Mistra-SWECIA

JOHN HASSLER AND PER KRUSELL

PROFESSORS OF ECONOMICS
IIES AT STOCKHOLM UNIVERSITY

REPORT

No 5

THE CLIMATE AND THE ECONOMY

EIGHT CLIMATE-ECONOMY PROPOSITIONS

1. Combustion of fossil fuels leads to a warmer climate, but how much warmer is uncertain.
2. According to available, albeit uncertain, estimates, the likely costs associated with climate change will be significant, though not catastrophic from a global perspective.
3. Available knowledge is insufficient to assess the risks of unlikely, though potentially major, global damages.
4. The geographic distribution of the costs caused by climate change is extremely uneven. Some, typically poor, countries and regions risk suffering catastrophic damages.
5. Combustion of fossil fuels creates an ”externality” that the market alone cannot manage. A global carbon tax is the best policy instrument for dealing with this market failure.
6. Given the available estimates, the optimal carbon tax is modest and is no higher than the tax levied in, e.g., Sweden.
7. The introduction of a global carbon tax does not threaten economic growth and welfare.
8. Taxes on conventional oil have, at best, a modest positive impact on the climate. Taxes on the use of coal and non-conventional fossil fuels with high extraction costs are effective, even if such taxes are not implemented everywhere.

THE CLIMATE AND THE ECONOMY

JOHN HASSLER AND PER KRUSELL | PROFESSORS OF ECONOMICS

The use of fossil fuels affects the climate and on average leads to non-negligible damages to the economy. A climate-economic model is required in order to analyse how our economies can be expected to affect the climate in the future and to assess various forms of climate policy. Such a model contains, as integrated parts, a climate model, a carbon-cycle model and an economic model. The purpose of this report is to describe key features of climate-economy models and to calculate how high an optimal carbon tax should be. This forms the basis for a discussion of the impact on economic growth and welfare.

Effects of climate change and measures to counteract these cannot be analysed properly without a comprehensive, long-term, and global perspective. For example, an EU carbon tax, has impacts in other parts of the world and the effects are extremely long-term.

We have been working on climate and economics for around five years. Our research specialisation is in the area of macroeconomics and long-term growth from a global perspective. As macroeconomists, we have a number of the skills and tools necessary for analysing climate-economic issues.

However, our knowledge is far from sufficient. Obviously, no climate-economic analysis can be carried out without an understanding of the scientific relationships that drive climate change. In the area of natural science we have no own expertise. For that reason, we have collaborated with scientists from institutions such as the Swedish Meteorological and Hydrological Institute and the universities of Lund and Stockholm, all within the interdisciplinary Mistra SWECIA project.

Unlike in the natural sciences, our climate-economic models must also describe human behaviour. Because humans think ahead and expectations about the future can affect what occurs in the present, economic models must incorporate forward looking, which places special demands on the analysis. In particular, available climate models must be radically simplified before they are integrated with an economic model. We received help in this simplification process from our scientific colleagues in the Mistra SWECIA project: they helped us appropriately condense knowledge from main-stream science.

In this report, we will describe a number of conclusions drawn from our work on climate-economics in the Mistra SWECIA project. We have also summarised our findings in eight points on page 1. For us, none of the points were obvious before we began our research project. Depending on individual background and knowledge, some of our readers are sure to consider some of the points trivial. However, the risk that all points will be considered trivial is small. Our conclusions are of course provisional and new research may very well lead us to revise some of them. Our climate-economic project represents work in progress and our own views on the topic may change as we learn more from our research as well as that of others. In the list above, we mark this uncertainty by referring to the points as ”propositions” rather than ”conclusions”.

THE SCIENTIFIC BACKGROUND

Carbon dioxide, like several other gases such as water vapour and methane, more easily passes through short-wavelength electromagnetic radiation, such as light, than more long-wave radiation, such as heat. With more CO_2 in the atmosphere, the balance between incoming and outgoing radiation therefore changes. This mechanism can easily be verified experimentally and the Swedish Nobel Prize Winner in Chemistry in 1903, Svante Arrhenius, described this so-called ”greenhouse effect” as early as 1896.[1] According to Arrhenius' experiments, higher CO_2 levels lead to a rise in temperature proportional to the percentage rise in the CO_2 level (a so-called log-linear relation). Even though the direct greenhouse effect is experimentally quantifiable, the final effect of higher CO_2 levels on the Earth's temperature is much more difficult to determine. This is due to the fact that the direct effect of more CO_2 leads to a series of feedback mechanisms that strengthen or weaken the direct effect. Via warming, more CO_2 leads to more water vapour in the atmosphere and to the release of methane, which strengthens the greenhouse effect. Reduced formation of ice in the Arctic reduces the Earth's ability to reflect incoming sunlight, which also strengthens the direct effect. A rise in the global average temperature might increase cloud formation, which can increase the reflection of incoming sunlight and thus weaken the direct effect. These are some examples of a large number of more or less well-known feedback mechanisms, some of which can be expected to strengthen the direct effect and some to weaken it. The difficulty of assessing the strength of the feedback mechanisms creates significant uncertainty surrounding the relation between climate change increased CO_2 concentrations. The most common view among scientists is that the feedback mechanisms in total reinforce the direct effect of higher CO_2 concentrations.

A reasonable approximation that we and others often use is that the best estimate of the effect of a doubling in the CO_2 concentration is a global average temperature rise of around 3 degrees Celsius. However, the analysis must take into account the fact that there is considerable uncertainty surrounding this figure.

The feedback mechanisms need not be of constant strength or independent of the CO_2 concentration. It is conceivable that some reinforcing mechanisms only begin to take effect after some warming has already occurred. If that is the case, the effects of higher CO_2 levels would not be so visible at first, but once a certain critical level has been reached, the rise in temperature would accelerate. It is possible that this would lead to so-called ”*tipping points*”. In other words, if the global temperature exceeds a certain critical level, a new higher equilibrium level arises so that, even without any further emissions, the Earth's climate ”tips over” to a new equilibrium with higher global temperatures. If such critical levels can be identified, it may be crucial not to exceed them.

To our knowledge, however, there is as yet no consensus among climate scientists as to how high these critical levels are and whether they even exist at all at a global level.[2] There are local tipping points that can lead to radical changes in local climatic conditions, but when aggregated globally, a smooth (log)linear relation still seems to be a reasonable approximation.

