Economics and Climate Change: Integrated Assessment in a Multi-Region World

John Hassler[1] and Per Krusell[2,3]

March, 2012

Abstract

This paper develops a model that integrates the climate and the global economy—an integrated assessment model—with which different policy scenarios can be analyzed and compared. The model is a dynamic stochastic general-equilibrium setup with a continuum of regions. Thus, it is a full stochastic general-equilibrium version of RICE, Nordhaus's pioneering multi-region integrated assessment model. Like RICE, our model features traded fossil fuel but otherwise has no markets across regions—there is no insurance nor any intertemporal trade across them. The extreme form of market incompleteness is not fully realistic but arguably not a bad approximation of reality. Its major advantage is that, along with a set of reasonable assumptions on preferences, technology, and nature, it allows a closed-form model solution. We use the model to assess the welfare consequences of carbon taxes that differ across as well as within oil-consuming and -producing regions. We show that, surprisingly, only taxes on oil producers can improve the climate: taxes on oil consumers have no effect at all. The calibrated model suggests large differences in views on climate policy across regions.

[1]Institute for International Economic Studies, Stockholm. University

[2]Institute for International Economic Studies, Stockholm. University, and University of Gothenburg

[3]The authors gratefully acknowledge research support from the European Research Council (Krusell), Mistra-SWECIA (Hassler and Krusell), the International Growth Centre (Hassler and Krusell) and Vetenskapsrådet (Hassler).

1 Introduction

Climate change engineered by human activity, i.e., by the burning of fossil fuel, is a pure externality with global scope. Moreover, the scope is long-run. I.e., when John drives his motorcycle to work, his gasoline use causes emissions of carbon dioxide into the atmosphere and consequent global heating, thus imposing a cost that not only hits all currently alive humans but also future generations.[4] How much should John should be restrained in his driving? The climate problem is a text-book case of public economics, but the key question—"how much?"—is a quantitative one. For example, although the scope is long-run, the stock of atmospheric carbon slowly depreciates over time so generations very far away are barely affected. More generally, for a reasonable quantitative analysis, the nature of the problem really requires (i) a global perspective, (ii) a dynamic perspective, and (iii) an equilibrium perspective. For the latter, one needs a fully micro-founded model, as any welfare calculations of different policy alternatives—say, different taxes or quotas—would require analysis of effects on global equilibrium prices. In this paper, we provide such a setup.

More importantly, in this paper we focus on a fourth element: regional heterogeneity. Impact studies suggest highly heterogeneous damages in different regions in the world. Moreover, policies are not adopted uniformly around the globe. The issue of heterogeneity, moreover, appears to be of first-order importance. First of all, damage estimates differ markedly across regions, both because the nature of the damages are different across regions (depending on geophysical region differences, differences in the industry structure, and so on) and because the response of temperature (and, more generally, of climate) to a given global mean temperature increase is very different in different regions. Second, it is apparent from the failures to agree on a global agenda to combat climate change that the views on climate change differs markedly across regions. One reason for this surely is that different regions suffer to varying degrees (and some areas of the world even gain from global warming), but another is that the costs of restricting the use of fossil fuel energy differ across regions. Third, fossil-fuel resources, in particular of oil, are unevenly spread over the world, further strengthening the need to study regional implications of various policy alternatives.

This paper takes a step toward a multi-region analysis. Specifically, we construct a heterogeneous-region dynamic stochastic general-equilibrium integrated-assessment (phew!) model, i.e., one integrating economics and climate change within an otherwise standard modern macroeconomic setting. The model is surprisingly easy to solve; it can be solved in closed form for the laissez-faire equilibrium, and a few simple but highly relevant alternative policies can be fully analyzed. With this model, we are able to capture heterogeneity in various forms. The general-equilibrium structure allows us to look at welfare effects of policy; we use permanent percentage consumption equivalents as a measure. The analysis here is only quantitative to a degree; the model has some short-comings that require a somewhat more elaborate structure to address quantitatively. However, in many ways we believe that the analysis herein is right in the middle of what one would more broadly consider a reasonable part of the parameter space, in particular for the degrees to which

[4] John also owns a bicycle, which he fortunately uses more often than the motorcycle.

different regions are in different positions on the climate-change issue.

We build on a structure which, apart from the key factor of heterogeneity, is a simple version of the model in Golosov et al. (2012). The assumptions there on preferences, technology, and nature conveniently allow simple analysis. They are (i) preferences are logarithmic (thus featuring risk aversion), (ii) capital can be accumulated but depreciates fully over a ten-year horizon, (iii) production is Cobb-Douglas in capital, labor, and energy, (iv) energy is produced from oil, which is costless to produce but of which there is a finite stock, (v) damages appear through a multiplicative factor on output whose logarithm is linear in the stock of carbon, and (vi) the carbon cycle is linear. As argued in Golosov et al. (2012), most of these assumptions are quantitatively reasonable. In the concluding section of this paper, however, we briefly suggest some potentially fruitful generalizations.

How do we introduce heterogeneity across regions? First, we assume that there is a continuum of different regions in the world economy, thus giving each region a negligible impact on world markets. In our calibrated section, we lump regions together into four groups—China, the U.S., Europe, and Africa—in an effort to draw out separate implications for regions within these four distinguishable groups. We also calibrate so as to match available damage estimates for these four groups. The world in our model is, moreover, divided into two sub-groups of regions: oil consumers and oil producers. Consumers in these two groups of regions have the same preferences but there is a stark but, we think, not all that unrealistic assumption that oil-consuming regions do not produce any oil at all; conversely, the only income of oil-producing regions is oil income. This assumption greatly simplifies the analysis. Moreover, damages differ across oil-consuming regions, as do capital stocks and productivity levels. Countries trade in oil, of course, and the world price of oil is endogenous. In fact, there will be a Hotelling-like formula for it, coming from the intertemporal utility maximization of oil producers. Another key assumption for tractability is that regions cannot trade in any other way: they cannot engage in either intertemporal trade or insurance. This assumption is not realistic, but we show that in our calibration, the differences across regions in their marginal products of capital, and hence in their intertemporal marginal rates of substitution of consumption, are closely aligned (though, of course, not literally identical).

We have several findings. First, we look at a version of our model with very limited heterogeneity: one where all oil-consuming regions are identical but where there is still a distinction between oil consumers and oil producers. This model allows us to look at two kinds of simple policy analyses. The first one is an experiment where oil producers are taxed on their profits, a tax that could be implemented internationally through tariffs on oil. The tax is assumed to be accompanied by a lump-sum transfer of the tax receipts back to oil producers, so that there is no net transfer. Thus, we are engineering a substitution effect only on oil producers, who do care about income effects since they have curved (logarithmic) utility. Here, if taxes on oil decline over time, oil producers postpone their oil production, which is an improvement from a world perspective as it postpones heating; the optimal rate of decline of taxes is about 0.3% per year, which involves a 18% percent decline in energy use now. The second experiment is a uniform (across regions) tax on oil purchases (imposed on firms using oil), again with lump-sum rebates so that there are no net transfers across regions. A quite surprising result here is that this experiment leads to no effect on oil use at all and, instead, only to a redistribution of world resources away from oil producers—even though the taxes imply no transfers between oil producers and oil consumers. The reason, again, is the role of income effects in oil countries. Lower demand for oil engineered by taxation of oil will lower the oil price, and although that means a substitution effect away from producing oil now, it also means that oil countries are poorer now. Since they are poorer, they want to increase production now, counteracting the substitution effect. Income and substitution effects cancel precisely under logarithmic utility and hence any taxation merely influencing demand will not have any effect at all in this model.

Second, we see that the differences in views among oil-consuming regions are striking quantitatively. Whereas China and the U.S. even would like to *subsidize* current oil production, Europe and Africa would like significant taxes. In terms of desired current oil use, China would like to actually have a 15% higher oil use than currently and the U.S. would like a 9% increase relative to status quo, whereas Europe would prefer to drop current oil use by 46% and Africa by even more: 60%. In our model, differences in initial capital stocks or TFP levels, however, are immaterial. Thus, it is the differential damage elasticities alone—the percentage GDP losses incurred from a unit increase in the atmospheric carbon concentration—that drive these differences. Third, the effect of carbon leakage is very strong in our model: a single region, such as the EU, could self-impose taxation on its oil use, but this would only have a redistributive effect on oil use—there would be a zero aggregate effect. Thus there would be perfect leakage and no change in the climate at all. The unilateral policy would (i) redistribute somewhat from oil producers to oil consumers, by lowering the oil price, and (ii) redistribute from itself toward other oil consumers; whether the first or second of these dominates depends on the size of region.

Section 2 first provides a brief connection to the literature. In Section 3 we then set up and analyze the basic model. Section 4 looks at taxation from an analytical perspective, whereas Section 5 solves the model and evaluates tax policies in a calibrated version of the model; the calibrated model can also be solved analytically but the purpose in this section is precisely a quantitative evaluation. Section 6 looks at three extensions—one studying uncertainty, in particular pertaining to a catastrophe scenario; another looking at carbon leakage due to policy differences across oil-consuming countries; and a third examining differences in energy intensities across regions—and Section 7 concludes.

