

NBER WORKING PAPER SERIES

OPTIMAL TAXES ON FOSSIL FUEL IN GENERAL EQUILIBRIUM

Mikhail Golosov
John Hassler
Per Krusell
Aleh Tsyvinski

Working Paper 17348
<http://www.nber.org/papers/w17348>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2011

We thank Bill Nordhaus, Jonas Nycander, Tony Smith, Sjak Smulders, and seminar participants at ESSIM, EIEF, EUI, IIES, Mistra-SWECIA, Yale, UCL, CREI, the Environmental Macro Conference at ASU, the EEA Annual Meeting (Glasgow), Fudan University (Shanghai), the Chinese University of Hong Kong, Bonn, Zurich, Carlos III, and REDg DGEM (Barcelona). Golosov and Tsyvinski thank NSF for support and EIEF for their hospitality; Krusell thanks ERC and Mistra-SWECIA for support and Hassler thanks Mistra-SWECIA and the Swedish Research Council for support. Tsyvinski acknowledges relationships with the National Science Foundation, Einaudi Institute of Economics and Finance, Guggenheim Memorial Fellowship. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2011 by Mikhail Golosov, John Hassler, Per Krusell, and Aleh Tsyvinski. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Optimal Taxes on Fossil Fuel in General Equilibrium
Mikhail Golosov, John Hassler, Per Krusell, and Aleh Tsyvinski
NBER Working Paper No. 17348
August 2011
JEL No. E13,E17,E6,H23,Q28,Q4,Q5

ABSTRACT

We analyze a dynamic stochastic general-equilibrium (DSGE) model with an externality through climate change from using fossil energy. A central result of our paper is an analytical derivation of a simple formula for the marginal externality damage of emissions. This formula, which holds under quite plausible assumptions, reveals that the damage is proportional to current GDP, with the proportion depending only on three factors: (i) discounting, (ii) the expected damage elasticity (how many percent of the output flow is lost from an extra unit of carbon in the atmosphere), and (iii) the structure of carbon depreciation in the atmosphere. Very importantly, future values of output, consumption, and the atmospheric CO₂ concentration, as well as the paths of technology and population, and so on, all disappear from the formula. The optimal tax, using a standard Pigou argument, is then equal to this marginal externality. The simplicity of the formula allows the optimal tax to be easily parameterized and computed. Based on parameter estimates that rely on updated natural-science studies, we find that the optimal tax should be a bit higher than the median, or most well-known, estimates in the literature. We also show how the optimal taxes depend on the expectations and the possible resolution of the uncertainty regarding future damages. Finally, we compute the optimal and market paths for the use of energy and the corresponding climate change.

Mikhail Golosov
Department of Economics
Yale University
Box 208268
New Haven, CT 06520-8269
and NBER
golosov@princeton.edu

Per Krusell
Institute for International Economic Studies
Stockholm University
106 91 STOCKHOLM
SWEDEN
and NBER
per.krusell@iies.su.se

John Hassler
Institute for International Economic Studies
Stockholm University
John.Hassler@iies.su.se

Aleh Tsyvinski
Department of Economics
Yale University
Box 208268
New Haven, CT 06520-8268
and NBER
a.tsyvinski@yale.edu

1 Introduction

In this paper we specify and solve a global economy-climate model—a so-called integrated assessment model (IAM)—based on a stochastic dynamic general-equilibrium (DSGE) approach. The main goal of the paper is to use this model to characterize optimal global policy qualitatively and quantitatively. The background for the work and for our particular approach is that there now is widespread consensus that human activity is an important driver of climate change (see, e.g., IPCC report, 2007). However, there are great remaining uncertainties, especially regarding how much global warming to expect, and how much damage the warming will cause. We take the climate problem as given here and use conventional estimates of how carbon emission drives global warming.

The effects of global warming and of policies to correct it are uncertain, are dynamic, and last for a long time. They also have potentially large general-equilibrium effects, and may crucially depend on technological change. These are the primary reasons why we use a standard DSGE model of the economy augmented with the damages coming from burning fossil fuel. The microeconomic foundations we use also allow welfare evaluation that obeys standard economic principles: given any path of policy, we can compute the effects on human welfare using the utility functions humans themselves use.

We first solve the social planner’s problem in our economy with the climate externality, which is entirely global in nature. The key finding is an analytical characterization and derivation of a simple formula for the marginal externality damage of carbon dioxide emissions. This social cost, expressed as a proportion of GDP, is a very simple function of the basic parameters of the model. At any given point in time, it is equal to the expected discounted value of the sum of future damages, weighted by carbon depreciation rates: at time $t + j$ only a fraction $1 - d_j$ of a unit emitted at t remains in the atmosphere (and, hence, affects climate). We approximate, on quite solid grounds, these rates to be constant (i.e., independent of the total stock of carbon in the atmosphere), but non-geometric: they feature (i) rapid initial depreciation and (ii) slow subsequent depreciation, though (iii) with an appreciable fraction of the initial unit remaining in the atmosphere forever. This formula is very convenient as it shows that only discounting, damages, and carbon depreciation parameters matter for the externality costs of emissions. More specifically, future values of output, consumption, and the stock of CO₂ in the atmosphere all disappear from the formula; and no knowledge about future technology, productivity, energy sources in use, or population is needed in order to calculate the social cost!

The formula we offer does require some assumptions, but we argue that these assumptions are quite reasonable. The intuition for the result also follows rather straightforwardly from understanding them. More specifically, we make four assumptions to derive our result: (i) period utility is logarithmic in consumption; (ii) current climate damages are proportional to output and are a function of the current atmospheric carbon concentration with a constant elasticity (a relationship that is allowed to vary over time/be random); (iii) the stock of carbon in the atmosphere is linear in the past and current energy use (as mentioned above); and (iv) the saving rate is constant. We discuss these assumptions in detail in the paper and, in particular, discuss the robustness of the results to alternative assumptions (such as

so-called threshold effects)—see Section 5.2.3.¹ Under our maintained assumptions, the level of output drops out of the formula as damages are proportional to output and the marginal utility, at which these damages are evaluated in the future, is inversely proportional to output. Quite importantly, our formula applies also when an alternative (clean) energy technology is available, as well as when there is deliberate investment into developing such a technology—when technology choice is endogenous.

The three key factors in our formula—damages, discounting, and depreciation—affect the social cost/optimal tax as follows. Higher expected damages raise the social cost (as a proportion of GDP). Our formula also reveals that it is the expected damage elasticity that enters the calculation: a certainty-equivalence result in terms of this elasticity, even though the utility function features risk aversion. A higher discount rate—we use a dynastic population structure—lowers it. The carbon-cycle parameters influence the optimal tax in the intuitively way as well. The longer CO₂ stays in the atmosphere, the higher is the optimal tax rate. As noted, technology growth does not matter for the social cost per GDP unit. This is because there are two opposing effects which cancel exactly under our specific assumptions: higher future technology increases damages, which are proportional to future TFP, but these damages are weighted by marginal utility, which is inversely proportional to future TFP.

We then characterize competitive market outcomes with and without taxes. As our model falls within the framework where Pigou’s general insights apply, we can easily construct a sequence of taxes implementing the fully optimal allocation. The optimal per-unit tax on emissions should, hence, be set to equal to the marginal externality cost of emissions.² The tax, as a proportion of output, is then given by the simple formula described above. Since stochastics are allowed in our analysis, and since a form of certainty equivalence applies, as discussed above, it is straightforward to see how future updates about the social cost will cause the tax to vary. A simple, and reasonable, case is that where no transition dynamics are expected for the damage elasticity parameter: the expected time path for the damage parameter is constant, but upward or downward adjustments can occur. We calibrate using two possible future outcomes—damages are high, or they are not so high—and the assumption that it takes time to find out whether they are high or low. Suppose, then, that society learns at some future date that the value of the damage elasticity is high (and will remain high). Then from that point on the marginal externality cost of emissions has gone up, as must the optimal tax rate. This outcome embodies the idea that there is an option value of waiting: not until we (potentially) find out that damages will be very high is it necessary to use very high tax rates, and before then only the expected value of the damage elasticity is relevant for the level of the tax.

The first part of the quantitative analysis is calculating the optimal tax using the simple formula we derived. As these taxes depend only on the discount factor, the damage elasticity, and the parameters of the carbon cycle and are independent of the other parameters of the

¹Our model is non-linear and none of the results rely on linearization techniques.

²It also follows that one should subsidize the removal of carbon dioxide from the atmosphere, if technologically feasible, at the same unit rate.

model, it is easy to evaluate them quantitatively. For the discount factor, we simply report results for different levels of discounting. For the damage elasticity, we follow Nordhaus and Boyer’s (2000) estimates entirely, including the uncertainty expressed there. They use a “bottom-up” approach, adding dollar measures of a large variety of effects of droughts, floods, and storms caused by climate change, along with changes in biodiversity, and so on. For the carbon-cycle (depreciation) parameters, we rely on estimates from Archer (2005). We calculate the optimal taxes both before and after we have learnt the long-run value of the damages. We express the tax per ton of emitted carbon at a yearly global output of 70 trillion dollars. It is useful to compare our quantitative results to two important and very influential policy proposals in the literature: that in Nordhaus (2000, 2007) and that in the Stern report (2007).³ These proposals amount to a tax of \$30 and \$250 dollars per ton of coal, respectively. A key difference between the two proposals is that they use very different (pure) subjective discount rates. Nordhaus uses a rate of 1.5% per year; he motivates this level mostly based on market measures. Stern, who adds an additional “moral” concern for future generations, uses the much lower rate of 0.1% per year. For these two values of the discount rate, the optimal taxes using our analysis are \$56.9/ton and \$496/ton, respectively. Thus, our calculations suggest significantly larger optimal taxes than do each of these studies. It is comforting that our damage estimate can be made consistent, quantitatively, with that computed by Nordhaus, even though his model is different in various ways—notably in his assumptions on energy supply, adaptation, and technology—and requires numerical solution. Nordhaus’s carbon cycle implies less carbon depreciation from the atmosphere in the short run and that all emitted carbon eventually ends up in the deep oceans. Adapting our carbon depreciation structure to replicate these features in Nordhaus’s model (as well as approximating his assumptions of a lag in the effects of a change in atmospheric carbon concentration on temperature and a higher utility-function curvature) yields an optimal tax rate close to the one he proposed; we discuss this comparison in detail in Section 5.2.3 of the paper. Our depreciation estimates, we believe, are entirely in line with updated natural science research, as represented by the 2007 IPCC report and, more specifically, Archer (2005): about 20% of the emitted carbon stays “forever” (i.e., for thousands of years), even though a significant fraction moves from the atmosphere to the biosphere and oceans relatively quickly.

Moreover, the consequences of learning can be dramatic. With a discount rate of 1.5%, the optimal tax rate if damages turns out to be moderate is \$25.3/ton but \$489/ton if they are what Nordhaus refers to as “catastrophic”. For the lower discount rate used by Stern, the corresponding values are \$221/ton and a whopping \$4,263/ton.

The second part of the quantitative analysis derives the energy use and the paths of the climate variables such as the global increase in the temperature. This part, unlike the calculations of the optimal taxes, requires additional model assumptions. It is, however, very easy to solve the model; it is “almost” possible to solve in closed form.⁴ We assume that

³Nordhaus’s estimate, moreover, is near the mode of a distribution of estimates which Tol (2008) derives in an effort to summarize over 200 estimates in the literature.

⁴The simplicity is due to the fact that the law of motion for fossil fuel contains no other endogenous

there are two sources of energy—oil and coal—and that these are perfect substitutes. Oil is more efficient than coal, extracted at zero cost, and is in finite supply. Coal is extracted using labor (at constant marginal costs) and is in infinite supply.⁵ This structure does deliver constancy of saving rates and, thus, is consistent with our optimal-tax formula. Our results are as follows. The no-tax market economy would empty our oil supplies in three decades, whereas optimal taxation would let oil be used for five decades. Coal use (setting in after the oil has run out) is twice as large in the market outcome as in the optimal outcome. That is, significant quantity adjustments are needed to correct for the externality and reach optimal allocations. We predict a market economy with a significant temperature increase around 2045, a rise that will peak at almost 7 degrees Celsius a little over 100 years from now. The optimal allocation dictates a significant delay in warming, by about two decades, and a peak of a little below 3 degrees Celsius. These predictions apply until we (possibly) learn that the damage elasticity is either higher or lower than what has been assessed so far; if updates arrive, these numbers have to be revised accordingly.

Our work is related to several literatures. The pioneering work in this area is due to William Nordhaus (for a recent description of his modeling, see Nordhaus and Boyer, 2000). In almost every way, the spirit of our modeling is entirely in line with the approach used by Nordhaus. His main framework is a computational model called RICE—Regional dynamic Integrated model of Climate and the Economy—or, in its earlier one-region version, DICE. Nordhaus’s work has been particularly pioneering in two areas. First, informed by modern climate and carbon-cycle modeling, he managed to summarize the key quantitative channels from the economy to the climate with a rather parsimonious, and mostly linear, dynamic system. This dynamic system is small enough that it can be embedded in a typical dynamic growth model. Second, Nordhaus did extensive work aimed at summarizing the damages from climate change. His modeling of these damages in RICE amounts to a multiplicative term on aggregate production which is a function of the average global temperature. That is, damages of various sorts are expressed as a production loss measure. The natural science part of Nordhaus’s model is close to the state-of-the-art specification used by climate scientists. Our natural-science model mostly follows that formulation; the key differences are in the depreciation structure of carbon emitted into the atmosphere, as discussed above, and in our assumption that the full temperature response to atmospheric carbon is immediate (Nordhaus uses slower temperature dynamics; see our detailed discussion in Section 5.2.3 below). The economic part of Nordhaus’s model is a bit less standard as, for example, it uses a finite horizon and a less than full microeconomic structure of energy supply. It is, therefore, somewhat difficult to compare it with standard macroeconomic models, and comparative statics and dynamics require full non-linear model solution and are therefore harder to interpret. Nordhaus is able to compute optimal tax rates, as those Pigou taxes that implement the first best, but cannot straightforwardly compare the optimal allocation to

variables.

⁵The assumption of infinite coal supply should be interpreted as follows: there is enough coal that, given that a cheap enough alternative technology will appear at some future date, coal will not command a Hotelling rent.

second-best alternatives, such as the market laissez-faire outcome or one with carbon taxes that are less than fully optimal. Our approach is firmly within the modern-macro tradition, thus relying on explicit microeconomic foundations both in terms of consumers and firms. A byproduct of being able to deal explicitly with a tax (which, in our setting, is equivalent to cap-and-trade policy), is that we do not need to include “abatement” in our analysis. Abatement is very often (e.g., in DICE) used as a stand-in variable for policy (taxes and other): the planner is assumed to be able to choose it, at some cost, in order to directly influence emissions. Finally, another difference is that we explicitly incorporate uncertainty, which can take a rich structure.⁶

The global economy-climate model that we construct in this paper can also be viewed as a natural extension of non-renewable resource models along the lines of Dasgupta and Heal (1974) and Stiglitz (1974), papers that appear to have been reactions to the oil shock in 1973 and focused mainly on conditions under which consumption could be sustainable. Compare to these papers, we include a climate externality, a carbon cycle, and consider technological change. Quite importantly, our model is also an extension in that we study a global competitive equilibrium with an externality, allowing us to discuss explicitly, with standard welfare analysis, how economic policy could and should be used to correct this externality. The prime purpose of the paper is indeed to characterize the optimal energy taxes in the global decentralized equilibrium economy.