Another significant source of uncertainty derives from what happens to emitted CO_2 over time. There are large amounts of carbon in plants, in the soil and in the oceans. Flows of carbon that are many times greater than current emission rates of carbon from combustion of fossil fuels occur continuously between these reservoirs. Feedback mechanisms whose effects are uncertain are also created here. For example, both higher CO_2 levels and climate change affect the ability of plants to grow and thereby store atmospheric carbon. Up until now, more than half of the fossil carbon we have emitted has left the atmosphere and has been stored in the oceans and as larger biomass. There are many indications that the ability of these reservoirs to store more carbon decreases as emissions accumulate.[3]

1. Arrhenius (1896).

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

3. See, for example, Matthews, Gillet, Stott and Zickfeld (2009).

CLIMATE DAMAGES

If climate change did not have an important effect on human welfare, climate economics would not be of interest as a subject. How can this impact be measured, and is it obvious that the climate damage is significant or that it even is a damage – could it not also mean an increase in human welfare?

These issues are central importance for us because our purpose is to find out how economic-political intervention in the market economy could improve the global outcome for humans.

Compared to the resources spent on natural-science research on climate change, research on the economic aspects of climate change is quite limited. It is therefore not surprising that knowledge here is relatively cursory and fragmented.

Two different approaches

The existing research can be divided into two distinct approaches – one that may be referred to as ”*bottom-up*” and the other as ”*reduced-form*”.

In the ”bottom-up” research, the idea is to first make a list of areas where climate change may be considered to have an impact on the economy. Obvious areas to be included are agriculture, floods, coastal erosion and health effects, for example.

Each of those areas is then studied separately in order to assess the anticipated costs of different amounts of climate change.

When this work is completed, the sum of all the climate damages studied can be calculated and expressed in the form of a damage function that indicates the intensity of damages globally (and for individual regions) for different levels of climate change. The latter is usually summarized in terms of the change in the global average temperature.

At this point it is important to note that the damages are not limited to effects that can be directly measured in monetary terms by using market prices. Climate change affects health, the value of leisure and many values that are not traded on any market.

However, in order for the full economic cost of climate change to be aggregated, all effects must be measured in the same unit, in monetary terms. Such an assessment takes place routinely in economic cost-benefit calculations: for example a life saved in road accidents is ascribed a certain monetary value although human lives are of course not traded in markets.

Ideally, one way to find out how climate change affects the global economy would be to conduct experiments in which individual countries are randomly ”allocated” different climate changes. A credible function between climate and economic effects could then be estimated.

Since many of the mechanisms that link the climate and economics take place at a global level, one may perhaps even need to experiment with a large number of different planets! Of course, this is not possible, but so-called ”natural experiments”, i.e., random variations in the climate, may instead be used under certain circumstances.

This forms the basis of the ”reduced-form” approach to measuring the effects of climate change on the economy. In this approach, historical climate fluctuations are studied and statistical methods are used to relate them to changes in economic outcomes. Is it the case, for example, that countries that have, by chance, suffered unusually hot weather for a certain period have also experienced poor GDP growth?

The advantage of the bottom-up approach is that it provides information on the direct mechanisms underlying the relationship between the economy and the climate. This likely produced more reliable results and it can possibly also provide an opportunity for extrapolation, i.e., analysis of larger climate changes than those so far observed. Of course, one major disadvantage of the bottom-up approach is the difficulty of ascertaining whether the list of conceivable mechanisms is complete. The focus is often on local effects. The costs of conflicts and global changes in animal and plant life, however,

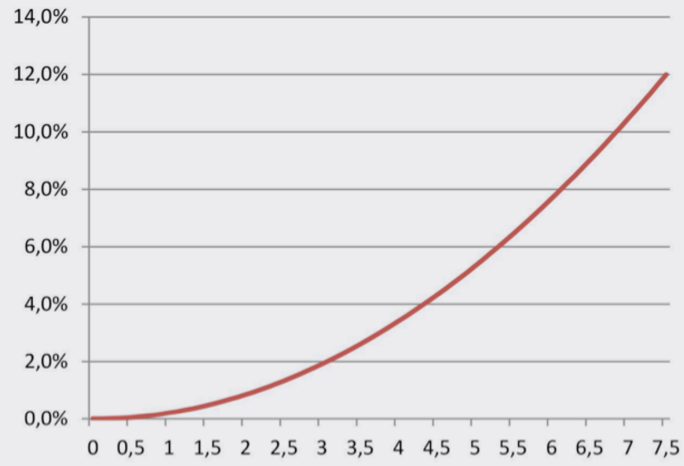

Figure 1. Climate damage as a percentage of GDP as a function of the increase in the global average temperature. Source: See notes 5 and 6.

are typically omitted for lack of underlying studies.

The use of the reduced-form approach should provide a greater chance of capturing all essential mechanisms. On the other hand, it does not provide a description of the actual mechanism behind the statistical relationships and it is hence more difficult to generalise and extrapolate. Of course, historical fluctuations in the climate do not include all conceivable future climate scenarios. Since the two methods have different strengths and weaknesses, they must be used as complements.

Climate studies

The pioneer in climate-economy modelling, William Nordhaus, also conducted some of the first studies on climate damage based on a bottom-up approach.

Seven different types of climate damage are specified[4]: i) agriculture, ii) other production sectors such as energy, forestry and tourism, iii) rise in sea level, iv) health, v) leisure and recreation, vi) migration and ecosystems, and vii) disasters. After a large number of studies on damage in these areas have been compiled, they are summarised in a function that describes the total global damage expressed as a percentage of GDP.

The latest version of the Nordhaus damage function is shown in Figure 1.[5] As we may observe, the curve becomes steeper as the temperature rises. This means that the damage increases more than proportionally to the temperature. This is a common result in the literature, but of course we have very little knowledge of the consequences of increases in temperature larger than those observed in historical data.

The so-called ”PAGE model” which, among other things, is used in the Stern report[6], also makes use of a damage function where damages increase more than proportionally. Unlike many other models, in this model an attempt is made to take into account uncertainty about what values the various parameters should have in order for the model to be as realistic as possible.

4. Nordhaus and Boyer (2000). Nordhaus developed the much-used DICE (Dynamic Integrated model of Climate and the Economy) and RICE (Regional Integrated model of Climate and the Economy) models. They are described and are generally available on the Nordhaus website http://aida.wss.yale.edu/~nordhaus/.

*5. This is included in Nordhaus' DICE 2013 model. Its form is extremely simple. The damage is equal to 1-1/(1+0,002131 T2) where T is the rise in global * average temperature.*

6. Stern (2007).

This is done by defining an interval from which the parameters in the model are then randomly drawn. For example, it is assumed that the damage will depend on an actual rise in temperature to a power whose value is randomly drawn from a given interval.[7] A large number of simulations are then executed with different parameters being drawn each time. This is a way of estimating the uncertainty in the model's predictions.