2 Connections to the literature

This paper is by no means the first one to provide an integrated analysis of climate and the economy. William Nordhaus (see, e.g., Nordhaus 1977, 1994, and Nordhaus and Boyer, 2000, as well as the overview in Nordhaus, 2011) has pioneered the area by building integrated assessment models around the neoclassical growth model, augmented essentially with (i) a carbon cycle, (ii) a set of climate equations mapping atmospheric carbon into temperature, (iii) an energy sector, and (iv) an abatement mechanism, allowing people to expend costly resources to limit emissions from a given amount of use of fossil fuel. Indeed, one of Nordhaus's main contributions was to incorporate up-to-date, but conveniently (for computational purposes) expressed, natural-science insights in economic models.[5] Nordhaus's models exist in both versions with a single region (labeled DICE—Dynamic Integrated model of Climate and the Economy) and with multiple regions (labeled RICE—Regional Integrated model of Climate and the Economy). By using a well-known economic setting and simplifying the climate model and the carbon cycle, Nordhaus's models offer a transparent framework for analyzing the interaction between the economy and the climate.

Following the increased interest in global warming from the 1990s, the literature on integrated assessment models has become quite broad; Kelly and Kolstad (1999) lists 21 integrated assessment models constructed already in the 1990s. Most of these setups, however, lack explicit modeling of endogenous economic responses to climate change and climate policy. Clarke et al. (2009) provides an overview of ten integrated assessment models used frequently in the climate science community. So far, these macroeconomic models have not received much attention from macroeconomists. A reason for this is likely that the models typically are large and non-transparent from the perspective of dynamic macroeconomic models with explicit markets and price formation. MERGE, for example, is a popular model that exists in many versions and Manne et al. (1995) reports that the economic part of the model itself contains 3,800 variables; another popular model, IGEM, contains 4,000 endogenous variables (Nordhaus, 2011). Needless to say, while gaining in significant realism in some dimensions, models with this large number of variables must be restricted in other dimensions, typically by limiting forward-looking, not allowing fully microeconomically founded market mechanisms, or imposing behavioral restrictions.

Our approach here is to drastically cut down on the detail with which the climate and energy sector are modeled in order to build a more transparent and easily communicated integrated assessment model that allows both dynamics and stochastics in a multi-region world. The setting is a quite close relative of Nordhaus's RICE model. In contrast, however, our model allows an analysis of the decentralized equilibrium and is therefore fully equipped to analyze different policy options in general equilibrium. In particular it is possible to make welfare comparisons between "status quo" and adopting the best possible policy instruments—or anything in between. In contrast, most of the existing literature looks at a planning problem and can therefore, at best, back out the marginal social cost of carbon—which would also be the optimal, first-best, tax (the so-called Pigou tax)—as a shadow value on the equation governing the carbon concentration in the atmosphere.

The present model setup, aside from allowing us to go quite far analytically, is also entirely in the "modern macroeconomic" tradition: frameworks that are usually formulated using recursive methods, based on the neoclassical growth model in a competitive-equilibrium decentralization, and allow model solution using recursive methods. In particular, the recursive approach is convenient for explicit modelling of uncertainty. Important earlier contributions within this class include Kelly and Kolstad (1999), who in addition study learning, and Leach

[5]Economists interested in learning more about the natural-science aspects of climate change are recommended to read the various reports from IPCC, the Intergovernmental Panel on Climate Change; see, e.g., IPCC (2007).

(2009); both of these use numerical model solution. As for explicit equilibrium models where prices in world markets are nontrivially determined and a function of climate policy, see the recent papers by Gars (2012) and Hémous (2012).

An important ingredient in our integrated assessment model is the extension of the neoclassical growth model to allow analysis of non-renewable resources. Here the classic reference is Dasgupta and Heal (1974). The present model takes a short-cut in only considering oil, a resource in finite supply and with low extraction cost (literally set to zero in the model). In contrast, some large-scale models in the literature have very elaborate energy supply structures, but then other shortcuts are necessary for computational reasons; for example, MIT's Joint Program on the Science and Policy of Global Change offers a very elaborate model with myopic agents and a less elaborate model with foresight but no uncertainty (see Paltsev et al., 2005, and Babiker et al., 2009). Relatedly, there is also a growing theoretical literature on the choice of R&D in renewable vs. fossil energy technologies. Here the available studies build on the endogenous-growth literature; an early contribution is Bovenberg and Smulders (1995) and a recent contribution is Acemoglu, Aghion, Bursztyn, and Hémous (2011). Empirical studies on directed technical change in the energy area include microeconomic studies, such as Popp (2002), and macroeconomic analyses, such as Hassler, Krusell, and Olovsson (2011). The present paper maintains the exogeneity of both energy supplies and any related technology. This is an important shortcoming for some purposes, especially for the study of policy aiming at a "technological resolution" of the problem of climate change, but for the assessment of welfare regional consequences, it is arguably less critical. Studies of the "Green Paradox", i.e., that early availability of renewable-energy technologies might precipitate the burning of fossil fuel, including Sinn (2008) and van der Ploeg and Withagen (2010), also do not necessarily endogenize technology but do look at the choice between what technologies to use. Finally, the present paper touches on the relevance of small-probability disasters; here, Weitzman's work is of essence (see, e.g., Weitzman, 2009).

3 Decentralized economy

The decentralized economy is autarkic across regions except for the existence of trade in oil. We first discuss the oil-using regions and then look at oil producers. Thereafter we derive the equilibrium outcome and discuss climate and related welfare evaluations.

3.1 Oil-using countries

A given country is characterized by a TFP level, which has an aggregate and an individual component, each of which originates in climate factors as well as economic/institutional factors. Let us summarize this variable, for now, by A_i. A_i moves over time and is stochastic. It also depends on an aggregate variable—an externality—denoted S, which also evolves over time, suppressed in the notation for the moment.

A country is also characterized by a level of capital. The labor input is abstracted from.

Allocations across consumers within a country are suppressed here—an interpretation is that there is a representative agent in each country. His/her preferences are

$$E_0 \sum_{t=0}^{\infty} \beta^t \log(c_t)$$

and the region's budget/resource constraint is $c_t + k_{t+1} = A_{it} k_t^{\alpha} e_t^{\nu} - p_t e_t \equiv \hat{y}_{it}$, where p is the price of oil, also potentially random, and $\hat{y}$ is defined as output net of oil costs—total value added in the country. Thus, a country chooses a stochastic process for capital and oil purchases to maximize the above objective subject to its resource constraint. Capital depreciation is set to 100%. This facilitates analytical solution of the model greatly and is not a terrible approximation for periods of 10 years or so or longer.

This delivers oil choices and resulting output according to

$$e_t = \nu^{\frac{1}{1-\nu}} A_{it}^{\frac{1}{1-\nu}} k_t^{\frac{\alpha}{1-\nu}} p_t^{\frac{-1}{1-\nu}} \quad \text{and} \quad y_{it} = A_{it} k_t^{\alpha} e_t^{\nu} = \nu^{\frac{\nu}{1-\nu}} A_{it}^{\frac{1}{1-\nu}} k_t^{\frac{\alpha}{1-\nu}} p_t^{\frac{-\nu}{1-\nu}}.$$

Straightforward guessing and verifying using the Euler equation implies that saving follows

$$k_{t+1} = \alpha\beta y_{it} = \alpha\beta \nu^{\frac{\nu}{1-\nu}} A_{it}^{\frac{1}{1-\nu}} k_t^{\frac{\alpha}{1-\nu}} p_t^{\frac{-\nu}{1-\nu}},$$

i.e., the saving rate out of net output is constant, despite productivity being random. Similarly, country i's consumption level is $(1-\nu-\alpha\beta)y_{it}$. Thus, as a function of oil prices—which are set internationally—and the productivity process—which is also externally determined—we can fully solve the country's problem.

3.2 Oil-producing countries

We assume that there are many oil producers operating under perfect competition. An oil producer, for simplicity, chooses only oil extraction; or, rather, it chooses how much oil to keep in the ground for next period: R_{t+1}. Extraction of oil is costless. The country maximizes the same kind of utility functions as that of oil consumers subject to a budget/resource constraint $c_t + p_t R_{t+1} = p_t R_t$, with $0 \leq R_{t+1} \leq R_t$.