Another difference that we obtain, relative to important parts of the literature and the policy discussions, regards the time path of taxes. As explained above, our estimate is that the per-unit tax on emissions should, as a fraction of GDP, be constant (unless new information about the parameters in the tax formula arrives). Interpreted in more standard terms, as a percentage value-added tax, it will fall over time, to the extent fossil-fuel prices are expected to grow faster than GDP—a feature of most available models, going back to Hotelling’s famous formula (Hotelling, 1931): it should rise at the rate of interest, which on average is above the rate of real GDP growth. Our prediction that taxes on value added, absent any future upward revision of damage estimates, should be declining over time is also a rather robust result.

Section 2 describes the model in generality as well as with more specialized assumptions. It also partially characterizes the solution to the planning problem; our proposition characterizing the marginal externality damage of emissions appears in Section 2.3. Section 3 then looks at a decentralized world economy, allowing for taxes on emissions, and demonstrates the optimal-tax result. In Section 4 we fully solve the model for two relevant special cases: that with oil only and that with coal only. Section 5 then contains our quantitative analysis (which relies on both oil and coal use). We discuss some obvious limitations of our work and conclude in Section 6.

⁶Nordhaus’s work generated much follow-up research; we comment on how our work relates to this research in the appropriate places in the text.

2 Optimal resource use

We begin by describing the general setting and the planning problem: how to optimally allocate resources over time, taking into account how the economy affects the climate. We first describe a rather general economy, in particular with regard to the carbon cycle and the different sources of fossil fuel and their extraction technologies. Later on, we specialize the analysis so as to arrive at transparent characterizations of the optimum. In the next section of the paper, we study the market solution and show how, using a simple tax/transfer scheme, the planning solution can be implemented.

2.1 The economy and the climate: a general specification

Time is discrete and infinite. The world is assumed to be inhabited by a representative household with the utility function

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

where the function U is a standard concave period utility function, C is consumption, and $\beta \in (0, 1)$ is the discount factor. The economy's constraint in the final-goods sector is

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

The left-hand side is resource use—consumption and next period's capital stock. The first term on the right-hand side, Y_t , is the output of the final good, possibly net of the resource cost of providing alternative sorts of energy. We do not introduce separate notation for such alternative energy sources. The production of fossil fuel is described explicitly below as taking place in a separate production sector. The second term, undepreciated capital, is entirely standard.⁷

Output in the final-goods sector can be described by an aggregate production function F_{1t} :

$$Y_t = F_{1t}(K_{1t}, N_{1t}, E_{1t}, S_t).$$

The arguments of F_1 include the standard inputs K_{1t} and N_{1t} (capital and labor in the final good sector), along with E_{1t} , which denotes fossil fuel energy in the final sector in period t . The sub-index t captures the possibility of technical change. Technical change can appear in a variety of ways here, e.g., an overall increase in productivity, a changed transformation technology across basic inputs (such as technical change saving on specific inputs), or a change in the “need” for fossil fuel (if a clean technology emerges). The possibility of endogenous technology is discussed in Section 6.

Finally, we also allow a climate variable S_t to affect output. This effect could in general be either positive or negative, and we use the word “damage” with this understanding. We focus on various sorts of damages that are all captured through the production function. We

⁷We assume zero adjustment costs for capital and that depreciation is geometric merely for simplicity.

specify later how F_{1t} depends on S . However, note that we view the climate to be sufficiently well represented by one variable only: S is to be read as the global concentration of carbon in the atmosphere, in excess of preindustrial levels. We argue that this is reasonable given available medium-complexity climate models used in the natural sciences. These imply that the current climate is quite well described by current carbon concentrations in the atmosphere (e.g., lags due to ocean heating are not so important).⁸ We do allow the damages, or the mapping from the atmospheric carbon concentration, to have a stochastic component but we suppress it here for notational convenience. We discuss the stochastic component in detail later in the paper. By making S_t an argument of production and not elsewhere—e.g., as a direct argument of utility or of the depreciation of capital—we are taking a stand on the nature of climate damages that is, arguably, more extreme than necessary. Aside from the notational advantage of damages only appearing in one place, there are two main reasons for this assumption: it allows the analysis to be simpler in certain important respects (see the discussion of the marginal damages in Section 4 below), and it makes our analysis closer to, and easier to compare with, Nordhaus’s RICE and DICE treatments.⁹

Turning to the origins of the fossil-fuel input, let the variable R_t denote the stock of remaining fossil fuel at the beginning of period t ($R_0 > 0$ is given). Letting $E_t \geq 0$ denote the amount of fossil fuel extracted, it follows a law of motion

$$R_{t+1} = R_t - E_t. \tag{1}$$

Fossil fuel can be extracted using technology F_{2t} represented as

$$E_t = F_{2t}(K_{2t}, N_{2t}, E_{2t}, R_t, R_{t+1}). \tag{2}$$

Here, K_{2t} , N_{2t} , and E_{2t} are the capital, labor, and energy inputs used in the extraction sector. The variables R_t and R_{t+1} are included as arguments in the extraction cost function in order to capture the costs of extracting fossil fuel. In general, the extraction costs depend on how much fossil fuel remains before and after extraction in a given period.¹⁰ Each period, the total available amount of each input is split costlessly across use in the two sectors:

$$K_{1t} + K_{2t} = K_t, \quad N_{1t} + N_{2t} = N_t, \quad E_{1t} + E_{2t} = E_t. \tag{3}$$

⁸Roe and Bauman (2011) analyze this and show that if the long-run sensitivity of the global mean temperature to the CO₂ concentration is high, it becomes important to take into account the temperature lags due to the fact that the ocean heats more slowly than the atmosphere. We discuss the implications of this mechanism in more detail below.

⁹As formulated above, the climate influences the economy by directly affecting output. This, of course, is a simplification in several ways, as the documented damages from climate change also include, among other factors, loss of life (which should appear through utility and makes labor input fall), deterioration in the quality of life (arguably also expressible with a more general utility function), and depreciation of the capital stock. How large these different damages are and exactly the form they take is highly uncertain. These damages should also include any resources used to prevent disasters and, more generally, to lessen the impact of climate change on humans and human activity (such as increased spending on air conditioning and on research aimed at adaptation and mitigation). The purpose of the present paper is not to push this particular frontier, and for this reason we choose to utilize Nordhaus’s formulation.

¹⁰It is possible to write net production at t as a function only of E_t and R_t without losing generality; we include all three for expositional symmetry only.

We do not include climate damages in the energy sector. We exclude it for the same reason we exclude it elsewhere: we prefer to stay close to Nordhaus’s treatment so as to be able to use his damage estimates and compare to his results. In general, of course, one would expect damages also in this sector.

In reality, energy should be thought of as a vector consisting of different energy sources, such as oil and coal (and non-fossil energy sources too). It is straightforward, but notationally cumbersome, to study the multi-dimensional case in full generality. None of the main analytical results in the paper depend on E_t being a scalar, and it is easy to extend the analysis when both R_t and E_t are arbitrary vectors. In the quantitative section we will explicitly consider a model with multiple energy sources.

It remains to describe how the climate evolves. Here, let us simply describe a rather general dependence:

$$S_t = L(E^t), \tag{4}$$

where $E^t \equiv \{E_{-T}, E_{-T+1}, \dots, E_t\}$. This is a general form for what natural scientists refer to as the carbon cycle. That is, the amount of carbon in the atmosphere at time t is a function of how much fossil fuel has been extracted in the past. Here, $t = -T$ represents some date at which human emissions began. Later, we assume that L has a simple form that we argue reasonably well approximates more complicated models of global carbon circulation.

2.2 Specializing some assumptions

In the section that follows we characterize the solution to the planner’s problem in the setup described above. We provide a sharp characterization of the optimal growth problem and the optimal carbon tax that implements it under the three assumptions that we discuss in this section.

2.2.1 Preferences

First, we assume that the utility function is logarithmic.

Assumption 1 $U(C) = \ln C$.

Logarithmic preferences are commonly used and rather standard. In the quantitative section of the paper, we discuss robustness and argue that at the long time horizons that we consider in this paper the implied values of risk aversion and intertemporal elasticity of substitution are within a reasonable range.

2.2.2 Damages

Second, we turn to modelling damages. As we discussed above, the purpose of this paper is not to break new grounds on modelling environmental damages. Rather, we use the formulation used by Nordhaus and assume that damages are multiplicative:

$$F_{1t}(K_{1t}, N_{1t}, E_{1t}, S_t) = (1 - D(\gamma_t S_t)) \tilde{F}_{1t}(K_{1t}, N_{1t}, E_{1t}).$$

Here, D is the damage function. It captures the mapping from the stock of carbon dioxide in the atmosphere, S_t , to economic damages measured as a percent of final-good output. As we discuss in detail below in Section 4, the $D(\gamma S)$ mapping can be thought of in two steps. One is the mapping from carbon concentration to climate (usually represented by global mean temperature): the *climate sensitivity*. The second is the mapping from the climate to damages. Both of these mappings are associated with significant uncertainty. For reasons summarized in Roe and Baker (2007) and also explored in Weitzman (2009) and Roe and Bauman (2011), because of there being feedback effects—a higher temperature can influence the nature of energy balance—the distribution of climate sensitivities can reasonably be thought of in terms of a distribution with quite fat tails.¹¹ As far as damages are concerned, we know relatively little, especially about the costs of significant climate change. We summarize this uncertainty by introducing a stochastic variable γ , which parameterizes how a given atmospheric carbon concentration translates to output damages. This parameter plays an important role in our calibration and for the quantitative results in this paper.

Unlike Nordhaus, we adopt a slightly different functional form for D by assuming exponential damages. As we show in the numerical section, our exponential-damage formulation approximates Nordhaus’s formulation well.

Assumption 2 *The production technology can be represented as*

$$F_{1t}(K, N, E, S) = (1 - D_t(S)) \tilde{F}_{1t}(K, N, E)$$

where $D_t(S) = 1 - \exp(-\gamma_t(S - \bar{S}))$ and where \bar{S} is the pre-industrial atmospheric CO_2 concentration.

2.2.3 The carbon cycle

Third, we describe the carbon cycle.

Assumption 3 *The function L is linear with the following depreciation structure:*

$$S_t = \sum_{s=0}^{t+T} (1 - d_s) E_{t-s} \tag{5}$$

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

Here $1 - d_s$ represents the amount of carbon that is left in the atmosphere s periods into the future. In RICE, Nordhaus also has a linear schedule but includes three stocks, all containing carbon, and a linear exchange of carbon between them. It is possible to

¹¹For example, the melting of ice reduces the earth’s capacity to reflect sunlight. Letting x denote the strength of this positive feedback, the long-run climate sensitivity depends on $\frac{1}{1-x}$. Symmetric uncertainty about x thus translates into a skewed and fat-tailed distribution of $\frac{1}{1-x}$.

show quantitatively that, for the kinds of paths considered by Nordhaus, a one-dimensional representation comes close to his formulation. We discuss the comparison with Nordhaus's carbon-cycle formulation in more detail below.

While for the rest of the analytical section we do not need to take any stand of a particular form of damages, it is useful to give an example of a tractable formulation of damages that we use in the quantitative section. As we discuss in detail in section 5, a reasonable approximation to the carbon cycle is a process under which a share φ_L of carbon emitted into the atmosphere stays there forever, a share $1 - \varphi_0$ of the remainder quickly exits the atmosphere into the biosphere and the surface oceans, and the remaining part decays at a geometric rate φ . Such a process implies that $1 - d_s$ can be represented as

$$1 - d_s = \varphi_L + (1 - \varphi_L) \varphi_0 (1 - \varphi)^s. \quad (6)$$

2.3 The planning problem

We now return to the general formulation in Section 2.1, state the planning problem, and characterize the solution to it in terms of some key relationships that will later be compared to market outcomes. Later, we will point out how our more specialized assumptions yield more specific results.

$$\max_{\{K_{t+1}, R_{t+1}, E_t, C_t, E_{1t}, E_{2t}, N_{1t}, N_{2t}, K_{1t}, K_{2t}, S_t\}_{t=0}^{\infty}} \mathbb{E} \sum_{t=0}^{\infty} \beta^t U(C_t)$$

subject to (1), (2), (3), (4),

$$C_t + K_{t+1} = F_{1t}(K_{1t}, N_{1t}, E_{1t}, S_t) + (1 - \delta) K_t$$

as well as the non-negativity constraints. In the problem above, we suppress uncertainty for notational convenience. As mentioned, damages are a function of the carbon concentration with a random element γ . In a full specification of the problem, histories of the random outcomes of γ , as defined by $\gamma^t \equiv \{\gamma_0, \gamma_1, \dots, \gamma_t\}$, along with associated probabilities $\pi(\gamma^t)$, would be spelled out explicitly instead of through the expectations operator.

Denote the constraint multipliers $\beta^t \lambda_{1t}$ (the shadow value of the final good at t), $\beta^t \lambda_{2t}$ (the shadow value of energy production at t), $\beta^t \zeta_t$ (the shadow value a unit of carbon in the atmosphere at t), and $\beta^t \mu_t$ (the shadow value of the natural resource at t), respectively. We do not introduce separate multipliers for the constraints summing up inputs across sectors but rather directly state the implied conditions.

First, thus, optimal allocation of inputs across sectors implies

$$\lambda_{1t} \frac{\partial F_{1t}}{\partial K_{1t}} = \lambda_{2t} \frac{\partial F_{2t}}{\partial K_{2t}}, \quad \lambda_{1t} \frac{\partial F_{1t}}{\partial N_{1t}} = \lambda_{2t} \frac{\partial F_{2t}}{\partial N_{2t}}, \quad \text{and} \quad \lambda_{1t} \frac{\partial F_{1t}}{\partial E_{1t}} = \lambda_{2t} \frac{\partial F_{2t}}{\partial E_{2t}}. \quad (7)$$

The first-order condition for K_{t+1} and C_t yield the standard consumption Euler equation

$$U'(C_t) = \beta \mathbb{E}_t \left[U'(C_{t+1}) \left(\frac{\partial F_{1,t+1}}{\partial K_{t+1}} + 1 - \delta \right) \right]. \quad (8)$$

Turning to the choice of energy use and emissions, combining the first-order condition for E_t and E_{1t} , one obtains

$$\frac{\partial F_{1t}}{\partial E_{1t}} \lambda_{1t} = \lambda_{2t} + \mu_t + \lambda_t^s, \quad (9)$$

where

$$\lambda_t^s = \mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \zeta_{t+j} \frac{\partial S_{t+j}}{\partial E_t}.$$

This first-order condition sets the marginal productive value of using fossil fuel (in period- t utils), on the left-hand side, equal to the marginal production cost, measured as the direct resource value of a unit of fossil fuel, λ_{2t} , plus the shadow value representing the scarcity of the stock, μ_t , plus the marginal damage cost, λ_t^s . The marginal damage cost, λ_t^s , which equals the expected present value of damages, including one in the present period t , plays a key role in our analysis. This is the externality damage that markets do not take into account, unless it is taxed accordingly. The first-order condition for S_t delivers

$$\zeta_t = -\frac{\partial F_{1t}}{\partial S_t} \lambda_{1t},$$

so that we can write the marginal damage cost as

$$\lambda_t^s = -\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \lambda_{1,t+j} \frac{\partial F_{1,t+j}}{\partial S_{t+j}} \frac{\partial S_{t+j}}{\partial E_t}. \quad (10)$$

It is useful later to express this damage cost in terms of period- t consumption. We thus define $\Lambda_t^s \equiv \lambda_t^s / U'(C_t)$ and have

$$\Lambda_t^s = -\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \frac{U'(C_{t+j})}{U'(C_t)} \frac{\partial F_{1,t+j}}{\partial S_{t+j}} \frac{\partial S_{t+j}}{\partial E_t}. \quad (11)$$

Equation (11) is the central equation of this paper.¹² In section 3, we show that Λ_t^s is equal to the optimal tax on the (emission) of the fossil fuel in competitive equilibrium. We now turn to analyzing the properties of Λ_t^s . Although the expression for Λ_t^s in general depends on the expectations about the future path of the endogenous variables and can be quite complicated, it simplifies significantly under the assumptions made in section 2.2. Under the assumptions 1, 2, and 3, expression (11) becomes

$$\Lambda_t^s = \mathbb{E}_t \sum_{j=0}^{\infty} \beta^j C_t \frac{Y_{t+j}}{C_{t+j}} \gamma_{t+j} (1 - d_j). \quad (12)$$

¹²The result that the marginal social damage (and, consequently, the optimal tax) is the discounted value of future marginal damages is straightforward and is well known (e.g., Perman et al. 2003, chapter 16). Withagen (1994) looks at an environment without physical capital and Grimaud et al. (2011) consider a growth model with capital and directed technological change. The key here, as we discuss below, is the simplification and analytical characterization of the components of this formula.

