

Economic Policy

62nd Panel Meeting
Hosted by the Banque Centrale du Luxembourg
Luxembourg, 16-17 October 2015

Climate Policy

John Hassler (Institute for International Economic Studies (IIES) and CEPR)
Per Krusell (IIES and CEPR)
Jonas Nycander (Stockholm University)

The organisers would like to thank the Banque Centrale du Luxembourg for their support.
The views expressed in this paper are those of the author(s) and not those of the supporting organization.

Climate policy

John Hassler*, Per Krusell† and Jonas Nycander‡

September 26, 2015

*Hassler: Institute for International Economic Studies (IIES) and CEPR.

†Krusell: IIES, CEPR, and NBER.

‡Nycander: Dept. of Meteorology and Bolin Centre for Climate Research, Stockholm University

1 Introduction

In light of our recent research, in this paper we present our views on how climate policy ought to be conducted. We summarize these views with six points:

1. *The optimal carbon tax is rather modest.* We judge an appropriate—from a global perspective—tax on carbon to be on the order of 25 cents per liter of gasoline.
2. *It's (almost) all about coal.* The (estimated) reserves of oil and natural gas are small relative to those of coal and would only increase global temperatures rather modestly. The using up of a large part of our coal reserves, in contrast, presents a major threat to our climate.
3. *A carbon tax is to be preferred over a quantity-based system.*
4. *The optimal tax on carbon does not appreciably harm growth.*
5. *Subsidies to green technology are beneficial for the climate only to the extent that they outcompete coal.* We also argue that they may not even be necessary if an optimal carbon tax is used.
6. *At this moment in time, we judge a carbon tax to be politically feasible.* One often hears that carbon taxes are politically infeasible; we argue that they are likely not.

Our paper is designed to explain and support these points. In fact, the bulk of the paper builds a background by reviewing the recent research using integrated assessment models, i.e., models which jointly describe the natural-science aspects of climate change with the economic ones. Although our discussion here contains some qualitative arguments, we place significant emphasis on quantitative conclusions from the literature. For this reason we will briefly summarize the integrated assessment models used and how their parameters are calibrated.

An economic model of climate change driven by the emission of carbon dioxide needs to describe three phenomena and their dynamic interactions: i) economic activity, ii) carbon circulation and iii) the climate. The economy is needed to model emissions and economic effects of climate change. The carbon circulation is needed to model how emissions translate into carbon dioxide concentrations at different points in time. Finally, one needs to understand how the climate is affected by the carbon dioxide concentration. Of course, all these systems are immensely complicated. In order to combine the mechanisms from natural science into an integrated model useful for conducting economic policy analysis, the different model blocks need to be expressed in a very simple form. The key complication is that, in a model with forward-looking economic agents, the outcome at any point in time depends on expectations about the subsequent future. Loosely speaking: the present depends on the future, a reverse causality that never arises in a natural science model, however dynamically complicated.

Next, in Section 2 and as a key background for the full integrated economy-climate model we will describe later, we describe a very simple, yet quantitatively reasonable, framework capturing the key natural-science mechanisms. This framework in particular is simple enough that it can be used in the forward-looking economy-climate model that we use for our policy

analysis. On another level, and perhaps more importantly, the description of the climate and carbon cycle determination serves to highlight the inescapable scientific mechanisms resulting in global warming. These mechanisms are not controversial per se but some quantitative aspects are uncertain; this will be highlighted in our presentation.

Section 3 goes over a key element behind the quantitative policy analysis, namely the measurement and modeling of damages from climate change. In Section 4 we then briefly discuss two simple integrated assessment models, one dynamic and one static. The static model captures most of the essence and builds very straightforwardly on the elements of a typical intermediate course in microeconomics; the dynamic model only adds marginally to this setting but allows a formal discussion of discounting, which has been discussed at length in this literature. Section 5 then goes over our policy messages and defends them based on the analysis in the earlier chapter and some additional, less formal arguments. Section 6 concludes.

2 The natural-science elements

We begin by discussing the object of interest—the climate—and then the determinants of the main climate driver, namely, the atmospheric carbon concentration. The damages from global warming also contain natural science elements, but these are discussed in the next section.

2.1 The climate

A natural definition of the global climate is the distribution of weather events, i.e., realizations of e.g., temperature, precipitation, wind speed and ice coverage, over time and space. Clearly, a complete description of the global climate is infeasible. However, it turns out that there is a key state variable describing the climate: the global mean surface temperature. In simulations with coupled climate models it turns out that the changes of many other climate parameters, for example precipitation and regional temperatures, are approximately proportional to the change of the global mean surface temperature. Furthermore, much of the empirical data on the economic effect of climate change focus on the temperature. We will therefore briefly describe how the emission of carbon dioxide affects the global mean temperature.

2.1.1 The energy balance and temperature

The earth is heated by incoming short-wave radiation from the sun, and cooled by outgoing long-wave (infrared) radiation. A simple energy balance model describes how the global mean temperature changes over time as a result of these energy fluxes. The incoming short-wave radiation is 340 Wm^{-2} when averaged over the surface of the earth, and one third of this is directly reflected back to space. In equilibrium, the resulting net short-wave radiation must be balanced by the outgoing long-wave radiation.

We now consider what happens after a change in the energy budget. We take as a starting point a pre-industrial equilibrium state in which the incoming and outgoing energy fluxes were equal, and the global mean temperature therefore constant. We analyze a positive perturbation of the energy budget by the amount F (measured in Wm^{-2} , and called *forcing*). Because of the

perturbation, the incoming energy flux is larger than the outgoing flux, which leads to increasing temperature. This is described by the equation

$$\frac{dT}{dt} = \sigma (F - \kappa T) \quad (1)$$

where $T(t)$ denotes the increase of the global mean temperature (measured in degrees °C) compared to the pre-industrial steady state. The forcing $F(t)$ is determined by other equations that describe the carbon cycle, and the term κT describes the fact that a higher temperature leads to a larger outgoing energy flux. The parameter σ determines how quickly the temperature changes due to a given imbalance in the fluxes. It is inversely proportional to the heat capacity of the climate system, which is dominated by the ocean. If F is constant, the solution to (1) with the initial condition $T(0) = 0$ (the pre-industrial state) is

$$T(t) = \frac{F}{\kappa} (1 - e^{-\sigma \kappa t})$$

Asymptotically, as $t \rightarrow \infty$, the climate approaches a new steady state:

$$T_{\infty} = \frac{F}{\kappa}. \quad (2)$$

Much climate research is devoted to determining the key parameters κ and σ . If the circulation and composition of the atmosphere would not change as the temperature changed, κ could be obtained from relatively simple radiation calculations, similarly as when calculating how black-body radiation depends on the temperature. This gives $\kappa = 3.2 \text{ Wm}^{-2}/^{\circ}\text{C}$, which would imply that a perturbation of the energy balance by 1 Wm^{-2} increases the equilibrium temperature T_{∞} by 0.3°C . Sometimes this simple mechanism is referred to as the ‘Planck feedback’. Due to various other feedbacks to be discussed below, κ is likely to be smaller than this value, i.e., the outgoing energy flux increases less with increasing temperature than what is implied by the Planck feedback. A given forcing then results in a larger temperature increase.

2.1.2 Carbon dioxide and the greenhouse effect

Now consider the reason why a higher CO_2 concentration changes the energy balance, i.e., implies a positive forcing. The atmospheric gases are transparent to the solar short-wave radiation, whose maximum intensity is in the visible wavelength range. The most abundant gases, which consist of molecules with one or two atoms (such as nitrogen and oxygen), are also transparent to the outgoing long-wave radiation. However, gases consisting of molecules with three or more atoms, such as carbon dioxide, water vapor and methane, strongly absorb long-wave infrared radiation, which leads to a positive forcing F . Gases with this property are called *greenhouse gases*. The mechanisms behind this are well understood and easy to verify experimentally. Even a small concentration of such gases has a large effect on the energy balance of the earth.

The effect of the CO_2 concentration on the energy balance is well approximated by the function

$$F = \frac{\eta}{\ln 2} \ln \left(\frac{S}{\bar{S}} \right), \quad (3)$$

where S and \bar{S} represent the actual and preindustrial atmospheric CO₂ concentrations, respectively. The present concentration is 400 ppm, and the preindustrial value is 280 ppm. The exact value of the parameter η is not known, but a value of 4.3 Wm⁻² may be used.¹ This means that a doubling of the CO₂ concentration leads to the forcing $F = 4.3 \text{ Wm}^{-2}$. Since the perturbation is related to the *relative* change in CO₂ concentration, the formula is valid regardless of the units used for the CO₂ concentration. We will use the unit GtC, billions of tons of carbon in the atmosphere as a whole. The present value of S is then approximately 840GtC, and the value of \bar{S} is 600GtC.

Combining (2) with (3), we find a relation between the CO₂ concentration and the steady state temperature:

$$T_{\infty} = \frac{\eta/\kappa}{\ln 2} \ln \left(\frac{S}{\bar{S}} \right).$$

The ratio η/κ is the heating that would arise in steady state after a doubling of the CO₂ concentration. Using the Planck feedback gives $\eta/\kappa \approx 1.3^{\circ}\text{C}$. This is a modest sensitivity, and very likely a too low estimate of the overall sensitivity of the global climate. The reason is that there are many other feedbacks. A higher temperature will increase the atmospheric water vapor concentration, which adds to the forcing from CO₂. A higher temperature will also change the size of the global ice cover and cloud formation, both having an effect on the energy budget. Formally, we can include the feedbacks in the energy budget by adding a term xT , giving

$$\frac{dT}{dt} = \sigma (F + xT - \kappa T),$$

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

$$T_{\infty} = \frac{\eta}{\kappa - x} \frac{1}{\ln 2} \ln \left(\frac{S}{\bar{S}} \right) \quad (4)$$

The coefficient $\lambda \equiv \eta/(\kappa - x)$ is called the equilibrium climate sensitivity and captures the response in the global mean temperature to a doubling of the CO₂ concentration.² Theoretically, we cannot rule out either $x < 0$ or $x \geq \kappa$. In the latter case, dynamics would be unstable and λ not well defined. This does not seem to be consistent with historical evidence. Also $x < 0$ is hard to reconcile with the observation that relatively small historical changes in the forcing appear to have had substantial effects on the climate. However, within these bounds a large degree of uncertainty remains. According to the latest judgements of the evidence, IPCC sets a likely range for λ to $3^{\circ}\text{C} \pm 1.5^{\circ}\text{C}$.

2.2 The carbon cycle

The global carbon circulation system is, of course, very complicated. However, we argue that a simple summary of how carbon depreciates is sufficient to give important insights into how the carbon circulation affects the economics of climate change. We base our summary description on IPCC (2007) and Archer (2005) who claim that

¹See Myhre et al. (1998).

²Natural scientists attach a different meaning to the word *equilibrium* than economists. A translation to the language of economics would be *steady state*.

- one share (about 50%) of the emitted CO₂ leaves the atmosphere quite quickly (within a few years to a few decades),
- another share (around 20–25%) stays very long (thousands of years) until CO₂ acidification has been buffered, whereas
- the remainder decays with a half-life of a few centuries.

We model this by specifying a carbon depreciation function $d(s)$ such that $1 - d(s)$ describes the share of the emitted carbon that remains in the atmosphere after s units of time. Targeting the above simple summary, we set

$$1 - d(s) = \varphi_L + (1 - \varphi_L) \varphi_0 (1 - \varphi)^{s/10}, \quad (5)$$

and the parameters $\{\varphi_L, \varphi_0, \varphi\} = \{0.2, 0.38, 0.023\}$ for s measured in years. Below is a graph of the function $1 - d(s)$.




Figure 1. Share of an emitted unit of carbon that remains in atmosphere against time since it was emitted.

