Abstract

We discuss climate change and resource scarcity from the perspective of macroeconomic modeling and quantitative evaluation. Our focus is on climate change: we build a very simple "integrated assessment model", i.e., a model that integrates the global economy and the climate in a unified framework. Such a model has three key modules: the climate, the carbon cycle, and the economy. We provide a description of how to build tractable and yet realistic modules of the climate and the carbon cycle. The baseline economic model, then, is static but has a macroeconomic structure, i.e., it has the standard features of modern macroeconomic analysis. Thus, it is quantitatively specified and can be calibrated to obtain an approximate social cost of carbon. The static model is then used to illustrate a number of points that have been made in the broad literature on climate change. Our chapter begins, however, with a short discussion of resource scarcity—also from the perspective of standard macroeconomic modeling—offering a dynamic framework of analysis and stating the key challenges. Our last section combines resource scarcity and the integrated assessment setup within a fully dynamic general equilibrium model with uncertainty. That model delivers positive and normative quantitative implications and can be viewed as a platform for macroeconomic analysis of climate change and sustainability issues more broadly.

Environmental Macroeconomics

John Hassler[*], Per Krusell[†], and Anthony A. Smith, Jr.[‡]

March 18, 2016

[*]Hassler: Institute for International Economic Studies (IIES) and CEPR.
[†]Krusell: IIES, CEPR, and NBER.
[‡]Smith: Yale University and NBER.

1 Introduction

In this chapter we discuss climate change and resource scarcity from the perspective of macroeconomic modeling and quantitative evaluation. Our focus is to build toward an "integrated assessment model" (IAM), i.e., a model that integrates the global economy and the climate in a unified framework. The chapter is not meant to be a survey of the rather broad field defined by interconnections between climate and economics. Rather, it has a sharp focus on the use of microeconomics-based macroeconomic models in this area, parameterized to match historical data and used for positive and normative work. Our understanding of the literature is that this approach, which is now standard macroeconomic in analyses (rather broadly defined), has not been dominant in the literature focused on developing IAMs, let alone anywhere else in the climate literature. We consider it a very promising approach also for climate-economy work, however, having contributed to it recently; in fact, the treatment we offer here is naturally built up around some of our own models and substantive contributions. Although there is a risk that this fact will be interpreted as undue marketing of our own work, it is rather that our climate-economy work from the very beginning made an effort precisely to formulate the IAM, and all the issues that can be discussed with an IAM, in terms of a standard macroeconomic settings and in such a way that calibration and model evaluation could be conducted with standard methods. Ex post, then, one can say that our work grew out of an effort to write something akin to a climate-economy handbook for macroeconomists, even though the kind offer to write an actual such a chapter arrived much later. At this point, with this work, we are simply hopeful that macroeconomists with modern training will find our exposition useful as a quick introduction to a host of issues and perhaps also as inspiration for doing research on climate change and sustainability. We do find the area of great importance and, at the same time, rather undeveloped in many ways.

One exception to our claim that IAMs are not microeconomics-based macroeconomic models is Nordhaus's work, which started in the late 1970s and which led to the industry standards DICE and RICE: dynamic integrated models of climate and the economy, DICE depicting a one-region world and RICE a multi-region world. However, these models remain the nearest thing to the kind of setting we have in mind, and even the DICE and RICE models are closer to pure planning problems. That is, they do not fully specify market structures and, hence, do not allow a full analysis of typical policies such as a carbon tax or a quota system. Most of the models in the literature—to the extent they are fully specified models—are simply planning problems, so a question such as "What happens if we pursue a suboptimal policy?" cannot be addressed. This came as a surprise to us when we began to study the literature. Our subsequent research and the present chapter thus simply reflect this view: some more focus on the approach used in modern macroeconomics is a useful one.

So as a means of abstract introduction, consider a growth economy inhabited by a representative agent with utility function $\sum_{t=0}^{\infty} \beta^t u(C_t, S_t)$ with a resource constraint $C_t + K_{t+1} = (1-\delta)K_t + F(K_t, E_t, S_t)$ and with a law of motion $S_{t+1} = H(S_t, E_t)$. The new variables, relative to a standard macroeconomic setting, are S and E. S, a stock, represents something that is affects utility directly and/or affects production, whereas E, a flow, represents an activity that influences the stock. To a social planner, this would be nothing but an augmented growth model, with (interrelated) Euler equations both for K and S. In fact, standard models of human capital accumulation map into this setup, with H increasing in both arguments and F increasing in S but decreasing in E.[1] However, here we are interested in issues relating to environmental management—from a macroeconomic perspective—and then the same setup can be thought of, at least in abstract, with different labels: we could identify S with, say, clean air or biodiversity, and E with an activity that raises output but lowers the stock S. Our main interest will be in the connections between the economy and the climate. Then, S_t can be thought of as the climate at t, or a key variable that influences it, namely, the stock of carbon in the atmosphere; and E_t would be emissions of carbon dioxide caused by the use of fossil fuel in production. The carbon stock S then hurts both utility (perhaps because a warmer climate makes people suffer more in various ways) and output. Thus, $u_2 < 0$, $F_2 > 0$, $F_3 < 0$, $H_1 > 0$, and $H_2 > 0$. The setting still does not appear fully adequate for looking at the climate issue, because there ought to be another stock: that of the available amounts of fossil fuel (oil, coal, and natural gas), which are depletable resources in finite supply. Indeed, many of our settings below do include such stocks, but as we will argue even the setting without an additional stock is quite useful for analyzing the climate issue. Furthermore, one would also think that technology, and technological change of different sorts, must play a role, and indeed we agree. Technology can enhance the production possibilities in a neutral manner but also amount to specific forms of innovation aimed at developing non-fossil energy sources or more generally saving on fossil-based energy. We will discuss these issues in the chapter too, including endogenous technology, but the exposition covers a lot of ground and therefore only devotes limited attention to technology endogeneity.

Now so far the abstract setting just described simply describes preferences and technology. So how would markets handle the evolution of the two stocks K and S? The key approach here is that it is reasonable to assume, in the climate case, that the evolution of S is simply a byproduct of economic activity: an externality. Thus, tracing out the difference between an optimal path for K and S and a laissez-faire market path becomes important, as does thinking about what policies could be used to move the market outcome toward the optimum as well as what intermediate cases would imply. Thus, the modern macroeconomist approach would be to (i) define a dynamic competitive equilibrium with policy (say, a unit tax on E), with firms, consumers, and markets clearly spelled out, then (ii) look for insights about optimal policy both qualitatively and quantitatively (based on, say, calibration), and perhaps (iii) characterize outcomes for the future for different (optimal and suboptimal) policy scenarios. This is the overall approach we will follow here.

We proceed in three steps. In the first step, contained in Section 2, we discuss a setting with resource scarcity alone—such as an economy with a limited amount of oil. How will markets then price the resource, and how will it be used up over time? Thus, in this section we touch on the broader area of "sustainability", whereby the question is how the economy manages a set of depletable resources. It appears to be a common view in the public debate

[1] See, e.g., Lucas (1988).

that markets do not carry this task out properly, and our view is that it really is an open question whether they do or not; indeed, we find this issue intriguing in itself, quite aside from any interest in the specific area of climate change. The basic market mechanisms we go through involve the Hotelling rule for pricing and then, coupled with a representative agent with preferences defined over time as in our abstract setting above and a specific demand for the resource (say, from its use in production), a dynamic path for resource use. As a preliminary exploration into whether our market-based analysis works, one can compare the models implications for prices and quantities and we briefly do. As a rough summary, it is far from clear that Hotelling-based pricing can explain our past data for depletable resources (like fossil fuel or metals). Similarly, it is challenging to account for the historical patterns of resource use, though here the predictions of the theory are arguably less sharp. Taken together, this suggests that it is not obvious that at least our benchmark theories of markets match the data, so it seems fruitful to at least consider alternatives. In Section 2 we also look at the case of fossil fuel in more detail and, in this context, look at (endogenous) technical change: we look at how markets could potentially react to resource scarcity by saving on the scarce resource instead of saving on other inputs. Thus, we apply the notion of "directed technical change" in this context and propose it as an interesting avenue for conducting further macroeconomic research within the area of sustainability more broadly. Finally, Section 2 should be viewed as a delivering a building block for the IAMs to be discussed later in the chapter, in particular that in Section 5.

In Section 4, we take our second step and develop a very simple, static integrated assessment model of climate change and the global economy. Despite its being simple and stylized, this baseline model does have a macroeconomic structure, i.e., it makes assumptions that are standard in modern macroeconomic analysis. Many of its key parameters are therefore straightforwardly calibrated to observables and thus, with the additional calibration necessary to introduce climate into the model, it can be used to obtain an approximate social cost of carbon. The static model is then used to illustrate a number of points that have been made in broad literature on climate change. None of these applications do full justice to the literature, of course, since our main purpose is to introduce the macroeconomic analyst to it. At the same time, we do offer a setting that is quantitatively oriented and one can imagine embedding each application in a fully dynamic and calibrated model; in fact, as far as we are aware, only a (minority) subset of these applications exist as full quantitative studies in the literature.

In our last section, Section 5, which is also the third and final step of the chapter, we describe a fully dynamic, stochastic IAM setting. With it, we show how to derive a robust formula for the (optimal) marginal cost of carbon and, hence, the appropriate Pigou tax. We show how to assign parameter values and compute the size of the optimal tax. The model can also be used as a complete setting for predicting the climate in the future—along with the paths for consumption, output, etc.—for different policy paths. We conclude that although the optimal-tax formula is quite robust, the positive side of the model involves rather strong sensitivity to some parameters, such as those involving different sources for energy generation and, of course, the total sizes of the stocks of fossil fuels.

Before transiting from discussing sustainability in Section 2 to climate modeling in Section 4, we offer a rather comprehensive introduction to the natural-science aspects of climate change. This section, Section 3, is important for explaining what we perceive as the basic and (among expert natural scientists) broadly agreed upon mechanisms behind global warming: how the climate is influenced by the carbon concentration in the atmosphere (the climate model) and how the carbon concentration evolves over time as a function of the time path for emissions (the model of the carbon cycle). This presentation thus offers two "modules" that are crucial elements in IAMs. These modules are extremely simplified versions of what actual climate models and carbon-cycle models in use look like. However, they are, we argue, decent approximations of up-to-date models. The reason why simplifications are necessary is that our economic models have forward-looking agents and it is well known that such models are much more difficult to analyze, given any complexity in the laws of motions of stocks given flows: they involve finding dynamic fixed points, unlike any natural-science model where particles behave mechanically.[2]

Finally, although it should be clear already, let us reiterate that this chapter fails to discuss many environmental issues that are of general as well as macroeconomic interest. For example, the section on sustainability does not discuss, either empirically or theoretically, the possible existence of a "pollution Kuznets curve": the notion that over the course of economic development, pollution (of some or all forms) first increases and then decreases.[3] That section also does not offer any theoretical discussion of other common-pool problems than that associated with our climate (such as overfishing or pollution). The sections on IAMs, moreover, does not contain a listing/discussion of the different such models in the literature; such a treatment would require a full survey in itself.

[2]The statement about the complexity of economic models does not rely on fully rational expectations, which we do assume here, but at least on some amount of forward-looking because any forward-looking will involve a dynamic fixed-point problem.

[3]See, e.g., Grossman and Krueger (1991) and Stokey (1998).

2 Limited natural resources and sustainability concerns

Climate change is a leading example within environmental economics where global macroeconomic analysis is called for. It involves a global externality that arises from the release of carbon dioxide into the atmosphere. This release is a byproduct of our economies' burning of fossil fuel, and it increases the carbon dioxide concentration worldwide and thus causes warming not just where the emission occurs. In two ways, climate change makes contact with the broader area of *sustainability*: it involves two stocks that are important for humans and that are affected by human activity. The first stock is the carbon concentration in the atmosphere. It exerts an influence on the global climate; to the extent warming causes damages on net, it is a stock whose size negatively impacts human welfare. The second stock is that of fossil fuels, i.e., coal, oil, and natural gas. These stocks are not harmful per se but thus can be to the extent they are burnt.

More generally, sustainability concerns can be thought of in terms of the existence of stocks in finite supply with two properties: (i) their size is affected by economic activity and (ii) they influence human welfare.[4] Obvious stocks are natural resources in finite supply, and these are often traded in markets. Other stocks are "commons", such as air quality, the atmosphere, oceans, ecosystems, and biodiversity. Furthermore, recently, the term "planetary boundaries" has appeared (Rockström et al., 2009, *Nature*). These boundaries represent other limits that may be exceeded with sufficient economic growth (and therefore, according to the authors, growth should be limited). This specific *Nature* article lists nine boundaries, among them climate change; the remaining items are (i) stratospheric ozone depletion, (ii) loss of biosphere integrity (biodiversity loss and extinctions), (iii) chemical pollution and the release of novel entities, (iv) ocean acidification, (v) freshwater consumption and the global hydrological cycle, (vi) land system change, (vii) nitrogen and phosphorus flows to the biosphere and oceans, and (viii) atmospheric aerosol loading. Thus, these are other examples of commons.

Aside from in the work on climate change, the macroeconomic literature has had relatively little to say on the effects and management of global stocks. The Club of Rome (that started in the late 1960s) was concerned with population growth and a lack of food and energy. The oil crisis in the 1970s prompted a discussion about the finiteness of oil (see, e.g., the 1974 *Review of Economic Studies* issue on this topic), but new discoveries and a rather large fall in the oil price in the 1980s appeared to have eliminated the concern about oil among macroeconomists. Similarly, technology advances in agriculture seemed to make limited food supply less of an issue. Nordhaus (1973, 1974) discussed a limited number of metals in finite supply, along with their prices, and concluded that the available stocks were so large at that point that there was no cause for alarm in the near to medium-run future. Thus, the concerns of these decades did not have a long-lasting impact on macroeconomics. Perhaps relatedly, so-called green accounting, where the idea is to measure the relevant stocks and count their increases or decreases as part of an extended notion of national economic

[4]Relatedly, but less relevant from the perspective taken in this section, there is theoretical work on sustainability, defining, based on a utility-function representation, what the term means: roughly, an allocation is sustainable if the indirect utility function of generation t is not be below that of generation $t-k$.

product, was proposed but has been implemented and used in relatively few countries.[5] Limited resources and sustainability are typically not even mentioned in introductory or intermediate undergraduate textbooks in macroeconomics, let alone in PhD texts. In PhD texts specifically on growth, there is also very little: Aghion and Howitt's (2008) growth book has a very short, theoretical chapter on the subject, Jones (2001) has a chapter in his growth book which mentions some data; Acemoglu's (2009) growth book has nothing.[6]

The purpose here is not to review the literature but to point to this broad area as one of at least potential relevance and as one where we think that more macroeconomic research could be productive. To this end, we will discuss the basic theory and its confrontation with data. This discussion will lay bare some challenges and illustrate the need for more work.

We will focus on finite resources that are traded in markets and hence abstract from commons, mainly because these have not been subject to much economic macroeconomic analysis (with the exception of the atmosphere and climate change, which we will discuss in detail below). Thus, our discussion begins with price formation and quantity determination in markets for finite resources and then moves on to briefly discuss endogenous technological change in the form of resource saving.

2.1 Prices and quantities in markets for finite resources

To begin with, let us consider the simplest of all cases: a resource e in finite supply R that is costless to extract and that has economic value. Let us suppose the economic value is given by an inverse demand function $p_t = D(e_t)$, which we assume is time-invariant and negatively sloped. In a macroeconomic context we can derive such a function assuming, say, that e is an input into production. Abstracting from capital formation, suppose $y_t = F(n_t, e_t) = An_t^{1-\nu}e_t^{\nu}$, with inelastic labor supply $n_t = 1$, that $c_t = y_t$, and that utility is $\sum_{t=0}^{\infty} \beta^t \log c_t$.[7] Let time be $t = 0, \dots, T$, with T possibly infinite. Here, the demand function would be derived from the firm's input decision: $p_t = \nu A e_t^{\nu-1}$.

2.1.1 The Hotelling result: the price equation in a baseline case

The key notion now is that the resource can be saved. We assume initially that extraction/use of the resource is costless. The decision to save is therefore a dynamic one: should the resource be sold today or in the future? For a comparison, an interest rate is needed, so let r_t denote the interest rate between $t-1$ and t. If the resource is sold in two consecutive periods, it would then have to be that on the margin, the owner of the resource is indifferent

[5] For example, in the United States, the BEA started such an endeavor in the 1990s but it was discontinued.

[6] The area of *ecological economics* is arguably further removed from standard economic analysis and certainly from macroeconomics. It is concerned precisely with limited resources but appears, at least in some of its versions, to have close connections Marx's labor theory of value, but with "labor" replaced by "limited resources" more broadly and, in specific cases, "energy" or "fossil fuel".

[7] In all of this section, we use logarithmic utility. More general CRRA preferences would only slightly change the analysis and all the key insights remain the same in this more general case.

between selling at t and at $t+1$:

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

This is the *Hotelling equation*, presented in Hotelling (1931). The price of the finite resource, thus, grows at the real rate of interest. The equation can also be turned around, using the inverse demand function, to deliver predictions for how the quantity sold will develop; for now, however, let us focus on the price. Thus, we notice that an arbitrage condition delivers a sharp prediction for the dynamics of the price that is independent of the demand. For the *price dynamics*, the demand is only relevant to the extent it may be such that the resource is not demanded at all at some point in time. For the *price level(s)*, however, demand is of course key: one needs to solve the difference equation along with the inverse demand function and the constraint on the resource to arrive at a value for p_0 (and, consequently, all its subsequent values). Here, p_t would be denoted the *Hotelling rent* accruing to the owner: as it is costless to extract it, the price is a pure rent. Thus, to the extent the demand is higher, the price/rent path will be at a higher level. Similarly, if there is more of the resource, the price/rent path will be lower, since more will be used at each point in time.

2.1.2 Prices and quantities in equilibrium: using a planning problem

Let us consider the planning problem implicit in the above discussion and let us for simplicity assume that $T=\infty$. Thus the planner would maximize $\nu \sum_{t=0}^{T} \beta^t \log c_t$ subject to $c_t = Ae_t^\nu$ for all t and $\sum_{t=0}^{T} e_t = R$.[8] This delivers the condition $\nu\beta^t/e_t = \mu$, where μ is the multiplier on the resource constraint, and hence $e_{t+1} = \beta e_t$. Inserting this into the resource constraint, one obtains $e_0(1+\beta+\dots) = e_0/(1-\beta) = R$. Hence, $e_0 = (1-\beta)R$ and the initial price of the resource in terms of consumption (which can be derived from the decentralization) will be $p_0 = A\nu\left((1-\beta)R\right)^{\nu-1}$. Furthermore, $p_t = A\nu\left((1-\beta)R\right)^{\nu-1}\beta^{(\nu-1)t}$; notice that the gross interest rate here is constant over time and equal to $\beta^{\nu-1}$.[9] We see that a more abundant resource translates into a lower price/rent. In particular, as R goes to infinity, the price approaches 0: marginal cost. Similarly, higher demand (e.g., through a higher A or higher weight on future consumption, β, so that the resource is demanded in more periods and will thus not experience as much diminishing returns per period), delivers a higher price/rent. Consider also the extension where the demand parameter A is time-varying. Then the extraction path is not affected at all, due to income and substitution effects canceling. The consumption interest rates will change, since the relative price between consumption and the resource must change. The equation for price dynamics applies just as before, however, so price growth is affected only to the extent the interest rate changes. The price level, of course, is also affected by overall demand shifts.

[8] For $\nu = 1$ this is a standard cake-eating problem.

[9] The Euler equation of the the consumer delivers $1+r_{t+1} = c_{t+1}/(c_t\beta) = e_{t+1}^\nu/(e_t^\nu\beta) = \beta^\nu/\beta = \beta^{\nu-1}$.

2.1.3 Extraction costs

More generally, suppose that the marginal cost of extraction of the resource is c_t in period t, and let us for simplicity assume that these marginal costs are exogenous (more generally it would depend on the amount extracted and the total remaining amount of the resource). The Hotelling formula for price dynamics becomes

$$p_t - c_t = \frac{1}{1 + r_{t+1}}(p_{t+1} - c_{t+1}).$$

Put differently, the Hotelling rent, which is now the marginal profit per unit, $p - c$, grows at the real rate of interest. This is thus the more general formula that applies. It is robust in a number of ways; e.g., allowing endogenous extraction costs delivers the same formula and the consideration of uncertainty reproduces the formula in expectation).[10] The discussion of determinants of prices and quantities above thus still applies, though the key object now becomes the marginal profit per unit. First, the general idea that more of the resource (higher R) lowers the price survives: more of the resource moves the price toward marginal cost, thus gradually eliminating the rent. Second, regarding the effects of costs, let us consider three key cases: one where marginal costs are constant (and positive), one where they are declining, and one where they are increasing. We assume, for simplicity, that there is a constant interest rate. A constant positive marginal cost thus implies that the price is rising at a somewhat lower rate initially than when extraction is costless, since early on the price is a smaller fraction of the rent (early on, there is more left of the resource). If the marginal cost of extraction rises over time—a case that would apply in the absence of technological change if the easy-to-extract sources are exploited first—the price will rise at a higher rate; and under the assumption of a falling marginal extraction cost, typically reflecting productivity improvements in extraction, prices rise more slowly. Quantity paths change accordingly, when we use an invariant demand function. With a faster price rise, quantities fall faster, and vice versa. In particular, when the future promises lower (higher) extraction costs, extraction is postponed (slowed down) and so falls less (more) rapidly.

2.2 Confronting theory with data

The Hotelling predictions are, in principle, straightforwardly compared with data. The ambition here is not to review all the empirical work evaluating the Hotelling equation for finite resources but merely to mention some stylized facts and make some general points.[11] As for prices, it is well known that (real) prices of metals fall at a modest rate over the "long run", measured as one hundred years or more; see, e.g., Harvey et al. (2010). The prices of fossil fuels (oil, coal, and natural gas) have been stable, with a slight net increase

[10]The case where the natural resource is owned by a monopolist produces a more complicated formula, as one has to consider marginal revenue instead of price and as the interest rate possibly becomes endogenous. However, the case of monopoly does not appear so relevant, at least not today. In the case of oil, Saudi oil production is currently only about 10% of world production.

[11]For excellent discussions, see, e.g., Krautkraemer (1998) and Cuddington and Nulle (2014).

over the last 40 or so years. The volatilities of all these time series are high, on the order of magnitude of those for typical stock-market indices.[12] When it comes to quantities, these time series have been increasing steadily, and with lower fluctuations than displayed by the corresponding prices. Are these observations broadly consistent with Hotelling's theory?

To answer this question, note that Hotelling's theory is mainly an arbitrage-based theory of prices and that quantity predictions involve more assumptions on supply and demand, such as those invoked in our planning problem above. To evaluate Hotelling's rule, we first need to have an idea of the path for extraction costs, as they figure prominently in the more general version of the theory. The situation is somewhat complicated by the fact that extraction occurs on multiple sites. For oil at least, it is also clear that the marginal costs differ greatly between active oil wells, for example with much lower costs in Saudi Arabia than in the North Sea. This in itself appears inefficient, as the less expensive oil ought to be extracted first in order to minimize overall present-value costs. We know of no study that has good measurements of marginal extraction costs going far back in time. Suppose, however, that productivity growth in the mining/extraction sector was commensurate with that in the rest of our economies. Then it would be reasonable to assume that the *relative* cost of extracting natural resources—and that is the relevant price given that we are referring to evidence on real prices—does not have any sharp movements upward or downward. Hence, the Hotelling formula, given a known total depletable stock of the resource, would imply an increasing price series, at a rate of a few percent per year, with a slightly lower growth rate early on, as explained above. This is clearly not what we see. It is, alternatively, possible that extraction costs have developed unevenly. Pindyck (1978) argues, for the case of oil, that lower and lower extraction costs explained a stable price path initially but that later extraction costs stabilized (or even increased), hence pushing prices up. In retrospect, however, although prices rose again in 1979 they did not continue increasing after that and rather fell overall; today, the oil price is back at a real price that is not terribly far from the pre-1973 level.

An proposed explanation for the lack of price growth in the data is a gradual finding of new deposits (of oil, metals, and so on). As explained above, the theory does predict lower prices for higher total deposits of the resource. However, it would then have to be that markets systematically under-predicted the successes of new explorations, and over very long periods of time.

Relatedly, it is possible that markets expect technological change in the form of the appearance of close substitutes to the resource in question. Consider a very simple case with a costless-to-extract raw material as in the baseline Hotelling model but where next period a perfect substitute, in infinite supply and with a constant marginal cost $\bar{p}$, appears with some probability. Then the arbitrage equation reads $p_t = \frac{1}{1+r_{t+1}} \left(\pi_{t+1}p_{t+1} + (1-\pi_{t+1})\bar{p}\right)$, where π_{t+1} is the probability of the perfect substitute appearing. Clearly, such uncertainty and potential price competition will influence price dynamics and will lead to richer predictions. However, we know of no systematic study evaluating a quantitative version of this kind of hypothesis and comparing it to data.

[12]There are also attempts to identify long-run cycles; see, e.g., Erten and Ocampo (2012).

A different view of the prices of natural resources (and commodities more generally) is the Prebisch (1950) and Singer (1950) hypothesis: that commodities have lower demand elasticities, so that when income rises, prices fall. Their hypothesis, thus, is in contrast with Hotelling's rule, since scarcity is abstracted from. Clearly, if one formulated a model with the Prebisch-Singer assumption and scarcity, as discussed above, the Hotelling formula would survive, and any demand effects would merely affect the level of the price path and not its dynamics.

In sum, although many authors claim that richer versions of the Hotelling model take its predictions closer to data, it seems safe to say that there is no full resolution of the contrast between the model's prediction of rising prices/profits per unit (at the rate of interest) and the data showing a stable or declining real price of the typical resource. Some would argue that markets are not fully rational, or not forward-looking enough: the power of the scarcity argument in Hotelling's seminal work is very powerful but relies crucially on forward-looking with a long horizon, to the extent there is a relatively large amount of the resource left in ground. It seems to us that this hypothesis deserves some attention, though it is a challenge even to formulate it.[13]

To evaluate quantities, as underlined above, a fuller theory needs to be specified. This leads to challenges as well, as we shall see. Here, we will simply look at an application, albeit a well-known one and one that is relevant to the climate context. In the context of this application, we will also discuss technological change as a means toward saving on a scarce resource.

2.3 An application: fossil energy

On a broad level, when a resource is in scarce supply, a key question is its substitutability with alternative resources. In this section we look at fossil energy and provide an outline of how one could go about looking at one aspect of scarcity in this market: the response of energy-saving, i.e., one of the ways in which markets can respond to a shortage. This analysis, like the rest of this chapter (that addresses climate change), is built on a quantitatively oriented macroeconomic model. It can also be regarded as one of the building blocks in the climate-economy model; indeed, the exhaustible-resource formulation in Section 5 coincides with the core formulation entertained here.

The starting point is the extension of basic growth theory to include energy; the standard reference is Dasgupta and Heal (1974), but noteworthy other contributions include those by Solow (1974) and Stiglitz (1974). One of the main concerns here was precisely sustainability, i.e., whether production functions (or various sorts) would allow future generations to be as well off as current generations. The Cobb-Douglas function was found to be an in-between case here; with more substitutability between energy and the other inputs, sustainability was possible. This line of work did not much address technical change, neither quantitatively nor theoretically. Clearly, much of the literature on scarce resources was written shortly after the oil-price hikes in the 1970s and it was not until the late 1980s that the theoretical

[13] See, e.g., Spiro (2014).

developments allowed technological change to be endogenized in market-based environments.

We build a similar framework to that in Dasgupta and Heal's work and formulate an aggregate production function with three inputs—capital, labor, and fossil energy—and we use it to account for postwar U.S. data. This analysis follows Hassler, Krusell, and Olovsson (2015) closely. We allow technical change in this production function in the form of capital/labor-saving and energy-saving and we consider three broad issues: (i) what substitution elasticity (between a capital-labor composite, on the one hand, and energy, on the other) fits the data best; (ii) measurement of the series for input-saving and to what extent they appear to respond to price movements (i.e., does energy-saving appear to respond to the price of fossil fuel?); and (iii) the model's predictions for future input saving and fossil-fuel dependence. The model focuses on energy demand, as derived from an aggregate production function, and all of the discussion can be carried out without modeling supply.

So consider an aggregate production function of the nested CES form

$$y = \left[(1-\nu)\left[Ak^{\alpha}l^{1-\alpha}\right]^{\frac{\varepsilon-1}{\varepsilon}} + \left[A^{e}e\right]^{\frac{\varepsilon-1}{\varepsilon}}\right]^{\frac{\varepsilon}{\varepsilon-1}},$$

with the obvious notation.[14] Here, we see that $\varepsilon \in [0,\infty]$ expresses the substitutability between capital/labor and energy. A is the technology parameter describing capital/labor-saving and A^{e} correspondingly describes energy-saving. If there is perfect competition for inputs, firms set the marginal product of each input equal to its price, delivering—expressed in terms of shares—the equations

$$\frac{wl}{y} = (1-\alpha)(1-\gamma)\left[\frac{Ak^{\alpha}l^{1-\alpha}}{y}\right]^{\frac{\varepsilon-1}{\varepsilon}} \tag{1}$$

and

$$\frac{pe}{y} = \gamma\left[\frac{A^{e}e}{y}\right]^{\frac{\varepsilon-1}{\varepsilon}}. \tag{2}$$

2.3.1 Accounting for input saving using U.S. data

Equations (1) and (2) can be rearranged and solved directly for the two technology trends A and A^{E}. This means that it is possible, as do Hassler et al., to use data on output and inputs and their prices to generate time paths for the input-saving technology series. This is parallel with Solow's growth-accounting exercise, only using a specific functional form. In particular, A^{e} can be examined over the postwar period, when the price of fossil fuel—oil in particular—has moved around significantly, as shown in Figure 12.

The authors use this setting and these data to back out series for A^{e} and A, conditional on a value for ε. With the view that the A and A^{e} series are technology series mainly, one can

[14]This production function introduces a key elasticity, along with input-specific technology levels, in the most tightly parameterized way. Extensions beyond this functional class, e.g., to the translog case, would be interesting not only for further generality but because it would introduce a number of additional technology shifters; see, e.g., Berndt and Christensen (1973).

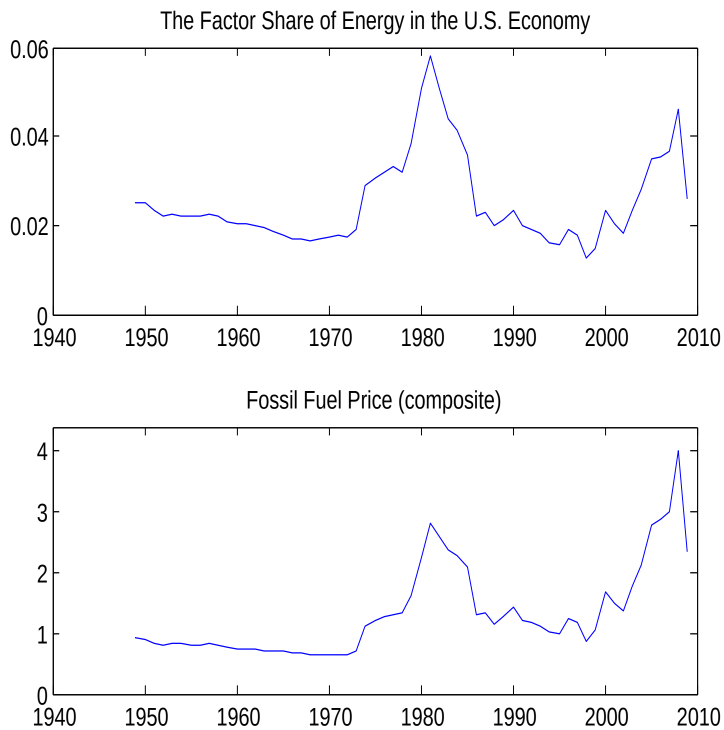

Figure 1: Fossil energy share and its price

then examine the extent to which the backed-out series for different ε look like technology series: are fairly smooth and mostly non-decreasing. It turns out that ε has to be close to zero for the A^e series to look like a technology series at all; if ε is higher than 0.2 or so, the implied up-and-down swings in A^e are too high to be plausible. On the other hand, for a range of ε values between 0 (implying that production is Leontief) and 0.1, the series is rather smooth and looks like it could be a technology series. Figure 2 plots both the A and A^e series. We see that A^e grows very slowly until prices rise; then it starts growing significantly. Hence, the figure does suggest that the scarcity mechanism is operative in a quantitatively important way. It is also informative to look at how the two series compare. A it looks like TFP overall, but more importantly it does seem to covary negatively in the medium run with A^e, thus suggesting that the concept of *directed technical change* may be at play. In other words, when the oil price rose, the incentives to save on oil and improve oil efficiency went up, and to the extent these efforts compete for a scarce resource that could alternatively be used for saving on/improving the efficiency of capital and labor, as a result the latter efforts would have fallen.

Hassler et al. (2015) go on to suggest a formal model for this phenomenon and use it, with a calibration of the technology parameters in R&D based on the negative historical association between A and A^E, to also predict the future paths of technology and of energy

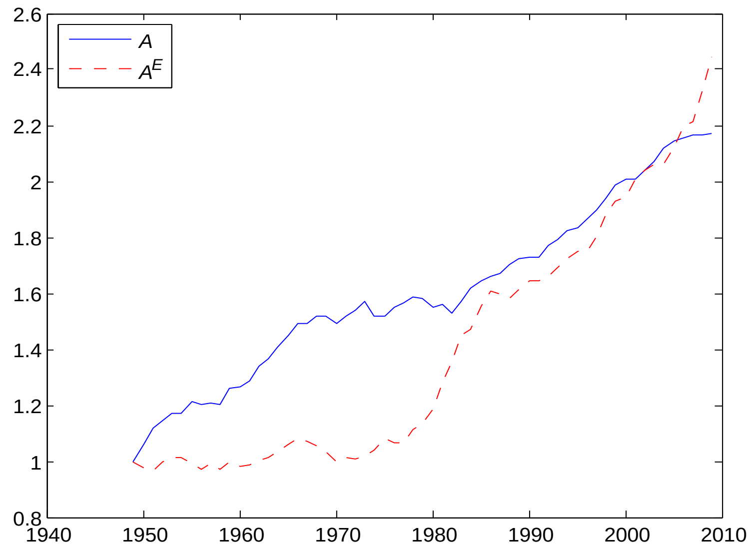

Figure 2: Energy- and capital/labor-saving technologies compared

dependence. We will briefly summarize this research below, but first it is necessary to formulate a quantitatively oriented dynamic macroeconomic model with energy demand and supply included explicitly.

2.3.2 A positive model of energy supply and demand with a finite resource

Using the simple production function above and logarithmic preferences, it is straightforward to formulate a planner's problem, assuming that energy comes from a finite stock. We will first illustrate with a production function that is in the specified class and that is often used but that does not (as argued above) fit the macroeconomic data: the Cobb-Douglas case, where $F(Ak^{\alpha}, A^{e}e) = k^{\alpha}e^{\nu}$, where a constant labor supply (with a share $1 - \alpha - \nu$) is implicit and we have normalized overall TFP including labor to 1. We also assume, to simplify matters, that (i) there is 100% depreciation of capital between periods (which fits a period of, say, 20 years or more) and that (ii) the extraction of energy is costless (which fits oil rather well, as its marginal cost is much lower than its price, at least for much of the available oil). For now, we abstract from technological change; we will revisit it later. Thus, the planner would maximize

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

subject to

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

and $\sum_{t=0}^{\infty} e_t = R$, with R being the total available stock. It is straightforward to verify that we obtain a closed-form solution here: consumption is a constant fraction $1-\alpha\beta$ of output and

$e_t = (1-\beta)\beta^t R$, i.e., energy use falls at the rate of discount. As energy falls, so does capital, consumption, and output. In fact, this model asymptotically delivers balanced (negative) growth at a gross rate g satisfying (from the resource constraint) $g = g^\alpha \beta^\nu = \beta^{\frac{\nu}{1-\alpha}}$. Capital is not on the balanced path at all times, unless its initial value is in the proper relation to initial energy use.[15] This model of course also generates the Hotelling result: p_{t+1} must equal $p_t(1+r_{t+1})$, where $1+r$ is the marginal product of capital and $1+r$ hence the gross real interest rate. Notice, thus, that the interest rate will be constant on the balanced growth path but that it obeys transition dynamics. Hence, even though energy use falls at a constant rate at all times, the energy price will not grow at a constant rate at all times (unless the initial capital stock is at its balanced-growth level): it will grow either faster or slower. Consumption, along with output and capital, goes to zero here along a balanced growth path, but when there is sufficient growth in technology (which is easily added in the model), there will be positive balanced growth. The striking fall in energy use over time would of course be mitigated by an assumption that marginal extraction costs are positive and decreasing over time, as discussed above, but it is not obvious that such an assumption is warranted.

