

**OPTIMAL TAXES ON FOSSIL FUEL
IN GENERAL EQUILIBRIUM**

BY

Mikhail Golosov, John Hassler, Per Krusell, and Aleh Tsyvinski

COWLES FOUNDATION PAPER NO. 1410



**COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281**

2014

<http://cowles.econ.yale.edu/>

OPTIMAL TAXES ON FOSSIL FUEL IN GENERAL EQUILIBRIUM

BY MIKHAIL GOLOSOV, JOHN HASSLER, PER KRUSELL,
AND ALEH TSYVINSKI¹

We analyze a dynamic stochastic general-equilibrium (DSGE) model with an externality—through climate change—from using fossil energy. Our central result is a simple formula for the marginal externality damage of emissions (or, equivalently, for the optimal carbon tax). 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. Thus, the stochastic values of future output, consumption, and the atmospheric CO₂ concentration, as well as the paths of technology (whether endogenous or exogenous) and population, and so on, all disappear from the formula. We find that the optimal tax should be a bit higher than the median, or most well-known, estimates in the literature. We also formulate a parsimonious yet comprehensive and easily solved model allowing us to compute the optimal and market paths for the use of different sources of energy and the corresponding climate change. We find coal—rather than oil—to be the main threat to economic welfare, largely due to its abundance. We also find that the costs of inaction are particularly sensitive to the assumptions regarding the substitutability of different energy sources and technological progress.

KEYWORDS: Climate change, optimal policy, optimal taxes.

1. INTRODUCTION

IN ORDER TO ASSESS THE ROLE OF ECONOMIC POLICY for dealing with climate change, we build a global economy-climate model—an integrated assessment model—using an approach based on stochastic dynamic general-equilibrium (DSGE) methods. Our first and main finding is an analytical characterization and derivation of a simple formula for the marginal externality damage of carbon dioxide emissions. The formula also serves as a prescription for the optimal level—from a global perspective—of the tax on carbon. The social cost/optimal tax, when expressed as a proportion of GDP, turns out to be a very simple function of a few basic model parameters. Quite strikingly, the parameters only involve assumptions on discounting, a measure of expected damages, and how fast emitted carbon leaves the atmosphere. Specifically, the stochastic values of future output, consumption, and the stock of CO₂ in the

¹We thank Lint Barrage, Jiali Cheng, 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, Beijing University, Bonn, Zurich, Carlos III, REDg DGEM (Barcelona), Oxford, Princeton, and Stanford. 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.

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 optimal-tax formula, which is very simple to derive and transparent to explain, does rely on some assumptions. They are: (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 emissions; and (iv) the saving rate is constant. We discuss these assumptions in detail in the paper and argue, based on numerical analysis of the model for more general economies, that our formula applies approximately for a much more general environment than the one we derive it for exactly.² This generality, and in particular the fact that the formula requires minimal assumptions on quantity determination, is valuable in comparing different policy instruments. A policy relying on quantity restrictions, such as a cap-and-trade system, by definition requires an estimate of the optimal amount of emissions. This estimate requires, and is highly sensitive to, a range of assumptions about which there is much uncertainty. These assumptions include general future technological progress, what sources of energy are available, population growth, and so on. In contrast, the tax formula requires very few assumptions—on discounting, expected damage elasticities, and carbon depreciation—and these assumptions, furthermore, matter in very direct and transparent ways.

Our formula is a discounted, expected sum of future damage elasticities, that is, percentage output responses from a percentage change in the amount of carbon in the atmosphere, caused by emitting a unit of carbon today. Discounting here involves discounting due both to time preferences and to the fact that carbon emitted into the atmosphere depreciates, or rather exits the atmosphere and is stored elsewhere (like in the biosphere or deep oceans) where it does not cause harm. The damages occur through global warming, which is produced by higher atmospheric carbon concentration, causing production shortfalls, poor health/deaths, capital destruction, and so on. We tabulate our optimal tax rate for different levels of discounting—a parameter one may have different views on. We use estimates for the remainder of our parameters from other studies, primarily from research by climate scientists on carbon depreciation and by economists on damage measurements. An important insight from our formula is that it only involves the *expected* damage elasticities. Even though the model builds on concave utility and, hence, risk has to be taken into account, the appropriate quantity in the optimal-tax formula does not involve

²We also provide a straightforward extension to the tax formula that covers utility functions with non-logarithmic preferences. This formula does not deliver the exact optimal tax rate but an approximation to it that we show is very close numerically.

any other moment than the expected damage elasticity. It is thus noteworthy that discussions of higher moments, such as fat tails (see, e.g., Weitzman (2009)), while relevant for other questions, are not relevant when it comes to computing the optimal carbon tax using the present model. Section 5.3.3 discusses how different forms of non-convexities than those entertained here can make fat tails matter.

Whereas our optimal-tax formula tells us what to do, it does not tell us what the cost is of not using the optimal tax. Thus, we also use our integrated assessment model to more generally determine the endogenous paths for all the quantity variables with and without policy. We can thus compare “business as usual” and the optimal outcome in welfare terms. Here, our main contribution is to offer a parsimonious and yet complete model that can be computed “almost” in closed form. Despite the parsimony, we argue that the model is quantitatively reasonable and, most importantly, flexible enough to allow for a variety of alternative assumptions. For example, we study how the assumptions on different sources of energy provision or technological change matter for the future climate path as well as for the path of consumption (or any other model variable).

Our integrated assessment model tells us, first, that whereas the optimal management of when to extract oil is only of marginal importance for outcomes, the policy toward coal is all the more important. The fundamental reason for this is that the stock of coal is so much larger than the stock of (cheap-to-extract) oil. According to our estimates, it is optimal to use up all the oil. Although the *laissez-faire* economy leads to an inefficient time path for oil use—oil is used up somewhat too quickly—this inefficiency is not quantitatively significant. For coal, in contrast, the stock is so much larger, and the *laissez-faire* allocation implies a much larger total out-take of coal than what is optimal. Thus, inefficient management of our coal resources leads to large welfare losses via significant global warming. Second, we learn from the integrated analysis that the assumptions on technology, whereas of second order for the optimal-tax formula, are very important for the quantitative results on climate and the potential welfare losses from not using optimal policy. This particularly concerns the degree of substitutability between different energy sources: if it is high, not taxing coal will imply a large surge in coal use, massive warming, and, hence, significant costs of inaction. Similarly, the evolution of alternative, “green,” technology is key, especially to the extent that it is a close coal substitute.

The pioneering work on integrated assessment modeling is due to William Nordhaus (for a description of his modeling, see Nordhaus and Boyer (2000)).³ 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

³Nordhaus’s work generated much follow-up research; we comment on how our work relates to this research in the appropriate places in the text.

model called RICE—Regional dynamic Integrated model of Climate and the Economy—or, in its earlier one-region version, DICE. Our natural-science model mostly follows that used by Nordhaus. The key differences are (i) in the depreciation structure of carbon emitted into the atmosphere, which we assume has a somewhat different time profile, with a large fraction of the initial emission remaining in the atmosphere for many thousand years (see Archer (2005)), and (ii) in our assumption that the full temperature response to atmospheric carbon is immediate (Nordhaus used slower temperature dynamics; see our detailed discussion in Section 5.3 below). The economic part of Nordhaus’s model, like ours, is a natural extension of non-renewable resource models along the lines of Dasgupta and Heal (1974) to incorporate a climate externality. However, Nordhaus’s model is not fully specified as a general-equilibrium model (especially on the energy supply side) and the computational methods for solving it make it difficult to fully analyze uncertainty. Thus, whereas his model delivers an optimal tax rate on carbon that can be compared to ours, it does not allow a comparison of the optimal allocation to second-best alternatives, such as the market *laissez-faire* outcome or one with carbon taxes that are less than fully optimal. Quantitatively, when we use Nordhaus’s calibration of the discount rate (1.5% per year, using market interest rates as a guide), we find that the optimal tax ought to be roughly twice that of his—Nordhaus’s value is \$30, whereas ours is \$57 per ton of coal. Stern (2007), in contrast, used a discount rate of 0.1% and concluded that a tax of \$250 dollars per ton of coal is optimal; for that discount rate, we find \$500 dollars to be the optimal tax.⁴ Our damage estimate can be made consistent, quantitatively, with that computed by Nordhaus: Section 5.3 of the paper shows how changing our assumptions toward his, especially as regards carbon depreciation and utility-function curvature, closes much of the gap between our estimates of the optimal tax rate. We also demonstrate that the consequences of updating the damage elasticities can be dramatic. With a discount rate of 1.5%, the optimal tax rate if damages turn out to be moderate is \$25.3/ton but \$489/ton if they are what Nordhaus referred to as “catastrophic.” For the lower discount rate used by Stern, the corresponding values are \$221/ton and a whopping \$4,263/ton.

It is important to point out that we show that our tax formula applies also when the economy can endogenously direct resources toward green technology. Acemoglu, Aghion, Bursztyn, and Hemous (2012) argued that this channel, and in particular subsidies to this activity, are key for dealing properly with climate change. Our analysis also argues for research subsidies, due to externalities in research as in their work. However, whether or not they should be

⁴Like Nordhaus and Stern, we restrict attention to exponential discounting. An extension to hyperbolic discounting, as in Karp (2005), is interesting but beyond the scope of the present paper. Iverson (2012) and Gerlagh and Liski (2012) recently used versions of the present model for this purpose.

directed toward clean technology depends on details of the model. Here, Acemoglu et al. (2012) took a specific, and arguably reasonable, position based on path-dependence implying that green subsidies are not only essential but also a very powerful instrument. A further and often-discussed possibility in policy discussions is carbon capture and storage (CCS). We do not specify such technologies in our parsimonious model, but our optimal-tax formula can be used to assess them. In particular, we conclude that—under Nordhaus’s preferred discount rate—carbon should be captured/stored if the costs of doing so are below \$60 per ton of carbon, but otherwise not.⁵

An interesting aspect of our general-equilibrium results is the prescription for the time path of optimal taxes. As explained above, our estimate is that the per-unit tax on emissions should be constant as a fraction of GDP (unless new information about the parameters in the tax formula arrives). Interpreted as a percentage value-added tax, whether it will grow or fall over time depends on the kind of fossil fuel we consider. For oil, to the extent its price will grow faster than GDP, a feature of the present in most available models, the value-added tax must then *fall* over time, thus encouraging postponed oil use.⁶ For coal, whose Hotelling rent we assume is zero, one might expect a falling price due to technological change, and hence a contrasting increasing path for its value-added tax.

The present work not only stands on Nordhaus’s and Hotelling’s shoulders, but also benefits from many early analytical insights using dynamic modeling of resource extraction. Formulas for the marginal damage externality have been derived in a variety of contexts. Uzawa (2003) considered a dynamic model without an exhaustible resource where pollution damages enter utility and showed that the optimal carbon taxes are proportional to income. Eyckmans and Tulkens (2003) derived 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 made some headway based on the exponential damage formulation. In various studies, Hoel also exploited implications of constant marginal damages (e.g., Hoel (2009)). The main contribution here relative to the earlier findings is to show that a simple formula is applicable—either exactly or approximately—for a very large set of models. Aside from Nordhaus’s integrated assessment models, there are also other computational models that derive a value for the social cost of carbon; see, in particular the FUND model (Tol (1997)) or the PAGE model (Hope (2008)).

There are also many studies of optimal extraction problems, with or without a stock externality as that considered here. Nordhaus’s setting builds on Das-

⁵With explicit CCS technologies added to the present model, there would be implications for CCS use from tax policy; differences in tax rates over time (relative to the marginal cost of CCS) would imply varying CCS use over time, as in Amigues, Lafforgue, and Moreaux (2012).