The advantage of formula (12) is that it expresses the costs of the externality only in terms of the exogenous parameters and the (endogenous) saving rate, along with initial output (since C_t can be written as a saving rate times output). Based on predictions from optimal growth models, we argue below that the effects of underlying trends (in technology, population, etc.) on the saving rate are rather small in for a broad range of parameter values. In particular, we show in sections 4 and 5 that several canonical models of resource extractions imply that the social planner optimally chooses to keep the ratio C_t/Y_t constant over time, regardless of initial conditions. Thus, even if the saving rate is not literally constant over the entire transition, the next proposition should provide a very useful benchmark for a large range of growth models with a climate externality.

Proposition 1 *Suppose Assumptions 1, 2, and 3 are satisfied and the solution to the social planner’s problem implies that C_t/Y_t is constant in all states and at all times. Then the marginal externality cost of emissions as a proportion of GDP is given by*

$$\hat{\Lambda}_t^s \equiv \Lambda_t^s/Y_t = \left[\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \gamma_{t+j} (1 - d_j) \right]. \quad (13)$$

This proposition provides a simple formula for studying both the quantitative and qualitative properties of the marginal externality cost of emissions. We also show in Section 3 below that the per unit extraction tax that implements the optimal allocation is precisely equal to the marginal externality cost of emissions: the Pigou tax. The marginal externality cost of emissions as a proportion of GDP is a very simple function of our basic parameters. The simplicity of the formula makes clear that—absent a dependence of the expected γ on time—the marginal externality cost of emissions as a proportion of GDP inherits the time path of GDP. Quite critically, future values of output, consumption, and the stock of CO₂ in the atmosphere all disappear from the formula. Thus, no knowledge about future technology, productivity or population is needed in order to calculate the marginal externality cost of emissions per GDP unit.¹³

The intuition for this important result is quite transparent. While damages are *proportional* to output, marginal utility is *inversely* proportional to output. Thus, whatever makes consumption or output grow (such as growth in TFP) will have exactly offsetting effects: damages will be higher, but due to decreasing marginal utility the value in terms of current consumption is not affected. We discuss natural departures from the result—say, if utility is not logarithmic—in our robustness section below.

¹³There are analytical results in the literature worth mentioning in this context, most of which are in a partial-equilibrium context. One can be found in Uzawa (2003), who considers a dynamic model in which pollution damages enter utility but in which there is no exhaustible resource; he shows that the optimal carbon taxes are proportional to income. Eyckmans and Tulkens (2003) derive a formula for damages in an environment with linear utility of consumption in which the emission to output ratio changes exogenously. Goulder and Mathai (2000) also make some headway based on the exponential damage formulation. In various studies, Hoel also exploits implications of constant marginal damages (e.g., Hoel, 2009).

Moreover, we see exactly how the different basic parameters matter. The higher expected damages raise the marginal externality cost of emissions as a proportion of GDP. A higher discount rate lowers it. The carbon-cycle parameters influence the optimal tax in the intuitive way as well: the longer the CO₂ stays in the atmosphere (through an increase in $1 - d_j$), the higher is the marginal damage cost.

This formula simplifies further if we assume that the expected time path for the damage parameter is constant, $\mathbb{E}_t [\gamma_{t+j}] = \bar{\gamma}_t$ for all j , and $1 - d_j$ is defined as in equation (6). In this case

$$\hat{\Lambda}_t^s \equiv \Lambda_t^s / Y_t = \bar{\gamma}_t \left(\frac{\varphi_L}{1 - \beta} + \frac{(1 - \varphi_L)\varphi_0}{1 - \varphi\beta} \right). \quad (14)$$

We now consider the optimality conditions for extraction assuming that the optimal R_t is interior.¹⁴ Combine the first-order condition for R_{t+1} with equation (9) to obtain

$$\left(\frac{\partial F_{1t}}{\partial E_{1t}} - \left(1 + \frac{\partial F_{2t}}{\partial R_{t+1}} \right) \frac{\lambda_{2t}}{\lambda_{1t}} \right) - \Lambda_t^s = \beta \mathbb{E}_t \left[\frac{U'(C_{t+1})}{U'(C_t)} \left(\frac{\partial F_{1,t+1}}{\partial E_{1,t+1}} - \left(1 - \frac{\partial F_{2,t+1}}{\partial R_{t+1}} \right) \frac{\lambda_{2,t+1}}{\lambda_{1,t+1}} \right) - \Lambda_{t+1}^s \right]. \quad (15)$$

This equation can be interpreted in terms of the choice between burning a marginal unit of fossil fuel at t or at $t + 1$. The left-hand side gives the net benefit of burning a unit at t : the productive value of energy less the production cost minus the damage cost. The production cost is given by $1 + \frac{\partial F_{2t}}{\partial R_{t+1}}$, a number greater than one since depleting the stock raises the cost, times $\lambda_{2t}/\lambda_{1t}$, capturing the relative value of energy in terms of consumption. The right-hand side gives the corresponding net benefit at $t + 1$, discounted to period- t utils. This equation characterizes efficient fossil-fuel use over time.

In the quantitative section of the paper we argue that for a realistic calibration of fossil-fuel supply it is useful to think of two regimes: (1) the regime in which a fossil-fuel source in sharply limited supply—oil—is mostly used, and (2) a regime in which mostly coal (or nuclear energy, if we appropriately modify the damage function), where the supply is large enough to at least last for a couple of hundred years. In the second case, the relevant point of transition into a “future” phase is likely not one where coal runs out but one where it is replaced by some alternative non-fossil technology. We assume that this occurs far enough into the future so as to not radically alter the medium-run implications of our model.

When the resource is not in limited supply, F_{2t} does not depend on R_t or R_{t+1} . This case leads to the condition

$$\left(\frac{\partial F_t}{\partial E_t} - \frac{\lambda_{2t}}{\lambda_{1t}} \right) - \Lambda_t^s = 0, \text{ for all } t.$$

This is a version of (15) which is not dynamic in the sense of relating energy use at t to that at $t + 1$, since the resource is not in limited supply.

¹⁴As long as no alternative technology for energy production becomes available, this condition is satisfied. The emergence of backstop technologies is the key alternative case. Whereas a backstop technology in use would not affect the formula for the marginal externality cost of emissions—unless it influences the saving rate—it will imply different conditions during at least one time period than in what follows.

In the oil regime (when the supply is limited) there is a trade-off between fossil-fuel energy use over time. We use equation (15) to develop an important interpretation of this trade-off. Re-write (15), together with the consumption Euler equation, as a portfolio-choice equation:

$$\mathbb{E}_t \left[U'(C_{t+1}) \frac{\frac{\partial F_{1,t+1}}{\partial E_{t+1}} - \left(1 - \frac{\partial F_{2,t+1}}{\partial R_{t+1}}\right) \frac{\lambda_{2,t+1}}{\lambda_{1,t+1}} - \Lambda_{t+1}^s}{\frac{\partial F_{1t}}{\partial E_t} - \left(1 + \frac{\partial F_{2t}}{\partial R_{t+1}}\right) \frac{\lambda_{2t}}{\lambda_{1t}} - \Lambda_t^s} \right] = \mathbb{E}_t \left[U'(C_{t+1}) \left(\frac{\partial F_{1,t+1}}{\partial K_{t+1}} + 1 - \delta \right) \right]. \quad (16)$$

This equation states that the accumulation of the two stocks—capital and (unused) fossil fuel—have to yield the same expected, marginal utility-weighted return. The return from saving a unit of fossil fuel is

$$\frac{\frac{\partial F_{1,t+1}}{\partial E_{t+1}} - \left(1 - \frac{\partial F_{2,t+1}}{\partial R_{t+1}}\right) \frac{\lambda_{2,t+1}}{\lambda_{1,t+1}} - \Lambda_{t+1}^s}{\frac{\partial F_{1t}}{\partial E_t} - \left(1 + \frac{\partial F_{2t}}{\partial R_{t+1}}\right) \frac{\lambda_{2t}}{\lambda_{1t}} - \Lambda_t^s},$$

taking into account the social costs of emitted carbon, Λ^s . We can think of this as a portfolio choice problem. The wealth should be split into capital, on the one hand, and, on the other, fossil-fuel resources left in the ground in such a way as to equalize social returns. When there is no uncertainty, the portfolio-choice equation can be written

$$\frac{\frac{\partial F_{1,t+1}}{\partial E_{t+1}} + \frac{\partial F_{2,t+1}}{\partial R_{t+1}} \frac{\lambda_{2,t+1}}{\lambda_{1,t+1}} - \Lambda_{t+1}^s}{\frac{\partial F_{1t}}{\partial E_t} - \frac{\partial F_{2t}}{\partial R_{t+1}} \frac{\lambda_{2t}}{\lambda_{1t}} - \Lambda_t^s} = \frac{\partial F_{1,t+1}}{\partial K_{t+1}} + 1 - \delta.$$

The total social return to leaving a unit of fossil fuel in the ground for one more period is equal to the return to capital. This is also a generalization of the original Hotelling formula in that it includes a damage cost Λ^s which is internalized by the planner but not by the market. In the market equilibrium, which we study below, the return from saving fossil fuel is different, as the Λ^s s does not appear—these costs are not internalized.

3 Market outcomes

The previous section characterized the solution to the social planner's problem and derived the expression for the emission externality Λ_t^s . In this section, we show that Λ_t^s is equal to the optimal tax on carbon emission that the government can impose in a competitive equilibrium to achieve the optimal allocations. The competitive equilibrium as defined here is also what underlies our quantitative analysis below comparing laissez-faire equilibria with those where taxes are set optimally.

3.1 Decentralized equilibrium

3.1.1 Consumers

A representative individual maximizes

$$\begin{aligned} & \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t U(C_t) \\ \text{s.t. } & \mathbb{E}_0 \sum_{t=0}^{\infty} q_t (C_t + K_{t+1}) = \mathbb{E}_0 \sum_{t=0}^{\infty} q_t ((1 + r_t)K_t + w_t N_t + \Pi_t + T_t), \end{aligned}$$

where r_t is the (net) rental rate of capital, w_t is the wage rate, and Π_t are the profits from the energy sector which (in general) are positive because ownership of the scarce resource has value. We assume perfect competition among energy firms. There is a large number of (identical) owners of the resource all producing at the same marginal cost. The assumption of perfect competition in oil is not ideal but serves as an important benchmark. For coal, this assumption is a better representation of reality. The assumption that all oil wells/coal mines are identical is for convenience only. It is straightforward but notationally cumbersome to consider heterogeneity across wells/mines and a gradual switching across them according to their relative profitability; our key results would not be affected. The consumption sector is competitive and has no rents from any scarce resource and, hence, delivers zero profits in an equilibrium. Denote by T_t government transfers. We assume that the present value of the transfers is equal to the present value of the tax revenues. We use probability-adjusted state-contingent prices of the consumption good. Let $Q_t(\gamma^t)$ denote the price of consumption goods in state γ^t at time t in terms of consumption goods at time 0. Define q from $\pi(\gamma^t)q_t(\gamma^t) \equiv Q_t(\gamma^t)$. As in the planning problem above, we suppress γ^t in our notation whenever possible.

There are two first-order conditions of interest here. The first condition is for consumption at different dates and states and is used to find expressions for date- and state-contingent prices:

$$q_t(\gamma^t) = \beta^t U'(C_t(\gamma^t)) / U'(C_0).$$

The second one, that for K_{t+1} , delivers (using the price characterization above):

$$U'(C_t) = \beta \mathbf{E}_t [U'(C_{t+1}) (1 + r_{t+1})]. \quad (17)$$

3.1.2 Producers

There are two types of firms. The consumption-good firm uses technology $F_{1t}(K_{1t}, N_{1t}, E_{1t}, S_t)$ to produce the consumption good. A representative firm in this sector solves

$$\max_{K_{1t}, N_{1t}, E_{1t}} F_{1t}(K_{1t}, N_{1t}, E_{1t}, S_t) - \delta K_{1t} - r_t K_{1t} - w_t N_{1t} - p_t E_{1t},$$

where p_t is the price of fossil fuel.

A representative, atomistic resource extraction firm owns a share of fossil-fuel resources. Denote a *per-unit* tax on the resource of τ . The problem of a representative resource extraction firm then is to maximize the discounted value of its profits:

$$\max_{\{K_{2t}, N_{2t}, E_{2t}, R_{t+1}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} q_t ((p_t - \tau_t)E_t - (\delta + r_t)K_{2t} - w_t N_{2t} - p_t E_{2t}),$$

subject to (1) and (2). The firm makes date- and state-contingent decisions as above. However, we suppress the stochastic shock for convenience.

Let η_t be a Lagrange multiplier on (2). Then the first-order conditions for the firms in the final-good and resource-extraction sectors imply that

$$\eta_t \frac{\partial F_{2t}}{\partial K_{2t}} = r_t + \delta = \frac{\partial F_{1t}}{\partial K_{1t}}, \quad (18)$$

$$\eta_t \frac{\partial F_{2t}}{\partial N_{2t}} = w_t = \frac{\partial F_{1t}}{\partial N_{1t}}, \quad (19)$$

and

$$\eta_t \frac{\partial F_{2t}}{\partial E_{2t}} = p_t = \frac{\partial F_{1t}}{\partial E_{1t}}. \quad (20)$$

Notice here that η can be interpreted as the relative value of inputs in different sectors. In equilibrium,

$$\eta_t = \frac{\partial F_{1t}}{\partial K_{1t}} / \frac{\partial F_{2t}}{\partial K_{2t}} = \frac{\partial F_{1t}}{\partial N_{1t}} / \frac{\partial F_{2t}}{\partial N_{2t}} = \frac{\partial F_{1t}}{\partial E_{1t}} / \frac{\partial F_{2t}}{\partial E_{2t}}, \quad (21)$$

and thus η is the equivalent of λ_2/λ_1 in the planner's problem.

