarithmic curvature), hence implying no trend change in hours in respond to trend changes in wages (again, based on seeing no trend in hours trend over the postwar period in the U.S.).[11]

Furthermore, the production function is usually assumed to be of the Cobb-Douglas form, i.e., $F(x,y) = x^{\alpha}y^{1-\alpha}$, implying that capital's share of income is constant at α and labor's constant at $1-\alpha$. As we have seen, there is a recent downward trend in the labor share, and some analyses therefore consider a slightly different production function—one with a constant elasticity of substitution at a value close to but not equal to one—but for our present purposes we choose to keep the Cobb-Douglas form.

With these functional-form choices, we obtain stationary versions of our key first-order conditions (1)–(2) by evaluating on the balanced growth path and using a small number of substitutions:

$$(1+g)^{\sigma} = \beta \left(1 + \alpha \left(\frac{k_0}{A_0 h_0} \right)^{\alpha-1} - \delta \right)$$

$$A_0 \left(\frac{k_0}{A_0 h_0} \right)^{\alpha} \left(k_0^{\alpha} (A_0 h_0)^{1-\alpha} + (1-\delta)k_0 - (1+g)k_0 \right) = \psi h_0^{\frac{1}{\theta}}.$$

As announced, these equations have two unknowns, k_0 and h_0, that can be easily solved for numerically.

Given our functional forms there are now 7 parameters to calibrate: β, σ, ψ, θ, δ, α, A_0, and g. We now list how they are selected, based on annual data:

1. $g = 0.02$ is set to match the average growth rate of output.

---

[10] It is perhaps not obvious that the MaCurdy function satisfies the general form. See Boppart and Krusell (2018) for details.

[11] The utility function curvature should be related to the time horizon, as there are arguments to suggest that long-run elasticities of substitution are higher than those in the short run. We take our choice to be a long-run elasticity.

2. σ and θ are selected so as to be consistent with the cited microeconomic estimates: $\sigma = 1$ (in most studies; or $\sigma = 1.7$, to obtain a fall in hours at an appropriate long-run rate, given the θ selected below) and $\theta = 0.5$, reflecting an empirically plausible Frisch elasticity.

3. ψ and A_0 can be normalized; we set them to 1.

4. $\alpha = 1/3$ from the data on the average capital share of output.

5. δ is selected so as to match the depreciation rates used in practice, say, when capital is written off from firms' balance sheets. A common value is somewhere between 5 and 10 percent per year; there is large variation across capital goods, and the composition of the aggregate capital stock changes over time. Alternatively, one can use aggregate investment and capital data and back out the depreciation rate used on average. So on a balanced path, and abstracting from population growth, we have $k(1+g) = (1-\delta)k + i$, where the left-hand side is the capital stock next period. Hence $i/k = g + \delta$. In the data, a number for i/k of a little less than 0.1 is roughly right—by definition $i/k = (i/y)(y/k)$, with i/y roughly 0.3 and y/k roughly 1/3.3, from the graphs in the previous section—and hence $\delta = 0.3/3.3 - 0.02 \approx 0.07$ seems a sensible value to select given that $g = 0.02$.

6. The model's parameters also need to be consistent with a saving rate of around 0.3 and a capital-output ratio of 3.3, as just stated. Since the two are connected through $(i/y)(y/k) = g + \delta$, let us derive an implication from $k/y = 3.3$. Given the Cobb-Douglas production function, $rk = \alpha y$, implying $r \approx 0.1$. r appears, naturally, in the optimality condition for saving. The Euler equation above thus implies a restriction on β. If we use MaCurdy preferences and $\sigma = 1$ for simplicity, we obtain $1 + g = \beta(1 + r - \delta)$, delivering $\beta \approx 0.99$.[12]

With this, we have a fully specified model matching all the stylized facts in Section 2.1 above.

A key takeaway from the above analysis is not only that it is possible to match our long-run facts based on a model with a coherent microeconomic structure but also that it is very difficult (if not impossible) to come up with an alternative structure that is also rooted in empirical microeconomic work. It is of course possible to come up with variants of the present, stripped-down setting, by considering consumer or firm/sector heterogeneity and a variety of frictions, but those amendments would have to obey the same long-run facts and will be hard to generate with fundamentally different basic components than those above. It is for this reason that we insist on building a long-run climate model around this foundation.

---

[12] Inserting the return on capital in the Euler equation without considering risk is not entirely appropriate. One could instead insert a number for the average riskfree rate from the graphs above in this equation and obtain thus the β directly. The connection between this rate and $r - \delta$ would then involve a risk premium.

2.2.3 Uncertainty

The model above is deterministic. In many contexts, especially that involving climate change, it can be important to explicitly incorporate uncertainty. In the macroeconomic business-cycle literature, uncertainty is at the core of many issues so we will now briefly review how uncertainty is incorporated into the analysis before we begin addressing climate issues.

The preferences of the consumer need to be extended to incorporate uncertainty. The von Neumann-Morgenstein expected-utility formulation does this and we simply use $\mathbb{E}_0\left[\sum_{t=0}^{\infty}\beta^t u(c_t, h_t)\right]$, where the variables are now interpreted as stochastic processes, in most contexts. Two issues are important to point out here, however.

One is that the typical business-cycle model used in applied contexts, such as at central banks, relies on sophisticated, stochastic versions of the present model (with a large number of different shocks hitting the economy), but they way these models are solved—usually, at least—makes uncertainty play a rather minor role. The reason is that they are solved by "linearization", i.e., by deriving a set of necessary conditions for equilibrium, including first-order conditions and constraints, and then Taylor-expanding them to first order around a balanced-growth/steady-state point and solving the obtained system for a set of linear rules that can then be simulated and analyzed. The issue is that this procedure makes all the coefficients of the solved-out linear rules independent of the shock variances. In particular, it makes the dynamics of the system identical to that of a deterministic model. Intuitively, linearization implies that the model will be equivalent to one that is linear-quadratic (LQ)—a model with a quadratic objective and linear constraints, whose first-order conditions and constraints will all be linear and hence easily solved—and it is well known that LQ models feature *certainty equivalence*, i.e., precisely that variances do not affect decision rule coefficients.

To the extent the modeler wants to incorporate uncertainty that is believed to have more fundamental importance on the decision rules, there are two typical ways forward. One is to still analyze the model locally but use higher-order Taylor terms (these are available in standard easy-to-use numerical packages such as DYNARE). The other is to solve the model globally with appropriate numerical methods. Non-linear global model solution is more application-specific and requires us to employ recursive methods—dynamic programming—as opposed to the sequence-based methods we used to describe the setting above. The reason is that recurrent uncertainty leads to an "expanding-tree" structure of states of the economy that, over a long time horizon, becomes untractable to analyze since the number of unknowns explodes. We discuss numerical methods in some more detail below.

Turning now to the second important point to discuss, models with time-separable preferences of the kind discussed so far are restrictive in that they do not allow us to separately model the elasticity of intertemporal substitution and risk aversion. Utility-function curvature—notably σ—regulates both these concepts, which are fundamentally different. In particular, a reasonably high willingness to substitute consumption over time should be possible to consider while allowing very high aversion to risk, but the former would mean a σ near one or below it whereas the latter would imply a much higher σ, such as 10 or above. Aversion to risk is indeed a commonly cited argument why climate policy should be strict. The two-parameter Epstein-Zin utility formulation is the most common way for applied macroeconomists to relax these modeling constraints. It is not time-additive and most easily described using recursive methods. Moreover, even beyond this point, one can consider ambiguity and ambiguity aversion, i.e., a departure from Savage's subjective probability assumptions. This too can be accomplished with recursive methods and one or more additional utility parameters to capture ambiguity aversion and its nature (see, e.g., Gollier, 2012).)

Both points above—global nonlinearities and high risk aversion with moderate intertemporal substitution elasticity—speak in favor of the use of recursive methods. Let us therefore specify our model in these terms. For simplicity, let us use time-additive preferences and that there is no growth, i.e., that $g = 0$, but consider an A that follows a two-state Markov process. For the planner the problem is simply one of finding a function v that solves

$$v(k, A) = \max_{k',h} u\left(F(k, Ah) + (1-\delta)k - k', h\right) + \beta E_A v(k', A')$$

for all (k, A). This problem delivers decision rules $k' = g^1(k, A)$ and $h = g^2(k, A)$, which can be simulated.

The competitive equilibrium is significantly more complicated to describe, since it involves the distinction between an individual consumer's behavior and the behavior of other consumers. Conceptually, even though the representative-agent model will have the feature that in equilibrium the individual's capital holdings and hours worked will coincide with those of all other consumers, they are distinct and must be separated in order to ensure price-taking behavior. Thus, let k be the individual's capital stock and K the aggregate/average capital stock—the two key *state variables* for the consumer. The evolution of the aggregate capital stock will be given by $K' = G^1(K, A)$ and aggregate hours worked by $H = G^2(K, A)$. An equilibrium is now defined by the functions v, G^1, G^2, g^1, g^2, R, and W such that

1. v solves

   $$v(k, A, K) = \max_{k',h} u\left((1 + R(K, A) - \delta)k + W(K, A)h - k', h\right) + \beta E_A v(k', A', K')$$

   for all (k, A, K), with the implied decision rules being $k' = g^1(k, A, K)$ and $h' = g^2(k, A, K)$;

2. $R(K, A) = F_1(K, AG^2(K, A))$ and $W(K, A) = AF_2(K, AG^2(K, A))$; and

3. $G^1(A, K) = g^1(K, A, K)$ and $G^2(A, K) = g^2(K, A, K)$.

This definition, in its steps, follows the sequence-based one in Section 2.2.2 above for the deterministic case; what is different is the last step, which ensures ensure *consistency*: that all agents behave identically. Walras's law already invoked here: the resource constraint holds and does not need to be stated.

The recursive equilibrium definition is much more complex in that it involves five nontrivial functions. In a solution based on linearization, the value function would be summarized as a quadratic function whereas the remaining functions would be linear.

An extension to a model with more consumers is now conceptually straightforward: the aggregate state variable will be the distribution of capital among agents instead of simply K, and G^1 will be a function mapping the distribution of wealth and A into a distribution next period.

2.3 Energy resources

Moving closer to the IAM, we now consider an addition to the basic macroeconomic framework that involves a natural resource: fossil fuel, such as oil, coal, or natural gas. This resource is an important input into an IAM since it, when burnt, is the source of carbon dioxide emissions and, hence, warming. In most purely macroeconomic studies, this resource is typically abstracted from, chiefly because enery commands a small share of overall costs (on the order of 5%) in any developed economy. The issue of finite energy resources and oil became topical in the mid 1970s and was central for about ten years, however, when the oil shocks hit, since these were substantial and, according to many, caused the protracted recession/productivity slowdown in the developed countries during the period.[13] Here, we will look, first, at the economy's demand for fossil fuel and then at its supply.

2.3.1 Energy demand

We model the demand side of fossil fuel/oil from the perspective of a production technology using this resource as an input. Fossil fuel is also used directly by consumers, notably for heating and driving vehicles, but we will follow Dasgupta and Heal (1974) here and think of all fossil-fuel needs as emanating from the production of GDP.[14] We will allow for technical change to be energy-saving as well as saving on other inputs partly because the issue of endogenous energy saving is an important one in the context of climate change. We will also briefly discuss other, non-fossil or "green", energy sources.

Thus, now consider a production function of consumption goods of the form $F(k, Ah, A_eE)$, where E itself is an energy aggregate and A_e is an energy-augmenting—energy-saving—technology parameter that will change over time, as does A. Here, it is again convenient

---

[13] A consensus later seemed to emerge where the oil shocks were not believed to be quantitatively relevant enough—precisely because oil commands such an insignificant cost share and the protracted recession was so significant. Recently, however, arguments have been put forth than a high short-run complementarity between oil and other inputs could, through non-linearities, generate a large recession from a large enough oil shock; see Baqaee and Farhi (2017).

[14] If the services from housing and consumer durables are included in GDP, this practice seems fully appropriate.

to take a quantitative-theory approach and ask what forms for this extended production function are quantitatively reasonable. Some basic facts were shown in Figures 2.19 and 2.20 above: large fluctuations in the real price of energy and in the energy share and a clear upward trend in the former and a much weaker trend in the latter. Hassler, Krusell, and Olovsson (2017) look more specifically at a fossil-fuel composite and have similar conclusions. Ideally, one would want to look at the energy share over a long period of time where there is a stronger price trend than that observed for the U.S., or alternatively across countries with very different costs of energy. We will base our structure here on the findings in Hassler, Krusell, and Olovsson (2017). They argue, based on long time-series data that a nested CES formulation of production with energy as an input works well: a Cobb-Douglas nesting of k and Ah and then a CES in this composite and A_eE. Moreover, they argue that on an annual frequency, the latter CES aggregator is indistinguishable from a Leontief, i.e., fixed-factor, function, because at this frequency, changes in fossil-fuel prices appear to drive their cost share up one for one. At a longer frequency, however, the oil/energy share is rather stable and, as we will see, one interpretation is that endogenous changes in the input-augmenting technology series A and A_e are plausible candidates behind this stability. For now, however, we will treat these series as exogenous. In sum, we assume $F(k, Ah, A_eE) = \min\left\{k^{\alpha}(Ah)^{1-\alpha}, A_eE\right\}$ for shorter-run modeling and $F(k, Ah, A_eE) = k^{\alpha}(Ah)^{1-\alpha-\nu}(A_eE)^{\nu}$, where ν is the energy cost share, for longer-run models (where the time period is, say, ten years or longer).