⁶This result goes back to Hotelling’s famous formula (Hotelling (1931)): the oil price net of extraction costs should rise at the rate of interest, which, on average, is above the rate of real GDP growth.

gupta and Heal (1974), who did not consider externalities. Withagen (1994) and Tahvonen (1997) studied the optimal depletion of fossil fuels under a utility damage, including with a backstop technology. These papers are forerunners to a large set of studies of different energy inputs and their optimal management; recent work includes that by van der Ploeg and Withagen (2012, 2014), which showed that the qualitative features of optimal paths may depend importantly on initial conditions. van der Ploeg and Withagen (2012, 2014), furthermore, argued—like we do here—that coal is the main threat to the climate. In the present paper, in terms of analytics, our functional-form assumption for the energy composite ensures interior solutions at all times. The “Herfindahl Principle,” that is, the prescription of using the cheapest resources first (Herfindahl (1967)), fundamentally holds in the model, but all the energy sources are used simultaneously; thus, initially the economy uses “mostly oil” and gradually moves away from oil into coal and a green energy source. We also consider a very simple backstop technology that is a perfect substitute for coal (see, e.g., the early work of Hoel (2009)). Popp (2006) also looked at a backstop technology, but considered endogenous R&D toward it. The literature on endogenous technology is somewhat more recent. Bovenberg and Smulders (1995, 1996) are early contributions, and Acemoglu et al. (2012) emphasized path-dependence. These are complementary contributions to the work here, which chiefly emphasizes that the first-best optimal carbon tax formula also holds in the presence of endogenous technical change; related results can be found in Grimaud, Lafforgue, and Magne (2011). How different policies interact is also discussed in the context of the Green Paradox (the idea that the future appearance of alternative energy technology speeds up the current extraction and use of fossil fuel); see Sinclair (1992) and Sinn (2008). We do not explicitly look at CSS (carbon capture and storage) here; for recent studies, see Gerlagh (2006) and van der Zwaan and Gerlagh (2009).

Section 2 describes the model in generality as well as with more specialized assumptions. It begins with the planning problem, for which it derives our key formula for the marginal externality damage of emissions, and then looks at decentralized outcomes. This section includes the derivation of the optimal-tax formula under our key set of assumptions, and also analyzes endogenous technical change. In Section 3, we specialize the assumptions further so as to fully solve the integrated assessment model. Sections 4 and 5 contain our quantitative analysis (which relies on both oil and coal use); Section 5.3 discusses the robustness of the results. We conclude in Section 6. An analysis of the sensitivity of the optimal carbon tax formulation comprises Supplemental Material (Barrage (2014)).

2. THE GENERAL MODEL

We begin by describing the general setting in Section 2.1. We then introduce a set of additional assumptions in Section 2.2 that are key in deriving our main

results. In Section 5.3 in the paper, we discuss in detail how the results would change if one would stay within the more general setting we start out with. In Section 2.3, we state the planning problem: how to optimally allocate resources over time, taking into account how the economy affects the climate. Based on the general setting, we derive an expression for the marginal externality damage and show how it simplifies considerably with our key assumptions. We then consider the decentralized economy in Section 2.4 and identify the optimal (Pigou) tax with the marginal externality damage. In Section 2.5, we consider a particular extension—endogenous technological change that is potentially directed toward “green energy”—and show that our formula for the optimal tax still applies.

2.1. *The Economy and the Climate: A General Specification*

We consider a version of the multi-sector neoclassical growth model with $I + 1$ sectors. Time is discrete and infinite. There is a representative household with the utility function

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

where U is a standard concave period utility function, C is consumption, and $\beta \in (0, 1)$ is the discount factor.

The production process consists of what we label a final-goods sector, denoted $i = 0$ and with output Y_t , and by I intermediate-goods sectors that produce energy inputs E_i , $i = 1, \dots, I$, for use in all sectors.

The feasibility 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. The second term is undepreciated capital.⁷

Output in the final-goods sector is described by an aggregate production function $F_{0,t}$:

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

The arguments of F_0 include the standard inputs $K_{0,t}$ and $N_{0,t}$ (capital and labor used in this sector), along with $\mathbf{E}_{0,t} = (E_{0,1,t}, \dots, E_{0,I,t})$ denoting a vector of energy inputs used in this final sector at t . The sub-index t on the production

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

function captures the possibility of technical change. This change can appear in a variety of ways, for example, as 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 way energy services are produced). This change can be either deterministic or stochastic.

Finally, we also allow a climate variable S_t to affect output. The effect of S_t on aggregate production 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 specify later how $F_{0,t}$ 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 amount of carbon in the atmosphere. 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 damages, or the mapping from the atmospheric carbon concentration, to have a stochastic component, but we suppress it here for notational convenience.⁹ We discuss the stochastic component in detail later in the paper. Our assumption that S_t affects production only is made mainly so as to make our analysis closer to, and easier to compare with, Nordhaus’s RICE and DICE treatments.¹⁰ Also, for an important special case, covered in Section 3 below, damages to utility, production, and to capital can all be aggregated into the form we consider here.

We now turn to the production of energy services, which are both inputs and outputs. We assume that each component of $\mathbf{E}_{0,t}$, $E_{0,i,t}$, is produced by its own technology $F_{i,t}$, which uses capital, labor, and a vector of energy inputs. Moreover, some energy sources i are in finite supply, such as oil. For any such energy source i , let $R_{i,t}$ denote its beginning-of-period stock at t , and let $E_{i,t}$ be the total amount extracted (produced) at t . Then the decumulation equation for any exhaustible stock i is

$$(1) \quad R_{i,t+1} = R_{i,t} - E_{i,t} \geq 0.$$

⁸Roe and Bauman (2011) showed that if the long-run sensitivity of the global mean temperature to the CO₂ concentration is higher than standard estimates, it becomes important to take into account the temperature lags. This occurs because the ocean heats more slowly than the atmosphere. We discuss the implications of this mechanism in more detail below.

⁹By normalizing the amount of air in the atmosphere to unity, we follow the convention of using the stock of atmospheric carbon and the atmospheric carbon concentration interchangeably.

¹⁰Documented 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 of modeling.

The production technology for energy from source i , exhaustible or not, is

$$(2) \quad E_{i,t} = F_{i,t}(K_{i,t}, N_{i,t}, \mathbf{E}_{i,t}, R_{i,t}) \geq 0.$$

The appearance of the stock in the production function here allows for the possibility that energy production involves an exhaustible resource whose production costs may depend on how much of the resource remains.¹¹ Moreover, by allowing F to have decreasing returns to scale, we can account for the possibility that some sources of energy, like wind, cannot grow without limits due, for example, to space constraints. This general formulation of production technology incorporates most cases that are typically considered in the literature.

We assume that sectors $i = 1, \dots, I_g - 1$ are “dirty” in the sense of emitting fossil carbon to the atmosphere. Sectors I_g, \dots, I are “clean,” or “green,” energy sources which are not associated with climate externalities. We normalize E_i for $i = 1, \dots, I_g - 1$ to be in the same units—one unit of E_i produces one unit of carbon content—and the relative energy efficiencies of different sources of energy are captured implicitly in the production functions.

Each period, the production factors are allocated freely across sectors:

$$(3) \quad \sum_{i=0}^I K_{i,t} = K_t, \quad \sum_{i=0}^I N_{i,t} = N_t, \quad \text{and} \quad E_{j,t} = \sum_{i=0}^I E_{i,j,t}.$$

The process for N_t is exogenous and can be either deterministic or stochastic.

We do not include climate damages in the energy sectors, mainly for comparison with Nordhaus’s treatment and so that we can use his damage estimates. Since damages to the energy sectors are expected to be a small part of the overall economy, this omission seems quantitatively unimportant.

Turning to the evolution of the climate, let us first just describe a general formulation. Let \tilde{S}_t be a function that maps a history of anthropogenic emissions into the current level of atmospheric carbon concentration, S_t . The history is defined to start at the time of industrialization, a date defined as $-T$:

$$(4) \quad S_t = \tilde{S}_t \left(\sum_{i=1}^{I_g-1} E_{i,-T}, E_{-T+1}^f, \dots, E_t^f \right),$$

where $E_s^f \equiv \sum_{i=1}^{I_g-1} E_{i,s}$ is fossil emission at s and we recall that $E_{i,s}$ is measured in carbon emission units for all i . Later, we will assume a simple form for \tilde{S}_t that we argue approximates more complicated models of global carbon circulation quite well.

¹¹We do not explicitly model a multitude of deposits with different extraction costs; see Kemp and Van Long (1980) for such an analysis.

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

The first special assumption is logarithmic utility:

ASSUMPTION 1: $U(C) = \log C$.

Logarithmic preferences are commonly used and rather standard. At long time horizons, the risk aversion and intertemporal elasticity of substitution implied by logarithmic curvature are probably not unreasonable. We discuss this issue in more detail in the quantitative section of the paper.

2.2.2. Damages

Second, we specialize our damage formulation. We follow Nordhaus and assume that damages are multiplicative:

$$F_{0,t}(K_{0,t}, N_{0,t}, \mathbf{E}_{0,t}, S_t) = (1 - D_t(S_t))\tilde{F}_{0,t}(K_{0,t}, N_{0,t}, \mathbf{E}_{0,t}).$$

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 in Section 3, the $D(S)$ mapping can be thought of in two steps. The first is the mapping from carbon concentration to climate (usually represented by global mean temperature). 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), climatic feedback mechanisms of uncertain strength imply that it is reasonable to think of the warming effect of a given atmospheric CO_2 concentration in terms of a distribution with quite fat tails.¹² Nordhaus explicitly modeled both steps in the mapping from the carbon concentration to damages. As we show in the numerical section, an exponential specification for $D(S)$ approximates Nordhaus's formulation rather well. Note that we generally allow D to depend on time, and, implicitly, on the state of nature in case there is a random element to damages. We parameterize this dependence through the following specification:

¹²For 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 $\frac{1}{1-x}$.

ASSUMPTION 2: *The production technology can be represented as*

$$F_{0,t}(K_{0,t}, N_{0,t}, \mathbf{E}_{0,t}, S_t) = (1 - D_t(S_t))\tilde{F}_{0,t}(K_{0,t}, N_{0,t}, \mathbf{E}_{0,t}),$$

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

Thus, the (possibly time- and state-dependent) parameter γ can be used to scale the damage function.

2.2.3. The Carbon Cycle

Third, we consider the following simplified carbon cycle:

ASSUMPTION 3: *The function \tilde{S}_t is linear with the following depreciation structure:*

$$(5) \quad S_t - \bar{S} = \sum_{s=0}^{t+T} (1 - d_s) E_{t-s}^f,$$

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 had a linear carbon depreciation schedule, but based his carbon cycle on three stocks, all containing carbon, and a linear exchange of carbon between them. It is possible to 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.¹³

While for the rest of the analytical section we do not need to take any stand of a particular form of the depreciation structure in (5), it may be useful to preview the formulation we use in the quantitative section. This structure amounts to a three-parameter family with (i) a share φ_L of carbon emitted into the atmosphere staying in it forever; (ii) a share $1 - \varphi_0$ of the remaining emissions exiting the atmosphere immediately (into the biosphere and the surface oceans), and (iii) a remaining share decaying at a geometric rate φ . That is, we use

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

¹³In structural and more elaborate carbon-circulation models, the depreciation structure is generally not independent of the amount of carbon in the atmosphere. This becomes a consideration if we consider extremely large pulses of carbon emissions; see Gars and Hieronymus (2012).

