

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

John Hassler and Per Krusell

Institute for International Economic Studies, Stockholm University

November, 2011

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 decent 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 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 use of gasoline causes emissions of carbon dioxide into the atmosphere and a consequent global heating, thus imposing a cost that not only hits all currently alive humans but also future generations.¹ The question is how much 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. As for the latter, it is very important to have a fully micro-founded model as any welfare calculations of different policy alternatives—such as different levels of

¹John also owns a bicycle, which he fortunately uses more often than the motorcycle.

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

We build a structure which, apart from the key factor of heterogeneity, is a simple version of the model in Golosov et al. (2011). 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. In the concluding section of this paper, we briefly review these assumptions—which ones appear critical for quantitative purposes and which ones do not. The appearance of, and endogenous development of, alternative clean energy is a key omission from the analysis, though the setting can be relatively straightforwardly extended to include it.

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. 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 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.

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.

This paper is by no means the first one to provide an integrated analysis of climate and the economy. Rather, the integrated-assessment literature, by now, is quite broad. It was pioneered by William Nordhaus, who built his DICE (single-region) and RICE (multiple-region) IAMs (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. Our setting is, as a matter of fact, a very close relative of Nordhaus’s RICE model, though one which does allow many more regions. His models, which need to be solved numerically, have been used and extended by many researchers and were a key input into the Stern report. They are not, strictly speaking, general-equilibrium models and are therefore not fully equipped to analyze different policy options in general equilibrium. Our approach here, aside from allowing us to go quite far analytically, is entirely in the “modern macroeconomic” tradition, thus relying on a fully micro-economically founded setting amenable for policy analysis. There are other general-equilibrium, DICE-like models in the literature but none using a large set of regions.²

In Section 2 we set up and analyze the basic model. Section 3 looks at taxation from an analytical perspective, whereas Section 4 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 5 looks at two extensions—one looking at carbon leakage due to policy differences across oil-consuming countries and the other at differences in energy intensities across regions—and Section 6 concludes.

2 Decentralized economy

We directly set up the decentralized economy, which 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.

2.1 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.

²References on one-region general-equilibrium models needed here.

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.

Preferences are

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

and the country'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.

This delivers oil choice according to

$$\nu A_{it} k_t^\alpha e_t^{\nu-1} = p_t,$$

or

$$e_t = \nu^{\frac{1}{1-\nu}} A_{it}^{\frac{1}{1-\nu}} k_t^{\frac{\alpha}{1-\nu}} p_t^{\frac{-1}{1-\nu}},$$

so that

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

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.

2.2 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

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

and the country's budget/resource constraint for consumption is

$$c_t + p_t R_{t+1} = p_t R_t,$$

with $0 \leq R_{t+1} \leq R_t$ and where we recall that p_t potentially is random.

This is a cake-eating problem with random returns; notice that we regard oil producers as not being endowed with anything but oil—they have no other production technology.

Defining the budget as

$$c_t + \tilde{R}_{t+1} = r_t \tilde{R}_t,$$

having defined r_t as $\frac{p_t}{p_{t-1}}$ and \tilde{R}_t as $R_t p_{t-1}$, 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.

2.3 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, we have

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

Thus,

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

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 $\mu > 1$ (for positive x s). 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.

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 is because of the

particular 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

$$p_t = \nu A_t k_t^\alpha ((1 - \beta)R_t)^{\nu-1},$$

net output will satisfy

$$\hat{y}_t = (1 - \nu)y_t = (1 - \nu)A_t k_t^\alpha ((1 - \beta)R_t)^\nu,$$

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. (2011). In Golosov et al. (2011), 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 its 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.

2.4 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 that 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. A linear depreciation schedule, with time-varying rates, appears to be a good approximation.

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

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

How these losses ought to be added up is far from obvious. We look at this issue next.