We turn next to the resource extracting firm's only dynamic choice. Combine the first-order conditions with respect to R_{t+1} and E_t to obtain

$$p_t - \tau_t - \eta_t \left(1 + \frac{\partial F_{2t}}{\partial R_{t+1}} \right) = \mathbb{E}_t \left[\frac{q_{t+1}}{q_t} \left(p_{t+1} - \tau_{t+1} - \eta_{t+1} \left(1 - \frac{\partial F_{2,t+1}}{\partial R_{t+1}} \right) \right) \right]. \quad (22)$$

This is our market version of Hotelling's formula. The original Hotelling rule, derived in Hotelling (1931), applied to a monopolistic resource owner but, essentially, stated that the price of the resource less its marginal production cost grows at the real rate of interest.¹⁵ Our formula, which has uncertainty and, therefore, is stated as an asset-price condition, extends the original rule. The (net-of-tax) market price of the resource less its marginal production cost at t is equal in expectation to the discounted value of that tomorrow. Without uncertainty, we obtain the original Hotelling's formula, since then q_{t+1}/q_t is not random and equals the inverse of the gross interest rate.

¹⁵Solow (1974) and Stiglitz (1974) derive an analogous condition for the case of perfect markets and no externalities, in which case the market implements the optimal extraction path. Finally, Sinn (2008) shows how to include an externality in the condition, arguing that this naturally leads to slower extraction than in *laissez-faire*.

3.2 Optimal taxation

It is straightforward now to derive the optimal tax formula that implements the first best allocation. Taxes implementing the optimal allocation have to be set to equal the externality. Private decision makers have to pay for the social cost of burning fossil fuel. Formally, substitute expressions for p_t and η_t from (20) and (21) into (22). Comparing the resulting expression with (15), it is immediate that if tax τ_t satisfies

$$\tau_t = \Lambda_t^s, \quad (23)$$

then the optimality condition (22) for the firm coincides with the optimal resource extraction equation (15) in the social planner's problem. It is straightforward to show that if the proceeds are rebated lump-sum to the consumers all necessary conditions for competitive equilibrium coincide with those in the solution to the planner's problem. We summarize this result in the following proposition.

Proposition 2 *Suppose τ_t are set as in (23) and the tax proceeds are rebated lump-sum to the representative consumer. Then the competitive equilibrium allocations coincide with the solution to the social planner's problem.*

The per-unit oil (or emission) tax is not the only way to implement the optimal allocation. Alternatively, one can impose a value-added (sales) tax on oil τ^v , so that the revenues of the energy producer become

$$(1 - \tau_t^v) p_t E_t$$

rather than $(p_t - \tau_t) E_t$. Under the sales tax the energy producer instead maximizes

$$\max_{\{K_{2t}, N_{2t}, E_{2t}, R_{t+1}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} q_t (1 - \tau_t^v) (p_t E_t - (\delta + r_t) K_{2t} - w_t N_{2t} - p_t E_{2t}).$$

Knowing optimal τ_t one can always find the equivalent value added tax: $\tau_t^v = \tau_t/p_t$. Since taxes are often thought of in sales (percentage) units, it is important to note here that a unit tax that, as in our benchmark case, is a constant times GDP will translate into a percentage sales tax which is *decreasing* over time so long as the price of oil rises faster than does GDP.¹⁶ In the typical optimal-growth solution, this is indeed the case.

4 Emissions in oil and coal regimes

We now use the assumptions in Section 2.2 above to further characterize the optimal tax policy and allocations. Throughout this section, we assume that Assumptions 1, 2, and 3 hold. We also assume a Cobb-Douglas specification for the production function,

$$Y_t = e^{-\gamma_t (S_t - \bar{S})} A_t K_{1t}^\alpha N_{1t}^{1-\alpha-\nu} E_{1t}^\nu, \quad (24)$$

¹⁶Sinn (2008) shows that optimal tax rates on fossil fuel in limited supply should fall over time under realistic assumptions about extraction costs.

and full depreciation of capital. We discuss the plausibility of these assumptions, and how the results change if we depart from them, below; for long enough time periods, we argue that the assumptions are not entirely implausible. For the source of energy, we use the two cases discussed above: (1) oil and costless extraction up to a finite amount, and (2) coal in infinite supply but with extraction costs in the form of a constant marginal labor cost. In the second case, equation (2) becomes

$$E_t = A_{ct} N_{2t}. \quad (25)$$

We argue in the quantitative section that these two cases provide useful insights into the behavior of a more realistic model with various sources of energy.¹⁷

To summarize, we consider the following planning problem (and associated laissez-faire equilibrium, where the damage is a production externality as described in Section 3):

$$\max_{\{K_{t+1}, R_{t+1}, E_t, C_t, S_t\}_{t=0}^{\infty}} \mathbb{E} \sum_{t=0}^{\infty} \beta^t \ln(C_t)$$

subject to constraints (1), (5), and (24) in the oil regime; and, in the coal regime, the constraints (5) and

$$Y_t = e^{-\gamma_t(S_t - \bar{S})} A_t K_t^\alpha \left(N_t - \frac{E_t}{A_{ct}} \right)^{1-\alpha-\nu} E_t^\nu.$$

Both models, which we refer to as our oil and coal regimes, are similar to the neoclassical growth model. As in the neoclassical growth model without endogenous energy determination, the saving rate is constant in the optimal solution:

$$K_{t+1} = \alpha\beta Y_t.$$

The proof is standard so we omit it and just state the implication of this result in the following proposition. Equilibrium energy use will also satisfy this condition, whether or not the tax rate is set optimally.

Proposition 3 *At the optimum of the oil and coal regimes, $C_t/Y_t = 1 - \alpha\beta$ for all t . The optimal (per-unit) tax on emission is given by (13).*

This proposition implies that in both of these cases the optimal tax is easy to compute as a function of exogenous variables. In the rest of this section we show the effect of taxes on emission and equilibrium allocations and compare the dynamics of emission to that in laissez-faire.

¹⁷Early studies along these lines include Chakravorty et al. (1997).

4.1 Emissions in the oil regime

In the oil regime when there are no extraction costs, a simple version of equation (22) holds:

$$p_t - \tau_t = \beta \mathbb{E}_t \left[\frac{C_t}{C_{t+1}} (p_{t+1} - \tau_{t+1}) \right].$$

Using equation (20) together with the Cobb-Douglas specification (24) we obtain that $p_t = \nu Y_t / E_t$. Since by Proposition 3 the ratio C_t / Y_t is constant, we obtain

$$\frac{\nu}{E_t} - \frac{\tau_t}{Y_t} = \beta \mathbb{E}_t \left[\frac{\nu}{E_{t+1}} - \frac{\tau_{t+1}}{Y_{t+1}} \right].$$

This formula allows us to make several observations. First, in a competitive equilibrium with zero taxes, $E_{t+1} = \beta E_t$, or $E_t = \beta^t (1 - \beta) R_0$. In each period, thus, independently of the initial conditions, energy consumption is falling over time at the rate of discount.¹⁸

Second, consider the case when $\tau / Y \equiv \hat{\tau}$ is constant. From Propositions 1 and 2 this is the case when, for example, the damage coefficient γ_t is constant, and the government uses the optimal tax on emission. In this case

$$E_{t+1} = \frac{\beta E_t}{1 - (1 - \beta) \frac{\hat{\tau}}{\nu} E_t}, \quad (26)$$

so that E falls more slowly, for all t , than under competitive equilibrium with no taxes, and asymptotically falls at a rate β . Because it falls more slowly, since energy use has to add up to the total available resource, we conclude that the initial use of energy is at a lower level. The intuition is as follows. The Hotelling equation shows that $p_t - \frac{\hat{\tau}}{\nu}$ has to grow at the rate of interest. Since costs do not grow, it is beneficial to postpone extraction. The price path involves a higher initial price but a growth rate that is lower, eventually converging to the rate of interest.

Third, consider a tax on energy that is not in per-unit terms but raised as a percentage of value, i.e., then tax τ^v defined in section 3.2. With no extraction costs, the after-tax profits in the oil-extracting sector in each period are $(1 - \tau_t^v) p_t E_t$. First, suppose τ_t^v is constant over time. Then there is no effect at all on allocations: a constant factor $1 - \tau^v$ drops out of the Euler equation. A constant value-added tax does not affect the intertemporal decisions of the firm, and hence has no effect on allocations, no matter how high this tax is. A value-added tax rate that is decreasing over time, however, is equivalent to the per-unit taxes discussed above. If the optimal allocations are such that the per-unit tax indexed to GDP is constant, value-added tax rates that implement the optimal allocation have to be decreasing over time in order to make the economy postpone extraction.¹⁹

¹⁸With preferences exhibiting constant relative risk aversion, the rate at which E_t falls depends on initial conditions. Asymptotically, the rate at which E_t falls is equal to β adjusted by a consumption growth factor.

¹⁹This is the result in Sinclair (1992, 1994), who shows that without extraction costs the optimal value added taxes must be decreasing.

Finally, note that in the competitive equilibrium with zero taxes all of the oil resource is used up, given that there are no extraction costs and that oil is essential in production. However, although this has been presumed in the discussion above, this need not be the case on the optimal plan. The marginal benefits from oil are declining due to decreasing marginal productivity. The marginal costs may either increase or not decline as fast, depending on the specific damage function used. Indeed, if it is not optimal to extract all the oil, the tax on oil will be equal to its price, $p_t = \tau_t = \Lambda_t^s$ so that all rents are taxes away from the resource extracting sector. To find out whether it is indeed optimal to extract all the oil, one solves the stochastic difference equation above, including the condition that $E_{t+1} = R_t - E_t \geq 0$, with a given R_0 . As we increase R_0 , the energy use goes up. For a high enough R_0 , the net social value of energy becomes zero, i.e., $\frac{\nu}{E} - \hat{\Lambda}^s = 0$ (or, similarly, for a given R_0 , the value is zero for a high enough marginal damage). At that point, it is optimal to not use all the oil and simply set E to solve this equation.

As we argued above, uncertainty plays an important role in evaluating future damages. Our analysis allows an easy incorporation of stochastic damages. The optimal energy sequence is given by the stochastic difference equation

$$\frac{\nu}{E_t} - \hat{\Lambda}_t^s = \beta \mathbb{E}_t \left[\frac{\nu}{E_{t+1}} - \hat{\Lambda}_{t+1}^s \right],$$

where $\hat{\Lambda}_t^s$ is defined in (14). Consider the special case where $\mathbb{E}_t [\gamma_{t+j}] = \bar{\gamma}_t$ for all j , i.e., where the expected time path for the damage parameter is flat. Then $\hat{\Lambda}_t^s$ is $\bar{\gamma}_t$ times a constant. The special case where the expected future time path is flat is useful as an example since we can think of the future as embodying two possibilities—damages are very high, or not so high at all—and that it takes some time to find out what γ will be. Suppose that society, at some future date, learns the value of γ , and that it will remain very high. Then from that point on we can compute the new damage $\hat{\Lambda}_{hi}^s$ and solve the difference equation under certainty:

$$\frac{\nu}{E_t} - \hat{\Lambda}_{hi}^s = \beta \left[\frac{\nu}{E_{t+1}} - \hat{\Lambda}_{hi}^s \right].$$

Similarly, one can solve this equation for a lower value, $\hat{\Lambda}_{lo}^s$. It is easy to see that E will begin at a lower level and fall more slowly when $\hat{\Lambda}^s$ takes the higher value: higher damages demand lower current energy use, i.e., a postponement of energy use to future times. The intuition, again, is that when the marginal damage cost is a constant times GDP, it is optimal to extract more later, and the higher this constant cost is, the stronger is the incentive to postpone.

Of course, the possibility of a high damage cost in the future will influence optimal energy use today. Our functional forms imply, however, that there is a certainty-equivalence of sorts at work here: current policy/optimal energy use will depend on the expected path for γ only. This is true despite our model featuring true risk aversion; γ , rather, can be interpreted as the percentage marginal damage cost, and it turns out that this parameter is key for policy. The optimal policy will feature high taxes on energy now if a high cost is perceived, but if

the perception later changes to an even higher cost, the policy of course will react and raise taxes then. In other words, there is an option value of waiting as part of optimal policy.

4.2 Emissions in the coal regime

We now characterize emissions in the coal regimes. We start with the social planner's problem. In the optimal plan the planner must set E_t so that

$$\frac{\partial F_{1t}}{\partial E_{1t}} - \frac{\lambda_{2t}}{\lambda_{1t}} - \Lambda_t^s = 0, \quad (27)$$

i.e., so as to make the net benefit of coal production equal to zero (since there is an infinite supply of coal and, thus, no coal rent). This equation is derived from equation (9) with the multiplier μ_t set to zero. Given that $\lambda_{2t}/\lambda_{1t} = \frac{\partial F_{1t}}{\partial N_{1t}}/\frac{\partial F_{2t}}{\partial N_{2t}}$, we can re-write (27):

$$\nu \frac{Y_t}{E_t} - (1 - \alpha - \nu) \frac{Y_t}{(N_t - \frac{E_t}{A_{ct}}) A_{ct}} = \Lambda_t^s.$$

This expression, in turn, delivers

$$\nu \frac{1}{E_t} - (1 - \alpha - \nu) \frac{1}{N_t A_{ct} - E_t} = \hat{\Lambda}_t^s, \quad (28)$$

where $\hat{\Lambda}_t^s$ is defined in equation (14). This is a second-order polynomial equation in E_t whose only feasible solution (with positive output in both sectors) gives optimal coal consumption as

$$E_t = \frac{N_t A_{ct} + \frac{1-\alpha}{\hat{\Lambda}_t^s}}{2} - \sqrt{\left(\frac{N_t A_{ct} + \frac{1-\alpha}{\hat{\Lambda}_t^s}}{2}\right)^2 - \nu \frac{N_t A_{ct}}{\hat{\Lambda}_t^s}}. \quad (29)$$

We can compare this emission level to that in the laissez-faire regime. Under laissez-faire, emission is determined by equation (14) in which $\hat{\Lambda}_t^s$ is set to zero in all periods. The laissez-faire emission E_t^{lf} , thus, is

$$E_t^{lf} = \frac{\nu A_{ct} N_t}{1 - \alpha - \nu \tau_t^v}. \quad (30)$$

Several observations about (29) and (30) are in order. In both the social planner's problem and in the competitive equilibrium with zero taxes, the quantity of coal consumption will be increasing in the productivity of coal extraction, the technology parameter A_{ct} , and the population level N_t . In particular, as long as the product $A_{ct} N_t$ grows, which is reasonable to assume, coal use increases over time. However, there is a fundamental difference in the behavior of emissions in the optimum and in the laissez-faire case. In the optimal allocation, coal extraction will remain finite. Specifically, it will converge to the value to which $\nu/\hat{\Lambda}_t^s$ converges. Intuitively, when $A_c N$ goes to infinity, the marginal production cost goes to zero (as a fraction of output), and since the remaining cost of using coal—the social cost, i.e.,

the climate externality—is a positive constant (fraction of output), coal use could not go infinity. Because of a diminishing marginal product of coal, if coal use were to increase without bound, its marginal benefits would be zero eventually, and would thus be lower than the damage costs. Hence, coal use must converge to a finite level.