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 subsequently be compared to market outcomes. Later, we will point out how our more specialized assumptions yield more specific results:

$$\max_{\{C_t, N_t, K_{t+1}, \mathbf{K}_t, R_{t+1}, \mathbf{E}_t, S_t\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t U(C_t)$$

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

$$(7) \quad C_t + K_{t+1} = F_{0,t}(K_{0,t}, N_{0,t}, \mathbf{E}_{0,t}, S_t) + (1 - \delta)K_t$$

as well as nonnegativity constraints.

Let $\beta^t \lambda_{i,t}$ be a Lagrange multiplier on the production constraint for sector i (equation (2)).¹⁴ Moreover, let $\beta^t \chi_{i,t}$ be the multiplier on the feasibility constraint for energy of type i (equation (3)) and $\beta^t \xi_{i,t}$ the multiplier associated nonnegativity constraint. Finally, let $\beta^t \mu_{i,t}$ be the multiplier on the decumulation equation for exhaustible resource i . The first-order condition with respect to $E_{i,t}$ can then be written, in terms of final consumption good at t , as

$$(8) \quad \frac{\chi_{i,t}}{\lambda_{0,t}} = \frac{\lambda_{i,t} + \mu_{i,t} + \xi_{i,t}}{\lambda_{0,t}} + \Lambda_{i,t}^s,$$

where

$$\Lambda_{i,t}^s = \mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \frac{\lambda_{0,t+j}}{\lambda_{0,t}} \frac{\partial F_{0,t+j}}{\partial S_{t+j}} \frac{\partial S_{t+j}}{\partial E_{i,t}}.$$

Since $\partial S_{t+j} / \partial E_{i,t} = 0$ for $i = I_g, \dots, I$ and, by construction,

$$\frac{\partial S_{t+j}}{\partial E_{i,t}} = \frac{\partial S_{t+j}}{\partial E_{i',t}} \quad \text{for } i, i' \in \{1, \dots, I_g - 1\},$$

we have that $\Lambda_{i,t}^s = 0$ for $i = I_g, \dots, I$ and that $\Lambda_{i,t}^s$ is independent of i for all $i \in \{1, \dots, I_g - 1\}$. Therefore, we refer to $\Lambda_{i,t}^s$ for the dirty sectors by Λ_t^s .

Equation (8) summarizes the costs and benefits of producing a unit of energy of type i . The benefit, on the left-hand side, is its use in production ($\chi_{it} = \frac{\partial F_{0,t}}{\partial E_{0,i,t}} \lambda_{0,t}$ in terms of output in sector 0, utility-weighted). The costs include (i) the cost of production (input use), $\lambda_{i,t} / \lambda_{0,t} = \frac{\partial F_{0,t} / \partial N_0}{\partial F_{i,t} / \partial N_i}$, being the amount lost in final-output units; (ii) the scarcity cost $\mu_{i,t} / \lambda_{0,t}$, which can only

¹⁴Note that since we allow uncertainty, $\lambda_{i,t}$ is a random variable.

be positive if the resource is an exhaustible one; and (iii) the marginal externality damage, Λ_t^s . This last cost can be written

$$(9) \quad \Lambda_t^s = \mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \frac{U'(C_{t+j})}{U'(C_t)} \frac{\partial F_{0,t+j}}{\partial S_{t+j}} \frac{\partial S_{t+j}}{\partial E_{i,t}},$$

and it will play a central role in our analysis. It captures the externality from carbon emission, and we show in the next section that it is exactly equal to the optimal Pigouvian tax. In general, Λ_t^s depends on the structural parameters of the model in complicated ways, both through its effect on C_{t+j} and the derivatives $\frac{\partial F_{0,t+j}}{\partial S_{t+j}} \frac{\partial S_{t+j}}{\partial E_{i,t}}$. The expression for Λ_t^s simplifies dramatically, however, when Assumptions 1, 2, and 3 are satisfied:

$$(10) \quad \Lambda_t^s = \mathbb{E}_t \sum_{j=0}^{\infty} \beta^j C_t \frac{Y_{t+j}}{C_{t+j}} \gamma_{t+j} (1 - d_j).$$

The advantage of formula (10) 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).

Expression (10) can be simplified even further if one assumes that the saving rate is constant. The tax formula in our main proposition below relies on saving-rate constancy. In Section 3, we provide a set of sufficient conditions on primitives delivering a constant saving rate. We also show, in Section 5.3, that the formula is a very good approximation when considering extensions to calibrated settings where saving rates are not constant. In the data, moreover, saving rates do not tend to vary so much over time, and long-run growth models are often specified so that C_t/Y_t is constant (see, e.g., Acemoglu (2009)).

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*

$$(11) \quad \Lambda_t^s = Y_t \left[\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \gamma_{t+j} (1 - d_j) \right].$$

This proposition provides a formula allowing us to discuss both the quantitative and qualitative properties of the marginal externality cost of emissions. 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 CO_2 in the atmosphere all disappear from the formula. Thus, no knowledge about future technology, productivity, or labor supply is needed to calculate the marginal externality cost of emissions per GDP unit.

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.

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 CO_2 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,

$$(12) \quad \Lambda_t^s / Y_t = \bar{\gamma}_t \left(\frac{\varphi_L}{1 - \beta} + \frac{(1 - \varphi_L)\varphi_0}{1 - (1 - \varphi)\beta} \right).$$

Finally, the planner needs to optimally manage any finite resource stock over time. Thus, for an energy source i that is exhaustible and has a binding constraint ($\mu_{i,t} > 0$), the first-order condition with respect to $R_{i,t+1}$ becomes $\mu_{i,t} = \beta \mathbb{E}_t(\lambda_{i,t+1} (-\frac{\partial F_{i,t}}{\partial R_{i,t+1}}) + \mu_{i,t+1})$. The first term on the right-hand side reflects the change in extraction costs as more is extracted. This equation is the core of Hotelling's famous formula: it equalizes the marginal value of extracting one unit today to the expected value of extracting it tomorrow.

2.4. Decentralized Equilibrium

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, first-best tax on carbon emission. 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.

2.4.1. *Consumers*

A representative individual maximizes

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

subject to

$$\begin{aligned} & \mathbb{E}_0 \sum_{t=0}^{\infty} q_t (C_t + K_{t+1}) \\ & = \mathbb{E}_0 \sum_{t=0}^{\infty} q_t ((1 + r_t - \delta)K_t + w_t N_t + T_t) + \Pi, \end{aligned}$$

where r_t is the (net) rental rate of capital, w_t is the wage rate, T_t is a government transfer, and Π are the profits from the energy sectors which (in general) are positive because ownership of the scarce resource has value. We use probability-adjusted state-contingent prices of the consumption good, where q_t denote Arrow–Debreu prices.

2.4.2. *Producers*

All output and input markets are assumed competitive. There are two types of firms: final-output firms and energy firms. A representative firm in the final-good sector solves

$$\begin{aligned} \Pi_0 \equiv & \max_{\{K_{0,t}, N_{0,t}, E_{0,t}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} q_t \left[F_{0,t}(K_{0,t}, N_{0,t}, E_{0,t}, S_t) \right. \\ & \left. - r_t K_{0,t} - w_t N_{0,t} - \sum_{i=1}^I p_{i,t} E_{0,i,t} \right], \end{aligned}$$

subject to nonnegativity constraints, where $p_{i,t}$ is the price of fuel of type i .

Consider first a representative, atomistic energy firm which owns a share of fossil-fuel resource i . Denote a *per-unit* tax on the resource of τ_i . The problem of this firm then is to maximize the discounted value of its profits:

$$\begin{aligned} \Pi_i \equiv & \max_{\{K_{i,t}, N_{i,t}, E_{i,t}, R_{i,t+1}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} q_t \left[(p_{i,t} - \tau_{i,t}) E_{i,t} \right. \\ & \left. - r_t K_{i,t} - w_t N_{i,t} - \sum_{j=1}^I p_{j,t} E_{i,j,t} \right], \end{aligned}$$

subject to nonnegativity constraints, the production constraint (2), and the decumulation constraint (1); firms make date- and state-contingent decisions, but we suppress the stochastic shocks for convenience. Firms producing clean energy (which may or may not involve non-fossil resources) solve similar problems. Total profits of all energy producers are $\Pi = \sum_{i=0}^I \Pi_i$.

The definition of competitive equilibrium is standard and we omit it.

2.4.3. Taxes

It is straightforward to show that the tax $\tau_{i,t} = \Lambda_t^s$ for $i = 1, \dots, I_g - 1$, $\tau_{i,t} = 0$ for $i \in \{1, \dots, I_g - 1\}$, and appropriate lump-sum rebates implement the solution to the social planner's problem. To see this, let the multipliers on the production, depletion, and nonnegativity constraints of the energy producer be $q_t \hat{\lambda}_{i,t}$, $q_t \hat{\mu}_{i,t}$, and $q_t \hat{\xi}_{i,t}$, respectively, with $\hat{\lambda}$, $\hat{\mu}$, and $\hat{\xi}$ all expressed in units of final consumption at t . The optimality conditions for labor inputs of the two kinds of firms become

$$\hat{\lambda}_{i,t} \frac{\partial F_{i,t}}{\partial N_{i,t}} = w_t = \frac{\partial F_{0,t}}{\partial N_{0,t}}.$$

The energy firm chooses i so that $\hat{\lambda}_{i,t} + \hat{\mu}_{i,t} + \hat{\xi}_{i,t} = p_{i,t} - \tau_{i,t}$. The optimality condition for energy input of type i in the final-output sector is $\frac{\partial F_{0,t}}{\partial E_{0,i,t}} = p_{i,t}$. Let us identify hat variables with their counterparts in the planning problem divided by λ_0 . We now see that the planner's optimality condition is identical to the condition here if, for all dirty technologies i , there is a uniform tax on all carbon energy inputs,

$$(13) \quad \tau_{i,t} = \Lambda_t^s \equiv \tau_t,$$

and $\tau_{i,t} = 0$ for $i \in \{1, \dots, I_g - 1\}$. It is immediate that the allocation of inputs across sectors in competitive equilibrium will also be the same. Finally, it remains to show that the energy firm manages the resource stock the same way the planner does. This is also immediate: the firm's intertemporal first-order condition directly produces the corresponding planner condition. Thus, in this model, there is no "sustainability problem": markets use the finite resource stocks optimally, so long as taxes are set so as to internalize the climate externality.

We summarize these findings in the following proposition.

PROPOSITION 2: *Suppose that τ_t is set as in (13) and that the tax proceeds are rebated lump-sum to the representative consumer. Then the competitive equilibrium allocation coincides with the solution to the social planner's problem.*

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

$$(1 - \tau_{i,t}^v) p_{i,t} E_{i,t}$$

rather than $(p_{i,t} - \tau_{i,t}) E_{i,t}$. Under the sales tax, the energy producer instead maximizes

$$\Pi_i \equiv \max_{\{K_{i,t}, N_{i,t}, E_{i,t}, \mathbf{E}_{i,t}, R_{i,t+1}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} q_t \left[p_{i,t} (1 - \tau_{i,t}^v) E_{i,t} - r_t K_{i,t} - w_t N_{i,t} - \sum_{j=1}^I p_{j,t} E_{i,j,t} \right].$$

Knowing the optimal τ_t , one can always find the equivalent value-added tax: $\tau_{i,t}^v = \tau_t / p_{i,t}$. Different energy inputs would now receive different value-added tax rates if their market values are different.

2.5. Endogenous Technical Change

So far we considered an environment in which technical change is exogenous. Here we argue that key parts of our analysis carry through also when technical change and, hence, growth is endogenous. Though the argument is rather general in nature, for ease of notation we use a model of endogenous technical change that builds on Romer (1986).