2.4.1 Assigning utility weights

In principle, one can assign utility weights to different regions and, given any set of weights, it will be relevant for the welfare evaluation whether or not the allocation is near the optimal allocation; in the almost-autarky equilibrium considered here there are strong restrictions on trade and insurance, and the extent to which these constrains a world planner is an open question. We consider some possible calculations below.

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

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

Thus, the same formula as in the Golosov et al. (2011) 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 across regions, consider again a given time period. Adding up all output losses across regions, the total is

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

i.e., a (net) *output-weighted* average of the damage coefficients γ_{it} . This calculation, however, relies on considering production in different countries as substitutable. Production is indeed if there are no restrictions on trade and if the world allocation is optimal. In the almost-autarkic allocation we consider, however, there are (implicit) such restrictions. Therefore, if one took a utilitarian perspective instead and 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 for a different reason. Here the output weights simply reflect the planner’s (subjective, and exogenous) weights placed on different regions.

2.4.2 Direct costs and benefits of interventions

In the specific model entertained here, consumption is a fixed share $1 - \alpha\beta - \nu$ of output in each region. Similarly, 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 , and 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_i s. 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.³

3 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.

3.1 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.⁴ 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

$$\left(-\tilde{R}_{t+1} + r_t \tilde{R}_t \right) \tau_t = T_t. \quad (1)$$

Now define the growth rate of net-of-tax rates as

$$\theta_t \equiv \frac{1 - \tau_t}{1 - \tau_{t-1}}$$

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

$$c_t = r_t \tilde{R}_t - \tilde{R}_{t+1},$$

³Below 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.

⁴If 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 $\tilde{R}_{t+1} = \beta \tilde{r}_t \tilde{R}_t$.

the Euler equation becomes

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

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

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

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.⁵ 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 $s_t = \beta\theta$.⁶ 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. 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

⁵In 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.

⁶It is also straightforward to analyze the case where a new permanent non-unitary value for θ is introduced next period with probability ρ . Then, we know that if the tax is introduced, $s_{t+j} = \beta\theta$ for $j \geq 1$. Thus,

$$s_t = \beta \left(1 - \rho + \rho \frac{1 - s_t}{1 - \beta\theta} \theta \right),$$

implying

$$s_t = \beta \frac{1 - \rho + \theta \frac{\rho}{1 - \beta\theta}}{1 + \beta\theta \frac{\rho}{1 - \beta\theta}}.$$

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

$$c_{pt} = \nu y_t \quad \text{and} \quad c_{ct} = (1 - \nu - \alpha\beta)y_t,$$

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!

Equilibrium utility is equal to

$$U = \sum_{t=0}^{\infty} \beta^t \log k_{t+1}$$

plus a constant (different constants for the two regions), since saving is proportional to consumption. Saving, in turn, satisfies

$$\log k_{t+1} = -\gamma S_t + \alpha \log k_t + \nu \log(1 - \beta\theta) + t\nu \log \theta$$

plus constants and variables unaffected by θ . Thus, we have

$$U = \sum_{t=0}^{\infty} \beta^t (-\gamma S_t + \alpha \log k_t + \nu \log(1 - \beta\theta) + t\nu \log \theta).$$

Noting that $U = (1/\beta) \sum_{t=0}^{\infty} \beta^t \log k_t - (1/\beta) \log k_0$, we thus have

$$(1/\beta) \sum_{t=0}^{\infty} \beta^t \log k_t = (1/\beta) \log k_0 - \gamma \sum_{t=0}^{\infty} \beta^t S_t + \alpha \sum_{t=0}^{\infty} \beta^t \log k_t + \nu \frac{\log(1 - \beta\theta)}{1 - \beta} + \nu \log \theta \sum_{t=0}^{\infty} \beta^t t.$$