It is important to note that the linear depreciation structure is a simplification. Specifically, the rate of depreciation as well as the share that stays thousands of years depends on the size of emissions. To give an illustration of how sensitive the depreciation is to the emission scenario, IPCC (2013) shows that while around 20% of an emitted pulse of 100–1000GtC remains in the atmosphere after 2000 years (in line with the summary above), almost 40% remains in the atmosphere after 2000 years of a pulse of 5000 GtC, i.e., a much more dramatic scenario.³ Even after 10000 years, the share in the atmosphere is above 20%. The half-life of the of the third share in the summary above is about twice as large in this scenario.⁴ Later, we will use the simple linear model of carbon depreciation, but it is important to note that the parameters may need to be adjusted if extreme emission scenarios are considered.

³See IPCC (2013) Chapter 6, Box 6.1.

⁴Archer et al. (2009).

3 Damages from climate change

Climate change is believed to significantly affect human welfare, and to do so for a long time to come. However, the assessment and, more generally, quantification of this belief, is a huge challenge. In fact, we believe that the measurement of “damages” from climate change is the single weakest element of the climate-economy package that scientists have to offer as a background for policymakers. Unfortunately, the missing knowledge includes both qualitative and quantitative parts. The qualitative parts involve the forms the damages take, and these forms are important for understanding how easy it might be to adapt to climate change. Depending on whether the bulk of the damages are due to a rising sea level or a “mere” temperature increase, very different responses are appropriate. The quantitative issue is, moreover, central: how much emissions should be reduced (or how high taxes should be) will naturally depend on the size of the damages. There may be “tipping points” in damages, so that small temperature increases are not so costly whereas increases above some threshold are close to disastrous; the existence of such thresholds are obviously crucial for policymakers. Damages, moreover, involve heterogeneous impacts across the world (and more generally across groups), making distribution policy a central tool. But as a summary statement, it is fair to say that we know very little—in fact, we have only recently begun to accumulate knowledge—in all these areas. This needs to be kept in mind throughout the policy discussion.

Damage measurements can perhaps usefully be put into two categories: top-down and bottom-up. The former looks at data on observable aggregates (such as output or mortality) and tries to relate them to climate, or merely temperature. An advantage of this approach is that one looks at aggregates, thus obtaining relatively broad-measured impacts. A disadvantage is that the method does not examine the mechanisms through which climate affects the aggregate: it does not indicate the specifics of the channel nor whether the link between the climate and the aggregate involves adaptation (and, to the extent there is adaptation, how costly adaptation was). The bottom-up approach looks at specific damages (say, output by narrowly defined sector or population segment), allowing a more careful study of mechanisms and adaptation. On the other hand, climate policy should be based on taking all damages into account, and thus coverage becomes a major issue. So far, there are relatively few studies of different disaggregated impacts, and they tend to be region-specific: the “world map of damages” is so far to a large extent full of uncharted territory. We will briefly go over the main results coming out of (a subset of) this literature, each one in turn. A key output of this discussion is a (qualitative and quantitative) formulation of a “damage function” that will be used in the integrated assessment model later.

3.1 Top-down studies

Researchers have used both cross-sectional and panel data. Cross-country regressions of outcome variables such as GDP on country-specific temperature suggest a clear negative relation, but of course omitted-variable bias can be important in such regressions (as for example institutional quality appears to be rather strongly correlated with temperature). Mendelsohn et al. (1994) argue, for the case of agriculture, that a regression using regions within a country, with a

fixed effect per country, allows a sensible control for institutions—under the assumption that the institutions within a country are very similar—and thus climate variability across regions within a country allows identification of a negative effect of temperature on output.

An influential study of temperature variations over time, in a broad cross-section of countries, is Dell, Jones, and Olken (2012). Their focus is more on short-term variations in temperature, and climate change is therefore arguably not captured well by these regressions. They find rather small effects of temperature increases on output, but they do find effects on the growth rate of output. This finding is potentially important, since growth-rate effects imply much more potent effects on human welfare. An additional finding in this study is that of heterogeneous losses from temperature change: the growth-rate losses are only observed in countries that are poor. A later study using shorter, though more disaggregated, data, is that in Krusell and Smith (2014) who find growth rate effects of the same sign, but these are statistically insignificant; moreover they find significant negative level effects on output of higher temperature and no heterogeneity in their estimates across poor and rich regions. Thus, these studies taken together clearly show negative effects of temperature increases of a magnitude similar to that found in Mendelsohn et al. (1994), but there is still significant uncertainty about the specific effects. A longer panel, thus potentially identifying climate, as opposed to temperature, change is that in Bluedorn et al. (2010). This paper finds rather weak, and statistically uncertain, effects of climate on current income, but non-monotonic effects historically.

3.2 Bottom-up studies

Nordhaus’s main calibration of his aggregate, as well as disaggregated, damage functions (we will discuss these functions below) is based on adding up detailed microeconomic estimates of the effects of temperature change. These damages take a variety of forms (mention a list) and amount to a total of 0.48 percent of output for a 2.5-degree warming. Ciscar et al. (2011, 2014) report detailed estimates for the European Union, covering a number of sectors in great detail. At a business-as-usual scenario leading to a 3.5-degree warming globally, damages in the EU due to climate change is estimated to 1.8% of GDP in the year 2080. This study is very ambitious in its coverage, but there are of course many other damages that it does not attempt to measure.

Nordhaus’s damage function, however, relies not only on the bottom-up estimate but also on survey evidence. Here, researchers (mostly from the natural sciences) were asked to estimate probabilities of various pre-specified events. This survey resulted in a probability of 6.8% that the damages from heating of 6 degrees centigrades are catastrophically large, defined as a loss of 30% of GDP. Nordhaus, moreover, calculates the willingness to pay for such a risk using a coefficient of relative risk aversion of 4 and adjusts the estimate up accordingly for use in an integrated assessment framework where uncertainty is not taken into account. Nordhaus thus adds the bottom-up information to that in the survey when selecting the parameters in his damage function. We will describe this damage function in the following section.

3.3 Damage functions

In this paper we will primarily discuss climate policy from a global perspective and not so much address issues of distribution. This is not because we do not believe they are important; on the

contrary, we do, both from the perspective of constructing an aggregate of the social cost of carbon and from a political-economy perspective.⁵ However, discussing heterogeneity carefully would necessitate an extension of the analysis which is hard to fit into the present paper. Hence, we now focus on how global damage functions used in the integrated assessment literature are usually modeled and calibrated.⁶

3.3.1 Nordhaus's damage function

Though damages appear in many places in the economy, Nordhaus early on adopted what has become the industry standard, namely a formulation where all damages appear in a factor that multiplies the aggregate production function of the economy. I.e., a damage is then expressed as lower total-factor productivity, TFP. To cut to the specifics, the multiplicative damage factor, D , that Nordhaus uses in his most recent work is

$$D(T) = 1 - \frac{1}{1 + 0.00267T^2} \approx 0.0267T^2. \quad (6)$$

This expression is increasing in T , global temperature, so that output is $(1 - D(T))Y$ net of damages (where Y defines output under no damages), no matter how (in what sectors, with what combination of production factors, etc.) output is produced. Also note that $D(T)$ is convex, so that the marginal damage factor is higher for higher temperatures.⁷ It is broadly believed that damages increase at a higher rate as the global temperature rises and Nordhaus's formulation is thus consistent with these beliefs. It should be added, however, that the size (and even presence) of the convexity has not been firmly established yet empirically.

It should be noted here that in terms of modeling, several other forms of damages (such as direct utility losses from higher temperature or higher depreciation of the capital stock) have the same analytical implications as the formulation with TFP damages. For a discussion of this equivalence, see Gars (2012).

3.3.2 A damage function expressed in terms of carbon concentration

It turns out, for the construction of a complete integrated assessment model, that a very valuable simplification can be achieved as follows: one can describe damages directly as a function of the level of atmospheric carbon concentration, rather than as a two-step function describing first how carbon concentration maps into temperature and then applying the damage functions above. The reason why this is a simplification is that the direct carbon-damage formulation can be calibrated with a functional form that is very analytically convenient. We noted above that there is a convexity in the temperature-damage relationship but a concavity in the carbon-temperature relationship (3) above and these two together imply that, over the range of carbon concentration

⁵For a discussion, see Hassler and Krusell (2012).

⁶Krusell and Smith (2015a) builds an integrated assessment model that focuses entirely on damages at a disaggregated level.

⁷In contrast, in our discussion of top-down damages above we referred to regression estimates as percentages of output without reference to the given temperature.

values that are empirically relevant, a linear-in-log relationship is a good approximation:

$$D(T(S)) \approx 1 - e^{-\gamma(S-\bar{S})}, \quad (7)$$

where γ is a constant. We refer to γ as an elasticity parameter because

$$\frac{1}{1-D} \frac{\partial D(T(S))}{\partial S} = \gamma.$$

Two different approximations have been made in equation (7). The first one is that T is treated as a unique function of S ; in effect, T is replaced by T_∞ . This means that the thermal inertia of the ocean is neglected, so that the temperature adjusts instantaneously to the CO_2 concentration. This approximation, which means that we model the effects of emissions on global warming as too large (over the transition period) appears acceptable since the focus is on the long-run. More specifically, a key purpose of the model is to analyze the social cost of carbon. Our assumption implies that damages from an emitted unit of carbon accruing near in time after the emission are somewhat exaggerated. Unless the subjective discount rate is high, this is not of large quantitative importance, however.⁸ The second approximation is that the functional form of $D(T_\infty(S))$ is simplified; the “error” here is thus relative to Nordhaus’s formulation but it is of second order over a large range of emission scenarios.

It will, quite naturally, turn out that the calibration of γ is key for determining the social cost of carbon. What is so useful about this functional form is that the marginal damage elasticity does not depend on the current level of carbon concentration, output, or any other variable and the marginal damage elasticity is the key element in calculating the optimal carbon tax, as we shall see below.

3.3.3 Remarks

Before proceeding, let us make a few remarks of caution against the backdrop of our initial point: that damage measurements is the area we know the least about.

An overall worry in damage measurement is that the historical range of climate variation—in any given region—for which there is useful economic data (on, say, output or mortality) is very limited compared to the increases in global temperatures that will likely result if a significant fraction of the remaining fossil fuels is used up. Over the last one hundred years, for example, we have seen the global temperature climb by about one degree centigrade but many projections forward involve increases of, say, five degrees. The only approach to damage measurements so far that allows any form of insight into what would occur at five or more degrees of warming is the cross-sectional top-down method referred to initially, since here one compares the outcome variables (like output) between regions in the world today with very different average temperatures. In fact, it is quite well known that there is a strong negative correlation between average temperature and GDP per capita, as depicted in Figure 2.

⁸See Golosov et al. (2014).

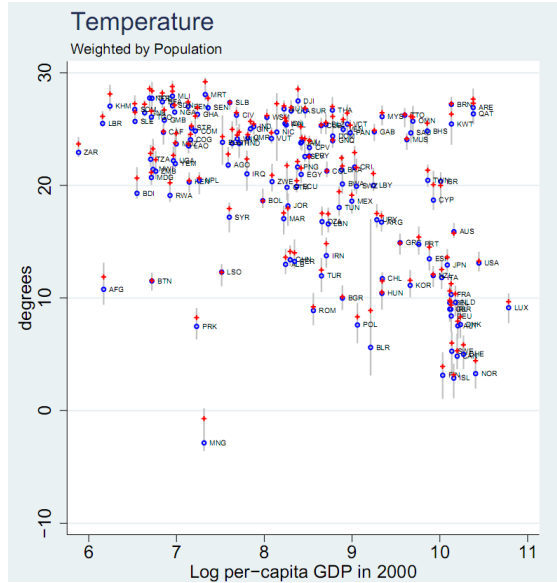


Figure 2. Average temperature 1950-59 (blue) and 1996-2005 (red) vs GDP (PPP). Reprinted from Dell, Olken and Jones (2012).