Turning to the energy aggregate E, we will simply assume a CES structure of three inputs: oil/natural gas, coal, and green sources. This structure could be nested as well but for simplicity we maintain one parameter for the elasticity of substitution across different energy sources. A meta-study (Stern, 2012) suggests that a reasonable elasticity value is close to unity, but there is much variation in the literature in terms of estimates. In our illustrative applications, we will look at special cases with either only oil or only natural gas.

2.3.2 Energy supply

The production of fossil fuel will be described in a very simple way in our benchmark model: either we consider oil alone or coal alone. The reason for the brief treatment is simply that these two cases are fundamentally different and bracket intermediate cases. Note, however, that the key approach here is to find formulations that are quantitatively reasonable. So consider oil in its "conventional" variety: available in land/desert-based areas like Saudi Arabia. This kind of oil, first of all, exists in limited supply: the total estimated reserves of conventional oil measured in carbon content is around 300GtC, of which a little less than half is in the middle east.[15] For conventional natural gas a reserve of 200GtC is estimated. The total use of fossil fuel is about 10Gt per year annually.

Second, conventional oil has very low marginal cost relative to its price; the cost of extracting Saudi oil is likely less than \$10 per barrel, with a price much higher than that.

---

[15] See, for example, EIA, World Energy Outlook 2015.

We will therefore adopt the approximation that oil is entirely costless to produce. With a resource in finite supply that is costless to produce, we obtain the famous Hotelling (1929) result from the "cake-eating" problem: if a stock of oil R_t can be sold today or at any point in the future and it can be produced and stored costlessly, then so long as the resource is sold at two consecutive points in time, the rate of price increase for the resource between these points in time must equal the real interest rate. The argument is arbitrage. With a positive marginal cost of production, the Hotelling formula becomes

$$\frac{p_{t+1} - mc_{t+1}}{p_t - mc_t} = 1 + rr_{t+1},$$

where p is the price of the resource, mc its marginal extraction cost, and rr the real interest rate. Here, $p - mc$ is often called the Hotelling rent—the value accruing to the seller over and above production cost due to the resource being in finite supply. Suppose now, for example, that the real interest rate is constant at rr and that the marginal cost rises at rate g_{mc}. Then this difference equation implies that the rate of price increase exceeds rr if and only if $\rho > rr$: a higher production cost in the future requires a compensating higher price increase to make the producer indifferent between production today and in the future.[16] It should be noted that our analysis here aims at long-run assessments. If one looks at short time horizons, oil production is very costly to vary (less so for unconventional supplies) so a much richer set of dynamics must be specified in order to understand equilibrium in the oil market (see Bornstein, Krusell, and Rebelo, 2018).

Notice, here, that we have discussed oil supply from the perspective of price-taking. Given the existence of OPEC, this assumption may not be innocuous. However, many analysts actually argue that price-taking is a good approximation, and certainly a better one than monopoly, given that OPEC controls less than 50% of the market, and we will therefore not consider monopoly power here.[17]

Turning to coal, the total estimated amount of coal is much larger. The amount of recoverable coal resources is in the order of 16,000GtC according to EIA, though only a part of this quantity is profitable to extract at current prices. Moreover, the marginal cost of coal is very close to its price (which is consistent with its available reserves being "nearly infinite"). Hence, we will, when modeling coal, approximate this production with that of a resource that is not in finite supply. Typically, we also will assume that its marginal cost is constant though possibly decreasing over time due to technical change.

When we also look at green energy, we will treat its production like we treat coal: with a constant marginal cost and its own rate of technical change.

2.3.3 Equilibrium

As announced, we focus on the special cases with coal only and oil only here.

---

[16] We obtain $p_{t+1}/p_t = 1 + r + mc(g_{mc} - rr)/p_t$.

[17] Monopoly power is also challenging to study since a monopolist in the oil market would view all macroeconomic variables in the world, today and in the future, as endogenous to its decisions; moreover, its profit-maximizing behavior under commitment would not generally be time-consistent.

Coal only Looking first at the simple case where only coal is used, suppose that its marginal cost is constant in terms of labor units. Thus labor is allocated across final-goods production and coal production to maximize $F(k, A(h - h_e), A_e A_c h_e)$ with respect to h_e—the part of overall labor h used for coal production. Here we have used $1/A_c$ to denote the marginal cost of coal in terms of labor. This outcome presumes either a planning solution where the use of coal does not involve externalities (such as that involving climate change) or a market allocation where coal production is not taxed—if it is taxed, a tax will enter the labor allocation first-order condition. This formulation, which abstracts from the use of capital and energy in the production of coal, is very convenient because it allows us to easily solve for h_e in terms of h and k. In the Cobb-Douglas case—where the energy share is constant—matters are even simpler: h_e will be a constant share of h that depends only on technology parameters. Thus, in this case, we are formally back in the optimal growth model above—after maximizing over coal energy we have merely added a constant in the production function.

Oil only The full model with costless-to-produce oil only amounts to adding a resource constraint to the growth model above—$\sum_{t=0}^{\infty} E_t = R_0$—aside from including E in the production function. We then obtain the canonical model of oil/finite resources in Dasgupta and Heal (1974), except for the presence of technology growth here. Let us briefly look at that model in its planning version and let us use a quantitative version without hours choice, as it will constitute a core of sorts in the climate model below.[18] Thus we have:

$$\max_{\{k_{t+1}, E_t\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \log\left(F(k_t, A_t, A_{et} E_t) + (1-\delta)k_t - k_{t+1}\right) \quad \text{s.t.} \quad \sum_{t=0}^{\infty} E_t \leq R_0.$$

The nature of growth paths for this economy depends critically on F.[19] Recall that in a calibrated version of the model we take F to be a low-substitutability CES function of the capital-labor composite and energy, if annual data are to be addressed, or a Cobb-Douglas in all inputs if the model is specified for long-run growth purposes.

Starting with the long-run model, when $F(k, A, E) = k^{\alpha} A^{1-\alpha-\nu} (A_e E)^{\nu}$ it is straightforward to show that the present model delivers exact balanced growth (under the right initial condition for k_0, and monotone convergence otherwise) where E grows at rate β, i.e., it goes to zero asymptotically, with output, consumption, and capital growing at the gross rate $(1+g)((1+g_e)\beta)^{\frac{\nu}{1-\alpha-\nu}}$, where g_e is the growth rate of A_e.[20]

We show in Krusell, Hassler, and Olovsson (2018) that the present model can deliver exact balanced growth also when formulated at the higher-frequency horizon when the CES has an elasticity of substitution less than one between capital-labor and energy, but

---

[18] With hours falling at some (small) exogenous rate we can simply reinterpret the growth of A as "net of hours falling". So abstracting from hours, in this sense, is without loss of generality.

[19] It is straightforward to consider the utility function to be the more general power function, but since $\log c$ is a focal point in applied work we use this formulation in most of our text.

[20] Along this path, of course, the marginal product of oil (its price in equilibrium) rises at the net rate of return on capital: its marginal product minus depreciation (the real interest rate).

only in a knife-edge case: when $\beta(1+g_e)(1+g)^{-1/(1-\alpha)} = 1$. When this expression is above (below) one, growth is not balanced in the usual sense; for example, energy share's cost share goes to zero (one). However, we also show that when the technology growth rates g and g_e are endogenous, then at least under certain assumptions the result is that energy's share again is robustly balanced in the long run: the economy looks like a Cobb-Douglas world, though with a share parameter that is a nontrivial function of other primitives.

An interesting feature of the long-run, Cobb-Douglas model here—where oil is a finite resource—is that oil use falls from time zero. That is, we do not obtain a rising path initially. Historically, oil use has been rising for a long time, and very steadily. Thus, one quantitative concern is whether the model ought not be altered, somehow, so as to match this rather basic fact. Interestingly, however, in the high-frequency version where oil and capital-labor has very low substitutability and where , this result obtains straightforwardly given certain initial conditions on k_0, A_0, and A_{e0}: if energy technology is, in some sense, at a high level relative to the capital-labor technology, adjusted for the initial size of the capital stock, then capital, not energy, is initially scarce and as capital is accumulated, energy follows along. Eventually, of course, as oil is finite, what factor is scarce in relative terms reverses, and oil use goes to zero at rate β.

3 The natural-science add-ons

In this section we cover the main natural-science modules needed in our IAM. Versions of these modules have been developed in Nordhaus's work; they summarize very complex natural-science mechanisms in a compact enough form that they can be feasibly used in a broad class of models, while still being quantitatively adequate. We keep this and the next section brief, however—much more extensive discussion can be found in Hassler, Krusell, and Smith (2016).

The overall logic of, and connection between, the modules should be clear: carbon dioxide is emitted by burning fossil fuel and atmospheric carbon dioxide—by virtue of being a greenhouse gas—causes warming. Carbon dioxide emissions quickly become global, i.e., their spatial spread is immediate from the perspective of modeling, but how long they remain in the atmosphere is a topic in itself. The "carbon cycle" describes this process and is described first. How atmospheric carbon then causes warming is dealt with in the climate module. We abstract from other aspects of climate than temperature.

3.1 The carbon-cycle module

A representation of the carbon circulation in IAMs is necessary in order to map emissions of carbon dioxide (CO_2) to a path of atmospheric CO_2 concentrations. In reality, carbon flows continuously between a number of carbon reservoirs (sinks) of which the atmosphere is one. The most important other reservoirs are the oceans and the biosphere. The interaction between these reservoirs is non-linear and implies that emitted CO_2 does not leave the atmosphere following a geometric path with a constant decay rate. This is in contrast to other greenhouse gases, like methane, for which a constant decay rate is a more reasonable approximation.

Carbon circulation can be modelled structurally, by defining the different sinks and the flows between them. A prototype example of this is the carbon-cycle model in the RICE/DICE model which contains three reservoirs, representing the atmosphere (S), the biosphere and upper layers of the ocean (S^{UP}), and the deep oceans (S^{LO}). The reason for separating the upper layers of the ocean from the deep ocean is that the gas exchange between the ocean surface and the atmosphere is much faster than that within the ocean as a whole. In the simplest structural model, the flows are modelled as proportional to the size of the respective source reservoir. In discrete time, this leads to a system of linear difference equations of the form

$$\begin{aligned} S_t - S_{t-1} &= -\phi_{12} S_{t-1} + \phi_{21} S_{t-1}^{UP} + E_{t-1}, \\ S_t^{UP} - S_{t-1}^{UP} &= \phi_{12} S_{t-1} - \left(\phi_{21} + \phi_{23}\right) S_{t-1}^{UP} + \phi_{33} S_{t-1}^{LO}, \\ S_t^{LO} - S_{t-1}^{LO} &= \phi_{23} S_{t-1}^{UP} - \phi_{33} S_{t-1}^{LO}. \end{aligned} \tag{3}$$

Here, E_t is emissions into the atmosphere, adding to S_t in the first equation. The flow of carbon from the atmosphere to the biosphere and upper layers of the ocean is given by $\phi_{12} S_{t-1}$. This term reduces S_t and increases S_t^{UP}, thus coming in with a negative sign in the first equation and a positive one in the second. The other terms have analogous interpretations.

An immediate implication of modeling the carbon-cycle as a linear system is that the ratios of the sizes of the different reservoirs will be restored in the long-run whenever emissions stop. If the model is calibrated so that the three reservoirs have realistic sizes, this will imply that all but a few percent of emissions will end up in the deep oceans within a few hundred years. This is not a realistic prediction. One remedy is to make the size of deep oceans smaller. Another is to use a non-structural approach and approximate atmospheric carbon depreciation using a sum of several geometric processes with different rates. The IPCC (2007) suggests

$$d\left(s\right) = a_0 + \sum_{i=1}^{3} \left(a_i e^{-\frac{s}{\tau_i}}\right), \tag{4}$$

with $a_0 = 0.217$, $a_1 = 0.259$, $a_2 = 0.338$, $a_3 = 0.186$, $\tau_1 = 172.9$, $\tau_2 = 18.51$, and $\tau_3 = 1.186$, where s and the τ_is are measured in years.[21] With this parametrization, 50% of an emitted unit of carbon has left the atmosphere after 30 years, 75% after 356 years and 21.7% stays for ever. A similar approach is used in Golosov et al. (2014).