We extend the competitive equilibrium in Section 2.4 by assuming that there is a large number of firms in each sector and that each firm in sector $i > 0$ has access to the technology

$$(14) \quad E_{i,t} = A_{i,t} F_{i,t}(K_{i,t}, N_{i,t}, \mathbf{E}_{i,t}, R_{i,t}, X_{i,t}),$$

where $X_{i,t}$ denotes an expenditure on an intermediate good that is produced one-for-one from final output. What is key here is that X both gives a private return to the firm and has an R&D-like spillover effect. We refer to $A_{i,t}$ as total-factor productivity in sector i and define production technology in sector 0 similarly.

We assume that total-factor productivity in sector i in period $t + 1$ is given by

$$(15) \quad A_{i,t+1} = G_i(A_{i,t}, \bar{X}_{i,t}),$$

where G_i is a differentiable, convex function increasing in both arguments and $\bar{X}_{i,t}$ is the *average* value of X across firms in sector i . Thus, when an individual firm chooses its X , it takes into account how it affects its output today but does

not internalize its impact on the production possibility frontier in the future. The equilibrium is inefficient and subsidies are required. Let $\tau_{i,t}^X$ be the subsidy applying to sector i . Each firm in sector $i > 0$ thus solves

$$\begin{aligned} \Pi_i \equiv & \max_{\{K_{i,t}, N_{i,t}, E_{i,t}, R_{i,t}, X_{i,t}\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} q_t [(p_{i,t} - \tau_{i,t}^X) E_{i,t} \\ & - r_t K_{i,t} - w_t N_{i,t} - (1 + \tau_{i,t}^X) X_{i,t}], \end{aligned}$$

subject to nonnegativity constraints and (14), along with the law of motion for $A_{i,t}$ given by (15). The maximization problem of a firm in sector 0 is defined analogously. Moreover, since all firms in the same sector make the same choices, the economy-wide feasibility constraint in sector 0 is given by

$$(16) \quad C_t + K_{t+1} + \sum_{i=0}^I X_{i,t} = A_{0,t} F_0(K_{0,t}, N_{0,t}, E_t, S_t, X_{0,t}) + (1 - \delta) K_t.$$

The rest of the definition of competitive equilibrium is the same as in Section 2.4.

The planning problem in this economy is

$$\max_{\{C_t, N_t, K_t, E_t, S_t, X_t, A_t, R_t\}_{t=0}^{\infty}} \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t U(C_t)$$

subject to (1), (3), (4), (14), (16),

$$(17) \quad A_{i,t+1} = G_i(A_{i,t}, X_{i,t}),$$

as well as nonnegativity constraints.

Let $\beta^t \zeta_{i,t}^*$ be the value of the Lagrange multiplier on (17) evaluated at the optimum. It is straightforward to show that the optimal subsidy satisfies

$$(18) \quad 1 + \tau_{i,t}^X = \frac{G_X^*(A_{i,t}^*, X_{i,t}^*)}{U'(C_t^*)} \zeta_{i,t}^*.$$

Thus, the solution to the social planner's problem is implemented by the pollution tax (13) together with the subsidies satisfying (18). This produces an analogue of Propositions 1 and 2 to this environment.

PROPOSITION 3: *Suppose Assumptions 1, 2, and 3 are satisfied and that 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 (11). Moreover, if τ_t are set as in (13) and $\tau_{i,t}^X$ as in (18), competitive equilibrium allocation coincides with the solution to the planning problem.*

It is straightforward to allow more general innovation technologies. The general idea is simply that when there are multiple externalities, a separate Pigouvian tax is required for each of the externality sources. In the above case, there is a climate externality, and hence a carbon tax set according to formula (13); and there is an R&D externality in each energy sector, and hence a subsidy for each R&D activity set according to (18). Note here that there is no presumption that especially high subsidies are needed for R&D toward green energy, so long as carbon is taxed at the optimal rate—unless the externality is stronger in this sector than in other sectors. Especially high subsidies in the energy sector may, however, be called for in a second-best situation when carbon emissions are not taxed at a high enough rate.

Finally, we briefly discuss technical progress that directly reduces the negative effect of carbon emission, such as adaptation to or direct ways of controlling climate change.¹⁵ Such a possibility can be introduced in our setting by letting γ_t in Assumption 2 be endogenous and depend on investments in adaptation and climate control technologies. Once again, the steps toward Proposition 1 remain unchanged, with the optimal tax formula (11) now coming from an evaluation of γ_t at its optimum amount. The interpretation of the formula remains the same, as γ_t still captures the expected damages from carbon emission.

3. COMPLETE CHARACTERIZATION

In this section, we provide a complete characterization for a version of the general multi-sector model discussed above. We will use this model in the quantitative analysis in Section 4. Throughout this section, we assume that Assumptions 1, 2, and 3 hold.

We assume that there are three energy-producing sectors. Sector 1 produces “oil,” which we assume is in finite supply but can be extracted at zero cost. This fossil-fuel source is to be thought of as the oil that is very cheap to extract (relative to its current market price). The constraint $E_{1,t} = R_{1,t} - R_{1,t+1}$ is an accounting equation for oil stocks and will bind at all times. Sectors 2 and 3, which we refer to as the “coal” and “green” sectors, respectively, produce energy using the technologies

$$(19) \quad E_{i,t} = A_{i,t}N_{i,t} \quad \text{for } i = 2, 3.$$

Coal is also in finite supply. However, we will assume that the parameters of the model are such that not all coal will be used up. Hence, it will not have a

¹⁵The range of such technologies is vast, from cheaper air condition units to measures for changing the radiative energy balance of the earth by, for example, emitting particles into the air, controlling cloud formation, or building giant parasols in space.

scarcity rent.¹⁶ Hence, we think of coal as a polar opposite kind of fossil fuel. In order to properly analyze the kind of oil reserves that are only possible to extract at high costs, and may not even be profitable to extract within the near term given current prices, one would want to extend the model to allow further categories of fossil fuel. As discussed in Section 5.3 below, such a generalization would only marginally affect the optimal tax rate on carbon emissions, but it will, of course, give richer, and possibly different, implications for quantity paths.

We chose these technologies to capture the stylized features of different energy sectors in a transparent way. In particular, oil (and natural gas) are relatively cheap to convert to energy, but they rely on exhaustible resources in limited supply. On the other hand, producing energy from coal and green sources is much more expensive, but relies either on an exhaustible resource in much larger supply (coal) or on no such resource (green energy).¹⁷

We also assume a Cobb–Douglas specification for producing final output:

$$(20) \quad Y_t = e^{-\gamma_t(S_t - \delta)} A_{0,t} K_t^\alpha N_{0,t}^{1-\alpha-\nu} E_t^\nu.$$

Here, E_t is an energy composite, defined as

$$(21) \quad E_t = (\kappa_1 E_{1,t}^\rho + \kappa_2 E_{2,t}^\rho + \kappa_3 E_{3,t}^\rho)^{1/\rho},$$

where $\sum_{i=1}^3 \kappa_i = 1$. The parameter $\rho < 1$ determines the elasticity of substitution between different energy sources, and κ measures the relative energy-efficiency of the different energy sources. In reality, coal is a “dirtier” energy source than oil: it produces more carbon emissions per energy unit produced. Since $E_{1,t}$ and $E_{2,t}$ are in the same units (carbon amount emitted) in the model, therefore, in a realistic calibration one should choose $\kappa_1 > \kappa_2$. We assume that $A_{i,t}$ and N_t are exogenous. Finally, we assume that there is full depreciation of capital—having in mind a time period of at least 10 years.

We now characterize the solution to the social planner’s problem. The first-order conditions for C_t and K_t yield

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

This condition together with the feasibility condition

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

¹⁶For some parameter values, this requires a backstop technology emerging at some distant point in the future; see the quantitative discussion below for further details. In the robustness section below, Section 5.3, we also discuss a setting where the backstop is introduced gradually by letting the emission content per unit of coal go to zero smoothly over time.

¹⁷Below, we show evidence that the existing amount of coal is at least an order of magnitude larger than the amount of conventional oil, that is, oil that can be extracted at low cost.

is satisfied if and only if the saving rate is constant at $\alpha\beta$, that is,

$$K_{t+1} = \alpha\beta Y_t \quad \text{for all } t.$$

This implies that the ratio C_t/Y_t is constant at $1 - \alpha\beta$, and Proposition 1 then applies directly. As a result, the optimal unit tax on emissions is given by (11).

The different sources of energy will all be used in strictly positive quantities due to the assumption $\rho < 1$, which amounts to an Inada condition. The first-order conditions for E_t and $E_{1,t}$ imply

$$(22) \quad \frac{\nu\kappa_1}{E_{1,t}^{1-\rho} E_t^\rho} - \hat{\Lambda}_t^s = \beta \mathbb{E}_t \left(\frac{\nu\kappa_1}{E_{1,t+1}^{1-\rho} E_{t+1}^\rho} - \hat{\Lambda}_{t+1}^s \right),$$

where

$$\hat{\Lambda}_t^s \equiv \frac{\Lambda_t^s}{Y_t}.$$

This expression is a version of Hotelling's formula corrected for the exogenous externality term $\hat{\Lambda}_t^s$.

The first-order conditions for $N_{i,t}$ imply

$$(23) \quad A_{2,t} \left(\frac{\nu\kappa_2}{E_{2,t}^{1-\rho} E_t^\rho} - \hat{\Lambda}_t^s \right) = \frac{1 - \alpha - \nu}{N_{0,t}}$$

and

$$(24) \quad A_{3,t} \frac{\nu\kappa_3}{E_{3,t}^{1-\rho} E_t^\rho} = \frac{1 - \alpha - \nu}{N_{0,t}}.$$

Now notice that, for a given value of $E_{1,t}$, the two equations (23) and (24), the labor resource constraint $\sum_{i=0}^3 N_{i,t} = N_t$, and (19) allow us to solve for $E_{2,t}$, $E_{3,t}$, and thus E_t . It is therefore possible, given any $R_{1,0}$, to guess on $E_{1,0}$ and solve for all other values recursively. To see this, we observe that, since all the energy levels in period 0 can be computed as a simple function of $E_{1,0}$, the Hotelling equation (22) delivers an equation in $E_{1,1}$ and E_1 . More to the point, it delivers $E_{1,1}$ as a function of E_1 . It can then be used to solve for all the energy levels in period 1, again using equations (23) and (24), now for period 1. This delivers the entire sequence of energy inputs and, hence, carbon concentrations, output, consumption, and investment. Whether $\sum_{t=0}^{\infty} R_{1,t} = R_{1,0}$ then needs to be verified and the initial guess on $E_{1,0}$ adjusted appropriately. Our construction is thus such that oil is used up asymptotically and coal is not used up and has zero Hotelling rent.

We summarize the properties of the optimal allocations in the following proposition:

PROPOSITION 4: *At the optimum, $C_t/Y_t = 1 - \alpha\beta$ for all t and the optimal (per-unit) tax on emission is given by (11). The optimum path of energy use satisfies (22) and*

$$E_{2,t} = E_t^{-\rho/(1-\rho)} \varepsilon_{2t},$$

$$E_{3,t} = E_t^{-\rho/(1-\rho)} \varepsilon_{3t},$$

where ε_{2t} and ε_{3t} are time-varying and given by

$$\varepsilon_{2t} \equiv \left(\frac{\nu \kappa_2 A_{2,t} N_{0,t}}{1 - \alpha - \nu + \hat{A}_t^\nu A_{2,t} N_{0,t}} \right)^{1/(1-\rho)},$$

$$\varepsilon_{3t} \equiv \left(\frac{\nu \kappa_3 A_{3,t} N_{0,t}}{1 - \alpha - \nu} \right)^{1/(1-\rho)}.$$

Here, $N_{0,t}$ is labor used in the final sector. This quantity is near 1 in quantitative applications—since coal production accounts for a very small part of GDP. The ε 's can also be expressed in terms of exogenous variables and E_t , and thus the model can be very easily simulated.