Figure 2.3.2, which is borrowed from Hassler et al. (2015), shows that, in the data, energy (defined as a fossil composite) rises significantly over time. In contrast, as we have just shown, the simple Cobb-Douglas model predicts falling energy use, at a rate equalling the discount rate. Suppose instead one adopts the model Hassler et al. (2015) argue fit the data better, i.e., a function that is near Leontief in k^α and e. Let us first assume that the technology coefficients A and A^e are constant over time. Then, there will be transition dynamics in energy use, for Ak^α has to equal $A^e e$ at all points in time. Thus, the initial value of capital and R may not admit balanced growth in e at all times, given A and A^e. Intuitively, if Ak_0^α is too low, e will be held back initially and grow over time as capital catches up to its balanced path. Thus, it is possible to obtain an increasing path for energy use over a period of time. Eventually, of course, energy use has to fall. There is no exact balanced growth path in this case. Instead, the saving rate has to go to zero since any positive long-run saving rate would imply a positive capital stock.[16] Hence, the asymptotic economy will be like one without capital and in this sense behave like in a cake-eating problem: consumption and energy will fall at rate β. In sum, this model can deliver *peak oil*, i.e., a path for oil use with a maximum later than at time 0. As already pointed out, increasing oil use can also be produced from other assumptions, such as a decreasing sequence of marginal extraction costs for oil; these explanations are complementary.

With exogenous technology growth in A and A^e it is possible that very different long-run extraction behavior results.[17] In particular, it appears that a balanced growth path with the property that $g_A g^\alpha = g_{A^e} g_e = g$ is at least feasible. Here, the first equality follows

[15] Initial capital then has to equal $(\alpha(R(1-\beta))^\nu)^{\frac{1}{1-\alpha}} \beta^{\frac{1-\alpha-\nu}{1-\alpha}}$.

[16] If the saving rate asymptotically stayed above $\underline{s} > 0$, then $k_{t+1} \geq \underline{s} A k_t^\alpha$. This would imply that capital would remain uniformly bounded below from zero. However, here, it does have to go to zero as its complement energy has to go to zero.

[17] An exception is the Cobb-Douglas case for which it is easy to show that the result above generalizes: e falls at rate β.

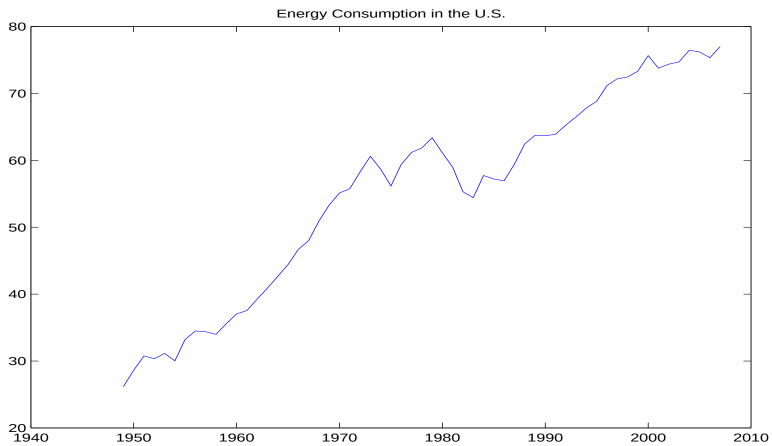

Figure 3: U.S. energy use

from the two arguments of the production function growing at the same rate—given that the production function F is homogeneous of degree one in the two arguments Ak^{α} and $A^e e$—and the second equality says that output and capital have to grow at that same rate. Clearly, if the planner chooses such asymptotic behavior, g_e can be solved for from the two equations to equal $g_A^{\frac{1}{1-\alpha}}/g_{A^e}$, a number that of course needs to be less than 1. Thus, in such a case, g_e will not generally equal β. A more general study of these cases is beyond the scope of the present chapter.

2.3.3 Endogenous energy-saving technical change

Given the backed-out series for A and A^e, which showed negative covariation in the medium run, let us consider the model of technology choice Hassler et al. (2015) propose. In it, there is an explicit tradeoff between raising A and raising A^e. Such a tradeoff arguably offers one of the economy's key behavioral responses to scarcity. That is, growth in A^e can be thought of as energy-saving technological change. In line with the authors' treatment, we consider a setup with directed technological change in the form of a planning problem, thus interpreting the outcome as one where the government has used policy optimally to internalize any spillovers in the research sector. It would be straightforward, along the lines of the endogenous-growth literature following Romer (1990), to consider market mechanisms based on variety expansion or quality improvements, monopoly power, possibly with Schumpeterian elements, and an explicit market sector for R&D. Such an analysis would be interesting and would allow interesting policy questions to be analyzed. For example, is the market mechanism not allowing enough technical change in response to scarcity, and does the answer depend on whether there are also other market failures such as a climate externality? We leave these interesting questions for future research and merely focus here on efficient outcomes. The key mechanism we build in rests on the following simple structure: we introduce one resource, a measure one of "researchers". Researchers can direct their efforts to the advancement of A and A^e. We look at a very simple formulation:

$$A_{t+1} = A_t f(n_t) \quad \text{and} \quad A^e_{t+1} = A^e_t f_e(1-n_t),$$

where $n_t \in [0,1]$ summarizes the R&D choice at time t and where f and f_e are both strictly increasing and strictly concave; these functions thus jointly demarcate the frontier for technologies at $t+1$ given their positions at t. Hence, at a point in time t, A_t and A^e_t are fixed. In the case of a Leontief technology, there would be absolutely no substitutability at all between capital and energy ex post, i.e., at time t when A_t and A^e_t have been chosen, but there is substitutability ex ante, by varying n_s for $s < t$. With a less extreme production function there would be substitutability ex post too but less so than ex ante.[18] Relatedly, whereas the share of income in this economy that accrues to each of the inputs is endogenous and, typically, varies with the state of the economy, on a balanced growth path the share settles down. As we shall see, in fact, the share is determined in a relatively simple manner.

The analysis proceeds by adding these two equations to the above planning problem. Taking first-order conditions and focusing on a balanced-growth outcome, this model rather surprisingly delivers the result that the extraction rate must be equal to β, regardless of the values of all the other primitives.[19] This means, in turn, that two equations jointly determining the long-run growth rates of A and A_e can be derived. One captures the technology tradeoff and follows directly from the equations above stating that these growth rates, respectively, are $g_A = f(n)$ and $g_{A^e} = f^e(1-n)$. The other equation comes from the balanced-growth condition that $A_t k_t^\alpha = A^e_t e_t$, given that F is homogeneous of degree one; from this equality the growth rates of A and A^e are positively related. In fact, given that e_t falls at rate β, we obtain n from $g_A^{\frac{1}{1-\alpha}} = g_{A^e}\beta$.

[18] The Cobb-Douglas case is easy to analyze. It leads to an interior choice for n that is constant over time, regardless of initial conditions and hence looks like the case above where the two technology factors are exogenous.

[19] The proof is straightforward; for details, see Hassler et al. (2015). It is thus the endogeneity of the technology levels in the CES formulation that makes energy fall at rate β; when they grow exogenously, we saw that energy does not have to go to zero at rate β.

One can also show, quite surprisingly as well, that the long run share of energy s_e in output is determined by $(1-s_e)/s_e = -\partial \log g_A/\partial \log g_{A^e}$.[20] In steady state, this expression is a function of n only, and as we saw above it is determined straightforwardly knowing β, α, f, and f^e. How, then, can these primitives be calibrated? One way to proceed is to look at historical data to obtain information about the tradeoff relation between g_A and g_{A^e}. If this relation is approximately log-linear (i.e., the net rates are related linearly), the observed slope is all that is needed, since it then gives $\partial \log g_A/\partial \log g_{A^e}$ directly. The postwar behaviors of A and A^e reported above imply a slope of -0.235 and hence a predicted long-run value of s_e of around 0.19, which is significantly above its current value, which is well below 0.1.

2.3.4 Takeaway from the fossil-energy application

The fossil-energy application shows that standard macroeconomic modeling with the inclusion of an exhaustible resource can be used to derive predictions for the time paths for quantities and compare them to data. Moreover, the same kind of framework augmented with endogenous directed technical change can be used to look at optimal/market responses to scarcity. It even appears possible to use historical data reflecting past technological tradeoffs in input saving to make predictions for the future. The presentation here has been very stylized and many important real-world features have largely been abstracted from, such as the nature of extraction technologies over time and space. The focus has also been restricted to the long-run behaviors of the prices and quantities of the resources in limited supply, but there are other striking facts as well, such as the high volatilities in most of these markets. Natural resources in limited supply can become increasingly limiting for economic activity in the future and more macroeconomic research may need to be directed to these issues. Hopefully the analysis herein can give some insights into fruitful avenues for such research.

[20] The authors show that this result follows rather generally in the model: utility is allowed to be any power function and production any function with constant returns to scale.

3 Climate change: the natural-science background

An economic model of climate change needs to describe three phenomena and their dynamic interactions. These are (i) economic activity; (ii) carbon circulation; and (iii) the climate. From a conceptual as well as a modeling point of view it is convenient to view the three phenomena as distinct sub-subsystems. We begin with a very brief description of the three sub-systems and then focus this section on the two latter.

The economy consists of individuals that act as consumers, producers and perhaps as politicians. Their actions are drivers of the economy. In particular, the actions are determinants of emissions and other factors behind climate change. The actions are also responses to current and expected changes in the climate by adaptation. Specifically, when fossil fuel is burned, carbon dioxide (CO_2) is released and spreads very quickly in the atmosphere. The atmosphere is part of the carbon circulation sub-system where carbon is transported between different reservoirs; the atmosphere is thus one such reservoir. The biosphere (plants, and to a much smaller extent, animals including humans) and the soil are other reservoirs. The oceans constitute the largest carbon reservoir.

The climate is a system that determines the distribution of weather events over time and space and is, in particular, affected by the carbon dioxide concentration in the atmosphere. Due to its molecular structure, carbon dioxide more easily lets through short-wave radiation, like sun-light, than long-wave, infrared radiation. Relative to the energy outflow from earth, the inflow consists of more short-wave radiation. Therefore, an increase in the atmospheric CO_2 concentration affects the difference between energy inflow and outflow. This is the *greenhouse effect*.

It is straightforward to see that we need at minimum the three sub-systems to construct a climate-economy model. The economy is needed to model emissions and economic effects of climate change. The carbon circulation model is needed to specify how emissions over time translate into a path of CO_2 concentration. Finally, the climate model is needed to specify the link between the atmospheric CO_2 concentration and the climate.

3.1 The climate

3.1.1 The energy budget

We will now present the simplest possible climate model. As described above, the purpose of the climate model is to determine how the (path of) CO_2 concentration determines the (path of the) climate. A minimal description of the climate is the global mean atmospheric temperature near the surface. Thus, at minimum we need a relationship between the path of the CO_2 concentration and the global mean temperature. We start the discussion by describing the *energy budget* concept.

Suppose that the earth is in a radiative steady state where the incoming flow of short-wave radiation from the sun light is equal to the outgoing flow of largely infrared radiation.[21]

[21] We neglect the additional outflow due to the nuclear process in the interior of the earth, which is in the order of one to ten thousands in relative terms when compared to the incoming flux from the sun; see the

The energy budget of the earth is then balanced, implying that the earth's heat content and the global mean temperature is constant.[22] Now consider a perturbation of this equilibrium that makes the net inflow positive by an amount F. Such an increase could be caused by an increase in the incoming flow and/or a reduction in the outgoing flow. Regardless of how this is achieved, the earth's energy budget is now in surplus causing an accumulation of heat in the earth and thus a higher temperature. The speed at which the temperature increases is higher the larger is the difference between the inflow and outflow of energy, i.e., the larger the surplus in the energy budget.

As the temperature rises, the outgoing energy flow increases since all else equal, a hotter object radiates more energy. Sometimes this simple mechanism is referred to as the 'Planck feedback'. As an approximation, let this increase be proportional to the increase in temperature over its initial value. Denoting the temperature perturbation relative to the initial steady state at time t by T_t and the proportionality factor between energy flows and temperature by κ, we can summarize these relations in the following equation:

$$\frac{dT_t}{dt} = \sigma \left(F - \kappa T_t \right). \tag{3}$$

The left-hand side of the equation is the speed of change of the temperature at time t. The term in parenthesis on the right-hand side is the net energy flow, i.e., the difference in incoming and outgoing flows. The equation is labeled the energy budget and we note that it should be thought of as a flow budget with an analogy to how the difference between income and spending determines the speed of change of assets.

When the right-hand side of (3) is positive, the energy budget is in surplus, heat is accumulated, and the temperature increases. Vice versa, if the energy budget has a deficit, heat is lost, and the temperature falls. When discussing climate change, the variable F is typically called *forcing* and it is then defined as the change in the energy budget caused by human activities. The parameter σ is (inversely) related to the heat capacity of the system for which the energy budget is defined and determines how fast the temperature changes for a given imbalance of the energy budget.[23]

We can use equation (3) to find how much the temperature needs to rise before the system reaches a new steady state, i.e., when the temperature has settled down to a constant. Such an equilibrium requires that the energy budget has become balanced, so that the term in parenthesis in (3) again has become zero. Let the steady-state temperature associated with a forcing F be denoted $T(F)$. At $T(F)$, the temperature is constant, which requires that the energy budget is balanced, i.e., that $F - \kappa T(F) = 0$. Thus,

$$T(F) = \frac{F}{\kappa}. \tag{4}$$

KamLAND Collaboration, (2011).

[22] We disregard the obvious fact that energy flows vary with latitude and over the year producing differences in temperatures over space and time. Since the outflow of energy is a non-linear (convex) function of the temperature, the distribution of temperature affects the average outflow.

[23] The heat capacity of the atmosphere is much lower than that of the oceans, an issue we will return to below.

Furthermore, the path of the temperature is given by

$$T_t = e^{-\sigma\kappa t}\left(T_0 - \frac{F}{\kappa}\right) + \frac{F}{\kappa}.$$

Measuring temperature in Kelvin (K), and F in Watt per square meter, the unit of κ is $\frac{W/m^2}{K}$.[24] If the earth were a blackbody without an atmosphere, we could calculate the exact value of κ from laws of physics. In fact, at the earth's current mean temperature $\frac{1}{\kappa}$ would be approximately 0.3, i.e., an increase in forcing by 1 W/m^2 would lead to an increase in the global temperature of 0.3 K (an equal amount in degrees Celsius).[25] In reality, various feedback mechanisms make it difficult to assess the true value of κ. One of the important feedbacks is that a higher temperature increases the concentration of water vapor, which is also is a greenhouse gas; another is that the polar ice sheets melt, which decreases direct reflection of sun light and changes the cloud formation. We will return to this issue below but note that the value of κ is likely to be substantially smaller than the blackbody value of 0.3^{-1}, leading to a higher steady-state temperature for a given forcing.

Now consider how a given concentration of CO_2 determines F. This relationship can be well approximated by a logarithmic function. Thus, F, the change in the energy budget relative to preindustrial times, can be written as a logarithmic function of the increase in CO_2 concentration relative to the preindustrial level or, equivalently, as a logarithmic function of the amount of carbon in the atmosphere relative to the amount in preindustrial times. Let S_t and $\bar{S}$, respectively, denote the current and preindustrial amounts of carbon in the atmosphere. Then, forcing can be well approximated by the following equation.[26]

$$F_t = \frac{\eta}{\log 2}\log\left(\frac{S_t}{\bar{S}}\right). \tag{5}$$

The parameter η has a straightforward interpretation: if the amount of carbon in the atmosphere in period t has doubled relative to preindustrial times, forcing is η. If it quadruples, it is 2η, and so forth. An approximate value for η is 3.7, implying that a doubling of the amount of carbon in the atmosphere leads to a forcing of 3.7 watts per square meter on earth.[27]

We are now ready to present a relation between the long-run change in the earth's average temperature as a function of the carbon concentration in the atmosphere. Combining

[24]Formally, a flow rate per area unit is denoted flux. However, since we deal with systems with constant areas, flows and fluxes are proportional and the terms are used interchangeably.

[25]See Schwartz et al. (2010) who report that if earth were a blackbody radiator with a temperature of $288K \approx 15$ degrees Celsius, an increase in the temperature of 1.1 K would increase the outflow by 3.7 W/m^2, implying $\kappa^{-1} = 1.1/3.7 \approx 0.3$.

[26]This relation was first demonstrated by the Swedish physicist and chemist and 1903 Nobel Prize winner in Chemistry, Svante Arrhenius. Therefore, the relation is often referred to as the *Arrhenius's Greenhouse Law*. See Arrhenius (1896).

[27]See Schwartz et al. (2014). The value 3.7 is,however, not undisputed. Otto et al. (2013) use a value of 3.44 in their calculations.

equations (4) and (5) we obtain

$$T(F_t) = \frac{\eta}{\kappa} \frac{1}{\log 2} \log \left(\frac{S_t}{\bar{S}} \right). \tag{6}$$

As we can see, a doubling of the carbon concentration in the atmosphere leads to an increase in temperature given by $\frac{\eta}{\kappa}$. Using the Planck feedback, $\eta/\kappa \approx 1.1°$C. This is a modest sensitivity, and as already noted very likely too low an estimate of the overall sensitivity of the global climate due to the existence of positive feedbacks.

A straightforward way of including feedbacks in the energy budget is by adding a term to the energy budget. Suppose initially that feedbacks can be approximated by a linear term xT_t, where x captures the marginal impact on the energy budget due to feedbacks. The energy budget now becomes

$$\frac{dT_t}{dt} = \sigma \left(F + xT_t - \kappa T_t \right), \tag{7}$$

where we think of κ as solely determined by the Planck feedback. The steady-state temperature is now given by

$$T(F) = \frac{\eta}{\kappa - x} \frac{1}{\ln 2} \ln \left(\frac{S}{\bar{S}} \right). \tag{8}$$

Since the ratio $\eta/(\kappa - x)$ has such an important interpretation, it is often labeled the *Equilibrium Climate Sensitivity (ECS)* and we will use the notation λ for it.[28] Some feedbacks are positive but not necessarily all of them; theoretically, we cannot rule out either $x < 0$ or $x \geq \kappa$. In the latter case, the dynamics would be explosive, which appears inconsistent with historical reactions to natural variations in the energy budget. Also $x < 0$ is difficult to reconcile with the observation that relatively small changes in forcing in the earth's history have had substantial impact on the climate. However, within these bands a large degree of uncertainty remains.

According to the IPCC, the ECS is "likely in the range 1.5 to 4.5°C", "extremely unlikely less than 1°C", and "very unlikely greater than 6°C".[29] Another concept, taking some account of the shorter run dynamics, is the *Transient Climate Response (TCR)*. This is the defined as the increase in global mean temperature at the time the CO_2 concentration has doubled following a 70-year period of annual increases of 1%.[30] IPCC (2013b, Box 12.1) states that the TCR is "likely in the range 1°C to 2.5°C" and "extremely unlikely greater than 3°C."

[28] Note that equilibrium here refers to the energy budget. For an economist, it might have been more natural to call λ the *steady-state climate sensitivity*.

[29] See IPCC (2013a, page 81) and IPCC (2013b, Box 12.1). The report states that "likely" should be taken to mean a probability of 66-100%, "extremely unlikely" 0-5%, and "very unlikely" 0-10%.

[30] This is about twice as fast as the current increases in the CO_2 concentration. Over the 5, 10, and 20 year-periods ending in 2014, the average increases in the CO_2 concentration have been 0.54, 0.54, and 0.48 percent per year, respectively. However, note that also other greenhouse gases, in particular methane, affect climate change. For data, see the Global Monitor Division of the Earth System Research Labroratory at the U.S. Department of Commerce.

3.1.2 Non-linearities and uncertainty

It is important to note that the fact that $\frac{1}{\kappa - x}$ is a non-linear transformation of x has important consequences for how uncertainty about the strength of feedbacks translate into uncertainty about the equilibrium climate sensitivity.[31] Suppose, for example, that the uncertainty about the strength in the feedback mechanism can be represented by a symmetric triangular density function with mode 2.1 and endpoints at 1.35 and 2.85. This is represented by the upper panel of Figure 4. The mean, and most likely, value of x translates into a climate sensitivity of 3. However, the implied distribution of climate sensitivities is severely skewed to the right.[32] This is illustrated in the lower panel, where $\frac{\eta}{\kappa - x}$ is plotted with $\eta = 3.7$ and $\kappa = 0.3^{-1}$.

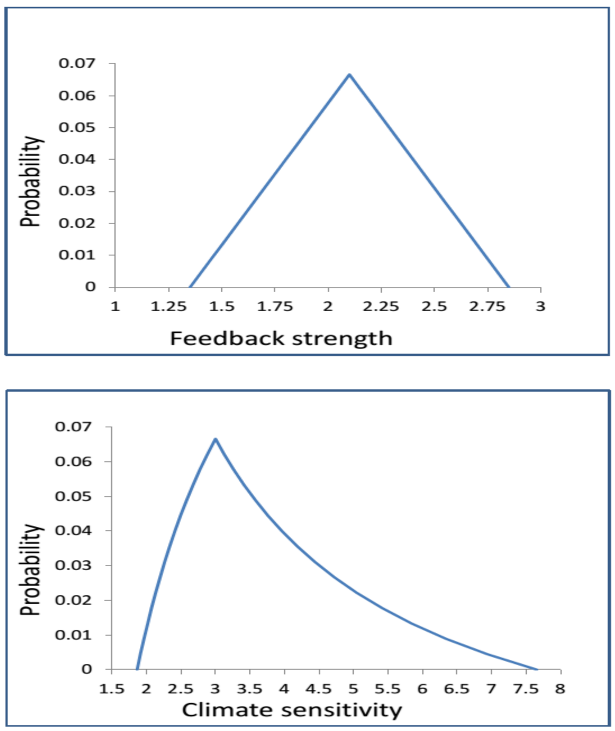

Figure 4: Example of symmetric uncertainty of feedbacks producing right-skewed climate sensitivity

The models have so far assumed linearity. There are obvious arguments in favor of relaxing this linearity. Changes in the albedo due to shrinking ice sheets and abrubt weaking of the Gulf are possible examples.[33] Such effects could simply be introduced by making x

[31] The presentation follows Roe and Baker (2007).

[32] The policy implications of the possibility of a very large climate sensitivity is discussed in Weitzman (2011).

[33] Many state-of-the-art climate models feature regional tipping points; see Drijfhouta at al. (2015) for a

in (7) depend on temperature. This could for example, introduce dynamics with so-called *tipping points*. Suppose, for example, that

$$x = \begin{cases} 2.1 \text{ if } T < 3^oC \\ 2.72 \text{ else} \end{cases}$$

Using the same parameters as above, this leads to a discontinuity in the climate sensitivity. For CO_2 concentrations below two times $\bar{S}$ corresponding to a global mean temperature deviation of 3 degrees, the climate sensitivity is 3. Above that tipping point, the climate sensitivity is 6. The mapping between $\frac{S_t}{\bar{S}}$ and the global mean temperature using equation (6) is shown in Figure 5.

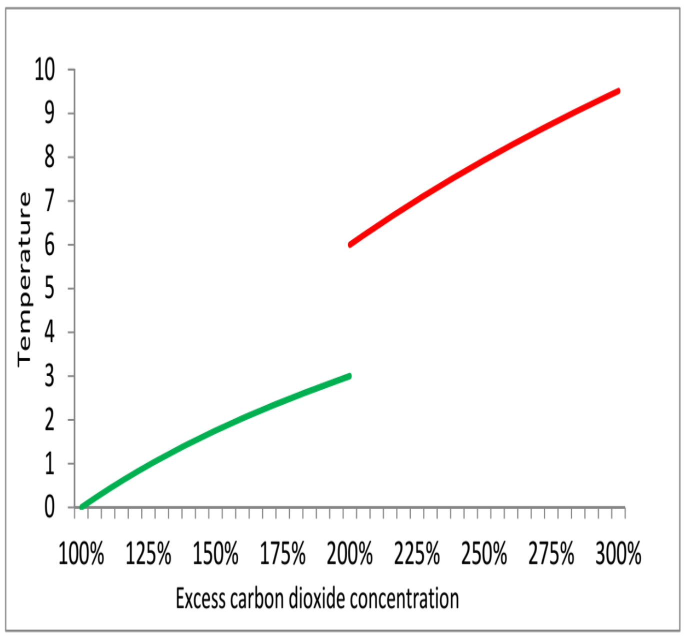

Figure 5: Tipping point at 3 K due to stronger feedback

It is also straightforward to introduce irreversibilities, for example by assuming that feedbacks are stronger (higher x) if a state variable like temperature or CO_2 concentration has ever been above some threshold value.

list. Currently, there is, however, no consensus on the existence of specific global tipping points at particular threshold levels; see Lenton et al. (2008), Levitan (2013), and IPCC (2013b, section 12.5.5).

3.1.3 Ocean drag

We have presented the simplest possible model of how the CO_2 concentration determines climate change. There are of course endless possibilities of extending this simplest framework. An example is to include another energy-budget equation. In equations (3) and (7), we described *laws of motion* for the atmospheric temperature, which heats much faster than the oceans. During the adjustment to a steady state, there will be net energy flows between the ocean and the atmosphere. Let T_t and T_t^L, respectively, denote the atmospheric and ocean temperatures in period t, both measured as deviations from the initial (pre-industrial) steady state. With two temperatures, we can define energy budgets separately for the atmosphere and for the oceans. Furthermore, allow for a variation in forcing over time and let F_t denote the forcing at time t. We then arrive at an extended version of equation (7) given by

$$\frac{dT_t}{dt} = \sigma_1 \left(F_t + xT_t - \kappa T_t - \sigma_2 \left(T_t - T_t^L \right) \right). \tag{9}$$

Comparing (9) to (7), we see that the term $\sigma_2 \left(T_t - T_t^L \right)$ is added. This term represents a new flow in the energy budget (now defined specifically for the atmosphere), namely the net energy flow from the atmosphere to the ocean. To understand this term, note that if the ocean is cooler than the atmosphere, energy flows from the atmosphere to the ocean. This flow is captured in the energy budget by the term $-\sigma_2 \left(T_t - T_t^L \right)$. If $T_t > T_t^L$, this flow has a negative impact on the atmosphere's energy budget and likewise on the rate of change in temperature in the atmosphere (the LHS). The cooler is the ocean relative to the atmosphere, the larger is the negative impact on the energy budget.

To complete this dynamic model, we need to specify how the ocean temperature evolves by using the energy budget of the ocean. If the temperature is higher in the atmosphere than in the oceans, energy will flow to the oceans, thus causing an increase in the ocean temperature. Expressing this as a linear equation delivers

$$\frac{dT_t^L}{dt} = \sigma_3 \left(T_t - T_t^L \right). \tag{10}$$

Equations (9) and (10) together complete the specification of how the temperatures of the atmosphere and the oceans are affected by a change in forcing.

We can simulate the behavior of the system once we specify the parameters of the system (σ_1, σ_2, σ_3, and κ all positive) and feed in a sequence of forcing levels F_t. Nordhaus and Boyer (2000) use $\sigma_1 = 0.226$, $\sigma_2 = 0.44$, and $\sigma_3 = 0.02$ for a discrete-time version of (9) and (10) defined as the analogous difference equations with a 10-year step. In 6 we show the dynamic response of this model to a constant forcing of $1W/m^2$ for $(\kappa - x)^{-1} = 0.81$. The lower curve represents the ocean temperature T_t^L, which increases quite slowly. The middle curve is the atmospheric temperature, T_t, which increases more quickly.

Clearly, the long-run increase in both temperatures is given by $\frac{1}{\kappa}$ times the increase in forcing, i.e., by 0.81 degrees Celsius. Most of the adjustment to the long-run equilibrium is achieved after a few decades for the atmosphere but takes several hundred years for the ocean temperature. Without the dragging effect of the oceans, the temperature increases

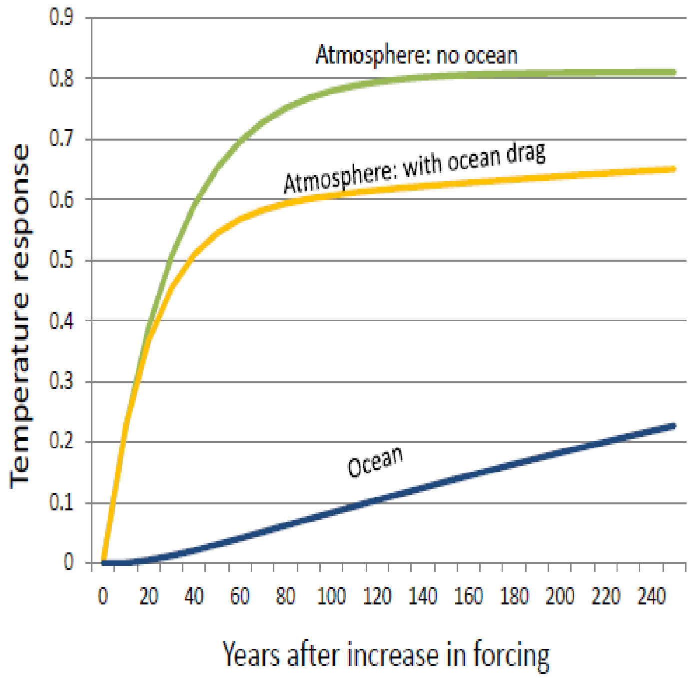

Figure 6: Increase in atmospheric and ocean temperatures after a permanent forcing of $1W/m^2$

faster, as shown by the top curve where we have set $\sigma_2 = 0$, which shuts down the effect of the slower warming of the ocean. However, we see that the time until half of the adjustment is achieved is not very different in the two cases.

3.1.4 Global circulation models

The climate models discussed so far are extremely limited in scope from the perspective of a climate scientist. In particular, they are based on the concept of an energy budget. Such models are by construaction incapable of predicting the large disparity in climates over the world. For this, substantially more complex general circulation models (GCMs) need to be used. Such models are based on the fact that the energy flow to earth is unevenly spread over the globe both over time and space. This leads to movements in air and water that are the drivers of weather events and the climate. These models exist in various degrees of complexity, often with an extremely large number of state variables.[34]

The complexity of general circulation models make them difficult to use in economics. In contrast to systems without human agents, such models do not contain any forward-looking

[34]See IPCC (2013b, chapter 9) for a list and discussion of GCMs.

agents. Thus, causality runs in one time direction only and the evolution of the system does not depend on expectations about the future. Therefore, solving such a complex climate model with a very large set of state variables may pose difficulties—in practice, because they are highly non-linear and often feature chaotic behavior—but not the kind of difficulties economists face when solving their dynamic models.

One way of modeling a heterogeneous world climate that does not require a combination of a very large state space and forward-looking behavior builds on *statistical downscaling*.[35] The output of large-scale dynamic circulation models or historical data is then used to derive a *statistical* relation between aggregate and disaggregated variables. This is in contrast to the actual nonlinear high-dimensional models because they do not feature randomness; the model output only looks random due to the nonlinearities. The basic idea in statistical downscaling is thus to treat a small number of state variables as sufficient statistics for a more detailed description of the climate. This works well due in part to the fact that climate change is ultimately driven by a global phenomenon: the disruption of the energy balance due to the release of green house gases, where CO_2 plays the most prominent role.

Let $T_{i,t}$ denote a particular measure of the climate, e.g., the yearly average temperature, in region i in period t. We can then estimate a model like

$$\begin{aligned} T_{i,t} &= \bar{T}_i + f\left(l_i, \psi_1\right) T_t + z_{i,t} \\ z_{i,t} &= \rho z_{i,t-1} + \nu_{i,t} \\ \operatorname{var}\left(\nu_{i,t}\right) &= g\left(l_i, \psi_2\right) \\ \operatorname{corr}\left(\nu_{i,t}, \nu_{j,t}\right) &= h\left(d\left(l_i, l_j\right), \psi_3\right). \end{aligned}$$

This very simple system, used for illustration mainly, explains downscaling conceptually. Here, $\bar{T}_i$ is the baseline temperature in region i. f, g, and h are specified functions parameterized by ψ_1, ψ_2, and ψ_3. $z_{i,t}$ is the prediction error and it is assumed to follow an AR(1) process. l_i is some observed characteristic of the region, e.g., latitude, and $d\left(l_i, l_j\right)$ is a distance measure. Krusell and Smith (2014) estimate such a model on historical data. The upper panel in Figure 7 shows the estimated function f with l_i denoting latitude. We see that an increase in the global mean temperature T_t has an effect on regional temperature levels that depends strongly on the latitude. The effect of a one degree Celsius increase in the global temperature ranges from 0.25 to 3.6 degrees. The lower panel in the figure shows the correlation pattern of prediction errors using d to measure Euclidian distance.

Now consider a dynamic economic model (where agents are forward-looking) with a small enough number of state variables that the model can be solved numerically. With one of these state variables playing the role of global temperature in the above equation system, one can imagine adding a large amount of heterogeneity without losing tractability, so long as the heterogeneous climate outcomes (e.g., the realization of the local temperature distribution) do not feed back into global temperature. This is the approach featured in Krusell and Smith

[35]See IPCC (2013b, chapter 9) for a discussion of statistical downscaling.

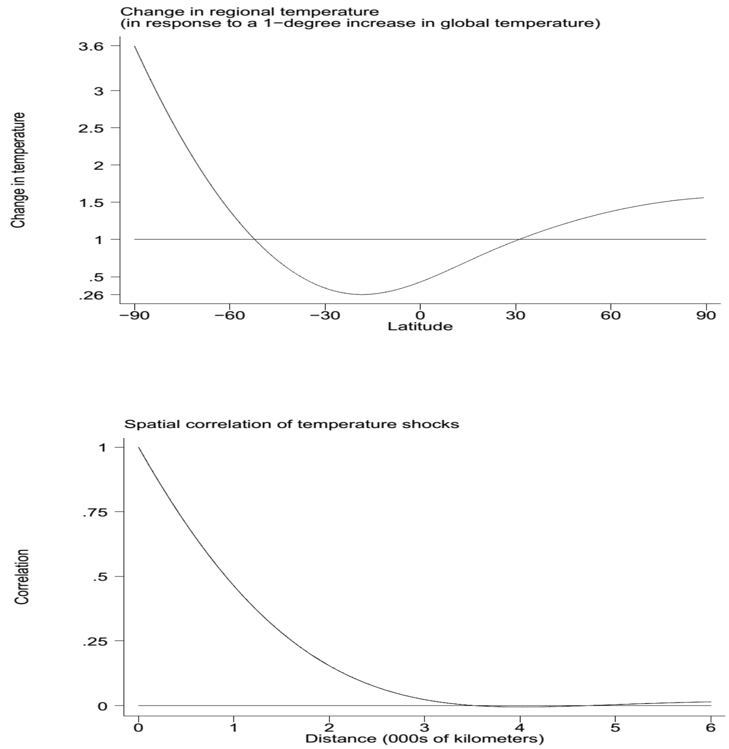

Figure 7: Statistical downscaling: regional climate responses to global temperature

(2015), whose model can be viewed as otherwise building directly on the models (static and dynamic) presented in the sections below in this chapter.[36]

3.2 Carbon circulation

We now turn to carbon circulation (also called the carbon cycle). The purpose of the modeling here is to produce a mapping between emissions of CO_2 and the path of the CO_2 concentration in the atmosphere. The focus on CO_2 is due to the fact that while other gases emitted by human activities, in particular methane, are also important contributors to the greenhouse effect, CO_2 leaves the atmosphere much more slowly. The half-life of methane is on the order of 10 years, while as we will see, a sizeable share of emitted CO_2 remains in the atmosphere for thousands of years.[37]

[36] Krusell and Smith (2015) actually allow some feedback, through economic variables, from the temperature distribution on global temperature but develop numerical methods that nevertheless allow the model to be solved.

[37] Prather et al. (2012) derive a half-life of methane of 9.1 years with a range of uncertainty of 0.9 years.

3.2.1 Carbon sinks and stores

The burning of fossil fuel leads to emissions of carbon dioxide into the atmosphere. The carbon then enters into a circulation system between different global reservoirs of carbon (carbon sinks) of which the atmosphere is one. In Figure 8, the carbon reservoirs are represented by boxes. The number in black in each box indicates the size of the reservoir in GtC, i.e., billions of tons of carbon. As we can see, the biggest reservoir by far is the intermediate/deep ocean, with more than 37,000 Gigatons of carbon. The vegetation and the atmosphere are of about the same size, around 600 GtC, although the uncertainty about the former is substantial. Soils represent a larger stock as does carbon embedded in the permafrost. Black arrows in the figure indicate pre-industrial flows between the stocks measured in GtC per year. The flows between the atmosphere and the ocean were almost balanced, implying a constant atmospheric CO_2 concentration.

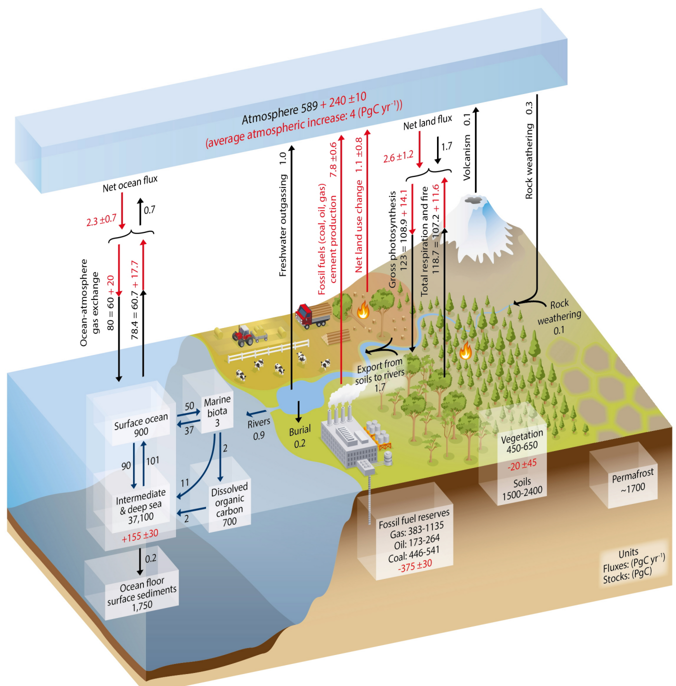

Figure 8: Global carbon cycle. Stocks in GtC and flows GtC/year. Source: IPCC (2013b, Figure 6.1).