This is a cake-eating problem with random returns; oil producers are not endowed with anything but oil—they have no other production technology. Defining the budget as $c_t + \tilde{R}_{t+1} = r_t \tilde{R}_t$, where $r_t \equiv \frac{p_t}{p_{t-1}}$ and $\tilde{R}_t \equiv R_t p_{t-1}$, the Euler equation becomes

$$\frac{1}{c_t} = \beta E_t \left[\frac{1}{c_{t+1}} r_{t+1} \right],$$

which is solved by $\tilde{R}_{t+1} = \beta r_t \tilde{R}_t$. I.e., we obtain, again as a function of prices that this country takes as given,

$$R_{t+1} = \frac{1}{p_t} \beta \frac{p_t}{p_{t-1}} R_t p_{t-1} = \beta R_t.$$

That is, despite the oil price being random, the oil producer chooses to extract a constant fraction $1-\beta$ of the remaining oil each period. Its consumption, moreover, equals $p_t E_t = p_t(1-\beta)R_t$. In sum, because logarithmic utility implies that income effects equal substitution effects, the price path for oil does not affect extraction: a high price in one period implies both that extraction should increase at that time (the substitution effect) and that extraction in other periods should increase (the income effect), with a net effect of no change at all.

It is important to realize that our autarky assumption matters here. If one allows oil producers to smooth consumption not by reallocating oil use over time but by investing (or de-investing) abroad, income effects become much weaker. In practice, oil countries do invest abroad; similarly, in North America and some other primarily oil-consuming regions, oil production is substantial. However, if this kind of intertemporal trade were allowed without restrictions, we would be in close to a one-region world, since rates of return would be perfectly equalized across regions. Moreover, oil prices would follow an exact Hotelling formula—the return from postponing extraction, i.e., the price rise for oil, would equal the real interest rate at all points in time. For the present model Hotelling pricing can occur, namely when there is no heterogeneity across oil-consuming regions (see the next section), but other outcomes are possible. We do not explore these in detail here; however, it is noteworthy that in the data, oil prices really do not appear to follow Hotelling's formula, thus giving our model some scope for addressing an important issue.[6] A realistic setting would allow imperfect possibilities for cross-country investment, but such a model is much harder to analyze. The autarky assumption and preferences where income and substitution effects cancel are, in some sense, a defining characteristic of the present model. They both encompass some realism—in incorporating some limits to intertemporal trade while allowing income effects that are often abstracted from—and lead to sharp results; that is a strength in terms of analytical insight but, in particular, it also allows a novel and arguably relevant perspective on the literature. At the same time, extensions away from these assumptions are important, especially perhaps as concerns intertemporal trade, since policy discussions often involve management of capital flows.

3.3 Equilibrium oil prices

The equilibrium oil price is given by market clearing: the amount of oil extracted equals the sum of all the oil consumed. Thus, with minimal algebra we obtain, for all dates and states,

$$p_t = \nu \left(\sum_i A_{it}^{\frac{1}{1-\nu}} k_{it}^{\frac{\alpha}{1-\nu}} \right)^{1-\nu} \left((1-\beta)R_t\right)^{\nu-1}.$$

Compared to a one-country world, or one in which there is full insurance across countries, the price of oil will be lower. Mathematically, this is because $(\sum x_i^{\mu})^{\frac{1}{\mu}} > \sum x_i$ whenever $\mu > 1$. Intuitively, the demand for oil is held back by a distribution of world resources that is not as skewed toward high marginal products of oil as it would otherwise be.

[6] For a discussion, see, e.g., Spiro (2012).

Notice that, even though the model is forward-looking, the general-equilibrium solution for prices can be obtained recursively—in a backward-looking fashion. This follows from the combination of preferences and technology used here; it would in general be more difficult.

If there is only one oil-consuming region (or a continuum of regions identical in all ways), the equilibrium price will be $p_t = \nu A_t k_t^\alpha \left((1-\beta)R_t\right)^{\nu-1}$, net output will satisfy $\hat{y}_t = (1-\nu)y_t = (1-\nu)A_t k_t^\alpha \left((1-\beta)R_t\right)^{\nu}$, consumption of the oil consumer and oil producer will be given by $(1-\nu-\alpha\beta)y_t$ and νy_t, respectively, and saving will be determined by $\alpha\beta y_t$. This allocation is actually *the same allocation* as the one in Golosov et al. (2012). In that paper, the oil-producing and oil-consuming regions are integrated and intertemporal markets automatically ensure that capital and oil give the same equilibrium return; here they do but only because of the special assumptions on preferences and technology. To show the return equalization formally, note that $p_{t+1}/p_t = (y_{t+1}/y_t)(R_t/R_{t+1}) = (y_{t+1}/y_t)/\beta$ and that the return to capital saved at t equals $\alpha y_{t+1}/k_{t+1}$—in both economies. Because $k_{t+1} = \alpha\beta y_t$, the return to capital becomes $(\alpha y_{t+1})/(\alpha\beta y_t) = p_{t+1}/p_t$.

If there are more than one oil-consuming regions, we have that the return to capital in region i equals $\alpha y_{i,t+1}/k_{i,t+1} = y_{i,t+1}/(\beta y_{it})$. Thus, since the oil expenses are a constant share of output, we have that

$$\text{MPK}_{i,t+1} = \frac{1}{\beta}\frac{e_{i,t+1}}{e_{it}}\frac{p_{t+1}}{p_t}.$$

Thus, a given country's return to saving will differ from the return on saving oil, p_{t+1}/p_t, to the extent that growth rate of its oil usage, $\frac{e_{i,t+1}}{e_{it}}$, differ from the growth rate of world oil usage, β. If a given region grows it oil usage faster than the world average, its marginal return to capital will be higher than in the rest of the world. Since the growth rate of oil usage in a country equals the growth rate of output divided by the growth rate of oil prices, the *relative* growth rate of oil usage of a given country equals its relative growth rate: faster-growing countries, have higher returns to capital.

3.4 Implications for climate

Now let us look more carefully at the climate effects of oil use. This is a "one-way investigation" in this simple model: oil use is determined by competitive producers, who under logarithmic utility will pump up a constant share of resources no matter what price path they face. (Prices in turn adjust to make demand adjust to this supply.) Of course, departures from this particular setting will break this link.

Let $A_{it} = e^{-\gamma_i S_t + Z_{it}}$. Here, $\gamma_i S_t$ is the damage from a global atmospheric carbon concentration of S_t; different regions/countries have different sensitivities to global climate changes, a fact strongly supported by the data. Z_{it} represents "other" productivity determinants, including technological and institutional change, but also temporary shocks like short-lived temperature variations. We will regard Z_{it} as exogenous. Parts of the recent growth literature argues that human-capital accumulation is more important than based on pure productivity accounting. That insight could potentially be incorporated here by allowing α to exceed the typically observed value for the share of income to physical capital; in the calibrated section below, we nevertheless use a standard value for α.

What is the link between emissions and damages? We assume a linear carbon cycle: $S_t = \sum_{v=0}^{t}(1 - d_{t-v})\sum_i e_{iv}$. I.e., a unit of emissions at time 0 leads to an increase in the atmospheric carbon concentration S_0 of $1 - d_0$ units—d_0 reflects immediate leakage out of the atmosphere (into the biosphere etc.). In period t, more of the carbon has disappeared; d_t increases over time, though it does not go to 1.[7] A linear depreciation schedule, with time-varying rates, appears to be a very good approximation and simpler than that in RICE.

Measured in units of the consumption good, the marginal per-period damage in region i caused by an increase in S_t is equal to $-\partial\hat{y}_{it}/\partial S_t = \gamma_{it}\hat{y}_{it}$. How these losses ought to be added up is far from obvious. We look at this issue next.

3.4.1 Assigning utility weights

In principle, one can assign social welfare weights to different regions and, given any set of weights, compute a social optimum. A critical issue in this evaluation, as we shall see in this section, is whether or not the so-obtained allocation is close to efficient in terms of production, intertemporal smoothing, and risk-sharing. If it is indeed nearly efficient, some weighting schemes are attractive, but if it is not, other weights may look more reasonable. In the almost-autarky equilibrium here there are restrictions on trade and insurance, but the quantitative extent of these restrictions still depend on parameter values.

Regardless of the issue of how regions ought to be weighted, one can look at the equilibrium effect of global emissions on a given region. Computing the response of an emissions increase of one unit at time zero, taking into account all the future effects, taking uncertainty into account, and translating into current value in consumption units as it is evaluated by this region on the equilibrium path, one obtains

$$-\frac{1}{u'(c_{i0})}E_t\left[\sum_{t=0}^{\infty}\beta^t\frac{\partial\hat{y}_{it}}{\partial S_t}(1-d_t)u'(c_{it})\right] = (1-\alpha\beta)\hat{y}_{i0}\sum_{t=0}^{\infty}\beta^t(1-d_t)E_t\left[\gamma_{it}\hat{y}_{it}\frac{1}{(1-\alpha\beta)\hat{y}_{it}}\right].$$

This becomes the simple expression

$$\hat{y}_{i0}\sum_{t=0}^{\infty}\beta^t(1-d_t)E_t\left[\gamma_{it}\right].$$

Thus, the same formula as in the Golosov et al. (2012) context obtains: the damage is proportional to current output in the region and, beyond that, only influenced by basic structural parameters governing carbon depreciation and damage elasticities.