From a natural-science perspective it may also be that five degrees of warming would involve potentially irreversible nonlinearities that imply that the mapping from carbon concentration to climate variables such as temperature (e.g., through feedback effects) becomes more convex. In this case our approximation to the damage function above will become worse.⁹

There are potential nonlinearities in damages. For example, humans appear to appreciate biodiversity intrinsically and biodiversity may involve tipping points. It may also be that ability to adapt to climate change is powerful within, but not beyond, a certain range of temperatures. The human body, for example, can handle (survive and be productive in) temperatures that are much higher than those in Europe today, but clearly there is a limit on temperature above which humans cannot survive or work productively.¹⁰ If more were known about these processes and it would be possible to map them into observables, one would adjust the aggregate damage function appropriately.

One reason why some argue that damages are bound to be limited is adaptation in the form of migration: humans can, if temperatures rise enough, always migrate to colder areas. However, migration is associated with costs. Obviously, moving New York city away from the coast (which may be needed if the sea level rises enough) would be very costly, but there are also other costs of migration, especially in poorer countries. Harari and La Ferrara document how migration caused by poor agricultural outcomes can cause armed conflict, perhaps by ethnic violence. If people need to move across borders, one can easily imagine political and military

⁹It may also be that the mapping from emissions to carbon concentration becomes more convex. That will not by itself case our damage function to be a poor approximation, but it will complicate the optimal carbon tax calculations, as we will show below.

¹⁰An path breaking study was Haldane (1905). See Sherwood and Huber (2010) for a more recent study related to climate change.

conflicts. It is very difficult to assess these costs. One approach (followed for example by Desmet and Rossi-Hansberg, 2015, who study migration in a theoretical model of climate and the economy) is to imagine that costs are U-shaped, i.e., that there is an ideal temperature for every location. This is an interesting way forward and could be combined with costs of “crossing borders”.¹¹

4 The integrated assessment model

The purpose of this section is not to go through the details of the typical integrated assessment models in the literature but rather to give one example of a model which can be understood based on basic microeconomics and which serves a useful framework for policy analysis. Any claims regarding how the model behaves here are substantiated in other papers, to which the reader is referred.

4.1 The dynamic economic model

Any economic model that is used for quantitative policy analysis of climate change should, in our view, have some basic properties. It should involve dynamics and long-run analysis. It should be similar, or a good approximation, to our standard frameworks from growth analysis—in this case the Solow model or, rather, versions of that model with (at least some degree of) optimizing saving.¹² It should allow for uncertainty. It should be based on microeconomic principles, so that standard welfare analysis can be conducted. These requirements are straightforwardly satisfied in a dynamic neoclassical model of the kind used in modern growth and business-cycle theory. While our model is not designed for business cycles, it could be altered to accommodate many views on business cycles were one to adapt it to short-run analysis—but that is not the purpose here.

There is a representative consumer (now a stand-in for the average world citizen) with a utility function of a single good that is consumed at different points in time. The utility function involves discounting, a key element in evaluating policy, as well as a need for smoothing consumption over time. We abstract from population growth here for simplicity.¹³ The consumption good is produced with an aggregate (world) production function of capital, labor, and energy, and it allows for technical change. We will assume that it has unitary elasticity across inputs. This does not appear restrictive in the case of capital and labor but may be restrictive when it comes to energy; in the short run, it seems much harder to substitute. However, over the longer run, technology choice is endogenous and the assumption of unitary elasticity is less inappropriate.¹⁴ Capital is accumulated in a standard Solowian manner, taking consumption and investment to be perfect substitutes (a questionable assumption for short-run analysis but a reasonable one for long-run applications such as this one).

¹¹This approach is also followed in the most recent work by Krusell and Smith (2015a).

¹²For evidence that optimizing saving at an aggregate level, compared to simple alternatives such as that entertained by Solow, is to be preferred, see Krusell and Smith (2015b).

¹³Population growth can be important in this context for some purposes but not, typically, for some key aspects of policy analysis, such as for the calculation of an optimal carbon tax.

¹⁴For a discussion, e.g., Hassler, Krusell, and Olovsson (2012)

In this exposition we assume, for simplicity, that the energy sector is a pure coal sector and that coal is produced using labor only—the same kind of labor as is used to produce consumption and investment goods. The sole reliance on labor in the coal industry makes for closed-form solutions but is not realistic; however, it is not a serious flaw in the quantitative analysis since the quantitative effects on the main variables of interest are limited given the small share of coal production in GDP. Moreover, we assume that the damages to TFP from climate change appear only in the consumption/investment sector. This simplifies the algebra and is not of quantitative significance, since the energy sector is a rather small part of GDP.

Thus, using standard notation, the utility function of the representative world consumer is

$$\mathbf{E}_0 \sum_{t=0}^{\infty} \left(\frac{1}{1+\rho} \right)^t u(c_t),$$

where t can be thought of as year t (and $t = 0$ is normalized as “now”), u is a strictly concave and increasing function of consumption, c , capturing a need to consumption-smooth as well as risk-insure, and ρ is the subjective discount rate. Using the damage function (7), the resource constraint for the consumption/investment good reads

$$c_t + k_{t+1} = e^{-\gamma_t(S_t - \bar{S})} A_t k_t^\alpha n_{1t}^{1-\alpha-\nu} E_t^\nu + (1 - \delta)k_t,$$

with k denoting capital, A an exogenous TFP component that is possibly growing over time, n_1 labor used in this sector, and E the energy (coal) input (and α and ν are exogenous share parameters); and that for coal reads

$$E_t = \chi_t n_{2t},$$

with n_2 denoting labor used in this sector and χ an exogenous productivity factor that like A may grow over time. Market clearing for labor occurs when $n_{1t} + n_{2t} = 1$ (we normalize labor to 1). Carbon in the atmosphere evolves according to a linear depreciation schedule, so that

$$S_t - \bar{S} = \sum_{s=-T}^t E_s (1 - d_{t-s}),$$

where $1 - d_{t-s}$ is given by (5). In this model, Greek letters are exogenous parameters of which γ_t may be random; in addition, A_t may be random as well.

It is straightforward to define what the socially optimal allocation is: a planner chooses sequences of consumption and energy subject to the above restrictions to maximize the stated objective function. One can similarly define a dynamic (stochastic) competitive equilibrium where all firms (including coal producers) make zero profits and consumers maximize their utility subject to budget constraints allowing saving. Importantly, in the market equilibrium, no agent takes the externality—how emissions E affect S and hence productivity—into account; the social planner, in contrast, does take this into account.

For this model, and much more general versions of it, it is straightforward (see Golosov et al., 2014) to derive the marginal social cost of emitting carbon at the optimal allocation, the OSCC (the Optimal Social Cost of Carbon). This formula says that the OSCC is the *appropriately*

discounted value of *current and future externality damages* caused by a current emission of a unit of carbon. *Appropriately discounted* involves both the discount rate ρ and any other element due to non-logarithmic curvature and consumption growth; in the case of logarithmic utility discounting involves only ρ , regardless of the rate of consumption growth. Computing the *current and future externality damages* involves two factors. First, one has to figure out, for any future date s periods after the emission, how much of the initial emitted unit is still in the atmosphere. The answer is given by the depreciation parameter for carbon s periods out. The second factor, to be multiplied with the first, is simply the marginal externality damage on production of carbon present at that future date s . Finally, sum these damages across all future dates (and states, in case there is uncertainty).

Golosov et al. (2014) show that, under assumptions that are viewed as quantitative reasonable in the macroeconomic literature on growth, the formula simplifies radically. In particular, it turns out that we can express the OSCC for emissions at time t —or, equivalently, the implied optimal period- t tax à la Pigou, τ_t —as

$$\tau_t = \gamma \hat{\delta} y_t, \tag{8}$$

where y_t is output of consumption and investment goods at t , γ is the (expected) damage elasticity parameter introduced in (7), and $\hat{\delta}$ is the appropriate combination of preference discounting and carbon depreciation.¹⁵ The formula reveals that the optimal tax in dollars per ton is proportional to global output. This may seem counterintuitive, but is explained by the fact that the damage (the externality) is proportional to global output. Thus, the externality brings the finite size of the earth into the model; without the externality the optimal prices and rent would be unaffected by a doubling of c , k and n . The constant of proportionality $\gamma \hat{\delta}$ in equation (8) is given by structural parameters, independent of variables such as production inputs, the atmospheric carbon concentration (now and later), technology (now and later), and so on.

Calibrating the model requires assigning values to all its parameters (preferences, technology, etc.). However, only a (rather small) subset of these parameters are needed to find the value for the OSCC. We shall look at a calibration below in order to gauge what an optimal tax ought to be.

What is the optimal level of carbon emissions? It turns out that even in the simple model, the answer depends on all details of the model; even in its most stripped-down form, this is evident. For example, if labor productivity in the coal sector is very high, optimal coal use is higher, simply because its private cost is lower. Hence, in order to determine the optimal quantity, the planner needs to know the cost structure in the coal industry. The static model below will illustrate.

¹⁵For details on the conditions under which the result obtains exactly, see the Golosov et al. (2014). The authors also show remarkable robustness of the formula to departures from the assumptions they state.

4.2 A simple static model

Now consider a conceptually much simpler model without a time dimension: utility is given by

$$u(c),$$

consumption from a static resource constraint by

$$c = e^{-\gamma(S-\bar{S})} A k^\alpha n_1^{1-\alpha-\nu} E^\nu,$$

where k is a fixed, exogenous factor, and energy (in the form of coal, again) from

$$E = \chi n_2,$$

with labor-market resources satisfying $n_1 + n_2 = 1$. Finally, carbon in the atmosphere is given by

$$S - \bar{S} = \phi E,$$

where ϕ represents the fraction of emissions ending up in the atmosphere, thus allowing depreciation within the static model. Clearly, this is the most straightforward static version of the dynamic model above. It should be pointed out that the simple model cannot formally be thought of as a steady state to, or a long-run outcome of, the dynamic model, but nevertheless the dynamic and static models are very similar and give rise to policy implications that have the same form. Of course, the parameters need to be reinterpreted; one can perhaps think of the static model as one of the outcome over a one hundred-year period (with no discounting during this period and infinite discount on any future after that, with ϕE reflecting the average carbon addition to the atmosphere during the century if E is emitted every year for a hundred years, and so on).

Solving for a competitive equilibrium with a unit tax τ on carbon (i.e., for every unit of E purchased, the firm has to pay τ units of consumption) really just involves setting the after-tax marginal product of labor to be the same across the two sectors. This condition gives, after a minor amount of algebra,

$$\frac{(1 - \alpha - \nu) A e^{-\gamma \phi E} k^\alpha \left(1 - \frac{E}{\chi}\right)^{-\alpha-\nu} E^\nu}{(\nu A e^{-\gamma \phi E} k^\alpha \left(1 - \frac{E}{\chi}\right)^{1-\alpha-\nu} E^{\nu-1} - \tau) \chi} = 1.$$

This is one equation in one unknown: coal use E . The equation has an easy solution when there are no taxes ($\tau = 0$); otherwise, it is a nonlinear equation in E that has to be solved numerically. In the case without taxes, the equation becomes

$$(1 - \alpha - \nu) E = \nu \left(1 - \frac{E}{\chi}\right) \chi \quad \Rightarrow \quad E = \chi \frac{\nu}{1 - \alpha}.$$

This equation is fairly simple: neither the TFP parameter A nor capital, k , or the damage and carbon depreciation parameters, γ and ϕ , end up mattering for the determination of coal use/the energy provision. What is important is productivity in the coal sector (χ) and the

relative cost share of energy in production ($\nu/(1 - \alpha)$). In contrast, what is socially optimal is given by

$$\gamma\phi + \frac{1 - \alpha - \nu}{\chi - E} = \frac{\nu}{E},$$

now naturally also involving both γ and ϕ .¹⁶ These two equations are not the same but we notice, going back to the equation determining market energy use for an arbitrary energy tax τ , that if the tax is set to satisfy

$$\tau = \gamma\phi A e^{-\gamma\phi E} k^\alpha \left(1 - \frac{E}{\chi}\right)^{1-\alpha-\nu} E^\nu = \gamma\phi c = \gamma\phi y$$

then the market allocation would be socially optimal. Thus, $\gamma\phi y$ is the Pigou tax in this case. This is intuitive: this amount is precisely the OSCC as given by the damage externality caused by coal use.