It is important to note that the validity of the structural and the non-structural models with constant parameters depend on the emission scenario. In particular, if emissions are very large, a larger share than 21.7% will remain in the atmosphere for a very long time.

---

[21] See IPCC (2007), table 2.14.

The carbon cycle is likely to be affected by other drivers than emissions. In particular, climate change will affect the ability of different reservoirs to hold carbon. Such mechanisms could be built into the structural model in (3) by letting the parameters be functions of, e.g., the global mean temperature. Both smooth feedbacks and drastic ones, causing thresholds and tipping points, can be included.

3.2 The climate module

The most important driver of climate change is emissions of CO_2 into the atmosphere—it is the factor that currently contributes the most to climate change, and its effects are long-lasting. The fundamental reason why atmospheric CO_2 affects the climate is that it changes the earth's energy budget. This budget is defined as the difference between the inflow of energy to, and outflow of energy from, earth, where both flows are averaged over time and space. Carbon dioxide has the property that it allows sunlight to pass through more easily than infrared radiation. Since most of the outflow of energy is in the form of infrared radiation, more CO_2 therefore implies a surplus in the energy budget and, thus, heat accumulates. This increases the temperature on earth, which in turn increases the outflow of energy until balance is restored at a new higher temperature. This basic mechanism was quantified already by Arrhenius (1896).

The simplest representation of this mechanism modelled in discrete time is

$$\begin{aligned} F_t &= \frac{\eta}{\log 2} \log\left(\frac{S_t}{\bar{S}}\right) \\ T_t - T_{t-1} &= \sigma\left(F_{t-1} - \kappa T_{t-1}\right). \end{aligned} \tag{5}$$

In the first equation, F_t is the perturbation in the energy budget relative to the pre-emission (pre-industrial) steady state, often called forcing, S_t is the amount of carbon in the atmosphere, and $\bar{S}$ is the pre-emission amount of atmospheric carbon. The parameter η determines the strength of the greenhouse effect. Specifically, a doubling of the CO_2 concentration, i.e., $\frac{S_t}{\bar{S}} = 2$, yields a forcing of η (W/m^2).

In the second equation, T_t is the global mean temperature deviation from the pre-emission steady state. The left-hand side is the change in the global mean temperature per unit of time. The two terms in parenthesis on the right-hand side come from the the energy budget, consisting of forcing and the term κT_t. The latter is a linear approximation of the so-called Planck feedback: the fact that hotter objects emit more heat radiation. The parameter σ determines the speed at which the temperature increases for a given surplus in the energy budget. Given a constant forcing F, the temperature has to reach $\frac{F}{\kappa}$ for a new steady state to arise. It is also immediate to see that a doubling of the atmospheric carbon concentration leads to a new steady state with a global mean temperature of $T = \frac{\eta}{\kappa}$. This value is often referred to as the (equilibrium) climate sensitivity. According to the IPCC, the climate sensitivity is "likely in the range 1.5 to 4.5°C", "extremely unlikely less than 1°C", and "very unlikely greater than 6°C".[22]

---

[22] The statement is taken from IPCC, 2013a, page 81 and IPCC, 2013b, Box 12.1. The report states that

The simplest model can be extended to include several energy budgets. The DICE/RICE model due to Nordhaus, for example, also includes a budget representing the flows of energy between the atmosphere and the oceans. Since the oceans have a much larger heat capacity than the atmosphere (more energy is required for a given increase in temperature), they will experience a slower increase in temperature for a given forcing. Additional forcing variables, like methane and particle emissions, can be added to the first equation and potentially be made contingent on, e.g., the temperature. It is also straightforward to allow non-constant parameters. For example, the feedback parameter κ could be made dependent on current and/or past temperature, capturing threshold and tipping-point effects. An example of a threshold effect would be to add the equation

$$\kappa = \begin{cases} \kappa_0 \text{ if } T_t < T_{th} \\ \kappa_1 \text{ else} \end{cases}$$

to (5), where T_{th} is the threshold temperature and $\kappa_1 < \kappa_0$. If instead the switch to κ_1 is permanent and occurs the first time period such that $T_t \geq T_{th}$, then we have an irreversible tipping point.

Clearly, there is a lot of uncertainty around all the parameters in this and similar models. Specifically, uncertainty about the parameter κ is important. First, uncertainty about κ implies uncertainty about the climate sensitivity. Second, since κ enters in the denominator of the expression for the climate sensitivity, a symmetric distribution for κ around a mean would imply that the distribution of the climate sensitivity is skewed to the right (Weitzman, 2011).

Energy-budget models usually do not have a geographic dimension. However, there is a systematic relation between regional climate change and the change in the global mean temperature. Statistical methods can be used to infer this from historical data or from advanced climate models. Such *statistical downscaling* can be used to predict regional climate change from the global mean temperature.

3.3 Constant Carbon-Climate Response

A highly tractable way of representing both carbon circulation and climate change jointly has been proposed by Matthews et al. (2009). They show that several of the dynamic and non-linear mechanisms described above tend to approximately cancel in a very convenient way. In concrete terms, a reasonable approximation to the dynamic relation between the global mean temperature and CO_2 emissions is that the temperature increase over any time period is proportional to the accumulated emissions over the same period. Furthermore, according to the approximation, the proportionality factor (denoted CCR) is independent of the length of the time period and of previous emissions.

$$T_{t+m} - T_t = \text{CCR} \sum_{s=t}^{t+m-1} E_s.$$

"likely" should be taken to mean a probability of 66-100%, "extremely unlikely" 0-5% and "very unlikely" 0-10%.

To obtain some understanding for this surprising result, first note that when oceans are included in the energy-budget model, there is a substantial delay in the temperature response of a given forcing. Second, if carbon is released into the atmosphere, a large share of it is removed quite slowly from the atmosphere. It happens to be the case that these dynamics approximately cancel each other out, at least if the time scale is from a decade up to a millennium. Thus, in the shorter run, the CO_2 concentration, and thus forcing, is higher, but this effect is balanced out by the cooling effect of the oceans.

Second, note that the Arrenhius law discussed in the previous chapter implies a logarithmic relation between CO_2 concentration and forcing. Thus, at higher CO_2 concentrations, an increase in the CO_2 concentration has a smaller effect on the temperature. On the other hand, existing carbon-cycle models tends to have the property that the storage capacity of the sinks diminish as more CO_2 is released into the atmosphere. These effects also approximately balance: at higher levels of CO_2 concentration, an additional unit of emissions increases the CO_2 concentration more but the effect of CO_2 concentration on temperature is lower by about the same proportion.

Matthews et al. (2009) argue that both model simulations and historical data suggest a best estimate of CCR of 1.5 degrees Celsius warming per 1,000Gt of carbon emissions. They also derive a 95% confidence interval for the CCR being between 1 and $2.1°C/1000GtC$.

4 Damages

We now describe how the economy is affected by climate change. This description thus closes the loop from the economy, which generates emissions that enter the carbon cycle and drive climate change, back to itself. Nordhaus (1994) pioneered the "bottom-up" approach to aggregating damages. His idea was to compile a large number of microeconomic studies on various consequences of climate change, e.g., negative effects on agriculture, coastal damages, lowered amenity values, worsened health, and low-probability catastrophic damages to the overall economy. In these studies, the common problem in environmental economics of valuing effects that have no, or very imperfectly measured, market prices is particularly salient.

Nordhaus (1994) constructed estimates of damages in 13 different regions of the world, allowing region- and mechanism-specific functional forms. Different ways of estimating these damages have been used; of particular interest is perhaps the "Ricardian" approach used in Mendelsohn, Nordhaus, and Shaw (1994) which estimates the relation between temperature and market prices of farm land across 3,000 U.S. counties with the idea that institutions are very similar across these locations but temperatures are not. All the different types of damages were then aggregated into region-specific (RICE) and a global (DICE) damage functions mapping the increase in global mean temperature over the pre-industrial level (T_t) into damages expressed as a share of current GDP. Given these estimates, a function $\Omega(T_t)$—representing the share of GDP that remains after climate damages—can be derived. Using a second-order approximation, this function is expressed as

$$\Omega\left(T_t\right) \equiv \frac{1}{1+\theta_1 T_t+\theta_2 T_t^2}, \tag{6}$$

with parameters θ_1 and θ_2 chosen so as to make the damage function approximate the sum of the underlying damage estimates. Obviously, higher-order terms can easily be included to increase the convexity of damages. Equally obviously, great care has to be taken when interpreting results that rely on extrapolations of the damage functions outside of the range over which it is estimated.

An alternative and complementary way of estimating the aggregate effects of climate on economic activity is to use reduced-form relations in data on economic outcomes and temperature. Here, both time variation and regional variation have been used to draw inferences about the effects of climate change. Regarding the former, Dell, Olken, and Jones (2014) summarizes the literature which uses natural variation in temperature and climate characteristics to identify effects on aggregate economic variables. They conclude that in poor countries, losses on the order of 1-2% per degree Celsius are typically found for output, labor productivity, and economic growth. These effects are identified using temporary changes in temperature and are arguably well identified short-run effects also of climate change. However, the authors caution against also inferring that the effects are permanent. This warning seems particularly relevant when it applies to growth rates of output and other economic variables.

There is also a systematic relation between geographic variation in temperature and economic output. Nordhaus (2006) uses data on output for 25 thousand 1 by 1 degree terrestrial grid cells and shows that there is a clear hump-shaped pattern between temperature and output per km^2. The peak of the hump, with the highest average output per km^2, was found at approximately 12 degrees Celsius.[23] Under the assumption that the relation between temperature and output is invariant it can used to infer the effects of climate change on global GDP. The estimates in Nordhaus (2006) indicate losses on the order of a few percent of GDP if the global mean temperature increases by three degrees. In contrast to the estimates using time variation, these effects do not suffer from the problem of being identified from short-run variation. More recent studies find hump-shaped patterns across regions for growth rates, indicating extremely strong long-run effects of climate on economic activity; see, e.g., Burke, Hsiang, and Miguel (2015).[24]

Let us finally describe a representation of the combined mapping from atmospheric CO_2 concentration via climate change to damages. Gosolov et al. (2014) show that an exponential damage function where the argument is the excess amount of carbon in the atmosphere $(S_t - \bar{S})$, rather than temperature, is a reasonably good approximation to simple climate models and damage functions as in (5) and (6). In their formulation, damages as a share of GDP before damages are given by

$$\Omega_s\left(S\right) = 1 - e^{-\gamma(S_t - \bar{S})},$$

---

[23] Interestingly, the relation between output per person and temperature is monotone and negative.

[24] These estimates are not consistent with a fairly stable distribution of GDP across regions and are hence hard to square with historical data.

where γ is a constant. This implies that the share of GDP lost per unit of carbon in the atmosphere is constant at γ.[25] This formulation is convenient for a number of applications. For example, under some additional assumptions, one can derive a very simple formula for the optimal tax rate on carbon; see Section 5.5.1 below.

5 A complete, quantitative IAM

We are now ready to formulate the first full macroeconomic model of climate change. It has one region only and a minimum of heterogeneity in other dimensions too. It is however, a framework that can be straightforwardly built on further, along the lines of the many branches of the macroeconomic literature—including consumer heterogeneity, multiple regions, and so on—and we will briefly look at examples of such cases below—without losing its quantitative anchoring in historical data: summary climate and growth facts. We will use a benchmark model with logarithmic utility, and hence there is risk aversion (to the extent there is uncertainty) but its level is moderate. Again, it is straightforward to incorporate higher curvature/Epstein-Zin preferences into this framework (see, e.g., Jensen and Traeger, 2014).

We first formulate a planning problem and then look at a competitive equilibrium. We also state a formula for the social cost of carbon based on Pigou (1920). At the end of this section, we discuss how to solve the model computationally.