The European Commission has conducted a bottom-up study for the EU. In the so-called ”PESETA project”[8], the effects of a variety of different kinds of climate change have been estimated for various different parts of Europe. The estimated damages include coastal damages, damage by floods and damage to agriculture, tourism and health. In looking at damages, the EU was divided into five regions and Sweden is included in the region of Northern Europe along with Finland and the Baltic States. For this region, it is actually estimated that the effects of climate change by 2080 will be positive, corresponding to approximately 0.5 per cent of total consumption for a rise in the average temperature of around 5 degrees in Europe and 3 degrees globally. For Southern Europe, the effect is estimated to be negative, corresponding to 1.5 per cent of annual consumption.

As far as the ”reduced form” approach is concerned, there are as yet very few studies taking a global perspective. One noted study[9] relates economic growth in 136 countries over the period 1950–2003 to changes in average annual temperature. The study differentiates between changes in level and growth rate of GDP and finds that there is a strong negative correlation between temperature and economic growth. A one-degree rise in temperature leads to lower growth of one percentage point per year, *but only in poor countries*.

The effect does not appear to diminish over time, though it is not possible to draw clear conclusions about very long-term effects from just over 50 years' of data on growth. However, these results are not undisputed. Preliminary results from research conducted by one of us (Per Krusell) with Anthony Smith at Yale University indicate significant effects on levels, but no effects on growth.

Constant damage per ton of carbon

As stated above, it appears reasonable that increase in climate damage should be more than proportional to temperature.

On the other hand, we know that the direct greenhouse effect caused by CO_2 implies that the temperature depends on the percentage increase in CO_2. This implies that a given increase in CO_2 concentration has a less than proportional effect on temperature. Under standard assumptions, such as those behind the Nordhaus damage function, the more than proportional relation between damages and temperature and the less than proportional relation between temperature CO_2 concentration lead to a proportional relation between damages and the CO_2 concentration. Thus, the overall effect on the economy expressed as a percentage loss of GDP of an extra ton of carbon in the atmosphere is approximately constant, independently of the temperature level.

Based on Nordhaus' calibration, we have used this insight to calculate that one extra gigaton (billion tons) of carbon reduces global GDP by 2.4 thousandths of a per cent each year. The current CO_2 concentration around 800 gigatons of carbon, which is approximately 200 gigatons higher than before the fossil fuel era began. Using the proportional relation between CO_2 concentration and damages, the 200 gigatons leads to damages of 0.5 per cent of global GDP or approximately 400 billion dollars per year.[10]

7. This is typically taken as a triangular distribution between 1.5 and 3, with 2 as a modal value.

8. Se http://peseta.jrc.ec.europa.eu.

9. Dell, Jones and Olken (2012).

10. Formally, this is an expected value because Nordhaus assumes a small probability of a "disaster" with much larger damages.

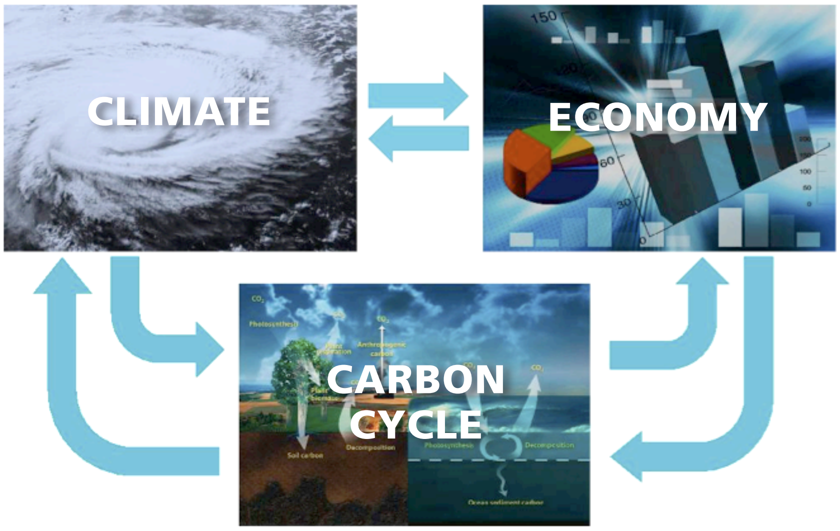

Figure 2. Interdependence in a climate-economic model.

CLIMATE-ECONOMIC MODELS

Two clear conclusions may be drawn from the above discussion: humans affect the climate through the use of fossil fuels and this impact on the climate leads to non-negligible damages to the global economy.

Climate-economic models are mainly used for two purposes: as forecasting tools, i.e. to predict how the economy and the climate are expected develop in the future, and as laboratories to evaluate how different policies affect this development.

A typical climate-economic model consists of three components: a climate model, a carbon cycle model and an economic model. The three components interact in a more or less complicated manner; see figure 2.

The carbon cycle determines the dynamic path of atmospheric carbon concentration and is fed by the emissions generated in the economic model. The carbon concentration, in turn, affects the energy balance in the climate model via the greenhouse effect and this determines how the climate evolves. The climate affects the economy by leading to damages of various sorts. Of course, these interactions can be made more or less complex, but we consider the circular interdependence just described as a reasonable minimum level of complexity for a climate-economy model. Such models are often called Integrated Assessment Models, IAMs.

Nevertheless, many models used to answer questions about climate and the economy are not integrated in this way and therefore lack the internal consistency that may be considered desirable. For example, it is common for scientific models to take one emissions pathway as given and to calculate the climate effects of that path, without taking into account how the implied climate would feed back on the economy and affect emissions, thus generating an emission path that differs from the one assumed.

Integrated climate model

We have developed a very simple but usable integrated climate model.[11] It describes the average global temperature in an equation based on Arrhenius' work referred to above.

The equation states that the temperature is proportional to the logarithm of the atmospheric CO_2 concentration with a constant of proportionality such that a doubling of the CO_2 concentration leads to a 3-degree increase in the global average temperature. As noted above, this proportionality is usually referred to as the climate sensitivity and in our framework, it can be allowed to be uncertain, with a possibility both of a higher and a lower level.

11. See Golosov, Hassler, Krusell and Tsyvinski (2014) and, for a regional analysis, Hassler and Krusell (2012).

The carbon cycle model

The carbon cycle model must, as a minimum, describe how the atmospheric CO_2 levels evolve over time given a particular carbon emission path.

In our work, we use a simple approximation such that a share (half) of emitted fossil carbon is assimilated by plants and the ocean surface within a couple of decades, another share (one fifth) remains in the atmosphere for thousands of years while the remaining proportion slowly trickles down to the deep oceans.[12]

We use more complex models, allowing feedbacks from the climate to the carbon cycle as described above, to check the validity of our simple approximation.