We will return to implementation below but it is important to note here that the OSCC is proportional to output through only the product of γ and ϕ , i.e., the damage elasticity and carbon depreciation. We saw above that another parameter matters as well in the dynamic analysis—discounting, as given by ρ there, but because the present analysis is static this parameter is not present—but we must note here that the nature of the solution is extremely similar across the static and the dynamic setting. Thus, both analyses point to the fact that a carbon tax requires relatively little information to implement. In contrast, a quantity regulation, hence going straight at what E ought to be, is more demanding—it requires knowledge also of details about coal production and how it impacts on GDP. Moreover, with slight (and reasonable) extensions of this setting, one realizes that the formula for the optimal tax is barely affected by population growth and other technology parameters (especially concerning green energy) that instead are crucial quantitatively when regulating quantities.

5 Policy analysis

We now provide a sequence of remarks on climate policy. The sequencing is put together in an order that, roughly speaking, has decreasing ties to the formal analysis above. For example, our first point is a quantitative evaluation of what the optimal carbon tax ought to be, whereas the final points have to do with political challenges in implementing different kinds of policies, a subject on which the above analysis is silent since it does not directly touch on politics (some insights from the formal analysis do, however, have implications for a political-economy analysis).

5.1 The optimal tax on carbon

As shown above and by Golosov et al. (2014), under assumptions that are viewed as quantitatively reasonable in the macroeconomic literature on growth, a simple formula for the OSCC can be derived. The formula deserves to be repeated here since we will now give it quantitative

¹⁶The planner's optimal choice comes from a condition that requires the marginal product of labor in each sector to be the same when one takes into account that a unit of labor in the coal sector adds damages to the consumption sector. It is thus very straightforward to derive.

content:

$$\tau_t = \gamma \hat{\delta} y_t$$

where y_t is global output of consumption and investment goods at t , γ is the (expected) damage elasticity parameter in (7), and $\hat{\delta}$ is the carbon duration using the subjective discount rate. Specifically,

$$\hat{\delta} = \sum_{s=0}^{\infty} \left(\frac{1}{1+\rho} \right)^s (1-d(s)).$$

We then obtain

$$\hat{\delta} = \frac{\varphi_L(1+\rho)}{\rho} + \frac{(1-\varphi_L)\varphi_0}{\varphi+\rho},$$

using the formulation of carbon depreciation in (5).¹⁷ Golosov et al. (2014) calibrated γ to 2.4×10^{-5} based on Nordhaus (2007). As discussed above, there is substantial uncertainty about the damage function. IPCC (2007), reports expected damages at 4 degrees heating to be 1-5% of GDP as shown in Figure 3.

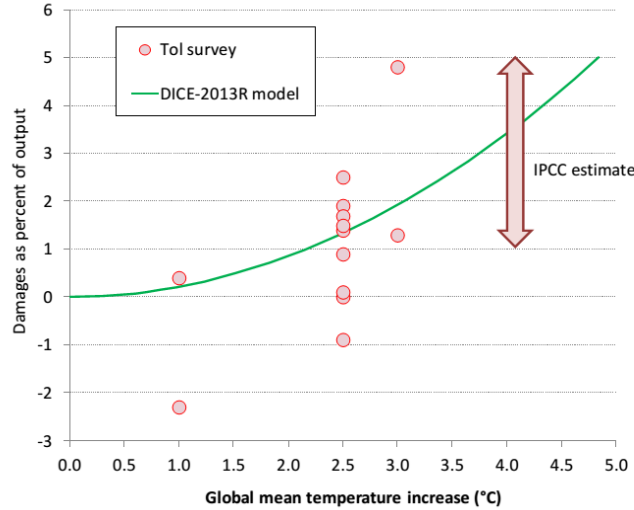


Figure 3. Global damage estimates. Dots are from Tol 2009. The solid line is the estimate from the DICE-2013R model. The arrow is from the IPCC (2007b) page 17.

Reprinted from Nordhaus (2013)

Let us consider a calibration of γ based on the upper end of this range. Using the Arrhenius equation (4) with a climate sensitivity of 3°C , we find that to obtain an increase in the global mean temperature by 4°C , the atmospheric CO_2 concentration is $600e^{\frac{4}{3}\ln 2} = 1512\text{GtC}$. Finally,

¹⁷The first term represents the duration of the part of emissions that is assumed to remain “forever” in the atmosphere. As the subjective discount rate approaches zero, this term approaches infinity. However, the assumption that a share φ_L of emissions remain in the atmosphere for a horizon that from an economic point of view is infinite does become less reasonable with a discount rate close to zero. In this case, the fact that over tens of thousands of years, also this share of emissions is slowly absorbed by the oceans starts to matter. Then a depreciation rate capturing this very slow process should be added to the denominator of the first term. For typical calibrations of subjective discount rates, say larger than 0.1% per year, however, the effect of this slow depreciation is quantitatively negligible.

we use the damage function $D(1512) = 1 - e^{\gamma(1512-600)} = 0.05$ to solve for γ , giving $\gamma = 5.624 \times 10^{-5}$. It may be noted that the calibration of γ in Golosov et al. (2014) corresponds to substantially lower damages, but still within the range reported by IPCC (2014), namely 2.2% at 4°C. Setting global GDP to 650 trillion euro per decade and using the carbon depreciation parameters described above, the formula for the optimal tax per GtC is given by the following expression.¹⁸

$$5.624 \times 10^{-5} \cdot 650 \times 10^{12} \left(\frac{0.2(1 + \rho)}{\rho} + \frac{(1 - 0.2)0.38}{0.023 + \rho} \right).$$

In Figure 4, we plot the optimal tax per ton of carbon against the subjective yearly discount rate. We also plot the tax for the more optimistic calibration of γ from Golosov et al. (2014).




Figure 4. Social Cost of Carbon (optimal tax) in current US\$ per ton of carbon as a function of the subjective discount rate for two calibrations of the damage elasticity.

To obtain some perspective on the tax, it may be helpful to note that a liter of gasoline contains about 0.64 Kg of carbon. A tax of, say 400 Euro/ton carbon, which is about the same level as the current tax in Sweden, therefore corresponds to 0.25 Euro per liter of gasoline.

Risk

An important finding in Golosov et al. (2014) is that, despite the presence of uncertainty and risk aversion, what matters for the optimal tax is only the expected value of the damage elasticity γ and not the degree of uncertainty about it.¹⁹ This is important since there is substantial uncertainty around all the underlying mechanisms determining the relation between CO₂ concentration and damages. It is, however, at least as important to realize the caveats to this result. To derive an approximately constant damage elasticity, we used the logarithmic relationship between CO₂ concentration and temperature in (4) and the moderately convex relationship between temperature and damages in (6). Deviations from these smooth relationships could make risk an important factor. Examples of such deviations would be thresholds in the climate system or carbon circulation. It could be the case that a particular temperature some

¹⁸Often, the carbon tax is expressed per mass unit of CO₂, i.e., including the oxygen.

¹⁹This is an exact result that involves a specific degree of risk aversion (that given by a logarithmic function) together with the exponential damage function.

positive feedbacks in the energy budget suddenly become much more potent. Similarly, the carbon circulation system could abruptly change when some level of concentration is passed.²⁰ Similarly, it may be the case that the quadratic damage function severely underestimates the convexity of damages as a function of the global mean temperature. In any of these cases, uncertainty starts to matter. Exactly contrary to our result that it is the mean that is important, Weitzman (2009) argues that it is the tails that are important. In Weitzman (2012), a damage function that becomes extremely convex at high levels of global warming due to a term with a power of 6.754 in temperature is used. This function is chosen to capture an assumption that at 12°C heating, the loss of GDP is 99%. Weitzman argues that it is hard to rule out a climate sensitivity of 12. In an example, he sets the probability of this event to around 1%. In this case, a mere doubling of the CO₂ concentration may have such dramatic consequences that they must be avoided, even when the costs are high. Also quite unlikely events may warrant forceful policy intervention if these events are sufficiently bad. A high carbon tax, or a very tight emission quota, can be an insurance against a catastrophe. Of course, in this case most likely, paying the insurance premium will *ex post* turn out to be of no value, but nevertheless it is worth it *ex ante*.

While we agree with the logic of Weitzman’s argument we also insist that policy be based on quantitative evaluations: it is not sufficient to refer to abstract arguments. There are many potential catastrophes of other forms that cannot be ruled out—a killer flue, super tsunamis, comets colliding with earth, giant volcano eruptions, etc.—and these could potentially be as far-reaching and damaging to mankind as CO₂ emissions. Although it is an extremely difficult task, our view is that to set policy and devote resources, policymakers must, with the help of scientists, try to quantify and weigh different such risks against each other (and against other areas of spending) in order to then be able to devise a desirable, cost-effective policy mix. So until measurements are available that go beyond the mere idea about how tail uncertainty can be disastrous (if it is large enough and risk aversion is high enough), we prefer to base our analysis on available scientifically based estimates, of course allowing robustness around them. The notion “prudence” is often mentioned in this context and, clearly, prudence is called for, but the question is how much. Another aspect of this issue is that, precisely because many suspect that nonlinearities and irreversible mechanisms exist and are central, much more research on it is needed.

Above, we have argued that the optimal tax is fairly modest and that it is independent of the emission scenario. Our trust in this depends on how difficult it will turn out to be to reduce emissions. If it turns out to be much more difficult than we expect, the optimal tax calculated with our formula may lead to more emissions and thus more global warming. At some level, our trust in the formula will then fade. The same thing applies if it turns out that the climate sensitivity is much higher than expected.²¹ The implication of this is that international negotiations on climate change should first establish a global carbon tax at a reasonable Pigouvian level. Most likely, such a carbon tax would be effective in the sense of

²⁰See, for example, Lenton et al. (2008) or, for a more popular scientific description, Levitan (2013).

²¹Note that a higher climate sensitivity would imply that the time until the steady state is reached for a given emission scenario increases. This means that more time for developing techniques for adaptation and carbon capture is allowed.

curbing climate change. However, if we obtain indications that this is not the case, and emission forecasts risk taking us into carbon concentrations and temperatures about whose consequences we have very little knowledge, more forceful measures should be considered. Without such indications, however, if the immediate focus is on very strong measures we fear that nothing will be achieved in the negotiations.

5.2 It's (almost) all about coal

A conventional oil or gas reserve is an asset with a positive value. As we all know, finding oil reserves has made countries rich. This reflects the fact that the average extraction cost for conventional oil and gas for a long time as been much lower than the price—the price has a rent component. As long as a tax does not eliminate this rent, i.e., does not drive the price net of taxes below the extraction cost, extraction remains profitable and is then likely to continue. Even quite high global carbon taxes are unlikely to eliminate the rent for a large share of existing conventional oil and gas reserves. These will therefore be exploited regardless of whether global carbon taxes are introduced or not.

