

FOIA MARKER

**This is not a textual record. This is used as an
administrative marker by the William J. Clinton
Presidential Library Staff.**

Collection/Record Group: Clinton Presidential Records
Subgroup/Office of Origin: Council of Economic Advisers
Series/Staff Member: Jeffrey Frankel
Subseries:

OA/ID Number: 13722
FolderID:

Folder Title:
[Pizer Climate Change Escape Clause]

Stack:	Row:	Section:	Shelf:	Position:
S	20	4	10	3

**Prices vs. Quantities Revisited:
The Case of Climate Change**

William A. Pizer

August 30, 1997

**Resources for the Future
1616 P Street, NW
Washington, DC 20036
Telephone 202-328-5000
Fax 202-939-3460**

Prices vs. Quantities Revisited: The Case of Climate Change

William A. Pizer

Abstract

Uncertainty about compliance costs causes otherwise equivalent price and quantity controls to behave differently. Price controls – in the form of taxes – fix the marginal cost of compliance, leading to uncertain levels of compliance. Meanwhile quantity controls – e.g., tradeable permits or quotas – fix the level of compliance, leading to uncertain marginal costs. This fundamental difference in the face of cost uncertainty leads to different welfare outcomes for the two policy instruments. Seminal work by Weitzman (1974) clarified this point and derived conditions under which one policy is preferred to the other.

Such a situation – where tax and permit policies have been proposed and cost uncertainty is pervasive – accurately describes the current debate concerning policies to reduce greenhouse gas emissions, specifically carbon dioxide generated by fossil fuel combustion. Despite the considerable attention given to the political, welfare and revenue differences between taxes permits, only anecdotal evidence has addressed the efficiency difference due to uncertainty about costs.

This paper simulates uncertainty in an integrated climate economy model in order to compare the efficiency of taxes and permits. The results indicate that an optimal tax policy generates gains which are *five times higher* than the optimal permit policy – a \$337 billion dollar gain versus \$69 billion. This result follows from the basic Weitzman intuition that relatively flat marginal benefits/damages favor taxes, a feature that drops out of standard assumptions about likely climate damages.

An alternative policy which uses an initial distribution of tradeable permits to set a target emission level, but then allows additional permits to be purchased at a fixed price, leads to an even higher welfare gain. Such a system would preserve the political benefits of a permit system – namely preferential distribution to large emitters – and the efficiency benefits of a tax system – e.g., a hedge against unexpectedly high costs. This combined permit-fee system may provide a more credible domestic alternative to a pure permit system which could prove too costly or a pure tax system which has little political support.

Key words:

JEL Classification No(s):

Contents

1 Introduction	1
2 Model Overview	3
2.1 Description	3
2.2 Key Modeling Assumptions	3
2.2.1 Trends	4
2.2.2 Damages	4
2.2.3 Costs	6
3 Single Period Policy Simulations	7
3.1 Weitzman's Result	7
3.2 Marginal Costs and Benefits	9
3.3 Welfare Consequences of Pure Tax and Permit Mechanisms	13
3.4 Combining Taxes and Permits	14
4 Dynamic Policy Under Uncertainty	16
4.1 Baseline Emissions	16
4.2 Optimal Permit Policy	18
4.3 Optimal Tax Policy	19
4.4 Combined Permit-Fee Mechanism	22
4.5 Costs Only Comparison	23
5 Conclusion	25
A Model Specification	26
A.1 Economic Behavior	26
A.2 Long-term Growth, Climate Behavior and Damages	28
A.3 Social Welfare	31
B Measuring Uncertainty	33
C Factors Which Complicate the Weitzman Analysis	34
C.1 Non-linearities and Non-additive Shocks	34
C.2 Correlation of Cost/Benefit Shocks	39
C.3 Truncation	39
C.4 Discounting	41

List of Figures

1	Deadweight Loss of Taxes Versus Permits	9
2	Distribution of Marginal Costs and Benefits in 2010	10
3	Expected Marginal Costs and Benefits in 2010	11
4	Histogram of Uncontrolled 2010 Emissions Levels	12
5	Welfare Consequences of Pure Tax and Permit Instruments in 2010	13
6	Net Expected Policy Benefits of Alternative Tax/Permit Policies in 2010	15
7	Simulated CO ₂ Emission Distribution vs IPCC Scenarios	17
8	Optimized Permit Level – Dynamic Policy	19
9	Optimized Tax Level – Dynamic Policy	20
10	Distribution of Emission Levels Under Alternative Policies	21
11	Optimal Initial Permit Distribution in a Combined Permit-Fee System	22
C.1	Distribution of $\frac{dq}{d(\max)}$	38
C.2	Distribution of Marginal Benefits in 2010	40
C.3	Distribution of Discount Factors in 2010	41

List of Tables

1	Comparison of Assumptions Across Models	5
2	“Cost-only” Comparison of Alternative Policies	24
B.1	Marginal distributions of uncertain economic parameters	35
B.2	Discrete distributions of uncertain climate/trend parameters	36
B.3	Description of fixed parameters	37

Prices vs. Quantities Revisited: The Case of Climate Change

William A. Pizer¹

1 Introduction

Seminal work by Weitzman (1974) drew attention to the fact that uncertainty about costs leads to a potentially important efficiency distinction between otherwise equivalent price and quantity controls. Despite this well-known observation and its relevance in the climate change policy context, most of the debate concerning carbon taxes versus permits has centered on political, legal and revenue concerns. This paper attempts to fill this gap by examining the efficiency properties of permits and taxes in the climate change arena.