By transforming carbon dioxide into organic substances, vegetation in the earth's biosphere induces a flow of carbon from the atmosphere to the biosphere. This is the photosynthesis. The reverse process, respiration, is also taking place in plants' fungi, bacteria, and animals. This, together with oxidation, fires, and other physical processes in the soil, leads to the release of carbon in the form of CO_2 to the atmosphere. A similar process is taking place in the sea, where carbon is taken up by phytoplankton through photosynthesis and released back into the surface ocean. When phytoplankton sink into deeper layers they take carbon with them. A small fraction of the carbon that is sinking into the deep oceans is eventually buried in the sediments of the ocean floor, but most of the carbon remains in the circulation system between lower and higher ocean water. Between the atmosphere and the upper ocean, CO_2 is exchanged directly. Carbon dioxide reacts with water and forms dissolved inorganic carbon that is stored in the water. When the CO_2-rich surface water cools down in the winter, it falls to the deeper ocean and a similar exchange occurs in the other direction. From the figure, we also note that there are large flows of carbon between the upper layers of the ocean and the atmosphere via gas exchange. These flows are smaller than, but of the same order of magnitude as, the photosynthesis and respiration.

3.2.2 Human influence on carbon circulation

Before the industrial revolution, human influence on carbon circulation was small. However, atmospheric CO_2 concentration started to rise from the mid-18th century and onwards, mainly due to the burning of fossil fuels and deforestation but also as a result of rising cement production.

In Figure 8, the red figures denote changes in the reservoirs and flows over and above pre-industrial values. The figures for reservoirs refer to 2011 while flows are yearly averages during the period 2000–2009. At the bottom of the picture, we see that the stock of fossil fuel in the ground has been depleted by 365 ± 30 GtC since the beginning of industrialization. The flow to the atmosphere due to fossil-fuel use and cement production is reported to be 7.8 ± 0.6 GtC per year. In addition, changed land use adds 1.1 ± 0.8 GtC per year to the flow of carbon to the atmosphere. In the other direction, the net flows from the atmosphere to the terrestrial biosphere and to the oceans have increased. All in all, we note that while the fossil reserves have shrunk, the amount of carbon in the atmosphere has gone from close to 600 to around 840 GtC and currently increases at a rate of 4 GtC per year. A sizeable but somewhat smaller increase has taken place in the oceans while the amount of carbon in the vegetation has remained largely constant.

We see that the gross flows of carbon are large relative to the additions due to fossil-fuel burning. Furthermore, the flows may be indirectly affected by climate change, creating feedback mechanism. For example, the ability of the biosphere to store carbon is affected by temperature and precipitation. Similarly, the ability of the oceans to store carbon is affected by the temperature. Deposits of carbon in the soil may also be affected by climate change. We will return to these mechanisms below.

3.2.3 The reserves of fossil fuel

The extent to which burning of fossil fuel is a problem from the perspective of climate change obviously depends on how much fossil fuel remains to (potentially) be burnt. This amount is not known and the available estimates depends on definitions. The amount of fossil resources that eventually can be used depends on estimates of future findings as well as on forecasts about technological developments and relative prices. Often, reserves are defined in successively wider classes. For example, the U.S. Energy Information Agency defines four classes for oil and gas. The smallest is *proved reserves*, which are reserves that geologic and engineering data demonstrate with reasonable certainty to be recoverable in future years from known reservoirs under existing economic and operating conditions. As technology and prices change, this stock normally increases over time. Successively larger ones are *economically recoverable resources*, *technically recoverable resources*, and *remaining oil and natural gas in place*.

Given different definitions and estimation procedures the estimated stocks differ and will change over time. Therefore, the numbers in this section can only be taken as indications. Furthermore, reserves of different types of fossil fuels are measured in different units, often barrels for oil, cubic meters or cubic feet for gas, and tons for coal. However, for our purpose, it is convenient to express all stocks in terms of their carbon content. Therefore non-trivial conversion must be undertaken. Given these caveats, we calculate from BP (2015) global proved reserves of oil and natural gas to be approximately 200 GtC and 100 GtC, respectively.[38] At current extraction rates, both these stocks would last approximately 50 years. Putting these numbers in perspective, we note that the atmosphere currently contains over 800 GtC. Given the results in the previous sections, we note that burning all proved reserves of oil and natural gas would have fairly modest effects on the climate.[39] Again using BP (2015), we calculate proved reserves of coal to around 600 GtC, providing more potential dangers for the climate.

Using wider definitions of reserves, stocks are much larger. Specifically, using data from McGlade and Ekins (2015) we calculate ultimately recoverable reserves of oil, natural gas and coal to close to 600 GtC, 400 GtC and 3000 GtC.[40] Rogner (1997) estimates coal reserves to be 3,500GtC with a marginal extraction cost curve that is fairly flat. Clearly, if all these reserves are used, climate change can hardly be called modest.

3.2.4 A linear carbon circulation model

A natural starting point is a linear carbon circulation model. Let us begin with a two-stock model as in Nordhaus and Boyer (2000). We let the variables S_t and S_t^L denote the amount of carbon in the two reservoirs, respectively: S_t for the atmosphere and S_t^L for the ocean. Emissions, denoted E_t, enter into the atmosphere. Under the linearity assumption, we assume that a constant share ϕ_1 of S_t flows to S_t^L per unit of time and, conversely, a

[38] BP (2015) reports proved oil reserves to 239,8 Gt. For conversion, we use IPCC (2006), table 1.2 and 1.3. From these, we calculate a carbon content of 0.846 GtC per Gt of oil. BP (2015) reports proved natural gas reserves to be 187.1 trillion m^3. The same source states an energy content of 35.7 trillion BtU per trillion m^3 equal to 35.9 trillion kJ. IPCC (2006) reports 15.3 kgC/GJ for natural gas. This means that 1 trillion m^3 natural gas contains 0.546 GtC. For coal, we use the IPCC (2006) numbers for antracite, giving 0.716 GtC per Gt of coal. For all these conversions, it should be noted that there is substantial variation in carbon content depending on the quality of the fuel and the numbers used must therefore be used with caution.

[39] As we will soon see, a substantial share of burned fossil fuel quickly leaves the atmosphere.

[40] See footnote 38 for conversions.

share ϕ_2 of S_t^L flows in the other direction implying

$$\begin{aligned}\frac{dS_t}{dt} &= -\phi_1 S_t + \phi_2 S_t^L + E_t, \\ \frac{dS_t^L}{dt} &= \phi_1 S_t - \phi_2 S_t^L.\end{aligned} \tag{11}$$

Equations (11) form a linear system of differential equations, similar to equations (9)–(10). However, there is a key difference: additions of carbon to this system through emissions get "trapped" in the sense that there is no outflow from the system as a whole, reflecting the fact that one of the characteristic roots of the system in (11) is zero.[41] This implies that if E settles down to a positive constant, the sizes of the reservoirs S and S^L will not approach a steady state, but will grow forever. If emissions eventually stop and remain zero, the sizes of the reservoirs will settle down to some steady-state values, but these values will depend on the amount of emissions accumulated before that. This steady state satisfies a zero net flow as per

$$0 = -\phi_1 S + \phi_2 S^L, \tag{12}$$

implying that

$$\frac{S}{S^L} = \frac{\phi_2}{\phi_1}$$

and that the rate of convergence is determined by the non-zero root $-(\phi_1 + \phi_2)$.

As we have seen above, CO_2 is mixed very quickly into the atmosphere. CO_2 also passes quickly through the ocean surface implying that a new balance between the amount of carbon in the atmosphere and the shallow ocean water is reached quickly.[42] The further transport of carbon to the deep oceans is much slower, motivating a third model reservoir: the deep oceans. This is the choice made in recent versions of the DICE and RICE models (Nordhaus and Sztorc, 2013), which use a three-reservoir linear system similar to (11).

3.2.5 Reduced-form depreciation models

Although the stock-flow model has a great deal of theoretical and intuitive appeal, it runs the risk of simplifying complicated processes too much. For example, the ability of the terrestrial biosphere to store carbon depends on temperature and precipitation. Therefore, changes in the climate may have an effect on the flows to and from the biosphere not captured in the model described above. Similarly, the storage capacity of the oceans depends (negatively) on the temperature. These shortcomings could possibly be addressed by including temperature and precipitation as separate variables in the system. Furthermore, also the processes involved in the deep oceans are substantially more complicated than what is expressed in the

[41] If we were to also define a stock of fossil fuel in the ground from which emissions are taken, total net flows would be zero. Since it is safe to assume that flows into the stock of fossil fuel are negligible, we could simply add an equation $\frac{dR_t}{dt} = -E_t$ to the other equations, which would thus capture the depletion of fossil reserves.

[42] This takes 1–2 years (IPCC, 2013b).

linear model. In particular, the fact that carbon in the oceans exists in different chemical forms and that the balance between these has an important role for the dynamics of the carbon circulation is ignored but can potentially be of importance.

An important problem with the linear specification (see, Archer, 2005, and Archer et al., 2009) is due to the so-called Revelle buffer factor (Revelle and Suess, 1957). As CO_2 is accumulated in the oceans, the water is acidified. This dramatically limits its capacity to absorb more CO_2, making the effective "size" of the oceans as a carbon reservoir decrease by approximately a factor of 15 (Archer, 2005). Very slowly, the acidity decreases and the pre-industrial equilibrium can be restored. This process is so slow, however, that it can be ignored in economic models. IPCC (2007, page 25, Technical Summary), take account of the Revelle buffer factor and conclude that "About half of a CO_2 pulse to the atmosphere is removed over a time scale of 30 years; a further 30% is removed within a few centuries; and the remaining 20% will typically stay in the atmosphere for many thousands of years". The conclusion of Archer (2005) is that a good approximation is that 75% of an excess atmospheric carbon concentration has a mean lifetime of 300 years and the remaining 25% remain several thousands of years.[43]

A way of representing this is to define a depreciation model. Golosov et al. (2014) define a carbon depreciation function. Let $1 - d(s)$ represent the amount of a marginal unit of emitted carbon that remains in the atmosphere after s periods. Then postulate that

$$1 - d(s) = \varphi_L + (1 - \varphi_L)\varphi_0 (1 - \varphi)^s . \tag{13}$$

The three parameters in (13) are easily calibrated to match the three facts in the above IPCC quote; we do this in Section 5. A similar approach is described in IPCC (2007, table 2.14). There,

$$1 - d(s) = a_0 + \sum_{i=1}^{3} \left(a_i e^{-\frac{s}{\tau_i}} \right), \tag{14}$$

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. 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 forever. It is important to note that this depreciation model is appropriate for a marginal emission at an initial CO_2 concentration equal to the current one (around 800 GtC). The parameters of the depreciation function should be allowed to depend on initial conditions and inframarginal future emissions. If emissions are very large, a larger share will remain in the atmosphere for a long time. To provide a measure for how sensitive the parameters are, note that of an extremely large emission pulse of 5,000 GtC, which is more than ten times the current accumulated emissions, around 40% remains after a thousand years, as opposed to half as much for a much smaller pulse.[44]

[43] Similar findings are reported in IPCC (2013, Box 6.1).
[44] See IPCC (2013b, Box 6.1).

3.2.6 A linear relation between emissions and temperature

As discussed above, it may be too simplistic to analyze the carbon circulation in isolation. The storage capacity of the various carbon sinks depends on how the climate develops. One might think that including these interactions would make the model more complicated. However, this does not have to be the case. In fact, there is evidence that various feedbacks and nonlinearity in the climate and carbon-cycle systems tend to cancel each other out, making the combined system behave in a much simpler and, in fact, linear way.[45] In order to briefly discuss this, let us defined the variable CCR_m (Carbon-Climate Response) as the change in the global mean temperature over some specified time interval m per unit of emissions of fossil carbon into the atmosphere over that same time interval

$$CCR_m \equiv \frac{T_{t+m} - T_t}{\int_t^m E_s ds}.$$

Given our previous discussions in this and the previous sections, one would think that this variable is far from a constant: the dynamic behavior of the climate and the carbon cycle will in general make the CCR_m depend on the length of the time interval considered. For example, since it takes time to heat the oceans, the temperature response could depend on whether the time interval is a decade or a century. Similarly, since also the carbon dynamics are slow, the extra CO_2 concentration induced by a unit of emission tends to be lower the longer the time interval considered. Furthermore, the the CCR_m might depend on how much emissions have already occurred; higher previous emissions can reduce the effectiveness of carbon sinks and even turn them into net contributors. The marginal effect on temperature from an increase in the CO_2 concentration also depends on the level of CO_2 concentration due to the logarithmic relation between CO_2 concentration and the greenhouse effect.

Quite surprisingly, Matthews et al. (2009) show that the dynamic and non-linear effects tend to cancel, making it a quite good approximation to consider the CCR_m as a constant, CCR, independent of both the time interval considered and the amount of previous emissions. Of course, knowledge about the value of CCR is incomplete but Matthews et al. (2012) quantify this knowledge gap and argue that a 90% confidence interval is between 1 and 2.5 degrees Celsius per 1000 GtC.[46] This means that we can write the (approximate) linear relationship

$$T_{t+m} = T_t + CCR \int_t^m E_s ds.$$

To get some understanding for this surprising result, first consider the time independence. We have shown in the previous chapter that when the ocean is included in the analysis, there is a substantial delay in the temperature response of a given forcing. Thus, if the CO_2 concentration permanently jumps to a higher level, it takes many decades before even half the final change in temperature has taken place. On the other hand, if carbon is released

[45] This subsection is based on Matthews et al. (2009).

[46] IPCC (2013a,b) defines the very similar concept, the Transient Climate Response to cumulative carbon Emissions (TCRE), and states that it is likely between 0.8 and 2.5 degrees Celsius per 1000 GtC for cumulative emissions below 2000 GtC.

into the atmosphere, a large share of it is removed quite slowly from the atmosphere. It happens to be the case that these dynamics cancel each other, 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 is balanced by the cooling effect of the oceans.

Second, for the independence of CCR with respect to previous emissions note that the Arrhenius law discussed in the previous chapter implies a logarithmic relation between CO_2 concentration and the temperature. 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 tend to have the property that the storage capacity of the sinks diminishes as more CO_2 is released into the atmosphere. These effects also 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 roughly the same proportion.

Given a value of CCR, it is immediate to calculate how much more emissions can be allowed in order to limit global warning to a particular value. Suppose, for example, we use a value of $CCR = 1.75$. Then, to limit global warming to 2 degrees Celsius, we cannot emit more than $(2/1.75) \times 1000 = 1140$ GtC, implying that only around 600 GtC can be emitted in the future. If, on the other hand, we use the upper limit of the 95% confidence interval ($CCR = 2.5$) and aim to reduce global warming to 2 degrees Celsius, accumulated emissions cannot be more than a total of 800 GtC of which most is already emitted.

3.3 Damages

In this section, we discuss how the economy is affected by climate change. Since economic analysis of climate change tends to rely on cost-benefit calculation, it is not only a necessary cornerstone of the analysis but arguably also a key challenge for climate economics. For several reasons, this is a very complicated area, however. First, there is an almost infinite number of ways in which climate change can affect the economy. Second, carbon emissions are likely to affect the climate for a very long time: for thousands of years. This implies that the quantitative issue of what weight to attach to the welfare of future generations becomes of key importance for the valuation. Third, global climate change can potentially be much larger than experienced during the modern history of mankind. Historical relations between climate change and the economy must therefore be extrapolated significantly if they are to be used to infer the consequences of future climate change. Fourth, many potential costs are to goods and services without market prices.

The idea that the climate affect the economy is probably as old as the economy itself, or rather as old as mankind. That the distribution of weather outcomes—the climate—affects agricultural output must have been obvious for humans since the Neolithic revolution. The literature on how the climate affects agriculture is vast and not reviewed here. It is also well known that in a cross-country setting, a hotter climate is strongly associated with less income per capita. Also within countries, such a negative relation between temperature and income per capita can be found (Nordhaus, 2006). However, Nordhaus (2006) also finds a hump-shaped relation between output density, i.e., output per unit of land area, and average temperature. This suggests that a method of adaptation is geographic mobility.

An overview is provided in Tol (2009). A more recent economic literature using modern methods emphasizing identification is now rapidly expanding. The focus is broad and climate change is allowed to have many different effects, including a heterogeneous effect on the economic productivity of different production sectors, effects on health, mortality, social unrest, conflicts, and much more. Dell, Jones, and Olken (2014) provide an overview of this newer literature.

Climate change thus likely has extremely diverse effects, involving a large number of different mechanisms affecting different activities differently. The effects are spatially heterogeneous and have different dynamics. Despite this, it appears important to aggregate the effects to a level that can be handled by macroeconomic models.[47]

3.3.1 Nordhaus's approach

Early attempts to aggregate the economic impacts of climate change were carried out in Nordhaus (1991).[48] Nordhaus (1992, 1993) constructed the path-breaking integrated assessment model named DICE, i.e., a model with the three interlinked systems—the climate, the carbon cycle, and the economy.[49] This is a global growth model with carbon circulation, and climate module, and a damage function. This very early incarnation of the damage function assumed that the economic losses from global warming were proportional to GDP and a function of the global mean temperature, measured as a deviation from the pre-industrial average temperature. Nordhaus's assumption in the first version of DICE was that the fraction of output lost was

$$D(T) = 0.0133 \left(\frac{T}{3} \right)^2 .$$

Nordhaus underlines the very limited knowledge that supported this specification. His own study (Nordhaus, 1991) studies a number of activities in the U.S. and concludes that these would contribute to a loss of output of 0.25 percent of U.S. GDP for a temperature deviation of 3 degrees Celsius. He argues that a reasonable guess is that the this estimate omits important factors and that U.S. losses rather are on the order of 1 percent of GDP and that the global losses are somewhat larger. Nordhaus (1992) cites Cline (1992) for an estimate of the power on temperature in the damage function but chooses 2 rather than the cited 1.3.

Later work (Nordhaus and Boyer, 2000) provided more detailed sectorial estimates of the damage function. Here, the aggregation includes both damages that accrue to market activities and those that could affect goods, services, and other values that are not traded. An attempt to value the risk of catastrophic consequences of climate change is also included. Obviously, this is an almost impossible task, given the little quantitative knowledge about tail risks. Nordhaus and Boyer use a survey, where climate experts are asked to assess the

[47] Macroeconomic modeling with large degrees of heterogeneity is developing rapidly, however. In the context of climate economy modeling, see e.g., Krusell and Smith (2015) for a model with nearly 20,000 regions.

[48] Other early examples are Cline (1992) and Fankhauser (1994) and Titus (1992).

[49] DICE stands for Dynamic Integrated Climate-Economy model.

probability of permanent and dramatic losses of output at different increases in the global mean temperature.

The latest version of DICE (Nordhaus and Sztorc, 2013) instead goes back to a more ad-hoc calibration of the damage function. Based on results in a survey in Tol (2009) and IPCC (2007) depicted in Figure 9, they postulate a damage function given by

$$D(T) = 1 - \frac{1}{1 + 0.00267T^2}. \tag{15}$$

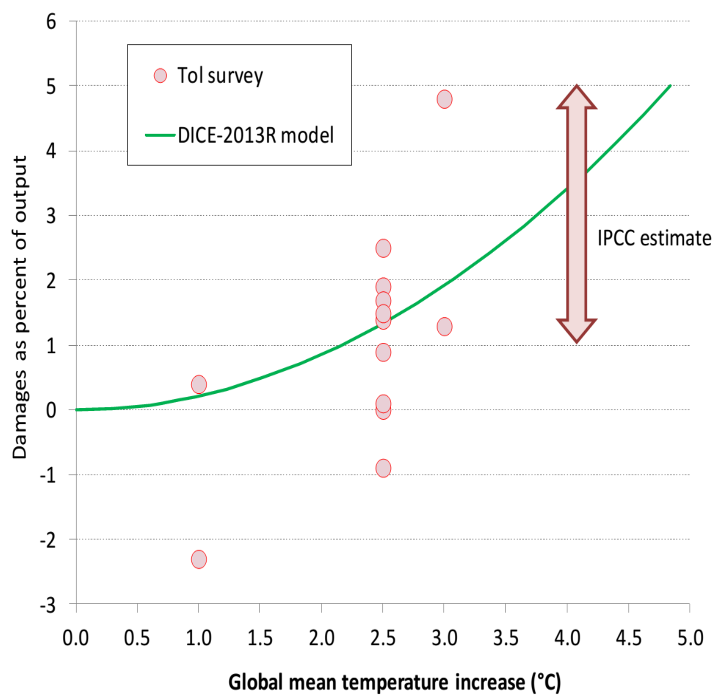

Figure 9: Global damage estimates. Dots are from Tol (2009). The solid line is the estimate from the DICE-2013R model. The arrow is from the IPCC (2007b, page 17). Reprinted from Nordhaus and Sztorc (2013).

The functional form in (15) is chosen so that damages are necessarily smaller than 1 but for the intended ranges of temperature, it may be noted that $1 - \frac{1}{1+0.00267T^2} \approx 0.023 \left(\frac{T}{3}\right)^2$.[50] Thus, the functional form remains similar to the first version of DICE but the estimated damages at three degrees have increased from 1.3 to 2.3% of global GDP.

Nordhaus has also developed models with multiple regions, RICE (Regional Integrated Climate-Economy model). The later versions of this model have different damage functions

[50]It is important to note that Nordhaus and Sztorc (2013) warn against using their damage function for temperature deviations over three degrees Celsius.

defined for 12 regions. Here the linear-quadratic function of the global mean temperature is appended with a threshold effect at a four-degree temperature deviation: at this level, the exponent on the temperature is increased to six. Separate account is also taken for sea-level rise, whose damages are described using a linear-quadratic function.

Similar aggregate damage functions are used in other global integrated assessment models; prominent examples are WITCH, FUND, and PAGE.[51]. Specifically, WITCH has quadratic, region-specific damage functions for eight global regions. FUND uses eight different sectorial damage functions defined for each of 16 regions. PAGE, which was used in the highly influential Stern report (Stern, 2007), uses four separate damage functions for different types of damages in each of eight regions. A special feature of the damage functions in FUND is that the exponent on the global mean temperature is assumed to be a random variable in the interval $[1.5 - 3]$.

3.3.2 Explicit damage aggregation

The damage functions described so far has only been derived to a limited degree from a "bottom-up approach" where explicit damages to particular regions and economic sectors are defined and aggregated. To the extent that such an approach has been used, the final results have been adjusted in an ad-hoc manner, often in the direction of postulating substantially larger damages than found in the explicit aggregation. Furthermore, the work has abstracted from general-equilibrium effects and simply added estimated damages sector by sector and region by region. Obviously this is problematic as the welfare consequences of productivity losses to a particular sector in a particular region depend on the extent to which production can move to other regions or be substituted for by other goods.

An example of a detailed high-resolution modeling of climate damages where (regional) general equilibrium effects are taken into account is the PESETA project, initiated by the European Commission.[52] Damages estimated are for coastal damages, flooding, agriculture, tourism, and health in the European Union. A reference scenario there is a 3.1-degree Celsius increase in the temperature in the EU by the end of this century relative to the average over 1961–1990. The resulting damages imply an EU-wide loss of 1.8 percent of GDP. The largest part of this loss is due to higher premature mortality in particular in south-central EU.[53] In the northern parts of the EU, welfare gains associated mainly with lower energy expenditures are approximately balanced by negative impacts in human health and coastal area damages.[54] Clearly, these effects are small relative to the expectations for economic growth over this period as well as compared to fears of dramatic impacts often expressed in the policy debate about climate change.

[51] See Bosetti et al. (2006), Tol (1995), and Hope et al., (1993) for descriptions of WITCH, FUND, and PAGE, respectively.

[52] See Ciscar et al. (2011) for a short description.

[53] France, Austria, Czech Republic, Slovakia, Hungary, Slovenia, and Romania.

[54] This area is defined by Sweden, Finland, Estonia, Lithuania, Latvia, and Denmark.

3.3.3 Top-down approaches

An alternative approach to the bottom-up approach is to estimate a reduced-form relation between aggregate measures like GDP, consumption, and investments and climate. The idea here is to associate natural historical variation in climate to changes in the aggregate variables of interest. Most of this work thus focuses on short-run changes in temperature as opposed to climate change. Examples of this approach are Dell, Jones, and Olken (2012) who examine how natural year-to-year variation in a country's temperature affects its GDP. Using data from 1950–2003, they find strong and persistent effects of a temporary deviation in temperature, with a point estimate of 1.4 percent of GDP per degrees Celsius—*but only in poor countries*. A similar result, but using global variation in the temperature, is reported by Bansal and Ochoa (2012). Krusell and Smith (2015), however, find that positive temperature shocks affect the level of GDP but not its rate of growth, and they do not find evidence of a difference between rich and poor countries.

Another approach is taken in Mendelsohn et al. (1994). Instead of attempting to measure a direct relation between climate and output, i.e., estimating a production function with climate as an input, the focus is here on agricultural land prices. They label this a *Ricardian approach*. The advantage of this is that adaptation, for example changed crops, can be taken into account. The finding is that higher temperature, except in the fall, is associated with lower land prices. However, the strength in this relation is lower than what is suggested by estimates based on traditional production function analysis. This indicates that the latter underestimates the potential for adaptation.

Burke et al. (2015) estimate empirical relations between economic activity and climate by assuming that local damages are a function not of global temperature but of local temperature. That is, heterogeneity here is built in not in terms of differences in responses to global temperature changes but simply through how local climates are very different to start with. If a region is very cold, warming can be beneficial, and if a region is very warm, further warming will likely be particularly detrimental. In line with Nordhaus (2006), a hump-shaped relation between economic activity and average yearly temperature is then estimated, with a maximum around 12-13 degrees Celsius. If this relation is taken as a causal relation from climate to productivity, it can be used to measure the long-run consequences of climate change. However, the use of the relation to evaluate long-run consequences precludes a study of short- and medium-run costs. This holds in particular for the costs of geographic reallocation of people, an area where little is known. In line with Burke et al. (2015), Krusell and Smith (2015) postulate a unique damage function of local temperature for a large number of regions and impose the condition that this function generate Nordhaus's estimated aggregate damages for warming of 1, 2.5, and 5 degrees Celsius. They find a somewhat lower ideal temperature than do Burke et al. but that the losses from having local temperatures far from the ideal value can be very large.

3.3.4 Remarks

The section on damage measurements in this chapter is short and does not do full justice to the literature. However, even a very ambitious survey would make clear that the research area of damage measurement is at a very early stage and provides frustratingly little guidance for cost-benefit analysis. On the one hand, most of the evidence points to rather limited aggregate global damages, at last for moderate degrees of climate change. On the other hand, it is not possible to rule out large damages, at least if climate change is more than moderate. After all, if the damages from climate change cannot be measured and quantified, how can we arrive at policy recommendations? There is no quick answer; much more research on this is clearly needed. In the absence of more solid evidence there is unfortunately ample room for extreme views—on both sides of the climate debate—to make claims about damage functions that support any desired action. We therefore prefer to proceed cautiously and to base our calibrations of damage functions on the evidence that, after all, has been gathered and put together. But before moving on to a description of the approach we take here, let us make some remarks about some mechanisms we will be abstracting from and that could nevertheless prove to be important.

One aspect of damages concerns the long run: is it possible that a warmer climate hurts (or helps?) long-run economic development, and might it even affect the growth rate of output? The work by Dell et al. (2012) as well as Burke et al. (2015) suggest such effects might be present on the local level, though without providing evidence on mechanisms. For an overall growth-rate effect on world GDP, there is as far as we know no evidence. Clearly, any growth effects—by naturally adding effects over time—will lead to large total effects, and that regions at different ends of the distribution would diverge in their levels of production and welfare, and it is not clear that our growth data support this conclusion. At the same time, the large implied effects make it all the more important to dig deeper and understand whether growth effects could actually be present. To be clear, our null hypothesis is that there are no effects on long-run growth rates of climate change.

Relatedly, it is common—following Nordhaus's lead—to describe damages as essentially proportional to GDP. This formulation, which to an important extent appears to be untested, has some important implications. One is that higher GDP ceteris paribus leads to higher damages. Another is that, since lower GDP means less to consume and consumption (typically, in macroeconomic models) is assumed to be associated with diminishing marginal utility, the welfare losses from a unit of damage measured in consumption units are lower the higher is GDP. Thus, if future generations will have higher GDP than we have today, there are two opposing forces: the total damages in consumption units will be higher but each of those units will hurt future generations less. As we shall see, under reasonable assumptions on utility, those two forces cancel, or roughly cancel. However, there are various ways to depart from Nordhaus approach. One is to assume that damages occur in consumption units but are not (linearly) proportional to GDP (e.g., our capital stock could be damaged). Another is to think of damages as occurring to specific consumption bundles that may not display the same degree of diminishing returns as consumption as a whole (examples include effects on leisure, health, or longevity). Damages can also occur in the form of changes in the distribution of resources and in other ways that are not easily thought of in terms of an aggregate damage function proportional to GDP.

Climate change can also lead to social conflict, as it changes the values of different activities and, more generally, “endowments”. One channel occurs via migration: if a region is hit hard by a changed climate and people migrate out, history tells us that the probability of conflict in the transition/destination areas will rise (see e.g., Miguel et al., 2004, Burke et al., 2009, Jia, 2014, and La Ferrara and Harari, 2014, and Burke et al., 2015, for an overview). At the same time, migration is also one of the main ways humans have to adapt to a changing climate. In fact, one view is that “populations can simply move toward the poles a bit” and hence drastically limit any damages from warmer weather; see Desmet and Rossi-Hansberg (2015) for an analysis that takes the migration mechanism seriously (see also Brock et al. (2014)). A related aspect is that climate change will have very diverse effects. It may be true that aggregated damages are small as a share of GDP and that those who lose a lot could be compensated by other, losing less or nothing at all. However, such global insurance schemes do not exist, at least not presently. The extent to which there are compensating transfers will likely to greatly impact any reasonable cost-benefit analysis of climate change and policies against it.

Tipping points are often mentioned in the climate-economy area and above we discussed some possible tipping points in the natural-science sections. Damages can also have tipping points in various ways and on some level a tipping point is simply a highly nonlinear damage function. One example leading to tipping points is the case of rising sea levels due to the melting of the ice caps. Clearly, some areas may become flooded and uninhabitable if the sea level rises enough, and the outcome is thus highly non-linear. This argument speaks clearly in favor of using highly non-linear damage functions on the local level, at least when it comes to some aspects of higher global temperatures. However, the sea-level rise equally clearly does not necessarily amount to a global nonlinearity in damages. Suffice it to say here that very little is known on the topic of global tipping points in damages. We will proceed with the null that a smooth convex aggregate damage function is a good starting point: we follow Nordhaus in this respect as well.

On an even broader level, let us be clear that different approaches are needed in this area. Bottom-up structural approaches like the PESETA project are very explicit and allow extrapolation, but they are limited to a certain number of factors and may miss important other mechanisms. Reduced-form micro-based approaches allow credible identification but may also miss important factors and general-equilibrium effects. Reduced-form aggregate approaches are less likely to miss mechanisms or general-equilibrium effects but necessarily involve a small number of observables and are much harder to interpret and extrapolate from. There is, we believe, no alternative at this point other than proceeding forward on all fronts in this important part of the climate-economy research area.

3.3.5 The operational approach: a direct relation

We now discuss a very convenient tool for the rest of the analysis in this chapter: a way of incorporating the existing damage estimates into our structural integrated-assessment models. In section 3.1.1, we have noted that the relation between the CO_2 concentration and the greenhouse effect is concave (it is approximately logarithmic). The existence of feedbacks is likely to imply an amplification of the direct effect, but in the absence of known global threshold effects, the logarithmic relation is likely to survive. Above we have also noted that that modelers so far typically have chosen a convex relation between temperature and damages: at least for moderate degrees of heating, a linear-quadratic formulation is often chosen. Golosov et al. (2014) show that the combination of a concave mapping from CO_2 concentrations to temperature and a convex mapping from temperature to damages for standard parameterizations imply an approximately constant marginal effect of higher CO_2 concentration on damages as a share of GDP. Therefore, they postulate

$$D(T(S)) = 1 - e^{-\gamma\left(S-\bar{S}\right)}, \tag{16}$$

where S is the amount of carbon in the atmosphere at a point in time and $\bar{S}$ is its pre-industrial level. This formulation disregards the dynamic relation between CO_2 concentration and temperature. It also disregards the possibility of abrupt increases in the convexity of the damage mapping and threshold effects in the climate system. These are important considerations, in particular when large increases in temperature are considered. However, the approximation provides a very convenient benchmark by implying that the marginal damage measured as a share of GDP per marginal unit of carbon in the atmosphere is constant and given by γ.[55] Measuring S in billions of tons of carbon (GtC), Golosov et al. (2014) show that a good approximation to the damages used to derive the damage function in DICE (Nordhaus, 2007) is given by (16) with $\gamma = 5.3 \cdot 10^{-5}$.

In Figure 10, we show an exponential damage function with this parameter. Specifically, the figure shows the implied damage function plotted against temperature using the relationship $T(S) = 3\frac{\ln S - \ln S_0}{\ln 2}$, i.e., using a climate sensitivity of 3 degrees. Comparing this damage function to the Nordhaus function as depicted in Figure 9 above, we see that the former is slightly less convex.[56] While the exponential damage function implies a constant marginal loss of 0.0053 percent per GtC, the quadratic formulation implies increasing marginal loss up to approximately four degrees Celsius. However, in the important range 2.5 to 5.0 degrees Celsius, the marginal loss is fairly constant within the range 0.0053 and 0.0059 percent per GtC.

[55] Output net of damages is $e^{-\gamma(S-S_0)}Y$. Marginal damages as a share of net-of-damage output then become $[d((1-e^{-\gamma(S-S_0)})Y)/dS]/e^{-\gamma(S-S_0)}Y = \gamma$.

[56] Reducing the exponent on temperature to 1.5 and increasing the constant in front of temperature to 0.0061 in (15) produces a damage function very close to the exponential one.

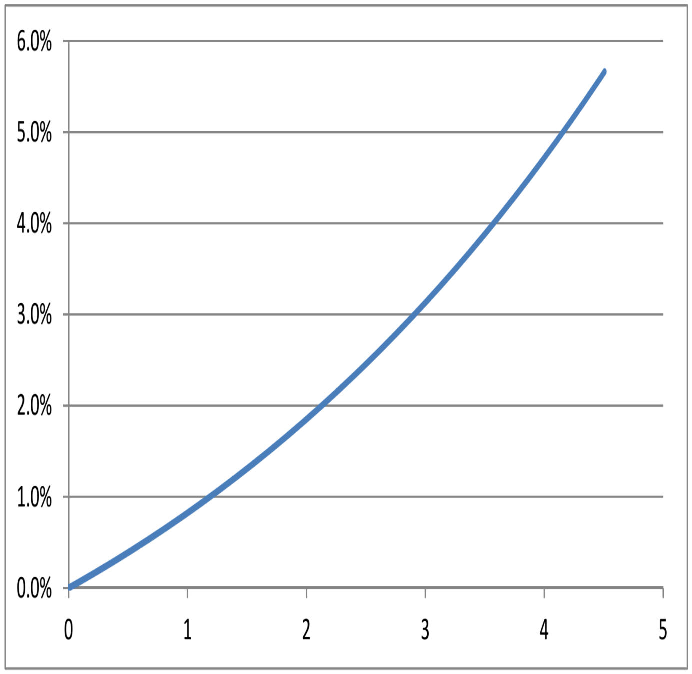

Figure 10: Damage function using $T(S) = 3\frac{\ln S - \ln S_0}{\ln 2}$ and $D(T(S)) = 1 - e^{-\gamma\left(S-\bar{S}\right)}$

4 A static global economy-climate model

Our discussion of integrated assessment models comes in two parts. The first part—in the present section—introduces an essentially static and highly stylized model, whereas the second part presents a fully dynamic and quantitatively oriented setup. The simple model in the present section can be viewed as a first step and an organizational tool: we can use it to formally discuss a large number of topics that have been studied in the literature. Moreover, for some of these topics we can actually use the model for a quantitative assessment, since it has most of the features of the macroeconomic structure in the later section. The model is thus a static version of Golosov et al. (2014) and it is also very similar to Nordhaus's DICE model.

We consider a world economy where the production of output—a consumption good—is given by

$$c = A(S)k^{\alpha}n^{1-\alpha-\nu}E^{\nu} - \zeta E.$$

Here, $A(S)$ denotes global TFP, which we take to be a function of the amount of carbon in the atmosphere, S. Moreover, we normalize so that S measures the excess carbon concentration, relative to a preindustrial average, $\bar{S}$. That is, the actual concentration is $S + \bar{S}$, whereas we will only need to use S in our modeling. The discussion in Section 10 allows us to use this notation and, moreover, to use a simple functional form that we argue is a decent approximation to the complex system mapping the amount of carbon in the atmosphere to temperature and then mapping temperature, with its negative impacts on the economy, to TFP. We will thus use

$$A(S) = e^{-\gamma S},$$

with $\gamma > 0$. Recall from the previous discussions that the map from S to T is logarithmic, so it features decreasing marginal impacts of increased atmospheric carbon concentration on temperature. The estimated mapping from T to TFP, on the other hand, is usually convex, so that the combined mapping actually can be described with the negative exponential function. Thus, damages are $(1 - e^{-\gamma S})k^{\alpha} n^{1-\alpha-\nu} E^{\nu}$, which is increasing and concave in S. (Note that we let energy, E, be capitalized henceforth, to distinguish it from Euler's number, e, used in the exponential damage function.) Though we argue above that this form for the damage function is a good one, it is straightforward to change it in this simple model, as we will below in one of our model applications. The exponential function is also useful because it simplifies the algebra and thus helps us in our illustrations. We will occasionally refer to γ as the *damage elasticity of output*.