Moving back to the issue of how losses could be added up—weighted—across regions, consider again a given time period. Adding up all output losses across regions, the total is

$$-\sum_i\frac{\partial\hat{y}_{it}}{\partial S_t} = \sum_i\gamma_{it}\hat{y}_{it}$$

[7] As in Golosov et al. (2012), we assume that a share φ_L of emissions stays in the atmosphere forever, that a share $1-\varphi_0$ of the remainder quickly and permanently exits the atmosphere, and that the remaining part decays at a geometric rate φ. These assumptions suffice to pin down the ds uniquely.

i.e., a (net) *output-weighted* average of the damage coefficients γ_{it}. This calculation, however, relies on production in different countries being perfectly substitutable. It is indeed if there are no restrictions on trade: if the world allocation is optimal. In the almost-autarkic allocation we consider, however, there are such restrictions. Therefore, social welfare weights across regions must enter the calculation. If one, for example, took a utilitarian perspective and simply added marginal utils in a region-weighted fashion one would obtain

$$-\sum_i \frac{\partial \hat{y}_{it}}{\partial S_t} \frac{1}{c_{it}} = \frac{1}{1-\alpha\beta} \sum_i \gamma_{it}$$

The key difference between this expression and that above is that the damage coefficients are not output-weighted here, implying a larger weight on (production-)poorer regions. Thus, how one perceives the damage costs depends on how close to a world-optimal consumption allocation we assume that we are (and how that optimum is defined). A utilitarian perspective would imply that the current world is VERY far from such an optimum. An alternative, consistent with the view that the current world is close to an optimum, is to use weights that are equal to the time-0 inverse of marginal utilities in our respective regions—this way, poorer regions' utils are scaled down in ways proportional to their consumption or wealth. A simple calculation along these lines delivers

$$-\sum_i \frac{1}{u'(c_{i0})} \frac{\partial \hat{y}_{it}}{\partial S_t} u'(c_{it}) = \sum_i (1-\alpha\beta)\hat{y}_{i0}\gamma_{it}\hat{y}_{it} \frac{1}{(1-\alpha\beta)\hat{y}_{it}} = \sum_i \gamma_{it}\hat{y}_{i0}.$$

Here, output weighting appears again—but now simply reflecting the planner's (subjective, and exogenous) weights placed on different regions.

Having said all this, our main approach will be to simply report, region by region, outcomes under different policy scenarios, thus avoiding how to rank outcomes overall.

3.4.2 Direct costs and benefits of interventions

In our specific model, consumption is a fixed share $1-\alpha\beta-\nu$ of output in each region and capital saving is a fraction $\alpha\beta$ of output. The logarithm of output, finally, is a fixed-weight sum of the logarithms of (i) TFP, which is endogenous due to the climate being endogenous, (ii) the capital stock, which in turn is a fixed-weight sum of past outputs, and (iii) the price of oil. The fixed weights, throughout, depend on the three parameters α, β, and γ. If these parameters are identical across regions, the welfare effects of interventions are particularly easy to study. Suppose a policy is undertaken which does not alter the basic workings of the model but affects the path of energy use and oil prices (along with output, capital accumulation, and consumption); we consider such policies in the next section. Then in terms of the impact on the changes in the logarithm of consumption, the impact will be identical for all countries *except for* that part of the impact that works through climate change. The reason for the "identical" part is that changes in oil prices will affect all countries equally, from the period of impact and dynamically forward, since capital accumulation is a fixed fraction of output at all times. The reason for the differential impact is simply that the climate sensitivities differ across region, as per by assumption through the γ_is. Thus, in the basic version of our model, countries with the same value of γ_i will be identically affected by climate policy, regardless of whether they are rich or poor in terms of capital or TFP.[8]

4 Taxation

Now let us introduce a (potentially time-varying) *ad valorem* tax on oil. Throughout, we assume that government simply transfers the entire amount of the tax back to the taxed party in a lump-sum manner. We will consider two taxation schemes. One taxes oil producers, and the other taxes oil consumers. These two cases turn out to have very different implications in terms of oil extraction and welfare. In both cases, the behavior of oil producers is crucial. As we saw above, their logarithmic utility means that equilibrium oil use will not depend on oil prices: high prices imply a substitution effect in favor of production but an income effect away from production, since a lower production is now necessary to maintain a smooth consumption level. Thus, taxes affecting oil demand will lower oil prices but not affect supply. Taxes on oil producers that are designed as pure substitution effects—by rebating any taxes raised—will be effective in affecting oil use.

4.1 Taxes on oil producers

We first look at taxes on oil producers. These taxes can either be implemented by the governments in the oil-producing countries themselves or it can be viewed as an import tax levied by the oil-consuming countries. In either case, the tax proceeds would be rebated back to the producers.[9] This policy will influence extraction, since there will be a substitution effect for producers, but no income effect.

The budget constraint in an oil-producing country now reads $c_t + p_t(1-\tau_t)R_{t+1} = p_t(1-\tau_t)R_t + T_t$, or $c_t + \tilde{R}_{t+1}(1-\tau_t) = r_t\tilde{R}_t(1-\tau_t) + T_t$, where τ_t is the tax rate on oil and T_t is the transfer. The assumption that tax revenues are rebated back implies that $(-\tilde{R}_{t+1} + r_t\tilde{R}_t)\tau_t = T_t$. Now define the growth rate of net-of-tax rates as $\theta_t \equiv (1-\tau_t)/(1-\tau_{t-1})$. The Euler equation then reads

$$\frac{1}{c_t} = \beta E_t\left[\frac{1}{c_{t+1}} r_{t+1}\theta_{t+1}\right].$$

Using $c_t = r_t\tilde{R}_t - \tilde{R}_{t+1}$, the Euler equation becomes

$$\frac{1}{r_t\tilde{R}_t - \tilde{R}_{t+1}} = \beta E_t\left[\frac{1}{r_{t+1}\tilde{R}_{t+1} - \tilde{R}_{t+2}} r_{t+1}\theta_{t+1}\right].$$

[8] Below we also consider cases where the consumption share is varying over time. For those extensions, a more elaborate discussion is needed than what is presented in the present section.

[9] If the taxes are used for other purposes, the problem is very simple to solve. Just define $\tilde{r}_t = r_t\frac{1-\tau_t}{1-\tau_{t-1}}$ and the solution is $\tilde{R}_{t+1} = \beta\tilde{r}_t\tilde{R}_t$.

Express the solution as $\tilde{R}_{t+1} = s_t r_t \tilde{R}_t$, which is equivalent to $R_{t+1} = s_t R_t$: s_t is the fraction of the oil resource at t saved for period $t+1$. Using this notation, the Euler equation becomes

$$\frac{s_t}{1-s_t} = \beta E_t \left[\frac{1}{1-s_{t+1}} \theta_{t+1} \right].$$

Defining $\hat{s}_t \equiv s_t/(1-s_t)$, this can be written

$$\hat{s}_t = \beta E_t \left[(1+\hat{s}_{t+1}) \theta_{t+1} \right] = E_t \left[\sum_{j=1}^{\infty} \beta^j \prod_{v=1}^{j} \theta_{t+v} \right].$$

Thus, one can solve for the s_t sequence in closed form given any sequence of tax rates. Thus, one obtains oil production as a function of tax rates only.

Looking at simple versions of this expression, we first see that a constant tax rate, which means $\theta_t = 1$ at all dates, implies that $s_t = \beta$, as in the laissez-faire case.[10] This means not only that a constant tax will not influence the extraction path but that it will have no effect at all on the equilibrium allocation.

Suppose, instead, that we introduce a tax sequence such that the growth rate of net-of-tax rates, θ, is constant but not equal to one: taxes grow, or shrink, over time. Then it is easy to verify that $s_t = \beta\theta$. For example, if tax rates decrease over time in such a way that $\theta > 1$, the stock of oil will fall more slowly over time, and the initial-period outtake of oil will be smaller. As we will show below, a $\theta > 1$ will raise welfare relative to laissez-faire.[11] In particular, we will have a total extraction at t equal to $R_t - R_{t+1} = (1-\beta\theta)(\beta\theta)^t R_0$. This taxation scheme will also influence the price path. From the perspective of a lower oil supply at time 0, it will raise the initial oil price. There will be a reinforcing effect by a lower damage at time 0, increasing productivity and therefore oil demand. As a result, capital accumulation is affected—it will likely fall, as less oil is used and net output hence falls—and this counteracts the impact effect: this channel puts downward pressure on oil demand. Over time these effects play out in a non-constant way: there will be transitional dynamics. Furthermore, energy use long enough into the future will be higher than before, reversing the impact effects. A full analysis of the resulting net effect over time on prices, and hence on utilities of oil-using and oil-producing countries, is non-trivial. However, let us briefly again consider the case without heterogeneity among oil consumers. The allocations will maintain the exact same form as before and hence $c_{pt} = \nu y_t$ and $c_{ct} = (1-\nu-\alpha\beta) y_t$, for producers and consumers, respectively, where $y_t = A_t k_t^{\alpha} ((1-\beta\theta) R_0 (\beta\theta)^t)^{\nu}$. Thus, both regions are affected the same way by the tax policy! It is, moreover, straightforward to solve explicitly for equilibrium utility and to show that an increase in θ above 1 is desirable.[12] Thus, oil extraction, as in Golosov et al. (2012) and similar models, ought to be delayed.