The opposite occurs in the laissez-faire regime. Equation (30) shows that as $A_{ct}N_t$ grows, coal use will be ever-growing and unboundedly large. Thus, there would be a radical difference, both in energy use, climate outcomes, and welfare, between the optimal allocation and the laissez-faire equilibrium. Of course, coal resources are finite, and ever-growing extraction is not possible. However, our view is that in, say, two hundred years one would expect a “backstop technology”, i.e., a non-fossil alternative that is economically profitable, to be available. At that time, unless the coal extraction technology is improving at a fast rate, there will still be coal left according to our estimates (for a discussion see Section 5), and thus there is really no coal scarcity, as assumed above. Moreover, if a viable backstop technology becomes available with probability one some time before the coal is exhausted (when its extraction follows formula (30)), the two models give identical predictions for extraction path as long as coal is used. In such a case, moreover, the radical differences between the optimum and the laissez-faire outcome become less radical.

4.3 Prices in the oil and coal regimes

What is the expected long-run price development for fossil fuels? The answer to this question can be of relevance for assessing the possibility that alternative energy technologies are developed by market forces alone, and for energy-saving technical change. The endogenous response of technology to price developments is a very important issue that we do not address here; there are many analyses of such issues, including the early Bovenberg and Smulders (1995, 1996) studies as well as recent studies like Acemoglu et al. (2010), Hartley et al. (2011), as well as Hassler, Krusell, and Olovsson (2011) are focused more singularly on these questions. Nevertheless, it is useful to note, for these reasons, what the price implications are in the models presented here.

Merely summarizing the above results, then, we note that oil prices must rise, in expectation, at the rate of interest (which, in turn, exceeds the growth rate of output). The fundamental reason for this is the rent originating in the finiteness of the supply, but the assumption of zero extraction costs plays a role as well. As we saw, with positive costs, prices do not exactly grow at the rate of interest; if marginal costs are positive and grow at a slower rate than the rate of interest, the price of the resource will grow at a slower rate than the rate of interest. Coal, in reality and as we represent it here, has a significant marginal cost of production and its price should therefore rise at a much slower rate. The fact that we assume that coal has no rent here plays a role as well but due to the high extraction costs it is likely a secondary factor; we conjecture that a model where coal would run out (in, say, 200-300 years) would deliver price implications very similar to those obtained here. The coal price develops according to $p_t A_{ct} = w_t$: it grows more *slowly* than wages (and output) to the extent there is technological change in coal extraction.

5 Quantitative analysis

We first discuss the calibration and then turn to the results.

5.1 Benchmark calibration

For our calculation of the marginal externality cost of emissions, since we can base our analysis on the derived closed-form formula, we only need to calibrate three sets of parameters: those involving the damage function (γ and its stochastic nature), the depreciation structure for carbon in the atmosphere (the ϕ s), and the discount factor. Thus, the remaining calibration—of the precise sources of energy, technology growth, etc.—is only relevant for the generation of specific paths of output, temperature, energy use, and so on or for discussion of robustness of the benchmark results.

5.1.1 Preferences and technology

We first calibrate the parameters relevant for the optimal tax formula, since this formula holds rather generally and, in particular, does not depend on the details of the energy supply. We assume logarithmic preferences and a Cobb-Douglas specification of technology (24). Each period is 10 years. We assume full depreciation of capital each period.

Logarithmic utility may not reflect everyone’s preferred assumptions regarding the elasticity of intertemporal substitution or risk aversion but it is commonly used and, even, standard. A long time period (10 years) should suggest lower curvature too; thus, even if one would argue for a higher curvature in an annual model, logarithmic curvature on the horizon studied here is probably not altogether unreasonable. A depreciation rate of 100% is too high, even for a 10-year period. We also solved versions of the model with a lower depreciation rate and the results did not change markedly. In particular, the key for the optimal tax formula to hold is a constant saving rate, and it continues to be very close to correct. Cobb-Douglas production is perhaps the weakest assumption. Though it has been widely used and defended, in particular in papers by Stiglitz (1974) and Dasgupta and Heal (1974), Hassler, Krusell, and Olovsson (2011) argue that, at least on shorter horizons, it does not represent a good way of modeling energy demand. In particular, a much lower input elasticity is called for if one wants to explain the joint shorter- to medium-run movements of input prices and input shares over the last half a century. However, on a longer horizon, Cobb-Douglas is perhaps a more reasonable assumption, as input shares do not appear to trend.²⁰ Finally, as for the discount rate, we do not aim to take a stand here but report results for a range of values.

We use standard values for α , β , and ν given by 0.3, 0.985¹⁰, and 0.03, respectively. The discount factor corresponds the one used by Nordhaus. When we report the optima tax rate, we do it as a function of the discount rate; thus, it is straightforward to read off the implications of a much smaller larger value, such as Stern’s choice of 0.999¹⁰.

²⁰Hassler, Krusell, and Olovsson (2011) also show that if technology is modeled as endogenous and potentially directed to specific factors, like energy or capital/labor, shares will settle down to robust intermediate values—and thus have the Cobb-Douglas feature—even if the input substitution elasticities are as low as zero.

5.1.2 The carbon cycle

The carbon emitted into the atmosphere by burning fossil fuel enters the global carbon circulation system, where carbon is exchanged between various reservoirs such as the atmosphere, the terrestrial biosphere, and different layers of the ocean. When analyzing climate change driven by the greenhouse effect, the concentration of CO_2 in the atmosphere is the key climate driver. We therefore need to specify how emissions affect the atmospheric CO_2 concentration over time. A seemingly natural way of doing this would be to set up a system of linear difference equations in the amount of carbon in each reservoir. This approach is taken by Nordhaus (2008) and Nordhaus and Boyer (2000), where three reservoirs are specified: *i.* the atmosphere, *ii.* the biosphere/upper layers of the ocean, and *iii.* the deep oceans. The parameters are then calibrated so that the two first reservoirs are quite quickly mixed in a partial equilibrium. Biomass production reacts positively to more atmospheric carbon and the exchange between the surface water of the oceans and the atmosphere also reach a partial equilibrium quickly. The exchange with the third reservoir is, however, much slower. Only a few percent of the excess carbon in the first two reservoirs trickles down to the deep oceans every decade.

An important property of such a linear system is that the steady-state shares of carbon in the different reservoirs are independent of the aggregate stock of carbon. The stock of carbon in the deep oceans is very large compared to the amount in the atmosphere and also relative to the total amount of fossil fuel yet to be extracted. This means that of every unit of carbon emitted now only a very small fraction will eventually end up in the atmosphere. Thus, the linear model predicts that even heavy use of fossil fuel will not lead to high rates of atmospheric CO_2 concentration in the long run.

The linear model sketched above abstracts from important mechanisms. One regards the exchange of carbon with the deep oceans; this, arguably, is the most important problem with the linear specification just discussed (see, Archer, 2005, and Archer et al., 2009). The problem is due to the Revelle buffer factor (Revelle and Suess, 1957): as CO_2 is accumulated in the oceans the water is acidified, which in turn limits the capacity of the oceans to absorb more CO_2 . This can reduce the effective “size” of the oceans as carbon reservoirs dramatically. Archer (2005) argues that the decrease can be as large as a factor 1/15. Very slowly, the acidity will then eventually decrease, and the pre-industrial equilibrium can be restored. This process is so slow, however, that we can ignore it in economic models. The IPCC 2007 report concludes that “*About half of a CO_2 pulse to the atmosphere is removed over a timescale 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*” and 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 year and the remaining 25% stays forever”. For our purposes, as shown above, what is key is the rate of depreciation of the atmospheric carbon concentration in excess of the pre-industrial level. Thus, rather than develop a nonlinear version of Nordhaus’s three-reservoir system, we just make direct assumptions on these depreciation rates, which we allow to change over time. From our perspective, thus, a simple, yet reasonable, representation of the carbon cycle is therefore

that we describe in equation (6), where (i) a share φ_L of carbon emitted into the atmosphere stays there forever; (ii) another share, $1 - \varphi_0$, of the remainder exits the atmosphere into the biosphere and the surface oceans within a decade; and (iii) a remaining part, $(1 - \varphi_L)\varphi_0$, decays (slowly) at a geometric rate φ . Thus, like Nordhaus, we use a linear specification, but one with a different interpretation and that implies qualitatively different dynamics. We show below that our formulation also has quantitative properties that are rather different. In fact, our formulation leads to significantly larger effects of human emissions on the climate.

We use the approximation of Archer (2005) to yield a half-life of carbon dioxide in the atmosphere of 30 periods. Hence, φ_L is set to 20%, as in the IPCC report. The remaining parameter φ_0 is set so that $d_2 = \frac{1}{2}$, giving $\varphi_0 = 0.393$. Thus, we have

$$\begin{aligned}\varphi &= 0.0228 \\ \varphi_L &= 0.2, \\ \varphi_0 &= 0.393.\end{aligned}$$

5.1.3 The damage function

We use an exponential damage function specified in Section 2.2.2 to approximate the current state-of-the-art damage function which is given in Nordhaus (2007). Nordhaus uses a proportional damage function specified as

$$1 - D_N(T_t) = \frac{1}{1 + \theta_2 T_t^2},$$

where T is the mean global increase in temperature above the preindustrial level, with $\theta_2 = 0.0028388$. The damage function D_N is, due to the square of temperature in the denominator, convex for a range of values up to some high temperature after which it is concave (naturally, since it is bounded above by 1).

For our purposes, however, we need to express the damage function in terms of the stock of atmospheric carbon, S_t . The standard assumption in the literature (say, as used in RICE) is to let the global mean temperature be a logarithmic function of the stock of atmospheric carbon:

$$T_t = T(S_t) = \lambda \ln \left(1 + \frac{S_t}{\bar{S}} \right) / \ln 2, \quad (31)$$

where $\bar{S} = 581$ GtC (gigatons of carbon) is the pre-industrial atmospheric CO₂ concentration. A standard value for the climate sensitivity parameter λ here is 3.0 degrees Celsius. Thus, we assume that a doubling of the stock of atmospheric carbon leads to a 3-degree Celsius increase in the global mean temperature. As noted above, there is substantial discussion and, perhaps more importantly, uncertainty, about this parameter, among other things due to imperfect understanding of feedback effects. Therefore, it is important to allow uncertainty, as we do in this paper.

In summary, to obtain a mapping from the carbon dioxide concentration in the atmosphere to damages as a percent of GDP, one needs to combine $D_N(T)$ and $T(S)$. This amounts to a composition of a convex and a concave function (for low values of S). In Figure

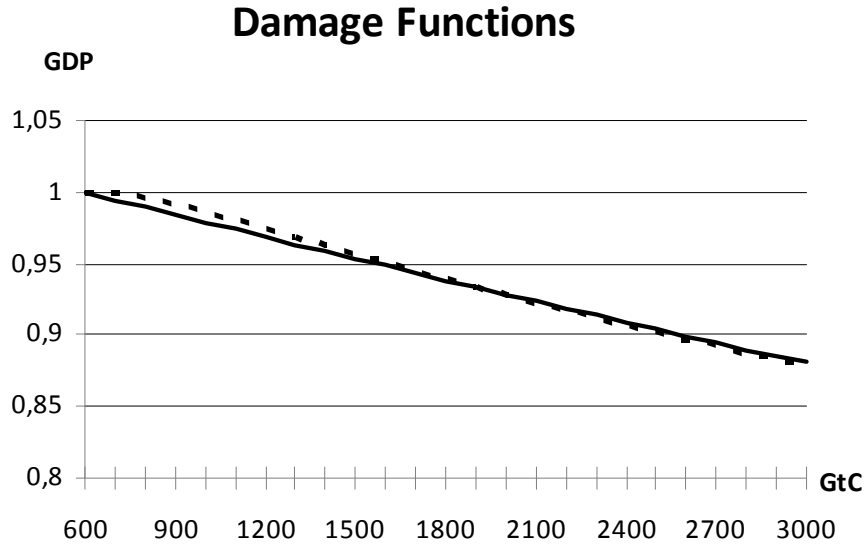


Figure 1: Net of damage function $1 - D(T(S))$; Nordhaus (dashed) and exponential (solid)

1, we show the mapping according to Nordhaus's calibration by plotting his $1 - D_N(T(S_t))$ (dashed) together with the damage function assumed in our analysis (solid): an exponential function with parameter $\gamma_t = 5.3 \times 10^{-5}$. The range of the x axis is from 600 GtC, which corresponds to preindustrial levels, to 3,000 GtC, which corresponds to the case when most of predicted stocks of fossil fuel are burned over a fairly short period of time.²¹ Nordhaus's formulation implies an overall convexity for a range of values of S , which our function does not exhibit. This convexity is not quantitatively large, however, and the two curves are quite close. We thus conclude that our exponential approximation appears rather reasonable.

To incorporate uncertainty in our analysis is straightforward. Many structures are possible; we simply assume that until some random future date there is uncertainty regarding the long-run value of γ . At that date, uncertainty is resolved and either it turns out that γ will be equal to γ^H or equal to γ^L , with $\gamma^H > \gamma^L$. The *ex-ante* probability of the high value is denoted p . Furthermore, we also assume that until the long-run value of γ is learned, the current value γ_t will equal $p\gamma^H + (1 - p)\gamma^L \equiv \bar{\gamma}$.²²

What are the sources of the specific damage parameters? When calibrating the damage function, Nordhaus (2000) uses a bottom-up approach by collecting a large number of studies on various effects of global warming. Some of these are positive, i.e., warming is beneficial, but most are negative. By adding these estimates up he arrives at an estimate that a 2.5-

²¹For a discussion of the estimates of the total stocks, see Section 5.1.4 below.

²²Pizer (1998) is an early study using uncertainty; Kelly and Kolstad (1999) also study Bayesian learning.

degree Celsius heating yields a global (output-weighted) loss of 0.48% of GDP.²³ Furthermore, he argues, based on survey evidence, that with a probability of 6.8% the damages from a 6-degree Celsius heating are catastrophically large, defined as a loss of 30% of GDP. Nordhaus, moreover, calculates the willingness to pay for such a risk and adds it to the damage function. Here, because our analysis allows uncertainty, we can proceed slightly differently. We thus directly use Nordhaus’s numbers to calibrate γ^H and γ^L . Specifically, we use the 0.48% loss at 3 degrees heating to calibrate γ^L (moderate damages) and the 30% loss at 6 degrees to calibrate γ^H (catastrophic damages). Using (31) we find that a 2.5- and a 6-degree heating occurs if S_t equals 1035 and 2324, respectively.

We thus calibrate γ^L to solve

$$e^{-\gamma^L(1035-581)} = 0.9952$$

and γ^H to solve

$$e^{-\gamma^H(2324-581)} = 0.70,$$

yielding $\gamma^L = 1.060 \times 10^{-5}$ and $\gamma^H = 2.046 \times 10^{-4}$. Using $p = 0.068$, we calculate an ex-ante (current) damage cost $\bar{\gamma}$ of 2.379×10^{-5} .

5.1.4 Energy

We model energy sources as a combination of oil and coal. We assume that coal and oil are perfect substitutes. In the very short run, this is obviously unreasonable—we do not drive our cars on coal. However, in the longer run, and for other energy uses, it seems quite plausible to assume a high degree of substitutability. For example, coal can be converted to liquid fuel that can be used to run existing combustion engines. Formally, we assume

$$E_t = E_{ct} + cE_{ot}, \tag{32}$$

where E_{ct} and E_{ot} are the period t consumption levels of coal and oil, respectively, and c is the efficiency of a unit of oil relative to that of a unit of coal. The fact that price per ton of oil is several times higher than coal can be interpreted as representing higher general efficiency of oil. We use the difference in price to calibrate c in the analysis below.