The basic distinction among instruments arises because taxes in a competitive, profit-maximizing world will tend to fix the marginal cost of abatement at the specified tax level. In a world where costs are uncertain, this will lead to a range of possible abatement levels and emission outcomes. In contrast, a permit system will precisely limit emissions but lead to a range of potential cost outcomes. This difference in emission and cost outcomes leads to a difference in social welfare when coupled with a model of reduction benefits.

In the case of climate change, much of the cost uncertainty arises due to uncertainty about the level of future emissions. The Intergovernmental Panel on Climate Change (1992) gives a range of CO₂ emission levels in 2025 of between 8.8 and 15.1 GtC. The cost of attaining a particular target, say the 1990 emission level of 7.4 GtC, will obviously fluctuate widely depending on the level of future uncontrolled emissions.

In addition to the baseline, however, there is considerable uncertainty about the cost of reducing emissions below the baseline. A study by Nordhaus (1993) reports that a \$30/tC tax might reduce emissions anywhere from 10 to 40%. While some models predict that a \$300/tC would virtually eliminate emissions, other models require \$430/tC. This wide range of reduction estimates only compounds the uncertainty about baseline to generate extreme uncertainty about the cost of a

¹Fellow, Quality of the Environment Division, Resources for the Future.

particular emission target.

With such motivation in mind, this paper adapts a popular integrated assessment model (Nordhaus 1994b) for analyzing alternative policies under uncertainty. In particular, the simulations are sped up using a technique developed in Pizer (1996) and uncertainty about a wide range of model parameters is introduced based on both Nordhaus (1994b) and Pizer (1996). Key modeling assumptions are discussed in the next section while specifics are covered in the appendix.

Before presenting the dynamic policy results in Section 4, Section 3 gives a simpler static analysis. This analysis, following Weitzman (1974), uses marginal cost and benefit estimates derived from the full dynamic model to examine optimal policy in a single year (2010). This simple static analysis provides the intuition for the full dynamic policy analysis and closely replicates the dynamic results in that year.

These results indicate that an optimal permit policy would begin with a 13 GtC target in 2010 and would rise gradually over time to accommodate some growth. This policy generates \$69 billion in expected net benefits (e.g., benefits-costs). Importantly, this gain is sensitive to correctly setting the target. Slightly lower targets lead to dramatic welfare losses. Meanwhile, the optimal tax policy starts at \$7.34/tC in 2010 and rises gradually over time. In contrast, this policy generates \$337 billion in net benefits and the gain is much less sensitive to miss-setting the tax level.

As an alternative to the pure tax and permit policies, a combined permit-fee policy is proposed (Roberts and Spence 1976; Weitzman 1978; McKibbin and Wilcoxon 1997). Such a mechanism would involve an initial distribution of tradeable permits with additional permits available at a specified price. This system turns out to be only slightly more efficient than a pure tax system but provides the political advantages of a permit system. Namely, permits can be distributed to existing emission sources to reduce the tax liability they would face under a pure tax system.

2 Model Overview

2.1 Description

The integrated climate economy model used in this analysis is based on a stylized representation of economic activity and climate behavior. This stylized approach, based on Nordhaus (1994b), is appropriate given this paper's focus on *uncertainty*. While additional detail might improve the results for a given set of assumptions, it is unlikely that such embellishments would affect the range of predicted outcomes or the insight concerning optimal policy choice.

The economic model is based on one-sector model of global economic activity. Global capital and labor are combined to produce a single output each year which is either consumed or invested in additional capital. A representative agent chooses an allocation of consumption across time which maximizes her utility.

Climate change enters the model through the emission of greenhouse gases caused by economic activity. These emissions accumulate in the atmosphere and lead to a higher equilibrium temperature. This higher temperature then causes damages by reducing output according to a simple quadratic damage function.

The opportunity to reduce the effect of climate change arises from the potential to reduce emissions of greenhouse gases by using a more expensive production technology. In particular, there is a cost function describing the reduction in output required to reduce emissions by a given fraction. This cost is born in the current period, while the consequences of reduced emissions persist far into the future due to the longevity of greenhouse gases in the atmosphere.

Appendix A describes the model in more detail.