5.1 The planning problem

We focus on the case of exogenous technical change and consider a long enough time horizon that a Cobb-Douglas production function is appropriate. Energy, in the present formulation, is a composite of several sources.

$$\max_{\left\{c_t,k_{t+1},E_t,E_{ot},E_{ct},E_{gt},h_{ct},h_{gt},S_t\right\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \log c_t$$

s.t.

$$c_t + k_{t+1} = e^{-\gamma_t S_t} k_t^{\alpha} \left(A_t(h - h_{ct} - h_{et})\right)^{1-\alpha-\nu} E_t^{\nu} + (1-\delta)k_t \quad \forall t,$$

$$E_t = \left(\kappa_o E_{ot}^{\rho} + \kappa_c E_{ct}^{\rho} + \kappa_g E_{gt}^{\rho}\right)^{1/\rho} \quad \forall t,$$

$$\sum_{t=0}^{\infty} E_{ot} \le R_0, \quad E_{ct} = A_{ct} h_{ct} \quad \forall t, \quad E_{gt} = A_{gt} h_{gt} \quad \forall t$$

and

$$S_t = \sum_{j=0}^{\infty} (1-d_j)(E_{o,t-j} + E_{c,t-j}) \quad \forall t.$$

---

[25] $\frac{\partial(\Omega_s(S)Y_t)/\partial S_t}{(1-\Omega_s(S))Y_r} = \gamma.$

Some comments are in order. First, in this formulation we opt for damages to be a function directly of the carbon concentration, hence bypassing temperature as a driver of this mechanism. An alternative we discussed above, and which will deliver similar quantitative conclusions, is to instead bypass carbon concentration and express damages as a function of temperature and temperature as the CCR function of total past emissions (undepreciated). Moreover, we let the damages be random through the dependence of γ on t. Second, the SEC energy composite contains the share parameters κ, which should be calibrated to reflect the relative efficiency with which the different energy sources are used in production. In particular, coal is "dirtier" than oil in that it gives rise to higher carbon emissions per unit of energy services. We let E_{ot} and E_{ct} have units measured in carbon, hence implying that $\kappa_o > \kappa_c$ would be satisfied in a calibration of these parameters. Third, notice that there is technical change in production of final output as well as in the production of coal and green energy. The energy-saving technology variable A_e is not needed here, since the production function is Cobb-Douglas (A_e can be viewed as a part of A).

We discuss how to solve this planning problem below. Note here that a solution entails whether or not to use up all the oil; oil is freely available but now—in contrast to above, where there was no negative effect of emissions on TFP—oil use has societal costs.

5.2 Market equilibrium

The definition of a sequential equilibrium closely follows that in Section 2.2.2. Here, a key point to notice, of course, is that no individual decision maker—consumer or firm—internalizes the effects of their decisions on aggregates. In particular, no consumer or firm internalizes their negative effect of the use of fossil fuel—a consumer selling E_o or a coal producer selling E_c—on TFP, via an increase in atmospheric carbon concentration (and hence warming).

An equilibrium is thus mathematically formulated as a set of stochastic sequences $\{c_t, k_{t+1}, E_t, E_{ot}, E_{ct}, E_{gt}, h_{ct}, h_{gt}, S_t, w_t, r_t, p_t, p_{ot}, p_{ct}, p_{gt}\}_{t=0}^{\infty}$ such that

1. $\{c_t, k_{t+1}, E_{ot}\}_{t=0}^{\infty}$ solves

$$\max_{\{c_t, k_{t+1}, E_{ot}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \log c_t$$

   s.t.

$$c_t + k_{t+1} = (1 - \delta + r_t)k_t + w_t h + p_{ot}(1 - \tau_t)E_t + T_t \quad \forall t$$

   and

$$\sum_{t=0}^{\infty} E_{ot} \leq R_0;$$

2. $r_t = \alpha y_t / k_t$, $w_t = (1 - \alpha - \nu)y_t/(h - h_{ct} - h_{gt})$, and $\nu p_t = y_t / E_t$, where $y_t = e^{-\gamma_t S_t} k_t^{\alpha} \left(A_t(h - h_{ct} - h_{et})\right)^{1-\alpha-\nu} E_t^{\nu}$ and $S_t = \sum_{j=0}^{\infty}(1 - d_j)(E_{o,t-j} + E_{c,t-j})$ $\forall t$;

3. $(p_t, p_{ot}, p_{ct}, p_{gt})$ satisfies

$$p_t = \left(\kappa_o^{\rho-1} p_{ot}^{\frac{-\rho}{1-\rho}} + \kappa_c^{\rho-1} p_{ct}^{\frac{-\rho}{1-\rho}} + \kappa_g^{\rho-1} p_{gt}^{\frac{-\rho}{1-\rho}}\right)^{-\frac{1-\rho}{\rho}} \qquad \forall t;$$

4. $p_{ct}(1-\tau_t)A_{ct} = w_t$ and $p_{gt}A_{gt} = w_t \ \forall t$;
5. $T_t = \tau_t(p_o t E_{ot} + p_{ct}E_{ct}) \ \forall t$; and
6. $E_t = \left(\kappa_o E_{ot}^{\rho} + \kappa_c E_{ct}^{\rho} + \kappa_g E_{gt}^{\rho}\right)^{1/\rho}$, $E_{ct} = A_{ct}h_{ct}$ and $E_{gt} = A_{gt}h_{gt} \ \forall t$.

This definition does not include a Hotelling price equation; instead, it is implied by consumer choice (as the consumer owns and manages the sale of oil over time). In particular, two different forms of saving must give the same expected, marginal-utility weighted return: $\mathbb{E}_t \frac{1+r_{t+1}-\delta}{c_{t+1}} = \mathbb{E}_t \frac{p_{o,t+1}/p_{ot}}{c_{t+1}}$. Hence a relative risk premium is involved here across the two kinds of risky assets capital and oil.

Instead of stating the profit-maximizing conditions for the different firms, we state the first-order conditions—as these are standard, and as in the definition in Section 2.2.2. The energy price index is also a result of profit maximization: it is derived from minimizing the costs $p_oE_o + p_cE_c + p_gE_g$ of producing one unit of E.

We see from the formulation of the problems that the consumer does not take into account how E_{ot} affects TFP (hence current and future prices); similarly, the first-order conditions from the coal firm's problem do not contain such effects either. However, the definition now includes taxes on fossil fuel, levied on the consumer for selling oil and on the coal producer for selling coal. Revenues are rebated back to the consumer as a lump sum.

In a laissez-faire (zero-tax) competitive equilibrium, all the oil will be used up because the consumer has no interest in selling less than the total amount, as p_{ot} (the marginal product of oil in producing output) will be positive at all times. The only way in which some oil will be left in the ground is by setting taxes so that $\tau_t = p_{ot}$ at all times. Then the consumer is indifferent as to how much to sell at each point in time (and one of the possible choices is consistent with the given p_{ot} at that time—in consistency with an equilibrium).

5.3 Model solution

The models just described are stochastic and non-linear and cannot, in general, be solved analytically. In economics there are different traditions in this regard: should one insist on models that can be solved analytically or is it acceptable to be aided by a computer? We first discuss this issue and then turn to some concrete comments regarding the "how-to" in our particular application.

5.3.1 Analytical vs. numerical model solution

The present section begins with a discussion of methods, chiefly with the purpose of providing arguments for the use of "complex" models solved with numerical methods. Then these methods are discussed briefly.

A very common view among economists is to insist that theoretical models be formulated in such a way that they have closed-form solutions, or in any case so that their properties can be ascertained analytically in theorem-proof style. The argument in favor of this approach is usually that the logic of the model is the central piece of the modeling: the "understanding" of the theory is key and cannot, according to this argument, be attained sufficiently precisely based on numerical solution. Economic models are abstractions and should not move beyond this stage.

In the approach to climate and economics pursued in this chapter, the goal is to formulate a model that can generate key features of the data. This cannot always be accomplished—rather, it is rare that it can—in models that allow analytical characterization. The idea is that if the most salient features of the data are captured correctly, the model can be used to interpret history, for prediction, or for policy analysis. The parameter selection in this procedure is often informal—and called calibration—but, conceptually, it is econometrics. An argument for the informality of the econometric procedure is that the model is not believed to be truth—many aspects of the real-world economy are abstracted from in the theory—and so the formal estimation loses an important part of its meaning (in particular the testing of hypotheses).

Which method is better suited for practical applications? Often, economic commentators and policymakers (perhaps especially those in the area of macroeconomics) express the idea that complex formal models are not useful, because they are still too simple and cannot be thought of as truth. The implication is that the preferred method for informing policy is to have a number of ideas at hand (perhaps in the form of separate, analytically tractable models) and then informally weigh the relative importance of these and, thus, arrive at a conclusion. We take issue with this view. Applied macroeconomic models, such as those discussed here, are indeed often complex enough that closed-form solutions are not available, while still being drastic simplifications of reality. We nevertheless insist that it is better to use explicit models, thereby transparently applying quantitative weights to the different theory components. The complex model should be viewed as incorporating the different theories. Calibration, or formal estimation when feasible, is the way in which the weighting scheme is formally implemented. The model's implications can still not be fully (or even half) believed, because important pieces of the theory can be missing, the parameter selection is fraught with errors, and so on. But at least this procedure is more transparent and, in particular, alternative viewpoints can be invited, in the form of introducing alternative mechanisms into the model, or even radically different models, so long as they are explicit and quantified. At the end, of course, a significant amount of subjective judgement has to be applied before actual decisions are made, or final views are formed. We believe, however, that the decision process is fundamentally more robust and less likely to deliver undesired outcomes if a quantitative model is used as in input, not least because it is then easier to carry out an evaluation ex post.

In the climate-economy area, an interdisciplinary field that in particular includes several natural-science areas, it would be particularly difficult to insist on having a set of small models and ideas and then informally weigh them together. When an engineer builds a bridge, a model is used: a quantitative model that abstracts from many aspects of reality but also captures what is key for the purpose at hand. We would not want to cross a bridge with a heavy truck if we knew that the bridge was constructed based on the informal weighing together of abstract arguments relating to the desired features of a bridge. The same goes for rocket launching and, closer to the subject here, meteorology. In all those areas, calculations are made based on explicit models, numerically solved. So, if nothing else to ease communication with climate scientists, the method here is what it has to be in order to have the credibility of a quantitatively grounded theory.

It should be stressed, of course, that complexity is not a goal in itself. In fact, the core model we analyze below is far simpler than the typical business-cycle model used in macroeconomics, because it turns out that a large degree of analytical tractability can be allowed while still accounting for the main historical facts. Some (highly relevant) extensions of this setting do necessitate advanced numerical methods, however, leaving us no alternative but to use these in such cases. The quantitative-theory route has also been Nordhaus's approach in formulating the DICE and RICE models: these models are non-linear and the non-linearities are believed to be important. In building these models, Nordhaus made an effort to summarize the appropriate natural-science modules so that the resulting IAM would be as compact as possible and so that it at least could be feasibly analyzed with numerical methods. The challenge was, in particular, that the number of state variables in typical climate and carbon-cycle models is very large and that dynamic economic models are forward-looking, thus requiring entirely different solution methods—fixed-point methods—than those used in the natural sciences.

5.3.2 How-to

There are three challenges to consider: (i) transition dynamics, (ii) non-linearities, and (iii) uncertainty. Transition dynamics, i.e., the need to study the economy's path toward a long-run steady state, or balanced growth path, is necessary in the climate-economy context because the use of fossil fuels must end and the task at hand is to analyze the path toward that long-run outcome. Non-linearities are inherent in the model formulation, although it is not clear in any given case how significant the departure from linearity (or log-linearity) will be. Uncertainty is relevant to the extent one models the global climate, or damages, this way—a reasonable approach given that there is significant uncertainty in several dimensions, as discussed above.

The numerical package that is most commonly used in macroeconomics is DYNARE. The main purposes of this use is monetary and fiscal policy analysis in the context of business-cycle fluctuations. DYNARE is potentially valuable in climate-economy applications too, because it has a module for solving for non-linear transition paths when there is no uncertainty. The procedure relies on convergence to steady state and backward solution: first a steady state is solved for and then the program simply finds a solution to a number of nonlinear equations: all the equilibrium conditions, including starting values for the state variables (in the case of the model above, K_0 and S_0). DYNARE is very convenient because it is highly automated: the researcher simply types in the equations and the parameter values, including initial conditions; the program does the rest (solves for steady state and then for dynamics).

Recurring uncertainty is more challenging to study because the number of variables explodes with the time horizon. If there is one random variables with n possible outcomes each period and the model is solved 100 periods forward, the number of variables to solve for in a one-dimensional economy like the one-sector neoclassical growth model (assuming the first random realization is one period from now) is $1 + n + n^2 + \cdots + n^{100}$—a prohibitively large number even for a coin-flip process. If the economy is (well approximated by) a linear system, however, then uncertainty can be handled with well-known methods from linear algebra: the analysis of stochastic differential (in continuous time) or difference (as here, in discrete time) equations, based on eigenvalue analysis, is straightforward and this method is also the core module on which DYNARE is built. The computation of equilibria for such economies is very fast and based on matrix manipulation. If the economy is not well approximated by a linear system, then DYNARE still offers solutions, namely higher-order approximations. The idea here is again to find a steady state in the absence of shocks and then to Taylor-expand—now to a higher order—the set of equilibrium conditions around that state. How well the resulting stochastic dynamic system behaves far from steady state is not known in any generality at all and it is therefore necessary in some applications—arguably the climate-economy case is one—to instead use "global methods", i.e., methods that do not rely on approximating the economy's behavior locally around a steady state.