In order to maintain tractability, along the lines of the above analysis, we make special assumptions about extraction costs and scarcity values. The extraction cost for oil is set to zero, since it is very small relative to the market price at least for the very large on-shore reserves in the Middle East. The marginal extraction cost for coal is set to a positive constant, which is allowed to vary over time, while at the same time coal is assumed to be in economically infinite supply. The latter assumption—no scarcity rent for coal—appears plausible because the reserves of coal are very large compared to those for oil. Rogner (1997) estimates global fossil fuel reserves, taking into account technical progress, to be over 5,000

²³Reduced-form estimates, e.g. those in Mendelsohn and Nordhaus (1994) exploiting interregional differences using cross-sectional data on temperatures and output from countries and regions within countries, suggest higher damages but are of the same order of magnitude. The regression coefficient on the “distance-from-equator” variable in Hall-Jones productivity regressions is a relative of the Mendelsohn-Nordhaus study.

GtC of oil equivalents, i.e., enough for several hundreds of years at current consumption levels. Of this, around 70% is coal and 16% is oil. This leads us to conclude that it is reasonable to proceed with a case where in the optimal allocation, alternative energy sources have replaced fossil fuel before fossil fuel reserves have been depleted. We do not model alternative energy here but, since we think of it as appearing rather far into the future, we do not view it as influencing the time horizon we are primarily looking at. To sum up, our assumptions coincide with those in the previous section, but here oil and coal coexist. Finally, we will employ a calibration such that in the optimal as well as the decentralized equilibrium, the oil will be extracted first and coal only later. The reason is well known and straightforward: it is better to use the cheapest resource first, in order to delay any additional extraction costs.

We state the maximization problem and provide characterization of this problem in the technical appendix to our paper.²⁴ It is straightforwardly analyzed; the main complication is in determining when the oil-to-coal switch will take place and whether both oil and coal will be used in the last period of oil use.²⁵ A key (and easily checked) feature of the combined oil-coal model is that C_t/Y_t again is constant for all t , so that the optimal tax on emissions remains given by (13).

The marginal coal extraction cost is given by $\frac{w_t}{A_t^c}$, where w_t is the wage. Normalizing the labor supply to unity, the wage, if no labor is used for coal extraction, is $w_t = (1 - \alpha - \nu) Y_t$.²⁶ The 5-year average of coal prices between 2005 and 2009 of \$74/ton.^{27,28} Under the assumptions that taxes currently are negligible and that the coal market is competitive, this price equals the marginal carbon extraction cost. We thus apply these equations for $t = 0$, taken in the middle of the previous decade. We use a yearly world GDP of 75 trillion US\$, thus amounting to 750 trillion US\$ for a model time period.²⁹ Given that a period in the model is a decade, this delivers, after solving for A_{c0} ,

$$A_{c0} = \frac{(1 - 0.3 - 0.03) \cdot 750 \cdot 10^{12}}{74 \cdot 10^9} = 6791/GtC.$$

Thus, a share $\frac{1}{6791}$ of a decade's labor supply is required for each gigaton of coal extracted.

To estimate the efficiency differences between oil and coal, we calculate the average price per ton of carbon content for oil relative to the coal price over the period 1990–2009.³⁰ The

²⁴The technical appendix is available online at <http://hassler-j.iies.su.se>.

²⁵With continuous time, the characterization of the switching is slightly easier. Then, the marginal value (price) of the last unit oil is determined by the extraction cost at the time of the switch to coal. Before that, oil use is determined by the Hotelling equation (26). Our discrete-time setting implies that either oil and coal are used together during the last period of oil use or they are not; both outcomes are generic.

²⁶We calibrate using current costs of coal extraction. Currently, a positive, though very small, share of world labor is employed in coal extraction. However, if instead only coal was used, the wage would be $(1 - \alpha) Y_t$, producing only a minor difference in the quantitative results.

²⁷U.S. Central Appalachian coal. Source: BP (2010).

²⁸The 10-year average over 2000–2009 is 58.8 dollars per ton.

²⁹For these calculations, we use PPP-adjusted GDP from the CIA World Fact Book (<https://www.cia.gov/library/publications/the-world-factbook/>).

³⁰We use spot prices of Dubai crude oil per barrel from BP (2010) and convert this to tons by multiplying

price per unit of coal is on average five times as high for oil as for coal: $c = 5$.

Another key parameter is the size of the oil reserve. BP (2010) reports that global proved reserves of oil are 181.7 gigatons. However, these figures only aggregate reserves that are economically profitable to extract at current economic and technical conditions. Thus, they are not aimed at measuring the total resource base taking into account in particular technical progress, and they do not take into account the chance that new profitable oil reserves will be discovered. Rogner (1997) instead estimates global reserves taking into account technical progress, ending up at an estimate of over 5000 gigatons of oil equivalents.³¹ Of this, around 16% is oil, i.e., 800 GtC. We take as a benchmark that the existing stock of oil is 400 GtC, i.e., somewhere well within the range of these two estimates. We also assume that there is growth in both extraction efficiency and population so that $A_t^c N_t$ grows at the rate 1% per year.

5.2 Results

We now show the implications of our calibrated model, beginning with the marginal externality damage of emissions.

5.2.1 The marginal externality damage and the optimal tax

Recall that the marginal externality damage of emissions—or, alternatively, the optimal tax on emissions—is characterized by Proposition 1. This tax depends only on the parameters β , γ and the φ s and is independent of the other parameters of the model and of initial conditions. We calculate the optimal taxes both before and after we have learnt the long-run value of γ . We use (14) and express the tax per ton of emitted carbon at a yearly global output of 70 trillion dollars. In Figure 2, we plot the three tax rates against the yearly subjective discount rate.

Two important and much-discussed policy proposals have been made so far: that in Nordhaus (2000) and that in the Stern report (Stern, 2007). These proposals amount to a tax of \$30 and \$250 dollar per ton coal, respectively. A key difference between the two proposals is that they use very different subjective discount rates. Nordhaus uses a rate of 1.5% per year, mostly based on market measures. Stern, who adds a “moral” concern for future generations, uses the much lower rate of 0.1% per year. In Figure 2, we plot the optimal taxes for a range of yearly subjective discount rate (the solid line is the ex-ante tax before the uncertainty is realized, the upper and lower dashed lines are, respectively, the optimal taxes for the high and low values of damages after the true value of damages is known). For these two values of the discount rate, the optimal taxes using our analysis are \$56.9/ton and \$496/ton, respectively. Thus, our calculations suggest a substantially larger optimal tax than both these studies. This difference is due to a number of factors; one factor is that our depreciation structure for carbon in the atmosphere, as calibrated,

by 7.33 barrels per ton and use a carbon content of 85%. For the coal price, we use the price of US Central Appalachian coal from BP (2010).

³¹By expressing quantities in oil equivalents, the difference in energy content between natural gas, oil, and various grades of coal is accounted for.




Figure 2: Optimal tax rates in current dollars/ton of emitted fossil carbon vs. yearly subjective discount rate.

implies that more carbon stays, and stays longer, in the atmosphere. Other factors include different utility-function curvatures and different temperature dynamics; we discuss all of these in detail in Section 5.2.3. Furthermore, we see that the consequences of learning are dramatic. With a discount rate of 1.5%, the optimal tax rate if damages turns out to be moderate is \$25.3/ton but \$489/ton if they are catastrophic. For the low discount rate, the corresponding values are \$221/ton and a whopping \$4,263/ton.

It is useful to relate these numbers to taxes used in practice. Here, we report that Swedish taxes on the private consumption of fossil fuel in fact are above our estimate based on the low discount rate (i.e., above \$500/ton). In 2010, the tax is 1.05 SEK per kilo emitted CO₂ (Swedish Tax Agency, 2010) corresponding to a tax per ton of carbon of \$US 600.³² A conclusion from this is that the Swedish consumption taxes, even from the perspective of Stern-level discounting, are actually too high. Furthermore, to justify the level of the Swedish carbon tax on private households, even if the intention of Swedish policymakers is to do what is best for the world, is harder still, because of “carbon leakage”, i.e., the idea that if one country lowers its demand for energy, world prices are reduced and other countries might demand more.³³ To take this into account, one would obviously need a model with more than one region. Finally, it turns out that our damages estimates/optimal taxes for standard discount rates are quite close to observed values of traded emission rights; we discuss the details in the concluding section.

It should be noted that the large difference in the assumed discounting between Nordhaus and Stern has implications for other aspects of the model too. If one uses Stern’s discount rate, in particular, it follows that the laissez-faire equilibrium gives much too little saving (and much too high a market interest rate), calling for major subsidies to saving as well. Stern’s view is not necessarily that capital accumulation is too low, but it is challenging to provide a theoretically consistent model where different discounting should be applied to different forward-looking decisions. One such model is that in Sterner and Persson (2008), who model the demand for environmental goods explicitly and assume a non-homotheticity in utility leading to trend growth in relative prices and implications for discounting that potentially can justify Stern’s position.

Does our analysis have implications for whether one (i.e., a global union of countries) should use taxes or quantities—cap and trade—for attaining the full optimum? In the model discussed, so long as there is no restriction either on tax rates or on quantity limits (they need to be allowed to vary over time and across states of nature), there is in principle no difference between tax and quantity measures. At the same time, our model reveals a new argument for taxes: the optimal tax formula does not, as long as the assumptions allowing us to derive it are met, require any specific knowledge about available stocks of fossil fuel,

³²A kilo of CO₂ contains 0.27 kilos of carbon. Using an exchange rate of 6.30 SEK/\$, this yields a tax of \$617.28/tC. Swedish taxes on industrial consumption of fossil fuel are substantially lower.

³³The most straightforward concrete illustration of this is a one-period model with an inelastically supplied amount of fossil fuel (say, because extraction costs are negligible): then whatever demand is withdrawn from one country will be exactly offset by an identical increase in other countries, and no single country can influence outcomes with unilateral tax increases.

technology or population growth rates, or more generally about anything beyond the three sets of parameters in the formula. Quantity restrictions, on the other hand, demand much more knowledge; in fact, they require knowledge of all the remaining aspects of the model. As we shall see below, it is not difficult to generate quantity paths once these assumptions are made, but so much is unknown about both total current (and yet-to-be-discovered) stocks as well as about different kinds of technological change to occur in the future that one would need to worry greatly about possible quantity misjudgments.

5.2.2 Implications for the future: climate, damages, and output

Given the assumptions made in Section 5.1.4 about oil and coal reserves, etc. (in addition to the assumptions underlying the optimal tax rates), we can now generate quantity paths very easily for all variables of interest both with and without optimal policy.³⁴ We leave the details of the model solution to the technical appendix and summarize the results here. The results reported refer to the case where the damage parameter γ remains, throughout time, at its expected level (significant adjustments apply in the two cases of a much higher, or a much lower, value of γ).

The use of fossil fuel in the optimal allocation and in *laissez faire* are depicted in Figure 3. Comparing the two outcomes, we find that optimal policy demands a rather significant delay of oil consumption. The no-tax market economy would finish the oil supplies in three decades, whereas optimal taxation would let oil be used for six decades. After the oil supply is used up, coal use increases due to technical progress. The increase is particularly rapid under *laissez faire*.

The paths for total damages are plotted in Figure 4 below. There are significant, though not enormous, gains from raising taxes to the optimal level. The gains in the short run are rather small, but start growing after around 50 years and reach 4.2% of GDP after a hundred years. In the very long run, a share $\varphi_L = 0.2$ of emitted fossil carbon remains in the atmosphere, implying a permanent damage of 2.8% and 1.1% percent of GDP.³⁵

Similarly, by using the relation between carbon concentration in the atmosphere and the temperature—using the functional forms above where T depends logarithmically on S —we can also compute the climate outcomes under the optimal and the market allocations. The results are summarized in Figure 5 below. We predict a market economy with a significant temperature increase that accelerates at the time when coal is replacing oil as the energy source. This acceleration is due to our assumption that coal is a substantially less efficient source of energy per unit of carbon than oil. The temperature will peak at 7.7 degrees Celsius a little over 100 years from now. The optimal allocation dictates a significant delay in warming due to the longer oil use, by about two decades, with a peak at 4.6 degrees Celsius. Note, however, that this is measured relative to the pre-industrial level. Relative to the model’s prediction for the current temperature, the increases are almost two degrees smaller.³⁶

³⁴As discussed above, the combined oil-coal model can be solved easily for the optimum, and it is equally straightforward to solve for the competitive equilibrium without taxes, or with suboptimal taxes.

³⁵Needless to say, the true permanent effect of global warming is very uncertain.

³⁶Standard models of climate change tends to over-predict the heating relative to current temperatures.




Figure 3: Fossil fuel use per year in Gt of oil equivalents in optimal allocation (lower curve) and *laissez faire* (upper curve)




Figure 4: Damages in percent of GDP in optimal allocation (lower curve) and *laissez faire* (upper curve).

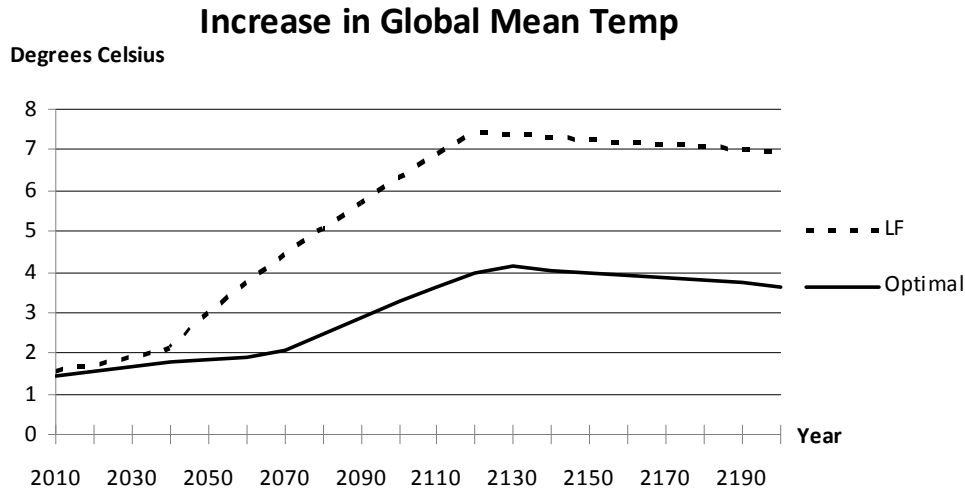


Figure 5: Increase in global mean temperature (Celcius) over pre-industrial levels in optimal allocation (lower curve) and *laissez faire* (upper curve)

Finally, we show the evolution of relative (net-of-damage) production of final-good output (GDP) in Figure 6 below. The optimal allocation involves short-run losses in GDP at around 1%, due to lower fossil fuel use. Of course, in the longer run, GDP regains (since GDP is measured net of damages) and is about the same percentage higher than under *laissez faire*, but in perpetuity. Output in the optimal allocation exceeds that in *laissez faire* from 2050. The first period coal is used in the optimal allocation, 2070, relative output falls slightly due to the fact that substantial labor resources then needs to be reallocated to coal extraction. However, GDP remains substantially above that in *laissez faire*. After the fossil fuel era is over, the CO₂ concentration very slowly decays at a rate given by the assumption half of the excess carbon is removed over three hundred years.