2.2 Key Modeling Assumptions

While the previous section gave a broad picture of the model and the appendix provides a detailed picture, it useful to summarize key assumptions in the model. These assumptions tend to drive the results in spite of the range of uncertainty considered. Table 1 compares this model with other

commonly referenced economy-climate models and is followed by a discussion of the major issues.

2.2.1 Trends

An important issue for modeling long-term phenomena is the treatment of trends. In particular, assumptions about exogenous population growth, productivity improvements and changes in energy efficiency/carbon content will determine the baseline of uncontrolled greenhouse gas emissions. This model borrows from the work of Nordhaus and Yohe (1983), Nordhaus (1994b) and Nordhaus and Popp (1997) to characterize the range of global population and carbon content trends. Those same sources provide estimates of a future productivity growth slowdown. The initial rate of productivity growth is based on econometric estimates using U.S. data (Pizer 1996).

CHECK WORLD BANK/UN DOCUMENTATION ON FORECASTS; REFERENCE NORDHAUS AND YOHE CITES; SUMMARIZE NORDHAUS/IPCC SECTION ON EMISSION TRENDS.

2.2.2 Damages

Damages are perhaps the least understood aspect of climate change and at the same time one of the most important. In this model, damages are modeled as a quadratic function of temperature change, in turn determined by GHG concentrations. The model is calibrated so that a three degree temperature change is equally likely to reduce GDP by 0, 0.4, 1.3, 1.6, or 3.2 percent. The warming due to a doubling of greenhouse gas concentrations (due to occur by the end of the next century) is similarly calibrated to be either 1.5, 2.2, 2.9, 3.7 or 4.4 degrees with equal probability.

Among these assumptions, there are two that tend to drive the final results – in particular by suggesting a relatively flat marginal benefit schedule. The first is that climate damages are a gradual phenomena (represented by a quadratic function). The fact that damages from increased GHG concentrations rise smoothly and gradually contributes to a relatively flat marginal benefit schedule (e.g., marginal damages from increased emissions). In contrast, a damage function that involved critical phenomena – such as the breaching of a concentration threshold generating dramatically higher damages – would lead to a much steeper or stepwise damage function (though uncertainty

Table 1: Comparison of Assumptions Across Models

Model	Population	Productivity	Emissions	Abatement Costs	Climate Change	Damages
This Paper	1995 growth rate of 1.24% declining geometrically 0.25% to 3.3% annually	1995 growth rate of % declining geometrically 0.2% to 2.5% annually	1995 growth rate (emissions/output) of -0.15 to -2.3% declining at the same rate as productivity growth	20% emission reductions from BAU baseline cost between 0.03% and 0.13% of GDP, 50% reductions cost 0.36-1.8%	50-80% of emitted CO ₂ remains in the atmosphere; CO ₂ doubling causes 1.5-4.4° rise in temperature	3° rise in temperature causes 0-3.2% loss of GDP
IPCC 1992						
Edmonds						
Mann & Richels						
Dowlatabadi						

about the threshold level would tend to smooth the expected damage function).

The second assumption – which is less controversial – is that damages are related to the *stock* of GHGs in the atmosphere and not the annual *flow*. This contrasts with traditional pollutants, such as particulates, SOx, NOx, etc., whose damages are related to annual flows.² Stock pollutants by their nature will have relatively flat benefit curves associated with reductions since the reductions in any period have a relatively small effect on the total stock. If the total stock does not change much, the marginal benefit cannot change much.

As an example of this stock pollutant phenomena, imagine a pollutant, like carbon dioxide, whose atmospheric stock decays naturally at a rate of less than 1% per year. At an equilibrium, it can be shown that the emissions in any year will equal to 1% of the total stock.³ If marginal benefits decline linearly from some initial level – say \$10/ton – to zero over the entire balance of the stock, the decline in marginal benefits over a 100% reduction in emissions in a single year will be only \$0.10 (1% of \$10).

In the absence of a compelling argument that climate damages exhibit sudden, catastrophic increases over a gradual increase in GHG concentrations, this fact that GHGs are a stock pollutant turns out to provide a strong rationale for emission taxes via the flatness of the marginal benefit schedule.

2.2.3 Costs

Reductions in GHG emissions from any given baseline are assumed to reduce global GDP according to a simple power rule,

$$\text{fractional reduction in global GDP} = b_1 (\text{fractional reduction in GHGs})^{2.887}$$

where the parameter b_1 takes on the values 0.027, 0.034, 0.069, 0.080, 0.133 with equal probability. That is, complete elimination of GHG emissions is forecast to cost between 2.7% and 13.3% of

²To some extent this is an arbitrary distinction. Any pollutant can be viewed as a stock pollutant when viewed on a suitably small time-scale. For conventional pollutants, this might be a period of a day or a single hour.