The inputs in production include capital and labor, which we take to be exogenously supplied in the static model. The production function is Cobb-Douglas in the three inputs. As for capital and labor entering this way, we just use the standard macroeconomic formulation. The substitution elasticity between the capital-labor composite and energy is also unitary here, which is not far from available estimates of long-run elasticities, and we think of the static model as a short-cut representation of a long-run model. The short-run elasticity is estimated to be far lower, as discussed in Section 2.3.

We also see that the generation of output involves a cost ζE of producing energy. We will discuss in detail below how how energy is generated but the simple linear form here is useful because it allows us to illustrate with some main cases. One of these cases is that when energy is only produced from oil. Much of the oil (say, the Saudi oil) is very cheap to produce relative to its market price, so in fact we can think of this case as characterized by $\zeta = 0$. Oil exists in finite supply, so this case comes along with an upper bound on energy: $E \leq \bar{E}$.

A second case is that when energy comes from coal. Coal is very different because its market price is close to its marginal cost, so here we can think of ζ as a positive deep parameter representing a constant marginal cost in terms of output units (and hence the cost of producing energy in terms of capital and labor, and energy itself, has the same characteristics as does the final-output good). Coal is also only available in a finite amount but the available amount here is so large that we can think of it as infinite; in fact, if we were to use up all the coal within, say, the next 500 years, the implied global warming will be so high that most analysts would regard the outcome as disastrous, and hence the presumption in this case is that not all of the amount will be used up (and hence considering the available amount to be infinite is not restrictive). In reality, the supply of fossil fuel is of course not dichotomous: a range of fuels with intermediate extraction costs exists (see the discussion above in Section 3.2.3).

A third case is that with "green energy", where a constant marginal cost in terms of output is also a reasonable assumption. Finally, we can imagine a combination of these three assumptions and we will indeed discuss such possibilities below, but it is useful to consider coal and oil first separately first.

Turning to the mapping between energy use and atmospheric carbon concentration, the different energy sources correspond to different cases. In the case of oil and coal, we will simply assume that $S = \phi E + \bar{S}$, where $\bar{S}$ is the part of carbon concentration that is not of anthropogenic origin. As constants in TFP do not influence any outcomes here, we normalize $\bar{S}$ to equal zero. The equation thus states that carbon concentration is increased by the amount of emissions times ϕ. The constant ϕ represents the role of the carbon cycle over the course of a model period—which we will later calibrate to 100 years—and captures the fraction of the emissions during a period that end up in the atmosphere. A explained in Section 3.2, the depreciation structure of carbon in the atmosphere, though nontrivial in nature, can be rather well approximated linearly. Emissions, in turn, are proportional to the amount of fossil fuel used.[57]

We consider a consumer's utility function that, for now, only has consumption as an argument. Hence, so long as it is strictly increasing in consumption the model is complete.

We will discuss outcomes in a market economy of this sort where the consumer owns the capital and supplies labor under price taking, just like in standard macroeconomic models. Firms buy inputs, including energy, in competitive markets and energy is produced competitively. Formally, we can think of there being two sectors where isoquants have the same shape but where in the consumption-goods sector firms solve

$$\max_{k,l,E} e^{-\gamma S} k^{\alpha} n^{1-\alpha-\nu} E^{\nu} - wn - rk - pE,$$

where we denote wages and rental rates by w and r, respectively, and where p is the price of energy; the consumption good is the numéraire. In the energy sector the firms thus solve

$$\max_{k,l,E} p \frac{e^{-\gamma S}}{\zeta} k^{\alpha} n^{1-\alpha-\nu} E^{\nu} - wn - rk - pE.$$

It is straightforward to show, because the Cobb-Douglas share parameters are the same in the two sectors and inputs can be moved across sectors without cost, that this delivers $p = \zeta$ (whenever energy is nontrivially produced, so in the coal and green-energy cases, $1/\zeta$ becomes the TFP in the energy sector relative to that in the final-goods sector). Note also that GDP, y, equals the production of the consumption good, since energy here is an intermediate input.[58]

[57] Constants of proportion are omitted and are inconsequential in this simple model. In a more general framework one must take into account how oil and coal differ in the transformation between the basic carbon content and the resulting emissions as well as how they differ in productive use. We discuss these issues below when we consider coal and oil jointly.

[58] We do not explicitly have a home sector demanding energy. We take GDP to include housing services and to the extent they can be thought of as produced according to the market production function, these energy needs are included, but other home energy needs (such as gasoline for cars) are simply abstracted from.

Note that in both of the above profit maximization problems firms do not choose S, i.e., they do not perceive an effect on TFP in their choice, even though $S = \phi E$ in equilibrium. This is as it should be: the climate damage from emissions are a pure, and global, externality. Markets fail to take this effect into account and optimal policy should be designed to steer markets in the right direction.

The associated planning problem thus reads

$$\max_E e^{-\gamma\phi E} k^\alpha n^{1-\alpha-\nu} E^\nu - \zeta E;$$

here, clearly, the externality is taken into account. In the case of oil, for which $\zeta = 0$ is assumed, there is an additional constraint for the planner, namely that $E \leq \bar{E}$.

We will now discuss the solution to this problem for the different cases, starting with the case of oil.

4.1 The case of oil

Here, $\zeta = 0$ and the energy-producing sector is trivial. Under laissez-faire, all of the oil is supplied to the market and its price will be given by its marginal product: $p \equiv \bar{p} = \nu e^{-\gamma\phi\bar{E}} k^\alpha n^{1-\alpha-\nu} \bar{E}^{\nu-1}$. To the extent $\bar{E}$ and $\gamma\phi$ are large, this will involve an allocation with large damages to welfare.

The planner, on the other hand, may not use up all the oil. It is straightforward to see that the solution to the planner's problem is a corner solution whenever $\bar{E} < \nu/(\gamma\phi)$: the planner would then, like the markets, use up all the available oil. Thus, there is a negative by-product of emissions but it is not, at its maximal use, so bad as to suggest that its use should be limited. (In fact, as we shall argue below, this is not an unreasonable conclusion for oil given a more general, calibrated structure.) If, on the other hand, $\bar{E} \geq \nu/(\gamma\phi)$, the solution is interior at an E that solves $E = \nu/(\gamma\phi)$.

4.1.1 Optimal taxes

What are the policy implications of this model? For a range of parameter values—for $\bar{E} < \nu/(\gamma\phi)$—no policy is needed. At the same time, taxes are not necessarily harmful: if we think of a unit tax on the use of oil (the firms, whose maximization problems are displayed above), so that users of oil pay $p + \tau$ per unit instead of p, all tax rates on oil less than $\bar{p}$ will deliver the optimal outcome (recall that the price of oil is a pure rent and the tax will therefore not affect the allocation). If the unit tax is exactly equal to $\bar{p}$, the market price of oil will be zero and oil producers are indifferent between producing or not. At this level there is still an equilibrium which delivers the optimal amount of oil, namely, when all producers choose to produce; otherwise, not enough oil is used.

So suppose instead that $\bar{E} > \nu/(\gamma\phi)$. Now a tax is needed, and the tax should be set so that $p = 0$; the price is zero at the socially optimal use of oil. Otherwise, no oil producer would restrict its production and the outcome would be $\bar{E}$. With a tax that is high enough that the price oil producers receive is zero, i.e.,

$$\tau = \nu e^{-\nu} k^{\alpha} n^{1-\alpha-\nu} \left(\frac{\nu}{\gamma\phi}\right)^{\nu-1},$$

there exists an equilibrium where precisely oil output is equal to $\nu/(\gamma\phi) < \bar{E}$.

4.1.2 Pigou and the social cost of carbon: a simple formula

A different way of getting at optimal policy here is to directly compute the optimal tax of carbon to be that direct damage cost of a unit of emission that is not taken into account by markets. This "marginal externality damage" is referred to in the literature as the *social cost of carbon*.[59] Moreover, the concept needs to be sharpened as the marginal externality damage can be computed at different allocations. We thus refer to the *optimal social cost of carbon* (OSCC) as the marginal externality damage of a unit of carbon emission evaluated at the optimal allocation. Let the optimal carbon amount be denoted E^*. Given Pigou's principle (Pigou, 1920), the OSCC is the way to think about optimal tax policy, so the tax to be applied is

$$\tau^* = \gamma\phi e^{-\gamma\phi E^*} k^{\alpha} n^{1-\alpha-\nu} (E^*)^{\nu},$$

since this is the derivative of the production function with respect to E where it appears as an externality, evaluated at E^*. The idea here is that this tax always allows the government to achieve the optimal outcome as a competitive equilibrium with taxes. To check that this is consistent with the brute-force analysis above, note first that for the case where $E^* = \bar{E}$, $\tau^* = \gamma\phi y^* < \nu y^*/\bar{E}$, where y^* is the optimal level of output. Thus, in equilibrium $p = \nu y^*/\bar{E} - \gamma\phi y^* > 0$, which is consistent with all oil being sold. For the case where $\bar{E} > \nu/(\gamma\phi)$, the optimal tax formula $\tau^* = \gamma\phi y^*$ implies, at the interior solution $E^* = \nu/(\gamma\phi)$, that $p + \tau^* = \nu y^*/E^* = \gamma\phi y^*$ so that $p = 0$. In other words, oil producers are indifferent between producing or not and E^* is therefore an optimal choice.

More generally, it is important to understand that Pigou pricing proceeds in two steps: (i) work out the optimal allocation, by solving the planning problem; and (ii) find the OSCC at this allocation and impose that tax. The first step is straightforward in principle but can be challenging if the planning problem is not convex, e.g., because the damage function is highly non-linear; in such a case, there may in particular be multiple solutions to the planner's first-order conditions. The second step has a potential difficulty if for a given tax there are multiple market equilibria. The simple baseline model here does not admit multiple equilibria for a given tax rate but such models are not inconceivable. One important case may be where there are coordination problems in which technology a society chooses—perhaps between a fossil and a green technology. We discuss such cases below.

The OSCC formula that we derived says that the optimal unit tax on carbon is proportional to the value of GDP at the optimal allocation, with a constant of proportionality

[59] The terminology is perhaps a little misleading since one might be led to think that the social cost is the sum of the private and the externality cost, i.e., the total cost. Instead "social" just refers to the part not taken into account by the market.

given by $\gamma\phi$. This result is an adaptation of the finding in Golosov et al. (2014) who derive the OSCC to be proportional to GDP in a much more general setting—a dynamic model that is calibrated to long-run data. The constant of proportionality in that model is also a function of other parameters relating to intertemporal preferences and the carbon cycle, both elements of which are dynamic modeling aspects. They also find this result to be very robust to a number of modeling changes. We shall review these results below but it is important to note already at this point that the core of the proportionality of the OSCC to output can be explained within the structure of the simple static model here.

4.1.3 Costs of carbon when taxes are not optimally set

Let us emphasize what the OSCC formula says and does not say. It tells us what the marginal externality cost of carbon is, provided we are in an optimal allocation. However, as there appear to be damages from global warming on net and very few countries have carbon taxes, the real world is not at an optimal allocation with respect to carbon use, and this fact suggests that there is another measure that might be relevant: what the marginal externality cost of carbon is today, in the suboptimal allocation. So let SCC, the *social cost of carbon*, be a concept that can be evaluated at any allocation, and suppose we look at the laissez-faire allocation.

One can, conceptually, define a SCC in more than one way. We will define it here as the marginal externality damage of carbon emissions *keeping constant behavior in the given allocation.* This is an important qualification, because if an additional unit of carbon is emitted into the atmosphere, equilibrium decisions will change—whether we are in an optimal allocation or not—and if the given allocation is not optimal, the induced changes in behavior will, in general, have a first-order effect on utility. Hence, an alternative definition would, somehow, take the induced changes in decisions into account. (If the allocation is optimal, these effects can be ignored based on an envelope-theorem argument.)

Let us thus compute the SCC for the case of our static model. Let us assume $\bar{E} > \nu/(\gamma\phi)$, so that there is excessive carbon use. Then the SCC, $\gamma\phi y$, is lower than the OSCC, $\gamma\phi y^*$. This is of course true since $y^* > y$ by definition: the planner's aim is precisely to maximize GDP in this simple model and laissez-faire markets fail to. Note also that the percentage difference between the two measures here is only a function of $\bar{E}$ and E^* and not of other indicators of the "size" of the economy, such as the amount of capital or labor.

Depending on the allocation we are looking at, the SCC may in general be higher or lower than the OSCC. There is also no presumption that the laissez-faire SCC have to be higher than the OSCC, which one might imagine if the marginal damages of emissions rise with the level of emissions. In the simple static model we just looked at here, however, the SCC is always be below the OSCC, because damages appear in TFP and are of a form that implies proportionality to output; the OSCC is chosen to maximize output in this setting, so the OSCC must then be higher than the SCC. In contrast, in our dynamic model in Section 5, although the SCC will be proportional to current output there too, the SCC will typically be above the OSCC. The reason there is that current output tends to be rising with higher current fossil use—it is primarily future output that will fall with current emissions, due to the incurred damages—implying that the SCC will be higher for higher levels of current emissions, and in particular the SCC will be higher than the OSCC since the latter dictates lower emissions. The comparison between the SCC and the OSCC is of practical importance: suppose we are in a laissez-faire allocation today, and that econometricians have measured SCC, i.e., damages from emissions based on our current allocation. Then this SCC measure is not of direct relevance for taxation; in fact, for the calibrated dynamic model, we would conclude that the optimal tax is below the econometricians' laissez-faire SCC estimates.

Most of the integrated-assessment literature on the social cost of carbon computes the cost as is indicated above, i.e., as a marginal cost at an optimal allocation and, more generally, comparisons between suboptimal and optimal allocations are rather unusual. The simple model here does allow such comparisons (as does the dynamic benchmark model described below). Thus define the *percentage consumption equivalent* as the value λ such that $u(c^*(1-\lambda)) = u(c)$, where c^* is the optimal consumption level and c any suboptimal level. Thus we can compute the laissez-faire value for λ in the simple model (i) to be 0, in the case where there is little enough carbon that all of it should be used ($\bar{E} > \nu/(\gamma\phi)$); and (ii), in the case where too much carbon is available, to satisfy

$$\begin{aligned} 1-\lambda &= \frac{e^{-\gamma\phi\bar{E}}k^{\alpha}n^{1-\alpha-\nu}\bar{E}^{\nu}}{e^{-\gamma\phi E^*}k^{\alpha}n^{1-\alpha-\nu}(E^*)^{\nu}} \\ &= e^{-\gamma\phi(\bar{E}-\frac{\nu}{\gamma\phi})}\left(\frac{\gamma\phi\bar{E}}{\nu}\right)^{\nu}. \end{aligned}$$

It is straightforward to verify that λ is increasing in $\bar{E}$ here. Note, however, that variables such as capital or labor do not enter, nor would the size of the population if it were introduced as a separate variable. So the "size" of the economy is not important for this measure.

4.2 The case of coal

Here, $\zeta > 0$ and we interpret E as coal. Laissez faire now always involves an interior solution for E and it is such that its (private) benefit equals its (private) cost $p = \zeta = \nu e^{-\gamma\phi E}k^{\alpha}n^{1-\alpha-\nu}E^{\nu-1}$. The planner chooses a lower amount of E: E^* is chosen so that the private benefit of coal minus its social cost equals its private cost:

$$-\gamma\phi e^{-\gamma\phi E^*}k^{\alpha}n^{1-\alpha-\nu}(E^*)^{\nu} + \nu e^{-\gamma\phi E^*}k^{\alpha}n^{1-\alpha-\nu}(E^*)^{\nu-1} = \zeta.$$

Notice here that when coal production becomes more productive (ζ falls), markets use more coal. The same is true for the planner, since the left-hand side of the above equation must be decreasing at an optimum level E^* (so that the second-order condition is satisfied): if ζ falls, the left-hand side must fall, requiring E^* to rise. Thus, technical improvements in coal production imply higher emissions.

4.2.1 Optimal taxes and the optimal social cost of carbon

Recall that, in the benchmark model, we think of coal as produced at a constant marginal cost in terms of output goods. Given that GDP, y, equals consumption or $e^{-\gamma\phi E}k^{\alpha}n^{1-\alpha-\nu}(E)^{\nu} -$

ζE, we can write the equation determining the optimal coal use as

$$-\gamma\phi(y^* + \zeta E^*) + \nu(y^* + \zeta E^*)/E^* = \zeta.$$

Hence, the optimal social cost of carbon, OSCC, is now $\gamma\phi y^*(1 + \zeta E^*/Y^*) = \gamma\phi y^*(1 + \frac{pE^*}{y^*})$. So it is not quite proportional to GDP (as it was in the case of oil) but rather to GDP plus firms' energy costs as a share of GDP. In practice, energy costs are less than 10% of GDP so a rule of thumb that sets the unit tax on coal equal to $\gamma\phi$ times GDP is still approximately correct.

4.2.2 Costs of carbon when taxes are not optimally set

What is the social cost of carbon at the laissez-faire allocation? It is $\gamma\phi(y + \zeta E)$, where y is laissez-faire GDP and E is laissez-faire carbon use, where we know that $y < y^*$ and $E > E^*$. Unlike in the case of oil, it is not clear whether this amount is smaller than the OSCC. The subtlety here is that the production of coal itself—an intermediate input—is hampered by a damage from climate change and thus the total externality from coal production is not just $\gamma\phi y$.

Consumption in the laissez-faire allocation is lower by a fraction λ that satisfies

$$1 - \lambda = \frac{e^{-\gamma\phi E} k^\alpha n^{1-\alpha-\nu} E^\nu - \zeta E}{e^{-\gamma\phi E^*} k^\alpha n^{1-\alpha-\nu} (\bar{E}^*)^\nu - \zeta E^*} = \frac{e^{-\gamma\phi E} k^\alpha n^{1-\alpha-\nu} E^\nu}{e^{-\gamma\phi E^*} k^\alpha n^{1-\alpha-\nu} (\bar{E}^*)^\nu} \frac{1-\nu}{1-\nu+\gamma\phi E^*},$$

where for the second equality we have used the equilibrium and planner's conditions, respectively. This expression is, unlike in the oil example, not explicit in terms of primitives. In general, it depends non-trivially on the size of the economy (of course, one can derive first-order conditions determining both E and E^* as a function of primitives but, for the latter, not in closed form).

4.2.3 Coal production only requires labor: our benchmark model

The case where coal is produced at a constant marginal cost in terms of output units is somewhat less tractable than the following alternative: coal production does not require capital and does not experience TFP losses from climate change. I.e., $E = \chi n_E$, where n_E is labor used in coal production and χ is a productivity parameter. This case is less realistic but given that energy production is a rather small part of firms' costs, it is convenient to use this specification for some purposes. In this case, we have output given as

$$y = e^{-\gamma\phi\chi n_E} k^\alpha (1 - n_E)^{1-\alpha-\nu} (\chi n_E)^\nu,$$

where total labor is now normalized: $n = 1$. In a laissez-faire allocation, we obtain that $n_E = \frac{\nu}{1-\alpha}$. The planner's allocation delivers optimal n_E^* from

$$-\gamma\phi\chi + \frac{\nu}{n_E^*} = \frac{1-\alpha-\nu}{1-n_E^*}.$$

It is straightforward to check that higher productivity in producing coal will increase emissions both in the laissez-faire allocation and in the optimal one.

Here, moreover, the social cost of carbon will be exactly proportional to GDP, as in the oil case: $\gamma\phi y^*$. The reason is that no indirect externality (through the production of fossil fuel) is involved in this case. Similarly, we can solve for laissez-faire measures of the cost of carbon and the welfare gap relative to the full optimum.

In what follows, when we focus on coal production or oil production that occurs at positive marginal cost, we will use this formulation since it allows for simpler algebra without forsaking quantitatively important realism.

4.3 Calibration

We will now calibrate the static model. This is of course heroic, given that so many aspects of the climate-economy nexus feature dynamics, but the point here is merely to show that the static model can be thought of in quantitative terms. It is also possible to compare the results here to those in the calibration of the fully dynamic model in Section 5.2.

So let the heroics begin by calling our model period 100 years. The benchmark model will have coal as the only source energy; as we will argue below, the stock of oil is rather small relative to the stock of coal, and we leave out renewables for now (in the dynamic model in the later section, we calibrate the production of energy services as using three sources: oil, coal, and green). We assume that coal is produced from labor alone as in the previous section, and the model thus has five parameters: γ, ϕ, α, ν, and χ. We thus need five observations to pin these down.

Output being a flow, we can straightforwardly set α and ν based on average historic data; we select 0.3 and 0.04, respectively (see Hassler et al., 2015). For the rest of the model parameters, let us relate the model's laissez-faire equilibrium to some other observables. We thus need to relate the equilibrium outcomes for the key variables—E, S, n_E, and y—to relevant data targets. A business-as-usual scenario with continuously increasing emissions can lead to increases of the temperature of around 4 degrees Celsius at the end of the century.[60] We interpret business as usual as our laissez-faire allocation. Let us use this information to find out the associated atmospheric concentration and emissions implied to generate this result, given our model. Arrhenius's formula gives

$$4 = \Delta T = \lambda \frac{\log \frac{S+\bar{S}}{\bar{S}}}{\log 2} = 3 \frac{\log \frac{S+600}{600}}{\log 2},$$

which allows us to solve for S as roughly 900 (GtC, in excess of the pre-industrial level 600). What are the corresponding emissions required? The model says $S = \phi E$. To select ϕ, use the estimated linear carbon depreciation formula in Section (3.2.5) above for computing the average depreciation from emitting a constant amount per decade. This amounts to a straight average of the consecutive depreciation rates and a value for ϕ of 0.48: the atmospheric carbon concentration rises by about one half of each emitted unit.

[60] Scenario RCP8.5 from IPCC's 5th Assessment Report.

To calibrate γ, let us take IPCC's upper estimate from Figure 9: at a warming of 4 degrees Celsius, they report a total loss of 5% of GDP. This is a flow measure and thus easy to map into our present structure. We thus need $e^{-\gamma S}$ to equal 0.95. This delivers $\gamma = 5.7 \cdot 10^{-5}$.

It remains to calibrate the parameter χ of the coal sector: its labor productivity. We can find it as follows. To reach 900 GtC, one needs to emit 900/0.48 units given the calculation above. In the model solution, $n_E = \nu/(1-\alpha)$. This means that $900/0.48 = \chi n_E = \chi \cdot 0.04/0.7$, which delivers a χ of approximately 32,813.

4.4 A few quantitative experiments with the calibrated model

We now illustrate the workings of the simple baseline model with coal with a few quantitative experiments. The chief purpose is to check robustness of the main results. Similar exercises could be carried out in all of the applications that follow (dealing with uncertainty, tipping points, tax-vs.-quota policy comparisons, and so on). We have left such quantitative analysis out for brevity but for each application it would be valuable to use the baseline calibration as discussed here, calibrate the new parameters relevant to the application, and then produce output in the form of tables and graphs. Indeed, such exercises appear ideal for teaching the present material.

Starting out from the calibrated benchmark, let us vary two of the parameters within reasonable ranges. We first look at the effect of the damage elasticity of output, varying it from a half of its estimated value to much higher ones. We see that a doubling of the damage elasticity a little more than doubles the GDP gap between laissez-faire and the optimum. For damages 10 times higher than the baseline estimate, the loss of GDP is almost a quarter of GDP.

Externality cost	$1 - \frac{y}{y^*}$
$\frac{1}{2}\gamma$	0.0037
γ	0.0177
2γ	0.0454
4γ	0.0983
6γ	0.1482
8γ	0.1954
10γ	0.2400

Turning to carbon depreciation, the robustness looks at a tighter range around the baseline calibration as compared to that for damages (the uncertainty about damages, after all, is much higher). Modest changes in carbon depreciation, as depicted in the table below, do nevertheless have some impact: a change of ϕ by 25 percentage point changes the output gap by about seven tenths of a percent and temperature by a little over half a degree.

1 - carbon depreciation	ΔT	$1-\frac{y}{y^*}$
0.75ϕ	3.2624	0.0107
0.95ϕ	3.8340	0.0164
ϕ	3.9658	0.0177
1.05ϕ	4.0938	0.0192
1.15ϕ	4.3388	0.0219
1.25ϕ	4.5707	0.0247

Finally, let us look at a more complete range of suboptimal taxes for the baseline calibration. The table and figures below illustrate by varying the tax, measured as a percent of GDP. Figure 4.4 below illustrates rather clearly that the model is more nonlinear for negative than for positive taxes: if the tax is turned into a sizeable subsidy the warming and output losses are substantial.

$(\tau/y)/(\tau^*/y^*)$	ΔT	$1-\frac{y}{y^*}$	n_E
-0.5	6.4084	0.0975	0.1294
0	3.9658	0.0177	0.0571
0.5	2.8365	0.0024	0.0353
1	2.2110	0	0.0254
2	1.5346	0.0035	0.0162

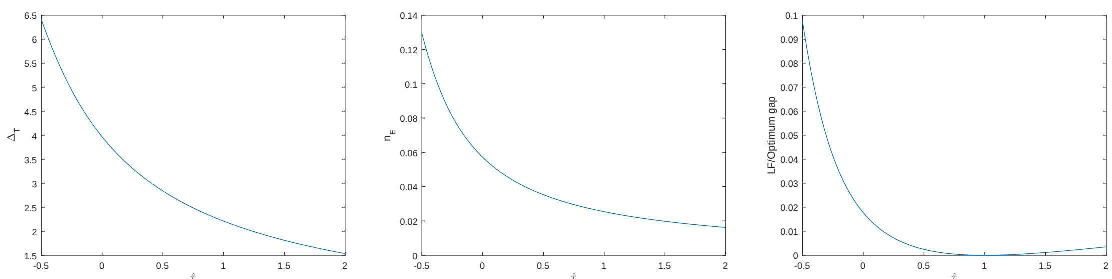

(a) Temperature change (b) Labor in coal sector (c) Output gap

Figure 11: Outcomes as a function of the tax-GDP rate, $\hat{\tau}$h

4.5 Summary: core model

We have built a simple static model which can be used to think about the key long-run aspects of carbon emissions and climate change. Though only a full dynamic, and much more complex, model can do the analysis of climate change full justice, our simple model does have some features that makes it quantitatively reasonable. The mapping from emissions to damages is described with a simple closed form but it captures the key features of this mapping in much more elaborate dynamic models, such as Nordhaus's DICE and RICE models. The role of fossil fuels in the economy is also described in a very rudimentary way but it too is the most natural starting point in dynamic quantitative models.

The simple model implies that the optimal social cost of carbon—the marginal externality damage at the optimal allocation—is proportional to GDP; this result is exactly true in some special cases of the model and approximately true otherwise. Also more generally, evaluated as a fraction of output, the (marginal) social cost of carbon (ignoring indirect effects on behavior of raising emissions) is independent of the allocation at which it is measured. This also means that the social cost of carbon is lower in the laissez-faire allocation than in the optimal allocation, because in the static model where damages appear to TFP optimal output by definition is higher than laissez-faire output. This feature will disappear in a dynamic model—where laissez-faire output tends to be higher (in the short run) than in the optimal allocation because less energy is used—and in a model where damages do not affect output, e.g., by affecting utility directly. We will of course look at these kinds of extensions below. Moreover, in the simple static model we formulated here, the utility loss from not using taxes to curb carbon use, expressed in percentages of consumption, is scale-independent.

Next, we use the simple model to address some issues that have featured prominently in the literature. These include the choice of policy instruments—in particular the comparison between price and quantity regulations (taxes vs. quotas)—along with extensions to consider utility damages, uncertainty, tipping points, technological change, and more.

4.6 Utility damages

We can, instead of or in addition to the damages to TFP, imagine that higher global temperatures affect welfare directly. This could occur in a variety of ways, through effects on health, the value of leisure, or more generally perceived life quality. Ignoring TFP damages for simplicity, consider first a utility function of a specific functional form:

$$u(c, E) = \log c - \gamma S,$$

where, again, $S = \phi E$ is carbon concentration in excess of the preindustrial level. Here, thus, atmospheric carbon concentration, and hence emissions, influence utility linearly, whereas consumption has decreasing marginal utility. This means that the value of one less unit of emissions in terms of consumption increases as the economy gets richer: $u_E/u_c = \gamma\phi c$. This implies, immediately, that the social cost of carbon in this economy is identical to that above: it is proportional to output. Thus, if the utility cost has the structure just assumed, the implications for how to tax carbon remain the same as in the more common case of TFP damages. In fact, we can now interpret the formulation with TFP damages as possibly coming from two sources: direct damages to TFP and utility damages.

With the remaining parts of the economy unchanged (except that we now view TFP as unaffected by emissions), we can solve for the laissez-faire equilibrium exactly as before. For sake of illustration, let us focus on coal and on the case where energy is produced linearly from labor. The social planner's problem is to solve

$$\max_{n_E} \log \left(k^{\alpha} (n - n_E)^{1-\alpha-\nu} (\chi n_E)^{\nu} \right) - \gamma \phi \chi n_E.$$

The problem simplifies to solving

$$\max_{n_E}(1-\alpha-\nu)\log(n-n_E)+\nu\log n_E-\gamma\phi\chi n_E.$$

The first-order condition gives $\frac{\nu}{n_E}=\frac{1-\alpha-\nu}{n-n_E}+\gamma\phi\chi$, which is the exact same equation as in the corresponding model with TFP damages.

What is the optimal tax/the OSCC in this model? The consumption-good firm's first-order condition for energy (assuming a unit tax τ) is $p+\tau=\nu k^{\alpha}(n-E/\chi)^{1-\alpha-\nu}E^{\nu-1}$, whereas the energy firm's first-order condition reads $p\chi=w$, with $w=(1-\alpha-\nu)k^{\alpha}(n-E/\chi)^{-\alpha-\nu}E^{\nu}$. This delivers $\frac{1-\alpha-\nu}{\chi}k^{\alpha}(n-E/\chi)^{-\alpha-\nu}E^{\nu}+\tau=\nu k^{\alpha}(n-E/\chi)^{1-\alpha-\nu}E^{\nu-1}$, from which we see that $\tau^*=\gamma\phi y^*$ is the optimal tax here as well.

More generally, the SCC at any consumption/energy allocation here can be obtained as $-u_E(c,E)/u_c(c,E)=\gamma\phi c$, and since consumption is GDP in the static model we again have that the SCC equals $\gamma\phi y$. We can, finally, define the utility loss in the laissez-faire allocation, measured in terms of a percentage consumption loss (i.e., from $u(c^*(1-\lambda),E^*)=u(c,E)$). We obtain $\log(1-\lambda)=\log\frac{c}{c^*}-\gamma\phi(E-E^*)$ and thus that $1-\lambda=e^{-\gamma\phi(E-E^*)}\frac{c}{c^*}$ which has the same form as before and, thus, is scale-independent.

4.7 Other damage functions

Our assessment in the section above on damages from climate change is that this is the sub-area in the climate-economy literature with the most striking knowledge gaps. Integrated assessment models differ to some extent in how they formulate damages as a function of climate (temperature) and how they parameterize their functions but the functional form used in Nordhaus's work (the DICE and RICE models) is the most common one. One possibility is that the overall damage *levels* are very different from the most common estimates in the literature, and another is that the functional-form assumptions are wrong. For this discussion, let us use the utility-damage formulation just outlined, and where we argued that $\log c-\gamma S$ is a formulation that is quantitatively close to that used by Nordhaus, given that this function should be viewed as a composition of the mapping from emissions to atmospheric carbon concentration and the mapping from the latter to damages. Let us therefore think about the choice of damage functions in terms of the more general formulation $\log c-\Gamma(S)$, with Γ being a more nontrivial function.[61] The function Γ, if truly described globally, should probably be increasing for positive values of S (since $S=0$ corresponds to the preindustrial concentration) and convex. For sufficiently low values of S (below 0) the function ought to be decreasing, since there is a reasonable notion of an "appropriate" climate: human beings could not survive if it is too cold either.

A concrete argument for a convex $\Gamma(S)$, rather than the linear one we use in our benchmark, is based on the arguments in Section 3.2.6 above: there appears to be an approximate reduced-form relationship between the global temperature and the unweighted cumulative amount of past anthropogenic emissions (since the industrial era began), which is *linear*.

[61] We maintain logarithmic curvature without loss of generality.

This was labeled the CCR (Carbon-Climate Response) formulation. Then take, say, Nordhaus's global damage function mapping temperature to output losses as given, and combine it with this approximate linear relationship. The resulting $\Gamma(S)$ must then be convex.[62]

With the more general damage function $\Gamma(S)$, all the above analysis goes through with the only difference being that $\Gamma'(S)$ now replaces γ earlier. Obviously, Γ could be calibrated so that $\Gamma'(S) = \gamma$ (with a standard calibration for γ) for current total emission levels, so the added insights here are about how the OSCC (and optimal tax) and the SCCs evolve as GDP evolves.

The SCC in this case becomes $\Gamma'(S)y$, where y again is GDP. Thus, to the extent Γ is convex, the optimal tax (as well as the SCC more generally) would not just be proportional to output but it would also increase with emissions; how much it would increase simply depends on the degree of convexity of Γ. Moreover, imagine an exogenous improvement in TFP. Such a shock would now increase the OSCC (the optimal tax) through two channels. The first channel was present before: a direct positive effect on y (leading to a higher tax by the same percentage amount). The second channel is an indirect effect via a higher demand for E. In terms of the decentralized economy, a higher TFP would, for a given tax, make firms demand a higher E, and since $\Gamma'(S)$ is increasing, this would then call for a further increase in the optimal tax rate.[63]

Similarly, the percentage consumption equivalent loss in welfare λ from remaining at laissez-faire can be computed from

$$\log(1-\lambda) = \log\frac{c}{c^*} - (\Gamma(S) - \Gamma(S^*)).$$

To the extent Γ is convex, this expression potentially increases faster in $S - S^*$ (and, more generally, depends on both these emission levels separately).

Now consider a highly non-linear damage function, and let us investigate whether such a case poses a difficulty for the Pigou approach to the climate problem. Consider the possibility that at a low level of emissions, so for a low S, the social cost of carbon is actually zero: $\Gamma'(S) = 0$. However, $\Gamma(S)$ is at the same time increasing rapidly for higher values of S, after which it again levels off and becomes flat: $\Gamma'(S) = 0$ also for high enough values of S. The latter amounts to a "disaster" outcome where more atmospheric carbon concentration actually does not hurt because all the horrible events that could happen have already happened given that S is so high. Here, though low emissions have a zero SCC, such low emissions are not what Pigou's formula would prescribe: they would prescribe that the SCC equal the net private benefits from emissions, and they are high for low emission levels. The net private benefits of emissions are, in particular, globally declining here (and, since damages appear

[62]Note, however, that the approximate linearity appears to be in somewhat of a conflict with Arrhenius's insight that the temperature change is proportional to the logarithm of the atmospheric carbon concentration (thus, a concave function). The conflict is not as strong as it seems, however. Our approximation that $\Gamma(S)$ is linear relies on a description of a carbon cycle that is rather realistic (e.g., has more complex dynamics) and that uses Arrhenius's formula, which still has widespread acceptance. The upshot of this really is that the just-mentioned convexity after all cannot be very strong.

[63]This discussion is a reminder that the optimal-tax formula $\tau^* = \Gamma'(S^*)y^*$ is not a closed form, since S^* and y^* are endogenous.

in preferences and not to production in the particular case under study, always positive). So instead, it is optimal to raise emissions to a point with a S^* such that $\Gamma'(S^*)$ is positive, perhaps one where Γ is increasing rapidly. The example shows that although a rapidly rising damage function in some sense poses a threat, the Pigou approach still works rather well. A key here is that for any given tax rate, the market equilibrium is unique; in the argument above, this manifested itself in the statement that the net private benefits from emissions are globally declining. They may not be, i.e., there may be multiple market equilibria, but such cases are unusual. We consider such examples in Section 4.14.1 below in the context of coordination problems in technology choice.

In conclusion, the model is well-designed also for incorporating "more convex" damage functions, and the qualitative differences in conclusions are not major nor difficult to understand. The key conclusion remains: more research on the determination and nature of damages—including the mechanisms whereby a warmer climate imposes costs on people—is of utmost importance in this literature, and integrated assessment modeling stands ready to incorporate the latest news from any such endeavors.

4.8 Tipping points

A tipping point typically refers to a phenomenon either in the carbon cycle or in the climate system where there is a very strong nonlinearity. I.e., if the emissions exceed a certain level, a more drastic effect on climate, and hence on damages, is realized. As discussed earlier in the natural-science part of the chapter, one can for example imagine a departure from the Arrhenius approximation of the climate model. Recall that the Arrhenius approximation was that the temperature increase relative to that in the pre-industrial era is proportional to the logarithm of the atmospheric carbon concentration (as a fraction of the preindustrial concentration), where the constant of proportionality—often labeled λ—is referred to as climate sensitivity. One way to express a tipping point is that λ shoots up beyond some critical level of carbon concentration. Another is that the carbon cycle has a non-linearity making ϕ a(n increasing) function of S, due to carbon sinks becoming less able to absorb carbon. Finally, we can imagine that damages feature a stronger convexity beyond a certain temperature point; for example, sufficiently high temperature and humidity make it impossible for humans and animals to survive outdoors.