The economic model

The economic model must describe how the economy develops over the long term because climate change is a slow-moving process. It must also be capable of describing the use of fossil fuels and how that use is affected over time by, for example, taxes or other policies.

Naturally, more or less complicated models can also be used here. In many cases, relatively detailed descriptions of the specific energy sector are used. The disadvantage of these models, which are built just to describe energy supply, is that they are not appropriate for modelling economic growth in a wider sense.

We have chosen to build on a straightforward extension of the simple and well-known Solow model, which has formed the basis for growth modelling for a long time. The extension mainly involves making the production function include fuel as a factor of production in addition to capital and labour. We also assume that the savings rate is determined by forward-looking households who maximize their individual welfare. One important aspect of the economic model is of course the damage function, as discussed above, which captures the loss of productivity caused by higher global temperatures.

OPTIMAL TAX

At least since Arthur Pigou's book on welfare economics[13], which was published in 1920, it is well known that an efficient way to deal with an externality (such as the climate damage caused by emissions) is to impose a tax corresponding to the damage that is caused. With the tax, private benefits will be correctly weighed against the external cost implying an efficient trade-off on a market that otherwise works properly.

In principle, the question of how fossil fuels should be taxed is therefore not different from the question of how to tax other goods with externalities. However, in practice things are slightly different than for many other pollutants because the externality is global and will remain present for a very long time.

Above, we have shown that a not unreasonable approximation of the effects of CO_2 is that an extra unit of CO_2 in the atmosphere creates damages that, expressed as a percentage of GDP, are independent of both the level of GDP and the global average temperature. Of course, this is no more than an approximation and new knowledge can lead to a revision of this result. Furthermore, we must allow considerable uncertainty concerning the value of this constant proportional damage. However, the approximation is valuable as a starting point for a transparent analysis of a reasonable level for a global carbon tax.

Above, we showed that one extra gigaton of carbon in the atmosphere can be expected to reduce global GDP by 2.4 thousandths of one per cent per year. We use the Greek letter γ for this constant. With a current global GDP of approximately 75,000 billion dollars, this damage is equivalent to 1.8 billion dollars per gigaton, or 1.8 dollars per ton of carbon.

12. We base this on IPCC (2007, page 25), which summarises as follows: "About half of a CO2 pulse to the atmosphere is removed over a time scale of 30 years; a further 30% is removed within a few centuries; and the remaining 20% will typically stay in the atmosphere for many thousands of years."

13. Pigou (1920).

If a ton of carbon emitted into the atmosphere stayed there for exactly one year and then disappeared, a tax of 1.8 dollars per ton would mean that the externality was internalised in the market. But since one ton of carbon emitted stays in the atmosphere for much longer than one year, future damage must also be included in the tax.

In order to tax emissions of fossil carbon correctly, we must therefore calculate the value of all damages, now and in the future, generated by a marginal emission today.

Two aspects must then be taken into consideration. Firstly, the length of time a ton of emitted carbon remains in the atmosphere is important. The longer it remains, the more damage it will cause. More specifically, we need to know how much of an emitted ton remains in the atmosphere at all future points in time.

Secondly, even if the damage per ton is constant, measured as a percentage of GDP, we must determine what value we should assign today to the loss of a given percentage point of GDP at a specific time in the future. This is necessary since the tax that leads to the internalization of future damages is to be applied today, when the emission occurs.

Depreciation of carbon

Regarding the first issue, we use the carbon cycle model that describes how an extra emitted ton of carbon affects the future path of the atmospheric CO_2 concentration. However, most carbon cycle models are so complicated that they do not lend themselves to being represented by a few parameters.

Therefore, to clarify how carbon depreciation affects the level of the optimal tax, we have instead chosen the approximation form described above.

We therefore assume firstly that a share φ_L (we use 20 per cent as our baseline) of an emitted ton stays forever.[14]

Secondly, we assume that a proportion $1 - \varphi_o$ (we use approximately 50 per cent) disappears immediately and thirdly, we assume that the remaining share $\varphi_o(1 - \varphi_L)$ trickles down to the deep ocean at a constant depreciation rate of φ per decade (we use 2.3 per cent, giving a half-life of 300 years).

As we will see, these parameters will be key determinants of the optimal tax on fossil carbon emissions.

Discounting

How to evaluate future welfare is a classic, multifaceted, and controversial issue. William Nordhaus[15] argues that market data should be used as a basis for such an evaluation. By observing how the market currently values a security that gives a specific payment in the future, for example a bond, we can calculate how the market participants value future consumption possibilities.

On the other hand, Nicholas Stern[16] argues that the market does not provide good guidance on how we should value future income and losses thereof because the question is rather one of moral values where we should not necessarily have the same view as the market participants.

Also, we need to evaluate consumption losses occurring hundreds of years into the future and financial claims with such a long duration are not easy to find on the market, if they even exist. So even if we are able to deduce from market data that people use a particular discount rate for the near future, it is by no means certain that the same rate should be used when we discount damages occurring several hundred years into the future.

Even if the choice of discounting involves moral issues that we will soon return to, there are certain aspects that can be analysed with some degree of objectivity. One of these consists of how we should look at the fact that future generations are likely to have different income and consumption levels than we have. From a welfare perspective, it is reasonable to consider it worse if a poor person loses resources than if that loss is suffered by a rich person. This is due to the standard assumption in welfare economics that the higher the consumption of an individual is, the lower is the value of a marginal consumption unit. Here, our assumption that damages are proportional to GDP will turn out to be quite helpful for the analysis.

So how is our evaluation of damages at a particular date in the future affected by income levels at that time? On the one hand, higher income levels (higher GDP) mean that the damage will be greater because, as Nordhaus and in line with other studies, we assume that the damage is proportional to GDP. This would lead to a higher valuation of the damage today.

On the other hand, higher income levels mean that

14. In reality, it is a question of thousands of years or more, which in economic calculations can be approximately expressed as "for ever".

15. Nordhaus (2008).

16. Stern (2007).

the welfare loss of a unit of lost consumption is lower since higher income means higher consumption and a lower value of a marginal consumption unit.

The two effects of higher future income are therefore pulling in opposite directions when it comes to the valuation of future consumptions losses and, using assumptions that are standard in macroeconomics, we can show that they in fact cancel each other out exactly. This means that the valuation of the damage caused by a ton of carbon in the atmosphere at a future date does not depend on income levels at that time.[17]

A formula for the optimal tax

We now have all the ingredients needed for calculating the optimal carbon tax. For each future point in time, we calculate the damage of an extra ton of carbon in the atmosphere. This is given by the constant γ times GDP at that point in time. We then multiply this number with the share of carbon emitted today that remains at that future point in time. This calculation yields the income loss at that future point in time caused by one ton of carbon emitted today.