4. PARAMETER SELECTION

We take each period to be 10 years. Since we can base our analysis on the derived closed-form optimal-tax expression, for our calculation of the marginal externality cost of emissions 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. All the parameters are summarized in Table I; the rest of this section discusses how the choices were made.

TABLE I
CALIBRATION SUMMARY

φ	φ_L	φ_0	α	ν	β	ρ	$10^4 \gamma^H$	$\frac{A_{2,t+1}}{A_{2,t}} = \frac{A_{3,t+1}}{A_{3,t}}$
0.0228	0.2	0.393	0.3	0.04	0.985 ¹⁰	-0.058	2.046	1.02 ¹⁰
$S_0(S_{1,0})$	p	R_0	κ_1	κ_2	$A_{2,0}$	$A_{3,0}$	$10^4 \gamma^L$	$\kappa_1 (\kappa_2)$
802 (684)	0.068	253.8	0.5008	0.08916	7,693	1,311	0.106	0.5429 (0.1015)

4.1. *Preferences and Technology*

We use the assumptions in Section 3: logarithmic preferences, Cobb–Douglas final-goods production, and full depreciation. Logarithmic utility is commonly used in the growth literature. In business-cycle models, curvatures are sometimes assumed to be higher, but with as long a time period as 10 years it is reasonable to have lower curvature, since presumably consumption smoothing can be accomplished more easily then. A depreciation rate of 100% is too high, even for a 10-year period. However, for somewhat lower depreciation rates and some implied movements in saving rates, the optimal-tax formula where saving rates are constant will still be a good approximation. 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 (2012) argued 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.¹⁸ 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 and ν given by 0.3 and 0.04, respectively. When we report the optimal 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 value, such as Stern’s choice of a discount rate of 0.1% per year, or that used by Nordhaus, which is 1.5% per year.

4.1.1. *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 CO₂ in the atmosphere is the key climate driver. We therefore need to specify how emissions affect the atmospheric CO₂ 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 was 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

¹⁸Hassler, Krusell, and Olovsson (2012) also showed 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.

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 reaches 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. The most important 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, Eby, Brovkin, Ridgwell, Cao, Mikolajewicz, and Tokos (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. 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. 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 that we believe better capture some of the long-run impacts from current emissions.¹⁹ We show below that our formulation also has quantitative properties

¹⁹Recently, Allen, Frame, Huntingford, Jones, Lowe, Meinshausen, and Meinshausen (2009) have noted, by studying an ensemble of simulations from carbon-cycle models, that the temperature effects can be captured very well by total cumulated emissions both at short- and long-run horizons and that such a formulation can be seen as an improvement over the kind of carbon concentration-based structures studied by Nordhaus. Our formulation here goes part of, though

that are rather different. In fact, our formulation leads to significantly larger effects of human emissions on the climate.

Our three-parameter formula amounts to $1 - d_s = \varphi_L + (1 - \varphi_L)\varphi_0\varphi^s$, where $1 - d_s$ denotes the total fraction of a unit emitted at time 0 that is left in the atmosphere at time s . We calibrate φ_L , φ_0 , and φ as follows. We set φ_L to 0.2, according to the estimate in the 2007 IPCC report that about 20% of any emission pulse will stay in the atmosphere for thousands of years.²⁰ According to Archer (2005), furthermore, the excess carbon that does not stay in the atmosphere “forever” has a mean lifetime of about 300 years. Thus, we impose $(1 - \varphi)^{30} = 0.5$, yielding $\varphi = 0.0228$. Third, again according to the 2007 IPCC report, about half of the CO₂ pulse to the atmosphere is removed after a time scale of 30 years. This implies $d_2 = 0.5$ in our formula, and $1 - \frac{1}{2} = 0.2 + 0.8\varphi_0(1 - 0.0228)^2$ thus gives $\varphi_0 = 0.393$.

Finally, we need to provide an initial condition for carbon concentration. Our process is equivalent to a recursive vector representation where S_1 denotes carbon that remains in the atmosphere forever, and S_2 , carbon that depreciates at rate φ . These assumptions imply that $S_{1,t} = S_{1,t-1} + \varphi_L E_t^f$ and that $S_{2,t} = \varphi S_{2,t-1} + \varphi_0(1 - \varphi_L)E_t^f$, with $S_t = S_{1,t} + S_{2,t}$. We calibrate so that time-0 (i.e., year-2000) carbon equals 802, with the division $S_1 = 684$ and $S_2 = 118$.²¹

4.2. 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). Our damage function has carbon concentration, S , as its argument, whereas other models of course express damages as a function of a climate indicator, such as global temperature. This mapping is typically modeled as convex. Our taking S as an input should be viewed as a composition of the typical damage function, with temperature as an argument, and another function mapping carbon concentration into temperature. Typical approximations used in climate science make the latter mapping concave—indeed logarithmic—so it is not clear whether the overall function mapping S into damages should be convex or concave. We chose the exponential form because it turns out to be a very good approximation of the composition of the two mappings used by

not all, the way toward such a formulation by obtaining stronger long-run impacts on temperature from current emissions. It would be interesting to alter the current model to allow Allen et al.’s formulation and use numerical methods for studying the implications for optimal tax rates and quantity paths.

²⁰Archer (2005) estimated an even higher fraction: 0.25.

²¹Note that S_1 here includes the pre-industrial stock of 581 plus 20% of accumulation emissions.

Nordhaus and many others. Nordhaus's damage function of global temperature is 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 pre-industrial 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).

Turning to the mapping from S to T , the standard assumption in the literature (say, as used in RICE) is to let the steady-state global mean temperature be a logarithmic function of the stock of atmospheric carbon:

$$(25) \quad T_t = T(S_t) = \lambda \log\left(\frac{S_t}{\bar{S}}\right) / \log 2,$$

where $\bar{S} = 581$ GtC (gigaton of carbon) is the pre-industrial atmospheric CO₂ concentration. A standard value for the climate sensitivity parameter λ here is 3.0 degrees Celsius. That means 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 Figure 1, we show the composition of the S -to- T and T -to-net-of-damages mappings, that is, $1 - D(T(S))$, as calibrated by Nordhaus (dashed) together




FIGURE 1.—Net-of-damages function $1 - D(T(S))$; Nordhaus (dashed) and exponential (solid).

with the net-of-damages 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 pre-industrial 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.²² The composition implied by Nordhaus's formulation is first concave, then convex; our function is approximately linear over this range. Overall, the two curves are quite close and we thus conclude that our exponential approximation is rather reasonable.

To incorporate uncertainty into 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, all uncertainty is resolved and it turns out that γ will either be equal to γ^H or 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 $p\gamma^H + (1 - p)\gamma^L \equiv \bar{\gamma}$.²³

What are the sources of the specific damage parameters? When calibrating the damage function, Nordhaus (2008) used a bottom-up approach by collecting a large number of studies on various effects of global warming. Some of these are positive, that is, warming is beneficial, but most are negative. By adding these estimates up, he arrived at an estimate that a 2.5-degree Celsius heating yields a global (output-weighted) loss of 0.48% of GDP.²⁴ Furthermore, he argued, based on survey evidence, that with a probability of 6.8% the damages from heating of 6 degrees Celsius are catastrophically large, defined as a loss of 30% of GDP. Nordhaus, moreover, calculated the willingness to pay for such a risk and added 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 (25), we find that a 2.5- and a 6-degree heating occurs if S_t equals 1,035 and 2,324, respectively. We thus calibrate γ^L to solve

$$e^{-\gamma^L(1,035-581)} = 0.9952$$

and γ^H to solve

$$e^{-\gamma^H(2,324-581)} = 0.70,$$

²²For a discussion of the estimates of the total stocks, see Section 4.3 below.

²³Pizer (1998) is an early study using uncertainty; Kelly and Kolstad (1999) also studied Bayesian learning.

²⁴Reduced-form estimates, for example, those in Mendelsohn, Nordhaus, and Shaw (1994) exploiting interregional differences using cross-sectional data on temperatures and output from countries and regions within countries, suggest damages that are higher but 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.

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} .

4.3. Energy

For the elasticity of substitution between the three sources of energy, we use a metastudy (Stern (2012)) of 47 studies of interfuel substitution. The unweighted mean of the oil-coal, oil-electricity, and coal-electricity elasticities is 0.95. The elasticity in what Stern defines as long-run dynamic elasticities is 0.72. These elasticities imply $\rho = -0.058$ and -0.390 , respectively. We use the former as a benchmark value. In addition, we will study a much higher elasticity, by setting $\rho = 0.5$, implying an elasticity of 2.

Another key parameter is the size of the oil reserve. BP (2010) reported that proven global reserves of oil amount to 181.7 gigaton. 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 estimated global reserves taking into account technical progress, ending up at an estimate of over 5,000 gigaton of oil equivalents.²⁵ Of this, around 16% is oil, that is, 800 gigaton. We take as a benchmark that the existing stock of oil is 300 gigaton, that is, somewhere well within the range of these two estimates.

To express fossil fuel in units of carbon content, we set the carbon content in crude oil to 846 KgC/ton oil. For coal, we set it to the carbon content of anthracite at 716 KgC/ton coal.²⁶

We have assumed that the scarcity rent for coal is negligible. This appears reasonable because, as noted, the reserves of coal are very large compared to those for oil. Rogner (1997) estimated that the coal supply is enough for several hundreds of years of consumption at current levels.

To calibrate κ_1 and κ_2 , we then use relative prices of oil to coal and oil to renewable energy, given by

$$(26) \quad \frac{\kappa_1}{\kappa_2} \left(\frac{E_{1t}}{E_{2t}} \right)^{\rho-1} \quad \text{and} \quad \frac{\kappa_1}{1 - \kappa_1 - \kappa_2} \left(\frac{E_{1t}}{E_{3t}} \right)^{\rho-1},$$

respectively. We use the average price of Brent oil over the period 2005–2009, which was \$70 per barrel (BP (2010)). One barrel is 7.33 metric tons. Using

²⁵The difference in energy content between natural gas, oil, and various grades of coal is accounted for by expressing quantities in oil equivalents.

²⁶IPCC (2006, Tables 1.2–1.3).

the carbon content of 84.6%, the oil price per ton of carbon is \$606.5. The 5-year average of coal price between 2005 and 2009 is \$74/ton. Using the carbon content of 71.6%, we obtain a price of \$103.35 per ton of carbon.^{27,28} Thus, the relative price of oil and coal in units of carbon content is 5.87.

It is a little more difficult to find a representative price of renewables since this is a quite heterogeneous source of energy. We, however, take unity as a reasonable value of the current relative price between green energy and oil. Finally, we use data on global energy consumption from IEA (2010).²⁹ Using these numbers and the benchmark value $\rho = -0.058$ in the expressions for the relative prices in (26) gives $\kappa_1 = 0.5008$, and $\kappa_2 = 0.08916$.