Appropriate global methods under uncertainty rely, as indicated above in Section 2.2.3, on recursive methods. Again taking the stochastic neoclassical growth model as an example, the unknowns are now functions such as the value function, which are infinite-dimensional objects, but there are computationally efficient methods for solving functional equations. In a non-trivial equilibrium setting (in particular where the equilibrium is not Pareto optimal, such as under externalities in the climate context), there are few theorems to rely on but there are a number of computational algorithms that have proven robust and fast in a large range of applications. The recursive equilibrium setting described in Section 2.2.3 involves a number of functions of several variables and to solve for all of these is challenging. However, the key functions are those determining the values of the endogenous variables, G^1 and G^2, and these can be solved for directly without solving for the remaining functions. To see this, consider, for simplicity, the case where hours are given exogenously by a value h. Then the only equation to derive is $G(A, K)$: that determining saving. Then after deriving the consumer's first-order conditions (together with an envelope condition) and the remaining equilibrium conditions, one arrives at

$$\begin{aligned} u'\left((1+R(K,A)-\delta)K + W(K,A)h - G(A,K)\right) = \beta \mathbb{E}\left[\left((1+R(G(A,K),A')-\delta)\right)\cdot\right. \\ \left.\cdot u'\left((1+R(G(A,K),A')-\delta)G(A,K) + W(G(A,K),A')h - G(A',G(A,K))\right)\right] \end{aligned}$$

for all (A, K), where R and W are known functions. Here we have one functional equation in one unknown function: G. The typical approach is then to select a grid in (A, K) space and apply an algorithm involving a starting guess on the function's value at the grid points and then an updating scheme. There are a host of such procedures, and others as well, and for most problems studied in macroeconomics they work well. The challenges arise particularly when the number of endogenous functions increases and when the state space increases—especially the endogenous state space. In the climate application above—the deterministic benchmark model—the endogenous state variables are capital, the stock of undepleted oil, and the atmospheric carbon concentration and the nontrivial decision functions involve saving and the energy use of the three different kinds and it is straightforward to derive the corresponding functional equations.

We will demonstrate, in Section 5.5, that a way to calibrate the present model based on long time periods actually leads to a number of closed-form solutions and few functional equations to solve. Thus, at least in important benchmark cases, it is possible to solve the model very easily without giving up on the model's ability to replicate history quantitatively.

5.4 The social cost of carbon

Before moving to a concrete application and concrete results—the next section—we want to highlight that one important result can be derived in a rather general formulation: the optimal-tax formula. This formula applies if a carbon tax can be implemented unrestrictedly and globally and the only model distortion is the externality occurring through carbon emissions. It also applies if there are other externalities, but these are taken care of optimally by separate taxes/transfers (e.g., a research spillover that is corrected by means of appropriate subsidies to research). The tax prescription can also be used to evaluate carbon capture and storage (CCS): if the cost of such a technology, per unit of carbon, is below the optimal tax, it should be adopted—otherwise not. The tax formula does not apply in the presence of other distortionary taxes, and indeed an important aspect is second-best taxation; we do not look into it here (see, e.g., Barrage, 2014, and Schmitt, 2014, who look at models very similar to that described here).

The optimal-tax formula says that the tax on carbon, per unit of carbon, should equal the externality damage. The externality damage has many components but can be expressed with a transparent formula:

$$\tau_t = \mathbb{E}_t \sum_{j=0}^{\infty} \underbrace{\beta^j \frac{u'(c_{t+j})}{u'(c_t)}}_{\text{discounting}} \cdot \underbrace{\left(-\frac{\partial Y_{t+j}}{\partial S_{t+j}}\right)}_{\text{damage/C in atm}} \cdot \underbrace{\frac{\partial S_{t+j}}{\partial E_{i,t}}}_{\substack{\text{C left in atm} \\ \text{per emitted unit}}} , \tag{7}$$

It is straightforward, but tedious, to use insert this formula into the equilibrium conditions and verify that the implied system replicates the first-order conditions of the planner. The externality damage of a unit of emissions of fossil fuel of type i (coal or oil) is an expected discounted sum of current and future damages. The discounting involves the marginal rate of substitution of consumption units over time: β and the marginal utilities u', which vary over time with consumption. The formula contains the derivative of damages with respect to carbon concentration as well as the atmospheric carbon depreciation patterns: the further into the future, the less is left of a unit of carbon emitted now. The formula can of course be generalized in a variety of ways; one is to include damages elsewhere than to TFP—then the formula is also a sum across the different places in which damages appear.

It is important to note that the optimal-tax formula, whose value in general depends on the allocation, is to be evaluated at the optimal allocation. One can label the expression on the right-hand side of equation (7) the social cost of carbon and evaluate it at any (equilibrium) allocation. It would then express the net social cost of a marginal change in emissions today. It would not equal the optimal tax unless the allocation is already optimal—which it would not be, for example, in a laissez-faire allocation.

5.5 A Mickey-Mouse model? Quantitative analytical IAMs

As argued above, the climate-economy modeling, by virtue of trying to construct quantitatively accurate long-run models, cannot be restricted to settings with analytical solutions. However, it turns out, as shown in Golosov et al. (2014), that a "Mickey-Mouse model" of the climate-economy interactions, despite its simplicity, provides a close quantitative fit to the historical data and to models that are significantly more complex. Thus, the Mickey Mouse model can be taken as more than an abstraction: it can be thought of as a quantitative model—a quantitative analytical IAM.

To be more concrete, consider the model in Section 5 above, with logarithmic utility (again, this remains the modal curvature choice for utility in the macroeconomic literature), Cobb-Douglas production, and full depreciation of capital from period t to period $t+1$. The full-depreciation rate is altogether inappropriate for the business-cycle context where a period is usually a quarter or a year. Not only is the depreciation rate much closer to zero than to one at such horizons, but its level also matters importantly for equilibrium outcomes. Here, however, if a time period is 10 years or longer a realistic depreciation rate is much closer to zero and there is no sense in which the model behaves very differently for almost full depreciation compared to literally full depreciation. The upshot of setting $\delta = 1$ is that the saving rate becomes constant: regardless of the path for TFP, whether stochastic or deterministic, a fraction $\alpha\beta$ of output is saved—this is easily verified both for the planner's problem and the market economy.

Moreover, the remaining energy choice can be handled rather straightforwardly. First, if the model has coal (and green energy) only, as we saw above, there is a closed-form solution in competitive equilibrium. The planner's solution in this case is not as simple, as it involves a dynamic first-order condition, but it can be solved, in the absence of uncertainty, using a the "guess" that the Pigou formula works—see Golosov et al. (2014). Second, if the model has oil (in addition to, or along with, coal and a green input), then a simple shooting algorithm can be used. These models are thus possible to solve even in Excel.

We will provide some of the outputs from such a model below. First, however, let us briefly revisit the optimal tax formula.

5.5.1 The Pigou tax in the quantitative analytical IAM

We now apply the optimal-tax formula in equation (7) to the analytical IAM. It is straightforward to show that

$$\frac{\tau_t}{y_t} = \left[\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \gamma_{t+j}(1-d_j)\right].$$

Here, again, γ_{t+j} is the damage coefficient in period $t+j$ and $1-d_j$ the fraction of atmospheric carbon remaining in j periods. The remarkable aspect of this formula is that the tax relative to output is a function only of exogenous parameters. That is, neither the level of output nor any energy variables, or the level of carbon concentration in the atmosphere, appears. Let us discuss the reason behind this result and its implications.

First, a unit of carbon emissions today will give a damage that is a constant percentage of output at each future date (though different constant percentage amounts depending on the time horizon). This is, first, because of linear atmospheric depreciation rates, i.e., the share of a unit of emissions that remains in the atmosphere j periods after it was emitted only depends on j; this assumption was found above to be a good approximation. Second, each unit of additional carbon in the atmosphere causes a constant damage in percent of output, an assumption we also found to be a good approximation. Stronger non-linearities in the carbon cycle and/or the climate parts of the model will break these results and the given level of carbon concentration will matter. Tipping points will give such effects but so long as no global tipping point can be discerned, which we argued above, our assumptions are at least a reasonable baseline. It will also be violated if we depart from an exponential damage function, such as under the assumption of a constant CCR and a standard damage function of temperature. However, the departure in this case will be minor in quantitative terms.

Furthermore, a constant percentage output loss in the future will correspond to a constant percentage loss in terms of current output. To understand this, note that a higher level of output in the future of a current emitted unit means that the loss is higher, since it is operates on a higher base. But on the other hand marginal utility is lower by the same proportion, given logarithmic utility. Moreover, when discounted to today, since constant saving rates mean that consumption-output ratios remain constant, we obtain the stated result. The result will be violated if $u(c)$ is a different power function, but so long as the power/curvature is not very different from that of the logarithmic function (where the power is zero), the stated formula will be a close approximation if β is replaced by β times the gross output growth rate to the corresponding power. Thus, we note that output does not appear in the formula and, hence, size effects are absent—such as those that would appear under population growth. Population growth, and size effects more generally, would call for higher energy use, but the margin-based calculation here makes the larger size will be weighed down by a correspondingly lower effect on marginal utility.[26] For the same reason, any issues relating to green vs. fossil energy are (approximately) irrelevant in computing the optimal tax on carbon.

Let us display results according to this simple formula. The optimal carbon tax is thus contained in Figure 5.1 for three scenarios: a high damage parameter, a modest one, and a probability-weighted mean (in solid).[27] The CO_2 depreciation parameters d_s are set using three parameters only, so that a fraction 0.2 of any unit emission stays in the atmosphere forever, another fraction 0.393 exits the atmosphere immediately, and the remaining fraction decays at a rate of 0.0228 per decade.

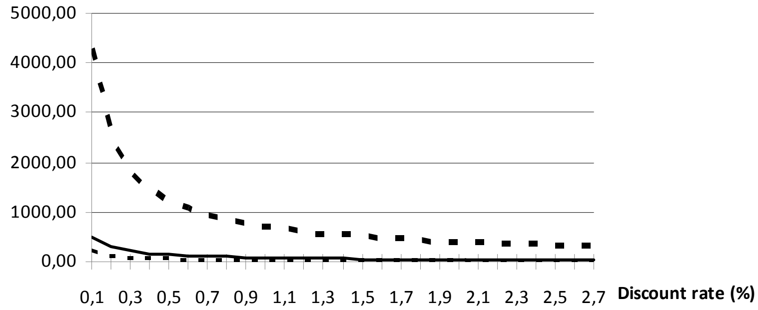

Figure 5.1: Optimal carbon tax as a function of the discount rate

We see from the figure that the optimal tax on carbon depends strongly on the discount rate: a rate of 0.1% (like that adopted by Stern in his 2006 review) gives a high carbon tax for the average damage scenario: about \$496/ton of carbon. This level is similar to (but not quite as high as) carbon tax rates in Sweden. Nordhaus has tended to focus on market rates, e.g., 1.5%, in which case the optimal tax is a magnitude higher: a little below \$60/ton.

The "catastrophe scenario" in terms of damages (where we envision a GDP flow loss of 30%—a very unlikely scenario) would roughly multiply all the taxes by a factor 10.

5.5.2 Quantitative results from the positive model

We now look at results from the calibrated economy. We will compare the optimal allocation to that obtained under laissez-faire. This comparison goes far beyond the computation of the optimal tax in the previous subsection, because it tells us the "total costs" of suboptimal policy, both in welfare terms and, for example, in terms of global temperature differences. We will also discuss the sensitivity to parameters. Rather than using

---

[26] Population growth may affect discounting directly, depending on the utility function assumed, i.e., by raising the effective β.

[27] The different γ's are 1.06, 2.38, and 20.5, all in percentage points per 1000 GtC and calculated using estimates from Nordhaus (2008).

an explicit stochastic structure here, we will follow Hassler, Krusell, and Olovsson (2018) where it is argued that the stochasticity itself is probably not key but rather the extreme outcomes. This point is also made in Weitzman (2011), but not by means of a quantitatively constructed example, which we will supply here.

The model is the analytical IAM with three carbon depreciation parameters: the parameter vector uses an interest-rate based β and an average damage coefficient estimate. α is set to 0.3 here and ν to 0.04, whereas the growth rates of both coal- and green-production technologies are set at 2% per year. The stock of conventional oil is calibrated to data and the κs and ρ to match estimates of the elasticity of substitution between different energy sources and current relative prices.[28]

Figure 5.2, first, shows that laissez-faire entails much higher fossil-fuel use than under the optimal allocation, even under Nordhaus's market-based discount rate.