By then calculating the discounted sum of all these future damages, we obtain the optimal tax. Given the assumptions we made above, we can show that that our calculations yield the following formula for the optimal tax:[18]

$$T_t = Y_t \gamma \left(\frac{\varphi_L}{\rho} + \frac{(1-\varphi_L)\varphi_0}{\rho + \varphi(1-\rho)} \right).$$

In the formula, T_t is the optimal tax per ton of carbon at time t, Y_t is the global GDP at time t, φ_L is the share of an emitted ton of carbon that remains in the atmosphere ”for ever”, $1 - \varphi_o$, is the share that disappears immediately and φ is the rate at which the rest of the carbon is absorbed by the deep oceans. The share of an emitted ton of carbon that is slowly absorbed by the deep ocean depths is therefore $(1\text{-}\varphi_L)\varphi_o$, which is the numerator in the second term in brackets.

Finally, ρ is what is usually referred to as the subjective discount rate. This is the part of the discounting that does not depend on consumption changing over time. It therefore measures the strength at which – all other things equal – we prefer to consume now relative to postponing consumption for the future. Specifically, ρ measures the rate at which our valuation of future welfare declines over time. If ρ is 1 per cent per year, that means that we value welfare one year out into the future 1 per cent less than current welfare.

As discussed above, there is a significant element here of moral and subjective consideration, in particular when we value the welfare of future generations. We therefore refrain from taking a definite stance on the issue of what a ”correct” value for ρ might be, though we note that the values normally used in the context of climate change as well as in other contexts are within the range of 0.1–2 per cent. This, however, is a very broad range; we can see that this by calculating how far into the future we need to go in order to value welfare half as much as we value current welfare. At a subjective discount rate of 0.1 per cent, the half-life of our evaluation is approximately 700 years, or around 20 generations. But at a rate of 2 per cent, the half-life is around 35 years, in other words one generation. At the latter rate, we hardly care at all about what happens in 700 years. Our valuation of welfare today is more than a million times higher.[19]

We can see from the formula that, all other things being equal, the tax is proportional to GDP. This means that it will rise in line with GDP. The tax is also proportional to the climate damage, γ. As we noted earlier, we are far from having a reliable estimate of the extent of climate damages; it may very well turn out that we need to update the value of γ in important ways. In that case, the tax should change by the same proportion as the change in γ.

Similarly, the values assigned to the parameters that control the carbon depreciation might change as new knowledge develops. For example, if a larger share of the emitted carbon remains in the atmosphere forever, the optimal tax should be raised. If a lower share of

17. The central assumption is utility is logarithmic in consumption. Then, marginal utility is inversely proportional to consumption, which balances with the assumption that damages are proportional to income. In order for this to be exact, consumption must also be proportional to income, but because the savings rate in most long-term growth models as well as in reality does not vary to any great extent, this is not of any great importance in quantitative terms. However, if the utility function has more curvature than the logarithmic function, a higher (lower) income growth to lead to larger (smaller) discounting and vice versa.

18. See Golosov, Hassler, Krusell and Tsyvinski (2014).

19. 1.02-700 is approximately 1/1,000,000.

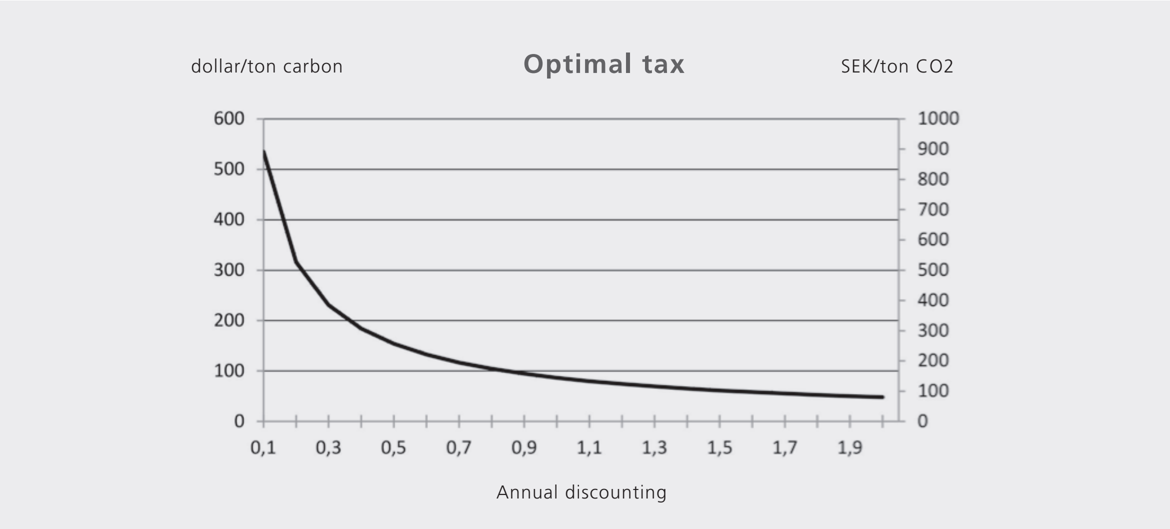

Figure 3. Optimal tax. Source: Our own calculations.

carbon disappears directly out of the atmosphere or the rate at which the atmospheric carbon is absorbed into the deep oceans falls, then the optimal tax also rises. In view of our previous discussion, these effects are expected, but we can now evaluate them in quantitative terms.

The level of the tax

Figure 3 shows the optimal tax for the parameter values we selected above and for a global GDP of 75,000 billion dollars. The x-axis represents the subjective discount rate. The left-hand side y-axis shows the tax in dollars per ton of carbon and the right-hand side axis shows the tax in Swedish kronor per ton of CO_2. Note that when a ton of coal is combusted, the carbon is combined with 2.67 tons of oxygen and produces 3.67 tons of CO_2. A tax of SEK 1,000 per ton of CO_2 is therefore equivalent to a tax of SEK 3,670 if it is instead calculated per ton of carbon.[20]

As expected, the tax is lower the less we care about the future, i.e., the higher is the discount rate. Compared to the annual damage that we calculated above at 1.8 dollars per ton, the tax is much higher for all discount rates displayed. The tax is thus mainly motivated by future damages and not by its short-run climate impact.

We also note that our formula generates results that are not far from those computed by Stern and Nordhaus. Even though they both use different models, the major difference in their respective analyses is in their choices of subjective discount rates.