³E.g., annual emissions will exactly equal the annual decline of the stock – 1%.

global GDP.⁴

More important than the particular numbers entering the cost function are the assumptions that (1) marginal cost rises more and more steeply as additional reductions are undertaken and (2) that the choice of emission level is an annual decision involving an annual cost function. Both of these points are crucial because they affect the relative slope of marginal costs (in turn affecting the difference in expected welfare between taxes and permits). If, for example, we believe that substantial reductions of GHGs will involve the development of new carbon-free technologies, it seems reasonable that the marginal cost of reducing emissions will eventually flatten. This will diminish the argument that marginal benefits are relatively flat compared to costs. Alternatively, many decisions to reduce emissions – such as those reductions resulting from investment in innovative research or in new capital – are made over horizons of decades rather than annually. In such cases, costs and benefits should be viewed over equally long intervals, again weakening the argument that marginal benefits are relatively constant.

3 Single Period Policy Simulations

Given the long-term nature of climate change, a policy to reduce GHG emissions inevitably involves decisions over periods of years or decades. However, understanding the differences between GHG tax and permit mechanisms under uncertainty is complicated when policies are viewed as paths for tax and permit levels, rather than single-period, single-value choices. With that in mind, this section presents results for a single-period policy analysis – reductions in GHG emissions in the year 2010. Since a single dimension captures the range of policies in this context, graphs can be used to view policy consequence. The next section will discuss the dynamic policy results.

3.1 Weitzman's Result

The analysis presented in Weitzman (1974) concerns the choice of policy instrument used to regulate a market where either political considerations or market failure requires government interven-

⁴See Nordhaus (1993) for further details.

tion. A price (tax) or quantity (permit) instrument is at the government's disposal and the question posed by Weitzman is which leads to the best welfare outcome, measured as net social surplus.⁵

With complete certainty concerning costs, price and quantity controls can be used to achieve the same outcomes. For every price, there is a profit-maximizing level of production where price equals marginal cost and for every level of production there is an associated marginal cost. Note that while uncertainty solely about *benefits* leads to uncertain welfare outcomes, it does not lead to uncertainty about the level or marginal cost of production – these are determined by the structure of costs. Therefore, the two instruments (which affect production) can be used to obtain exactly the same set of potential welfare outcomes when costs are known, though the welfare outcomes may be uncertain due to unknown benefits.

The interesting case arises when *costs* are not certain. In this case, fixing the marginal cost through a price instrument leads to an uncertain level of production. Correspondingly, fixing the level of production with a quantity instrument leads to an uncertain marginal cost.

Weitzman's basic result was that price instruments would be favored when the marginal benefits were relatively flat and quantity instruments would be favored when the marginal costs were relatively flat. In particular, he derived an expression for the difference in expected welfare between the price and quantity instrument:

$$\Delta = \frac{\sigma^2}{2C''^2}(B'' + C'') \quad (1)$$


where σ^2 is the variance of the shocks to the marginal cost schedule, C'' is the slope of the (linear) marginal cost schedule and B'' is the slope of the (linear) marginal benefit schedule.⁶ Since the marginal benefit schedule is assumed to be downward sloping ($B'' < 0$), the price instrument is preferred when $|B''| < |C''|$ and the quantity instrument is preferred when $|B''| > |C''|$.

The intuition behind this result is that when marginal benefits are relatively flat ($|B''| < |C''|$), the ex post (after cost uncertainty is resolved) optimal price is relatively constant. In contrast,

⁵ Here and throughout it is assumed that the quantity instrument is an *efficient* quantity instrument; e.g. a tradeable permit system with negligible transaction costs.

⁶ This result is derived for the case of linear marginal costs and benefits, where uncertainty enters as small shifts to each curve. The uncertainty about costs is assumed to be independent of the uncertainty about benefits.


Figure 1: Deadweight Loss of Taxes Versus Permits



when marginal benefits are relatively steep ($|B''| > |C''|$), the ex post optimal quantity is relatively constant. Therefore, fixing the ex ante price when the optimal ex post price is fairly constant and fixing the ex ante quantity when the optimal ex post quantity is fairly constant leads to smaller deadweight losses. This is shown visually in Figure 1.

3.2 Marginal Costs and Benefits

While there are a number of strong assumptions in the Weitzman analysis which are inappropriate for examining climate change policy, it remains a sensible starting point. How do the marginal costs and benefits of GHG emission reductions in the year 2010 compare in terms of relative slopes? In order to answer this question, the integrated assessment model of uncertainty outlined in Section 2 and described in detail in Appendix A is used to compute two quantities. First, the welfare associated with different levels of emissions in 2010 is computed in net present (2010) value terms for several thousand randomized trials. Marginal benefits are computed by numerically differentiating the derived schedule of benefits. Second, the cost associated achieving different levels of emissions in 2010 is computed based on the model's cost function for the same set of


Figure 2: Distribution of Marginal Costs and Benefits in 2010^a

^aMarginal costs are based on dollar value (\$2010) of lost global GDP in order to reduce emissions at or below the indicated level. In those cases where uncontrolled emissions are below the indicated level, the marginal cost is zero. Marginal benefits are based on the dollar value (\$2010) of the net present value of forgone damages at the given emission levels. These forgone damages hold constant all future emissions at their baseline level. Values are expressed in \$2010 and, due to different discount rates across states of nature, will not be weighted equally when balancing costs and benefits.

randomized trials. This cost is similarly numerically differentiated to obtain marginal costs in the year 2010. In some trials and for some levels of emissions, the cost is zero since the given emission level may be higher than actual uncontrolled emissions.