The parameter $A_{2,t}$, which determines the cost of extracting coal, is calibrated to an average extraction cost of \$43 per ton of coal (as reported by IEA (2010, p. 212)). Thus, a ton of carbon costs \$43/0.716, since the carbon content of coal is 0.716. In the model, the cost of extracting a ton of carbon in the form of coal is given by $\frac{w_t}{A_{2,t}}$, where w_t is the wage. We normalize the total labor supply to unity. The current share of world labor used in coal extraction and green energy production is very close to zero. Using the approximation that it is literally zero, the wage is given by $w_t = (1 - \alpha - \nu)Y_t$.³⁰ Using a world GDP of \$700 trillion per decade, we thus have the cost of a gigaton of carbon (our model unit) as $w_t/A_{2,t} = (1 - \alpha - \nu)Y_t/A_{2,t}$, which becomes $43 \cdot 10^9/0.716 = 0.66 \cdot 700 \cdot 10^{12}/A_{2,0}$. This yields $A_{2,0} = 7,693$. Thus, to extract one gigaton of carbon in the form of coal, a share $\frac{1}{7,693}$ of the labor supply of a decade is needed. The calibration of $A_{3,0}$ is derived by noting that $A_{3,0}/A_{2,0}$ is equal to relative price between coal and green energy, implying $A_{3,0} = 7,693/5.87 = 1,311$, since we calibrate the prices of oil and green to be equal and the relative price of oil in terms of coal to be 5.87.

We finally assume that there is growth in both extraction efficiency and the efficiency of green technologies, so that $A_{2,t}$ and $A_{3,t}$ both grow at a rate of 2% per year.³¹

²⁷The numbers refer to U.S. Central Appalachian coal. Source: BP (2010).

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

²⁹Primary global energy demand in 2008 was 3.315 Gtoe (gigaton of oil equivalents) of coal, 4.059 of oil, 2.596 of gas, and $0.712 + 0.276 + 1.314 = 2.302$ of nuclear, hydro, and biomass/waste/other renewables. Using the IPCC tables quoted above, we find that the ratio of energy per ton between oil and anthracite is $\frac{42.3}{26.7} = 1.58$, so one ton of oil equivalents is 1.58 tons of coal. We express the amount of oil and coal in carbon units by multiplying by the carbon contents 84.6 and 71.6%, respectively. Source: IEA (2010).

³⁰This is thus a slight overestimate, as labor used in the production of final output in the model is not 1; it is a little over 0.97.

³¹The stated assumptions do not imply that coal use goes to zero; hence, coal would have scarcity value. If, however, a competitive close and renewable substitute for coal is invented over the next couple of hundred years, coal will have zero scarcity value. Such a scenario seems rather likely, and we prefer it over one where coal is exhausted and has a positive scarcity rent.




FIGURE 2.—Optimal tax rates in current dollars per ton of emitted fossil carbon versus yearly subjective discount rate.

5. QUANTITATIVE RESULTS

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

5.1. *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. We calculate the optimal taxes both before and after we have learned the long-run value of γ . We use (12) 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.

To relate our optimal tax to available estimates, consider the much-discussed policy proposals in Nordhaus (2008) and 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 used a rate of 1.5% per year, mostly based on market measures. Stern, who added a “moral” concern for future generations, used the much lower rate of 0.1% per year. In Figure 2, the solid line is the ex ante tax before the uncertainty is realized, and 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 significantly larger optimal tax than computed in both these studies. This difference is due to a number of factors. One is that our depreciation structure for carbon in the atmosphere, as calibrated, 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 5.3. Furthermore, we see that the consequences of learning are dramatic. With a discount rate of 1.5%, the optimal tax rate if damages turn out to be moderate

is \$25.3/ton, but \$489/ton if damages turn out catastrophic. For the low discount rate, the corresponding values are \$221/ton and a whopping \$4,263/ton.

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, in particular, one uses Stern's discount rate, then it follows that the laissez-faire equilibrium generates too little saving (and too high a market interest rate), calling for 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 modeled the demand for environmental goods explicitly. They assumed 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, 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 there is significant uncertainty about both the total current (and yet-to-be-discovered) stocks as well as technological developments that one would need to worry greatly about possible quantity misjudgments.

5.2. Implications for the Future: Climate, Damages, and Output

Given the assumptions made in Section 4.3 about fossil fuel reserves, in addition to the assumptions underlying the optimal tax rates, we can now generate quantity paths—optimal and suboptimal ones—for the different energy sources and thus also for climate and damages. Solving the model is very easy due to the fact that the saving rate is constant and that the law of motion for energy use can be easily simulated with a guess only on initial energy use. The results reported refer to the case where the damage parameter γ remains at its expected level throughout time (significant adjustments, of course, 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. The model's predictions of current fossil use under laissez-faire are close to actual use. For coal, the model predicts a yearly use of 4.5 GtC




FIGURE 3.—Fossil fuel use: optimum versus laissez-faire.

during the coming decade, compared to an actual value of 3.8 GtC. For oil, the model predicts a yearly use of 3.6 GtC, which is close to the actual value for 2008 of 3.4 GtC.

Comparing the two outcomes, we find that optimal policy leads to a much smaller use of fossil fuel. The no-tax market economy would have a continuous increase in fossil fuel use, whereas optimal taxation would imply an almost flat consumption profile.

A notable feature of our results is that the difference between the two paths for fossil fuel is almost entirely driven by differences in coal use. In Figures 4 and 5, we plot coal use and oil use in the optimal versus the laissez-faire allocations. Although the carbon tax is the same for oil and coal, its effects are very different. Coal grows quickly in the laissez-faire allocation but very slowly if optimal taxes are introduced: the tax reduces coal use immediately by 46%, and 100 years from now, the laissez-faire coal use is seven times higher than optimally. In 100 years, accumulated coal use will have risen to 340 GtC in the optimal allocation and 1200 GtC in the laissez-faire allocation; after 200 years, the accumulated optimal outtake will have risen to a little below 900 GtC, and under laissez-faire coal use increases quickly, leading to a scarcity rent unless a




FIGURE 4.—Coal use: optimum versus laissez-faire.




FIGURE 5.—Oil use: optimum versus laissez-faire.

backstop appears some time within the coming 200 years.³² The two curves for oil, on the other hand, are very close to each other: they never differ by more than about 6%. The optimal and laissez-faire paths for green energy are even more similar, since they are not affected by taxes in any of the regimes (the difference is never larger than 1.1%).

The paths for total damages are plotted in Figure 6. There are significant, though not enormous, gains from raising taxes to the optimal level. The gains in the short run are small, but grow over time. One hundred years from now, damages are at 2.2% of GDP in the laissez-faire regime rather than 1.1% in the optimal allocation. At the end of the simulation period (2200), damages in laissez-faire have grown to over 10%, while they are only 1.5% in the optimal allocation.

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 7. Under laissez-faire, temperatures will have increased by 4.4 degrees




FIGURE 6.—Total damages as a percent of GDP: optimum versus laissez-faire.

³²In the robustness analysis conducted, we verified that optimal coal use over the next 50 years is only marginally affected by the precise timing of the emergence of the backstop, so long as it appears in 100 years or later.




FIGURE 7.—Increases in global temperature: optimum versus laissez-faire.

Celsius 100 years from now, while the optimal use of fossil fuels leads to a heating of 2.6 degrees, that is, about half. At the end of the simulation period, the climate on earth is almost 10 degrees Celsius warmer without policy intervention, while the optimal tax limits heating to about 3 degrees. Note, however, that these temperature increases are measured relative to the pre-industrial climate; relative to the model's prediction for the current temperature, the increases are about $11/2$ degrees smaller.³³

Finally, we show the evolution of relative (net-of-damage) production of final-good output (GDP) in Figure 8. The optimal allocation involves negligible short-run losses in GDP. Output net of damages in the optimal allocation exceeds that in laissez-faire already from 2020. To understand this finding, recall (i) that using less coal implies less labor used in coal energy production and (ii) that oil consumption is not much affected by the optimal tax. After 100 years, GDP net of damages is 2.5% higher in the optimal allocation, and in year 2200, the difference is almost 15%.




FIGURE 8.—Net output: optimum versus laissez-faire.

³³Standard models of climate change tend to overpredict the heating relative to current temperatures. Our model overpredicts the current temperature by around 1 degree Celsius. A common explanation for this is that anthropogenic aerosols lead to a cooling effect, temporarily masking the full impact of greenhouse gases (see Schwartz, Charlson, Kahn, Ogren, and Rodhe (2010)).

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 are much higher (e.g., because feedback effects turn out stronger than expected), 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.

5.3. Robustness

We now discuss how our results are affected by changes in the benchmark model. We focus first on the formula for the optimal tax, equation (11), which we argue is quite robust to a number of generalizations of our benchmark formulation. We then turn to the results on the future path of quantities, including the path for climate.

5.3.1. The Tax Formula

Our exact tax formula relies on the assumption of a constant saving rate. We demonstrated in the previous section that this assumption can be justified in a model with logarithmic utility, Cobb–Douglas production, and 100% depreciation of capital, presuming that energy is produced either for free (“oil”) or at a constant returns to labor (“coal/wind”). We now discuss departures from these assumptions, beginning with the depreciation-rate assumption.

The depreciation rate of capital does not appear directly in the optimal-tax formula; it is straightforward that the formula can be written

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

in this case. Thus, a lowering of the depreciation rate to what we consider a reasonable value, 0.65 per decade, will only change the time path of $\frac{1-s_t}{1-s_{t+j}}$ from 1 to different values. Figure 9 depicts the carbon tax-GDP ratio for a 65% depreciation rate of capital under a number of different assumptions about the growth rate of TFP (in this figure, all other parameters are like in the benchmark calibration). In order to facilitate the comparison with our benchmark, we report the tax-GDP ratio relative to the predictions of the formula. The TFP growth rates considered include 0 and 1.5% per year as well as a path taken from Nordhaus (2008) where the growth rate of TFP is assumed to decline slowly over time from 1.4% to 0.3% per year. We adjust the initial value of capital in some of the cases (graphs labeled “recal” in Figure 9) in order to keep the net initial return on capital the same as in the benchmark. The experiments




FIGURE 9.—Optimal tax/GDP ratio, relative to case with constant saving rate.

conducted are reported on in detail—including how they are computed—in a note available online: Barrage (2014).

Two things stand out from the figure. First, the transition dynamics for the optimal tax rate are negligible (recall that the rate is constant over time under our benchmark assumptions). Second, after the transition is essentially over, the tax rate is the same as in the exact formula.

Let us now turn to variations in the utility function and consider the more general power-function class (i.e., a constant relative risk aversion/elasticity of intertemporal substitution). It is again straightforward to generalize the tax formula; it becomes

$$(27) \quad \hat{A}_t^s = \mathbb{E}_t \sum_{j=0}^{\infty} \hat{\beta}_j \frac{1-s_t}{1-s_{t+j}} \gamma_{t+j} (1-d_j),$$

where

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

and σ is the coefficient of relative risk aversion (curvature increases in σ). Clearly, a change of σ away from 1 implies: (i) different de facto discounting, as given by β_j , since marginal utility will shrink at a rate that is not equal to the rate of consumption growth; and (ii) transitional dynamics in the saving rate. We consider values of σ of 0.5, 1.5, and 2. If the curvature is higher than for logarithmic utility (above 1), a high growth rate implies a higher effective discounting of future damages, as is evident from the expression for $\hat{\beta}$ above, resulting in a lower optimal tax-GDP ratio. With significant growth in consumption, a change in σ between, say, 1 and 2 will influence the tax rate significantly because of the discounting effect. However, our formula is still useful: the computations in Barrage (2014) show that the simple and highly intuitive generalization of it provides a very good approximation to the optimal tax rate. The idea behind the approximation is to assume a constant consump-




FIGURE 10.—Optimal tax/GDP ratio, relative to approximation formula.

tion growth at net rate g , so that $\hat{\beta}_j$ becomes $(\beta(1+g)^{1-\sigma})^j$, thus delivering a closed-form expression

$$\gamma \left(\frac{\phi_L}{1 - \beta(1+g)^{1-\sigma}} + \frac{(1 - \phi_L)\phi_0}{1 - (1 - \phi)\beta(1+g)^{1-\sigma}} \right).$$

To operationalize the formula in terms of exogenous parameters, we set g equal to the growth rate of TFP times $(1 - \alpha)^{-1}$. Figure 10 shows the optimal tax-GDP ratio relative to the approximation formula for the three different cases of risk aversion, assuming a TFP growth rate of 1.3% per year (producing an approximated growth rate of consumption of 1.9% per year).