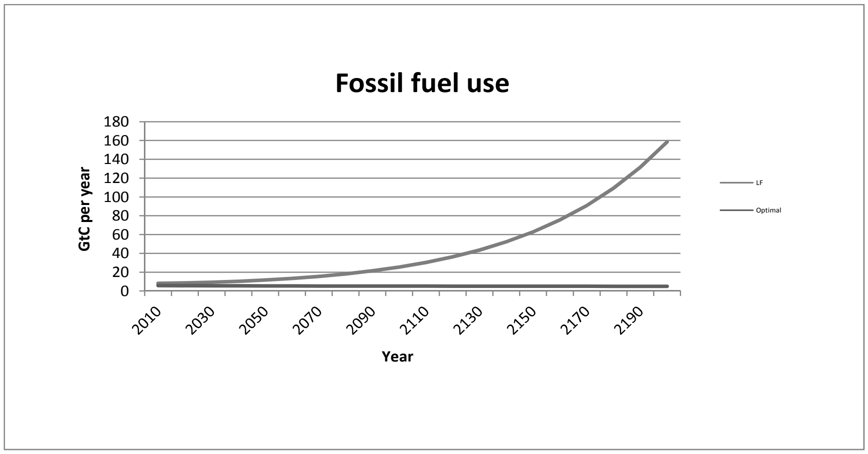

Figure 5.2: Fossil-fuel use, laissez faire vs. optimal allocation

Figure 5.3 shows how the path of (conventional) oil use ought to be altered from the perspective of our quantitative model: barely at all. Oil use should be postponed somewhat but it should be used up. The reason for this is simply that the net societal benefits to oil, given how cheap it is to produce and how efficient it is as an energy source; the postponement occurs because the timing of the damages, taking into account discounting, imply an optimal delay.

Figure 5.4, in contrast, shows how coal (including non-conventional oil) ought to be altered: very significantly. An optimal tax, even at the level implied by strong discounting, will make much of the coal stock unprofitable.

---

[28] The specific numbers are $\kappa_o = 0.543, \kappa_c = 0.102, \kappa_g = 0.357$ and $\rho = -0.058$.

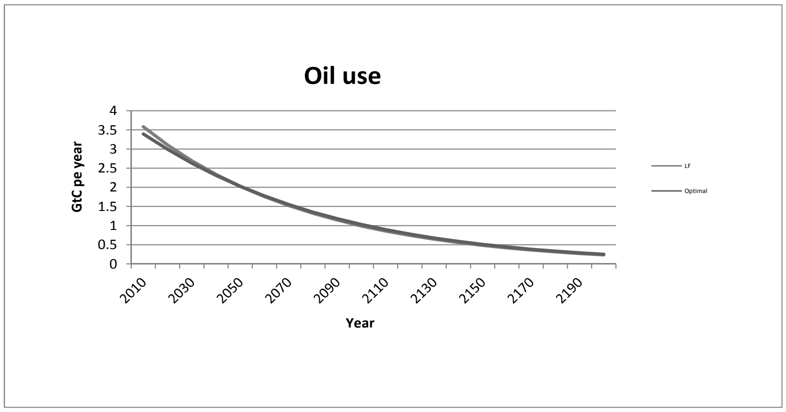

Figure 5.3: Oil use, laissez faire vs. optimal allocation

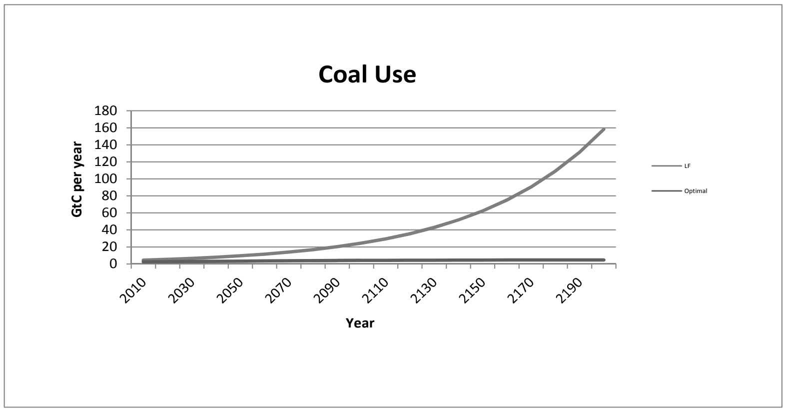

Figure 5.4: Coal use, laissez faire vs. optimal allocation

The climate damages from an optimal vs. a laissez-faire policy are depicted in Figure 5.5. Here, three lines are displayed for each of the optimal and laissez-faire outcomes, much like in the optimal-tax graph of the previous section: these cover “very high”, “average”, and “very low” damages. In Hassler et al. (2018) an argument is put forth that it is not risk aversion per se that we need to worry about but rather the extreme (but low-probability) outcomes, and for this purpose it is useful to graph these outcomes, along with the knowledge of associated probabilities. The IPCC states that the climate sensitivity likely is in the range 1.5 to 4.5 degrees Celsius. The endpoints of this range are used as extremes, although neither higher nor lower values can be ruled out. Hassler et al. (2018) use a metastudy by Nordhaus and Moffat (2017) to compute a similar range of economic sensitivities to climate change. The former also show that the upper endpoints (high climate and economic sensitivity) can be expressed as a damage elasticity of $\gamma = 10.4$. The lower endpoints, on the other hand, correspond to a damage elasticity of $\gamma = 0.27$. The moderate level is the the same as the moderate one used in the benchmark above.[29]

We see that under the mean damage scenario the optimal damages stay below a couple of percentage points of GDP in flow terms, whereas "business as usual" will imply rapidly rising costs up to levels that are over 10 percent of GDP (so a magnitude higher than today) by the end of the 23rd century. If we contemplate the very high damage scenario then the flow effects are very large: by the end of the period considered, comparable to (and even higher than) a perpetual Great Depression.

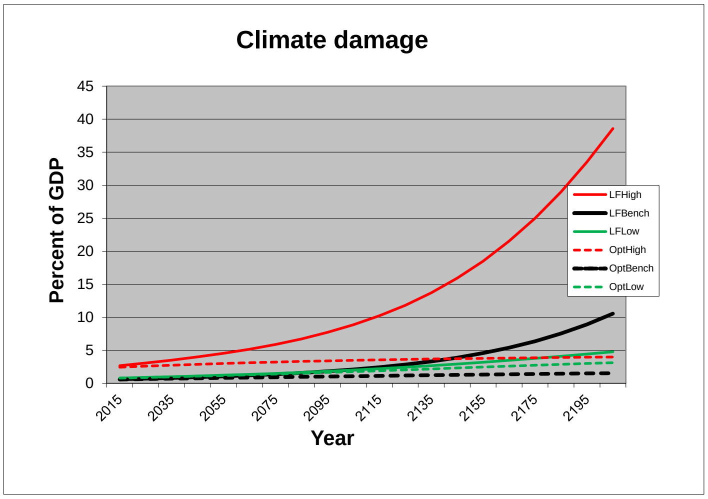

Figure 5.5: Climate damage, laissez faire vs. optimal allocation

The implications for temperature, finally, can be found in Figure 5.6, similarly for the three different damage scenarios. The temperature paths are obtained by using the model-implied carbon concentration rates and then applying Arrhenius's logarithm formula. The highest damage scenario is produced from a very high climate sensitivity; the economic damages per unit of carbon underlying this scenario also very high, actually tempering the temperature rise since lower output means less carbon use.

---

[29] See Hassler et al. (2018) for details.

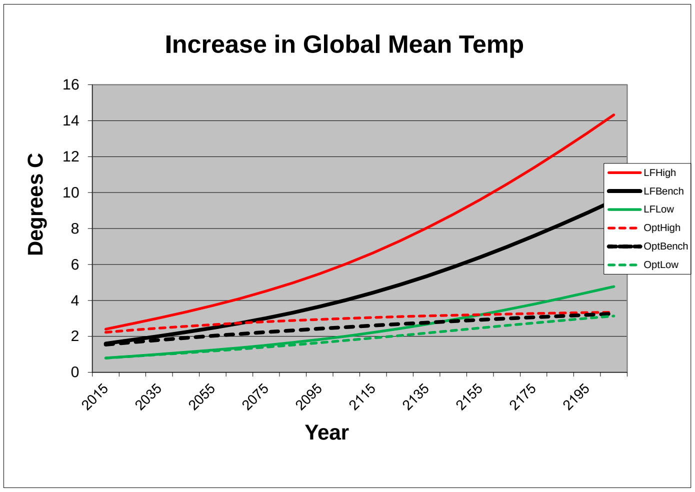

Figure 5.6: Increase in global temperature, laissez faire vs. optimal allocation

The theory's predictive features, despite being produced by our very simple model, are broadly in line with those from other models. The present treatment of temperature outcomes is particularly simple, since it rests on immediate equilibrium: there is no slow build-up, as in DICE or more elaborate climate models, e.g., because the ocean heat changes slowly. In this sense, the model overpredicts temperature increases early on. The model also does not feature tipping points, chiefly because global tipping points have not been pinned down quantitatively; the present analysis should perhaps be viewed as probability averaging over all the possible values of tipping points.[30]

6 Extensions

The literature on economics and climate change has covered many more aspects than those discussed above. Arguably the two most important ones are endogenous technical change and multiregion analysis. We now briefly discuss these, without introducing much formal detail and instead emphasizing the key focus of each of these extensions and how they could be incorporated into our model setting.

---

[30] For studies formally modeling tipping points, see, e.g., Lemoine and Traeger (2014) and Cai, Judd, and Lontzek (2015).

6.1 Endogenous technical change

For the optimal tax on carbon, it is not obvious—following the arguments above—that taking the energy market into account is critical, at least not if there are instruments to ensure that R&D in the energy sector is handled efficiently. Indeed the optimal-policy view is that one can separate the analysis in two: ensuring that there is a tax to internalize the climate damage from emissions and ensuring that there are subsidies to ensure that any spillovers in the research sector are internalized. This way of thinking even suggests that if the energy technology spillovers are stronger for fossil technologies than for non-fossil technologies, there should be a higher subsidy to fossil technology development! However, as soon as one contemplates the possibility that an optimal carbon tax cannot be implemented globally, let alone in most of our large economic regions, R&D policy in the energy sector can become very important and, indeed, take center stage.

In this short discussion, we will not have the ambition to arrive at a quantitative model; hence, we will also not be able to provide useful quantitative guidance in this important policy arena. We will instead (i) indicate how the model above can be, and has been, usefully extended to incorporate technology choice and (ii) make some comments about challenges and open questions in this area. Before addressing these issues, however, let us make reference to Hassler et al. (2018b), which looks at climate, damage, and welfare outcomes for different scenarios for technical change occurring in an exogenous fashion: the outcomes for technology are quantitatively very important for the future of the climate. One point is perhaps obvious: with a powerful green alternative, fossil fuels can be abandoned soon at low cost, and may even be abandoned without any policy intervention. However obvious the point is, it comes out clearly and quantitatively in a model essentially of the kind considered above when different relative growth rates of the coal- and green-energy production technologies are entertained.

What are the technology potentials of green vs. fossil energy production? That is a key question and the literature so far has not had much to say on it. Fundamentally, of course, predicting the future course of technical change is very hard. There have been some attempts to formalize these questions in macroeconomic settings. The most well-known piece is Acemoglu et al. (2012), which uses endogenous directed technical change à la Acemoglu (1998), a paper that in turn applied the market-based endogenous-growth analyses from the seminal contributions by Romer (1990) and Aghion and Howitt (1992). Acemoglu et al. (2012) is stylized: it is not a quantitative paper. It looks at "clean" vs. dirty goods and has endogenous technologies for each of these. The negative effects of dirty goods are modeled in a reduced-form way in such a way as to generate a "true disaster" outcome if the atmospheric carbon concentration rises high enough. Hence, the focus of the technology policy analysis is to discuss how such a disaster outcome can be avoided. The paper offers a nice and potentially important analytical insight, which is that a temporary subsidy to green technology can be sufficient for the purpose of avoiding the disaster. The reason is that the accumulation equations for technology (i) have the usual unit roots and (ii) are not interdependent. Hence, if a temporary policy can push the green technology ahead of the fossil one, it is possible—under conditions of sufficient substitutability between clean and dirty goods—that dirty goods lose importance over time and vanish from the economy.