It is important to reiterate that the paths estimated above assume constant damage coefficients equalling the appropriate expected values calibrated above. Clearly, to the extent damages, for example because feedback effects are stronger than expected, are much higher, the above paths would need to be adjusted upward (and a similar adjustment downward is of course possible too). Similarly, the effects of adopting Stern's proposed discount factor instead of Nordhaus's would also be major in terms of the difference between the optimum and the *laissez-faire* outcomes.

We have also calculated the optimal and *laissez-faire* allocations under the alternative

Our model over-predicts the current temperature by around one degree. A common explanation for this is that anthropogenic aerosols lead to a cooling effect, temporarily masking the full impact of greenhouse gases (see, Schwartz et al., 2010).

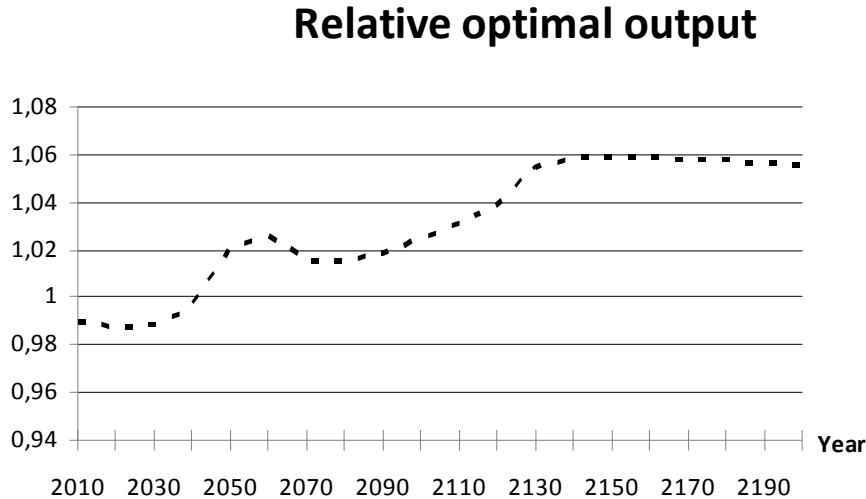


Figure 6: Final-good output (GDP) in optimal allocation relative to that in *laissez faire*

assumption that there is 600 Gt of oil reserves, i.e., 50% more than in the baseline calibration. A key difference is that with more oil, the difference between the two allocations becomes larger. Specifically, the initial drop in output due to the introduction of the optimal CO₂ tax (that we note is independent of the oil reserves if expressed as a share of GDP) is slightly larger with more oil (1.30% rather than 1.05%). Furthermore, the local peak in relative output that occurs at the end of the oil period is larger, at over 104.6% rather than at 102.6%.

5.2.3 Robustness

We now discuss how our optimal taxation results are affected by changes in the benchmark model. First, suppose the utility function is in the more general power-function class (i.e., constant relative risk aversion and constant elasticity of intertemporal substitution). In particular, consider the version of equation (13), with saving rates inserted, when the coefficient of relative risk aversion is σ :

$$\hat{\Lambda}_t^s = \mathbb{E}_t \sum_{j=0}^{\infty} \hat{\beta}_j \frac{1 - s_t}{1 - s_{t+j}} \gamma_{t+j} (1 - d_j), \quad (33)$$

with

$$\hat{\beta}_j \equiv \beta^j \left(\frac{c_t}{c_{t+j}} \right)^{\sigma-1}.$$

With maintained other assumptions there are two key implications to consider of higher curvature: (i) de-facto discounting, as given by β_j , changes since marginal utility will shrink

at a rate that is not equal to the rate of consumption growth; and (ii) there are transitional dynamics in the saving rate. Moreover, through all these effects, all the remaining elements of the calibration (technology growth, etc.) re-enter and affect the marginal externality damage. If there is consumption growth and curvature is higher than for log, discounting goes up, making for smaller externalities of current emissions. Regarding movements in the saving rate, transitional dynamics in the capital stock are slowed down with higher utility curvature than log. This makes the saving rate increase over time when capital starts below steady state—slower adjustment to steady state implies slower capital accumulation, and hence initially lower saving rates.³⁷ Thus, if the economy starts below steady state and there is higher curvature than log, the movements in saving rates will increase the externality, thus counteracting the discounting effect. Unless there is much more than log curvature, these effects will not produce a large quantitative departure from the results based on the benchmark model.

Second, if there is less than 100% depreciation, which is reasonable to assume for 10-year time horizons, transition is also slowed down, and saving rates will change over time. Equation (33) above for the marginal externality damage continues to hold; maintaining logarithmic utility, the only new implication then occurs due to non-constant saving rates. If the economy starts below the steady state, the saving rates will decline over time. The quantitative magnitude of this effect depends on how far from the steady state, or balanced growth path, the economy is initially. In typical macroeconomic applications, and apart from high-frequency variations in capacity utilization, the general presumption appears to be that we are close to a balanced growth path. From the perspective of the world economy, this is arguably also true, though here one would think of different countries (or, more generally, subregions) being on different relative levels due to differences in policy, institutions, and so on; see Acemoglu (2009) and Jones (2002) for discussions. The distribution of countries has remained fairly constant over the last 60 or so years. Transitional dynamics in the factors making countries different would generate non-constancy of world saving rates, but there is no consensus on the nature of such dynamics, to the extent it is even expected to occur.³⁸

Third, looking at conditions on technology under which saving rates would not be constant, suppose the production function is not Cobb-Douglas between energy and the other factors (capital and labor). Less than unitary input substitutability here does appear realistic; as illustrated in Hassler, Krusell, and Olovsson (2011), on a medium to high frequency, the share of fossil fuel in costs is highly correlated with its price (but there does not appear to be a long-run trend in the share). How much saving-rate movements would be implied by such a setting? Consider extreme complementarity within the class of constant substitution

³⁷With long-run growth there is also an effect of curvature on the balanced-growth saving rate—it falls in curvature.

³⁸For example, if differences in GDP across countries were due to differences in saving rates alone, one could hypothesize further that such differences are due to, say, differential taxation, broadly defined, of capital income/entrepreneurial activity. However, differences in saving rate account for only about 20% of the total income differences; a similar share is accounted for by human-capital differences, and the remainder by total-factor productivity differences. What factors influence the accumulation of these other sources of income, and the associated transitional dynamics, is even less clear.

elasticity: output is Leontief in $AK^\alpha l^{1-\alpha}$ and $A_E E$.³⁹ Again, there is a balanced growth path (whose properties depend on how fast A and A_E grow), and the question is how saving rates vary during the transition. Here, it appears reasonable to assume that K (or A , relative to A_E) is low initially, as this will deliver increasing energy use over time, something which we have observed over a long period of time.⁴⁰ This implies high initial saving rates. Hassler, Krusell, and Olovsson (2011) solve such a model and the implied transition dynamics are rather quick. Thus, an extension even to the extreme Leontief formulation is unlikely to give very different optimal-taxation results than implied by the formula based on constant saving rates.

Different assumptions regarding the extraction of fossil fuels, or energy production more generally, can also lead to time-varying saving rates. However, since the share of fossil fuel is less than 5% of total output, extensions in this direction will only change the optimal-taxation results marginally.

What are the consequences of different formulations of how/where the damages occur? Here, one can imagine a variety of alternative formulations, and speculating about all of these goes beyond the scope of the present analysis. One formulation that is commonly considered (say, in van der Ploeg and Withagen, 2010), is an additive damage in utility: $U(C, S) = \log C - V(S)$. Under this assumption, the marginal externality damage of emissions would become

$$\hat{\Lambda}_t^s = (1 - s_t) \mathbb{E}_t \sum_{j=0}^{\infty} \beta^j V'(S_{t+j})(1 - d_j),$$

and thus the computation of the damage would require knowing the implications for the future path of carbon in the atmosphere S_{t+j} , something which is not required with our formulation. Under the assumption that V is linear, however, the formula would again be in closed form as a function of deep parameters only (except for the appearance of the initial, endogenous saving rate). Linearity is arguably not too extreme a simplification, since the composition of a concave S -to-temperature mapping with a convex temperature-to-damage function may be close to linear. Other utility-function generalizations, such as that by Sterner and Persson (2008) discussed above, would change our formula more fundamentally.

Allowing technologies for carbon capture is straightforward. If such technologies are used at the source of emissions, the tax rate should apply to emissions rather than fossil fuel use. The fact that the tax rate reflects the social cost of emission implies that it also reflects the social value of removing CO_2 from the atmosphere. Capturing CO_2 directly from the atmosphere should thus be subsidized at the rate of the optimal tax rate.

More broadly, the model here regards the world as one region. Realistically, one would want to have a model which aggregates explicitly over regions. Will such a model feature an aggregation theorem, allowing a one-region representation? Different contributions to the macroeconomic literature on inequality—between consumers and between firms—suggest

³⁹Two technology factors are now relevant; in the Cobb-Douglas case, they factorize into one.

⁴⁰A model with unitary elasticity features energy use that is declining over time, at least in the case of oil with zero extraction costs.

that whereas there will not be exact aggregation, at least if intertemporal and insurance markets are operating with some frictions, there may well be approximate aggregation; see Krusell and Smith (1998, 2006), Angeletos and Calvet (2006), Angeletos (2007), and Covas (2006). However, to our knowledge, there are no calibrated medium- to long-run models of the world economy in the literature, and the extent to which approximate aggregation would hold in such a model is an open question.

The above discussion focuses on the robustness of the optimal-tax formula. How would our predictions for output, the temperature, energy use, and so on change if one considered the generalizations just discussed? These predictions are more sensitive to our assumptions. First, unlike our basic tax formula, they require knowledge of how all variables in the model develop: all exogenous parameters matter. For example, what happens to energy use over time depends critically on the details of how the supply of energy is modeled, fossil fuel-based and other, including any technological change that would influence it. We expect that many of the predictions, however, will not be so sensitive to assumptions, like those for damages and temperature; these will inherit the same uncertainty as are contained in our tax formula, namely that over γ and the φ s.

Before concluding, let us relate our results to the state-of-the art analysis conducted by Nordhaus (2007).⁴¹ In particular, we will compare to his calculation of the optimal CO₂ tax.⁴² Nordhaus reports an optimal tax of \$27 for 2005 that should rise to \$42 in 2015. Nordhaus uses a subjective discount rate of 1.5% per year at which our tax formula yields a tax rate of \$56 dollars. However, we should note that Nordhaus uses a CRRA utility function with an intertemporal elasticity of substitution of $\frac{1}{2}$. He calibrates the subjective discount rate to yield a net return of capital of 5.5%. For log utility, as we use, he reports that the subjective discount rate should be 3% to match the 5.5% capital return with only negligible effects on the optimal tax rate. Thus, taking into account the different utility functions, a more reasonable comparison is with our prescription for 3% subjective discounting. For this discount rate, our formula yields \$32, which makes our results even more in line. However, a closer inspection implies that there are a number of countervailing effects behind this similarity.

⁴¹A review of the many, rather comprehensive, studies with various degrees of integration between the climate and the economy is beyond the scope here; many of these are extremely detailed and realistic in their focus than our present analysis. The paper by Leach (2007) is a particularly close relative of the current work—a numerically solved DGE model in the spirit of DICE. Weyant (1996) gives a detailed assessment and Weyant (2000) summarizes the main commonalities and differences behind the most widely used models. A more recent comprehensive analysis (Clarke et al., 2009) is an overview of the EMF 22 International Scenarios of the ten leading integrated assessment models used to analyze the climate actions proposed in the current international negotiations. Specifically, they discuss the impact on the climate and the costs of the three policy initiatives: (1) the long-term climate target, (2) whether or not this target can be temporarily overshoot prior to 2100; (3) and assessment of such impacts depending on when various regions would participate in emissions mitigation. For the US economy, Jorgenson et al. (2008) examine the effect on the U.S. economy of predicted impacts in key market activities using a computable general equilibrium model with multiple sectors. McKibbin and Wilcoxon (1999) is another important multi-country, multi-sector intertemporal general-equilibrium model that has been used for a variety of policy analyses.

⁴²Details of this comparison are available upon request.

First, we deal with uncertainty in different ways. Nordhaus uses a “certainty-equivalent damage function”, i.e., he optimizes under certainty. If we use the same approach, and calibrate our exponential damage function to match Nordhaus’s damage function directly, our optimal tax rates are higher by more than a factor two.

Second, there are important differences in the modeling of the carbon cycle. Specifically, while we assume that almost a half of the emissions are absorbed by the biosphere and the upper layers of the ocean within 10 years, Nordhaus assumes away such a within-period absorption completely. Using Nordhaus’s carbon cycle would again lead to higher tax rates in our model; how much would depend on the subjective discount rate (at 3% discounting, we would need to adjust tax rates upward by a factor 1.5).

Nordhaus, finally, uses a more complicated climate model, where, in particular, the ocean creates a drag on the temperature; in contrast, we assume an immediate impact of the CO₂ concentration on temperature.⁴³ Of course, this biases our estimate upwards, and more so the larger is the discount rate. It can be shown that by adjusting our carbon depreciation structure d_s in a very simple way, we can approximate the temperature response of a CO₂ emissions that follows from Nordhaus assumptions rather well. By doing so, we take into account the differences in assumptions on the carbon cycle as well as on the dynamic temperature effects of emissions. A good fit is achieved by lagging the response by one period (setting $d_0 = 1$) and then multiplying $1 - d_s$ by $\frac{1}{2}$ for all s . Using this adjusted depreciation structure in combination with a damage function that approximates the one used by Nordhaus, we obtain an optimal tax of \$37.6, which is almost identical to the one calculated by Nordhaus.

Although our model is non-linear, it does not incorporate so-called threshold effects. Thresholds refer to literal discontinuities (or “very strong non-linearities”) in some of the model relationships. Such phenomena are usually argued to occur in the feedback between the carbon cycle and the climate. E.g., it has been argued that if the global temperature rises enough, it could trigger a large amount of “new” additional greenhouse gas emissions, such as leakage from methane reservoirs near the surface of the arctic tundra. We follow Nordhaus, however, in not explicitly incorporating strong non-linearities. In our model, the appearance of this kind of non-linearity would be in the damage function: as the elasticity γ depending explicitly on atmospheric carbon S . As we showed above, Nordhaus’s damage function mapping S to damages—which we approximate rather closely—does have some convexity but this convexity is weak and, for higher levels of S , turns into a concavity. The difficulty of incorporating a non-convexity is not an analytical one—it can be rather straightforwardly analyzed using our setting, with some more reliance of numerical methods—but rather a quantitative one: at what levels of S does a nonlinearity appear, and what is its nature? There does not appear to be anything near a consensus among scientists on these issues, let alone on the issue of whether threshold effects are at all relevant. Therefore, Nordhaus’s

⁴³Recent work by Roe and Bauman (2011) show that it is important to take this drag effect into account if the climate sensitivity (λ in our analysis) is high, but much less so for more moderate values like the ones we have used. When dealing with an uncertain climate sensitivity including very large but unlikely values, this may be a relevant concern.

approach seems reasonable at the present level of scientific understanding of the links between the carbon cycle and the climate. Finally, one must be reminded that several aspects of our model have elements that are often mentioned in the context of threshold effects; one is the fact that a significant fraction of emissions stay forever in the atmosphere (a feature motivated by the acidification of the oceans) and another is our explicit consideration of a probabilistic catastrophe scenario (a very high γ).