This expression can easily be solved out for $\sum_{t=0}^{\infty} \beta^t \log k_t$, which determines utility. Carrying this out, it is clear that the part of utility that is influenced by θ is

$$-\gamma \sum_{t=0}^{\infty} \beta^t S_t + \nu \frac{\log(1 - \beta\theta)}{1 - \beta} + \nu \log \theta \frac{\beta}{(1 - \beta)^2},$$

where we have also used a well-known formula for $\sum_{t=0}^{\infty} \beta^t t$. It is possible to show that the first of these terms is strictly increasing in θ and that the remaining terms are a strictly concave function of θ with a maximum at $\theta = 1$. It follows from these facts that utility is maximized for a θ strictly above 1: there is no first-order loss from raising θ but there is a first-order

gain. To see that the first term is increasing in θ , note that since $S_t = \sum_{s=0}^t E_s(1 - d_{t-s})$ we can reorder the E terms to observe that

$$\sum_{t=0}^{\infty} \beta^t S_t = \sum_{t=0}^{\infty} \beta^t E_t x$$

plus a constant, where $x \equiv \sum_{t=0}^{\infty} \beta^t (1 - d_t) > 0$. Hence, since $E_t = (1 - \beta\theta)(\beta\theta)^t R_0$, the precise depreciation structure is not important for how θ influences this expression: it suffices to determine how θ influences $(1 - \beta\theta) \sum_{t=0}^{\infty} (\beta^2\theta)^t$. This quantity equals $(1 - \beta\theta)/(1 - \beta^2\theta)$, whose derivative with respect to θ is $-\beta(1 - \beta)/(1 - \beta^2\theta)^2 < 0$. Thus the effect on utility of increasing θ is positive from this perspective: it delays the damages, and discounting makes this beneficial. Finally, to see how the two last terms are influenced by θ we need to examine the behavior of $\log(1 - \beta\theta) + \frac{\beta}{1-\beta} \log \theta$. Its derivative is equal to

$$-\frac{\beta}{1 - \beta\theta} + \frac{\beta}{1 - \beta} \frac{1}{\theta}.$$

Each of these terms are decreasing in θ and at $\theta = 1$ the expression is zero.

3.2 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.,

$$\frac{1}{c_{it}} = \beta \frac{r_{i,t+1}}{c_{i,t+1}}$$

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,

$$c_t + k_{t+1} = (1 - \hat{\nu}_t) y_{it}.$$

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

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

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 for oil thus reads

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

so that

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

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. In order to delay extraction, one would need to use a path of increasing (decreasing) tax rates on oil to the extent the oil producer has more (less) than logarithmic curvature in utility.

3.3 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.⁷ 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 5.1, 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.

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

4.1 Calibration

We use as the same carbon cycle as in Golosov et al., 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. 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 3.91%, i.e., 2.61 times larger, whereas in OECD Europe, they are 2.83%, i.e., 1.89 times larger. The corresponding

⁷Recall that this result comes about because oil producers have logarithmic utility.

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 .

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.⁸

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. 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-region setup in Golosov et al.—and adopt the parameter values from that study. Those parameter values imply a damage of 3.1% of output for a doubling of the CO₂ concentration.

4.2 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 3.1—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., 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 5 different regions we consider—the four oil-consuming regions and the oil-producing region. The table illustrates

⁸Caselli and Feyrer (2007) study the dispersion of the marginal products of capital across countries and argue that the dispersion is not large.

how the faster the rate of decline in taxes implies that oil use is delayed, hence leading to a lower current oil use.

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%

Table 1: initial oil use relative to laissez faire

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 consequently 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 .

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%

Table 2: compensating variations and oil use relative to laissez faire

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.⁹

The analysis summarized in Table 2 reveals that the compensating variations are fairly small quantitatively, as are the reductions in damages that can be achieved with this policy. Is this a robust finding? The explanation for the finding in the present context is a combination of assumptions: fossil fuel is a necessary input in production, it is not costly to extract, and it exists in finite supply. This means (i) that all fossil fuel will eventually be used and (ii) that tax policy can only change the time profile of oil use and not the total amount consumed. A more realistic model should feature both oil and coal, where the latter exists in large but finite supply, though with relatively high extraction costs.¹⁰ With an additional assumption

⁹The one-region world is not an exact aggregation of the heterogeneous-region world, and hence the optimal θ s differ slightly.