[10] In fact, this result will obtain also if $E_t \theta_{t+v} = 1$ for $v > 0$ if the growth rates of the net-of-tax rates are also independent over time.

[11] In a simple setting, this point was made already by Sinclair (1992).

[12] The derivation, which does take more than one line of algebra, is available upon request.

4.2 Taxes on oil consumers

Suppose all oil-consuming regions use a tax of τ_t on every dollar of oil purchases. Considering the decentralized equilibrium behavior in the region, firms will then maximize $A_{it}k_{it}^{\alpha}e_{it}^{\nu} - r_{it}k_{it} - (1+\tau_t)p_t e_{it}$ by choice of k_{it} and e_{it}. This implies that $(1+\tau_t)p_t e_{it} = \nu y_{it}$, i.e., that $p_t e_{it} = \hat{\nu}_t y_{it}$, where $\hat{\nu}_t \equiv \nu/(1+\tau_t)$.

Consumers satisfy their Euler equation, i.e., $1/c_{it} = \beta r_{i,t+1}/c_{i,t+1}$, and the resource constraint reads $c_t + k_{t+1} = A_{it}k_{it}^{\alpha}e_{it}^{\nu} - p_t e_{it}$, since the tax revenues are rebated back. Thus, we obtain $c_t + k_{t+1} = (1-\hat{\nu}_t)y_{it}$. Guessing that $k_{i,t+1} = s_{it}y_{it}$, the Euler equation becomes

$$\frac{s_{it}}{1 - s_{it} - \hat{\nu}_t} = \alpha\beta \frac{1}{1 - s_{i,t+1} - \hat{\nu}_{t+1}}.$$

This equation is solved by $s_t = \alpha\beta$ whenever $\tau_t = \tau$, i.e., whenever taxes are constant over time. Otherwise, one needs to resort to numerical solution to obtain the saving rates. Consumption is a higher fraction of output the higher is the tax on energy, since consumption equals $(1-\alpha\beta-\hat{\nu}_t)y_{it}$ and $\hat{\nu}_t$ decreases in τ_t.

To find equilibrium energy use, note that

$$e_{it} = \nu^{\frac{1}{1-\nu}} A_{it}^{\frac{1}{1-\nu}} k_{it}^{\frac{\alpha}{1-\nu}} \left((1+\tau_t)p_t\right)^{\frac{-1}{1-\nu}}$$

for region i. The oil producer's problem is unaffected by the taxes in the oil-consuming regions: total oil production will equal $(1-\beta)R_t$ in period t. Market clearing thus implies

$$p_t = \nu \left(\sum_i A_{it}^{\frac{1}{1-\nu}} k_{it}^{\frac{\alpha}{1-\nu}} (1+\tau_t)^{\frac{-1}{1-\nu}} \right)^{1-\nu} \left((1-\beta)R_t\right)^{\nu-1}.$$

Region i's output, finally, equals

$$y_{it} = \nu^{\frac{\nu}{1-\nu}} A_{it}^{\frac{1}{1-\nu}} k_{it}^{\frac{\alpha}{1-\nu}} (1+\tau_t)^{\frac{-\nu}{1-\nu}} p_t^{\frac{-\nu}{1-\nu}}.$$

What are the welfare benefits of taxation here? Suppose that all regions use the same taxes. Then the equilibrium price of oil will be $1/(1+\tau_t)$ times what it was in laissez-faire, so energy use and output is the same as in the laissez-faire allocation in all regions, as is capital accumulation. The only difference is that, since the share of output paid for the oil input is decreasing in the tax rate, there is more left over for consumption the higher is the tax rate. From the perspective of oil consumers, the "best" tax rate on oil is infinity at all dates, leaving oil producers with a zero price on oil and thus zero resources to consume at all times. In sum, taxes used in oil-consuming regions will lead to a redistribution of resources but will not affect total energy use at all. Thus, the effects on the climate are nil.

The result that taxes in the oil-consuming regions cannot influence the path of oil extraction is a result of assuming that oil producers have logarithmic utility functions. If they had non-logarithmic utility, changes in oil demand would, through the effect on the equilibrium oil price, be able to influence oil extraction. With logarithmic utility, an increase in the oil price at some date will introduce an income effect (extract less at this date to save more oil for later dates) and a substitution effect (extract more at this date because it pays more) but the two effects cancel exactly. With more (less) than logarithmic curvature, a decreased oil price today—relative to those in the future—will lead to higher (lower) extraction today. Thus, oil consumers can, by changing the domestic tax rates on oil over time, influence the extraction path. To delay extraction, one would need a path of increasing (decreasing) tax rates on oil to the extent the oil producer has more (less) than logarithmic utility curvature.

4.3 Summing up

We learn from the above analysis that (*ad-valorem*) taxes on oil producers have several convenient features. One is that total oil production can be influenced, so long as the tax rates on oil change over time; a constant tax on oil has no effect at all on the equilibrium. Looking at the case of no heterogeneity among oil consumers, another convenient feature is that they affect the utility of oil producers and oil consumers the same way—consumption changes by the same percentage amount in the two groups, whichever time period is being considered. A third convenient feature is that the equilibrium can again be solved analytically. Considering a decreasing tax path such that the net-of-tax rate grows at a constant rate, one obtains simple formulas for welfare and it is possible to prove that there is a path of this form that improves everyone's welfare.

Taxes on oil consumers, on the other hand, have no effect at all on oil extraction: they only affect demand, and therefore only the equilibrium oil price, which we know has no effect on extraction.[13] From the perspective of correcting the climate externality, this kind of tax is therefore not useful. This result does depend on our assumption that utility has logarithmic curvature in consumption; in other cases, taxes that change over time will be able to influence oil supply, but the effect will be small unless curvature is far from logarithmic. Taxes—even constant ones, in this case—do influence the equilibrium allocation: they have, first and foremost, distributionary effects. If all oil-consuming regions coordinate, a tax will allow these regions to capture surplus that the oil-producing regions would otherwise have obtained: the equilibrium price of oil falls. From the perspective of oil consumers, an infinite tax would be best. If oil-consuming regions do not cooperate, a case we look at in Section 6.2, there will be carbon leakage from high-tax to low-tax countries causing redistribution effects among them, and total damages may be affected as well. In the next section we will consider specific quantitative examples with taxation levied on producers.

5 Quantitative examples

Let us now calibrate the model in a stylized way. As we will discuss below in a little more detail, in some respects—in particular, in its treatment of the energy supply—the model here is too simple and therefore the welfare gains from optimal carbon policy are made look

[13] Recall that this result comes about because oil producers have logarithmic utility.

smaller than they really ought to be. However, the present model does allow us to show the divergent views on this policy among regions.

5.1 Calibration

We use as the same carbon cycle as in Golosov et al. (2012), but allow heterogeneity in productivity and climate damages. We use the exponential damage function as calibrated there but we assume that there is heterogeneity in the damage parameter; we use Nordhaus's calibration in RICE 2007 for the specific values. Nordhaus uses a bottom-up approach and thus adds seven different sorts of damages that have been estimated in separate applied studies: damages in agriculture, damages due to sea-level rise, damages in sectors other than agriculture, health damages, non-market amenity impacts, impacts on human settlements and ecosystems, and catastrophe damages (for more specifics on catastrophe estimates, see Section 6 below). Moreover, estimates are obtained in a larger number of regions and then aggregated. In RICE, a 2.5% increase in the global mean temperature leads to an output-weighted loss of 1.5%. In Africa, the damages are estimated at 3.91%, i.e., 2.61 times larger, whereas in OECD Europe, they are 2.83%, i.e., 1.89 times larger. The corresponding ratios for China and the U.S. are 0.15 and 0.3, respectively. We use these ratios to calibrate the region-specific damage parameters γ_i. Before proceeding, we should point out that there really is little firm knowledge about what damages really are, be it in total or by region; unfortunately, we know much less about them than about how human activity causes global warming and other forms of climate change. Here, mainly in order to show how the model can be used, we choose numbers that at least are standard in the literature.

Furthermore, we assume that productivity is permanently twice as high in the richer as in the poorer countries. This assumption does not matter, however, for our welfare measures, as indicated above: costs and benefits from changing oil extraction are all proportional to the level of output in this model. Similarly, the levels of the capital stocks do not matter. We select these so that the marginal productivity of capital is not too different across regions and smooth over the initial periods, i.e., neither showing sharp increases nor decreases.[14]

We set the current amount of fossil fuel in ground (R_0) to match a laissez-faire increase in the global mean temperature of 4 degrees Celsius, which matches the IPCC 2007 laissez-faire scenario. The implied calibrated value for R_0 is 1,963 GtC. Atmospheric carbon depreciation is linear, again as in Golosov et al. (2012), and is thus calibrated with three parameters: 20% of any emitted unit never depreciates (stays forever in the atmosphere), 60.7% depreciates immediately, and the remainder (which neither depreciates immediately nor stays forever) decays at a rate of 2.28% per decade. We set the capital share α to 0.3 and the fossil fuel share, ν, to 0.04, a figure which is in between the values of the U.S. and the EU. The common growth rate of productivity is assumed to be 2% per year, again a figure that does not affect the welfare comparisons. The subjective discount rate, finally, is set to 1.5% per year.