These two calculations result in a distribution of marginal benefits and marginal costs at different levels of emission. Figure 2 attempts to summarize these distributions by showing how different quantiles of marginal costs and benefits vary over the range of emissions considered. Keeping in mind that 1990 GHG emissions were around 8.5 GtC, the left-hand panel indicates that achieving 1990 emission levels in 2010 would involve a marginal cost of between zero and \$180/ton – a very wide range. This large variation occurs for two reasons: (1) marginal costs are assumed to rise steeply given the specified cost function, and (2) the baseline emissions in 2010 are not known with certainty. This panel essentially depicts a distribution of fairly steep curves whose horizontal intercept is unknown.

The right panel, in contrast, indicates rather constant but unknown benefits. As suggested earlier, the fact that damages due to GHGs depend on their *stock* in the atmosphere rather than


Figure 3: Expected Marginal Costs and Benefits in 2010^a

emissions in any one year – and the fact that GHGs remain in the atmosphere for a very long time – makes marginal benefits insensitive to the level of annual emissions. For example, in 1995 the concentrations of GHGs in the atmosphere was 760 GtC and annual emissions were around 10 GtC. Therefore the difference in concentrations between no reductions and 100% reductions in any one year is small – making it unlikely that the marginal benefit of the first ton of reductions is much different than the marginal benefit of the last ton.

Taking expectations across these marginal costs and benefits yields the schedules shown in Figure 3. Under the assumptions made by Weitzman, the optimal permit level is simply emission level where expected marginal benefits equal expected marginal costs and the optimal tax level is similarly the marginal benefit at that intersection. Thus P^* and Q^* indicate the optimal tax and permit policies in 2010 for controlling GHG emissions.

Visually, it is apparent that expected marginal benefits are flatter than expected marginal costs. Calculating the slopes at the intersection, $B'' = -0.0012$ and $C'' = 5.4$. Further, setting $\sigma^2 = \text{var}(MC_{\text{emission} = 11,000})$ yields $\sigma^2 = 270$. This allows a rough calculation of the welfare gain of taxes over permits using (1): $\Delta = \frac{270}{2 \cdot 5.4^2} (-0.0012 + 5.4) \approx \25 billion. Discounting this to 1995 (the base year of the model) with a 6% discount rate generates an estimated gain of \$10 billion

Figure 4: Histogram of Uncontrolled 2010 Emissions Levels




*GHG emissions include CO₂ and CFCs. 1990 emissions are based on 7.4 GtC of CO₂ and 1.1 GtC of CFCs (p. 71, Nordhaus 1994b). The 1990 emissions also correspond to data from IPCC(1992), p. 12, using a factor of 1300 to convert tons of CFCs into tons of carbon equivalent global warming potential (GWP) (Lashof and Ahuja 1990). The optimal 2010 permit level is discussed in the text.

from using taxes instead of permits – just in the year 2010.

In addition to the welfare difference, it is interesting to compare the optimal quota in 2010 to the range of actual emissions. Figure 4 shows 1990 emissions (a target under consideration), the optimal 2010 permit level, and the distribution of uncontrolled emissions in 2010. Roughly half of the emission scenarios involve uncontrolled emissions *below* the optimal permit level (versus around 3% below the 1990 level). That is, implementing the optimal permit policy would result in non-binding targets half the time.

Why is the optimal permit policy so loose? Intuitively, uncertainty about baseline emissions is large relative to reductions and a large amount of reductions is costly. Therefore committing to an emission target which is almost surely below all forecasts (such as 1990 emissions levels) will involve extremely high costs in the event of high growth and high baseline emissions. These high costs, based even on the highest estimated benefits, lead one to prefer a less stringent (higher) target.

Figure 5: Welfare Consequences of Pure Tax and Permit Instruments in 2010



3.3 Welfare Consequences of Pure Tax and Permit Mechanisms

While the previous analysis based on Figure 3 provides important intuition and a rough approximation of the welfare consequences of taxes and permits, it fails to capture several important failures of the Weitzman assumptions. While the intuition behind these failures and their individual consequences is discussed in Appendix C, in this section their net effect is summarized. In particular, the net welfare gains of alternative tax and permit policies are shown in Figure 5

These results simply confirm the intuition in Figure 3. Namely the welfare gain from the optimal tax instrument, around \$2.5 billion, is much larger than the gain from the optimal permit instrument, around \$0.3 billion. Although the rough calculation using Weitzman's formula suggested a difference of \$10 billion, there are many reasons why Weitzman's result is not exactly right in this context (as discussed in Appendix C).