Notice that all these examples simply amount to a different functional form for damages than that assumed above (whether damages appear to TFP or to utility). Thus, one can proceed as in the previous section and simply replace the total damage γS by a damage function $\Gamma(S)$, where this function has a strong nonlinearity. One could imagine many versions of nonlinearity. One involves a kink, whereby we would have a linear function $\gamma_{lo} S$ for $S \leq \underline{S}$ and $\gamma_{hi} S$ for $S > \underline{S}$, with $\gamma_{lo} << \gamma_{hi}$. A second possibility is simply a globally more convex (and smooth) function Γ. One example is Acemoglu et al. (2013), who assume that there is something labeled "environmental quality" that, at zero, leads to minus infinity utility and has infinitely positive marginal utility (without quantitative scientific references). One can also imagine that there is randomness in the carbon cycle or the climate, and this kind of randomness may allow for outcomes that are more extreme than those given by a simple (and deterministic) linear function γS. Finally, the $\Gamma(S)$ function could feature an irreversibility so that it attains a higher value if S ever has been above some threshold, thus even if S later falls below this threshold.

As discussed in the previous subsection, the formulation with a tipping point does not change the analysis of the laissez-faire equilibrium. It does, however, alter the social planner's problem. In particular, in place of γ as representing the negative externality of emissions in the planner's first-order condition we now have $\Gamma'(S)$ and this derivative may be very high. It is still possible to implement the optimum with a carbon tax, though it will no longer just be proportional to the optimal level of GDP and may respond nonlinearly to any parametric change, as discussed above. Suppose, for example, that γ becomes "infinite" beyond some $\underline{S}$. Then, from the perspective of a government choosing the optimal tax rate on carbon emissions, the objective function would have highly asymmetric payoffs from the tax choice: if the tax rate is chosen to be too low, the damage would be infinite, and more generally changes in the environment (such as increases in the capital stock or labor input, which would increase the demand for energy) would necessitate appropriate increases in the tax so as to avoid disaster.

Overall, in order to handle tipping points in a quantitative study based on an integrated assessment model one would need to calibrate the nonlinear damage function. In terms of our first example, how would one estimate $\underline{S}$? As we argued in the natural-science sections above, our interpretation of the consensus is that whereas a number of tipping points have been identified, some of which are also quantified, these are tipping points for rather local systems, or systems of limited global impact in the shorter run. To the extent there is a global (and quantitatively important) tipping point, there does not appear to be a consensus on where it would lie in S space. Therefore, at this point and in waiting for further evidence either on aggregate nonlinearities in the carbon cycle or climate system or in how climate maps into economic costs, we maintain a linear formulation (or, in the case damages appear in TFP, in the equivalent exponential form). Performing comparative statics on γ is, of course very important, and we return to it below.

4.9 Uncertainty

It is possible to analyze uncertainty in a small extension of the simple benchmark model. Suppose we consider a pre-stage of the economy when the decisions on emissions need to be made—by markets as well as by a fictitious planner. We then think of utility as of the expected-utility kind, and we begin by using a utility formulation common in dynamic macroeconomic models: $u(c) = \log c$. Thus, the objective is $E(\log(c))$. Uncertainty could appear in various forms, but let us simply consider a reduced-form representation of it by letting γ, the damage elasticity of output, be random. That is, in some states of nature emissions are very costly and in some they are not. Recall that the uncertainty can be about the economic damages given any temperature level or about how given emissions influence temperature.

For the sake of illustration, we first consider the simplest of cases: γ is either high, γ_{hi}, or low, γ_{lo}, with probabilities π and $1 - \pi$, respectively. The emissions decision has to be made—either by a planner or by actors in decentralized markets—ex ante, but there is no "prior period" in which there is consumption or any other decisions than just how high to make E. We consider the case of coal here, and with coal production requiring labor only, without associated TFP damages.

Looking at the planning problem first, we have

$$\max_E \pi \log \left(e^{-\gamma_{hi}\phi E} k^\alpha \left(1 - \frac{E}{\chi}\right)^{1-\alpha-\nu} E^\nu \right) + (1-\pi) \log \left(e^{-\gamma_{lo}\phi E} k^\alpha \left(1 - \frac{E}{\chi}\right)^{1-\alpha-\nu} E^\nu \right).$$

Save for a constant, this problem simplifies to

$$\max_E -(\pi\gamma_{hi} + (1-\pi)\gamma_{lo})\phi E + (1-\alpha-\nu)\log\left(1 - \frac{E}{\chi}\right) + \nu \log E.$$

A key feature of this maximization problem is that the damage elasticity appears only in expected value! This means that the solution of the problem will depend on the expected value of γ but not on any higher-order properties of its distribution. This feature, which of course holds regardless of the distributional assumptions of γ, will not hold exactly if coal/oil is produced with constant marginal cost in terms of final output (as in our very first setting above), but approximately the same solution will obtain in any calibrated version of the model since the fossil-fuel costs are small as a fraction of output.

Notice that the "certainty equivalence" result obtains here even though the consumer is risk-averse. However, it obtains for logarithmic utility only. If the utility function curvature is higher than logarithmic, the planner will take into account the variance in outcomes: higher variance will reduce the choice for E.[64] Formally, and as an example, consider the utility function $c^{1-\sigma}/(1-\sigma)$ so that the planner's objective is

$$\mathbf{E}_\gamma \frac{\left(e^{-\gamma E} k^\alpha \left(1 - \frac{E}{\chi}\right)^{1-\alpha-\nu} E^\nu \right)^{1-\sigma}}{1-\sigma}.$$

Since E is predetermined, we can write this as

$$\frac{\left(k^\alpha \left(1 - \frac{E}{\chi}\right)^{1-\alpha-\nu} E^\nu \right)^{1-\sigma}}{1-\sigma} \mathbf{E}_\gamma e^{-\gamma E(1-\sigma)}.$$

Assume now that γ is normally distributed with mean $\bar{\mu}$ and variance σ_μ^2. Then we obtain the objective

$$\frac{\left(e^{-\Gamma(E)} k^\alpha \left(1 - \frac{E}{\chi}\right)^{1-\alpha-\nu} E^\nu \right)^{1-\sigma}}{1-\sigma},$$

[64]The asset pricing literature offers many utility functions that, jointly with random processes for consumption, can deliver large welfare costs; several of these approaches have also been pursued in the climate-economy literature, such as in Barro (2013), Gollier (2013), Crost and Traeger (2014), and Lemoine (2015).

with

$$\Gamma(E) = -\bar{\gamma}E + \frac{\sigma_\mu^2 E^2 (1-\sigma)}{2}.$$

Thus, the objective function is a monotone transformation of consumption, with consumption determined as usual in this model except for the fact that the damage expression γE is now replaced by $\Gamma(E)$, a convex function for $\sigma > 1$ (higher curvature than logarithmic). To the extent that the variance σ_μ^2 is large and σ is significantly above 1, we thus have uncertainty play the role of a "more convex damage function", as discussed above. We see that the logarithmic function that is our benchmark does apply as a special case.

4.9.1 The Dismal Theorem

In this context let us briefly discuss the so-called *Dismal Theorem* derived and discussed by Weitzman in a series of papers (e.g., Weitzman, 2009; see also the discussion in Nordhaus, 2009). Weitzman provides conditions under which, in a rather abstract context where governmental action could eliminate climate uncertainty, expected utility is minus infinity in the absence of appropriate government action. Thus, one can (as does Weitzman) see this as an argument for (radical) government action. His result follows, very loosely speaking, if the uncertainty has fat enough tails, the risk aversion is high enough, and the government is able to entirely eliminate the tail uncertainty, but the details of the derivation depend highly on specifics. In our present context, a normal distribution for γ is clearly not fat-tailed enough and the only way for the government to shut down tail risk is to set E to zero. However, imagine that the economy has an amount of free green energy, denoted $\tilde{E}$, i.e., the production function is $e^{-\gamma E} k^\alpha \left(1 - \frac{E}{\chi}\right)^{1-\alpha-\nu} (\tilde{E} + E)^\nu$; then setting $E = 0$ still allows positive output. Now imagine that γ has a distribution with fat enough tails, i.e., one allowing infinitely high values for γ and slowly decreasing density there. Then expected utility will become infinite if σ is large enough.[65]

The Dismal Theorem is not connected to data, nor applied in a quantitatively specified integrated assessment model. It relies fundamentally on a shock structure that allows infinitely negative shocks (in percentage terms), and our historical data is too limited to allow us to distinguish the shape of the left tail of this uncertainty in conjunction with the shape of marginal utility near zero; at this point, it seems hard enough to be sure of the mean of the shocks.

[65] A simpler, reduced form setting is that where consumption is given by a t distribution (which has fatter tails than the normal distribution), representing some risk which in this case would be labeled climate risk. Then with power utility, $u(c) = c^{1-\sigma}/(1-\sigma)$, and if σ is high enough, the marginal utility at zero goes to infinity fast enough that expected utility is minus infinity. This point was original made by Geweke (2001). If the government can shut down the variance, or otherwise provide a lower bound for consumption, it would then be highly desirable.

4.10 Taxes vs. quotas

In the discussion above we have been focusing on a tax as the obvious candidate policy instrument. Indeed the damage externality is a pure externality for which the Pigou theorem applies straightforwardly. What are alternative policies? The Coase theorem applies too as well but it does not seem possible in practice to define property rights for the atmosphere (into which emissions can then be made, in exchange for a payment to the owner). What about regulating quantities? Indeed the "cap-and-trade" system, which is a quota-based mechanism, has been the main system proposed in the international negotiations to come to a global agreement on climate change. A cap-and-trade system is indeed in place in Europe since 2005.[66] There is a debate on whether a tax or a quota system is better, and here we will only allude to the main arguments. Our main purpose here, instead, is to make a few basic theoretical points in the comparison between the two systems. These points are also relevant in practice.

Before proceeding to the analysis, let us briefly describe the "-and-trade" part, which we will not subject to theoretical analysis. If a region is subject to a quantity cap—emissions cannot exceed a certain amount—the determination of who gets to emit how much, among the users of fossil energy in the region, must still be decided on. The idea is then to allocate *emission rights* and to allow trade in these rights. The trading, in theory at least, will then ensure that emissions are made efficiently. The initial allocation of emission rights can be made in many ways, e.g., through grandfathering (giving rights in proportion to historic use) or auctions. To analyze the trading system formally we would need to introduce heterogeneity among users, which would be straightforward but not yield insights beyond that just mentioned.

The first, and most basic, point in comparing quotas and taxes is that, if there is no uncertainty or if policies can be made contingent on the state of nature, both instruments can be used to attain any given allocation.[67] If a tax is used, the tax applies to all users; if a quota is used, regardless of how the initial emission rights are used, the market price of an emission right will play the role of the tax: it will impose an extra cost per unit emission and this cost will be the same for all users, provided the market for emission rights works well.

Second, suppose there is uncertainty and the policy cannot be made state-contingent. This is a rather restrictive assumption—there is no clear theoretical reason why policies could not change as the state of nature changes—but still an interesting one since it appears that political/institutional restrictions of this sort are sometimes present. To analyze this case, let us again consider uncertainty and an ex-ante period of decisions. To capture the essence of the restriction we assume that the only decision made ex ante is the policy decision. A policy could be either a unit tax or a quantity cap. We assume that the quantity cap is set so that it is always binding ex post, in which case one can view the government as simply

[66] The European Union Emission Trading System (EU ETS) was launched in 2005 covering about half the CO_2 emissions in the union (Ellerman and Buchner, 2007).

[67] This statement requires a qualification for taxes in the (rather unusual) cases for which a Pigou rule is not sufficient, as discussed already.

choosing the level of emissions ex ante.

The choice between a tax and a quota when there is uncertainty (or private information on the part of "the industry") has been studied extensively in the environmental literature since Weitzman (1974) and similar analyses are available in other parts of economics (e.g., William Poole, 1970). One can clearly provide conditions under which one policy or the other is better, along the lines of Weitzman's original paper. Weitzman considered a cost and a benefit of a pollutant, each of which depended on some random variable, and the two random variables were assumed to be independent. He then showed that what instrument would be best depended on the relative slopes of the marginal benefits and cost curves. Follow-up papers relaxed and changed assumptions in a variety of directions, but there appear to be no general theorems that apply in the climate-change application to conclude decisively in one way or the other. In fact, we know of no quantitatively parameterized dynamic model that looks at the issue so what we will do here is simply provide a straightforward example using our simple static model and then discuss a couple of separate, and we believe important, special cases.

For our example, we use one type of uncertainty only: that of the cost of producing fossil fuel, χ. With the calibrated model and a uniform distribution around the calibrated value for χ we obtained the ex-ante utility levels for a range of taxes and for a range of emissions, both committed to before the randomness is realized. Figure 12 shows the results: a range of tax values around the optimal tax outperform the optimal quota. In this case, the pre-committed tax rate is a fixed value. If it could be set as a proportion of output, which is ruled out now by assumption since the tax cannot be state-contingent but output will be, it would be fully optimal also ex post, since the best tax ex post is always a fraction $\gamma\phi$ of output. Apparently, the ex-post randomness of output is not significant enough to overturn this result. It is straightforward to look at other types of shocks. Shocks to γ deliver more similar welfare outcomes for (optimal and pre-committed) taxes and quotas.

Now suppose that we consider a case of a tipping point and that the uncertainty is coming from energy demand (through, say, a separate, exogenous and random TFP factor) or from the cost of coal production (through χ). If the tipping point is known to be $\underline{E}$, and $\Gamma(E)$ is equal to zero for $E < \underline{E}$ but positive and very high otherwise, what is then the best policy from an ex-ante perspective? Clearly, a policy with an emissions cap would simply be set at $\underline{E}$, a cap that may or may not bind ex post: if the demand for energy is low, or the cost of producing it is high, the ex-post market solution will (efficiently) be to stay below $\underline{E}$, and otherwise the cap will (efficiently) bind. A tax will not work equally well. One can set the tax so that the economy stays below the tipping point, but in case the energy demand is low, or its production costs are high, ex post, output will be inefficiently depressed. Thus, when we are dealing with *asymmetric* payoffs of this sort (relative to the amount of emissions), a quantity cap is better.

The previous example would have emissions rights trading at a positive price sometimes and at a zero price otherwise. Thus, the system with a quantity cap leads to a random cost for firms of emitting carbon dioxide (beyond the price the firms pay the energy producers). Variations in the supply of emissions rights, decided on by regulatory action, influence the

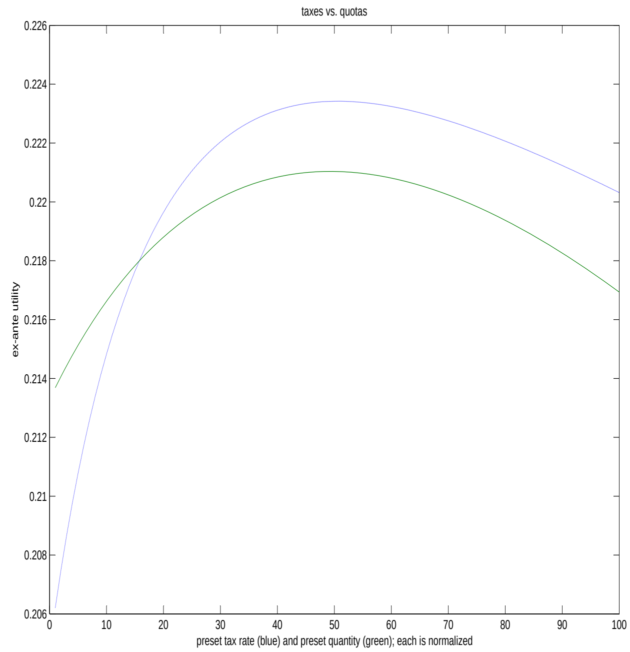

Figure 12: Utility from pre-committing to a unit tax (blue, with the tax on the x-axis) or a quantity cap (green, with the quantity cap on the x-axis)

price of the trading rights as well. The experience in Europe since the cap-and-trade system illustrates these points well: carbon prices have fluctuated between over 30 euro and virtually zero since the system started. Such fluctuations are observed also in other regions with cap-and-trade systems (e.g., New Zealand). Clearly, since optimal carbon pricing should reflect the social cost of carbon, such fluctuations are only efficient if the social cost of carbon experiences fluctuations. Damages from carbon emissions are likely not experiencing large fluctuations, but our assessments of how large they are of course change over time as scientific knowledge accumulates. The recent large drops in the price of emission rights can therefore be viewed as problematic from a policy perspective.

A cap-and-trade system could be augmented with a "central emission bank" that would have as its role to stabilize the price of emission rights by trading actively in this market, hence avoiding the large and inefficient swings observed in the EU system. Notice, however, that we would then be very close in spirit to a tax system: a tax system would just be a completely stable (provided the chosen tax is stable) way of implementing a stable price of emissions for firms.[68]

[68] This and other issues in this policy discussion are covered in Hassler et al. (2016).

4.11 Carbon taxation in the presence of other distortionary taxes

Suppose the government needs to raise revenue and needs to do this in a distortionary manner; the most common example would involve labor taxation and it is also a form of taxation that can be studied in the baseline model here by the addition of valued leisure. How, then, will the optimal carbon tax change? For example, suppose preferences are $\log c + \psi \log l$, where l is leisure, so that the labor input in the final-goods sector would be $1 - n_E - l$ (and, as before, n_E in the coal sector). Suppose also that the government has a distortionary tax on labor income, τ_l. Taxes are used to pay for an exogenous amount G of consumption good (that does not enter agents' utility). Lump-sum taxation is ruled out (but lump-sum transfers are not), and thus the setup mimics a typical second-best situation in public finance.[69]

Consider first a planning solution where the government is unrestricted and can just mandate quantities. Thus, it maximizes

$$\log\left(e^{-\gamma\phi\chi n_E}(1 - n_E - l)^{1-\alpha-\nu}(\chi n_E)^{\nu} - G\right) + \psi \log l$$

by choice of n_E and l. This delivers two first-order conditions. One is familiar from our baseline model:

$$-\gamma\phi\chi_E - \frac{1-\alpha-\nu}{1-n_E-l} + \frac{\nu}{n_E} = 0.$$

The other is the standard macro-labor condition

$$-\frac{1}{c} \cdot \frac{(1-\alpha-\nu)y}{1-n_e-l} + \frac{\psi}{l} = 0,$$

which says that the marginal utility of consumption times the marginal product of labor has to equal the marginal utility of leisure (in the expression, of course, y denotes $e^{-\gamma\phi\chi n_E}(1 - n_E - l)^{1-\alpha-\nu}(\chi n_E)^{\nu}$ and $c = y - G$). These two first-order conditions can be solved for first-best levels of n_E and l given any G.

Now consider in contrast a competitive equilibrium which is laissez-faire with regard to the taxation of carbon and which only uses labor taxes to raise revenue. Then, the two conditions above would be replaced, first, by the laissez-faire condition for coal

$$-\frac{1-\alpha-\nu}{1-n_E-l} + \frac{\nu}{n_E} = 0$$

and, second, a distorted macro-labor condition

$$-\frac{1}{c} \cdot \frac{(1-\alpha-\nu)y(1-\tau_l)}{1-n_E-l} + \frac{\psi}{l} = 0,$$

with the additional constraint that the government budget balances: $\tau_l(1-\alpha-\nu)y/(1 - n_E - l) = G$. These three conditions now determine n_E, l, and τ_l and do not deliver the first

[69] One can also consider an alternative assumption: there is no need to raise revenue ($G = 0$), there is an exogenous tax rate on labor income, $\tau > 0$, and any tax revenues are rebated back lump-sum.

best. In particular, one can think of two "wedges" defining different departures from the first best: the externality wedge due to climate damages and the tax wedge on labor supply (these are defined as the differences between the left-hand sides of the above equations with taxes and the corresponding ones from the first-best first-order conditions).

Now suppose we increase the carbon tax marginally from 0. Then (i) the climate wedge would become smaller and (ii) because τ_l falls—the government budget now reads $\tau_l(1-\alpha-\nu)y/(1-n_E-l)+\tau\chi n_E = G$ so that $\tau > 0$ allows a lower τ_l—the labor wedge would fall as well. Hence relative to a laissez-faire situation from the perspective of coal, introducing coal taxation involves a *double dividend*: it diminishes the climate externality and it reduces the labor distortion. This is an often-discussed point in the climate literature; for example Jorgenson et al. (2013a, 2013b) argue that the double dividends are quantitatively important for the U.S. and China, respectively.[70] Of course, the extent to which labor taxes can be reduced depends on the size of the coal tax base.

What, then, will the best level of carbon taxation be? Will carbon taxes be higher than in the absence of distortionary labor taxation? It would be straightforward to derive an answer in the present model by maximizing consumer welfare—with the same objective as that used by the planner—subject to the macro-labor first-order condition above, $\tau\chi/y-\frac{1-\alpha-\nu}{1-n_E-l}+\frac{\nu}{n_E}=0$ for the market's marginal condition for coal, and the government's budget constraint. One can derive a marginal condition for the planner's choice of τ which involves the setting of a weighted combination of wedges to zero; this condition can be solved numerically, together with the other equations, for the endogenous variables. The final level of taxes in this second-best solution is hard to characterize in terms of primitives but some intuition can perhaps be gleaned. If the use of coal is complementary with labor (which it is in the Cobb-Douglas formulation of production), on the margin the reduction of coal will hurt labor supply because it lowers the marginal product of labor. This speaks for a second best with a coal tax that is lower than in the absence of distortionary labor taxation. If coal were instead complementary with leisure (say because people burn coal to heat their homes when not working), this effect would go in the opposite direction on the margin. However, exactly how all these effects play out depends on the details of preferences and technology. For recent work on these issues that in addition also addresses distortions due to capital taxation, see Schmitt (2014), who pursues this approach in a dynamic model closely related to the setup here, and Barrage (2015), who looks at a closely related setting and uses a primal approach to taxation.[71]

4.12 A more detailed energy sector

We set out with a stylized description of energy production using either oil, coal, or some green alternative. In practice it is not either or; rather, these sources can all be used and are partially, but not fully, substitutable. Some integrated assessment models include very

[70] One can also identify a third dividend from introducing coal taxation: the reduce in local pollution from the burning of coal, a factor which appears of first-order relevance particularly in China.

[71] As is typically the case, in dynamic analyses it makes a difference whether the government has commitment or not; Schmitt considers cases without commitment.

complex energy systems (e.g., WITCH or MERGE; the latter is described in Manne et al., 1995). One way to incorporate multiple energy sources explicitly is to keep one kind of energy as an input into production but let this energy itself be produced from an array of sources, including fossil fuel. Thus, consider the CES technology

$$E = \left(\kappa_o E_o^\rho + \kappa_c E_c^\rho + (1 - \kappa_o - \kappa_c) E_g^\rho\right)^{\frac{1}{\rho}},$$

where E_i is the energy produced from source i, with $i = o$ representing oil (and natural gas), $i = c$ representing coal, and $i = g$ representing energy generated without fossil fuel.[72] This description is still stylized but it allows us to look into some interesting issues. The parameter $\rho \in (-\infty, 1]$ regulates the (constant) elasticity of substitution between the different energy sources.[73] The κ_is are share parameters regarded as exogenous in all of our analysis. We continue to think about the production of oil, coal, and green energy as in the previous discussion.

It is straightforward to check that the social cost of carbon is still γy with this formulation. Thus, this extension is not interesting from the perspective of optimal policy. Its value, instead, is to deliver a much richer view of what the cost is of remaining at laissez-faire, or in any case far from the optimum, because this cost turns out to crucially depend on the elasticity of substitution between the different kinds of energy.

First, and just for illustration, let us look at the case where there is just oil and coal, i.e., where there is no green energy. Clearly, then, if the degree of substitutability between oil and coal is very low, the difference between laissez-faire and the optimum is small. Consider the extreme case: a Leontief function, i.e., $\rho = -\infty$. Then if the total stock of oil is small enough that the optimum involves using it all, the laissez-faire and optimal allocations are identical. With some more substitutability, the laissez-faire allocation is not optimal, because coal use should be reduced given the externality and its unlimited supply (recall its constant marginal cost in terms of labor). However, the difference is still limited. In practice, however, oil and coal are rather good substitutes, so let us instead (again, for illustration only) consider the opposite extreme case: perfect substitutability ($\rho = 1$). Then the level of coal is determined very differently: laissez faire is far from the optimum (provided γ is large). Thus, in this case there will be significant total losses from government inaction.

According to available estimates, the remaining amount of (low-cost) oil left is quite limited, in particular in comparison with the amount of remaining coal, so oil is not of key importance for climate change.[74] What is of importance, however, is the substitutability with green energy. So, second, let us consider fossil fuel (interpreted as coal) versus green energy. In a meta-study, Stern (2012) reports a long-run elasticity of substitution of 0.95, as an average of oil-coal, oil-electricity, and coal-electricity elasticity measures. Thus, this unweighted average is close to a Cobb-Douglas specification. In this case, there can be a rather

[72] It would be natural to consider a slight extension of this formulation with a nested CES between a composite of oil and coal, on the one hand, and green energy on the other. Thus, oil and coal would form a separate CES aggregate and one could consider the quantitatively reasonable case with a high degree of substitutability between oil and coal and a lower one between the oil-coal composite and green energy.

[73] The elasticity is $1/(1 - \rho)$.

[74] See McGlade and Ekins (2015) for supply curves of different types of fossil fuel.

significant difference between the optimum and laissez-faire; relatedly, price incentives, or the effects of imposing a tax, are large if there is a non-taxed good that is a close substitute.[75] However, it is conceivable that green technology in the future will be a very good substitute with fossil fuel. Considering a higher elasticity than the unitary Cobb-Douglas elasticity is therefore a relevant robustness check. In this case, the difference between the optimum and laissez-faire is rather large. For example, Golosov et al. report, using a calibrated dynamic counterpart of the model here, that an elasticity of 2 leads laissez-faire coal use 100 years from now to rise to levels that imply exhaustion of all the coal deposits and would likely have catastrophic consequences for the climate. In contrast, in the optimum, coal use in 100 years is *lower* than it is today, and the climate as a result is rather manageable.

By definition, in the case of green energy vs. fossil fuel, the observation that a high elasticity of substitution leads to large welfare losses from not imposing a carbon tax (or a quota) at the same time means that there is a large potential social benefit from climate change action. A closely related implication is that there are, in such a case, strong incentives—high social payoffs—from doing research to come up with green alternatives. We turn to this issue in Section 4.14.

4.13 The substitutability between energy and other inputs

What aspects of the above analysis are influenced by the nature of the production function? We have assumed a Cobb-Douglas structure in part for simplicity and part because the energy share, though having gone through large swings over shorter periods of time, has remained fairly stable over the longer horizon (recall Figure 12 in Section 2). It is nevertheless necessary to also discuss departures from unitary elasticity. In this discussion, we will maintain the assumption of a unitary elasticity between the capital and labor inputs, thus confining attention to a different elasticity between the capital-labor composite, on the one hand, and energy on the other.

Consider the aggregate production function $e^{-\gamma S}F(Ak^{\alpha}n^{1-\alpha}, A_E E)$, where F is CES and A and A_E are technology parameters, thus maintaining the assumption that damages appear as decreases in TFP. The social cost of carbon with this formulation will then obey the same structure as before, i.e., the marginal externality damage of fossil fuel (through increased emissions E) is $\gamma\phi y$. What is different, however, is the difference between the laissez-faire allocation and the optimum or, expressed differently, the consumption equivalent cost of a suboptimal allocation. Consider oil, i.e., a fossil fuel with zero extraction costs in a finite supply $\bar{E}$. Assume that it is not optimal to use all of the oil, and let us simply examine the two extreme cases: Leontief and perfect substitutability.

We begin with the Leontief case. Here, output is given by $e^{-\gamma\phi E}\min\{Ak^{\alpha}n^{1-\alpha}, A_E E\}$. I.e., there is no substitutability between the capital-labor composite and oil. In laissez-faire, oil use is $\bar{E}$. It is easy to show from the planner's first-order condition that $E^* = 1/(\gamma\phi)$ in this case.[76] Recall from Section 4.1.3 that, under Cobb-Douglas, the optimal allocation is

[75]The Cobb-Douglas case is very similar to the case with only coal considered above.

[76]This holds so long as there is an interior solution, i.e., if $1/(\gamma\phi) < Ak^{\alpha}n^{1-\alpha}/A^{E}$. Note that there is

$E^* = \nu/(\gamma\phi)$ and that the ratio of optimal to laissez-faire output is $e^{\gamma\phi(\bar{E}-\nu/(\gamma\phi))}\left(\frac{\nu}{\gamma\phi\bar{E}}\right)^{\nu} > 1$. Now we obtain $e^{\gamma\phi(\bar{E}-1/(\gamma\phi))}\frac{1}{\gamma\phi\bar{E}}$. Because $-\nu + \nu\log\nu$ is decreasing we therefore conclude that in the Leontief case, the difference between the optimal and the laissez-faire allocation is smaller than under unitary elasticity. The fall in energy use is smaller, and this effect dominates the stronger impact on output of any given fall in energy.

Under perfect substitutability, we have output given by $e^{-\gamma S}\left(Ak^{\alpha}n^{1-\alpha} + A_E E\right)$ and we assume that capital and labor are in use. Now the planner's first-order condition leads to $E^* = 1/(\gamma\phi) - Ak^{\alpha}n^{1-\alpha}/A_E$, which (as for the unitary-elasticity case) is a smaller amount than in the Leontief case. It is also possible to show that the wedge between optimal and laissez-faire output in this case is smaller than in the Leontief case.

In sum, we see that the energy use can be different than in the case with unitary elasticity between energy and other inputs. With production functions with very low substitution elasticity between energy and other inputs, energy use will dictate that energy use in the optimum fall more, but there is also a corresponding gain in a higher TFP. There does not, perhaps surprisingly, therefore appear to be a very strong effect on the net gap between optimal output and laissez-faire output as the elasticity of substitution between inputs is varied. This is comforting given that the Cobb-Douglas formulation is much easier to handle analytically.

4.14 Green technology and directed technical change

The existence of the green technology was taken as given above; green technologies of various sorts—versions of water and wind power—have of course existed since before the industrial revolution. These technologies have been improved and there are also new sources of electricity production that do not involve fossil fuels, such as nuclear power and solar power.[77] A central issue of concern in the area of climate change is the further development of these technologies and research toward new ones. In the macroeconomically oriented literature on climate change, various models have been developed, with early papers by Bovenberg and Smulders and others (see, e.g., Bovenberg and Smulders, 1995). More recently, Acemoglu et al. (2012) provided a setting of directed technical change and made the point that there may be *path dependence* in R&D efforts toward the development of different energy technologies. We will now use the simple model to illustrate these facts and some other points that have been made in the literature.

A static model cannot fully do justice to the much more elaborate dynamic settings where many of the arguments in this part of the literature have been developed. It does, however, allow us to make a number of basic points. One simplification in our analysis here is that

abundance of capital and labor now: on the one hand, the market uses oil to the point where $E = Ak^{\alpha}n^{1-\alpha}$, so that there is excessive oil. On the other hand, the planner may want to decrease the oil use if the just stated inequality holds, so that from the planner's perspective, there is an abundance of capital and labor instead.

[77]Nuclear power is problematic from an environmental perspective too but we do not discuss this issue here.

we will not explicitly describe a decentralized R&D sector.[78] We will distinguish between two different kinds of technological developments: new techniques for the efficient use of energy ("energy saving") and new techniques for the production of energy. We begin with the latter.

4.14.1 Energy production

We will mostly abstract from the determination of the overall efforts toward technological developments, which one could model as well (say as a tradeoff between these activities and using labor directly in production), and simply assume that there is an R&D input available in fixed supply; we set the total amount to 1 without loss of generality. The use of this input can be *directed* toward either improving the productivity in producing energy from fossil sources, m_c, or from green sources, m_g, with the constraint that $m_c + m_g = 1$. E.g, we can think of this choice as one between improving the drilling/extraction technologies for North Sea oil and technological improvements in the cost-efficiency of solar-based units. The most straightforward setting would maintain the production function in a two-energy-input form:

$$e^{-\gamma E_c} k^{\alpha} n^{1-\alpha-\nu} \left(\lambda_c E_c^{\rho} + (1-\lambda_c) E_g^{\rho}\right)^{\frac{\nu}{\rho}},$$

with the production of energy given by

$$E_c = \chi_c n_c \quad \text{and} \quad E_g = \chi_g n_g$$

with $n + n_c + n_g = 1$. Along the lines indicated above, for given values of χ_c and χ_g, this model is straightforwardly solved either for the optimum or for a laissez-faire allocation.

A very simple way of modeling research into making energy production more efficient can now be expressed as follows:

$$\chi_c = \bar{\chi} m_c \quad \text{and} \quad \chi_g = \bar{\chi} m_g,$$

with $m_c + m_g = 1$. (If $\lambda_c = 1/2$, this setting is now entirely symmetric.)

A decentralized version of this model would have no agent—either the producer or the user of fossil fuel—take into account the negative externality. However, notice that there are increasing returns to scale in producing energy: double n_c, n_g, m_c, and m_g, and E_c and E_g more than double. A decentralized equilibrium here would then have a much more elaborate structure of varieties within each energy type, either with variety expansion à la Romer or fixed variety but creative destruction à la Aghion-Howitt (1992), monopolistic competition with profits, and then perfectly competitive research firms producing new varieties (in the Romer case) or product improvements (in the Aghion-Howitt case). We will not spell the variety structure out, but we will make the assumption that the aggregation across varieties is identical for fossil fuel and green energy, e.g., implying identical markups across these two energy sectors. Finally, there would normally (in dynamic models) also be spillovers, mostly

[78] We could have developed such a version even in our static model but it would have complicated notation without adding much of significance.

for tractability, but they are not needed here.[79] We will, however, discuss spillovers below because there are substantive issues surrounding them.

A decentralized model such as that just described delivers equilibrium existence despite the technological nonconvexity but we omit the description of it for brevity; see Romer (1990) for the basic variety-expansion structure and Acemoglu (2009) for a more recent description of a range of endogenous-growth models and many of their uses. Monopolistic competition would distort the allocation, in the direction of under-provision of energy, which itself would be beneficial for counterbalancing the climate externality and thus to some extent relieve the government of the pressure to tax fossil fuel. In the laissez-faire equilibrium, in the case of symmetry between fossil fuel and energy, the markets will produce whatever the total energy composite is in an efficient manner.[80] Denoting this level E, the laissez-faire allocation will minimize $n_c + n_g$ subject to

$$E_c^\rho + E_g^\rho \geq E^\rho, \quad E_c = n_c \bar{\chi} m_c, \quad E_g = n_g \bar{\chi} m_g, \quad \text{and} \quad m_c + m_g = 1.$$

The solution to this problem depends critically on ρ. So long as $\rho < 1/2$, i.e., so long as the two sources of energy are poor enough substitutes, the solution is to set $n_g = n_e$ and $m_c = m_g = 1/2$; it is straightforward to compute the implied total labor use. If, on the other hand, $\rho > 1/2$, then the outcome is to set either $n_c = m_c = 0$ or $n_g = m_g = 0$, i.e., a corner solution obtains, with another easily computed labor use. So if the energy inputs are substitutable enough, there are multiple equilibria. The multiplicity is knife-edge in this case since we assumed full symmetry. However, the essential insight here is not multiplicity but rather sensitivity to parameters, as we will now elaborate on.

Suppose now, instead, that we change the setting slightly and assume

$$\chi_c = \bar{\chi}_c m_c \quad \text{and} \quad \chi_g = \bar{\chi}_g m_g,$$

i.e., we assume that there are two separate constants in the two research production functions. Then, in the case where ρ is high enough, there will be full specialization but the direction of the specialization will be given by the relative sizes of $\bar{\chi}_c$ and $\bar{\chi}_g$. If the former is higher, the energy will be produced by fossil fuel only; if the latter is higher, the energy will be produced by green energy only. If the economy experienced a small change in these parameters switching their order, we would have a complete switch in the nature of the energy supplies. Crucially, now, note that we can think of $\bar{\chi}_c$ and $\bar{\chi}_g$ as given by historical R&D activities. Then we can identify the kind of *path dependence* emphasized in Acemoglu et al. (2012). These authors argued that temporary efforts, via subsidies/taxes, to promote the research on "clean goods"—those produced using green energy—would have permanent effects on our energy supplies by managing to shift our dependence on fossil fuel over to a

[79] The reason they improve tractability is that if the researchers' output does not give the researcher herself dynamic gains, the R&D decision becomes static.

[80] The assumption of symmetry across the two energy sectors, and hence identical markups, is an important assumption behind this result.

dependence on green energy.[81] This can be thought of, in terms of this model, as having managed to make $\bar{\chi}_g > \bar{\chi}_c$ by past subsidies to green R&D. Acemoglu et al. used a dynamic model with details that differ from those here—among other things, they assumed much stronger convexities in damages so that a switch to green energy was necessary or else utility would be minus infinity—but this is the gist of their argument.