¹⁰Such a setting is analyzed in Golosov et al. (2011).

that an economically viable substitute for fossil will eventually be developed, such a model would imply that policy can affect the total amount of fossil fuel extracted. The more would be extracted, the higher are damages, and the higher is the value of an active carbon policy. How much would be extracted would critically depend both on how much coal there is—at an economically viable extraction cost—and how soon a clean “backstop” technology is available. There are reasonably reliable estimates of the former, discussed in Golosov et al. (2011), but how soon effective clean technologies appear is anybody’s guess. In sum, our simple setting here is not robust to reasonable extensions of the setting, but it should be straightforward to add the missing features. With a potential for a large total fossil-fuel extraction, there is also a scope for stronger disagreements on policy across regions of the world.

We also report the global transitions for the different policy alternatives considered in our experiments. In Figure 1, we plot global output relative to laissez faire for three values of θ : 1.03, 1.075, and 1.098.




Figure 1: output paths relative to laissez faire for different policies

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 2 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 .




Figure 2: global temperature paths for different policies

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 3 displays the yearly interest rates in the different regions for $\theta = 1.03$.




Figure 3: interest rates across regions, globally optimal policy

There are significant differences, 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.¹¹ Figure 4 shows the interest rates for the laissez-faire economy.

¹¹See, e.g., Castro (2005) for a model of this sort.

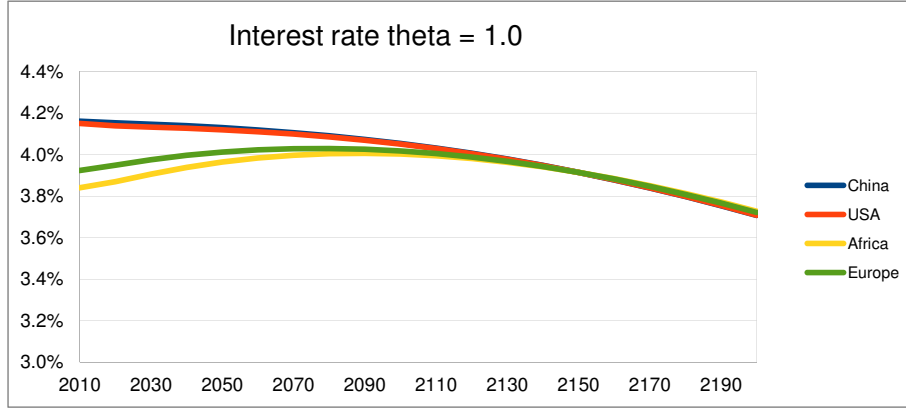


Figure 4: interest rates across regions, laissez faire

5 Extensions

We look at two extensions: the possibility that taxes differ across oil-consuming regions and that energy intensities differ across regions.

5.1 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 3.2 above.

The Euler equation reads

$$\frac{1}{c_{it}} = \beta \frac{r_{i,t+1}}{c_{i,t+1}},$$

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 3.2, 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 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

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

implying

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

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 4.2), 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.

5.2 Differences in energy intensities across time and countries

There are large differences in the amount of energy used per unit of output produced in different countries. In this model, we can think of this as being captured by differences in ν , our Cobb-Douglas technology parameter. Moreover, there are systematic differences across levels of development in energy intensity, with low-income countries using much more energy to produce each unit of output. Thus, one might also want to consider ν to be dependent on time, or technological development.