We also study a benchmark global economy without regional heterogeneity—the one-

[14]Caselli and Feyrer (2007) study the dispersion of the marginal products of capital across countries and argue that the dispersion is not large.

region setup in Golosov et al. (2012)—and adopt the parameter values from that study. Those parameter values imply a damage of 3.1% of output for a doubling of the CO_2 concentration.

5.2 Model predictions

Consider first the case of no heterogeneity. The laissez-faire allocation implies, as calibrated, a temperature increase that peaks at 4 degrees with peak damages at 4.7% of output. For this economy, we look for optimal tax policy within the class of ad-valorem tax rates on producers—with compensating lump-sum transfers, as in Section 4.1—such that the net of the tax, i.e., $1 - \tau_t$, grows at a constant rate θ. I.e., we maximize over θ. The optimal value for θ in our one-region calibration turns out to be 1.03, implying an initial drop of 18% of oil use. This is broadly consistent with the optimal policy found in Golosov et al. (2012), where optimal policy was allowed to be time-varying. Since the model is specified with a period being a decade this means that the optimal ad-valorem tax rate should fall by slightly less than 0.3% per year. The gain from implementing this tax path, expressed as a compensating variation—a permanent percentage increase in consumption—is 0.09%. The implied temperature peaks at 3.8 degrees but is somewhat delayed relative to laissez-faire and the maximal damage is now 4.4%.

Turning to the case of heterogeneity, the results are quite different: there are substantial differences in optimal policies from the perspective of the different regions. We illustrate this by looking at different values of θ and their implications. Table 1 displays a range of values for θ and the associated initial levels of oil use relative to the laissez-faire outcome. The values are chosen to be the optimal choice from the perspective of the regions we consider—the four oil-consuming regions and the oil-producing region. The table shows how the faster the rate of decline in taxes implies that oil use is delayed, hence implying lower oil use now.

Table 1: initial oil use relative to laissez faire

gross growth rates, net-of-tax on oil: θ	0.975	0.985	1.038	1.075	1.098
initial oil use	115%	109%	76.9%	54%	39.7%

Looking at the implications for heating, the laissez-faire allocation in the heterogeneous-region model leads to very similar heating to that in the homogeneous case, but the damages are very differently distributed. They peak at 1.44% for the U.S. and a mere 0.7% for China. However, they are as high as 11.8% and 8.7% for Africa and Europe, respectively. Consequently, heterogeneity implies that different values of θ are optimal from the perspective of different regions. Table 2 shows the compensating variations for the different values of θ.

Table 2: welfare gains and oil use rel. to laissez faire

θ	0.975	0.985	1.038	1.075	1.098
China	0.05%	0.05%	-0.33%	-0.43%	-0.61%
US	0.01%	0.02%	-0.27%	-0.98%	-1.82%
Europe	-0.42%	-0.25%	0.51%	0.74%	0.59%
Africa	-0.61%	-0.37%	0.85%	1.51%	1.68%
oil producers	-0.21%	-0.12%	0.14%	-0.08%	-0.55%

The low-damage regions would prefer no increases in the tax at all; in fact, their most preferred values of θ are 0.985 (U.S.) and 0.975 (China), which would thus increase heating in the short run. The corresponding figures for Africa and Europe are a contrasting 1.098 and 1.075, respectively, implying a significant delay in fossil-fuel use. As for oil producers, they gain if and only if the world as a whole gain—given the preferences and technology assumed here—and hence their most preferred θ is near the value that was optimal in the one-region world, and it is also right in between the optimal values for China and the U.S., on the one hand, and Europe and Africa, on the other.[15]

The analysis summarized in Table 2 reveals that the compensating variations are fairly small quantitatively: policy cannot reduce damages much. Is this finding robust? Its source is a combination of assumptions—fossil fuel is a necessary input in production, it is not costly to extract, and it is in finite supply. This means that all fossil fuel will eventually be used and that tax policy can only change the time profile of oil use and not the total amount consumed. A more realistic model would feature both oil and coal in parallel use, where the latter exists in large but finite supply, though with relatively high extraction costs. Such a setting, also including renewable energy, is studied in Golosov et al. (2012) and briefly discussed in Section 7. A key difference between that setting and the one here is that the total welfare costs of suboptimal climate policy rise significantly: it becomes necessary to sharply reduce the use of coal, of which there is abundance relative to the total oil supply.

We also report transition paths; in Figure 1 (top left), we plot global output relative to laissez faire for three values of θ: 1.03, 1.075, and 1.098.

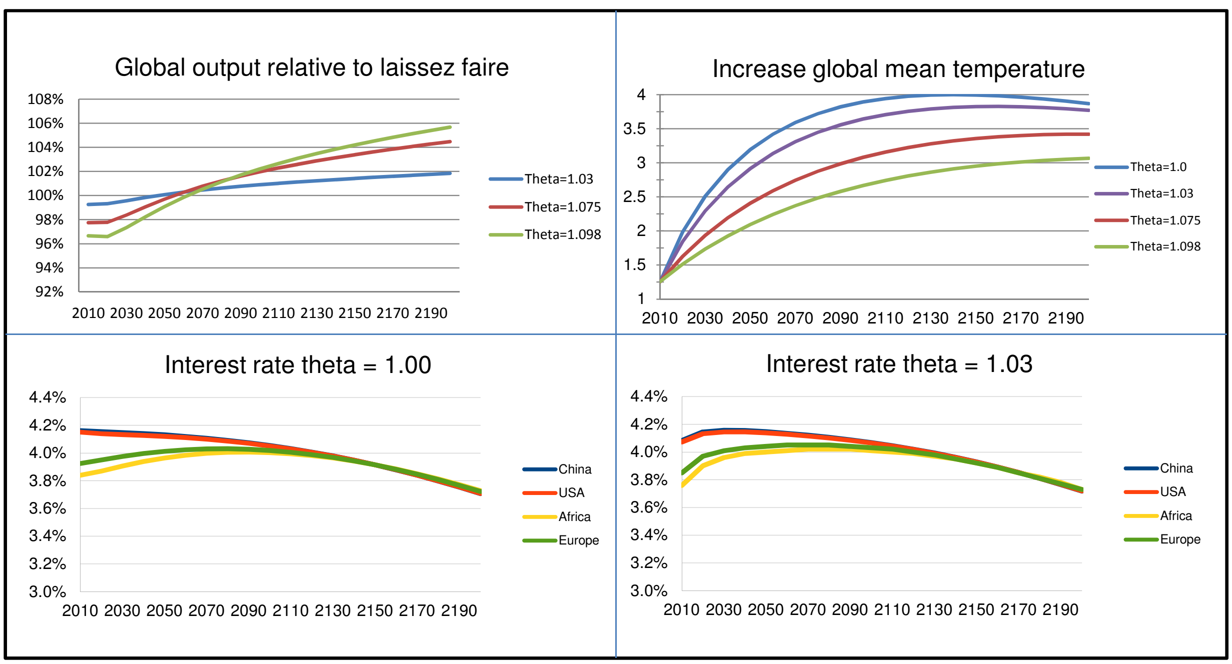

Figure 1

[15]The one-region world is not an exact aggregation of the heterogeneous-region world, and hence the optimal θs differ slightly.

We see that a stronger policy stance, i.e., a higher θ, leads to lower current output but higher long-run output. Break-even in terms of output—when output starts exceeding its laissez-faire path—occurs in the 2050s for the weaker policy stance and about 10 years later for the stronger stance. For the world as a whole, as explained above, with the discounting assumed here, the weaker stance is optimal, whereas the stronger stance is optimal for Europe and Africa.

Figure 1 (top right) shows the temperature development under laissez faire (which, again, is approximately the path preferred by China and the U.S.) under the same three values for θ.

Temperature divergence across the different policy paths comes about slowly, since temperature is a function of the carbon stock, to which we only make minor additions in any given decade. The differences across policy paths are at most a little less than 2 degrees Celsius, and these differences are reached in the latter half of the current century and then only fade away rather slowly.

Figure 1 (bottom left and right) displays the yearly interest rates in the different regions for $\theta = 1$ (laissez-faire) and $\theta = 1.03$, respectively.

There are significant differences across regions, but the differences are not very large. Thus, one might expect that an extended model where the regions are allowed to trade in a risk-less bond—a model where the marginal returns to capital would not be equated either, but would come closer together—will not have very different implications.[16]

6 Extensions

We look at two extensions: probabilistic, catastrophic climate change and the possibility that taxes differ across oil-consuming regions.