One can question whether the substitutability is strong enough for the path-dependence argument to apply. For example, Hart (2013) argues that there are strong complementarities in research across dirty and clean technologies. These complementarities could, in practice, take the form of external effects/spilliovers. For example, research into improving electric cars can be helpful for improving the efficiency of cars running on gasoline or diesel, and whether these complementarities are fully paid for or not in the marketplace is not obvious. A way of expressing this formally within our simple framework is a further generalization of our framework as follows:

$$\chi_c = \bar{\chi}_c m_c^{\zeta} m_g^{1-\zeta} \quad \text{and} \quad \chi_g = \bar{\chi}_g m_g^{\zeta} m_c^{1-\zeta}.$$

To the extent ζ is not too much higher than 1/2 here, there are strong complementarities in technology development and path dependence would not apply. Hart (2013) argues this is the relevant case, but it would be hard to argue that the case is settled. Aghion et al. (2014), furthermore, show that there is empirical support for persistence, though whether these effects are strong enough to generate the kind of path dependence emphasized in Acemoglu et al. (2012) is still not clear.

Turning, finally, to the planning problem in these economies, it is clear that the planner faces a tradeoff between the forces discussed here and the climate externality generated by fossil fuel. The setting is rather tractable and it is straightforward to determine the optimal mix of energy supplies. We leave out the detailed analysis for brevity.

4.14.2 Energy-saving

Research into alternative (green) energy supplies is definitely one way of decreasing our fossil-fuel use. Another is energy-saving. To formalize this idea, let the energy composite be written in a somewhat more general way, again emphasizing two energy sources (c and g) only:

$$E = \left(\lambda_c \left(A_c E_c\right)^{\rho} + \left(1-\lambda_c\right)\left(A_g E_g\right)^{\rho}\right)^{\frac{1}{\rho}}.$$

The technology factors A_i here indicate the "efficiency" with which different energy sources are used. Note, parenthetically, that there is a direct parallel with how we treated energy vs. a capital-labor composite in Section 2 above. Now the A_is introduce asymmetry between the different energy sources through another channel, and moreover we can think of them as being chosen deliberately. One interpretation of these choices is temporary decisions to save

[81] In their analysis, the authors use a notion of two kinds of goods, one clean and one dirty, with labels deriving from the energy source used to produce them. The setting we use here, with an energy composite relevant for the whole economy, is of course also an abstraction but we prefer it because it lends itself more easily to calibration and comparison with data.

on energy, e.g., by directing effort toward closing windows or making sure machines don't run unnecessarily. Another interpretation emphasizes research toward energy efficiency that are of a permanent nature. One example is the development of more fuel-efficient cars; another is to develop methods for using less jet fuels when airplanes land. In parallel with our treatment of energy production, we then add the equations

$$A_c = \bar{A}_c m_c^{\zeta} m_g^{1-\zeta} \quad \text{and} \quad A_g = \bar{A}_g m_g^{\zeta} m_c^{1-\zeta},$$

again with the constraint $m_c + m_g = 1$.[82] With this structure as well, market allocations may end up with specialization for a range of parameter configurations, as will the solution to the planning problem, and path dependence is again possible.

An important concern in the modeling of energy saving or the efficiency of producing energy is that there is a natural upper limit to efficiency. For example, light produced with LED has almost reached the efficiency limit and the same is true for electrical engines. However, this does not mean that we are close to maximal energy efficiency in the production of transportation services. For the transportation example it is less appropriate to capture efficiency through A_g; rather, improvements come about through increasing general energy efficiency (say, a coefficient in front of E in the overall production function). The limits to efficiency are normally not made explicit in economic models but arguably should be in quantitative applications.

4.14.3 Are subsidies for green technology needed?

To attain the optimal allocation, the planner will of course need to tax the use of fossil fuel. What other taxes and subsidies might be necessary? To the extent there is monopoly power, and the energy sources undersupplied, subsidies are needed. Should the green R&D sector be subsidized? Following Pigou's principle, it should be to the extent there are positive spillovers. So in the absence of technology spillovers in the green R&D sector, there would actually be no reason to subsidize. Moreover, if there are spillovers but they are identical for the two sorts of energy, it is not clear that green technology should receive stronger subsidies than should fossil-fuel technology, so long as fossil fuel is taxed at the optimal rate.

In a second-best allocation, of course, matters are quite different. Suppose no coal tax is used. Then subsidies to the production of green energy, or to the development of new green technologies, would be called for. In political debates, subsidies to the development of green technology appear to be quite popular, and our analysis is in agreement with this view insofar as an optimal (global) carbon tax is not feasible. In practical policy implementation, though less so in debates, it also appears that coal subsidies are popular, perhaps not as per-unit instruments but as support in the construction of plants. A study (Hassler and Krusell, 2014) in fact claims that the average global tax on carbon is set at about the right magnitude but with the wrong sign—owing to large subsidies for coal production across the world.

[82]One can also state these constraints using other functional forms, such as $(\bar{A}_c A_c)^{\zeta} + (\bar{A}_g A_g)^{\zeta} \leq A^{\zeta}$. It is an empirical matter what formulation works best, and it is probably fair to say that the literature is so far silent on this issue.

The view expressed in Acemoglu et al. (2012) appears to contrast somewhat with ours. They argue, based on their model of path dependence, that subsidies to green technology are necessary for attaining an optimum and that carbon taxes would not suffice. They obtain this result not only because their model features strong intertemporal spillovers to R&D but also because they make assumptions such that if the "clean good" does not take over from the "dirty good", the climate damages will be infinitely costly (thus, they have strong non-convexities in their damage function, a tipping point of sorts). Moreover, their model has a second-best structure with spillovers and very limited patent lives. How can we understand this result from the perspective of Pigou taxation? Recall that we pointed out that Pigou taxation may not work if there are multiple market equilibria, and the kind of setting Acemoglu et al. describe has a feature of this kind. The simplest parallel in our static model is the coal-green setup we described in Section 4.14.1. There, we looked at a planning problem with a choice between two energy sources. So suppose that $\bar{\chi}_c = \bar{\chi}_g = \chi$ there, and let us imagine a market allocation where the labor productivity of coal and green energy production, χm_c and χm_g, respectively, derive from variety expansion in patent efforts (m_c and m_g) driven by monopoly profits for intermediate, specialized goods. Suppose, moreover, that there are no research spillovers in this setting: this assumption is perhaps natural in a static model (but less so in a dynamic one). In this framework, then, there would be two equilibria if ρ, the parameter guiding the key energy elasticity, is high enough. Suppose, moreover, that damages are to preferences, as in Section 4.6, and with highly nonlinear features, as discussed in Section 4.7: the marginal damages are first zero for a range of low emission levels, then high and positive, and then again zero in a "disaster zone". Suppose, moreover, that if the economy ends up using coal, emissions will end up in the disaster zone. Then the Pigou procedure would amount to finding the optimal solution—that with green technology only—and an associated tax on carbon that is zero, since the marginal damage at zero emissions is zero. So here Pigou's procedure is highly problematic, since there are now two market outcomes given a zero tax on carbon, and one of them is a disaster outcome! Thus another instrument would be needed to select among the two market outcomes, and one option would be a large enough subsidy to green technology creation to rule out an equilibrium where markets engage in the research on coal technologies.[83]

4.14.4 Green technology as a commitment mechanism

Some argue that future decision makers cannot be trusted to make good decisions and that, therefore, to the extent we can affect their decisions with irreversible decisions made today, we should. Why would future decision makers not make good decisions? One reason is based on time-inconsistent discounting, as discussed above: the current decision maker may have lower discount rates between any two future cohorts than that between the current and next cohort, and if this profile of decreasing discount rates is shared by future cohorts—updated by the appropriate number of cohorts—then profiles are time-inconsistent. In particular, from the perspective of the current cohort, future cohorts look too impatient. Since future

[83]With monopolistic competition, one would in general also need to encourage production to prevent under-supply for those technologies that end up being patented.

carbon taxes cannot literally be committed to today, then, the current cohort is restricted and appears to not be able to attain its preferred outcome.[84] Another conceptually distinct reason for disagreements is that politicians (and possibly the voters who support them) may be "myopic"; Amador (2003) shows that rationality-based dynamic voting games in fact can lead to reduced forms characterized by time-inconsistent preferences of politicians.[85] Finally, Weitzman (1998) provides further arguments for falling discount rates based on the idea that the true future discount rate may be uncertain.

If current decision makers cannot decide directly on the future use of fossil fuels, they may be able to at least influence outcomes, for example by investing in green technology that, ex post, will tilt the decision makers in the future in the right direction. To illustrate, consider a model where production is given by

$$e^{-\gamma\phi\chi_E n_E}(1-n_E-n_g)^{1-\alpha-\nu}(\chi_E n_E+\chi_g n_g)^{\nu}.$$

$E=\chi_E n_E$ is coal-produced energy and $E_g=\chi_g n_g$ is green energy; we make the assumption, only for obtaining simpler expressions, that these two energy sources are perfect substitutes. Now assume that there is an ex-ante period where an irreversible decision can be made: that on n_g. The cost is incurred ex post, so only the decision is made ex ante. Moreover, it is possible to increase n_g ex post but not decrease it: it is not possible to literally reverse the first decision.[86] Finally, assume that the ex-ante decision maker perceives a different damage elasticity than the ex-post decision maker (they have different γs, with the ex-ante value higher than the ex-post value): this captures, in a simple way, the intertemporal disagreement.

We make two further simplifying assumptions, for tractability. First, we take the ex-post decision maker to perceive a damage elasticity of exactly 0 and the ex-ante decision maker to use the value $\gamma>0$. Second, we assume that $\chi_E>\chi_g$, i.e., that—climate effects aside—the coal technology is a more efficient one for producing energy, regardless of the level at which the two technologies are used (due to the assumption of perfect substitutability). How can we now think about outcomes without commitment?

It is clear that the ex-post decision maker sees no reason to use the green technology at all. Facing a given amount of n_g that he cannot decrease (and will not want to increase), the level of n_E will be determined by the first-order condition

$$\frac{1-\alpha-\nu}{1-n_E-n_g}=\frac{\nu\chi_E}{\chi_E n_E+\chi_g n_g}. \tag{17}$$

This expression delivers a linear (affine) and decreasing expression for n_E as a function of n_g: $n_E=h(n_g)$, with $h'<0$ and independent of n_g.

[84] Karp (2005), Gerlagh and Liski (2012), and Iverson (2014) analyze optimal taxes in the presence of time-inconsistent preferences.

[85] See also Azzimonti (2011) for a similar derivation.

[86] We may think of this setup as a reduced-form representation for a case when an ex-ante investment in capital or a new technology makes it profitable to use at least n_g units of labor in green energy production, even if it the emission reduction is not valued per se. In a dynamic model, the cost of this investment would at least partly arise ex ante, but this is not of qualitative importance for the argument.

What is the implied behavior of the ex-ante decision maker without commitment? She will want to maximize

$$e^{-\gamma\phi\chi_E h(n_g)}(1-h(n_g)-n_g)^{1-\alpha-\nu}(\chi_E h(n_g)+\chi_g n_g)^{\nu}$$

by choice of n_g, a decision that delivers a second-order polynomial equation as first-order condition, just like in the baseline case (though now with somewhat more involved coefficients in the polynomial). Does this first-order condition admit the first best outcome of the ex-ante decision maker? Such a first best would amount to the solution of the two first-order conditions

$$\gamma\phi\chi_E + \frac{1-\alpha-\nu}{1-n_E-n_g} = \frac{\nu\chi_E}{\chi_E n_E + \chi_g n_g} \tag{18}$$

and

$$\frac{1-\alpha-\nu}{1-n_E-n_g} = \frac{\nu\chi_g}{\chi_E n_E + \chi_g n_g} \tag{19}$$

which result from taking derivatives with respect to n_E and n_g, respectively. It is easy to see that these cannot deliver the same solution as the problem without commitment. For one, equation (19) and equation (17) cannot deliver the same values for both n_E and n_g, since they differ in one place only and $\chi_E > \chi_g$. Thus, we are in a second-best world where the ex-ante decision maker uses her instrument but cannot, without an additional instrument, obtain her first-best outcome. Moreover, total energy use and/or total labor used to produce energy will be lower with the ex-ante decision on green energy than in the absence of it, comparing equations (17) and (18). This model is stylized and it would appear that the specific predictions could change when moving to a more general setting. However, the second-best nature of the setting would remain.

4.14.5 The Green Paradox

The Green Paradox, a term coined by Sinn (2008), refers to the following logical chain. Decisions to subsidize green technology so as to speed up the research efforts in this direction will, if these efforts are successful, lead to better and better alternatives to fossil fuel over time. This, in turn, implies that fossil-fuel producers have an incentive to produce more in advance of these developments, given that their product is more competitive now than it will be in the future. As an extreme example, imagine that cold fusion is invented but takes one year to implement, so that one year from now we have essentially free, green energy in the entire economy. Then owners of oil wells will produce at maximum capacity today and, hence, there will be much higher carbon dioxide emissions than if cold fusion had not been invented. Hence the "paradox": green technology (appearing in the future) is good but therefore bad (in the short run).

Our static model fully cannot express the Green Paradox, of course, since the essence of the paradox has to do with how events play out over time. Consider therefore a very simple two-period version of the model that allows us to think about how the intertemporal decision for oil producers depends on the availability of green technology. We assume that consumers' preferences are linear so that the gross interest rate is given by $1/\beta$. We assume that fossil fuel is (free-to-produce) oil and that $\rho = 1$, so that oil and green energy are perfect substitutes. We also assume that there is no green technology in the first period. A simplified production function thus reads $e^{-\gamma\phi_1 E_1} k^\alpha E_1^\nu$ for period 1 and $e^{-\gamma\phi_1(E_2+\phi_2 E_1)} k^\alpha (E_2 + E_g)^\nu$ for period 2; for simplicity, we also abstract from the costs for producing green energy and set E_g to be exogenous, with $n = 1$ in both periods). Here, ϕ_1 and ϕ_2 allow us to capture a carbon depreciation process that does not occur at a geometric rate, a feature we argued is realistic. Our notation reveals that capital cannot be accumulated in this example, but we will comment on accumulable capital below.

Given this setting, the price of oil in period 1 is given by $p_1 = \nu e^{-\gamma\phi_1 E_1} k^\alpha E_1^{\nu-1}$ and in period 2 it is given by $p_2 = \nu e^{-\gamma\phi_1(E_2+\phi_2 E_1)} k^\alpha (E_2 + E_g)^{\nu-1}$. All of the available oil, $\bar{E}$, will be used up in the laissez-faire allocation and so oil use in the two periods will be given by the Hotelling condition, a condition we derived and analyzed in Section 2 above: $p_1 = \beta p_2$. Recall that this equation expresses the indifference between producing a marginal unit of oil in period 1 and in period 2. This condition implies that E_1 can be solved for from $e^{-\gamma\phi_1 E_1} E_1^{\nu-1} = \beta e^{-\gamma\phi_1(\bar{E}-E_1(1-\phi_2))} (\bar{E} - E_1 + E_g)^{\nu-1}$. Clearly, this equation has a unique solution and comparative statics with respect to E_g shows that more green energy in period 2 makes E_1 rise and E_2 fall. Hence the Green Paradox.

Is the move of emissions from period 2 to period 1 bad for welfare? The negative externality (SCC) of emissions in period 1 is $\gamma\phi_1(y_1 + \beta\phi_2 y_2)$ and the present value of the corresponding externality in period 2 is $\gamma\phi_1\beta y_2$. In the absence of a green technology in period 2 ($E_g = 0$) it is easy to show that $y_2 < y_1$ in the laissez-faire allocation and, hence, at least for a range of positive values of E_g, the externality damage is higher for early emissions. Intuitively, emissions in period 2 have two advantages. One is that they hurt the economy only once: emissions in period 1 will, except for the depreciated fraction $1 - \phi_2$, remain in the atmosphere—a significant factor given calibrated carbon-cycle dynamics—and hence also lower second-period TFP. The second advantage of emissions in the future is that their negative effect is discounted (to the extent we assume $\beta < 1$). Note, finally, that the possibility of accumulating physical capital would not change any of these conclusions: with more green energy in the second period, capital accumulation with rise somewhat to counteract the initial effect, and it would work toward an increase in p_2, but this mechanism would not overturn our main observation.

Can the future appearance of green technology also make overall welfare go down in the laissez-faire allocation? This is much less clear, as an additional unit of E_g (for free) has a direct positive welfare effect.[87] However, now consider competitive production of green energy under laissez-faire, at a unit labor cost χ_g. Here, a second-best argument would suggest that there is a negative "induced externality" of green energy production: since the economy is far from the optimum, and emissions in period 1 would be detrimental, any additional unit of E_g would have a negative side-effect on welfare. Hence, a(t least

[87] If there are strong non-linearities, like a threshold CO_2 concentration level above which climate damages are catastrophic, then the introduction of a green technology in the second period could make laissez-faire welfare fall.

a small) *tax* on green energy production would be desirable! The reason for this perhaps counterintuitive effect—aside from the Green-Paradox logic—is that the total amount of fossil fuel used will still be $\bar{E}$: green technology, in this setting, will not curb the use of fossil fuel, only change the timing of emissions (in the wrong direction).

The previous example points to counterintuitive policy implications: green technology should be discouraged. However, aside from the assumptions that make the Green Paradox relevant, this result also relies on second-best analysis. In the social optimum, green technology should not be taxed (nor subsidized): there is, simply, no externality from producing green technology in this model. If green technology is developed in an R&D activity, then support of this activity (relative to other activities) may be called for, but only if there is an R&D externality to green technology development that is, in the appropriate sense, larger than the corresponding one for fossil-fuel technology developments. Hence, the optimum (in this economy, where oil is free to produce) involves fossil-fuel taxes but no net support to green technology.

Is the Green Paradox empirically relevant? The key assumption that leads to the paradox is that the accumulated use of fossil fuel is the same under laissez-faire as in the optimal allocation. In this case, sub-optimality only comes from the speed at which the fossil reserves are used. That all reserves are used also in the optimal allocation is arguably reasonable when it comes to conventional oil with low extraction costs (e.g., Saudi oil). However, it is not reasonable for non-conventional reserves and coal. Here policy, including subsidies to the development of future green energy production, can and should affect how much fossil resources are left in ground. So suppose, instead, that we focus on fossil fuel in the form of coal and that we maintain our assumption that the marginal cost of coal is constant (in terms of labor or some other unit). Then an increase in E_g would lead to a lower demand for coal and hence have an impact on coal use: it would clearly induce lower coal production in the second period. Lower coal use, in turn, has a positive externality on the economy. Moreover, coal use in period 1 is not affected. Hence, the conclusion here is the opposite one: green energy has a positive effect on the economy (beyond its direct positive effect, to the extent it comes for free). In addition, relative to a laissez-faire allocation it would be beneficial to subsidize, not tax, green energy production. Which case appears most relevant? We take the view that the latter is more relevant. The argument has two parts. First, the intertemporal reallocation of emissions emphasized in the Green-Paradox argument, though logically coherent, is not, by our measure, quantitatively important. The main reason is that the total amount of oil is rather small and its effect on climate is limited, and a reallocation of emissions due to oil over time is of second-order importance compared to being able to control the cumulated (over time) emissions. Second, if the fossil fuel is costly to extract then there would be lower emissions, as argued above, and in terms of the total amount of fossil fuel available, most of it is costly to extract (most of it is coal). Coal is produced at a price much closer to marginal cost and the Hotelling part of the coal price appears small. This argument, moreover, is quantitatively important given the large amounts of coal available.

4.15 Regional heterogeneity

Nordhaus's basic DICE model is a one-region integrated assessment model, but there are by now several calibrated models in the literature with more than one region. His own RICE (R for Regional) model was perhaps the first multi-region model and it had 7 regions, defined by geographic and economic indicators; Krusell and Smith (2015) have developed a model at the extreme end of heterogeneity, treating one region as a 1-by-1-degree square with land mass on the global map. Regional models can serve a variety of purposes and we first briefly discuss the chief purposes. We then use a multi-region version of our basic model as an illustration; in particular, we use a simple version of Hassler and Krusell (2012) and look at some extensions.

A major purpose for looking at regional heterogeneity comes from recognizing that damages are very different in different parts of the world; some regions, such as Canada and most of Russia, are even expected to gain from a warmer climate. Thus, using a multi-region IAM as a simulation device, one can trace out the heterogeneous effects of climate change under different policy scenarios. Even if there is no agreement on a social welfare function for the world, surely policymakers are very interested in this heterogeneity.

Another purpose of a multi-region IAM is to look at the effects of regionally heterogeneous policies. Suppose the Western world adopts a strict carbon tax and the rest of the world does not. How effective will then the western policies be in combatting climate change, and what will its distributional consequences be?

Relatedly, one of the key concepts in policymakers' studies of climate change is *carbon leakage*. The idea here is simply that when carbon is taxed at higher rates in some regions than in others, the decreases in carbon use in the high-tax regions will presumably be (partially, or fully) offset by increases in carbon use in other regions. *Direct carbon leakage* would for example occur if the oil shipments are simply redirected away from low-tax to high-tax regions. But there can also be *indirect carbon leakage* in that the other factors of production (capital and/or labor) can move to where carbon taxes are lower—and hence carbon will be used more there as a result. Differential policies can also affect outcomes through trade (see, e.g., Gars, 2012, and Hémous, 2013). Finally, when there is R&D in the development of fossil-fuel and green technologies, differential policies in this regard come into play as well (Hémous, 2013, looks at this case as well).

Still another important aspect of a multi-region IAM is its potential for discussing *adaptation* to climate change through the migration of people (along with other production factors).[88] Adaptation is not just important in practice but it is important to think about from a theoretical and quantitative perspective since the damages from climate change really are endogenous and depend on how costly it is to migrate. If migration were costless, significant warming would potentially be less detrimental to human welfare since there are vast areas on our continents that are too cold today but, with significant warming, inhabitable. There is very little research on this issue so far (Brock et al., 2014, and Desmet and Rossi-Hansberg, 2015, are promising exceptions) but we believe it is an important area for future research and

[88] For a recent example, see Krusell and Smith (2015), who allow for the migration of capital.

one with much potential. Empirical research on the costs of migration is also scant, but some work does exist (Feng, Shuaizhang, Alan B. Krueger, and Michael Oppenheimer, "Linkages among climate change, crop yields and Mexico-US cross-border migration," PNAS, 2010, 107 (32), 1425714262, and, for a study of conflict in this context, see Harari and La Ferrara, 2014; see also the review Burke et al., 2015).

4.15.1 A two-region IAM with homogeneous policy: oil

Our simple model is easily extendable to include another region (or more). Let us look at a series of simple cases in order to illustrate some of the main points made in the literature.[89] Let us first look at heterogeneous damages, so assume that production in region 1 is $e^{-\gamma_1 E} k_1^\alpha n_1^{1-\alpha-\nu} E_1^\nu$ whereas production in region 2 is $e^{-\gamma_2 E} k_2^\alpha n_2^{1-\alpha-\nu} E_2^\nu$. Energy is coming from fossil fuel only, and let us first assume that it is (costless-to-produce) oil available at a total amount $\bar{E}$ in a third region of the world, which supplies the oil under perfect competition (the third region thus plays no role here other than as a mechanical supplier of oil). Let us also for simplicity start out by assuming that the two regions are homogeneous in the absence of climate damages, so that $k_1 = k_2 = k$ and $n_1 = n_2 = n$. It is easy to work out a laissez-faire equilibrium for this world and we can look at different cases, the first of which is that when neither capital nor labor can move. Thus, the only trade that occurs takes the form that the oil-producing region sends oil to the two other regions and is paid in consumption goods; regions 1 and 2 do not interact, other than by trading in the competitive world oil market. All of the oil will be used and the equilibrium oil distribution will now be determined by the following condition:

$$e^{-\gamma_1 \bar{E}} E_1^{\nu-1} = e^{-\gamma_2 E} E_2^{\nu-1},$$

i.e., by ($E_1 + E_2 = \bar{E}$ and)

$$\frac{E_1}{E_2} = e^{\frac{\gamma_2 - \gamma_1}{1-\nu} \bar{E}}.$$

Thus, the relative use of oil is higher in the country with lower climate damages.[90] Suppose that region 1 experiences stronger damages. Clearly, then, region 1 is worse off and the damage has a small "multiplier effect" to the extent that its energy used is curbed: more energy is used in region 2. In other words, we would see lower TFP in region 1 but lower activity there also because of reduced energy use. Consumption is a fraction $1-\nu$ of output, with the remainder sent to the third, oil-producing region.

If we also allow capital to move—but maintain that the populations cannot move—the output effect will be somewhat strengthened as capital will also move to region 2 to some extent. If half of capital is owned by each region, this makes region 1 gain, however, because

[89]It should be noted, however, that there are very few examples of multi-region IAM that are studied in full general equilibrium. Thus the number of formal results from the literature is therefore very limited relative to the number of informal conjectures.

[90]Of course this result depends on damages occurring to TFP; if they affect utility, oil use is identical in the two regions.

its GNP will rise even though its GDP will fall. In the real world, there are moving costs and cultural and other attachments to regions, so full and costless migration is probably not an appropriate assumption even in the long run (as the static model is supposed to capture a longer-run perspective).

Suppose now that regions 1 and 2 consider a common tax τ on carbon and suppose that this tax is collected in each country and redistributed back lump-sum to the local citizens. Would such a tax be beneficial? To regions 1 and 2, yes. The analysis depends on the size of the tax but suppose the tax is low enough that firms are not sufficiently discouraged from using oil that the total amount of oil use is lowered. Then the relative energy uses in the two regions will still satisfy the equations above and the levels will not change either. The price of oil, p, will satisfy

$$p = \nu y_1/E_1 - \tau,$$

the first term of which is independent of the tax size (for a small enough tax). Hence, country i's consumption will now be $y_i - (p+\tau)E_i + \tau E_i = (1-\nu)y_i + \tau E_i$ so that consumption is strictly increasing in τ for both regions. Thus, the two regions can use the tax to shift oil revenues from the oil-producing region to its own citizens, without affecting output at all.[91] When the tax is high enough that p reaches zero, the level of production responds to taxation: as producers now receive nothing for their oil, they are indifferent as to how much to supply. At that tax level, the total energy supply will still be given by $\bar{E}$ and the equations above, but now consider a slightly higher tax, still with a zero price of oil. Then the total amount of energy E is then lower and is determined from

$$\tau = \nu e^{-\gamma_1 E} k^{\alpha} n^{1-\alpha-\nu} E_1^{\nu-1} \quad \text{and} \quad \frac{E_1}{E - E_1} = e^{\frac{\gamma_2-\gamma_1}{1-\nu}E}.$$

It is straightforward to show, if the γs are not too far apart, that these two equations imply a lower E and E_1 as τ is raised and that E_1/E_2 will rise. Now, for each region there would be an optimal τ and there would be a conflict between these two values. Generally, the region with a higher climate externality would favor a higher tax.

4.15.2 A two-region IAM with homogeneous policy: coal

These discussions all refer to the case of oil, i.e., a free-to-extract fossil fuel. Suppose we instead look at coal, and assume that coal is domestically produced: it costs $1/\chi_i$ units of labor per unit, as in most of our analysis above. We also assume that the transport costs for coal are inhibitive so that there is no trade at all. The only connection between the regions is thus the climate externality. In the absence of taxes the world equilibrium is then determined independently of the externality and according to

$$\frac{1-\alpha-\nu}{\chi_i - E_i} = \frac{\nu}{E_i}$$

for $i = 1, 2$.

[91] This argument is of course unrelated to any climate externality; the climate is unaffected by the taxation.

Now the reason to tax in order to transfer resources away from a third region and to the home country is no longer applicable; the only reason to tax is the climate externality. As in the oil case, let us assume that any tax on coal is lump-sum transferred back to domestic consumers. What is then the best outcome for each of the two regions? The two countries can, in principle, act in a coordinated fashion so as to maximize overall welfare—by maximizing world output—and then choose a point on the Pareto frontier by the use of transfers. World output is maximized by setting the tax equal to the marginal damage externality in the world, i.e., $\gamma_1 y_1 + \gamma_2 y_2$. Thus, the social planner chooses E_1 and E_2 to solve

$$\gamma_1 e^{-\gamma_1(E_1+E_2)} k_1^{\alpha} \left(1 - \frac{E_1}{\chi_1}\right)^{1-\alpha-\nu} E_1^{\nu} + \gamma_2 e^{-\gamma_2(E_1+E_2)} k_2^{\alpha} \left(1 - \frac{E_2}{\chi_2}\right)^{1-\alpha-\nu} E_2^{\nu} =$$

$$e^{-\gamma_1(E_1+E_2)} k_1^{\alpha} \left(1 - \frac{E_1}{\chi_1}\right)^{1-\alpha-\nu} E_1^{\nu} \left(\frac{1-\nu-\alpha}{\chi_1 - E_1} - \frac{\nu}{E_1}\right) =$$

$$e^{-\gamma_2(E_1+E_2)} k_2^{\alpha} \left(1 - \frac{E_2}{\chi_2}\right)^{1-\alpha-\nu} E_2^{\nu} \left(\frac{1-\nu-\alpha}{\chi_2 - E_2} - \frac{\nu}{E_2}\right).$$

The first line represents the global damage externality (which is also the optimal tax on coal); it has to be set equal to the net benefit of emissions in each of the two regions (the following two lines). The allocation will have lower E_1 and E_2 amounts (provided, at least, both γs are positive) than in the laissez-faire allocation.

Suppose, however, that the regions cannot use transfers to arrive at a Pareto-optimal allocation. Then an optimal allocation would be obtained by maximizing a weighted value of the utilities of consumers in the two regions. Often, macroeconomic models adopt the utilitarian approach. Assuming, as in a benchmark case above, logarithmic utility of consumption, and a utilitarian social welfare function, we would then need to solve

$$\max_{E_1,E_2} \log\left(e^{-\gamma_1(E_1+E_2)} k_1^{\alpha} \left(1 - \frac{E_1}{\chi_1}\right)^{1-\alpha-\nu} E_1^{\nu}\right) + \log\left(e^{-\gamma_2(E_1+E_2)} k_2^{\alpha} \left(1 - \frac{E_2}{\chi_2}\right)^{1-\alpha-\nu} E_2^{\nu}\right).$$

This problem delivers two simple first-order conditions:

$$\gamma_1 + \gamma_2 = \frac{1-\nu-\alpha}{\chi_1 - E_1} - \frac{\nu}{E_1} = \frac{1-\nu-\alpha}{\chi_2 - E_2} - \frac{\nu}{E_2}.$$

It is easy to see from these two equations the only parameters that influence emissions in country i are parameters specific to that country plus the damage elasticity parameter of the other country. Suppose now that we try to back out what tax on coal in country i would be necessary to attain this allocation. From the firm's first-order condition we obtain

$$\tau_i = e^{-\gamma_1(E_1+E_2)} k_i^{\alpha} \left(1 - \frac{E_i}{\chi_i}\right)^{1-\alpha-\nu} E_i^{\nu} \left(\frac{1-\nu-\alpha}{\chi_i - E_i} - \frac{\nu}{E_i}\right).$$

Let us now evaluate the right-hand side at the utilitarian optimum as given by the previous equations. This delivers

$$\tau_i = (\gamma_1 + \gamma_2)\, e^{-\gamma_1(E_1+E_2)} k_i^{\alpha} \left(1 - \frac{E_i}{\chi_i}\right)^{1-\alpha-\nu} E_i^{\nu}.$$

Does this imply a uniform tax across countries? The answer is no. We obtain, in particular, that

$$\frac{\tau_1}{\tau_2} = \left(\frac{k_1}{k_2}\right)^{\alpha} \left(\frac{1-\frac{E_1}{\chi_1}}{1-\frac{E_2}{\chi_2}}\right)^{1-\alpha-\nu} \left(\frac{E_1}{E_2}\right)^{\nu} = \frac{y_1}{y_2}.$$

Clearly, this expression is not 1 in general. It depends on the ratio of capital stocks (note that E_1 and E_2 do not) and the expression involving the Es and χs is also not equal to 1 in general: it is above (below) 1 if χ_1 is above χ_2. In the latter case, the richer country imposes a larger tax on carbon. Note, however, that we obtain a common tax rate, i.e., a *common tax on coal per output unit.*

We have learned from the above analysis (i) that the Pareto optimum involves a globally uniform tax on coal (along with some chosen lump-sum transfers across regions) but (ii) the utilitarian optimum assuming no transfers across regions does not, and instead prescribes—in the benchmark case we look at—a tax that is proportional to the country's output. It is straightforward to go through a similar exercise with population sizes differing across regions; in this case, the optimal tax rate in region i is equal to the region's per-capita income times the world's population-weighted γs.

4.15.3 Policy heterogeneity and carbon leakage

International agreements appear hard to reach and it is therefore of interest to analyze policy heterogeneity from a more general perspective. So suppose region 1 considers a tax on its fossil fuel but knows that region 2 will not use taxes. What are the implications for the output levels of the two regions and for the climate implied by such a scenario? We again begin the analysis by looking at the case of oil, and we start off by assuming that neither capital nor labor can move across regions.

In a decentralized equilibrium, oil use in region 1 is given by

$$p + \tau = \nu e^{-\gamma_1(E_1+E_2)} k_1^{\alpha} n_1^{1-\alpha-\nu} E_1^{\nu-1}$$

and in region 2 it is given by

$$p = \nu e^{-\gamma_2(E_1+E_2)} k_2^{\alpha} n_2^{1-\alpha-\nu} E_2^{\nu-1}.$$

Thus, we can solve for E_1 and E_2 given $E_1 + E_2 \leq \bar{E}$. Clearly, we must have $p > 0$—otherwise, region 2 would demand an infinite amount of oil—and so we first conclude that $E_1 + E_2 = \bar{E}$: there is no way for one country, however large, to influence total emissions. What the tax will do is change energy use across regions: region 1 will use less and region 2 more. Moreover, in utility terms region 1 is worse off and region 2 better off from this unilateral tax policy. This example illustrates direct (and full) carbon leakage: if one region taxes oil, oil use will fall in this region but there will be an exact offset elsewhere in the world.

In the coal example, the situation is rather different. The laissez-faire allocation is now given by

$$\tau_1 = e^{-\gamma_1(E_1+E_2)} k_1^{\alpha} \left(1 - \frac{E_1}{\chi_1}\right)^{1-\alpha-\nu} E_1^{\nu} \left(\frac{1-\nu-\alpha}{\chi_1 - E_1} - \frac{\nu}{E_1}\right)$$

and

$$0 = \frac{1-\nu-\alpha}{\chi_2 - E_2} - \frac{\nu}{E_2}.$$

We see that coal use in region 2 now is independent of the tax policy in region 1.[92] It is easy to show that region 1's coal use will fall and that, at least if both γs are positive and locally around $\tau_1 = 0$, welfare will go up in both regions. There will be an optimal tax, from the point of view of region 1's utility, and it is given by the SCC (computed ignoring the negative externality on region 2), i.e., $\gamma_1 y_1$.

If one allows capital mobility, as in Krusell and Smith (2015), there will be indirect carbon leakage. In the case of oil, a tax in region 1 would act as a multiplier and tilt the relative oil use more across regions, i.e., increase the leakage. In the case of coal, whereas there is no leakage when capital cannot flow, there is now some leakage: the lower use of coal will decrease the return to capital in region 1 and some capital will then move to region 2, in turn increasing emissions there. We thus see that the extent of leakage depends on (i) how costly fossil fuel is to extract and (ii) to what extent other input factor flow across regions.[93]

It would be straightforward to apply this model, and even dynamic versions of it as they can allow closed-form analysis, for a range of qualitative and quantitative studies. A recent example is Hillebrand and Hillebrand (2016), who study tax-and-transfer schemes in a dynamic multi-region version of the model.

4.15.4 More elaborate regional models

Multi-region models of the sort discussed here can be applied rather straightforwardly, and without much relying on numerical solution techniques, in a number of directions. However, some extensions require significant computational work. One example is the case where the intertemporal cross-regional trade is restricted; a specific case is that where there are shocks and these shocks cannot be perfectly insured. Krusell and Smith (2014, 2015) study such models and also compare outcomes across different assumptions regarding such trade; in their models with regional temperature shocks, the model is similar to that in Aiyagari (1994), with the Aiyagari consumers replaced by regions, and where the numerical methods

[92] Our particular assumptions on how coal is produced explains why there is no effect at all on coal use in region 2: the costs and the benefits of coal are both lowered by the same proportion as a result of the tax in region 1. With coal produced with a constant marginal cost in terms of output (as opposed to in terms of labor), there would be a small effect on region 2's coal use.

[93] We did not consider the case where coal is costless to trade and potentially produced in a third region but it is straightforwardly analyzed.

borrow in part from Krusell and Smith (1998). The Krusell and Smith (2015) model has regions represent squares that are 1 by 1 degree on the map; Nordhaus's G-Econ database with population and production on that level of aggregation can then be used to calibrate the model. Thus, the calibration makes the initial model output distribution match that in the data, and the marginal products of capital are assumed to be equal initially—these two restrictions are made possible by choosing TFP and capital-stock levels for each region. There is also heterogeneity in two aspects of how regions respond to climate change. One is that for any given increase in global temperature, the regional responses differ quite markedly according to certain patterns, as discussed in Section 3.1.4 above; Krusell and Smith use the estimates implied by a number of simulations of advanced climate models to obtain region-specific parameters. These estimated "climate sensitivities" are plotted by region on the global map in left panel of Figure 4.15.4 below.

A second element is differences in damages from climate change across regions. In the latest version of their work and as mentioned in Section 3.3.3, Krusell and Smith use the assumption that there is a common, U-shaped damage function for all regions defined in terms of the local temperature, i.e., there ideal temperature is the same at all locations. This common damage function has three parameters which are estimated to match, when the model is solved, the aggregate (global) damages implied by Nordhaus's DICE damage function for three different warming scenarios (1, 2.5, and 5 degrees of global warming). The estimates imply that an average daily temperature of 11.1 degrees Celsius (taken as a 24-hour average) is optimal.