5.2.1 Regional decisions given world prices

Adopting the firm’s optimal energy purchasing strategy, we obtain This delivers oil choice according to

$$\nu_{it} A_{it} k_t^\alpha e_t^{\nu_{it}-1} = p_t,$$

yielding

$$e_{it} = \nu_{it}^{\frac{1}{1-\nu_{it}}} A_{it}^{\frac{1}{1-\nu_{it}}} k_t^{\frac{\alpha}{1-\nu_{it}}} p_t^{\frac{-1}{1-\nu_{it}}},$$

and output satisfying

$$y_{it} = \nu_{it}^{\frac{\nu_{it}}{1-\nu_{it}}} A_{it}^{\frac{1}{1-\nu_{it}}} k_t^{\frac{\alpha}{1-\nu_{it}}} p_t^{\frac{-\nu_{it}}{1-\nu_{it}}}.$$

The Euler equation, assuming a guess of $k_{i,t+1} = s_{it} y_{it}$, reads

$$\frac{1}{(1 - \nu_{it} - s_{it})y_{it}} = \frac{\beta \alpha A_{i,t+1} k_{i,t+1}^{\alpha-1} \nu_{i,t+1}^{\frac{\nu_{i,t+1}}{1-\nu_{i,t+1}}} A_{i,t+1}^{\frac{\nu_{i,t+1}}{1-\nu_{i,t+1}}} k_{i,t+1}^{\frac{\alpha \nu_{i,t+1}}{1-\nu_{i,t+1}}} p_{t+1}^{\frac{-\nu_{i,t+1}}{1-\nu_{i,t+1}}}}{(1 - \nu_{i,t+1} - s_{i,t+1})y_{i,t+1}}.$$

This becomes

$$\frac{k_{i,t+1}}{\alpha \beta y_{it}} = \frac{\nu_{i,t+1}^{\frac{\nu_{i,t+1}}{1-\nu_{i,t+1}}} A_{i,t+1}^{\frac{1}{1-\nu_{i,t+1}}} k_{i,t+1}^{\frac{\alpha}{1-\nu_{i,t+1}}} p_{t+1}^{\frac{-\nu_{i,t+1}}{1-\nu_{i,t+1}}}}{y_{i,t+1}} \frac{1 - \nu_{it} - s_{it}}{1 - \nu_{i,t+1} - s_{i,t+1}},$$

so that

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

A special case is that where regions permanently have different intensities, a case where s_{it} becomes equal to $\alpha\beta$ for all countries.

5.2.2 Equilibrium

The equilibrium oil price is again given by market clearing: the amount of oil extracted equals the sum of all the oil consumed. Thus, we have

$$(1 - \beta)R_t = \sum_i \nu_{i,t}^{\frac{1}{1-\nu_{i,t}}} A_{it}^{\frac{1}{1-\nu_{i,t}}} k_{it}^{\frac{\alpha}{1-\nu_{i,t}}} p_t^{\frac{-1}{1-\nu_{i,t}}}.$$

This equation no longer has a closed-form solution for p_t , but it is at least not forward-looking—it is one equation in one unknown.

5.2.3 A preliminary assessment, and oil vs. coal

The data on energy shares suggests relatively minor differences across the major regions considered in the quantitative analysis above. Spending on energy in Europe, gross of taxes, is a slightly larger share of GDP in the U.S. than in Europe and China. It would seem relevant to include these differences in the empirical evaluation of the model. However, a more relevant difference from this perspective is the fact poorer countries use more coal than oil. Oil is about five times more expensive per carbon unit, presumably because it provides more efficient/useful energy services. With a similar share of GDP spent on the sum of oil and coal across countries but a much larger share of coal in poorer countries, we thus see that emissions per GDP unit will be significantly higher in poor countries. To model these differences in a more satisfactory manner, one would like to not only have both oil and coal as energy inputs, but they would need to enter production in such a way that poorer countries choose more coal than oil.¹² To conclude, adding differences in energy shares makes the model solution only slightly more complicated and will at least marginally change the quantitative analysis. However, from the perspective of climate change it would seem necessary to at the same time generalize the structure to allow differences in coal and oil use across countries, an enterprise that is beyond the scope of the present analysis.