Stern uses 0.1 per cent per year which, using our formula, gives a tax of SEK 890 per ton of CO_2. This is close to the tax advocated by Stern. With a discount rate of 1.5 percent per year, our formula indicates that the tax should be SEK 102 per ton of CO_2, or USD 57 per ton of carbon, which is approximately consistent with Nordhaus's analysis.[21]

It is also interesting to compare our results with the Swedish carbon tax, which is levied at SEK 1,100 per ton. This is an even higher number than that implied by the very low discount rate advocated by Stern. Finally, compare the optimal tax to the price of emission rights in the EU, which plays the role of a tax on using carbon. This price has recently fallen to just over EUR 4 per tonne, which is also considerably lower than the optimum tax even under an exceptionally high discount rate.

The price of emission rights has fluctuated significantly during the period for which the system has existed in Europe. These fluctuations have nothing to do with changes in the estimate of the harmful effects of CO_2. This illustrates a disadvantage of using quanti-

*20. The molecular weight of carbon is 12 and the molecular weight of oxygen is 16. The ratio between the weight of a CO2 molecule and a carbon atom is thus (12+2*16)/12, which is approximately 3.67.*

21. Ironically, we note that the global subsidies for fossil fuel consumption in 2011 amounted to 523 billion dollars (World Energy Outlook 2012 Factsheet). With global emissions of over 9 gigatons per year, this means that the use of fossil fuel is subsidised by about as much as it should be taxed. The politicians have found the right number, but the wrong sign!

ty restrictions instead of direct taxes. If quantity restrictions are used, the policymaker must not only estimate the harmful effects of CO_2, but also the cost of reducing its use. The latter varies over the business cycle and is different over the short and long term. An optimal quantity restriction will therefore vary over time even if the negative effects of CO_2 remain constant.

Determining the optimal variation over time is quite difficult and policies can go seriously wrong if we are not sure about the costs of reducing carbon emissions. In our opinion, this is a serious disadvantage of using quantity restrictions. It should be noted, however, that if tipping points or critical levels for atmospheric CO_2 concentrations or temperature are identified, the argument for quantity restrictions would be strengthened. Under such circumstances, the value of not exceeding the thresholds could become very large, and since it may be very difficult to determine the relation between taxes and emission quantities, a direct quantity restriction may be preferable.

Climate damage and growth

Most climate-economy models are based on the assumption that climate damage affects the level of GDP but not its rate of growth. It is easy to include growth effects in our formula for the optimal tax, and if they are included they have a considerable effect on the calculation, particularly if they are permanent or at least long-term.

We have assumed above that CO_2 emissions have increased the quantity of carbon in the atmosphere by around 200 gigatons and that, as a result, global GDP will be 0.5 per cent lower. With a discounting rate of 1 per cent, only one hundredth of the effect on permanent growth is required to justify the same tax.

In other words, if the current extra 200 gigatons lowered growth permanently by one tenth of the level effect, i.e. by 0.05 per cent per year, the tax would need to be 10 times what we calculated. It is easy to realize that it is extremely difficult to measure such small effects. This significantly increases the uncertainty in the calculation of the optimal tax.[22]

The importance of the initial CO2 concentration

In our calculations, we have assumed that the speed at which excess carbon disappears from the atmosphere does not depend on the CO_2 concentration in the atmosphere and oceans. Particularly, if CO_2 concentrations are high, this assumption may be problematic in that high concentrations can lead to a lower capacity to absorb more carbon.[23]

However, in a noted article in the journal Nature[24], the authors show that this effect is offset by the fact that a higher CO_2 concentration has a diminishing effect on the energy balance and that the warming of the atmosphere is slowed down by the fact that the ocean warms up slowly.

The gist of this argument is that the increase in the average global temperature in both the short and the long term is roughly proportional to the accumulated quantity of carbon emitted. The anticipated effect, according to the article, is that the temperature rises by 1.5 degrees per 1,000 gigatons of emissions.[25]

So far we have emitted approximately 500 gigatons, which would thus lead to a rise in temperature of around 0.75 degrees. This result can also be used as the basis for a calculation of the optimum tax, given an assumption concerning the relationship between damage and temperature. This calculation is simple if we assume that an increase in the temperature produces the same damage to the economy, no matter how high the temperature is.

If we instead assume that the damage to the margin is greater the higher the temperature rises, the optimal tax, unlike in the above calculation, will depend on the forecasts we make for future emissions. If the use of fossil fuels in the future is high, future marginal cost will be high, thus justifying a higher tax already today – and vice versa.

22. It also shows that other measures that may have long-term effects on growth, such as taxes and education policy, can have very large long-run impacts on welfare.

23. Extensive research into this is being carried on as part of our Mistra SWECIA project.

24. Matthews, Gillet, Stott and Zickfeld (2009).

25. The authors also indicate a 95 per cent confidence interval of 1 to 2.1 degrees per 1,000 gigatons of carbon.

THE OPTIMAL CARBON TAX – A THREAT TO GROWTH AND WELFARE?

The introduction of a carbon tax at a reasonable level has consequences for the economy. Higher energy prices reduce GDP because energy is an important input in the production process.

In the short term, energy use is almost proportional to GDP and the only way to radically reduce the use of energy is to switch off machines and consume less of such services as transports. Because of this, it is easy to get the impression that energy taxes could constitute a threat to welfare and growth. We consider this to be a fallacy and believe that the introduction of a reasonable global carbon tax would have very limited effects on economic growth.

The fact that energy use is proportional to GDP is not a good approximation in the slightly longer term. Instead, there are strong indications that there is great potential for replacing energy with other factors of production and new technologies and thereby increase the quantity of GDP produced per unit of energy.

In the short run (over a few years), the share of GDP that is spent on energy varies substantially following changes in energy prices. In a longer-term perspective, however, that share is quite stable at a few per cent. This indicates that the potential of increasing energy efficiency if prices go up, for example because of a tax, is high and that this can take place without a large loss of production.

Also, a comparative international perspective indicates that there are substantial possibilities for substitution. Figure 4 shows energy efficiency for a number of countries at around the same level of development and with the same climate. The measure is GDP in dollars per unit of energy (all forms of energy are converted to kilograms of oil equivalent). The blue bar shows the average energy efficiency of each country over the period 2003-2011.