The right panel of Figure 4.15.4 displays the model's predicted laissez-faire outcomes in year 2200. We see large gains in percent of GDP in most of the northern parts of the northern hemisphere and large losses in the south. Overall, the damage heterogeneity is what is striking here: the differences across regions swamp those obtained for any comparisons over time of global average damages. The results in this figure of course rely on the assumption that the damage function is the same everywhere so that warming implies gains for those regions that are too cold initially and losses for those that are too warm. This, however, seems like a reasonable assumption to start with and, moreover, is in line with recent damage-function estimates using cross-sectional data: see Burke et al. (2015). These results at the very least suggest that the returns from further research on heterogeneity should be rather high.

We already mentioned Hémous's (2014) work on the R&D allocation across regions, emphasizing the importance of understanding the determinants and consequences of the regional distribution of R&D and of trade in goods with different carbon content.[94] Another very promising and recent line of research that we also made reference to above is that on endogenous migration pursued in Brock et al. (2014) and Desmet and Rossi-Hansberg (2015). The latter study, which is an early adopter of the kind of damage-function assumption (for both agriculture and manufacturing) used in the later study by Krusell and Smith (2015), assumes free mobility and that there is technology heterogeneity across regions, with operative region-to-region spillovers. The model structure used by Desmet and Rossi-Hansberg

[94] See also Acemoglu, Aghion, and Hémous, 2014.

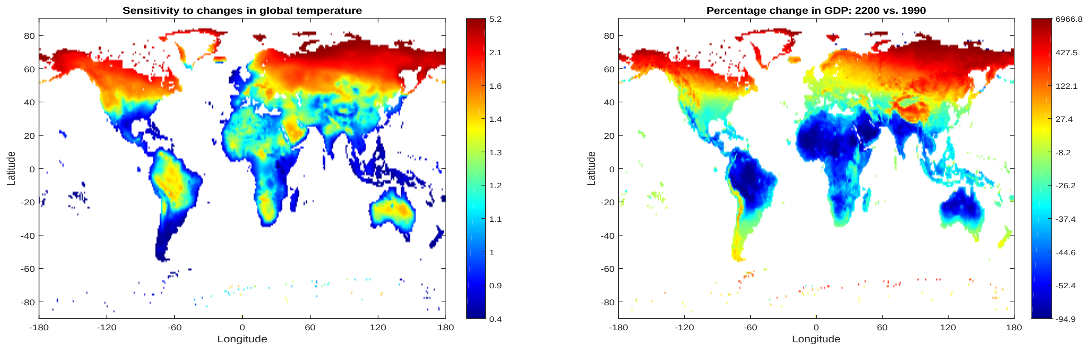

Figure 13: Left: temperature increases for global warming of 1 degree; Right: simulation of Krusell and Smith (2015) model, future % GDP losses under laissez-faire

is particularly tractable for the analysis of migration, as it uses indifference conditions to distribute agents across space. In contrast, models where location is a state variable (in a dynamic sense) and moving is costly are much more difficult to characterize, as moving then is a highly multidimensional and nonlinear problem both with regard to state and control variables. Stylized two-region models like those studied herein and in Hémous's work can perhaps be solved for endogenous migration outcomes but full dynamics are probably very challenging to solve for.

5 Dynamic IAMs

Even though the static IAM setting analyzed in the previous section is useful in many ways, its value in quantitative evaluations is limited: climate change plays out very slowly over time—the dynamics of the carbon cycle especially—and the intertemporal economics aspects involving the comparison between consumption today and consumption far out in the future are therefore of essence. Thus, a quantitatively oriented integrated assessment model of economics and climate change needs to incorporate dynamics. In addition, there are some conceptual issues that cannot be properly discussed without a dynamic setting, such as time preferences.

To our knowledge, the first steps toward modern integrated assessment model appear in Nordhaus (1977). A little over a decade later, Nordhaus developed a sequence of dynamic models, all in the spirit of the simple model above, but formulated in sufficient complexity that numerical model solution is required. The core, one-region version of Nordhaus's model is DICE: a Dynamic Integrated Climate-Economy model, described in detail in Nordhaus and Boyer (2000). In one respect, almost all the dynamic IAMs, including Nordhaus's, are more restrictive than the setting in our previous section: they focus on a planning problem, i.e., on characterizing optimal allocations. That is, decentralized equilibria without carbon policy, or with suboptimal carbon policy, are rarely analyzed, let alone explicitly discussed in dynamic models.[95] In our present treatment, we insist on analyzing both optima and suboptimal equilibria, in large part because the quantitative assessments of the "cost of inaction" cannot be computed otherwise.

In what follows we will discuss a general structure for which we define the social cost of carbon and, under some additional assumptions, can derive a simple and directly interpretable formula for the tax. It is a straightforward extension of the results from the static model above. This material is contained in Section 5.1. In Section 5.2 we then make further assumptions, relying also on the finite-resource modeling from Section 2, and simplify the general structure so as to arrive at an easily solved, and yet quantitatively reasonable, model that can be used for positive as well as normative analysis. Throughout, the discussion follows Golosov et al. (2014) rather closely.

5.1 The social cost of carbon in a general dynamic model

We now focus on how the SCC is determined in a dynamic setting that is reasonably general. For this, we use a typical macroeconomic model with a representative (for the global economy, at this point) agent, as in Nordhaus's DICE model, a production structure, and a specification of the climate system as well as the carbon cycle.

The representative agent has utility function

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t U(C_t),$$

where U is a standard, strictly concave utility function of (the one and only) consumption good C and where $\beta \in [0,1)$ is the discount factor. The resource constraint for the consumption good is more broadly a constraint for the final good, because like in most of the macroeconomic literature we treat consumption and investment as perfect substitutes. The constraint thus reads

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

which involves a typical capital accumulation specification with geometric depreciation at rate δ and where Y denotes global output. Global output, in turn, is generated from

$$Y_t = F_{0,t}(K_{0,t}, N_{0,t}, \mathbf{E}_{0,t}, S_t).$$

Here, "0" represents the sector producing the final good. The function F_0 is assumed to display constant returns to scale in the first three inputs. $N_{0,t}$ is labor used in this sector and $\mathbf{E}_{0,t} = (E_{0,1,t}, \ldots, E_{0,I,t})$ denotes a vector of different energy inputs. We use a subindex t on the production function to indicate that there can be technical change over time (of various sorts and deterministic as well as stochastic). S, finally, is atmospheric carbon concentration, and it appears in the production function because it causes damages—through the effect of S on the climate (in particular through the temperature).

[95]For an exception, see, e.g., Leach (2007).

In our formulation here, as discussed above, we adopt the common assumption that damages only appear in the production function. Moreover, they only appear in the time-t production function through atmospheric carbon concentration at t, thus subsuming the mapping from S to temperature and that from temperature to output loss in one mapping. As we already argued, these assumptions are convenient in that they map neatly into Nordhaus's DICE model. We should remind the reader that the inclusion of only S_t in the damages at t captures a lack of dynamics; as we pointed out, this should still be a reasonable approximation to a more complex setting where, conceptually, one would include past values of S in the production function at t as a way of capturing the full dynamics. An extension to include such lagged variables is straightforward but would not greatly change the results as the temperature dynamics are rather quick.

Turning to energy production, we assume that there are $I_g - 1$ "dirty" energy sources (involving fossil fuel), $i = 1, \ldots, I_g - 1$, and a set of green sources, $i = I_g, \ldots, I$. Each component of $\mathbf{E}_{0,t}$, $E_{0,i,t}$ for $i = 1, \ldots, I$, is then produced using a technology $F_{i,t}$, which uses the three inputs capital, labor, and the energy input vector. Some energy sources, such as oil, may be in finite supply. For those i in finite supply, $R_{i,t}$ denotes the beginning-of-period stock at t and $E_{i,t}$ the total amount extracted (produced) at t. Thus, the exhaustible stock i evolves as

$$R_{i,t+1} = R_{i,t} - E_{i,t} \geq 0. \tag{20}$$

Production for energy source i, whether it is exhaustible or not, is then assumed to obey

$$E_{i,t} = F_{i,t}\left(K_{i,t}, N_{i,t}, \mathbf{E}_{i,t}, R_{i,t}\right) \geq 0. \tag{21}$$

The resource stock appears in the production function because the production costs may depend on the remaining resource stock. Notice, also, that S_t does not appear in these production functions: we assume that climate change does not cause damages to energy production. This, again, is a simplification we make mainly to adhere to the TFP damage specification that is common in the literature, but it also simplified formulas and improves tractability somewhat. Given that the energy sector is not so large, this simplification should not be a major problem for our quantitative analysis.

To close the macroeconomic part of the model, we assume that inputs are allocated across sectors without costs, again a simplifying assumption but one that appears reasonable if the period of analysis is as long as, say, 10 years. Thus we have

$$\sum_{i=0}^{I} K_{i,t} = K_t, \qquad \sum_{i=0}^{I} N_{i,t} = N_t, \quad \text{and} \quad E_{j,t} = \sum_{i=0}^{I} E_{i,j,t}. \tag{22}$$

We assume that the sequence/process for N_t is exogenous.

Finally, we let the carbon cycle generally be represented by a function $\tilde{S}_t$ as follows:

$$S_t = \tilde{S}_t\left(E^f_{i,-T}, E^f_{-T+1}, \ldots, E^f_t,\right). \tag{23}$$

Here, T periods back represents the end of the pre-industrial era and $E^f_s \equiv \sum_{i=1}^{I_g-1} E_{i,s}$ is fossil emission at s and we recall that $E_{i,s}$ is measured in carbon emission units for all i.

When we specialize the model, we will adopt a very simple structure for $\tilde{S}_t$ that is in line with the discussion in the section above on the carbon cycle.

We are now ready to state an expression for the SCC. Using somewhat abstract (but obvious) notation, and denoting the social cost of carbon at time t, in consumption units at this point in time, by SCC_t, we have

$$\text{SCC}_t = \mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \frac{U'(C_{t+j})}{U'(C_t)} \frac{\partial F_{0,t+j}}{\partial S_{t+j}} \frac{\partial S_{t+j}}{\partial E_t^f}. \tag{24}$$

Before we discuss this equation, let us emphasize—as we pointed out in the context of the static model—that this expression amounts to keeping decisions fixed as emissions are increased incrementally. I.e., this concept of the social cost of carbon does not correspond to a policy experiment (where presumably induced changes in decisions would add indirect damage effects, positive or negative). Golosov et al. (2014) derive this equation as part of an optimal allocation but then the interpretation really is that the right-hand side equals the OSCC_t.

Equation (24) is easily interpreted. First, $\frac{\partial S_{t+j}}{\partial E_t^f}$ captures the carbon cycle dynamics: it tells us how much the atmospheric carbon content j periods ahead is increased by a unit emission at t. That amount of increase in S_{t+j} then changes final output in period $t+1$ by $\frac{\partial F_{0,t+j}}{\partial S_{t+j}}$ per unit. The total effect (the multiplication of these two factors), which is presumably negative, is the marginal damage in that period in terms of the final output good arising from a unit of emission at t. To translate this amount into utils at $t+j$ one multiplies by $U'(C_{t+j})$, and to bring the utils at $t+j$ back to time-t utils one multiplies by β^j: utility discounting. The division by $U'(C_t)$ then translates the amount back into consumption units at t. Finally, since one needs to take into account the effect of emissions at all points in time $t, t+1, \ldots$, one needs the infinite sum.

Conceptually, thus, equation (24) really is straightforward. However, in its general form it is perhaps not so enlightening. A key result in Golosov et al. (2014) is that with some assumptions that the authors argue are weak, one can simplify the formula considerably and even arrive at a closed-form expression in terms of primitive parameters. We present the assumptions one by one.

Assumption 1 *$U(C) = \log C$.*

Logarithmic utility, both used and relaxed in our static model, is very often used in macroeconomic models and seems appropriate as a benchmark. It embodies an assumption about the intertemporal elasticity of consumption but obviously also about risk aversion.

Assumption 2

$$F_{0,t}\left(K_{0,t}, N_{0,t}, \mathbf{E}_{0,t}, S_t\right) = \exp\left(-\gamma_t S_t\right) \tilde{F}_{0,t}\left(K_{0,t}, N_{0,t}, \mathbf{E}_{0,t}\right),$$

where we have normalized so that S is the atmospheric CO_2 concentration in excess of that prevailing in pre-industrial times, as in the above section, and where γ can be time- and state-dependent.

This assumption was discussed in detail in Section 10: we argue that it allows a good reduced-form approximation to the most commonly used assumptions on the S-to-temperature and the temperature-to-damage formulations in this literature.

Assumption 3

$$S_t = \sum_{s=0}^{t+T} (1 - d_s) E^f_{t-s} \tag{25}$$

where $d_s \in [0, 1]$ for all s.

A linear carbon cycle was also discussed Section 3.2.4 on carbon circulation above and argued to be a good approximation. The linear structure was also simplified further there, and we will use that simplification below.

Assumption 4 *C_t/Y_t does not depend on time.*

This assumption, which is tantamount to that used in the textbook Solow model, is not an assumption on primitives as we usually define them. However, it is an assumption that can be shown to hold exactly for some assumptions on primitives—as those that will be entertained below—or that holds approximately in a range of extensions; see Barrage (2014). Major changes in saving behavior away from this assumption are needed to drastically alter the quantitative conclusions coming out of our SCC formula.

Now given these four assumptions only a minor amount of algebra suffices to arrive at a formula for the SCC, as well as for the optimal tax on carbon. It is

$$\text{SCC}_t = Y_t \left[\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \gamma_{t+j} (1 - d_j) \right]. \tag{26}$$

As can be seen, this formula is a straightforward extension of that arrived at for the static economy. As in the static economy, the formula for the tax as a fraction of output is a primitive: there, simply γ; here, a present value of sorts of future γs. Note, of course, here as well as for the static model, that if one needs to assign a specific value to the optimal tax, one would strictly speaking need to evaluate output at its optimal level, and the optimal level of output is not expressed in closed form here (and may be cumbersome to compute). However, given our quantitative analysis below, we note that the optimal tax rate does not alter current output so much. Hence, a good approximation to the optimal tax rate is that given by the expression in brackets in equation (26) times current output.[96]

In the static economy, we assumed a Cobb-Douglas form for output, as we will in the next section as well for our positive analysis. However, Cobb-Douglas production is apparently not necessary for the result above. What is true is that Cobb-Douglas production, along

[96]In the dynamic model, this approximation would overstate the exact value of the tax since optimal output in the short run will be lower than laissez-faire output. In the static model with TFP damages, the reverse inequality will hold.

with logarithmic utility and 100% depreciation for capital, are very helpful assumptions for arriving at a constant C/Y ratio (Assumption 4), but we also know that an approximately constant C/Y ratio emerges out of a much broader set of economies.

We note that, aside from the damage parameter γ, utility discounting and carbon depreciation now matter very explicitly as well. This is quite intuitive: it matters how long a unit of emitted carbon stays in the atmosphere and it also matters how much we care about the future. As for how γ appears, note that the formula is an expectation over future values—as in the static model, a certainty equivalence of sorts applies—but that one could also imagine γ as evolving over time, or incorporating different amounts of uncertainty at different points in time.[97] Of course, suppose more information is revealed about γ as time evolves, the optimal tax will evolve accordingly (as, e.g., in a specification where γ is assumed to follow a unit-root process).

A final expression of our SCC is obtained by (i) assuming that $\mathbb{E}_t\left[\gamma_{t+j}\right] = \bar{\gamma}_t$ for all j (as for example for a unit root process) and (ii) letting the $1-d_j$s be defined by equation (13) above (which we argued gives a good account of the depreciation patterns). Then we obtain

$$\mathrm{SCC}_t/Y_t = \bar{\gamma}_t \left(\frac{\varphi_L}{1-\beta} + \frac{(1-\varphi_L)\varphi_0}{1-(1-\varphi)\beta} \right). \tag{27}$$

Here, the expression inside the parenthesis on the right-hand side can be thought of as the *discount-weighted duration of emissions*, an object that is stationary by assumption here.

A remarkable feature of the formula for the SCC as a fraction of output as derived here is that it depends on very few parameters. In particular, no production parameters appear, nor do assumptions about technology or the sources of energy. In contrast, we will see in the positive analysis below that such assumptions matter greatly for the paths of output, the climate, energy use, and the total costs of suboptimal climate policy. These are obviously important as well, so we need to proceed to this analysis. However, for computing what optimal policy is, straightforward application of the formula above works very well, and in some sense is all that is needed to optimally deal with climate change. To compute the optimal quantity restrictions is much more demanding, because then precisely all these additional assumptions are made, and to predict the future of technology (especially that regarding energy supply) is extremely difficult, to say the least. Section 5.2.3 below calibrates the key parameters behind the formula above and Section 5.2.4 then displays the numerical results for the social cost of carbon.

5.2 A positive dynamic model

The positive dynamic model will be a straightforward extension of the static model in Section 4 in combination with the basic model from Section 2.3.2 (without endogenous technical change).

[97] Learning (about γ or the natural-science parameters) could also be introduced formally, as in the planning problem studied by Kelly and Kolstad (1999).

Thus we assume a production function that is Cobb-Douglas in capital, labor, and an energy input, along with TFP damages from climate:

$$Y_t = e^{-\gamma_t S_t} A_t K_t^{\alpha} N_{0t}^{1-\alpha-\nu} E_t^{\nu}. \tag{28}$$

Here, we maintain the possibility that γ changes over time/is random.

There are three energy-producing sectors, as in one of the extensions of the static model. Sector 1 thus produces "oil", which is in finite supply and is extracted at zero cost. The accounting equation $E_{ot} = R_t - R_{t+1}$ thus holds for oil stocks at all times. The second and third sectors are the "coal" and the "green" sectors, respectively. They deliver energy using

$$E_{i,t} = \chi_{it} N_{it} \text{ for } i = c, g. \tag{29}$$

Here, $N_t = N_{0t} + N_{ct} + N_{gt}$. We will focus on parameters such that coal, though in finite supply, will not be used up; hence, its Hotelling premium will be zero and there will be no need to keep track of the evolution of the coal stock. [98] This specification captures the key stylized features of the different energy sectors while maintaining tractability. In practice, oil (as well as natural gas) can be transformed into useable energy quite easily but these resources are in very limited supply compared to coal. Coal is also more expensive to produce, as is green energy.

Here, energy used in production of the final good, E_t, then obeys

$$E_t = \left(\kappa_o E_{ot}^{\rho} + \kappa_c E_{ct}^{\rho} + \kappa_g E_{gt}^{\rho}\right)^{1/\rho} \tag{30}$$

with $\sum_{i=o,c,g} \kappa_i = 1$. As before, $\rho < 1$ regulates the elasticity of substitution between different energy sources; the κs are share parameters and also influence the efficiency with with the different energy sources are used in production. In addition, coal is "dirtier" than oil in that it gives rise to higher carbon emissions per energy unit produced. With E_{ot} and E_{ct} in the same units (of carbon emitted), the calibration therefore demands $\kappa_o > \kappa_c$.

The variables A_t, χ_{it}, and N_t are assumed to be exogenous and deterministic. Population growth is possible within our analytically tractable framework but we abstract from considering it explicitly in our quantitative exercises below, since A and N play the same role.[99] Our final assumption, which is key for tractability, is that capital depreciates fully between periods ($\delta = 1$). This is an inappropriate assumption in business-cycle analysis but much less so when a model focusing on long-run issues; a model period will be calibrated to be 10 years.

[98] This will, under some specifications, require that a *back-stop technology* emerge at a point in the future, i.e., a technology that simply replaces coal perfectly at lower cost.

[99] We formulate the utility function in terms of total consumption, and we do not adjust discounting for population growth. One might want to consider an alternative here, but we suspect that nothing substantial will change with this alternative.

5.2.1 Solving the planner's problem

For brevity, we do not state the planner's problem; it is implicit from the description above. The first-order conditions for C_t and K_t yield

$$\frac{1}{C_t} = \beta \mathbb{E}_t \frac{\alpha}{C_{t+1}} \frac{Y_{t+1}}{K_{t+1}}.$$

Together with the resource constraint

$$C_t + K_{t+1} = Y_t$$

we then obtain an analytical solution for saving as $K_{t+1} = \alpha\beta Y_t$ for all t. It follows that C_t/Y_t is equal to $1-\alpha\beta$ at all times, and we have therefore demonstrated that Assumption 4 is verified for this economy. A byproduct of our assumptions here, then, are that the formula for the optimal carbon tax, equation (26), holds exactly.

What is the planner's choice for the energy inputs, and what is the resulting effect on atmospheric carbon concentration and, hence, the climate? First, we assume that $\rho < 1$, and from this Inada property we then conclude that the energy choices will be interior at all times. Looking at the first-order conditions for E_t and E_{ot}, we obtain

$$\frac{\nu\kappa_o}{E_{ot}^{1-\rho} E_t^{\rho}} - \frac{\mathrm{SSC}_t}{Y_t} = \beta \mathbb{E}_t \left(\frac{\nu\kappa_o}{E_{o,t+1}^{1-\rho} E_{t+1}^{\rho}} - \frac{\mathrm{SSC}_{t+1}}{Y_{t+1}} \right), \tag{31}$$

where SSC_t/Y_t is, again, defined equation (26). This equation expresses Hotelling's formula in the case where there is a cost of using carbon: the damage externality (thus, playing a similar role to an extraction cost).

Looking at the other two energy source, by choosing $N_{i,t}$ optimally we obtain

$$\chi_{ct} \left(\frac{\nu\kappa_c}{E_{ct}^{1-\rho} E_t^{\rho}} - \frac{\mathrm{SCC}_t}{Y_t} \right) = \frac{1-\alpha-\nu}{N_t - \frac{E_{ct}}{\chi_{ct}} - \frac{E_{gt}}{\chi_{gt}}} \tag{32}$$

and

$$\chi_{gt} \frac{\nu\kappa_g}{E_{gt}^{1-\rho} E_t^{\rho}} = \frac{1-\alpha-\nu}{N_t - \frac{E_{ct}}{\chi_{ct}} - \frac{E_{gt}}{\chi_{gt}}}. \tag{33}$$

From the perspective of solving the model conveniently, it is important to note now that SSC_t/Y_t is available in closed form as a function of primitives: the remaining system of equations to be solved is a vector difference equation but only in the energy choices. I.e., the model can be solved for energy inputs first, by solving this difference equation, and then the rest of the variables (output, consumption, etc.) are available in the simple closed forms given above.

To solve the vector difference equation—to the extent there is no uncertainty—is also simple, though in general a small amount of numerical work is needed.[100] A robust numerical

[100] Solving the model with only coal or only green energy is possible in closed form.

method goes as follows. With any given value for E_{ot}, the equations (32) and (33) can be used to solve for E_{ct} and E_{gt}, and thus E_t. The solution is nonlinear but well defined. For any given initial stock of oil R_0, one can now use a simple shooting algorithm. The "shooting" part is accomplished by (i) guessing on a number for E_{o0}; (ii) deriving the all the other energy inputs at time 0; (iii) using the Hotelling equation (31), which is stated in terms of E_{o1} and E_1, to obtain E_{o1} as a function of E_1; (iv) combining this relation between E_{o1} and E_1 with equations (32) and (33) evaluated for period 1 to obtain all the energy choices in period 1; and (v) going back to step (iii) to repeat for the next period. The so-obtained path for all energy inputs in particular delivers a path for oil extraction. To check whether the fired shot hits the target involves simply checking that the cumulated oil use exactly exhausts the initial stock asymptotically. If too much or too little is used up, adjust E_{o0} appropriately and run through the algorithm again.

If there is uncertainty about γ that is nontrivial and does not go away over time, one needs to use recursive methods, given the nonlinearity of the vector difference equation. It is still straightforward to solve, however, with standard versions of such methods.

5.2.2 Competitive equilibrium

It is straightforward to define a dynamic (stochastic) general equilibrium for this economy as for the static model. All markets feature perfect competition. Firms in the final-goods sector make zero profits, as do firms in the coal and green-energy sectors. In the oil sector, there is a Hotelling rent, and hence profits. These profits are delivered to the representative consumer, who otherwise receive labor and capital income and, to the extent there is a tax on fossil fuel, lump-sum transfers so that the government budget balances. When taxes are used, we assume that they are levied on the energy-producing firms (oil and coal). The consumer's Euler equation and the return to capital satisfying the first-order condition for capital from the firm's problem deliver the constant saving rate $\alpha\beta$. The energy supplies (or, equivalently, the labor allocation) is then given by a set of conditions similar to those from the planning problem. Assuming that the carbon tax in period t is set as an exogenous fraction of output in period t, we then obtain from the energy producers' problems

$$\frac{\nu\kappa_o}{E_{ot}^{1-\rho}E_t^{\rho}} - \tau_t = \beta\mathbb{E}_t\left(\frac{\nu\kappa_o}{E_{o,t+1}^{1-\rho}E_{t+1}^{\rho}} - \tau_{t+1}\right), \tag{34}$$

$$\chi_{ct}\left(\frac{\nu\kappa_c}{E_{ct}^{1-\rho}E_t^{\rho}} - \tau\right) = \frac{1-\alpha-\nu}{N_t - \frac{E_{ct}}{\chi_{ct}} - \frac{E_{gt}}{\chi_{gt}}}, \tag{35}$$

and

$$\chi_{gt}\frac{\nu\kappa_g}{E_{gt}^{1-\rho}E_t^{\rho}} = \frac{1-\alpha-\nu}{N_t - \frac{E_{ct}}{\chi_{ct}} - \frac{E_{gt}}{\chi_{gt}}}. \tag{36}$$

Since this vector difference equation is very similar to the planner's vector difference equation, it can be solved straightforwardly with the same kind of algorithm. The laissez-faire allocation is particularly simple to solve.

5.2.3 Calibration and results

In the spirit of quantitative macroeconomic modeling, the calibration of our model parameters is critical. Also in this part, we follow Golosov et al. (2014) in selecting parameter values. The calibration is important to review in some detail here, as calibration of this class of models is not standard in the macroeconomic literature. Given our assumptions, two parameters are easy to select: we assume that α and ν are 0.3 and 0.04, respectively; the value for the capital share is standard in the macroeconomic literature and the energy share is taken from the calibration in Hassler et al. (2015).

Discounting

As will be clear from our results, the discount factor matters greatly for what optimal tax to recommend. We do not take stand here but rather report our results for a range of values for β. Nordhaus's calibrations start from interest-rate data; interest rates should mirror the interest rate, if markets work, so to set $1/\beta - 1 = 0.015$ is then reasonable. Stern, in his review on climate change, takes a very different view and uses what is essentially a zero rate: $1/\beta - 1 = 0.001$. A view that sharply differs from the market view can be motivated on purely normative grounds, though then there may be auxiliary implications of this normative view: perhaps capital accumulation should then be encouraged more broadly, e.g., using broad investment/saving subsidies. Persson and Sterner (2007), however, argue informally that it is possible to discount consumption and climate services—to the extent the latter enter separately in utility—at different rates.

A third and, we think, interesting argument for using a lower discount rate is that it is reasonable to assume that discounting is time-inconsistent: people care about themselves and the next generation or so with rates in line with observed market rates but thereafter, they use virtually no discounting. The idea would be that I treat the consumption of my grand-grand-grand children and that of my grand-grand-grand-grand children identically in my own utility weighting. If this is a correct description of people's preferences, and if people have commitment tools for dealing with time inconsistency, we would see it in market rates, but there are not enough market observations for such long-horizon assets to guide a choice of discount rates. Hence, it is not easy to reject a rate such as 0.1% (but, by the same token, there is no market evidence in favor of it either). If people have no commitment tools for dealing with time inconsistency, observed market rates today would be a mix of the short- and long-run rates (and very heavily weighted toward present-bias), thus making it hard to use market observations to back out the longer-run rates. These arguments can be formalized: it turns out that the present model—if solved with a simplified energy sector (say, coal only)—can be solved analytically also with time-inconsistent preferences (see Karp, 2005, Gerlagh and Liski, 2012, and Iverson, 2014).

The carbon cycle

We calibrate the carbon cycle, as indicated, with a linear system implying that the carbon depreciation rates are given by equation (13). Thus with the depreciation rate at horizon j given by $1 - d_j = \varphi_L + (1 - \varphi_L)\,\varphi_0\,(1 - \varphi)^j$, we have to select three parameter: φ_L, φ_0, and φ. Recall the interpretation that φ_L is the share of of carbon emitted into the atmosphere that stays there forever, $1 - \varphi_0$ the share that disappears into the biosphere and the surface oceans within a decade, and the remaining part, $(1-\varphi_L)\,\varphi_0$, decays (slowly) at a geometric rate φ. We set φ_L to 0.2, given the estimate in the 2007 IPCC report that about 20% any emission pulse remains in the atmosphere for several thousand years.[101] Archer (2005), furthermore, argues that the excess carbon that does depreciate has a mean lifetime of about 300 years. Thus, we set $(1-\varphi)^{30} = 0.5$, implying $\varphi = 0.0228$. Third, the 2007 IPCC report asserts that about 50% of any CO_2 emission pulse into the atmosphere has left the atmosphere after about 30 years. This means that $d_2 = 0.5$ so that $1 - \frac{1}{2} = 0.2 + 0.8\varphi_0(1 - 0.0228)^2$, and hence $\varphi_0 = 0.393$. Finally, to set the initial condition for carbon concentration we showed above that the assumed depreciation structure is consistent with the existence of two "virtual carbon stocks" S_1 (the part that remains in the atmosphere forever) and S_2 (the part that depreciates at rate φ), with $S_{1,t} = S_{1,t-1} + \varphi_L E_t^f$ and $S_{2,t} = \varphi S_{2,t-1} + \varphi_0(1-\varphi_L)E_t^f$, and $S_t = S_{1,t} + S_{2,t}$. We choose starting values so that time-0 (i.e., year-2000) carbon equals 802, with the division $S_1 = 684$ and $S_2 = 118$; the value of S_1 comes from taking the pre-industrial stock of 581 and adding 20% of accumulation emissions.[102]

Damages

Turning to the calibration of damages, recall that we argued that for a reasonable range of carbon concentration levels the exponential TFP expression $e^{-\gamma S}$ is a good approximation to the composed S-to-temperature and temperature-to-TFP mappings in the literature. It remains choose γ, deterministic or stochastic. Here, in our illustrations, we will focus on a deterministic γ and only comment on uncertainty later. Following the discussion in the damage section above and Golosov et al. (2014), with S measured in GtC (billions of tons of carbon), an exponential function with parameter $\gamma_t = 5.3 \times 10^{-5}$ fits the data well.

Energy

Turning, finally, to the energy sector, we first need to select a value for ρ, which guides the elasticity of substitution between the energy sources. Stern (2012) is a metastudy of 47 studies of interfuel substitution and reports the unweighted mean of the oil-coal, oil-electricity, and coal-electricity elasticities to be 0.95. Stern's account of estimates of "long-run dynamic elasticities" is 0.72. In terms of our ρ, the implied numbers are -0.058 and -0.390, respectively, and the former will constitute our benchmark.

As for the different energy sources, for oil we need to pin down the size of the oil reserve. According to BP (2010), the proven global reserves of oil are 181.7 gigatons. However, these figures only refer to reserves that are economically profitable to extract at current conditions. Rogner (1997), on the other hand, estimates the global reserves of potentially extractable oil, natural gas, and coal taken together to be over 5,000 gigaton, measured as oil equivalents.[103] Of this amount, Rogner reports around 16% to be oil, i.e., 800 gigatons. We use a benchmark that is in between these two numbers: 300 gigatons. To express fossil fuel in units of carbon content, we set the carbon content in crude oil to be 846KgC/ton oil. For coal, we set it

[101] Archer (2005) argues for a slightly higher number: 0.25.

[102] These number include the pre-industrial stock and, hence, do not strictly follow the notation above, where S_t denotes the concentration in excess of pre-industrial levels.

[103] The difference in energy content between natural gas, oil, and various grades of coal is accounted for by expressing quantities in oil equivalents.

to the carbon content of anthracite, which is 716KgC/ton coal.[104] As for coal, as implied by Rogner's (1997) estimates, the coal supply is enough for several hundreds of years of consumption at current levels, and hence we have assumed the scarcity rent to be zero.

To calibrate κ_o and κ_c we use relative prices of oil to coal and oil to renewable energy, given by

$$\frac{\kappa_o}{\kappa_c}\left(\frac{E_{ot}}{E_{ct}}\right)^{\rho-1} \quad \text{and} \quad \frac{\kappa_o}{1-\kappa_o-\kappa_c}\left(\frac{E_{ot}}{E_{gt}}\right)^{\rho-1},$$

respectively. The average price of Brent oil was \$70 per barrel over the period 2005–2009 (BP, 2010); with a barrel measuring 7.33 metric tons and a carbon content of 84.6%, the oil price per ton of carbon is then \$606.5. As for coal, its average price over the same period is \$74/ton. With coal's carbon content of 71.6%, this implies a price of \$103.35 per ton of carbon.[105] The implied relative price of oil and coal in units of carbon content is 5.87.

As for renewables/green energy, there is substantial heterogeneity between different such sources. With unity as a reasonable value of the current relative price between green energy and oil, we employ data on global energy consumption to finally pin down the κs. Primary global energy use in 2008 was 3.315 Gtoe (gigaton of oil equivalents) of coal, 4.059 of oil, 2.596 of gas, and $0.712 + 0.276 + 1.314 = 2.302$ of nuclear, hydro, and biomass/waste/other renewables. Based on the IPCC tables quoted above, the ratio of energy per ton between oil and anthracite is then $\frac{42.3}{26.7} = 1.58$, implying that one ton of oil equivalents is 1.58 tons of coal.[106] With these numbers and the value for ρ of -0.058, we can finally use the equations above to back out $\kappa_o = 0.5008$ and $\kappa_c = 0.08916$.

The parameters χ_{ct}, which determines the cost of extracting coal over time, are set based on an average extraction cost of \$43 per ton of coal (see IEA, 2010, page 212). Thus, a ton of carbon in the form of coal costs \$43/0.716. The model specifies the cost of extracting a ton of carbon as $\frac{w_t}{\chi_{ct}}$, where w_t is the wage. The current shares of world labor used in coal extraction and green energy production is very close to zero, so with total labor supply normalized to unity we can approximate the wage to be $w_t = (1-\alpha-\nu)\,Y_t$. With world GDP at \$700 trillion per decade and a gigaton of carbon (our model unit) costing $w_t/\chi_{ct} = (1-\alpha-\nu)\,Y_t/\chi_{ct}$ to produce delivers $43 \cdot 10^9/0.716 = 0.66 \cdot 700 \cdot 10^{12}/\chi_{c0}$ and hence $\chi_{c0} = 7{,}693$. This means, in other words, that a share $\frac{1}{7{,}693}$ of the world's labor supply during a decade is needed to extract one gigaton of carbon in the form of coal. The calibration of χ_{g0} comes from using the fact that χ_{g0}/χ_{c0} equals the relative price between coal and green energy, thus delivering $\chi_{g0} = 7{,}693/5.87 = 1{,}311$ since the prices of oil and green are assumed to be equal and the relative price of oil in terms of coal is 5.87. Lastly, we posit growth in both χ_{ct} and χ_{gt} at 2% per year.[107]

[104] IPCC (2006), table 1.2–1.3.

[105] BP (2010) gives these estimates for U.S. Central Appalachian coal.

[106] The amounts of oil and coal in carbon units is obtained by multiplying by the carbon contents 84.6 and 71.6%, respectively.

[107] Under our calibration, coal use does not go to zero, which contradicts it being a finite resource. Strictly speaking, one should instead, then, solve the model under this assumption and the implication that coal would have scarcity value. But we consider it quite likely that a competitive close and renewable substitute

5.2.4 Results

We begin by reporting what our model implies for the optimal tax on carbon. Given our calibration, and expressed as a function of the discount rate, we plot the tax per ton of emitted carbon in Figure 14, given annual global output of 70 trillion dollars.[108]

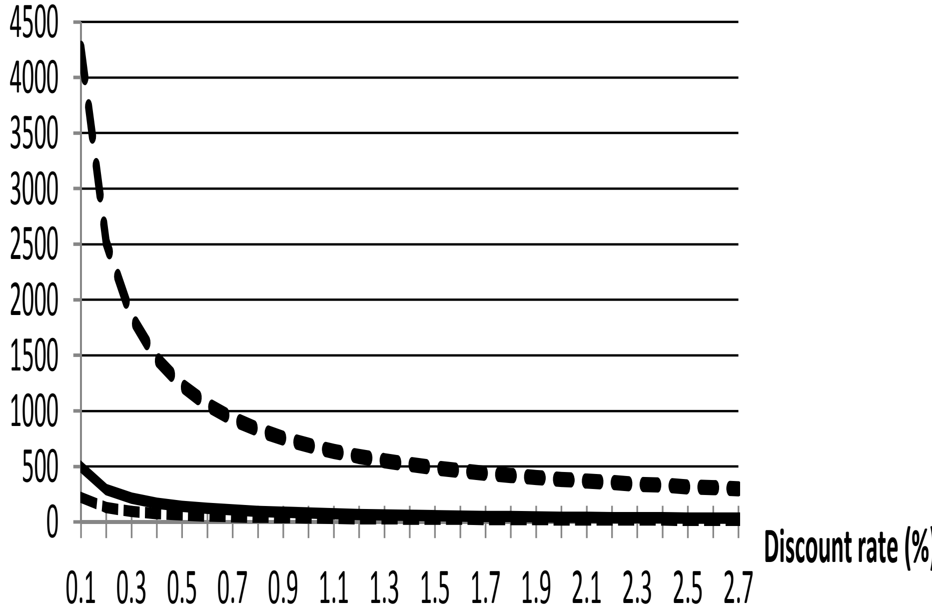

Figure 14: Optimal tax rates in current dollars per ton of emitted fossil carbon vs. yearly subjective discount rate

Figure 14 displays our benchmark as a solid line along with two additional lines representing two alternative values for γ, the higher one of which represents a "catastrophe scenario" with losses amounting to about 30% of GDP and the lower one representing an opposite extreme case with very low losses. The numbers in the figure can be compared to the well-known proposals in Nordhaus and Boyer (2000) and in the Stern review (Stern, 2007), who suggest a tax of \$30 and \$250 dollar per ton of carbon, respectively. As already pointed out, these proposals are based on very different discount rates, with Nordhaus using 1.5% per year and Stern 0.1%. For these two discount-rate values, the optimal taxes using our analysis are \$56.9/ton and \$496/ton, respectively, thus showing larger damages than in these studies. There are a number of differences in assumptions between the model here and those maintained in, say, Nordhaus's work; perhaps the most important one quantitatively is that we calibrate the duration of carbon in the atmosphere to be significantly higher.