We should note that if the carbon tax is set to reflect the damages caused by emissions, exploiting these reserves is socially efficient if it is privately profitable. That is, in such a case the social value of using the reserves is higher than keeping them unexploited also when the externalities are included in the calculations. The current estimate of existing conventional oil and gas reserves indicates that indeed these reserves are not large enough to pose a substantial threat to the global climate. Current estimates of oil and gas resources indicate a stock of 300GtC.²²

Currently, the amount of carbon in the atmosphere is about 840GtC. Assuming, fairly conservatively, that half of carbon emissions stay in the atmosphere for an economically long horizon, emissions of 300GtC would lead to an increase the carbon concentration of 18%. Using equation (4) with a climate sensitivity of 3°C, this leads to an increase of the global mean temperature of $3 \frac{\ln 1.18}{\ln 2} \approx 0.7$ °C. This is certainly not trivial, but neither does it appear to be a major threat.

For coal, the situation is very different. First, no countries become rich by finding coal reserves. This is due to the fact that extraction costs are close to market prices—rents are negligible. This in turn means that a relatively small reduction in the price before taxes makes a lot of coal extraction unprofitable. In contrast to oil, a carbon tax therefore has the potential to have a large effect on extraction.

Second, coal reserves are substantially larger than oil and gas reserves. Official global coal reserves are 640GtC.²³ However, it is likely that this is a substantial underestimate. Since coal is priced close to extraction cost, the value of searching for new coal mines is limited. In fact, Rogner (1997), estimates coal reserves to be 3,500GtC with a marginal extraction cost curve that is quite flat.²⁴

²²BP (2015) reports global oil reserves to 239,8Gt. Using a carbon content of 0.775 this is 196GtC. The same source reports natural gas reserves to 187.1 trillion m³. Using a carbon content of 0.574 kg/m³, this is 107 GtC. At current extraction rates, both these stocks would last approximately 50 years.

²³BP (2015) reports 891Gt of coal. Using a carbon content of 0.716, this corresponds to 638GtC.

²⁴Rogner estimates coal reserves to 3,400Gtoe, which is approximately 3500GtC.

Thus, the conclusion is that a carbon tax is unlikely to have a large effect on the use of conventional oil but that this is not a major problem. On the other hand, a carbon tax is likely to have a large effect on coal use and limiting coal use is therefore of utmost importance. As an example, the Swedish CO₂ tax is SEK 4,110 (430 euros or, in USD, \$485) per ton of carbon. The average price of oil in 2014 was approximately \$100 per barrel, corresponding to \$733 per ton of oil, and the average price of coal in northwestern Europe the same year was \$75 per ton. In percent of the fuel price, the Swedish CO₂ tax was thus 55% for oil and 460% for coal.²⁵ Clearly such a tax makes coal use uneconomical while gasoline use has not collapsed in Sweden. Also a more modest tax, of say \$100 per ton of carbon, would likely have very large effects on coal use but only modest ones on oil use.

The fact that the coal price is fairly close to the extraction cost in combination with the fact that coal supplies are fairly evenly spread over the major regions of the world explains why global trade in coal is limited. Figure 5 shows coal production for the major regions of the world in the left panel. Consumption is shown in the right panel. As we can see, the two panels are fairly similar implying that trade between the regions is limited. This contrasts sharply with oil, where global trade is very important. As is seen in Figure 6, oil production and consumption do not correspond to each other. These differences between coal and oil also have important policy implications. Given segmented markets, a reduction in coal use in one region of the world is not likely to affect the price and the use in other regions. The market for oil is not segmented in this way, implying that a reduction in demand in one region is likely to affect the world market price negatively and thus increase consumption in other regions.




Figure 5. Coal production and consumption. Reprinted from BP (2015).

²⁵We use a carbon content of 0.846 for oil and 0.716 for coal.




Figure 6. Oil production and consumption. Reprinted from BP (2015).

5.3 Taxes vs. cap-and-trade

Climate change driven by emission of greenhouse gases is an almost perfect example of an externality. The benefits from using fossil fuel are private to the emitter, but since CO_2 quickly mixes in the atmosphere, it has global effects that are independent of who emits and where the emission takes place. Such a text book case of an externality implies that unregulated markets will not lead to a socially efficient level of emissions. In theory, the market failure due to the externalities induced by CO_2 emissions can be solved by quantity restrictions as well as with Pigouvian taxes.²⁶ To illustrate this point, consider a simple static case when emissions have private benefits and social costs. The private benefits represent the net value to the user of burning fossil fuel in excess of the costs associated with producing the fuel, for example due to extraction and refining. These benefits accrue to the user (consumer surplus) and to the producer (profits) in shares that are determined by market conditions. Here, we are only concerned with the sum of these benefits, not how they are split. In an unregulated and efficient market, profit opportunities are exploited and this will imply that marginal private benefits will be driven down to zero. Emissions also have social costs. These are assumed to be external, i.e., they are neither borne by the emitter nor the fuel producer. Instead, they are borne by a large number of agents spread around the world. The market transaction will therefore be undertaken as if there were no social costs.²⁷

Let us now depict this graphically. In Figure 7, the blue solid curve represents the marginal private value of emissions and the red solid curve the social cost. In an unregulated market, emissions will be Q^{LF} since all private gains will be exploited. However, the social optimal

²⁶Pigou (1920) was first to show how taxes can solve market failures caused by externalities.

²⁷It should be said that an increasing number of consumers, especially in rich countries such as Sweden, appear to be willing to “internalize” the externalities by imposing restrictions on their own behavior such as reducing their fossil-fuel demands. However, this group is, and arguably will be for the foreseeable future, negligible viewed from the global perspective.

emission quantity is Q^* , where the marginal private benefits are equal to the marginal social costs. The optimal allocation can be implemented by either a tax per unit of emissions equal to τ^* or a quantity restriction (a quota) at Q^* . A way of implementing the quantity restriction is to require all emitters to purchase an emission permit that is provided in the quantity Q^* . If these permits are traded, their unit price will be τ^* . Note that in this case, information about the position of both curves is required to find Q^* as well as τ^* .




Figure 7. Marginal private value of emissions (blue) and marginal social cost (red).

Now, recall the discussion in Section 4, where we concluded that a robust and reasonably accurate approximation of the optimal tax is independent of the emission quantity (because a concavity in the emission-to-carbon concentration mapping cancels with a convexity in the carbon concentration-to-damage mapping, thus making the emissions-damage relationship linear). With this result, a graphical representation of marginal costs and benefits of emissions instead can be depicted as in Figure 8. There, the curve representing the marginal social cost of emissions is horizontal. Now, it is no longer the case that information about the position of both curves is necessary in order to find τ^* . In fact, no information about the private value of emissions is needed. Equivalently, changes in the curve representing the marginal private value of emissions, e.g., from the solid to the dashed line in Figure 8, will change Q^* but not τ^* .




Figure 8. Marginal private value and marginal social cost (constant).

The finding that the marginal cost curve is flat and its implication that τ^* is independent of the private benefits of emissions point to important benefits of using taxes rather than quotas.

First, less information is required. This is a value in itself and can also make it easier to come to agreements about the right policy. Individual countries have incentives to misreport their own marginal value of emissions since in a multi country version of the analysis, the optimal allocation of emission quotas between countries depends on the individual marginal values. Such an incentive does not exist if the policy instrument is a tax.

Second, the marginal value of emissions varies over time, in part due to changes in economic activity, which shifts the blue curve in Figure 8. As the global financial crisis hit the economies of the world, for example, it shifted to the left, and since the stipulated quota Q^* did not shift, this led to a collapse of the price of emission rights in Europe. This is illustrated by the new price $\tau^{*'}$ in Figure 8, far below the marginal social cost. The gap is caused by the failure to predict the appropriate Q^* . With the price collapse regulation itself collapsed.²⁸ The high volatility of emission rights have been observed in the past also for other pollutants such as e.g., sulfur dioxide.²⁹

It should be pointed out that for the climate, what matters is not so much whether the price of carbon fluctuates but what its average price is—since the climate is so slow-moving. However, unpredictability of prices is really undesirable for business, i.e., from a pure economics perspective. A case in point is the massive investments in coal made by the Swedish (state-owned) company Vattenfall that turned out to be extremely unsuccessful after the fall in energy prices, and in the value of the free emissions allowances that were obtained with these

²⁸In principle, one could think of a European Emission Trading Central Bank with the responsibility to stabilize the price of emission rights at τ^* . A simpler implementation of the optimal policy would, however, be to have a tax, which wouldn't have to vary as energy demand and thus emission values vary with the level of economic activity.

²⁹See Green et al. (2007) and Nordhaus (2007).

investments because of grandfathering, so much so as to be a contender for the worst business deal in Swedish history. To the extent that different emission permit markets are not perfectly integrated, the trading system can also lead to large discrepancies in the emission price faced by different agents. This will lead to unnecessarily high costs of achieving a given amount of emission reduction.³⁰

Third, an efficient transition to a less fossil fuel dependent economy requires long-term policy commitments. However, predicting the future value of emissions is difficult since, for example, the speed at which alternative energy sources are developed as well as GDP growth rates are hard to predict. Unconditional commitments to a path of emissions are therefore hard to make credibly, especially for economies that change rapidly, and if they were made, they could turn out to be quite sub-optimal. Conditional commitments, where emission paths would be made depend explicitly on the factors that drive emission values, seem too complicated to be a way forward.

To introduce and commit to keeping a reasonable level for a carbon tax has much fewer of these problems. We also think that a conditioning of the future tax on new information about the global flow cost γ , carbon duration D , and global GDP, is likely to be credibly implementable.

5.4 The optimal tax will not be very harmful

A common question asked by practitioners and the informed public is whether the proposed carbon tax will harm growth, or welfare, greatly, while of course having benefits in the form of a more agreeable climate. One could attempt to answer this question with an empirical study that convincingly would identify the effects of carbon taxation on growth. One could alternatively use (quantitatively restricted) theory to make predictions. We will briefly discuss both. Our overall conclusion is a “no”, given our (admittedly tentative) evidence.

In terms of empirical evidence, let us first specify the questions we want to answer. One question concerns a single, and perhaps small, country adopting a tax when other countries do not. Another question involves the effects of a world-wide common tax. The first of these questions—which includes the phenomenon of “leakage” of economic activity and energy use from high-tax to low-tax countries—could potentially be answered empirically, but we are not aware of any studies that convincingly establish causality. The second question seems even harder to answer. However, let us make some observations.

First, from our own perspective—that of Sweden—let us point out that a carbon tax on a high level has been in place now for over twenty years (detailed reference about details and changes in tax rate over time?) and that there is no indication that Sweden has done particularly poorly over the same period. This, of course, can be due to other counteracting factors, such as a wave of deregulation and lower tax rates on income. Nevertheless, it is hard to imagine that our carbon taxes really have been very detrimental to Swedish growth.³¹

Second, there is ample evidence that there is major scope for straightforward energy-saving

³⁰ An example of such an inefficiency is that the Swedish government in 2014 sold unused emission rights to Merrill Lynch at a price of 0.03 SEK/kgCO₂ at the same time as it taxed Swedish emission at the rate 1.08 SEK/kgCO₂. Unfortunately, we believe that this example is not an outlier.

³¹ The total energy bill in the economy is on the order of magnitude of 5% of GDP.

measures (such as cheap insulation, “closing windows and shutting off machines”, etc.) and technological advances directed at saving energy. For examples, the following chart illustrates how the energy efficiency, measured by GDP in U.S. dollars per unit of energy, is very different even across developed economies: for example, it differs by a factor of five between Iceland (low efficiency) and Switzerland (high efficiency) as seen in Figure 9. These differences, we suspect, likely reflect differences in energy prices across countries (and these energy prices in turn surely reflect supply factors as well as policy). Thus, there appears to be great scope for energy saving. A second empirical argument for energy saving is contained in Hassler, Krusell, and Olovsson (2014) who use postwar U.S. postwar data to back out a time series for energy-saving technical change; we reproduce the graph below in Figure 10. As can be seen below, this series was essentially flat until the oil price shocks hit and then started growing substantially and persistently. The effects, moreover, are quantitatively large.