As we see in the figure, there is an extremely large variation in energy efficiencies, more than a factor of

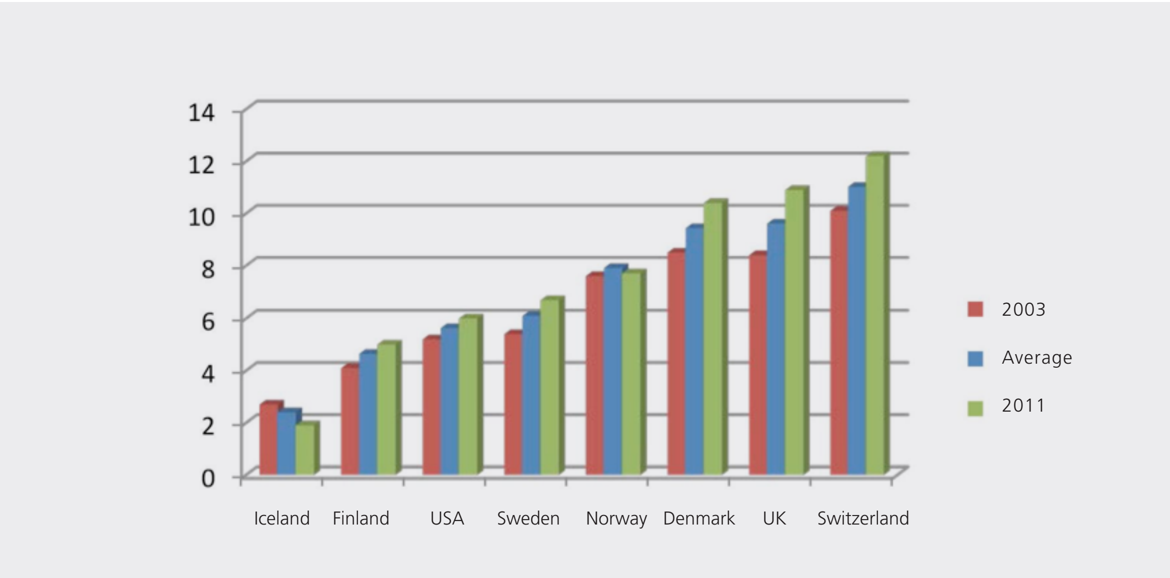

Figure 4. Energy efficiency GDP (in purchasing power-adjusted dollars in 2005) per unit of energy (kilograms of oil equivalent). Source: World Bank

1:5 between the least and the most energy-efficient nations. We also note significant differences within the Scandinavian countries including Finland. Denmark is more than twice as energy-efficient as Finland, and Sweden is 30 per cent more efficient than Finland, but more than 20 per cent less efficient than Norway. We believe that these differences likely reflect differences in energy prices caused by differences in access and policy.

How effective is policy?

As we have discussed above, there are strong arguments supporting the assertion that the costs of climate change are very unevenly distributed across the world. This means that different countries should be expected to have different views on what constitutes a reasonable carbon tax.

In theory, these different views could be eliminated through a properly designed international transfer scheme. Then all countries would have the most to gain from the introduction of a global optimal tax as calculated above. However, such a system is unlikely to be introduced for several reasons, among them political ones. It is therefore important to analyse the effect of taxes levied in some individual regions but not in others.

Also, it is of interest to analyse how the use of fossil fuels would be affected by the introduction of global or regional carbon taxes. A closely related question is what happens to the use of fossil fuels if technological development increases energy efficiency in the economy.

Conventional oil: finite supply

To analyse the issues just raised, it is key to take supply factors into consideration. The price of conventional oil is set on the world market. It has little to do with the producers' extraction costs, but reflects the fact that there is a limited quantity of conventional oil.

In a market like this, supply is finite in a long-term perspective and will not be affected unless the tax is set so high that it is not profitable to extract oil.

In such a supply situation, a global tax at a realistic level does not affect the price for the consumer, but is borne by the producers in the form of lower profits. Thus, consumption is not significantly affected.

This also applies to technological development which, like a tax, reduces demand – the demand curve shifts inwards. If supply is inelastic, i.e. not price-sensitive, such shifts in demand will only affect the price received by the producers and not the quantity used.

Our reasoning is illustrated in figure 5. A tax or a new energy-saving technology shifts demand inwards/downwards from line D1 to D2. Because the supply (line S) is vertical, the equilibrium quantity is unaffected and the price falls to a level where demand equals the quantity that prevailed before the demand shift.

In this case, taxes or technological development therefore do not affect total use. Although the figure represents a static calculation, a similar argument can be made for the long-run. As long as the price of oil is not pushed down to the extraction cost at any given time, all the oil will be used up sooner or later. With taxes that vary over time, the time profile of oil use can be changed and there is some, albeit relatively modest, economic value in postponing the use of oil through falling taxes.

When supply is inelastic as described above, taxes in only some regions of the world have even smaller effects and the same applies to region-specific technological development, such as the introduction of energy-efficient cars.

If one region, say the EU, raises taxes or introduces energy-efficient cars, it will lead to a slightly lower price on the world market. The lower price will lead to increases in demand that exactly offset the reduced demand for oil in the region that introduced the tax.

One might think that quantity restrictions would lead to different consequences, but that is not the case. If a region reduces its demand with the aid of quantity restrictions, this also leads to a reduction in the price at the world market that implies that global demand remains unchanged.

The discussion so far has centred on conventional oil, which we defined as a commodity with a limited supply and low extraction costs in relation to its price. However, for other fossil fuels, in particular coal, the situation is quite different.

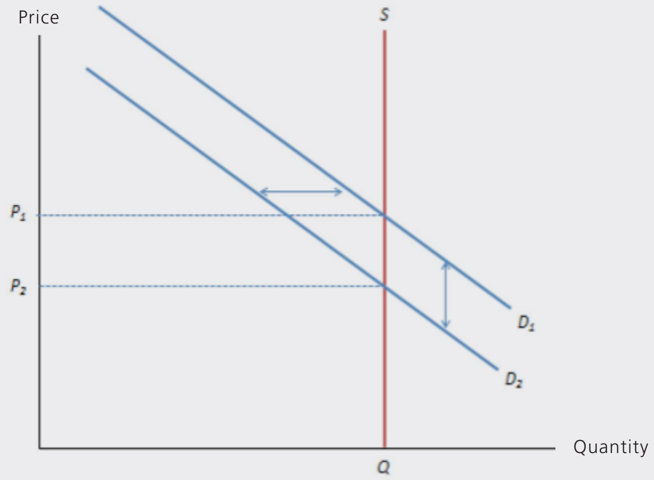

Figure 5. Supply and demand when the supply is not price-sensitive.

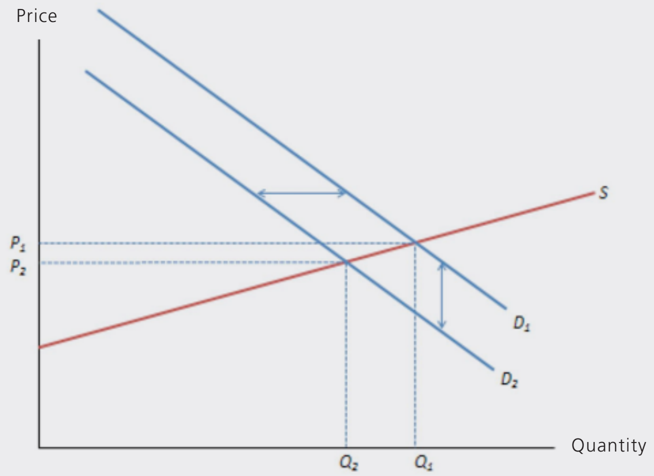

Figure 6. Supply and demand when the supply is price-sensitive.