Quantitative-theory analysis of endogenous directed technical change is most easily carried out by (i) first identifying the possible ways in which technology matters from the perspective of a model that maps to the available data and (ii) then estimating the necessary functional forms specifying the extent to which, and how, technology is subject to choice. The first of these is more straightforward than the second. Thus, rather than relying on abstract notions like dirty and clean goods—a classification that does not line up in an obvious way with the national accounts—one could simply identify energy-saving technology, on the demand side, and energy-producing technology, on the supply side. These are already expressed in the model above: alongside the standard capital/labor-saving technology A, A_e corresponds to general energy saving (and of course allows several notions of energy saving, both by simply reducing wasted energy and switching toward goods that are less energy-intensive), and A_c and A_g are the technologies for producing coal and green energy, respectively. It is, under some additional functional-form assumptions possible to back these technology variables out from the relevant aggregate data, in a manner similar to growth accounting. In particular, Hassler, Krusell, and Olovsson (2017) use a nested CES and arrive at estimates for A and A_e. These can then, in turn, be related to each other and one can obtain some information about possible past tradeoffs, at least. It turns out that there is a significant and negative medium-run correlation between A and A_e, suggesting a tradeoff like that considered in the literature on directed technical change. Whether these variables develop under a unit root-like structure or are subject to decreasing returns is another important question and one that is far from settled so far. It would be interesting to consider similar aggregate-style estimates of A_c vs. A_g. A number of microeconomic studies are available too, e.g., Aghion et al. (2017) recently and earlier those discussed in Popp (2002), though it is a challenge so far to translate scattered results from very specific studies into parametric assumptions for our aggregate IAM.

The specific modeling of directed can take different forms. One is that in Acemoglu et al. (2012), which likes most of the endogenous R&D literature rests on modeling competitive patent races (in the quality and/or variety dimensions). There, the dynamic returns to innovation are usually modeled as spillovers—standing on the shoulders of giants for free. Another one is that in Hassler et al. (2018b), which explores more of a reduced-form setting with spillovers: at each point in time any firm, entirely under perfect competition, can choose any technology pair (A_t, A_{et}) from a technology frontier given by the past decisions of others (so under full dynamic spillovers here as well). I.e., the constraint reads $G(A_t/A_{t-1}, A_{et}/A_{e,t-1}) = 0$, where G thus describes the frontier. Under appropriate assumptions on G and $F(Ak^{\alpha}, A_e E)$, a competitive equilibrium exists here, having all firms choose the same technology at each point in time. Here, the arguments of G are growth rates; they could more generally be formulated in levels, or at least without long-run growth effects. Interestingly, when F is CES and G is log-linear, the reduced-form production function—after maximizing over technology—is Cobb-Douglas in the basic inputs k^{α} and E. This general formulation can thus capture a low short-run substitution elasticity (once technology is fixed) and higher (in the special parametric case, unitary) long-run substitution.

In terms of policy prescriptions, although research spillovers generally call for interventions, the direction of these interventions is not clear. If the overall research efforts are given and society can only choose their direction, then it may well be that no policy should be undertaken: subsidizing A, in the above example, is good from the perspective of internalizing the externality to A accumulation, but it is by the same token harming the spillover in the A_e direction. Indeed, in the benchmark case of Hassler, Krusell, and Olovsson (2017), no subsidy at all is optimal. Of course, if a climate externality is present (it is not in that model) and is not appropriately handled by a carbon tax, then subsidies to A_e innovation would be called for.

6.2 Multi-region modeling

When climate damages are concerned, as briefly mentioned in Section 4 above, what stands out is the heterogeneity in damages rather than a high global average. For many regions, damages are negative, i.e., a warmer climate is expected to improve human welfare, and the differences in outcomes between regions vulnerable to climate change and those benefitting from is are enormous by most measures. Thus, in some sense, what is really needed is an IAM that allows us to study regional impacts. Another reason for using a framework with many regions is that policy analysis in practice is far more difficult as policy is not coordinated (or at least not well, as evidenced by repeated summits without much concrete progress). Thus, positive analysis of policy combinations like "EU and China uses a carbon tax and the rest of the world does not" really demand a global IAM with explicit regions with different policies.

Early on, Nordhaus developed RICE for this reason, and other researcher developed other multi-region settings. Here let us discuss how the present core model can be, and has been, extended to allow multi-regional analysis. First, Krusell and Smith (2016) constructed what is essentially a many-region version of the above setup, with "many" as in around 20,000, i.e., regions defined by 1-by-1 degree squares on the global map. Each region then runs its own production function and has its own energy supply—say, coal—and the only global market is the market for saving: there is a global general equilibrium determining the world real interest rate. Versions of the model also allow idiosyncratic (region-specific) climate/weather shocks that the region may want to insure against. The model is solved based on techniques from the macroeconomic literature on consumer heterogeneity and saving under idiosyncratic risks; here one region takes the place of one consumer in that literature (see Aiyagari, 1994, and Krusell and Smith, 1998). Thus, it builds on dynamic programming, as described above, where K, the aggregate state variable, is now replaced by some representation of the distribution of world capital across regions. Migration is not allowed in the baseline version, but it is straightforward to allow frictionless migration across neighboring regions (as in regions within a country or collection of countries such as the EU). In the multiregion model, countries differ in three dimensions: TFP (permanent differences are assumed and the calibration of these differences is based on matching initial output-per-capita differences across countries), differences in initial capital stocks (calibrated so that the marginal product of capital is the same initially across countries), and differential sensitivities to global temperature (for example, further away from the equator the responses to a given degree of global warming are stronger, i.e., there is more than one degree of warming; calibration of these sensitivities is based on "statistical downscaling", i.e., regressions based on simulations of large climate models). Finally, in this work, damages are assumed to occur in the form of a TFP drag, as above, that simply uses a common, U-shaped function of local temperature. I.e., warming makes matter better (worse) for regions with local temperatures to the left (right) of the U's minimum, since there the drag is declining (rising) with local warming. Clearly, for a region that starts far to the right of the U's minimum, further warming can be very damaging. Indeed, when the world economy based on this IAM is solved and simulated forward, the effects of global climate change, while moderate on average, are disastrously bad for some regions and positive for many other regions, with only slight majority losing. The regions that stand much to lose are typically developing economies that are already at a very low standard of living in relative terms. This kind of IAM thus highlights the distributional impacts and the needs for policies that take this heterogeneity into account. Further work along these lines appears urgent, in particular when it comes to better understanding the vulnerabilities of the developing world and the potentials for adaptation policy there.

A closely related approach is that in Hassler and Krusell (2012) and, in a more recent rendition, Hassler et al. (2018b). The idea here is again that an analytical setting offers a near-equivalent setting to the more elaborate one just discussed. The key simplification here is financial autarky: there is no world market for loans. In the Krusell and Smith (2016) approach just discussed, a world market for loans does make a difference compared to autarky, but the difference is rather minor. Hence, this is a case for autarky as a good approximation. Financial autarky, along with logarithmic utility, Cobb-Douglas production, and full depreciation again delivers constant saving rates. Hassler and Krusell (2012) also considers a world market for oil (Krusell and Smith, 2016, do not) using the simplifying assumption that there is a separate region of oil countries and that these countries cannot invest abroad. It is easy to show that if the oil economies only have income from oil and have logarithmic utility of consumption, they will simply run down the oil supply at a constant rate (equal to their discount rate). Hence, oil quantities in the world are entirely supply-driven (the income and substitution effects from oil-price changes cancel for the producer). The prices are demand-driven: the price of oil has to equal the marginal product of oil in each country (taking taxes into account, if such taxes are levied). Given that the global oil quantity is determined from supply, the equilibrium price is easily solved for period by period. The model is a very convenient vehicle for studying a range of issues, such as policy differences across countries (in Hassler and Krusell, 2012) and endogenous technical change and international R&D spillovers. Given its tractability, versions with many regions can be considered as well, allowing special focus on developing countries and their special features.[31]

A first-order issue in the climate area not yet studied much in the literature is migration. Allowing for migration, on the one hand, can limit the damages from climate change: given enough time—and climate change is expected to be slow—moving toward cooler areas is one way to adapt. Hence models building on an absence of migration can overstate the negative impacts of climate change. On the other hand, pressures for migration can lead to high costs in the form of social and, perhaps, armed conflict, as populations move across regions and borders. Solving models that allow for migration at a cost is a challenge; very promising approaches forward include Desmet and Rossi-Hansberg (2015) and Brock, Engstrom, and Xepapadeas (2014).

6.2.1 Leakage

One of the important policy issues in this area concerns the possible "leakage" of economic activity that would occur if a uniform global policy is not adopted and instead only a subset of countries/regions impose significant restrictions or taxes on carbon use. Let us now briefly discuss this point and only from a conceptual perspective. From a quantitative perspective, it seems to us too early to be able to draw firm conclusions. There are some specific microeconomic studies on the topic, but the body of empirical evidence is quite scarce; and as indicated above, though holding much promise in general and for this topic in particular, the quantitative-theory literature looking at country/region heterogeneity is still in its infancy.[32]

Leakages can take on different forms. One is that, for a high carbon-tax area, production moves abroad. Or, put differently, costs matter in an industry competing across borders, and so even if no firms literally move across borders, sales or output levels are expected to react as cost structures change. This point is well known and the quantitative effects depend on transportation costs, the energy share, the existence of substitute goods or services, the returns to scale, and so on. Thus, a range of specific factors that differ by industry will matter and we are far from an overall, macroeconomic assessment of these effects.

Another form of leakage that is perhaps less discussed is that occurring on the supply side of the energy market. This effect is present also if the use of energy is not subject to trade across borders, so long as the energy itself can move. Take transportation services as an example. As a given country raises its tax on diesel and gasoline, at some point the transportation structure will shift away from fossil-based transportation. In many countries, in fact, such a policy is high on the public-policy agenda; in Sweden, broad "political commitments" have already been made to reduce and eliminate fossil-based transportation going forward, according to a specific time schedule. What are the likely effects of these commitments, beyond the possible effect on transportation costs to the industries in the countries where the policy is undertaken? Given that gasoline and diesel

---

[31] In ongoing work we explore a global model of this sort, paying special attention to agriculture and adaptation in countries particularly vulnerable to climate change.

[32] For the microempirical approach, see, e.g., Fowlie, Reguant, and Ryan (2016) and the discussion therein.

are traded in world markets and imported, typically from far away at relatively low costs, the effects are—if a set of other countries do not impose similar taxes—that a share of the fossil fuels will be used elsewhere. Reduced demand in a set of countries will in general lead to a reduction in the world market equilibrium price and a lower equilibrium production of fuel. If the global supply elasticity is low, the world market price has to fall so as to (largely) compensate the decrease in demand in the taxing countries by increased consumption elsewhere and at a later time. Under the assumption—the one entertained in our simple model above—that there is an amount of oil in finite supply and that its production costs are negligible, the long-run supply elasticity of oil is zero and leakage of climate policy is complete. In the case of coal, the situation is quite different: here, the marginal cost coal is close to its price and moving coal across space is costly. This implies a large supply elasticity. Hence, leakage is much less of a concern in the case of coal. In sum, fossil-fuel leakage of this sort can be very important to take into account also for understanding the impact on the climate and not just for understanding the effects on local industry. The key features to look for in terms of a quantitative assessment is the transportation costs of the fuel and the marginal cost of its production relative to its price, since these factors are important for the supply elasticity. Another potentially important consideration is that the long-run supply elasticity may increase due to endogenous technical change. Suppose that a set of countries manages to reduce the world market price of oil price by introducing climate policies. This would reduce the incentive to develop technologies for extracting marginal fossil-fuel reserves, hence reducing the leakage from climate policy.

7 Concluding remarks

This chapter has focused rather singularly on the construction of a global integrated assessment model on the basis of what can be labeled quantitative theory. The text therefore started with a long, motivating section displaying the facts and pointing to the need for a framework that can account for these facts. The next section introduced the neoclassical growth model developed by Solow and extended, with optimal-saving and optimal-work decisions by households and profit-maximizing input demands by firms in a market context, so as to also address price facts, such as those on cost shares. This involved a fair amount of theory and methods discussion, after which the natural-science elements were briefly introduced and then the final integrated assessment model presented and analyzed. Section 6 discussed two important extensions—one to endogenous technical change in the energy sector and another to multi-region models—but many other topics were left without discussion. For example, an important policy issue is whether quotas (along with trade in permits) can have advantages over taxes. In models where quotas and taxes can be changed over time and in response to various shocks, the model would treat the two instruments as identical; in a richer, real-world context, they are of course not.[33] Relatedly, there is a very important political-economy element in this area: how are policies chosen endogenously, and what are good institutions from the perspective of the challenges of global warming? These issues are also not discussed here, in particular because quantitative-theory models of endogenous policy are rare even in the context of standard macroeconomic policy (such as tax choice). Thus, the focus on a quantitative-theory approach has been a guiding principle throughout the chapter, and this approach was briefly defended in the introduction as well as in Section 5.3 of the paper.