---

for coal is invented over the next couple of hundred years, in which case our solution would work well as an approximation.

[108]The graphs are taken from Golosov et al. (2014).

The figure reveals that, to the extent the catastrophe scenario—which comes from a hypothesis Nordhaus entertained in a survey study—might actually materialize, there will be dramatic consequences on the level of the optimal tax: we see that the tax is roughly multiplied by a factor 20.

5.2.5 Positive implications

Fossil fuel use in the optimal allocation and in the *laissez-faire* allocation are shown in Figure 15. We base our results in this section on the discount rate 1.5%.

Looking at the comparison between the optimum and laissez faire, we see a markedly lower use of fossil fuel in the optimum.[109] In the laissez-faire scenario, there would be a continuous increase in fossil fuel use, but in the optimum the consumption of fossil fuel is virtually flat.

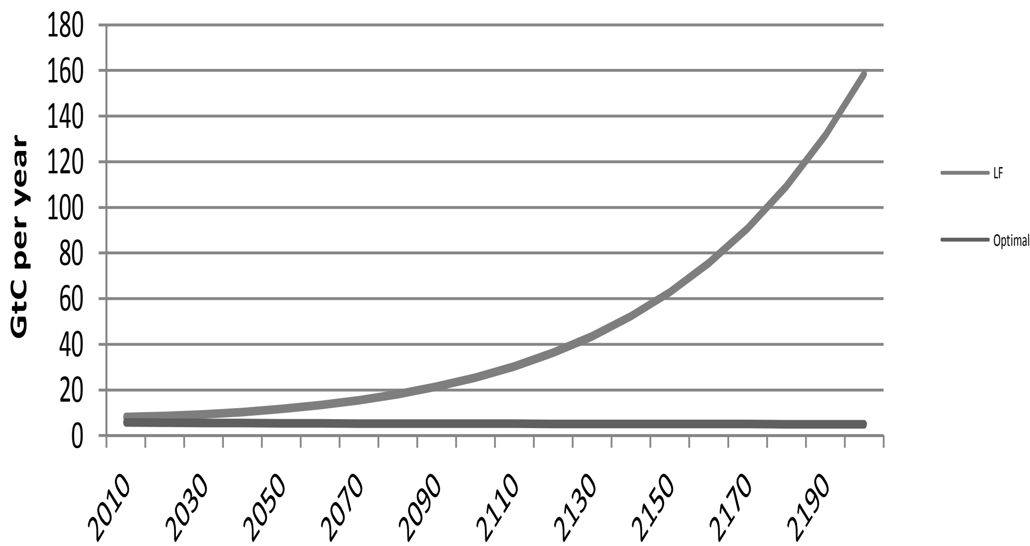

Figure 15: Fossil fuel use: optimum vs. laissez faire

It is important to realize that the difference between the fossil-fuel use in the optimum and in laissez faire is almost entirely coming from a lower coal use in the former. In Figures 16 and 17, we look separately at coal use and oil use in the optimal vs. the laissez-faire allocations. Although the tax on carbon is identical for oil and coal in the optimal allocation, its effects are very different: coal use is simply curbed significantly—the whole path is shifted down radically—but oil use is simply moved forward slightly in time. With optimal taxes, coal use would fall right now to almost half; a hundred years from now, laissez-faire coal use would

[109]The model predicts coal use in laissez faire of 4.5 GtC during the coming decade; it is currently roughly 3.8 GtC. It predicts oil use of 3.6 GtC, which is also close to the actual value for 2008 or 3.4 GtC.

be seven times higher than optimally. Green energy use is very similar across the optimum and laissez-faire allocations.

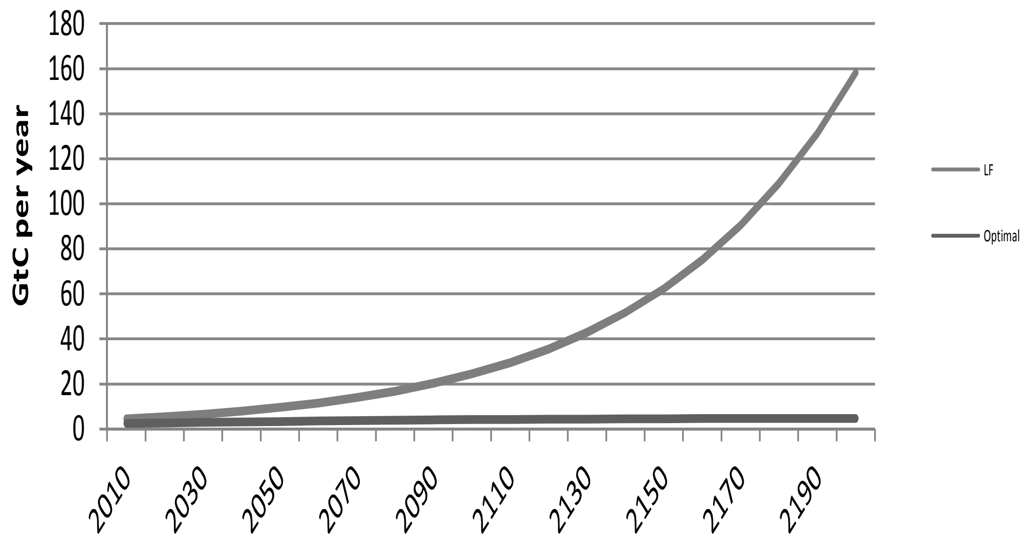

Figure 16: Coal use: optimum vs. laissez faire

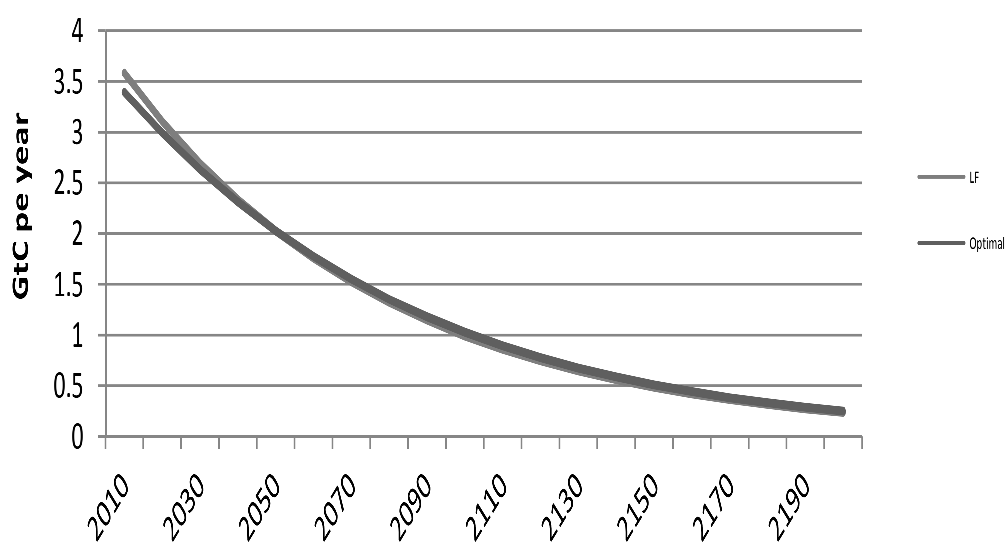

Figure 17: Oil use: optimum vs. laissez faire

Total damages are shown in Figure 18 below. We note large, though not gigantic, gains from moving from laissez faire to the optimum allocation. The gains grow over time, with damages at a couple of percent of GDP in the laissez-faire allocation, thus about double its optimal value at that time. In 2200, the difference is a factor of six.

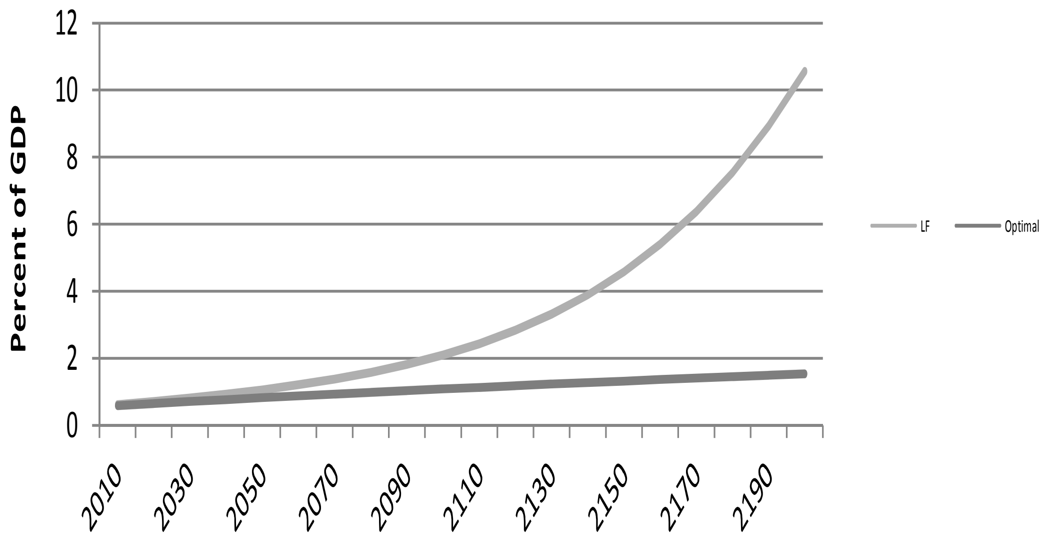

Figure 18: Total damages as a percent of global GDP: optimum vs. laissez faire

We can also back out the path for global temperature in the two scenarios, using the known mapping from S to temperature. Figure 19 illustrates that laissez faire is associated with a temperature rise of 4.4 degrees Celsius a hundred years from now; in the optimum, heating is only 2.6 degrees. Toward the end of the simulation period, however, due to massive coal use, laissez faire predicts increased heating by almost 10 degrees Celsius; the optimum dictates about 3 degrees.

Finally, Figure 20 displays the evolution of the (net-of-damage) production of final-good output (GDP). The intertemporal trade-off is clear here, but not as striking as one might have guessed: the optimal allocation involves rather limited short-run losses in GDP, with optimal output exceeding that of laissez faire as early as 2020. 100 years later, GDP net of damages is 2.5% higher in the optimum and in year 2200, it is higher by almost 15%.

5.2.6 Discussion

How robust are the quantitative results in Section 5.2.4? First, the tax formula appears remarkably robust. The point that only three kinds of parameters show up in the formula is a robustness measure in itself; e.g., no details of the fossil-fuel stocks, production technologies, or population matter. Strictly speaking, these features begin mattering once one or more of the main assumptions behind the formula are not met, but they will only matter indirectly, e.g., insofar as they influence the consumption-output path, and if their impact here is minor, the formula will be robust. In a technical appendix to the Golosov et al. (2014)

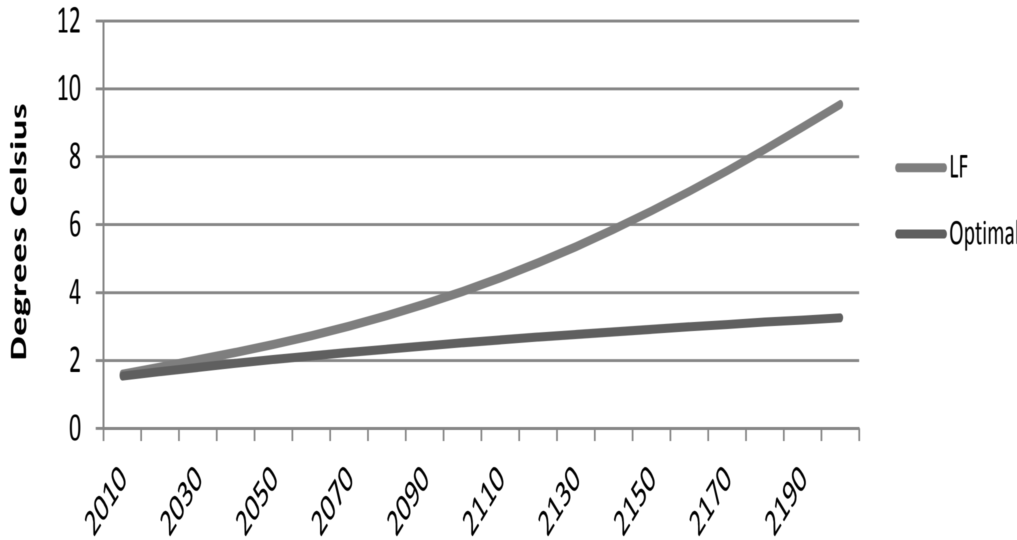

Figure 19: Increases in global temperature: optimum vs. laissez faire

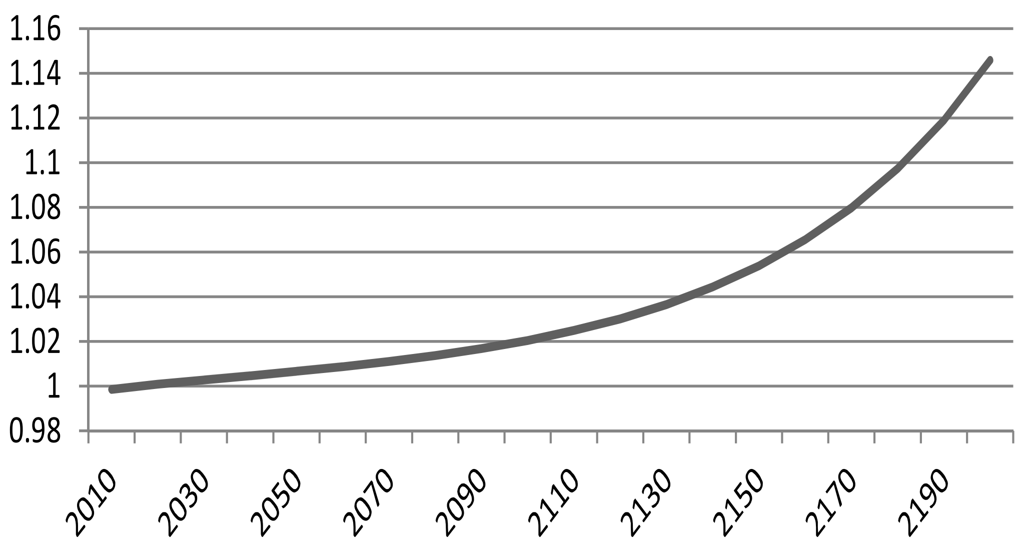

Figure 20: Net output: optimum vs. laissez faire

paper, Barrage (2014) considers a version of the model where not all of the assumptions are met. In particular, this version of the model has more standard transitional dynamics (with a calibration in line with the macroeconomic literature). For example, the assumption that the consumption-output ratio is constant will not hold exactly along a transition path, but the departures almost do not change the results at all. Also, at least U.S. data show very minor fluctuations in this ratio so to the extent a model delivers more drastic movements in the consumption-output ratio it will have trouble matching the data. Higher curvature in utility also delivers very minor changes in the tax rate, with the correction that discounting now involves not just β but also the consumption growth rate raised to $1 - \sigma$, where $\sigma = 1$ gives logarithmic curvature and $\sigma > 1$ higher curvature.

Second, when it comes to the positive analysis—e.g., the implications for temperature and damages under different policy scenarios—the message is quite different: many of the assumptions can matter greatly for the quantitative results. Perhaps the best example of non-robustness is the example considered in Golosov et al. (2014): the elasticity of substitution between energy sources was raised by setting $\rho = 0.5$, i.e., assuming an elasticity of 2 instead of one slightly below one. If the different energy sources are highly substitutable, coal can easily be used instead of oil, making the laissez-faire allocation deliver very high coal use. On the other hand, taxes are now more powerful in affecting the use of different energy sources. This means, in particular, that the difference in outcomes between an optimal tax and laissez-faire is very large compared to the benchmark, where the different energy sources are less substitutable. Hence, the substitutability across energy sources is an example of an area where more work is needed. Relatedly, we expect that the modeling of technical change in this area—energy-saving, as in Section 2.3.3 above or making new energy resources available—will prove very important.

A number of straightforward extensions to the setting are also possible and, in part, they have been pursued by other researchers.[110] One is the inclusion of damages that involve growth effects; Dell et al. argue that such effects may be present.[111] It is easy to introduce such damages to the present setting by letting the TFP term read $e^{-\gamma_l S + \gamma_g S t}$, where γ_l regulates level effect of carbon concentration S, and γ_g the damages to the growth rate of output; the baseline model admits closed-form solution. As already pointed out, the baseline model can also accommodate time-inconsistent preferences rather easily.[112]

Finally, the discussion of dynamic integrated assessment models here is based entirely on the simple baseline model in Golosov et al. (2014) not because it is the only model of this sort, or even the most satisfactory one in some overall sense; rather, this model has been chosen, first, because it is the model with the closest links to standard macroeconomic settings (with forward-looking consumers, dynamic competitive equilibrium with taxes, and so on). Second, the baseline model in Golosov et al. admits highly tractable analysis (with closed-form solutions) and hence is very well suited for illustrations; moreover, for the optimal carbon tax it gives a very robust formula that is also quantitatively adequate. The model is also useful for positive analysis but here it is important to point out that many other approaches can offer more realistic settings and, at least from some perspectives, do a better job at prediction. It would require a long survey to review the literature and such an endeavor

[110] E.g., Rezai and van der Ploeg (2014).

[111] See Moyer et al. (2013).

[112] Such cases have been discussed by Karp (2005) and, in settings closely related to the model here, Gerlagh and Liski (2012) and Iverson (2014) show that it is possible to analyze the case without commitment relatively straightforwardly; lack of commitment and Markov-perfect equilibria are otherwise quite difficult to characterize.

is best left for another paper; perhaps the closest relative among ambitious, quantitative settings is the WITCH model, which also builds on forward-looking and, among other things, has a much more ambitiously specified energy sector.[113]

[113] See Bosetti et al. (2006).

References

Acemoglu, Daron (2009), *Introduction to Modern Economic Growth*, Princeton, NJ: Princeton University Press.

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

Acemoglu, Daron, Philippe Aghion, and David Hémous (2014), "The Environment and Directed Technical Change in a North-South Model", *Oxford Review of Economic Policy*, 30 (3), 513–530.

Aghion, Philippe and Peter Howitt (1992), "A Model of Growth Through Creative Destruction", *Econometrica*, Vol. 60, No. 2, pp. 323-351

Aghion, Philippe and Peter Howitt (2008), *The Economics of Growth*, Cambridge, MA: MIT Press.

Aghion, Philippe, Antoine Deschezlepretre, David Hémous, Ralf Martin, and John Van Reenen (2014), "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*, Vol. 109, No. 3, pp. 659–684.

Archer, D. (2005), "The Fate of Fossil Fuel CO_2 in Geologic Time", *Journal of Geophysical Research*, 110.

Archer, D., M. Eby, V. Brovkin, A. Ridgwell, L. Cao, U. Mikolajewicz, and K. Tokos (2009), "Atmospheric Lifetime of Fossil Fuel Carbon Dioxide", *Annual Review of Earth and Planetary Sciences*, 37, 117–134.

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.

Amador, Manuel (2003), "A political model of sovereign debt repayment", mimeo.

Azzimonti, Marina (2011), "Barriers to Investment in Polarized Societies", *American Economic Review*, 101(5): 2182–2204.

Bansal, Ravi and Marcelo Ochoa (2011), "Welfare Costs of Long-Run Temperature Shifts", NBER Working Paper 17574.

Barrage, Lint (2014), "Optimal Dynamic Carbon Taxes in a Climate-Economy Model with Distortionary Fiscal Policy", working paper.

Barrage, Lint (2014), "Sensitivity Analysis for Golosov, Hassler, Krusell, and Tsyvinski (2014): 'Optimal Taxes on Fossil Fuel in General Equilibrium'", *Econometrica Supplemental Material*, 82, `http://www.econometricsociety.o`

Barro, R.J. (2013), "Environmental protection, rare disasters, and discount rates", NBER Working Paper 19258.

Berndt, Ernst R. and Laurits R. Christensen (1973), "The Translog Function and the Substitution of Equipment, Structures, and Labor in U.S. Manufacturing 1929–68", *Journal of Econometrics* 1 (1): 81–113.

Bosetti, Valentina, Carlo Carraro, Marzio Galeotti, Emanuele Massetti, and Massimo Tavoni (2006), "WITCH—A World Induced Technical Change Hybrid model", *Energy Journal* 27, 13–37.

Bovenberg, Lans and Sjak Smulders (1995), "Environmental Quality and Pollution-Augmenting Technological Change in a Two-Sector Endogenous Growth Model", *Journal of Public Economics*, 57.

BP (2010, 2015), in *BP Statistical Review of World Energy*, `http://bp.com/statisticalreview`.

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, Edward Miguel, Shanker Satyanath, John A. Dykema and David B. Lobell (2009), "Warming Increases the Risk of Civil War in Africa", *Proceedings of the National Academy of Sciences*, 106:49, 20670–20674.

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

Ciscar, Juan-Carlos, Ana Iglesias, Luc Feyen, László Szabó, Denise Van Regemorter, Bas Amelung, Robert Nicholls, Paul Watkiss, Ole B. Christensen, Rutger Dankers, Luis Garrote, Clare M. Goodess, Alistair Hunt, Alvaro Moreno, Julie Richards, and Antonio Soria (2011), "Physical and economic consequences of climate change in Europe", *Proceedings of the National Academy of Sciences*, 108:7, 2678–2683.

Cline, William R. (1992), *Economics of Global Warming*, Institute for International Economics, Washington, D.C.

Crost, B. and C.P. Traeger (2014), "Optimal CO_2 Mitigation under Damage Risk Valuation", *Nature Climate Change* 4: 631–636.

Cuddington, John and Grant Nülle (2014), "Variable long-term trends in mineral prices: The ongoing tug-of-war between exploration, depletion, and technological change", *Journal of International Money and Finance*, vol. 42, issue C, 224–252. Dasgupta, Partha and Geoffrey Heal (1974), "The Optimal Depletion of Exhaustable Resources", *Review of Economic Studies*, 41.

Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken (2012), "Temperature Shocks and Economic Growth: Evidence from the Last Half Century", *American Economic Journal: Macroeconomics*, 4:3, 66–95.

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.

Drijfhouta, Sybren, Sebastian Bathiany, Claudie Beaulieu, Victor Brovkin, Martin Claussen, Chris Huntingford, Marten Scheffer, Giovanni Sgubin, and Didier Swingedouw (2015), "Catalogue of abrupt shifts in Intergovernmental Panel on Climate Change climate models", *Proceedings of the National Academy of Sciences*, vol. 112, no. 43.

Ellerman, A. Denny and Barbara K. Buchner (2007), "The European Union Emissions Trading Scheme: Origins, Allocation, and Early Results", *Review of Environmental Economics and Policy*, 1:1, 66–87.

Erten, Bilge and José Antonion Ocampo (2012), "Super-cycles of commodity prices since the mid-nineteenth century", DESA Working Paper No. 110.

Feng, Shuaizhang, Alan B. Krueger, and Michael Oppenheimer, "Linkages among climate change, crop yields and Mexico-US cross-border migration," *PNAS*, 107 (32).

Fankhauser, Samuel (1994), "The economic costs of global warming damage: A survey", *Global Environmental Change,* 4, 301–309.

Gars, Johan (2012), *Essays on the Macroeconomics of Climate Change*, Ph.D. Thesis, IIES Monograph series No. 74, Institute for International Economic Studies, Stockholm University.

Gerlagh, Reyer and Matti Liski (2012), "Carbon Prices for the Next Thousand Years", CESifo Working Paper Series No. 3855.

Geweke, John (2001), "A note on some limitations of CRRA utility", *Economic Letters*, 71, 341–345.

Gollier, Christian (2013), "Evaluation of long-dated investments under uncertain growth trend, volatility and catastrophes", Toulouse School of Economics TSE Working Papers 12-361.

Golosov, Michael, John Hassler, Per Krusell, and Aleh Tsyvinski (2014), "Optimal Taxes on Fossil Fuel in Equilibrium", *Econometrica*, Vol. 82, No. 1, 41–88.

Grossman, Gene and Alan Krueger (1991), "Environmental impacts of a North American Free Trade Agreement", NBER Working Paper 3914.

Hart, Rob (2013), "Directed technical change and factor shares", *Economics Letters*, Volume 119, Issue 1, 77–80.

Harvey, David, Neil Kellard, Jakob Madsen, and Mark Wohar (2010), "The Prebisch-Singer Hypothesis: Four Centuries of Evidence", *The Review of Economics and Statistics*, Vol. 92, No. 2, 367–377.

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 and Per Krusell (2014), "The Climate and the Economy", Mistra-SWECIA Nr. 5.

Hassler, John, Per Krusell, and Jonas Nycander (2016), "Climate Policy", forthcoming in *Economic Policy.*

Hassler, John, Per Krusell, and Conny Olovsson (2012), "Energy-Saving Technical Change", NBER Working Paper No. 18456.

Hassler, John, Per Krusell, and Conny Olovsson (2015), "Will we need another mad max? Or will energy-saving technical change save us?", working paper.

Hémous, David (2013), "Environmental Policy and Directed Technical Change in a Global Economy: The Dynamic Impact of Unilateral Environmental Policies?", working paper.

Hillebrand, Elmar and Marten Hillebrand (2016), "Optimal Climate Policies in a Dynamic Multi-Country Equilibrium Model", working paper.

Hope Chris, John Anderson, and Paul Wenman (1993), "Policy analysis of the greenhouse effect: An application of the PAGE model", *Energy Policy,* 21, 327–338.

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

IEA (International Energy Agency), (2010), "World Energy Outlook", OECD/IEA Paris.

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

IPCC (Intergovernmental Panel on Climate Change) (2006), Guidelines for National Greenhouse Gas Inventories, Vol 2 Energy.

IPCC (2006), "2006 IPCC Guidelines for National Greenhouse Gas Inventories", Volume 2: Energy, Task Force on National Greenhouse Gas Inventories, Intergovernmental Panel on Climate Change.

IPCC (Intergovernmental Panel on Climate Change) (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 (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.

Iverson, Terry (2014), "Optimal Carbon Taxes with Non-Constant Time Preference", working paper.

Jia, Ruixue (2014), "Weather Shocks, Sweet Potatoes and Peasant Revolts in Historical China", *Economic Journal*, 124:575, 92–118.

Jones, Charles I. (2001), *Introduction to Economic Growth*, New York, NY: W.W. Norton.

Jorgenson, Dale W., Richard J. Goettle, Mun S. Ho, and Peter J. Wilcoxen (2013a), *Double Dividend*, Cambridge, MA: MIT Press.

Jorgenson, Dale W., Jing Cao, and Mun S. Ho (2013b), "The Economics of Environmental Policies in China", in Chris P. Nielson and Mun S. Ho (eds.), *Clearer Skies Over China*, Cambridge, MA: MIT Press.

KamLAND Collaboration (2011), "Partial radiogenic heat model for Earth revealed by geoneutrino measurements", *Nature Geoscience*, Vol 4:9, September, 647–651.

Karp, Larry (2005), "Global warming and hyperbolic discounting", *Journal of Public Economics*, 89(2): 261–282.

Kelly, David and Charles Kolstad (1999), "Bayesian Learning, Growth, and Pollution", *Journal of Economic Dynamics and Control*, 23:4, 491–518.

Krautkraemer, Jeffrey A. (1998), "Nonrenewable Resource Scarcity", *Journal of Economic Literature*, Vol. 36, No. 4, 2065–2107.

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 (2014), "A Global Economy-Climate Model with High Regional Resolution", working paper.

Krusell, Per and Anthony Smith (2015), "Climate Change Around the World", working paper.

La Ferrara, Eliana, and Mariaflavia Harari (2014), "Conflict, Climate and Cells: A disaggregated analysis, IGIER, Working Paper.

Leach, Andrew (2007), "The Welfare Implications of Climate Change Policy", *Journal of Economic Dynamics and Control*, 57.

Lemoine, Derek (2015), "The climate risk premium: how uncertainty affects the social cost of carbon", working paper.

Lenton, Timothy M., Hermann Held, Elmar Kriegler, Jim W. Hall, Wolfgang Lucht, Stefan Rahmstorf and Hans Joachim Schnellnhuber, (2008), "Tipping elements in the Earth's climate system", *Proceedings of the National Academy of Sciences*, vol. 105, no. 6, 1786–1793.

Levitan, Dave (2013), "Quick-Change Planet: Do Global Climate Tipping Points Exist?", *Scientific American* 25.

Lewis, N., Marotzke, J., Stevens, B., Myhre, G., and Shindell, D. (2013), "Energy budget constraints on climate response", *Nature Geoscience*, 6 (6), 415–416.

Lucas, Robert E., Jr. (1988), "On the Mechanics of Economic Development", *Journal of Monetary Economics*, 3–42.

Manne, Alan, Robert Mendelsohn, and Richard Richels (1995), "MERGE: a model for evaluating regional and global effects of GHG reduction policies", *Energy Policy*, 23:1, 17–34.

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

Matthews, H. Damon, Susan Solomon, and Raymond Pierrehumbert (2012), "Cumulative carbon as a policy framework for achieving climate stabilization", *Philosophical Transactions A, Royal Society*, 370, 4365–4379.

McGlade, Christophe and Paul Ekins (2015), "The geographical distribution of fossil fuel unused when limiting global warming to 2 C", *Nature*, 517.

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.

Miguel, Edward, Shanker Satyanath, and Ernest Sergenti (2004), "Economic Shocks and Civil Conflict: An Instrumental Variables Approach", *Journal of Political Economy*, 112:4, 725-753.

Moyer, Elisabeth J., Mark D. Woolley, Michael J. Glotter, and David A. Weisbach (2013), "Climate Impacts on Economic Growth as Drivers of Uncertainty in the Social Cost of Carbon", working paper.

Nordhaus, William D. (1973), "World Dynamics: Measurement Without Data", *Economic Journal*, Volume 83, No. 332.

Nordhaus, William D. (1974), "Resources as a Constraint on Growth", *American Economic Review*, Volume 64, No. 2.

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 D. (1991), "Economic approaches to greenhouse warming", in *Global warming: Economic policy approaches*, ed. R.D. Dornbush and J.M. Poterba, 33–68. Cambridge, MA: MIT Press.

Nordhaus, William D. (1992), "An Optimal Transition Path for Controlling Greenhouse Gases," *Science*, vol. 258, 1315–1319.

Nordhaus, William D. (1993), " "Rolling the 'DICE': An Optimal Transition Path for Controlling Greenhouse Gases," *Resource and Energy Economics*, vol. 15, 27–50.

Nordhaus, William D. (2006), "Geography and macroeconomics: New data and new findings", *Proceedings of the National Academy of Sciences*, vol. 103, no. 10, 3510–3517

Nordhaus, William D. (2007), "To Tax or Not to Tax: The Case for a Carbon Tax", *Review of Environmental Economics and Policy.*

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

Nordhaus, William D. (2009), "An analysis of the Dismal Theorem", working paper.

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

Nordhaus, William D. and Paul Sztorc (2013), "DICE 2013R: Introduction and User's Manual", mimeo, Yale University.

Otto, A., Otto, F.E.L., Allen, M.R., Boucher, O., Church, J., Hegerl, G., Forster, P.M., Gillett, N.P., Gregory, J., Johnson, G.C., Knutti, R., Lohmann, U., Lewis, N., Marotzke, J., Stevens, B., Myhre, G., and Shindell, D. (2013), "Energy budget constraints on climate response", *Nature Geoscience*, 6 (6), 415–416.

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

Pindyck, Robert S. (1978), "The Optimal Exploration and Production of Nonrenewable Resources", *Journal of Political Economy*, Vol. 86, No. 5, 841–861.

Poole, William (1970), "Optimal Choice of Policy Instruments in a Simple Stochastic Macro Model, *Quarterly Journal of Economics*, 84:2, 197–216.

Prather, Michael J., Christopher D. Holmes and Juno Hsu (2012), "Reactive greenhouse gas scenarios: Systematic exploration of uncertainties and the role of atmospheric chemistry, *Geophysical Research Letters*, 39:9.

Prebisch, Raul (1950), "The economic development of Latin America and its possible problems", reprinted in *Economic Bulletin for Latin America* 7 (1), 1–22, 1962.

Revelle, Roger and Hans Suess (1957), "Carbon Dioxide Exchange between Atmosphere and Ocean and the Question of an Increase of Atmospheric CO_2 during Past Decades", *Tellus*, 9, 18–27.

Rezai, Armon and Frederick van der Ploeg (2014), "Robustness of a simple rule for the social cost of carbon", *Economics Letters*, 132, 48–55.

Rockström, Johan, Will Steffen, Kevin Noone, sa Persson, F. Stuart Chapin, Eric F. Lambin, Timothy M. Lenton, Marten Scheffer, Carl Folke, Hans Joachim Schellnhuber, Björn Nykvist, Cynthia A. de Wit, Terry Hughes, Sander van der Leeuw, Henning Rodhe, Sverker Sörlin, Peter K. Snyder, Robert Costanza, Uno Svedin, Malin Falkenmark, Louise Karlberg, Robert W. Corell, Victoria J. Fabry, James Hansen, Brian Walker, Diana Liverman, Katherine Richardson, Paul Crutzen, and Jonathan A. Foley (2009), "A safe operating space for humanity", *Nature* 461, 472–475.

Roe, Gerard H. and Marcia B. Baker (2007), "Why Is Climate Sensitivity So Unpredictable?", *Science,* 26, Vol. 318, no. 5850, 629–632.

Rogner, Hans-Holger (1997), "An Assessment of World Hydrocarbon Resources", *Annual Review of Energy and the Environment*, 22, 217–262.

Romer, Paul M. (1990), "Endogenous Technological Change", *Journal of Political Economy*, vol. 98, no. 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.

Schwartz, Stephen, Robert Charlson, Ralph Kahn, John Ogren, and Henning Rodhe (2010), "Why Hasn't Earth Warmed as Much as Expected?", *Journal of Climate,* 23.

Schwartz, S. E., R. J. Charlson, R. Kahn, and H. Rodhe (2014), "Earth's Climate Sensitivity: Apparent Inconsistencies in Recent Assessments", *Earth's Future*, 2.

Sinn, Hans-Werner (2008), "Public Policies against Global Warming: A Supply Side Approach", *International Tax and Public Finance*, 15:4, 360–394.

Solow, Robert (1974), "Intergenerational Equity and Exhaustible Resources", *Review of Economic Studies*, 41, 29–45.

Singer, H.W. (1950), "U.S. foreign investment in underdeveloped areas: The distribution of gains between investing and borrowing countries", *American Economic Review Papers and Proceedings*, 40, 473–485. Spiro, Daniel (2014), "Resource prices and planning horizons", *Journal of Economic Dynamics and Control*, Vol. 48, 159–175.

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

Stern, Nicholas (2007), *The Economics of Climate Change: The Stern Review*, Cambridge and New York: Cambridge University Press.

Sterner, Tomas and Martin Persson (2008), "An Even Sterner Review: Introducing Relative Prices into the Discounting Debate", *Review Environmental Economics and Policy*, 2:1, 61–76.

Stiglitz, Joseph (1974), "Growth with Exhaustible Natural Resources: Efficient and Optimal Growth Paths", *Review of Economic Studies*, 41, 123–137.

Stokey, Nancy (1998), "Are There Limits to Growth?", *International Economic Review*, 39, 1–31.

Titus, James G. (1992), "The costs of climate change to the United States", In: *Global Climate Change: Implications, Challenges and Mitigation Measures,* S.K. Majumdar, (ed.), Pennsylvania Academy of Science, Easton, Pennsylvania, 384–409.

Tol, Richard S.J. (1995), "The damage costs of climate change toward more comprehensive calculations", *Environmental and Resource Economics,* 5, 353–374.

Tol, Richard S.J. (2009), "The Economic Effects of Climate Change", *Journal of Economic Perspectives,* 23:2, 29–51.

Weitzman, Martin L. (1974), "Prices vs. Quantitites", *Review of Economic Studies*, 41(4), 477–491.

Weitzman, Martin L. (1998), "Why the far-distant future should be discounted at its lowest possible rate", *Journal of Environmental Economics and Management*, 36(3): 201–208.

Weitzman, Martin L. (2009), "On Modeling and Interpreting the Economics of Catastrophic Climate Change", *Review of Economics and Statistics*, 91, 1–19.

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