Figure 9. Energy efficiency. GDP US\$ (2005 PPP) per unit of energy (kg oil equivalent). Source: Worldbank, World Development Indicators Online.

In terms of quantitative theory there is also relatively little work. On the effects of a global tax, there is some work, but the model-based analysis presented above is one of a relatively small set of examples where comparisons between laissez-faire and an optimum (where the climate externality is managed perfectly) is carried out using model parameters that allow a reasonable match between the model and historical data (thus representing theory that is “quantitative”). Though we will not belabor the exact arguments here, the findings are that the effects on growth are relatively minor. One should keep in mind, of course, that in part this is because the climate damage is estimated to not be disastrously high in our benchmark economy, and thus the welfare gains are limited from taxing carbon optimally, hence also implying limited costs on the economy.

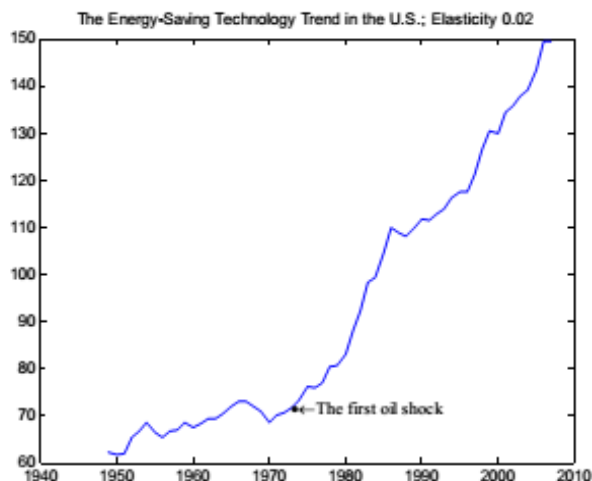


Figure 10. Energy efficiency in the U.S. Source: Hassler, Krusell and Olovsson, (2012)

A channel that is not present in the model of the previous section is that running via endogenous technology. Suppose, for instance, that a tax induces energy saving that attracts research capital (human and financial). Then less of this research capital will be used elsewhere, and hence although the economy will adapt well to energy scarcity (or high taxes on energy), it will do less well in the development of other growth-enhancing innovation. Some preliminary calculations based on Hassler, Krusell, and Olovsson (2012) suggest that the growth effects of fossil-fuel scarcity could be quite noticeable: in the long run, this scarcity would lead to a drop in the growth rate of consumption by somewhere around 1 percentage point. However, fossil-fuel scarcity cannot be miraculously eliminated, so that growth loss is not a matter of choice. An optimal tax would increase the scarcity somewhat, marginally lowering the growth rate further. Moreover, if R&D is carried out efficiently—if appropriate subsidies are applied to the R&D where technology spillovers are present—then by the argument above, the losses from the tax cannot be larger than the damages from suboptimally fast climate change, and are hence limited. However, a second-best analysis—due to pre-existing inefficiencies in the research sector—would change these results; how they would change would depend on the details of these inefficiencies.

On the question of leakage, there is virtually no work. Work in progress by Krusell and Smith (2015a) offer the methods but, as of the time of writing the present piece, no firm conclusions.

In sum, with the caveats stated, we are not aware of any evidence pointing to large output or welfare losses from an optimal tax on carbon, and the evidence that is available rather indicate rather limited losses.

5.5 Green technology is great but it needs to outcompete coal

The externalities associated with CO₂ emissions create a market failure that motivates market intervention. If the externality were the only market failure, a tax would be sufficient to solve the problem. In reality, there are, of course, many other market failures or imperfections. Of

particular interest are market imperfections associated with technological development. The market for ideas can hardly be efficient since, on the one hand, ideas are reproducible at zero marginal cost so using them should entail a price of zero. On the other, if the price of ideas is zero, there are no incentives to produce new ideas. Therefore, an efficient production of new ideas, e.g., for new energy sources or higher energy efficiency, requires subsidies and/or monopoly rents to the firms that carry out R&D. An important issue is to what extent climate change motivates changes in the policies vis-a-vis R&D. Specifically, should green technologies and energy efficiency be subsidized more than they would have been in the absence of the emission externality from fossil fuel? A number of arguments are put forward in this context.

The Green Paradox

In a series of contributions, Hans-Werner Sinn (see, e.g., Sinn, 2008, 2012) argues that subsidies to the development of green alternatives to fossil fuel has the malign consequence of speeding up the exploitation of fossil fuel reserves. The logic is straightforward. Owners of fossil fuel reserves must make sure that they are sold before new technology or other regulation make them worthless. If subsidies to green alternatives are expected to be successful in the sense of making economical alternatives to fossil fuel arrive earlier in the market, this will speed up the rate of extraction of fossil fuel reserves in fixed supply. Not only will the subsidies then be powerless in the sense of not affecting fossil fuel extraction accumulated over time, but they will also speed up extraction, which typically reduces welfare also for given accumulated emissions.

The logic behind the argument is impeccable. However, it is important to realize that it applies only to fossil fuels that (i) exist in finite supply and (ii) are expected to be all used up. A perfect example of such a resource, for which thus Sinn's argument is valid, is conventional oil with extraction costs close to zero, e.g., Saudi oil where the market price mostly consist of rents. For such a resource, the accumulated supply over time is inelastic.

In order to reduce the accumulated supply, policy must be sufficiently powerful to completely eliminate rents, not only in the future but in all periods. A policy strong enough to make the extraction of Saudi oil profitless already today is inconceivable and almost certainly not socially beneficial. In this case, policy in the form of taxes, as well as subsidies to green technology, can only affect the time path of extraction. If green technologies will ever make it unprofitable to extract and sell the Saudi oil, the Saudis are likely to make sure they sell all of it before that happens.

Consider instead a resource that has a flat extraction cost curve and exists in such large quantities that no rents exists. For such a resource, the accumulated supply is elastic and will respond to changes in policy and technology. The absence of rents implies that the extraction decision is static and not forward-looking. The market equilibrium will imply that extraction is positive whenever the market price is sufficiently high to cover extraction costs, regardless of expectation about the future. In this case, there is no green paradox. A policy that affects the future time at which green technologies replace fossil ones has no impact on current extraction rates.

Whether the green paradox is a problem or not is therefore an empirical question, determined by the supply side of the market. Clearly, fossil fuel exists in many forms. At one side of the spectrum, we have conventional oil with low extraction costs and high rents. At the other, we

have coal where extraction costs are close to the price so there are very limited rents. Above we argued that the supply of conventional oil with low extraction costs is quite small. In fact, this is the reason for the high rents. Given that the supply is small enough not to be a key problem from the point of view of climate change, neither is the time path of its extraction.

Coal, on the other hand, exists in much larger supplies and for coal, the green paradox does not apply, at least as long as no economically important rent arise. To the extent that such rents exist a successful policy to reduce accumulated coal extraction must limit or eliminate these rents. Policy should attempt to drive the value of coal deposits to zero. The conclusion is that from a climate perspective it is only meaningful to subsidize green technology that can replace coal, for example solar cells and wind power (which replace coal in the production of electricity), but not technology that can only replace oil, for example ethanol for cars.

To the extent that reserves of non-conventional oil and gas are large enough as to be a problem, a similar argument applies to these reserves as to coal. As a consequence, the value of investing in developing technologies to extract these reserves should therefore be reduced so as to not make them profitable.

Path dependence and technology ladders

Acemoglu et al. (2012) argue that not only are subsidies to the development of green technology valuable, but they are in fact key for addressing the market failure of CO₂ emissions and quantitatively much more important than Pigouvian taxes. Acemoglu et al. base their argument on the idea that technological developments can be seen as a ladder—new ideas build on and improve upon previous ideas. Furthermore, they assume that the technology ladder for production and consumption using green energy sources are distinct from their fossil counterpart. If, in particular, most of the R&D resources have been spent on developing the fossil technology ladder, the green technology will be very far behind. In this case, it is unlikely that R&D focusing on green technologies will be profitable without subsidies. Under these conditions, Acemoglu et al. show that even a temporary (but perhaps large) subsidy toward green technology would suffice in producing a permanent technology shift away from fossil fuel. The argument rests precisely on the assumption that technology choice (green vs. dirty) is inherently path-dependent, as just explained.

The logic of the argument is of course coherent but there is little evidence in practice on just how strong technological path dependence is. It is not enough that there is some of it, and there is ample evidence that technologies based on green energy can apply many innovations that came about in research directed at fossil-based energy. More generally, our argument is that if an optimal carbon tax is adopted, green subsidies may only be necessary to the extent that they involve research externalities. The Acemoglu et al. counterargument would be that there are strong “non-convexities”, whereby the energy sector is stuck on an inappropriate path; thus, green subsidies would be necessary for changing paths.³² As we said, however, it is far from clear that such non-convexities are really present. Having said that, we must point out that in the absence of an optimal carbon tax, green subsidies are of course highly relevant and desirable, and given the lack of international agreements in the climate area, that unfortunately

³²Formally, Pigou-based analysis rests on marginal conditions. With non-convexities, one can imagine multiple local optima and, hence, a role for additional policies in order to select globally among multiple equilibria or steady states.

appears to be the situation we are currently in.

A main point about green technology is that, from the climate perspective, it is not sufficient that an increasing fraction of our energy use come from green sources. Subsidizing green energy technology implicitly subsidizes energy consumption, and although this is expected to lead to an increase in the production of green energy, the implied decrease in fossil energy production may be quite small, if there is any. To prevent harmful climate change, we primarily need to make sure that only a small fraction of the coal reserves are actually used. From a green-technology perspective, what this simply means is that green technology must not only be good: it must outcompete coal. This point is, we think, not sufficiently emphasized in the debate; one instead obtains an impression that green technology itself is a solution, regardless of how it affects coal. Our point here is that green technology is *only* beneficial for the climate if it “kills coal”. Of course, this is implicit in Acemoglu et al.’s analysis as well. Their point is in fact that a temporary subsidy to green energy may suffice to outcompete coal—a position on which we think the jury is still out—but the broader point is really the central one.

Non-constant discounting and green investments

A standard assumption in economics is that the subjective discount rate is constant. This assumption simplifies the analysis but is questionable. First, there is evidence that individuals use non-constant discount rates when comparing their own welfare at different horizons.³³ Second, there is little empirical reason to assume that individuals discount the welfare of future generations in exactly the same as their own. A subjective discount rate in the range 1-2 percent per year is often used and this rate is not obviously inconsistent with actual returns on financial and real assets and other macroeconomic variables.³⁴ However, the implications of discount rates at this level applied to welfare far out into the future appear hard to justify on moral grounds. With a subjective discount rate of 2% per year, welfare 35 years from now are valued half as much as today’s. An implication of a constant interest rate is that if welfare 35 years from now is valued half as much as today’s, the same two-to-one ratio also applies for any two points in time that are 35 years apart. For example, welfare 500 years from now is twice as valuable as welfare 535 years from now. Whether such an assumption is appropriate is far from clear: in fact, we know very little about how people discount welfare of future generations far out in time. The implication that relative valuation of welfare at two different points in time only depends on the amount of time between them may seem unreasonable for many people. While a discount rate of 2% for the near future may be reasonable, the value of welfare 500 or 535 years out may be regarded as approximately the same. This would imply that the subjective discount rate is not constant but rather falling as the horizon is extended.³⁵

A falling discount rate may seem like a minor change in our analytical environment. However, there is a substantial complication, because it makes any policy plans for the future *time-inconsistent*. In particular, what, say, an optimal carbon tax at a particular future time t is depends on how distant this period is, so that if a decision can be made on this period- t tax at some earlier point in time, the decision will be made differently. In concrete terms, if this

³³See, e.g., Laibson (1997).