Coal: price-dependent supply

Compared with the quantities of conventional oil, there are enormous quantities of coal. The price of coal therefore primarily reflects the producers' extraction costs. If a country discovers conventional oil, it may become rich from it, but not if it discovers coal.

Also, because coal is relatively expensive to transport in relation to its price, most of it is consumed in the region where it is produced.

Under such supply conditions, the effect of carbon taxes is quite different, whether they are global or regional. Because the price of coal is close to the extraction cost, there is not much scope for reducing the price if a tax is introduced. Therefore, a tax does not need to be particularly high for a large part of today's coal mining to cease to be profitable. A tax of 50 dollars per ton of coal would be approximately as high as current extraction costs and could therefore have a significant effect on the use of coal. Because the coal market is not global, a tax in Europe should not lead to an increase in coal use in China, for example.

The analysis is the same for reductions in demand due to technological development that leads to increased energy-efficiency or to green energy competing against coal. The lower demand reduces the price and this leads to less coal mines being profitable and thus to less coal use.

The market equilibrium is illustrated in figure 6. The difference relative to figure 5 is that the supply here is assumed to be price-dependent. Lower prices will now mean that it is not profitable to sell as much. A tax or technological development leads to a fall in demand in the same way as in figure 5, but because the supply is price-dependent, the equilibrium quantity falls from Q1 to Q2.

The conclusion from our reasoning is therefore that it is difficult or impossible to affect the use of conventional oil, whereas the use of coal can definitely be affected by carbon taxes or other policy instruments.[27]

It is important to note that the remaining quantities of conventional oil are not of great significance for climate change. We have previously noted that the quantity of carbon in the atmosphere has increased from around 600 gigatons to 800 gigatons. The quantity of conventional oil still remaining to be extracted is estimated at approximately 200 gigatons. Even if all this oil is used, it will have a relatively modest impact on the climate.[28]

The quantity of coal reserves is estimated to be at least 20 times that of conventional oil, which makes coal a considerable threat to the climate. It is therefore cause for concern that so much of the policy against the climate threat is directed at futile attempts to influence the use of conventional oil and its end products, such as gasoline and aircraft fuel.

New fossil fuels

So far we have discussed coal and conventional oil, which may be seen as extremes in terms of their supply structure. High oil prices in combination with technological development have led to the emergence of fossil fuels between these extremes in recent years.

Some examples include deep offshore oil, tar sands and hydraulic fracturing (fracking).

From a climate perspective, there is reason to view the emergence of these new fossil fuels with great concern since there are potentially large quantities of these resources.

On the other hand, for the moment these new sources of fossil fuel expensive to extract, which means that their use can be affected by taxes. Unlike conventional oil, a reasonable carbon tax and the abolition of subsidies should mean that some of these technologies will not be used on a large scale.

27. The only possible solution would be for someone to buy up existing oil sources and undertake never to extract the oil in them.

*28. If we use our rough estimates from above and assume that half the carbon from the combusted oil is assimilated rapidly by the ocean surface and plants and that the climate sensitivity is 3 degrees, the additional warming would be 3*ln(900/800) = 0.35 degrees.*

WE THANK OUR COLLEAGUES WITHIN MISTRA-SWECIA
AND PARTICULARLY MARKKU RUMMUKAINEN FOR HIS COMMENTS.

THIS REPORT IS A TRANSLATION OF SNS ANALYS NO 14 2013.

REFERENCES

Arrhenius, Svante (1896), "On the Influence of Carbonic Acid in the Air upon the Temperature of the Ground", Philosophical Magazine and Journal of Science, vol. 41, no. 5, pages 237–276.

Dell, Melissa, Benjamin F. Jones and Benjamin A. Olken (2012), "Temperature Shocks and Economic Growth: Evidence from the Last Half Century", American Economic Journal: Macroeconomics, vol. 4, no. 3, pages 66–95.

Golosov, Mikhail, John Hassler, Per Krusell and Aleh Tsyvinski (2014), "Optimal taxes on fossil fuel in general equilibrium", Econometrica, in the process of being published.

Hassler, John and Per Krusell (2012), "Economics and Climate Change: Integrated Assessment in a Multi-Region World", Journal of the European Economic Association, vol. 10, no. 5, pages 974-1000.

IPCC (2007), Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change, Cambridge and New York: Cambridge University Press.

Lenton, Timothy M., Hermann Held, Elmar Kriegler, Jim W. Hall, Wolfgang Lucht, Stefan Rahmstorf and Hans Joachim Schnellnhuber (2008), "Tipping elements in the Earth's climate system", Proceedings of the National Academy of Sciences, vol. 105, no. 6, pages 1786-1793.

Levitan, Dave (2013), "Quick-Change Planet: Do Global Climate Tipping Points Exist?", Scientific American, 25 March 2013. Matthews, H. Damon, Nathan P. Gillet, Peter A.

Stott and Kirsten Zickfeld (2009), "The proportionality of global warming to cumulative carbon emissions", Nature, vol. 459, pages 829–832, 11 June.

Myndigheten för samhällsskydd och beredskap [The Swedish Civil Contingencies Agency] (2012), "Evaluation of Risk. Economic evaluation of health risks at present and in future", https://www.msb.se/RibData/Filer/ pdf/26246.pdf.

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

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

Pigou, Arthur (1920), The Economics of Welfare, London: Macmillan.

Stern, Nicholas (2007), The Economics of Climate Change: The Stern Review, Cambridge and New York: Cambridge University Press.

JOHN HASSLER
AND PER KRUSELL

John Hassler and Per Krusell are both Professors of Economics at the Institute for International Economic Studies at Stockholm University.

Mistra-SWECIA

Mistra-SWECIA is an interdisciplinary research programme on climate, impacts and adaptation. The programme brings together researchers in the fields of climate science, economics, life sciences and social sciences. Since 2008, Mistra-SWECIA has deepened our understanding of how research on climate change, its impacts and economic implications can support stakeholders.

Mistra SWECIA is funded by the Swedish Foundation for Strategic Environmental Research. The programme involves SMHI, SEI, Lund University and Stockholm University, SMHI is the programme host. □

Read more about Mistra-SWECIA at www.mistra-swecia.se/en

MISTRA SWECIA
CLIMATE, IMPACTS & ADAPTATION

WWW.MISTRA-SWECIA.SE