6 Conclusions

In this paper we have formulated a DSGE model of the world, treated as a uniform region inhabited by a representative consumer dynasty, where there is a global externality from emitting carbon dioxide, a by-product of using fossil fuel as an energy input into production. We showed that, under quite plausible assumptions, the model delivers a closed-form formula for the marginal externality damage of emissions. Due to standard Pigou reasoning—if a tax is introduced that makes the user internalize the externality, the outcome is optimal—the formula also expresses the optimal tax on carbon emissions. We evaluated this formula quantitatively and found results that are about twice the size of those put forth by Nordhaus and Boyer (2000). The difference between our findings are due to a variety of differences in assumption, e.g., the carbon depreciation structure. However, it is possible to arrive at estimates that are very close to Nordhaus’s by making appropriate adjustments to carbon depreciation rates, the discount rates, utility-function curvatures, and lags in temperature dynamics. Stern (2007) arrives at much higher estimates; if we simply adjust our subjective discount rate down to the level advocated in his report, we obtain an optimal tax rate that is about twice of his.

Our estimate, for a discount rate of 1.5% per annum, is that the marginal externality damage cost is a little under \$60 per ton of carbon; for a discount rate of 0.1%, it is about \$500 per ton. We also argued that the optimal-tax computation relying on our closed form is likely robust to a number of extensions. Put in terms of projections for future taxes, our optimal-tax computation robustly implies a declining value-added tax on fossil energy use.⁴⁴ To relate our estimates to actually implemented carbon taxes, consider Sweden, where the tax on private consumption of carbon actually exceeds \$600 per ton. Though industrial carbon use is subsidized relative to private consumption in Sweden, these rates are very high from a world-wide perspective. Whether they are also too high, even with Stern’s discounting assumption and even if Swedish policymakers truly take the whole world’s utility into account, because our high taxes may induce higher fossil-fuel use elsewhere (“carbon leakage”) is an interesting issue. This issue, however, requires a more elaborate model for a meaningful evaluation. We may also relate our findings to the price of emission rights in the European Union Emission Trading System, in operation since 2005 and covering large CO₂ emitters in the EU. After collapsing during the great recession of 2008-09, the price has hovered around 15 Euro per ton CO₂, at an exchange rate of 1.4 dollar per Euro

⁴⁴It should also be pointed out that we have in mind a tax on emissions; energy use based on clean energy should not be taxed, and any negative emissions should be subsidized.

corresponding to US\$77 per ton carbon.^{45,46} This price is more in line with the optimal tax rates we find for standard discount rates.

Based on further assumptions about fossil fuel stocks and their extraction technologies and about important sources of output growth, such as TFP growth, we then computed paths for our key variables for a laissez-faire market economy and compared them to the optimal outcome. In the optimal outcome, oil extraction and use is prolonged over time—it will be exhausted over a twice as long period as under laissez-faire—and coal use is about half of that without taxes. The temperature increase will be limited to 4.5 degrees Celsius under the optimal outcome and will be over three degrees higher under laissez-faire. Total damages will rise over time and amount to a maximum of close to 7% of GDP, in a bit over 100 years from now; the optimal outcome generates at most a 2.5% GDP loss. These numbers all refer to an estimate of the damage elasticity—how much an extra unit of CO₂ in the atmosphere will decrease output in percentage terms—that is the baseline considered in Nordhaus and Boyer (2000). It is well known, however, that the damages may turn out to be much higher, either because a given carbon concentration will influence temperatures more (see, e.g., Roe and Baker, 2007, or Weitzman, 2009) or because the damages implied by any additional warming will be higher; but, of course, they can be lower, too. These numbers, and our optimal-tax prescription, should be revised up or down as more accurate measures of the damage elasticity become available; until then, it is optimal to keep it at our prescribed level.

As already mentioned, our tax formula has the very important feature that little about the economy needs to be known to compute the tax rate: one neither needs information about the precise sources of energy—fossil or not—nor about the future paths of population growth and technical change (energy-specific or other). Quantity restrictions that would implement the optimum, i.e., a “cap-and-trade” system, are equally good *in principle*: if the entire model is known. That is, to compute optimal quantity restrictions, one would critically need to know the many details that go into computing the endogenous variables in our model, for example the available stocks of fossil fuels, their extraction costs, and technological change in alternative energy technologies. Since our optimal-tax formula does not depend on these assumptions, we believe to have uncovered an important advantage of using taxes over using quantity restrictions. Other pros and cons of taxes and quantity restrictions, we believe, remain.

It is also important to realize that our optimal-tax prescription holds whether or not energy technology is provided endogenously. In terms of formal analysis, endogenous technology choice—not formally spelled out in this version of our paper—simply amounts to more model equations and more first-order conditions, the outcome of which might influence consumption and output, as well as what sources of energy are in use at different points in time. But since none of these variables appear in our central formula, the formula remains intact. An implication of this is that if taxes are set according to our formula, there is no a

⁴⁵The price of EU emissions allowances can be found on the homepage of the European Energy Exchange, <http://www.eex.com/>.

⁴⁶A ton of CO₂ contains 0.273 tons of carbon, implying a conversion factor of $0.273^{-1} = 3.66$.

priori need to subsidize alternative (“clean”) technology relative to other kinds of technology, at least not from the perspective of climate change. Such subsidization—and a possible Green Paradox (Sinn, 2008)—would of course be relevant policy issues if the optimal carbon tax cannot be implemented for some reason. Moreover, it seems reasonable that technology accumulation in general, and that for green technology in particular, ought to be subsidized, since there are arguably important externalities associated with R&D. It is far from clear, however, that there should be favorable treatment of green R&D in the presence of an optimal carbon tax. An argument in favor of this has been proposed in important recent work: Acemoglu et al. (2010) show that green-technology R&D should be favored even under an optimal carbon tax; the reason is a built-in path dependence where reliance on fossil energy eventually would lead to a disaster, motivating early efforts to switch to alternatives. We conjecture that if the present model were to be enhanced with a choice between green and fossil energy technologies, then it would be optimal to subsidize both, and rather symmetrically, given that an optimal carbon tax has been adopted.⁴⁷ Of course, this is not to say that it is feasible to implement the optimal tax: for this, worldwide agreement is needed. As a general conclusion, no general insights are yet available here, and further research in this area should be quite valuable.

Finally, it should be clear from our discussions of the model throughout the text that many extensions to the present setting are desirable. One advantage of the simplicity/tractability our model offers is precisely that extensions come at a low cost. Work in several directions along the lines of the present setting is already in progress (see Krusell and Smith, 2009, for multi-regional modeling, Hassler, Krusell, and Olovsson, 2011, for some productivity accounting and an examination of endogenous technology, and Gars, Golosov, and Tsyvinski, 2009 for a model with a back-stop technology).

7 References

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

Acemoglu, Daron, Philippe Aghion, Leonardo Bursztyn, and David Hemous (2010), “The Environment and Directed Technical Change”, *American Economic Review*, forthcoming.

Angeletos, Marios (2007), “Uninsured Idiosyncratic Investment Risk and Aggregate Saving”, *Review of Economic Dynamics*, 10:1.

Angeletos, Marios and Laurent Calvet (2006), “Idiosyncratic Production Risk, Growth and the Business Cycle”, *Journal of Monetary Economics*, 53:6.

Archer, David (2005), “The Fate of Fossil Fuel CO₂ in Geologic Time”, *Journal of Geophysical Research*, 110(C9).

Archer, David et al. (2009), “Atmospheric Lifetime of Fossil Fuel Carbon Dioxide”, *Annual Review of Earth and Planetary Sciences*, 37.

⁴⁷See also Saint-Paul (2002,2007); in the latter paper it is argued that optimal subsidies should be higher for environmental innovation even if Pigou tax on emissions is used.

- 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.
- Bovenberg and Smulders (1996), "Transitional Impacts of Environmental Policy in an Endogenous Growth Model", *International Economic Review*, 37:4.
- BP (2010), *BP Statistical Review of World Energy*, June, <http://bp.com/statisticalreview>.
- Chakravorty, Ujjayant, James Roumasset, and Kinping Tse (1997), "Endogenous Substitution among Energy Resources and Global Warming", *Journal of Political Economy*, 105:6.
- Clarke, Leon et al. (2009), "International Climate Policy Architectures: Overview of the EMF 22 International Scenarios", *Energy Economics*, 31.
- Covas, Francesco (2006), "Uninsured Idiosyncratic Production Risk with Borrowing Constraints", *Journal of Economic Dynamics and Control*, 30:11
- Dasgupta, Partha and Geoffrey Heal (1974), "The Optimal Depletion of Exhaustible Resources", *Review of Economic Studies*, 41.
- Eyckmans, Johan and Henry Tulkens (2003), "Simulating Coalitionally Stable Burden Sharing Agreements for the Climate Change Problem", *Resource and Energy Economics*, 25.
- Gars, Johan, Michael Golosov and Aleh Tsyvinski (2009), "Carbon Taxing and Alternative Energy", mimeo, Yale University.
- Goulder, Lawrence and Koshy Mathai (2000), "Optimal CO₂ Abatement in the Presence of Induced Technological Change", *Journal of Environmental Economics and Management*, 39.
- Grimaud André, Gilles Lafforgue, and Bertrand Magné (2011), "Climate Change Mitigation Options and Directed technical Change: A Decentralized Equilibrium Analysis", *Resource and Energy Economics*, 33:4.
- Hartley, Peter, Kenneth Medlock, Ted Temzelides, and Xinya Zhang (2011), "Energy Sector Innovation and Growth: An Optimal Energy Crisis", mimeo, Rice University.
- Hassler, John, Per Krusell, and Conny Olovsson (2011), "Energy Saving Technical Change", mimeo, Institute of International Economic Studies, Stockholm University.
- Hoel, Michael (2009), "Climate change and carbon tax expectations", CESifo Working Paper Series 2966.
- Hotelling, Harald (1931), "The Economics of Exhaustible Resources", *Journal of Political Economy*, 39:2.
- 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.
- Jones, Charles (2002), *Introduction to Economic Growth*, W. W. Norton.

- Jorgenson, Dale, Richard Goettle, Mun Ho, and Peter Wilcoxon (2008), "The Economic Costs of a Market-Based Climate Policy", White Paper, Pew Center on Global Climate Change.
- Kelly, David and Charles Kolstad (1999), "Bayesian Learning, Growth, and Pollution", *Journal of Economic Dynamics and Control*, 23:4.
- Krusell, Per and Anthony Smith (1998), "Income and Wealth Heterogeneity in the Macroeconomy", *Journal of Political Economy*, 106.
- Krusell, Per and Anthony Smith (2006), "Quantitative Macroeconomic Models with Heterogeneous Agents", Eds Richard Blundell, Whitney Newey, and Torsten Persson, *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress, Econometric Society Monographs*, 41, Cambridge University Press.
- Krusell, Per and Anthony Smith (2009), "Macroeconomics and Global Climate Change: Transition for a Many-Region Economy", working paper.
- Leach, Andrew (2007), "The Welfare Implications of Climate Change Policy", *Journal of Economic Dynamics and Control*, 57.
- McKibbin, Warwick and Peter Wilcoxon (1999), "The Theoretical and Empirical Structure of the G-Cubed model", *Economic Modelling*, 16:1.
- Mendelsohn, Robert, William Nordhaus, and Dai Gee Shaw (1994), "The Impact of Global Warming on Agriculture: A Ricardian Approach", *American Economic Review*, 84:4.
- Nordhaus, William (2007), "To Tax or Not to Tax: The Case for a Carbon Tax", *Review of Environmental Economics and Policy*.
- Nordhaus, William (2008), *A Question of Balance: Weighing the Options on Global Warming Policies*, Yale University Press, New Haven, CT.
- Nordhaus, William and Joseph Boyer (2000), *Warming the World: Economic Modeling of Global Warming*, MIT Press, Cambridge, MA.
- Perman, Roger et al. (2003), *Natural Resource and Environmental Economics*, Chapter 16, Pearson Higher Education.
- Pizer, William (1998), "The Optimal Choice of Climate Change Policy in the Presence of Uncertainty", *Resource and Energy Economics*, 21.
- Revelle, Roger and Hans Suess (1957), "Carbon Dioxide Exchange between Atmosphere and Ocean and the Question of an Increase of Atmospheric CO₂ during Past Decades", *Tellus*, 9.
- Roe, Gerard and Marcia Baker (2007), "Why Is Climate Sensitivity So Unpredictable?", *Science*.
- Roe, Gerard and Yoram Bauman (2011), "Should the Climate Tail Wag the Policy Dog?", mimeo, University of Washington, Seattle.

- Rogner, Hans-Holger (1997), "An Assessment of World Hydrocarbon Resources", *Annual Review of Energy and the Environment*, 22.
- Saint-Paul, Gilles (2002), "Environmental Policy and Directed Innovation in a Schumpeterian Growth Model", working paper.
- Saint-Paul, Gilles (2007), "Quels instruments pour une politique environnementale?", working paper.
- 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.
- Sinclair, Peter (1992), "High Does Nothing and Raising is Worse: Carbon Taxes Should Be Kept Declining to Cut Harmful Emissions", *Manchester School*, 60.
- Sinclair, Peter (1994), "On the Optimal Trend of Fossil-Fuel Taxation", *Oxford Economic Papers*, 46.
- Sinn, Hans-Werner (2008), "Public Policies against Global Warming: A Supply Side Approach", *International Tax and Public Finance*, 15:4.
- Solow, Robert (1974), "Intergenerational Equity and Exhaustible Resources", *Review of Economic Studies*, 41.
- Stern, Nicholas (2007), *The Economics of Climate Change: The Stern Review*, Cambridge University Press, Cambridge, UK.
- Sterner, Tomas and Martin Persson (2008), "An Even Sterner Review: Introducing Relative Prices into the Discounting Debate", *Review Environmental Economics and Policy*, 2:1.
- Stiglitz, Joseph (1974), "Growth with Exhaustible Natural Resources: Efficient and Optimal Growth Paths", *Review of Economic Studies*, 41.
- Swedish Tax Agency (2010), "Skatter i Sverige – Skattestatistisk Årsbok 2010" (in Swedish), <http://www.skatteverket.se>.
- Tol, Richard (2008), "The Social Cost of Carbon: Trends, Outliers and Catastrophes", *Economics: The Open-Access, Open-Assessment E-Journal*, 2, <http://dx.doi.org/10.5018/economics-ejournal.ja.2008-25>.
- Uzawa (2003), *Economic Theory and Global Warming*, Cambridge University Press.
- Van der Ploeg, Frederick and Cees Withagen (2010), "Is There Really a Green Paradox?", CESifo Working Paper Series 2963.
- Weitzman, Martin (2009), "On Modeling and Interpreting the Economics of Catastrophic Climate Change", *Review of Economics and Statistics*, 91.
- Weyant, John et al. (1996), "Integrated Assessment of Climate Change: An Overview and Comparison of Approaches and Results", in Eds. Bruce, J., Lee, H.P., Haites, E.F. Eds., *Climate Change 1995. Economic and Social Dimensions of Climate Change: Contribution*

of Working Group II to the Second Assessment Report of the Intergovernmental Panel on Climate Change, Cambridge University Press, Cambridge.

Weyant, John (2000), “An Introduction to the Economics of Climate Change”, prepared for the Pew Center on Global Climate Change, see

http://www.pewclimate.org/globalwarming-in-depth/all_reports/economics_of_climate_change

Withagen, Cees (1994), “Pollution and Exhaustibility of Fossil Fuels”, *Resource and Energy Economics*, 16.