6.1 Probabilistic, catastrophic climate change

Nordhaus's calibration is partly based on catastrophic damages; Nordhaus (2000) uses survey evidence based on questions to scientists and economists about hypothetical extreme scenarios. One given especial weight is a probability, estimated to be of 6.8%, of catastrophically large damages, defined as a loss of 30% of GDP, occurring through 6-degree Celsius heating. We consider such a possibility here, including Nordhaus's more disaggregated assessments of catastrophe damages, in the context of our model and make welfare calculations for the value of contingent policy should catastrophe occur.

Nordhaus assumes two possibilities, a "moderate" and a "catastrophic" effect of climate change. The latter, as per the survey answers, occurs with a probability of 6.8% with a heating of 6 degrees. Using a climate sensitivity of ($\lambda =$) 3 degrees per doubling of the CO_2 concentration and the formula $\Delta T = \lambda \frac{\ln S - \ln 600}{\ln 2}$, describing global heating as a function of how carbon concentration differs from its preindustrial level of about 600, we obtain $6 = 3\frac{\ln S_H - \ln 600}{\ln 2}$ and $2.5 = 3\frac{\ln S_L - \ln 600}{\ln 2}$ as solving for the two carbon concentration levels S_H

[16]See, e.g., Castro (2005) for a model of this sort.

and S_L corresponding to catastrophic and moderate outcomes, respectively. These values are then used to calibrate values for the damage parameters γ_{iH} and γ_{iL} specific to region. Nordhaus's calibration is summarized in Table 3 below:

Table 3: damages, different scenarios

	no catastrophe	catastrophe
China	-0.3%	22.1%
US	0.01%	22.1%
Europe	0.92%	44.2%
Africa	3.52%	22.1%

Thus, he assesses catastrophe in Europe (and India, which we do not discuss here) as twice as costly as in each of the other regions (thus reversing the Europe-Africa ranking from the moderate case, where Africa's damages exceed those for Europe). Equipped with these numbers, which appear to be rather rough guesses, we can compute the values for γ_{iL} and γ_{iH} from $e^{-\gamma_{iL}(1069-600)} = 1 - D_{iL}$ and $e^{-\gamma_{iH}(2400-600)} = 1 - D_{iH}$, where the damage parameters (the Ds) are obtained from Table 3. This results in γs as tabulated in Table 4:

Table 4: regional damage parameters

	no catastrophe	catastrophe
China	-6.39×10^{-6}	1.39×10^{-4}
US	2.13×10^{-7}	1.39×10^{-4}
Europe	1.97×10^{-5}	3.24×10^{-4}
Africa	7.64×10^{-5}	1.39×10^{-4}

We assume that the uncertainty is a one-time event: in the first period of the model, which is calibrated to last for a decade, the damages are assumed to be given by the moderate ones, i.e., by the γ_{iL}s. In the second period, with a probability of $\rho = 6.8\%$, the γs change to the γ_{iH}s. The change (or lack thereof) is assumed to then remain permanent. In terms of policy, we assume a status-quo policy of $\theta = 1$ in the first period and thereafter unless the catastrophe materializes, in which case the policy is changed to a $\theta > 1$; for that case we consider two values—one optimal from the European perspective and one optimal from the perspective of the other regions.

It is straightforward to calculate the implied saving rate for oil in the first period. We know that once the tax is introduced, $s_t = \beta\theta$ for all t. For the first period, $s_0 = \beta(1 - \rho + \rho\theta(1 - s_0)/(1 - \beta\theta))$, implying

$$s_0 = \beta \frac{1 - \rho + \theta \frac{\rho}{1-\beta\theta}}{1 + \beta\theta \frac{\rho}{1-\beta\theta}},$$

displaying how the oil producers take the probability of disaster into account—because that event implies a different, and higher, subsequent return on saving oil for the future.[17]

[17] If there are shocks in every period, it is also possible to solve for savings rates rather easily. For the present purposes, however, a one-time shock captures the main aspects we wish to highlight.

We find that the optimal θ from the perspective of China, Africa, and the US is 1.099; for Europe, it is 1.14. Thus, in either case, oil consumption falls drastically when optimal tax is implemented: by 45.6% and 58.8%, respectively. One can consequently compute the values of using the "contingency tax" (compensating variations as permanent consumption shift, compared to never using taxes) preferred by the different regions. For China, Africa, and the US, this value is 0.092% if their preferred tax is used but negative, -0.004%, if Europe's preferred tax is used upon catastrophe. In contrast, for Europe, the value is 0.428% if the other regions' tax is used and 0.587% if their own preferred tax is used. These numbers are relatively low, in part because the probability of disaster is low, and in part because the losses in the disaster scenario are "only" of the magnitude of those during the Great Depression. Weitzman (2009) argues for much more extreme adverse events. Another possibility to consider would be higher risk aversion or even ambiguity aversion; see, e.g., Gollier and Weitzman (2010). These can be feasibly studied also in the simple model here. Learning is another possible extension; here the relevant reference is Kelly and Kolstad (1999).

6.2 Different taxes on oil in different regions: carbon leakage

Because our model is a multi-region world, it is in principle straightforward to analyze tax policies that differ across regions. This is of interest not only because of the presently very large differences in carbon taxation—and the apparent infeasibility to come to a world agreement of uniform taxes/quotas—but also because it allows a transparent analysis of *carbon leakage*. Carbon leakage occurs when one country taxes its own emission but other countries do not: in that case, at least some of the reduction in emissions will "leak out" and cause an increase in the other countries. How important are these effects? We do not pursue a quantitative analysis here but we at least offer a simple qualitative model which, when appropriately (and relatively easily) amended, can be made quantitative.

Therefore, suppose region i uses a tax rate τ_{it} per dollar of oil purchases. Firms in region i maximize $A_{it}k_{it}^{\alpha}e_{it}^{\nu} - r_{it}k_{it} - (1+\tau_{it})p_t e_{it}$, implying that $(1+\tau_{it})p_t e_{it} = \nu y_{it}$ so that $p_t e_{it} = \hat{\nu}_{it} y_{it}$, where $\hat{\nu}_{it} \equiv \nu/(1+\tau_{it})$—the same kind of notation as used in Section 4.2.

The Euler equation reads $1/c_{it} = \beta r_{i,t+1}/c_{i,t+1}$, with an associated resource constraint $c_t + k_{t+1} = A_{it}k_{it}^{\alpha}e_{it}^{\nu} - p_t e_{it}$ (tax revenues are, again, rebated lump-sum). It follows that $c_t + k_{t+1} = (1-\hat{\nu}_{it})y_{it}$. The saving rate will be a fraction of output: $k_{i,t+1} = s_{it}y_{it}$. This implies an Euler equation that turns into

$$\frac{s_{it}}{1 - s_{it} - \hat{\nu}_{it}} = \alpha\beta \frac{1}{1 - s_{i,t+1} - \hat{\nu}_{i,t+1}}.$$

As in Section 4.2, this equation can only be solved explicitly if the tax rate does not change over time. In this case, we obtain $s_{it} = \alpha\beta$, i.e., equal saving rates everywhere, even if $\tau_{it} = \tau_i$ is not the same across regions. However, as before, it is straightforward to solve this difference equation for any sequence of tax rates, as it does not involve any other endogenous variables than the saving rate. In the case where the tax is constant over time, consumption increases in the level of this tax: consumption equals $(1 - \alpha\beta - \hat{\nu}_i)y_{it}$ at t, and $\hat{\nu}_i$ decreases in τ_i, all in line with our previous findings.

To find a world equilibrium, we use the fact that

$$e_{it} = \nu^{\frac{1}{1-\nu}} A_{it}^{\frac{1}{1-\nu}} k_{it}^{\frac{\alpha}{1-\nu}} ((1+\tau_{it})p_t)^{\frac{-1}{1-\nu}}$$

for region i. As in the case of equal tax rates across oil-consuming countries, the oil producer's problem has the same solution no matter what the taxes are: total oil production will still equal $(1-\beta)R_t$ in period t. It remains to find the price that clears the oil market. It satisfies

$$p_t = \nu \left(\sum_i A_{it}^{\frac{1}{1-\nu}} k_{it}^{\frac{\alpha}{1-\nu}} (1+\tau_{it})^{\frac{-1}{1-\nu}} \right)^{1-\nu} ((1-\beta)R_t)^{\nu-1}.$$

Output in region i is then

$$y_{it} = \nu^{\frac{\nu}{1-\nu}} A_{it}^{\frac{1}{1-\nu}} k_{it}^{\frac{\alpha}{1-\nu}} (1+\tau_{it})^{\frac{-\nu}{1-\nu}} p_t^{\frac{-\nu}{1-\nu}}.$$

What are the new lessons relative to the case with uniform world taxation? As seen from the formulas, a single region deciding to tax oil consumption domestically—while no other regions imposes a tax on oil—will be able to decrease the equilibrium oil price, though less for any given tax increase, of course. Now all the other regions will benefit from the lower oil price, without having the negative side effect of a higher tax rate. For the individual country, the amount of oil use will be lower, and thus output lower as well, since $p_t(1+\tau_{it})$ is now higher than under laissez-faire. Unless the country is "large", it is therefore unlikely that the unilateral domestic tax increase will deliver higher utility. Thus, we have carbon leakage out of the specific region, with the global amount of carbon consumption left unaffected: there is "perfect leakage". Thus, any given oil-consuming region will prefer to have other regions raise taxes on oil consumption. Some sort of coordination among oil consumers is thus necessary in order for them to become better off.