From this perspective, the reader may wonder about the tension between the approach advocated here and the view given in Pindyck's critique of integrated assessment modeling (Pindyck, xyz). It would require a long discussion to address all the relevant points of contact here. In sum, our view is that the IAMs have much to offer, both in terms of accomplishments and future promise. In fact they are, to us, the only game in town. They are flexible tools that allow arguments to play out quantitatively. If one doesn't like the conclusions, the natural next step is to suggest model amendments, adopt the associated quantitative discipline, and examine the implied results. IAMs do not close the door to obtaining either a call for very urgent and significant action to combat climate change or to the conclusion that no significant action is needed. The quantification is what allows us to obtain an answer. Moreover, the answer is by no means trivial. For the optimal social cost of carbon, we argue, subject to some qualifications, that the key parameters are three (governing utility discounting, damages, and carbon depreciation)—and their respective impacts are nonlinear. For the quantitative effects of no action, or partial action, many more aspects of the IAM become crucial, such as the future of technology, the substitutability between different energy sources, population growth, and so on. Constructing an explicit IAM will, more than anything else, force transparency in this area and this is our main argument behind why we advocate their construction and practical use. Finally, IAMs can be used to answer questions like "what is the most efficient way of making sure that a given temperature target is not exceeded?", which represents a more modest view on what we can say about what is optimal in a broader sense.

Finally, let us connect back to the introduction where we argued that climate change should be regarded as a first-order issue for macroeconomists and ask whether there might be other environmental challenges that of equal, or even greater, importance from a global perspective. At this point in time, it is hard to argue that another, equally important challenge can be identified, not because we know that such challenges do not exist but rather because the evidence and scientific studies of this evidence (in natural as well as social sciences) are very incomplete compared to what we know about the climate. One concern is the broad issue of sustainability of the world's resources, defined to include natural non-renewable resources in finite supply as well as water supply and other strains on our planet. We find this area potentially very important, at the same time as it seems difficult to refer to direct, quantitative evidence of scarcity that is as "binding" as is that on our climate (for an early discussion that concluded that we were far from binding constraints at least in some dimensions, see Nordhaus, 1974). Moreover, it is much less

---

[33] Hassler, Krusell, and Nycander (2016) discusses practical policy a bit more.

clear what the market failures are in these contexts than in the climate area. Having said this, and as a final personal note, this topic is high on our own research agenda.

References

Acemoglu, Daron (1998), "Why do new technologies complement skills? directed technical change and wage inequality", *Quarterly Journal of Economics* 113:4, 1055-1089.

Acemoglu, Daron, Philippe Aghion, Leonard Bursztyn, and David Hémous (2012), "The Environment and Directed Technical Change", *American Economic Review*, 102 (1), 131–166.

Aghion, Philippe and Peter Howitt (1992), "A Model of Growth Through Creative Destruction", *Econometrica* 60:2, 323–351.

Aghion, Philippe, Antoine Deschezlepretre, David Hémous, Ralf Martin, and John Van Reenen (2017), "Carbon Taxes, Path Dependency and Directed Technical Change: Evidence from the Auto Industry", forthcoming in *Journal of Political Economy*.

Aiyagari, Rao (1994), "Uninsured Idiosyncratic Risk and Aggregate Saving", *Quarterly Journal of Economics* 109:3, 659–684.

Arrhenius, Svante (1896), "On the influence of carbonic acid in the air upon the temperature of the ground", *Philosophical Magazine and Journal of Science* 41:5, 237–276.

Baqaee, David, and Emmanuel Farhi (2017), "The macroeconomic impact of microeconomic shocks: beyond Hultens theorem", working paper.

Barrage, Lint (2014), "Optimal dynamic carbon taxes in a climate-economy model with distortionary fiscal policy", working paper.

Bergeaud, Antonin, Gilbert Cette, and Rémy Lecat (2016), "Productivity trends in advanced countries between 1890 and 2012", *Review of Income and Wealth*, 62:3, 420-444.

Boppart, Timo and Per Krusell (2018), "Labor supply in the past, present, and future: a balanced-growth perspective", working paper.

Bornstein, Gideon, Per Krusell, and Sergio Rebelo (2018), "Lags, costs, and shocks: an equilibrium model of the oil industry", working paper.

Brock, William, Gustav Engström, and Anastasios Xepapadeas (2014), "Spatial climate-economic models in the design of optimal climate policies across locations", *European Economic Review* 69, 78–103.

Burke, Marshall, Solomon Hsiang, and Edward Miguel (2015), "Climate and conflict", *Annual Review of Economics* 7, 577–617.

Cai, Yongyang, Kenneth Judd, and Thomas Lontzek (2015), "The social cost of carbon with economic and climate risks", forthcoming in the *Journal of Political Economy*.

Dasgupta, Partha and Geoffrey Heal (1974), "The Optimal Depletion of Exhaustable Resources", *Review of Economic Studies*, 41.

Dell, Melissa, Benjamin Jones, and Benjamin Olken (2014), "What do we learn from the weather? The new climate-economy literature", *Journal of Economic Literature* 52:3, 740–798.

Desmet, Klaus and Esteban Rossi-Hansberg (2015), "On the spatial economic impact of global warming", *Journal of Urban Economics* 88, 16–37.

Fowlie, Meredith, Mar Reguant, and Stephen Ryan (2016), "Measuring leakage risk", working paper.

Gennaioli, Nicola and Andrei Shleifer (2010), "What comes to mind", *Quartlery Journal of Economics* 125, 1399-1433.

Gollier, Christian (2012), *Pricing the planets future: the economics of discount- ing in an uncertain world*, Princeton University Press, Princeton.

Golosov, Michael, John Hassler, Per Krusell, and Aleh Tsyvinski (2014), "Optimal Taxes on Fossil Fuel in Equilibrium", *Econometrica* 82:1, 41–88.

Gordon, Robert (2012), "Is U.S. economic growth over? Faltering innovation confronts the six headwinds", NBER Working Paper No. 18315.

Hassler, John and Per Krusell (2012), "Economics and Climate Change: Integrated Assessment in a Multi-Region World", *Journal of European Economic Association* 10:5, 974–1000.

Hassler, John, Per Krusell, and Jonas Nycander (2016), "Climate policy", *Economic Policy* 31:87, 503-558.

Hassler, John, Per Krusell, and Conny Olovsson (2017),"Directed technical change as a response to natural-resource scarcity", working paper.

Hassler, John, Per Krusell, and Conny Olovsson (2018), "The Consequences of uncertainty: climate sensitivity and economic sensitivity to the climate", forthcoming, *Annual Review of Economics*.

Hassler, John, Per Krusell, and Conny Olovsson (2016), "Environmental macroeconomics", in *Handbook of Macroeconomics*, eds. John Taylor and Harald Uhlig, Volume 2B, 1893–2008.

Hassler, John, Per Krusell, Conny Olovsson, and Michael Reiter (2018b), "Integrated Assessment in a Multi-Region World with Multiple Energy Sources" (joint with John Hassler, Conny Olovsson, and Michael Reiter).

Heathcote, Jonathan, Kjetil Storesletten, and Giovanni Violante (2014), "Consumption and labor supply with partial insurance: an analytical framework", *American Economic Review*, 104:7, 2075–2126.

Herrendorf, Berthold, Richard Rogerson, and Akos Valentinyi (2014), "Growth and structural transformation", Chapter 6 of the *Handbook of Economic Growth*, Volume 2B, Elsevier.

Hotelling, Harold (1931), "The Economics of Exhaustible Resources", *Journal of Political Economy* 39(2): 137–175.

IEA (2015), Key World Energy Statistics 2015, International Energy Agency.

IPCC (2007), "Climate Change 2007: The Physical Science Basis", Eds. Bert Metz, Ogunlade Davidson, Peter Bosch, Rutu Dave, and Leo Meyer. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change.

IPCC (2013a), Stocker, T.F., D. Qin, G.-K. Plattner, L.V. Alexander, S.K. Allen, N.L. Bindoff, F.-M. Bréon, J.A. Church, U. Cubasch, S. Emori, P. Forster, P. Friedlingstein, N. Gillett, J.M. Gregory, D.L. Hartmann, E. Jansen, B. Kirtman, R. Knutti, K. Krishna Kumar, P. Lemke, J. Marotzke, V. Masson-Delmotte, G.A. Meehl, I.I. Mokhov, S. Piao, V. Ramaswamy, D. Randall, M. Rhein, M. Rojas, C. Sabine, D. Shindell, L.D. Talley, D.G. Vaughan and S.-P. Xie, (2013), *Technical Summary.* In: *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change* [Stocker, T.F., D. Qin, G.-K. Plattner, M. Tignor, S.K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P.M. Midgley (eds.)], Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.

IPCC (2013b), Stocker, T.F., D. Qin, G.-K. Plattner, M. Tignor, S.K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P.M. Midgley, *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change,* Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.

Jensen, Svenn and Christian Traeger (2014), "Optimal climate change mitigation under long-term growth uncertainty: stochastic integrated asssessment and analytic findings", *European Economic Review* 69, 104–125.

Kaldor, Nicholas (1957), "A model of economic growth", *Economic Journal* 67(268), 591-624.

Karabarbounis, Loukas and Neiman (2014), "Capital depreciation and labor shares around the world: measurement and implications", NBER Working Paper No. 20606.

Kongsamut, Piyabha, Sergio Rebelo, and Danyang Xie (2001), "Beyond balanced growth", *Review of Economic Studies* 68:4, 869-882.

Krusell, Per and Anthony Smith (1998), "Income and Wealth Heterogeneity in the Macroeconomy", *Journal of Political Economy*, 106, 867–896.

Krusell, Per and Anthony Smith (2016), "Climate Change Around the World", working paper.

Lemoine, Derek and Christian Traeger (2014), "Watch your step: optimal policy in a tipping climate", *American Economic Journal: Economic Policy* 6:1, 137–166.

MaCurdy, Thomas (1981), "An empirical model of labor supply in a life-cycle setting", *Journal of Political Economy*, 89:6, 1059–1085.

Matthews, H. Damon, Nathan P. Gillet, Peter A. Stott and Kirsten Zickfeld (2009), "The proportionality of global warming to cumulative carbon emissions", *Nature* 459, 829–833.

Mendelsohn, Robert, William Nordhaus, and Dai Gee Shaw (1994), "The impact of global warming on agriculture: a ricardian approach", *American Economic Review* 84:4, 753–771.

Nordhaus, William (1974), "Resources as a constraint on growth", *American Economic Review* 64:2, 22–26.

Nordhaus, William D. (1977), "Economic Growth and Climate: The Carbon Dioxide Problem", *American Economic Review Papers and Proceedings*, Vol. 67, No. 1, pp. 341–346.

Nordhaus, William (1994), *Managing the Global Commons: The Economics of Climate Change*, MIT Press, Cambridge.

Nordhaus, William (2006), "Geography and macroeconomics: New data and new findings", *Proceedings of the National Academy of Sciences* 103:10, 3510–3517.

Nordhaus, William (2008), *A Question of Balance: Weighing the Options on Global Warming Policies*, Yale University Press, New Haven, CT.

Nordhaus, William D. and Joseph Boyer (2000), *Warming the World: Economic Modeling of Global Warming*, Cambridge, MA: MIT Press.

Nordhaus, William, and Andrew Moffat (2017), "A survey of global impacts of climate change: replication, survey methods, and a statistical analysis", Cowles Foundation Discussion Paper No. 2096.

Pigou, Arthur (1920), *The Economics of Welfare*, London: MacMillan.

Pindyck, Robert (2013), "Climate change policy: what do the models tell us?", *Journal of Economic Literature* 51:3, 860-872.

Piketty, Thomas (2013), *Capital in the 21st century*, Ed. du Seuil, Paris.

Piketty, Thomas and Emmanuel Saez (2006), "The evolution of top incomes : a historical and international perspective", *American Economic Review* 96:2, 200–205.

Popp, David (2002), "Induced Innovation and Energy Prices", *American Economic Review* 92, 160–180.

Rachel, Lukasz and Thomas Smith (2015), "Secular drivers of the global real interest rate", Bank of England Working Paper No. 571.

Romer, Paul M. (1990), "Endogenous Technological Change", *Journal of Political Economy* 98:5, S71–S102.

Schmitt, Alex (2014), "Beyond Pigou: climate change mitigation, policy making and distortions", Ph.D. Thesis, IIES Monograph series No. 85, Stockholm University.

Solow, Robert (1956), "A contribution to the theory of economic growth", *Quarterly Journal of Economics*, 70: 1, 65–94.

Solow, Robert (1957), “Technical change and the aggregate production function”, *Review of Economics and Statistics* 39:3, 312–320.

Stern, David I. (2012), “Interfuel Substitution: A Meta-Analysis”, *Journal of Economic Surveys*, 26, 307–331.

Stern, Nicholas (2006), *The Economics of Climate Change: The Stern Review*, Cambridge and New York: Cambridge University Press.

Weitzman, Martin L. (2011), “Fat-Tailed Uncertainty in the Economics of Catastrophic Climate Change”, *Review of Environmental Economics and Policy*, 5:2, 275–292.

Williamson, Jeffrey (1995), “The evolution of global labor markets since 1830 background evidence and hypotheses”, *Explorations in Economic History* 32:2, 141–196.