Figure 5 also indicates a large risk associated with setting an emission target incorrectly – specifically setting a target too low. While the net benefits of a tax are positive for a wide range of values, from zero to \$20/ton C, the net benefits of a permit system rapidly become negative as the target falls below 11 GtC. At the proposed 1990 emission level, 8.5 GtC, the welfare gain is below –\$10 billion. This result is a consequence of reductions becoming extremely expensive in the high-emission states of the world. Thus, an important conclusion from this graph is that

low targets could be prohibitively costly. While adopting a more optimistic view about the cost of reductions might weaken this conclusion, it seems unlikely that the uncertainty surrounding even such optimism would allow a complete reversal.


3.4 Combining Taxes and Permits

Not long after Weitzman's original article, several authors suggested using a combined fee and permit policy (Weitzman 1978; Roberts and Spence 1976) to regulate a market. That is, a producers have a choice of either obtaining a permit or paying a tax-like fee. Such a system works like a permit system by fixing emissions as long as the marginal cost (e.g., the price of the permit) lies below the fee level and works like a tax system by fixing marginal cost when marginal cost rises to the level of the fee. When the fee is set high, such a combined mechanism functions like a pure permit system (since the fee is never paid) and when the number of permits is set low, it functions like pure tax mechanism (since permits are never used).

Even ignoring efficiency concerns, such a mechanism has considerable appeal in the climate change context as pointed out by McKibbin and Wilcoxon (1997). Relative to a pure carbon tax, it effectively exempts payment of taxes on the volume of emissions in 1990. This would lower the cost to industry by over \$10 billion in the U.S. alone.⁷ While it would not necessarily stabilize emissions, the policy would reduce emissions below their baseline and quite possibly be more effective than an aggressive policy which is later abandoned as too costly. The fact that governments would receive some revenue also provides an incentive mechanism for monitoring which is absent in a pure permit system. Finally, as McKibbin and Wilcoxon highlight, it would be relatively easy to add new countries to an international agreement involving such a mechanism. In contrast to a pure permit system involving complex international trades, a permit and fee system could be adopted domestically by non-participating countries without international negotiations. Countries such as India and China, which are unlikely to participate initially in a climate change agreement, would have the additional incentive of fee revenue to enact such a policy in the future.

⁷ As pointed out in other work, losing the revenue from such taxes could also have important welfare consequences (Goulder, Parry, and Burtraw 1996).

Figure 6: Net Expected Policy Benefits of Alternative Tax/Permit Policies in 2010



*The optimal permit level when taxes are not used (13 GtC) does not become apparent until the tax rate is set roughly twice as high as shown in the figure (around \$50/ton carbon), otherwise the tax is still used in some states of nature. From the figure, it is evident that the optimal permit level as a function of the tax is increasing. At \$25/ton carbon, however, the optimal permit level is only 11 GtC. The net expected benefit when the tax is no longer used and the permit level is 13 GtC is \$0.3 billion, roughly one-tenth of the benefit from a straight tax or combination tax/permit policy.

Importantly, the efficiency gains of such a combined mechanism are large relative to a pure permit scheme. As shown in Figure 6, the gains from a combined permit-fee mechanism are roughly the same as a pure tax mechanism, but roughly ten times the gain from a pure permit mechanism. It is interesting that the global optimum – a 5 GtC permit volume and \$7/ton C tax – is quite close to the proposal suggested by McKibbin and Wilcoxon.⁸

4 Dynamic Policy Under Uncertainty

Up to this point the analysis has focused on the costs and benefits of different policies *in a single year*. The problem of climate change, however, is spread out over decades if not centuries. Policies to combat climate change are therefore likely to be in place for many years. It is not immediately obvious whether the results comparing different instruments in a single year are immediately applicable to a multiperiod policy.


In this section, optimal paths for taxes, permits and combined permit-fee mechanisms are explored. In addition to considering the net welfare consequences of these policies, results concerning the range of climate outcomes and compliance costs (ignoring climate benefits) are presented.

4.1 Baseline Emissions

Before exploring different dynamic policies, it is useful to report on the model's ability to replicate the results of other models under business-as-usual assumptions. In particular, the IPCC (1992) gives six forecasts of future emission scenarios for the major greenhouse gases based on a range of assumptions concerning global population and productivity growth. Figure 7 shows the IPCC forecasts against the range produced by the model presented in this paper for the most important greenhouse gas, carbon dioxide.

The IPCC forecasts tend to fall between the 25th and 75th percentile in 2025 and between the 5th and 75th in 2050 and 2100. The median of the IPCC forecasts remains quite close to the 50th

⁸They advocated 1990 emission levels as the permit volume coupled with a \$10/ton C tax. 1990 global controllable GHG emissions were 8.5 GtC.