¹²This could be accomplished in several ways. One is to make the two inputs imperfectly substitutable and have different complementarity properties with other inputs; oil could, for example, be more complementary with capital than coal would be. Another is to make them perfect substitutes but simply assume that the relative efficiency of coal is lower in poorer countries, an assumption only a step away from the setting in Golosov et al. (2011), which uses perfect substitutability.

6 Concluding remarks

Emission of greenhouse gases produces a global externality and calls for a global policy. So far, however, it has proven very difficult to reach international agreements on such a 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 models also shows that taxation of fossil fuel 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.

Our model is very stylized and some of its underlying assumptions are motivated by convenience rather than by an aim to maximize realism. Let us therefore conclude by a discussion of where we think we need to generalize the analysis. We have used logarithmic utility and noted that the finding that fuel taxation in oil-consuming rather than oil-producing countries has no effect is sensitive to this, although the basic mechanism of counteracting income and substitution effects would remain with more general preferences and is unlikely to radically alter our conclusion. Other preferences, as well as departures from Cobb-Douglas production functions and full depreciation, would imply different transitional dynamics but would not likely change the key conclusions in the paper (see Golosov et al. (2011) for a more thorough discussion on this). There is, however, a number of assumptions that have the potential of changing the results substantially.

First and perhaps most obviously, we have quite limited knowledge about the long-run consequences of climate change. Substantial efforts by natural as well as social scientists are devoted toward advancing our knowledge in this respect. If they come along, new insights can easily be incorporated into our analysis.

Second, we have assumed that fossil fuel is (i) necessary for production, (ii) exists in finite supply, (iii) can be extracted at zero costs, and (iv) is sold under perfect competition. These assumptions are highly relevant for the qualitative as well as quantitative results here. The assumptions are obviously debatable a priori and they also have some implications that are difficult to reconcile with historical data. In particular, they imply that aggregate fossil fuel use should follow the famous Hotelling rule, which implies falling quantities and increasing prices, a pattern notoriously at odds with reality. We therefore need to improve our understanding of how the fossil fuel market functions. In other work, we have explored some alternative assumptions to the ones used in this paper (Golosov et al., 2011, Hassler, Krusell, and Olovsson, 2010 and 2011).

Third, we have treated technological change as exogenous and following a smooth trend. Although we have argued in other work that this may be of little consequence for the issue of how large optimal fossil fuel taxes ought to be, it is obvious that normative issues regarding other policies, like R&D subsidies, as well as predictions about future fossil fuel use depend on technological development and how it interacts with the policy considered in this paper.

Finally, we have assumed that regions cannot insure each other or transfer resources between them. Implicitly, we have also assumed that agents within regions have access to perfect markets, allowing each region to be inhabited by a representative agent. Although these assumptions have some amount of similarity with reality, in ongoing work we are relaxing these stark assumptions, thus allowing a more realistic description of how agents interact within and across regions.

References

*Apologies: we will add a large number of references later
We simply wanted to make sure the submission is not delayed more*

- Castro, R. (2005), “Economic Development and Growth in the World Economy”, *Review of Economic Dynamics* 8:1, 195–230.
- Golosov, M., J. Hassler, P. Krusell, and A. Tsyvinski (2011), “Optimal Taxes on Fossil Fuel in General Equilibrium”, working paper.
- Hassler, John, Per Krusell, and Conny Olovsson (2010), “Oil Monopoly and the Climate” (joint with John Hassler and Conny Olovsson), *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.
- Nordhaus, William (2007), “To Tax or Not to Tax: The Case for a Carbon Tax”, *Review of Environmental Economics and Policy*.
- Nordhaus, William (2008), *A Question of Balance: Weighing the Options on Global Warming Policies*, Yale University Press, New Haven, CT.
- Nordhaus, William and Joseph Boyer (2000), *Warming the World: Economic Modeling of Global Warming*, MIT Press, Cambridge, MA.
- Stern, Nicholas (2007), *The Economics of Climate Change: The Stern Review*, Cambridge University Press, Cambridge, UK.