As we see, the deviations are fairly modest. It should also be reiterated that the formula gives quite a large response in the tax-GDP ratio to σ at a growth rate of output as high as 1.9%, since the effective discount factor changes substantially with σ .³⁴ For example, with $\sigma = 0.5$, the tax rate doubles compared to the logarithmic case, and with $\sigma = 2$, it is roughly cut in half.

Consider now different assumptions in the production technology for output and energy. It is important to first note, with reference to equation (27), that there are no direct effects of changing production technologies on the optimal-tax formula: the effects occur entirely through changes in the saving rate over time and, to the extent the utility function does not have logarithmic curvature, through making the consumption growth rate nonconstant. Suppose, then, that 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 at least on short horizons.³⁵ Consider extreme complementarity within the class of constant substitution elasticity: output is Leontief in $AK^\alpha N^{1-\alpha}$ and E . Again, there is a balanced growth path (whose properties depend on how fast A and A_E grow). It appears reasonable to assume that K (or

³⁴As the growth rate of output approaches zero, the approximated optimal tax becomes invariant to σ .

³⁵As illustrated in Hassler, Krusell, and Olovsson (2012), 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).

A) is low initially, delivering increasing energy use over time, something which we have observed over a long period of time. This implies high initial saving rates, but as shown in Hassler, Krusell, and Olovsson (2012), 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, but since the share of fossil fuel is less than 5% of total output, extensions in this direction will only change the optimal-taxation results very 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

$$(28) \quad \hat{\Lambda}_t^s = (1 - s_t) \mathbb{E}_t \sum_{j=0}^{\infty} \beta^j V'(S_{t+j}) (1 - d_j),$$

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 that 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 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.

5.3.2. *Quantity Predictions*

We now turn to the robustness of the results on quantities: how would our predictions for output, the temperature, energy use, and so on change if one considered the generalizations just discussed? These predictions are much more sensitive to our assumptions. Unlike our basic tax formula, they require knowledge of how all variables in the model develop: all exogenous parameters matter. For example, how energy use evolves 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. This does not, of course, mean that all parameters of the model have large effects on quantities. In particular, in the robustness exercises for the tax formula discussed above, the obtained quantity predictions are all quite similar to those obtained for the benchmark case at least for the first 100 years.³⁶

One parameter that does matter greatly for the quantity predictions is the elasticity of substitution between the different sources of energy. We illustrate this by considering a much higher elasticity. In the benchmark, the elasticity is 0.95. Let us instead consider an elasticity of 2. In this case, the introduction of a carbon tax is much more urgent since the difference between coal use in *laissez-faire* and in the optimal allocation is rather dramatic. One hundred years from now, coal use in *laissez-faire* will have increased by almost a factor of 20 relative to today. In the optimal allocation, in contrast, coal use will always be lower than it is today. In the optimal allocation, moreover, output net of damages in 2110 are 4.8% higher in the optimal allocation. After 200 years, the difference is 40%. In fact, with a high elasticity, the optimal policy implies that the temperature starts to decline in the middle of the next century by making fossil fuel use negligible.³⁷

The intuition for the sensitivity of the results to the elasticity of substitution can be understood as follows. If energy sources are highly substitutable, coal can easily substitute for oil, which makes the *laissez-faire* allocation involve significant coal use. On the other hand, the optimal tax (which is independent of the elasticity) has a much stronger impact on the allocation when energy sources are highly substitutable. Thus, the social gains from introducing the optimal tax—or the costs of not doing so—are much larger in the case of high substitutability. Our results here are in line with Acemoglu et al. (2012), who

³⁶Since a gradual lowering of the carbon content of coal-burning emissions is considered in many of the robustness exercises in Barrage (2014), the predictions for temperature are more favorable in the longer term.

³⁷With a high elasticity, assumptions about the technology trends also become more critical.

used a high degree of substitutability between clean and dirty energy. There, if one of the sources is more effective at energy production, it will dominate the market.

5.3.3. *Comparisons With DICE and RICE*

Before concluding, let us relate our results to the state-of-the-art analysis conducted by Nordhaus (2007).³⁸ In particular, we will compare to his calculation of the optimal CO₂ tax.³⁹ Nordhaus reported an optimal tax of \$27 for 2005 that should rise to \$42 in 2015. Nordhaus used 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 used a utility function with higher curvature (an elasticity of intertemporal substitution of 1/2). He calibrated the subjective discount rate to yield a net return of capital of 5.5%. For logarithmic utility, which we use, he reported 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 difference in utility functions used in our studies, it is perhaps more reasonable to make comparisons if we adopted a subjective discounting of 3%. For this discount rate, our formula yields \$32, which brings the two sets of results even closer together. Alternatively, the use of our approximation formula (28) delivers \$32 as well if the growth rate is set to 1.5% per year and the elasticity is 1/2. However, a closer inspection implies that there are a number of countervailing effects behind this similarity.

First, we deal with uncertainty in different ways. Nordhaus used a “certainty-equivalent damage function,” that is, he optimized 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 of 2.

³⁸A 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 compared to 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, Davidson, Dowlabathi, Edmonds, Grubb, Parson, and Fankhauser (1996) gave a detailed assessment and Weyant (2000) summarized the main commonalities and differences behind the most widely used models. A more recent comprehensive analysis (Clarke, Edmonds, Krey, Richels, Rose, and Tavoni (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 discussed 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, and (3) assessment of such impacts depending on when various regions would participate in emissions mitigation. For the U.S. economy, Jorgenson, Goettle, Mun, and Wilcoxon (2008) examined 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.

³⁹Details of this comparison are available upon request.

Second, there are important differences in the modeling of the carbon cycle. Specifically, while we assume that almost half of the emissions are absorbed by the biosphere and the upper layers of the ocean within 10 years, Nordhaus assumed 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 of 1.5).

Nordhaus, finally, used 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 CO_2 concentration on temperature.⁴⁰ Of course, this biases our estimate upward, 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 CO_2 emissions in a way that follows those Nordhaus assumed rather precisely. 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 nonlinear, it does not incorporate so-called threshold effects or “tipping points.” These refer to literal discontinuities (or very strong nonlinearities) in some of the model relationships, implying sharply changing local dynamics and steady-state multiplicity. For example, 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. In our model, this kind of non-linearity could appear 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. Non-convexities can be rather straightforwardly analyzed using our setting, with some more reliance of numerical methods. The real challenge is a quantitative one: at what levels of S does a nonlinearity appear, and what is its nature, including its dynamics? A tipping point could also occur in the model of carbon depreciation: if the temperature becomes sufficiently high, some carbon reservoirs may switch rather abruptly from net absorption to net

⁴⁰Recent work by Roe and Bauman (2011) showed 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.

emission. We follow Nordhaus, however, in not explicitly incorporating strong nonlinearities. 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 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 formulate 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 show 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 evaluate this formula quantitatively and find results that are about twice the size of those put forth by Nordhaus and Boyer (2000). The differences between our findings are due to a variety of differences in assumption, for example, 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) arrived 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 the size 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 argue 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.⁴¹

To relate our estimates to actually implemented carbon taxes, consider Sweden, where the tax on private consumption of carbon actually exceeds \$600

⁴¹It 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.

per ton.⁴² Though industrial carbon use is subsidized relative to private consumption in Sweden, these rates are very high from a worldwide 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 2008–2009, the price has hovered around 15 Euro per ton CO₂, at an exchange rate of 1.4 dollar per Euro corresponding to US\$77 per ton carbon.^{43,44} 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 compute paths for our key variables for a laissez-faire market economy and compare them to the optimal outcome. In the optimal outcome, coal extraction is much lower than in laissez-faire. The use of oil and green energy is, however, almost identical in the two allocations. The temperature increase will therefore be much smaller if the optimal tax is introduced. Total damages in laissez-faire will rise over time and amount to over 2% of GDP 100 years from now and close to 10% in the year 2200. In the optimal allocation, in contrast, they grow only slowly to reach 1.4% 200 years from now. 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 needs information neither about the precise sources of energy—fossil or not—nor

⁴²In 2010, the tax was 1.05SEK per kilo of emitted CO₂ (Swedish Tax Agency (2010)). A kilo of CO₂ contains 0.27 kilos of carbon. Using an exchange rate of 6.30SEK/\$, one obtains a tax of \$617.28/tC.

⁴³The price of EU emissions allowances can be found on the home page of the European Energy Exchange, <http://www.eex.com/>.

⁴⁴A ton of CO₂ contains 0.273 tons of carbon, implying a conversion factor of $0.273^{-1} = 3.66$.

about the future paths of population growth and technical change (energy-specific or other). Quantity restrictions that would implement the optimum, that is, 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 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. (2012) showed 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.⁴⁵ 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

⁴⁵See also Saint-Paul (2002, 2007); in the latter paper, it was argued that optimal subsidies should be higher for environmental innovation even if Pigou tax on emissions is used.

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 (2012) for some productivity accounting and an examination of endogenous technology, and Gars, Golosov, and Tsyvinski (2009) for a model with a backstop technology).

REFERENCES

- ACEMOGLU, D. (2009): *Introduction to Modern Economic Growth*. Princeton, NJ: Princeton University Press. [53]
- ACEMOGLU, D., P. AGHION, L. BURSZTYN, AND D. HEMOUS (2012): "The Environment and Directed Technical Change," *American Economic Review*, 102 (1), 131–166. [44-46,79,84]
- ALLEN, M., D. FRAME, C. HUNTINGFORD, C. JONES, J. LOWE, M. MEINSHAUSEN, AND N. MEINSHAUSEN (2009): "Warming Caused by Cumulative Carbon Emissions Towards the Trillionth Tonne," *Nature*, 458, 1163–1166. [64]
- AMIGUES, J.-P., G. LAFFORGUE, AND M. MOREAUX (2012): "Optimal Timing of Carbon Capture Policies Under Alternative CCS Cost Functions," Working Paper 12-318, Toulouse School of Economics. [45]
- ANGELETOS, M. (2007): "Uninsured Idiosyncratic Investment Risk and Aggregate Saving," *Review of Economic Dynamics*, 10 (1), 1–30. [79]
- ANGELETOS, M., AND L. CALVET (2006): "Idiosyncratic Production Risk, Growth and the Business Cycle," *Journal of Monetary Economics*, 53 (6), 1095–1115. [78]
- ARCHER, D. (2005): "The Fate of Fossil Fuel CO₂ in Geologic Time," *Journal of Geophysical Research*, 110, C09S05. [44,64,65]
- ARCHER, D., M. EBY, V. BROVKIN, A. RIDGWELL, L. CAO, U. MIKOLAJEWICZ, AND K. TOKOS (2009): "Atmospheric Lifetime of Fossil Fuel Carbon Dioxide," *Annual Review of Earth and Planetary Sciences*, 37, 117–134. [64]
- BARRAGE, L. (2014): "Sensitivity Analysis for Golosov, Hassler, Krusell, and Tsyvinski (2014): 'Optimal Taxes on Fossil Fuel in General Equilibrium'," *Econometrica Supplemental Material*, 82, http://www.econometricsociety.org/ecta/supmat/10217_extensions.pdf. [46,76,79]
- BOVENBERG, L., AND S. SMULDERS (1995): "Environmental Quality and Pollution-Augmenting Technological Change in a Two-Sector Endogenous Growth Model," *Journal of Public Economics*, 57, 369–391. [46]
- BOVENBERG, L., AND S. SMULDERS (1996): "Transitional Impacts of Environmental Policy in an Endogenous Growth Model," *International Economic Review*, 37 (4), 861–893. [46]
- BP (2010): "BP Statistical Review of World Energy," June, available at <http://bp.com/statisticalreview>. [68,69]
- CLARKE, L., J. EDMONDS, V. KREY, R. RICHEL, S. ROSE, AND M. TAVONI (2009): "International Climate Policy Architectures: Overview of the EMF 22 International Scenarios," *Energy Economics*, 31, S64–S81. [80]
- COVAS, F. (2006): "Uninsured Idiosyncratic Production Risk With Borrowing Constraints," *Journal of Economic Dynamics and Control*, 30 (11), 2167–2190. [79]
- DASGUPTA, P., AND G. HEAL (1974): "The Optimal Depletion of Exhaustable Resources," *Review of Economic Studies*, 41, 3–28. [44,46,63]
- EYCKMANS, J., AND H. TULKENS (2003): "Simulating Coalitionally Stable Burden Sharing Agreements for the Climate Change Problem," *Resource and Energy Economics*, 25, 299–327. [45]
- GARS, J., AND J. HIERONYMUS (2012): "The Marine Carbon Cycle in an Integrated Assessment Model," Report, Stockholm University. [51]
- GARS, J., M. GOLOSOV, AND A. TSYVINSKI (2009): "Carbon Taxing and Alternative Energy," Report, Yale University. [85]