Figure 7: Simulated CO₂ Emission Distribution vs IPCC Scenarios

*Lines indicate the distribution of CO₂ emission paths generated by the model. These reflect controllable carbon equivalent GHG emissions, scaled by the fraction due to CO₂ (e.g., 86.6%; see p. 71, Nordhaus 1994b). Circles (o) indicate 1992 IPCC CO₂ emission scenarios (p. 12, IPCC 1992; pp. 101-112; Pepper et al. 1992); letters in right margin refer to individual scenarios.

percentile of the simulated emission levels. This suggests that, relative to the IPCC forecasts, the model in this paper predicts a similar central tendency but with a larger spread. In particular, this paper suggests that future uncontrolled emissions could be much higher than all six of the IPCC forecasts. Is this unreasonable?

First, it is important to recognize that the IPCC scenarios do not have a probabilistic interpretation. They are subjectively developed scenarios combining a large number of alternative assumptions in six particular combinations. Their high emission scenarios, for example, alternatively combine high population growth with lower per capita productivity growth, and vice versa. Further, the underlying growth forecasts themselves have no probabilistic interpretation.⁹

Second, even analyses that are well-grounded in probability often underestimate the probability of extreme events. This phenomena has been documented in everything from the measurement of

⁹See Pepper et al. (1992) for further details. Note that early Census population forecasts using similar non-probabilistic techniques often grossly misforecast population. For example, the forecast range given in 1966 (#381, Table 1) lies completely above the actual population reported in 1989 (#1045, U.S. Department of Commerce, Bureau of the Census).

physical constants to the forecast of future energy demand (Shlyakhter and Kammen 1992). Such results suggest that forecasts with "thicker tails" are likely to be a more realistic description of likely outcomes.

Finally, the time horizon under consideration – over one hundred years – makes any forecast based on historical data somewhat dubious. It is implicitly assumed that historical trends will continue with no structural change. When in the history of the world have such periods of stability existed? For all these reasons, the wider range of uncontrolled emission forecasts produced in this paper is likely a more accurate description of emission outcomes than the range of IPCC scenarios.

4.2 Optimal Permit Policy

To compute optimal policies, the paths of alternative permit, tax and combined permit-fee systems were parameterized with six parameters describing stringency in 2010 (the first year of implementation), 2020, 2040, 2070, 2110 and 2160. Policies in intervening years are taken smooth interpolations,¹⁰ except the stringency in 2160 which is allowed to be discontinuous and is held fixed through the end of the simulation (2245). The length of the simulation as well as the spacing of the policy parameters was subjectively chosen to make the policy evaluation in the 2000-2100 interval as accurate as possible, especially the early 2000-2050 period.


The resulting optimized permit policy, which limits global greenhouse gas emission to a specified level, is shown in Figure 8. The policy is not known with certainty due to sampling error – only 8,000 states of nature were used to estimate the policy (this could be reduced with additional computing power).¹¹ Interestingly, the optimal permit level of 13 GtC in 2010 is roughly the same as the optimal permit level determined in the static analysis. This is not so surprising given the fact that initial emission reductions will not substantially affect GHG concentrations for many years, at which point the discounted benefits in 2010 will be small.¹²

¹⁰Policies are interpolated to ten-year intervals using a cubic spline; annual policies are linearly interpolated from the ten-year values. The combined permit-fee system is always interpolated linearly due to sampling error.

¹¹The stringency was optimized over eight sets of 1,000 states of nature, taking on average 30 minutes to converge. These eight sets were then averaged and the standard deviation among the eight estimates was used to compute a standard error for the average.

¹²An important assumption in the underlying DICE model is that damages are continuous and have been occurring

Figure 8: Optimized Permit Level - Dynamic Policy



An important observation is that the proposal to reduce emissions to their 1990 levels (roughly 8.5 GtC) is far below the optimal permit level in these simulations. Further, the optimal permit level rises in the future to accommodate growth in population and per capita productivity. While this kind of policy might appear to ignore the possibility that a carbon-free technology becomes not only commercially available but dominant in the next century, it conversely accommodates the possibility that such a technology never gets off the ground.

When the optimal policy is implemented it improves welfare on average by \$69 billion (\$1989).¹³

4.3 Optimal Tax Policy

While the permit policy requires that emissions in every state of nature be at or below the specified level, a tax policy effectively fixes the marginal cost of emission reductions. Consumers and producers facing a tax on greenhouse gas emissions will reduce emissions until the cost of reducing emissions further equals the cost of paying the tax instead. If costs are particularly low, emissions


since the beginning of industrialization. Thus we are already experiencing the consequences of global warming and current emission reductions are predominately concerned with reducing damages for roughly the next 30 years.