³⁴However, the low return on safe government bonds is hard to reconcile with subjective discount rates substantially above zero.

³⁵See Arrow et al. (2014) for a discussion on whether declining discount rates should be used in policy analysis.

tax rate is not set until period t it will likely end up being set at a much lower rate, as the decision makers at that point in time will be more “impatient” about the future than we are today about their future. That is, they will attach a smaller relative weight on the welfare 35 years after time t than we do, and hence if we planned on a high tax in period t , they will have an incentive to reverse this decision as time t arrives. Iverson (2014) shows that the formula in Golosov et al. (2014) can actually be extended to allow for non-constant subjective discount rates. He focuses on the case when there is no possibility to commit, so that taxes are set every period for the current period only. In his model, every decision maker would thus like to bind future decision makers to set higher taxes due to the time-inconsistency but cannot do this directly; hence, only indirect, and imperfect, channels remain as means to exert this influence. The relatively higher concern for future generations, however, leads to higher taxes than in the case of constant discount rates. Gerlagh and Liski (2012) also show that declining discount rates have a quantitatively important effect on the tax rate that obtains in equilibrium. The latter authors also emphasize that the tax is higher than what is implied by the Pigouvian principle. The reason is that the equilibrium is inefficient also for other reasons than the climate externality. In such a case, the carbon tax can be useful for other purposes than dealing with climate externality. Specifically, due to falling discount rates, the equilibrium will feature that current generations would like to influence future generations to leave more resources for even more distant generations. Since the effects of carbon emissions, and thus of carbon taxes, are very long-lived, it turns out that this gives a quantitatively important argument for higher carbon taxes.³⁶

Another case of time inconsistency arises in a political setting where groups with different views on the benefits of a carbon taxes take turn in office. Clearly, the group with a greater concern for the damages from climate change acknowledge that if it is currently in power, it may be replaced in the future by another group which is less concerned. A model with these feature is analyzed in Schmitt (2014).

Regardless of whether it is due to non-constant discount rates or political preference heterogeneity, the resulting time-inconsistency provides an argument for investments in green technology, both in green capital and in the development of new green technologies. If a current policymaker believes that future policymakers will set carbon taxes that are too low, the negative consequences of this may be reduced by investing in green technology. Since it is a state variable, the amount of green capital and green technology will hence affect future decisions and can be used strategically. Schmitt (2014) analyzes this mechanism in the case of political heterogeneity but, as far as we know, an analysis of how investments in green technology can mitigate the negative consequences of time-inconsistency due to non-constant discounting is still rather unexplored.

³⁶In general, the Pigouvian principle only holds exactly when the climate externality is the sole source of inefficiency. The existence of other distortionary taxes, for example on labor, implies that the optimal tax may deviate from the Pigouvian principle. Intuitively, if the carbon tax makes distortions on the labor market worse, the optimal carbon tax should be lower than otherwise. Our assessment, however, is that these other second-best considerations are of minor importance. See Bovenberg and de Mooij (1994) for a static analysis and Barrage (2014) and Schmitt (2014) for a dynamic one.

5.6 Politics

The above analysis ignores political aspects entirely: they assume that a “benevolent planner” can simply implement any policy. In practice, there are arguably “political frictions”. One example may be undue influence by certain interest groups that happen to be well organized. Another is the lack of commitment power in politics. There are many other possibilities too. We cannot possibly cover all of these aspects here and we will simply make a small set of comments, none of which is very closely related to the formal analysis above. We will discuss the complications in arriving at an international agreement, and also comment on two often-heard assertions: the ideas that a tax is politically infeasible and that green-technology subsidies are a better way forward.

5.6.1 It is easier to agree on a tax than on quotas

Ever since the climate negotiations began in the early 1990s, the approach has been to set national quotas for CO₂ emissions and to complement this with cap-and-trade to make it economically efficient. This approach has failed to produce anything close to a global agreement. There is a limited agreement within Europe, thus covering a minor fraction of the global emissions; that is all. Moreover, both India and China (the world’s largest emitter) have clearly declared that they will not commit to any binding quotas. There are promises about future quantity cuts, e.g., in negotiations between the U.S. and China, but these clearly fail the purpose since (i) they are far out in the future and (ii) they involve significant growth in fossil-fuel use until then, at which point the level will be frozen, but no cuts are intended. Given this failure of the current quantity-based approach, we think it is high time to revisit the tax suggestion.

Why has it been impossible to agree on global quantity restrictions? In our view, there are two basic reasons. The first one is the huge uncertainty about the economic consequences of such an agreement. These consequences, in particular, depend very critically on the general economic development. Thus, to predict the consequences for one country, we essentially need to predict the position of the blue curve in Figure 8 for that country several decades into the future.

In the financial markets such uncertainty has a price, measured as the higher return demanded from a risky investment than from a safe one. The corresponding price is less clear in politics, but the principle is the same, and the effect is that the negotiating parties want a safety margin. In the Kyoto protocol this safety margin was in most cases so large that the quotas turned out to exceed business-as-usual. They were therefore meaningless.

We saw in section 5.3 that the uncertainty in setting the appropriate tax is much smaller than in setting the appropriate quantity. For climate negotiations, the important point is that the economic consequences of a carbon tax are much more predictable than of quotas, which makes an agreement easier.

The second reason why it is easier to agree on a tax than on quotas is connected with the logic of zero-sum games: it is very hard to split a cake. Each country has a clear interest in getting as large a piece—as large a quota—as possible. The externality, i.e., the damages caused by climate change, does not affect this consideration at all if there is already a fixed global limit, and only very little otherwise, since the climate change caused by the country’s own emissions

is small.

The situation is very different when negotiating about the global level of a carbon tax. The interest of each country is then determined by both the burden caused by the tax (which in this case at least stays in the country) and the benefit from mitigating climate change, assuming that all other countries have the same tax. In the words of Weitzman (2014), who analyzed this situation theoretically, “the externality is internalized” in the tax negotiations.

A parallel to congestion tolls in major cities is perhaps useful here. The congestion externality is similar in nature to that arising from carbon use. The different countries needing to agree on a common tax, or a quantity constraint for each country, are represented by the different boroughs around the city in the congestion case. So imagine the cap-and-trade system for congestion externalities: one has to agree on how many rides to the city will be allowed by each borough. How easy would it be to agree on how many rides each borough is allowed? On the other hand, road tolls in the form of unit prices have indeed been put in place and shown to work rather well in many cities, as it is far less challenging to agree on a common price.³⁷

The congestion example, and the fact that some boroughs expect growth and hence a greater future need for metropolitan transportation, also illustrate the effect of uncertainty, and why countries like China or India refuse to enter into a quantity agreement: how many emission rights will they need, under different growth scenarios (which for natural reasons differ greatly), and how much will they cost? There is no straightforward answer to these questions. The optimal carbon tax is very different in this regard: it is easily evaluated—it should be proportional to global GDP—and predicted. This is also why boroughs that might experience high growth in the future would be reluctant to adopt a quantity-based metropolitan transit regulation, even if there is trade in transit permits: it might just be very costly to them. It is therefore not a surprise that they have opted for the unit toll.

Just like with the negotiations surrounding cap-and-trade, one would like a carbon tax to be adopted globally. What are the problems if not all countries/regions agree on a tax? The current knowledge suggests that even a partial implementation of taxes may be very helpful for the climate. The reason is that much of the fossil fuel—coal—is available only locally, due to high transportation costs. Therefore, worries about leakage, which are justified when it comes to oil and natural gas, are actually less relevant for the bulk of the fossil fuel reserves. This point, of course, also applies to a cap-and-trade system.

5.6.2 “A tax is politically infeasible”

As the story goes, and this may well be true, the cap-and-trade system was once adopted because the obvious choice—a carbon tax—was not politically feasible. The cap-and-trade system was perhaps easier for politicians to adopt because it allowed grand-fathering: pre-existing emitters would obtain emission permits for free to a very large extent, thus significantly limiting their opposition to cap-and-trade relative to a tax system. Moreover, many may have been skeptical

³⁷The road toll example can be made to liken the case of climate change even more if one imagined that the boroughs were represented by the following leaders: Barack Obama (or Donald Trump?), Vladimir Putin, Xi Jinping, . . . It is not hard to understand why an agreement on ride quotas would not work in this more elaborate example. Despite the heterogeneity of these borough leaders, however, we do think that there would be a chance of an agreement of a unit toll price even between such leaders.

toward the effectiveness a tax: does it really work (will it really change behavior), and isn't it safer to directly regulate quantities? Finally, the cap-and-trade system has also tended to come with a notion that severe quantity restrictions will be implemented, only not right now but later.

An often heard argument is that in the U.S., a tax is impossible to implement. We think this is a substantial exaggeration of the problem. High taxes are definitely possible in the U.S.; in particular, property taxes are higher there than for the average OECD country and much higher than in, say, Sweden. An explanation for this is that property taxes are not administered by the federal government but by local governments. There is nothing that prevents a solution where the carbon tax is also collected by local governments and that these local governments use the tax proceeds as they see convenient. International agreements should focus on the tax rate to be applied on fossil fuel, as opposed to on what the revenues are used for.

It has been mentioned that grand-fathering may have been a major reason behind the political feasibility of the cap-and-trade system. Is grand-fathering desirable? From a basic economics perspective, grand-fathering of a given set of quotas just amounts to one particular way of distributing the wealth corresponding to the market value of the rights, and in this sense it may not be worse than any other distribution of this wealth (one would be that the government auction it off, thus making the wealth end up in the hands of the tax payers). There are two potential problems, however. One occurs when there is an expectation that more rights will be issued, and possibly grand-fathered, in the future, because such beliefs will increase the incentives to use carbon now. Governments fundamentally have problems with commitment power: they would like to promise that they will not do this in the future, but may not be able to stick to such a promise.

Another potential political challenge with grand-fathering, to the extent that cap-and-trade were to be introduced globally, might also arise. If the emission quotas for each country were to be decided from historic emissions and a politically negotiated guess about the future development, one can easily imagine that such limits would be much more restrictive for countries with unexpectedly high growth than for slow-growing countries, and for countries with traditionally high energy taxes (which have already implemented the easiest energy saving measures) than for those with low energy taxes. In the resulting trade, emissions allowances would be sold from countries with easy limits to countries with restrictive limits. A likely outcome might be that allowances will be sold by American and Russian companies to European companies. This could result in substantial capital flows. Moreover, the money would flow to countries that are more wasteful of energy and, in some cases, also richer than the European countries. Because the general public might well view these cash flows as punishing the righteous (those who have already been serious about energy saving) and rewarding the sinners, it is hard not to imagine great opposition to the system if it indeed becomes reality and if the "and-trade" part becomes operative. A domestically collected carbon tax, on the other hand, does not generate any capital flows between countries. Thus, the view that cap-and-trade is politically more feasible than a carbon tax is superficial, and rests on the simple fact that the public has not yet seen and understood the effects of cap-and-trade.

It appears that the political feasibility of cap-and-trade and carbon tax have an opposite

time dependence. A tax is difficult to introduce, but as time goes by it will be seen that it has no drastic effects on either the personal or the national economy. People, and the economy in general, will adjust by relying less on fossil fuel, and this itself will decrease the public resistance to the tax. This can be seen from the fact that a carbon tax (and its relative, the gasoline tax) is less controversial in countries where it is already high, such as Sweden, than where it is low, such as the U.S.