- GERLAGH, R. (2006): "ITC in a Global Growth-Climate Model With CCS: The Value of Induced Technical Change for Climate Stabilization," *The Energy Journal*, 27 (Special I), 223–240. [46]
- GERLAGH, R., AND M. LISKI (2012): "Carbon Prices for the Next Thousand Years," Working Paper 3855, CESifo. [44]
- GOULDER, L., AND K. MATHAI (2000): "Optimal CO₂ Abatement in the Presence of Induced Technological Change," *Journal of Environmental Economics and Management*, 39, 1–38. [45]
- GRIMAUD, A., G. LAFFORGUE, AND B. MAGNÉ (2011): "Climate Change Mitigation Options and Directed Technical Change: A Decentralized Equilibrium Analysis," *Resource and Energy Economics*, 33 (4), 938–962. [46]
- HASSLER, J., P. KRUSELL, AND C. OLOVSSON (2012): "Energy-Saving Technical Change," Working Paper 18456, NBER. [63,77,78,85]
- HERFINDAHL, O. (1967): "Depletion and Economic Theory," in *Extractive Resources and Taxation*, ed. by M. Gaffney. Madison, WI: University of Wisconsin Press. [46]
- HOEL, M. (2009): "Climate Change and Carbon Tax Expectations," Working Paper 2966, CESifo. [45,46]
- HOPE, C. W. (2008): "Discount Rates, Equity Weights and the Social Cost of Carbon," *Energy Economics*, 30 (3), 1011–1019. [45]
- HOTELLING, H. (1931): "The Economics of Exhaustible Resources," *Journal of Political Economy*, 39 (2), 137–175. [45]
- IEA (INTERNATIONAL ENERGY AGENCY) (2010): "World Energy Outlook," OECD/IEA Paris. [69]
- IPCC (INTERGOVERNMENTAL PANEL ON CLIMATE CHANGE) (2006): "Guidelines for National Greenhouse Gas Inventories, Vol. 2 Energy," IPCC. [68]
- IVERSON, T. (2012): "Optimal Carbon Taxes With Non-Constant Time Preference," Working Paper. [44]
- JORGENSEN, D., R. GOETTLE, M. HO, AND P. WILCOXEN (2008): "The Economic Costs of a Market-Based Climate Policy," Working Paper, Pew Center on Global Climate Change. [80]
- KARP, L. (2005): "Global Warming and Hyperbolic Discounting," *Journal of Public Economics*, 89, 261–282. [44]
- KELLY, D., AND C. KOLSTAD (1999): "Bayesian Learning, Growth, and Pollution," *Journal of Economic Dynamics and Control*, 23 (4), 491–518. [67]
- KEMP, M., AND N. VAN LONG (1980): "On Two Folk Theorems Concerning the Extraction of Exhaustible Resources," *Econometrica*, 48 (3), 663–673. [49]
- KRUSELL, P., AND A. SMITH (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106, 867–896. [78]
- (2006): "Quantitative Macroeconomic Models With Heterogeneous Agents," in *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress*. Econometric Society Monographs, Vol. 41, ed. by R. Blundell, W. Newey, and T. Persson. Skatteverket: Cambridge University Press, 298–340. [78]
- (2009): "Macroeconomics and Global Climate Change: Transition for a Many-Region Economy," Working Paper. [85]
- LEACH, A. (2007): "The Welfare Implications of Climate Change Policy," *Journal of Economic Dynamics and Control*, 31, 151–165. [80]
- MCKIBBIN, W., AND P. WILCOXEN (1999): "The Theoretical and Empirical Structure of the G-Cubed Model," *Economic Modelling*, 16 (1), 123–148. [80]
- MENDELSON, R., W. NORDHAUS, AND D. G. SHAW (1994): "The Impact of Global Warming on Agriculture: A Ricardian Approach," *American Economic Review*, 84 (4), 753–771. [67]
- NORDHAUS, W. (2007): "To Tax or Not to Tax: The Case for a Carbon Tax," *Review of Environmental Economics and Policy*, 1 (1), 26–44. [65,80]
- (2008): *A Question of Balance: Weighing the Options on Global Warming Policies*. New Haven, CT: Yale University Press. [63,67,70,75]
- NORDHAUS, W., AND J. BOYER (2000): *Warming the World: Economic Modeling of Global Warming*. Cambridge, MA: MIT Press. [43,63,82,83]

- PIZER, W. (1998): "The Optimal Choice of Climate Change Policy in the Presence of Uncertainty," *Resource and Energy Economics*, 21, 255–287. [67]
- POPP, D. (2006): "ENTICE-BR: The effects of Backstop Technology R&D on Climate Policy Models Energy Economics," *Energy Economics*, 28 (2), 188–222. [46]
- REVELLE, R., AND H. SUESS (1957): "Carbon Dioxide Exchange Between Atmosphere and Ocean and the Question of an Increase of Atmospheric CO₂ During Past Decades," *Tellus*, 9, 18–27. [64]
- ROE, G., AND M. BAKER (2007): "Why Is Climate Sensitivity so Unpredictable?" *Science*, 318 (5850), 629–632. [50,83]
- ROE, G., AND Y. BAUMAN (2011): "Should the Climate Tail Wag the Policy Dog?" Report, University of Washington, Seattle. [48,50,81]
- ROGNER, H.-H. (1997): "An Assessment of World Hydrocarbon Resources," *Annual Review of Energy and the Environment*, 22, 217–262. [68]
- ROMER, P. (1986): "Increasing Returns and Long-Run Growth," *Journal of Political Economy*, 94 (5), 1002–1037. [57]
- SAINT-PAUL, G. (2002): "Environmental Policy and Directed Innovation in a Schumpeterian Growth Model," Working Paper. [84]
- (2007): "Quels instruments pour une politique environnementale?" Working Paper. [84]
- SCHWARTZ, S., R. CHARLSON, R. KAHN, J. OGREN, AND H. RODHE (2010): "Why Hasn't Earth Warmed as Much as Expected?" *Journal of Climate*, 23, 2453–2464. [74]
- SINCLAIR, P. (1992): "High Does Nothing and Raising Is Worse: Carbon Taxes Should Be Kept Declining to Cut Harmful Emissions," *Manchester School*, 60, 41–52. [46]
- SINN, H.-W. (2008): "Public Policies Against Global Warming: A Supply Side Approach," *International Tax and Public Finance*, 15 (4), 360–394. [46,84]
- STERN, D. I. (2012): "Interfuel Substitution: A Meta-Analysis," *Journal of Economic Surveys*, 26, 307–331. [68]
- STERN, N. (2007): *The Economics of Climate Change: The Stern Review*. Cambridge, U.K.: Cambridge University Press. [44,70,82]
- STERNER, T., AND M. PERSSON (2008): "An Even Sterner Review: Introducing Relative Prices Into the Discounting Debate," *Review Environmental Economics and Policy*, 2 (1), 61–76. [71, 78]
- STIGLITZ, J. (1974): "Growth With Exhaustible Natural Resources: Efficient and Optimal Growth Paths," *Review of Economic Studies*, 41, 123–137. [63]
- SWEDISH TAX AGENCY (2010): "Skatter i Sverige—Skattestatistisk Årsbok 2010" (in Swedish), available at <http://www.skatteverket.se>. [83]
- TAHVONEN, O. (1997): "Fossil Fuels, Stock Externalities, and Backstop Technology," *Canadian Journal of Economics*, 30 (4a), 855–874. [46]
- TOL, R. (1997): "On the Optimal Control of Carbon Dioxide Emissions: An Application of FUND," *Environmental Modeling and Assessment*, 2, 151–163. [45]
- UZAWA, H. (2003): *Economic Theory and Global Warming*, Cambridge: Cambridge University Press. [45]
- VAN DER PLOEG, F., AND C. WITHAGEN (2010): "Is There Really a Green Paradox?" Working Paper 2963, CESifo. [78]
- (2012): "Too Much Coal, Too Little Oil," *Journal of Public Economics*, 96 (1–2), 62–77. [46]
- (2014): "Growth, Renewables and the Optimal Carbon Tax," *International Economic Review* (forthcoming). [46]
- VAN DER ZWAAN, B. AND R. GERLAGH (2009): "Economics of Geological CO₂ Storage and Leakage," *Climate Change*, 93, 285–309. [46]
- WEITZMAN, M. (2009): "On Modeling and Interpreting the Economics of Catastrophic Climate Change," *Review of Economics and Statistics*, 91, 1–19. [43,50,83]
- WEYANT, J. (2000): "An Introduction to the Economics of Climate Change," prepared for the Pew Center on Global Climate Change, available at http://www.pewclimate.org/globalwarming-in-depth/all_reports/economics_of_climate_change. [80]

WEYANT, J., O. DAVIDSON, H. DOWLABATHI, J. EDMONDS, M. GRUBB, E. A. PARSON, AND S. FANKHAUSER (1996): "Integrated Assessment of Climate Change: An Overview and Comparison of Approaches and Results," in *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*, ed. by J. Bruce, H. P. Lee, and E. F. Haites. Cambridge: Cambridge University Press. [80]

WITHAGEN, C. (1994): "Pollution and Exhaustibility of Fossil Fuels," *Resource and Energy Economics*, 16, 235–242. [46]

Dept. of Economics, Princeton University, 111 Fisher Hall, Princeton, NJ 08544, U.S.A.; golosov@princeton.edu,

IIES, Stockholm University, SE-106 91, Stockholm, Sweden; john@hassler.se, IIES, Stockholm University, SE-106 91, Stockholm, Sweden; per.krusell@iies.su.se,

and

Dept. of Economics, Yale, 28 Hillhouse Avenue, New Haven, CT 06511, U.S.A.; a.tsyvinski@yale.edu.

Manuscript received August, 2011; final revision received July, 2013.