The irrelevance result—taxes in the oil-consuming regions cannot influence oil production—does, as pointed out earlier, depend on logarithmic utility functions assumed for oil producers. What about the generality of the result that there is perfect leakage? It does, of course, rely on the same assumption, i.e., logarithmic utility for oil producers. Of course, if energy can be produced from coal instead (as discussed in Section 5.2), coal having a positive marginal cost so that the total amount of coal being extracted depends on demand, and thus on taxes. There would then be less than perfect leakage, since lower world demand would imply lower world supply, and what then would become important is the nature of the marginal cost curve.

7 Concluding remarks

So far, it has proven very difficult to reach international agreements on a global emissions policy. In order to achieve progress on this front, it is central to understand the distributional consequences of climate change as well as of policies aimed at mitigating the problem. In this paper, we constructed a simple and transparent model allowing us to quantify how key features of heterogeneity between different regions of the world affect their preferences over different policy options. We show that in absence of international transfer mechanisms, Pareto-improving policies to curb climate change may not exist, making it easy to understand the difficulty of coming to international agreements.

Our analysis also shows that the taxation of oil use in oil-importing countries may be a completely toothless weapon against the threat of climate change unless the tax receipts are transferred to oil-producing countries. On the other hand, the tool of taxing oil producers proportionately and giving the proceeds back in a lump-sum manner is an interesting experiment since this scheme proposes no net redistribution across oil consumers and oil producers; thus, it may appear less difficult to implement politically than some other policies.

Our model is very stylized and some of its underlying assumptions are motivated by convenience rather than by an aim to maximize realism. As discussed in Golosov et al. (2012) in the context of a one-region model, though, many of these assumptions can likely be relaxed without significantly altering the main results. These include preferences, the technology for producing output and capital's depreciation structure, the carbon cycle, how damages appear, and the market structure in the energy market. The multi-region model developed here, however, introduces new assumptions (regarding the nature of the different regions and how they interact) and sharpens others. Let us therefore briefly mention an incomplete list with some extensions that we think would be especially fruitful.

First, consider the supply of energy. We focus on oil here: a source of energy that is efficient and very inexpensive to produce. There are strong reasons to introduce coal into the analysis: coal is much more abundant—by the current estimates—and hence constitutes a larger potential threat to the climate. Coal is also more costly to produce per unit of energy provided, which makes taxes more potent: a tax on consumers would now have an effect on fossil fuel use. Relatedly, one ought to introduce "clean" (non-fossil) energy, both because it is increasing in importance and because it may offer a way of resolving our climate problems—if it becomes inexpensive enough to produce. In particular, R&D efforts aimed at improving clean technology would be central for the analysis, and they are likely straightforward to analyze in our simple model. In Golosov et al. (2012), we argue that a highly tractable setting would have a CES technology in oil, coal, and clean energy, along with an assumption that coal and clean energy are produced linearly from labor—with productivity levels that could evolve over time (and, potentially, also be endogenous). Many interesting questions can be addressed with such a setting, such as (i) how to interpret the striking heterogeneity in intensities with which the three inputs are used in different parts of the world and (ii) how different taxes/subsidies across the three inputs would influence both overall efficiency and the distribution of welfare across regions. An extension of the present model to the study of technology transfer across regions in the context of clean innovation also appears to be within reach.

Second, to buy tractability, we have assumed that oil producers cannot save abroad. This assumption is critically behind the stark result here that taxes on oil consumers are ineffective. It would therefore be very interesting to relax it. A straightforward relaxation would necessitate numerical model solution, but a two-region world—producers and consumers of oil—would still be an interesting setting to look at and rather easily analyzed. Thirdly, and relatedly, the autarky assumptions here are stark. It would be straightforward model oil consumers in a richer way, especially by looking at how some of them trade intertemporally and may even insure each other (say, the EU, or the U.S. states) either with explicit markets or transfer mechanisms.

Fourth, regarding how we formulate the preferences in different regions, we take a rather standard approach. Two reasonably straightforward extensions would involve introducing time-inconsistent preferences (arguably particularly relevant in intergenerational contexts) and ambiguity aversion.

Fifth, it would be very illuminating to introduce stronger nonlinearities in the analysis, capturing irreversibilities in damages and/or tipping points (for a discussion of tipping points, see, e.g., Lenton et al., 2008). Such extensions are more demanding but because the rest of the model is so tractable, such an analysis does appear to be within reach and the implications would be easy to identify. Relatedly, one can look at fat-tail uncertainty, an issue which would be even easier to study.

Finally, the simplicity of the present setup may even allow tractable political-economy modeling. As argued above, policy disagreement across regions is at the heart of the climate problem and the comparison of different institutional settings for policymaking in this arena would be very interesting to study further.

References

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

Babiker, Mustafa, Angelo Gurgel, Sergey Paltsev, and John Reilly (2008), "A Forward Looking Version of the MIT Emissions Prediction and Policy Analysis (EPPA) Model", The MIT Joint Program on the Science and Policy of Global Change Report Series, No. 161.

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

Caselli, Francesco and James Feyrer (2007), "The Marginal Product of Capital", *Quarterly Journal of Economics*, 122:2, 535–568.

Castro, Rui (2005), "Economic Development and Growth in the World Economy", *Review of Economic Dynamics*, 8:1, 195–230.

Clarke Leon, Cristoph Böhringer, and Tom Rutherford (2009), "International, U.S. and EU climate change control scenarios: Results from EMF 22," *Energy Economics*, 31 (Suppl 2), pp. S63-S306.

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

Gars, Johan (2012), "Indirect Effects of Climate Change", working paper, Institute for International Economic Studies, Stockholm University.

Gollier, Christian, and Martin Weitzman (2010), "How Should the Distant Future be Discounted When Discount Rates are Uncertain?", *Economics Letters*, 107:3, 350–353.

Golosov, Michael, John Hassler, Per Krusell, and Aleh Tsyvinski (2012), "Optimal Taxes on Fossil Fuel in General Equilibrium", NBER WP:17348.

Hassler, John, Per Krusell, and Conny Olovsson (2010), "Oil Monopoly and the Climate", *American Economic Review Papers and Proceedings* 100:2, 460–464.

Hassler, John, Per Krusell, and Conny Olovsson (2011), "Energy-Saving Technical Change", mimeo, Institute of International Economic Studies, Stockholm University.

Hémous, David (2012), "Environmental Policy and Directed Technical Change in a Global Economy: Is There a Case for Carbon Tariffs?", mimeo, Harvard University.

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

Kelly, David L. and Charles D. Kolstad (1999), "Integrated Assessment Models For Climate Change Control," In International Yearbook of Environmental and Resource Economics 1999/2000: A Survey of Current Issues, ed. Henk Folmer and Tom Tietenberg. Cheltenham, UK: Edward Elgar.

Leach, Andrew (2007), "The Welfare Implications of Climate Change Policy", *Journal of Environmental Economics and Management*, 57:2, 151–165.

Lenton, Timothy, Hermann Held, Elmar Kriegler, Jim Hall, Wolfgang Lucht, Stefan Rahmstorf, and Hans Joachim Schellnhuber (2008), "Tipping elements in the Earth's climate system", *Proceedings of the National Academy of Sciences*, 105:6, 1786–1793 .

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

Nordhaus, William, (1977), "Economic Growth and Climate: The Carbon Dioxide Problem," The American Economic Review, 67, 1, 341–346.

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

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

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

Nordhaus, William (2011), “Integrated Economic and Climate Modeling”, Cowles Foundation Discussion Paper No. 1839.

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

Paltsev, Sergey, John M. Reilly, Henry D. Jacoby, Richard S. Eckaus, James McFarland, Marcus Sarofim, Malcolm Asadoorian, and Mustafa Babiker (2005), “The MIT Emissions Prediction and Policy Analysis (EPPA) Model: Version 4”, The MIT Joint Program on the Science and Policy of Global Change Report Series, Report No. 125.

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

Sinclair, Peter (1992), “High Does Nothing and Raising is Worse: Carbon Taxes Should Be Kept Declining to Cut Harmful Emissions”, *Manchester School*, 60.

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

Spiro, Daniel (2012), “Resource Extraction, Capital Accumulation and Time Horizon”, working paper, Institute for International Economic Studies, Stockholm University.

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

van der Ploeg, Frederick and Cees Withagen (2010), “Is There Really a Green Paradox?”, CESifo Working Paper Series 2963.

Weitzman, Martin (2009), “On Modeling and Interpreting the Economics of Catastrophic Climate Change”, *Review of Economics and Statistics*, 91.