Cap-and-trade, on the other hand, may seem easy to introduce, when few understand how it works, but will generate much more resistance when its effects become visible. In the long term, which is what matters for climate change, it is therefore less politically feasible than a carbon tax.

5.6.3 “It is better to subsidize green technology”

One can hope that green technology, particularly through improvements in the production of energy not based on fossil fuel, will be the way forward. We certainly hope, and believe, that green technology will eventually come to rescue. There are several challenges, however, in implementing green subsidies.

One problem is that the implementation of green subsidies can be mismanaged. Suppose, for example, that the setup is one where politicians try to identify the best green technologies for the future and thus direct subsidies only to selected companies/technologies. Not only may politicians not be well equipped to identify which technologies are most promising, but there are strong incentives for private actors to misrepresent the cost and technology structures. The research area is indeed typically fraught with informational asymmetries and contracting problems. The Swedish system with so-called green certificates, giving subsidies to any non-fossil based energy, is one promising way forward where no judgment calls are necessary.

From a policymaker’s perspective, an attractive feature of green subsidies is that, just like for the cap-and-trade system based on an implementation using grand-fathering, there does not appear to be much political, or lobby-based, opposition. Those who receive money (green-tech companies or green-energy users) are happy about the subsidies and presumably the general public would only have to pay slightly higher taxes on other items and, beside, gain on net due to a better climate. However, the apparent absence of opposition is actually very worrying. In particular, if the coal industry felt truly threatened by the alternative technologies to be developed, they would be vocal—as they have been in other contexts. Our worry is that they do not feel threatened because the green-tech subsidies are unlikely to be effective enough to compete with coal. If this is true, nothing is accomplished, as we have argued above: coal is the major threat to our climate. Thus, the reaction of the coal industry to any proposal, be it cap-and-trade or carbon taxes or green-tech subsidies, is a litmus test for the desirability of the policy: the more they protest, the better is the policy proposal. A carbon tax at a relatively modest level would clearly be a serious threat to many coal-based energy producers as the price of coal-produced energy appears to be close to marginal cost.

Finally, an increasingly common claim about green-tech subsidies is that, because they are a way of the future, they present the possibility of a double dividend: they also create jobs and know-how that will pay off beyond any climate dividend that may result (if coal is outcompeted).

This is an interesting thought but, unfortunately, we know of no convincing study supporting these claims and therefore, awaiting evidence, view it as another example of wishful thinking. After all, as we have argued, no matter how many jobs are created from subsidizing green technology, if it does not outcompete coal it has no benefit from the climate perspective. (Note: perhaps a note is warranted here on the currently low energy prices is called for. These are indeed problematic for some energy producers, including coal-based ones. It is unclear how much of the low prices are due to lower world demand, increased supplies of new fossil fuels, and alternative technologies, respectively.)

6 Conclusions

We have put forth a number of policy recommendations in this essay and we have, most importantly, tried to build them, as far as we have been able to, on the results from the climate-economy research field. We have argued, among other things, in favor of a carbon tax and that this tax likely (i) will not need to be too high and (ii) will not have severe consequences for growth. The arguments we have given in favor of a carbon tax instead of cap-and-trade are economic and political. There are also other important arguments of more practical and legal nature that have been discussed by Nordhaus (2007). For example, cap-and-trade is more vulnerable to corruption, and typically only covers large industrial sources of emissions.

We have also argued that subsidies to green technology are important, but that to the extent that they do not outcompete most of the existing coal reserves they actually do little for the climate. Our future climate rests (almost) only on the outcome for coal: if the very large available deposits of coal are used up, our climate will change drastically (even according to conservative estimates about the natural-science mechanisms), whereas the climate is likely not much affected if we use up the conventional reserves of oil and gas. The policy issue, to us, is therefore mainly about preventing the coal use. So we simply need to “grab the bull by the horns”, the bull being represented by present and future coal producers, whether private or government-controlled. Put this way, policy measures to support green technology are good from a climate perspective only to the extent they make the bull rage. This is, unfortunately, the unpleasant and unavoidable arithmetic of the climate problem.

Having said all this, we need to add a final caveat, and one that is very important: many of our quantitative statements are quite uncertain, because the knowledge in this field is still very limited. Thus, whereas we believe that it is crucial that policy be guided by best-practice science, one has to recognize the presence of significant uncertainty in the scientific findings. We have pointed out that the uncertainty is particularly large when it comes to measuring the damages from climate change. One of the important aspects of this measurement is the great heterogeneity of outcomes for different regions of the world as the earth heats up, but even the average effect is not well understood and measured. We would argue that damage measurement, both in the form of pure empirical work and theory work toward understanding the nature of damages and possible adaptation mechanisms to them, is absolutely crucial going forward. Revisions of damage measurements, moreover, have direct consequences for the optimal carbon tax, and are thus amply motivated from a policy perspective. We of course also have to add

that there are many question marks in the natural science components, even though we believe that they are far better understood than the damages from warming.

7 References

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

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

Archer, David, Michael Eby, Victor Brovkin, Andy Ridgwell, Long Cao, Uwe Mikolajewicz, Ken Caldeira, Katsumi Matsumoto, Guy Munhoven, Alvaro Montenegro, and Kathy Tokos (2009), "Atmospheric Lifetime of Fossil Fuel Carbon Dioxide", *Annual Review of Earth and Planetary Sciences*, 37, 117–134.

Arrow, Kenneth J., Maureen L. Cropper, Christian Gollier, Ben Groom, Geoffrey M. Heal, Richard G. Newell, William D. Nordhaus, Robert S. Pindyck, William A. Pizeryy, Paul R. Portney, Thomas Sterner, Richard S. J. Tol, and Martin L. Weitzman (2014), "Should Governments Use a Declining Discount Rate in Project Analysis?", *Review of Environmental Economics and Policy*, 8:2, 145–163.

Barrage, Lint (2014), "Optimal Dynamic Carbon Taxes in a Climate-Economy Model with Distortionary Fiscal Policy", Working Paper, University of Maryland.

Bluedorn, John C., Akos Valentinyi, and Michael Vlassopoulos (2010), "The Long-Lived Effects of Historic Climate on the Wealth of Nations", mimeo, Southampton University.

Bovenberg, A. Lans and Ruud A. de Mooij (1994), "Environmental Levies and Distortionary Taxation", *American Economic Review*, 84:4, 1085–1089.

BP (2015), BP Statistical Review of World Energy June 2015, <http://www.bp.com/>.

Ciscar, J.C., Iglesias, A., Feyen, L., Szabo, L., Van Regemorter, D., Amelung, B., Nicholls, R., Watkiss, P., Christensen, O., Dankers, R., Garrote, L., Goodess, C., Hunt, A., Moreno, A., Richards, J., Soria, A. (2011), "Physical and Economic Consequences of Climate Change in Europe", *Proceedings of the National Academy of Sciences*, 108:7, 2678–2683.

Ciscar J.C., Feyen L., Soria A., Lavallo C., Raes F., Perry M., Nemry F., Demirel H., Rozsai M., Dosio A., Donatelli M., Srivastava A., Fumagalli D., Niemeyer S., Shrestha S., Ciaian P., Himics M., Van Doorslaer B., Barrios S., Ibáñez N., Forzieri G., Rojas R., Bianchi A., Dowling P., Camia A., Libertà G., San Miguel J., de Rigo D., Caudullo G., Barredo J.-I., Paci D., Pycroft J., Saveyn B., Van Regemorter D., Revesz T., Vandyck T., Vrontisi Z., Baranzelli C., Vandecasteele I., Batista e Silva F., Ibarreta D. (2014), "Climate Impacts in Europe. The JRC PESETA II project", European Commission.

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

Desmet, Klaus and Esteban Rossi-Hansberg (2015), "On the spatial economic impact of global warming", *Journal of Urban Economics*, 88, 16–37.

Gars, Johan (2012), "Essays on the Macroeconomics of Climate Change", Ph.D thesis,

Institute of International Economic Studies, Stockholm University.

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

Golosov, Michael, John Hassler, Per Krusell, and Aleh Tsyvinski (2014), "Optimal Taxes on Fossil Fuel in General Equilibrium", *Econometrica* 82:1, 41–88.

Haldane, John S. (1905), "The Influence of High Air Temperatures: No. 1.", *The Journal of Hygiene*, 5:4, 494–513.

Harari, Mariaflavia and Eliana La Ferrara (2012), "Conflict, Climate and Cells: A Disaggregated Analysis," IGIER Working Paper 461.

Hassler, John and Per Krusell (2012), "Economics and Climate Change: Integrated Assessment in a Multi-Region World", *Journal of the European Economic Association*, 10:5, 974–1000.

Hassler, John, Per Krusell, and Conny Olovsson (2012), "Directed, Energy-Saving Technical Change", Working Paper.

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

IPCC (2007b), "Climate Change 2007: Impacts, Adaptation and Vulnerability," Eds. Martin Parry, Osvaldo Canziani, Jean Palutikof, Paul van der Linden, Clair Hanson, Contribution of Working Group II to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change.

IPCC (2013), "Climate Change 2013: The Physical Science Basis", Eds. Thomas F. Stocker, Dahe Qin, Gian-Kasper Plattner, Melinda M.B. Tignor, Simon K. Allen, Judith Boschung, Alexander Nauels, Yu Xia, Vincent Bex and Pauline M. Midgley, Working Group I Contribution to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change.

Krusell, Per and Anthony A. Smith, Jr. (2014), "Regional Damages From Temperature Fluctuations", Working Paper.

Krusell, Per and Anthony A. Smith, Jr. (2015a), "Climate Change Around the World", Working Paper.

Krusell, Per and Anthony A. Smith, Jr. (2015b), "Is Piketty's "Second Law of Capitalism" Fundamental?", *Journal of Political Economy*, 123:4, 725–748.

Laibson, David (1997), "Golden eggs and hyperbolic discounting", *The Quarterly Journal of Economics*, 112:2, 443–478.

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

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

Mendelsohn, Robert, William H. Nordhaus, and Daigee Shaw (1994), "The Impact of Global Warming on Agriculture: A Ricardian Analysis", *American Economic Review*, 84:4, 753–771.

Myhre, Gunnar, Eleanor J. Highwood, Keith P. Shine, and Frode Stordal (1998), "New estimates of radiative forcing due to well mixed greenhouse gases", *Geophysical Research Letters*,

25:14, 2715–2718.

Nordhaus, William D. (2007), “To Tax or Not to Tax: The Case for a Carbon Tax”, *Review of Environmental Economics and Policy*, 1:1, 26–44.

Nordhaus, William D. and Paul Sztorc (2013), “DICE 2013R: Introduction and User’s Manual”, retrived from <http://aida.wss.yale.edu/~nordhaus/homepage/>.

Schmitt, Alex (2014), “Beyond Pigou: Climate Change Mitigation, Policy Making and Distortions”, Ph.D thesis, Institute of International Economic Studies, Stockholm University.

Sherwood, Steven C., and Matthew Huber (2010), “An adaptability limit to climate change due to heat stress, *Proceedings of the National Academy of Sciences*, 107, 9552–9555.

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

Sinn, Hans-Werner (2012), *The Green Paradox. A Supply Side Approach to Global Warming*, MIT Press, Cambridge, Mass.

Tol, Richard S.J. (2009), “The Economic Effects of Climate Change,” *Journal of Economic Perspectives*.

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

Weitzman, Martin (2012), “GHG Targets as Insurance Against Catastrophic Climate Damages”, *Journal of Economic Public Theory*, 14:2, 221–244.

Weitzman, Martin (2014), “Can Negotiating a Uniform Carbon Price Help to Internalize the Global Warming Externality?”, *Journal of the Association of Environmental and Resource Economists*, 1:1/2, 29–49.