¹³This is the net present value of benefits minus costs and has an associated sampling error of \$16 billion. This can be compared to annual global output in 1995 which was \$24 trillion.

- to the IPCC Scientific Assessment; eds. J.T. Houghton, B.A. Callander and S.K. Varney. Cambridge: Cambridge University Press, 1992.
- Pizer, W. A. (1996). *Modeling Long-Term Policy Under Uncertainty*. Ph. D. thesis, Harvard University.
- Pizer, W. A. (1997). Optimal choice of policy instrument and stringency under uncertainty: The case of climate change. Working Paper #97-17, Resources for the Future.
- Roberts, K. (1980). Interpersonal comparability and social choice theory. *Review of Economic Studies* 47, 421-439.
- Roberts, M. J. and M. Spence (1976). Effluent charges and licenses under uncertainty. *Journal of Environmental Economics and Management* 5, 193-208.
- Shlyakhter, A. I. and D. M. Kammen (1992). Estimating the range of uncertainty in future development from trends in physical constants and predictions of global change. working paper 92-06, Center for Science and International Affairs, John F. Kennedy School of Government, Harvard University.
- Stavins, R. N. (1996). Correlated uncertainty and policy instrument choice. *Journal of Environmental Economics and Management* 30(2), 218-232.
- U.S. Department of Commerce, Bureau of the Census (various issues). *Current Population Reports*. Series P-25, Population estimates and projections.
- Weitzman, M. L. (1974). Prices vs. quantities. *Review of Economic Studies* 41(4), 477-491.
- Weitzman, M. L. (1978). Optimal rewards for economic regulation. *American Economic Review* 68(4), 683-691.

Figure 9: Optimized Tax Level - Dynamic Policy

(dashed lines indicate 95% confidence interval)




may be completely eliminated and the marginal cost of the last ton of reductions will lie below the specified tax rate.

Figure 9 shows the tax policy which maximizes expected welfare. The initial tax of \$7.35 is close to, but slightly lower than, the tax computed in the static analysis. Unlike the optimal permit policy, which rises and relaxes in stringency in order to accommodate growth, the optimal tax policy becomes *more stringent* in the future. This occurs because, unlike a constant permit level, a constant tax automatically leads to proportionally higher emissions as the economy grows. While some increase in emissions is desirable as the economy grows, a proportional increase is not – therefore the tax must increase in stringency while the permit system must be relaxed.

An important difference between tax and permit policies is that taxes do not provide a strict limit on emissions. This, however, should not be viewed strictly as a weakness: relaxing the level of emissions when costs turn out to be particularly high is desirable if the benefits of reduction are fairly constant. Figure 10 shows the distribution of resulting emissions with and without the tax policy. Note that the optimal permit policy is *not* binding over 75% of the time (it lies above the 75th quantile of baseline emissions). Meanwhile the optimal tax policy leads to emissions above

Figure 10: Distribution of Emission Levels Under Alternative Policies




the optimal permit level between 5 and 25% of the time.

Finally, if welfare improvements are averaged across states of nature, the typical welfare improvement amounts to \$338 billion – compared to \$69 for the permit instrument (the standard error in this case is \$21 billion). This represents an improvement five times higher than that obtained with the permit system – a smaller ratio than the factor of eight determined in the static analysis.

Why might permits appear slightly better when viewed over a longer horizon? As pointed out by Weitzman (1974) and discussed later by Stavins (1996), negative correlation between marginal costs and benefits leads to a preference for quantity instruments (e.g., permits). In the case of a stock pollutant which is regulated by an emissions target, such correlation is likely over long periods of time due to uncertainty about baseline emissions. When baseline emissions are consistently low, the pollutant stock is lower in the future making marginal damages lower. At the same time, lower emissions make any single period target cheaper to obtain. This assumes that baseline emissions in one period are correlated with baseline emissions in other periods – a likely consequence when uncertainty about future growth is modeled as “low growth” versus “high growth” scenarios.

Figure 11: Optimal Initial Permit Distribution in a Combined Permit-Fee System



4.4 Combined Permit-Fee Mechanism

An alternative to the pure tax and permit systems discussed so far is a combined system of permits and fees. Such a system involves an initial allocation of permits coupled with a price at which additional permits are sold by the government. The initial distribution involves a quantity mechanism while the subsequent permit sales involve a price mechanism. Since the initial distribution can be set to an arbitrarily small number of permits and the subsequent sale price can be set arbitrarily high, the combined permit-fee system can be made to resemble either the pure tax or permit policy.

To explore this possibility, the subsequent sale price was set to the optimized tax level and welfare was then re-optimized over alternative levels for the initial permit distribution.¹⁴ The resulting optimal permit level is shown in Figure 11. There is a much higher degree of sampling error in the estimation of the optimal permit level. This occurs because the marginal effect of the policy is derived only from those states of nature where baseline emissions are above the permit level but marginal costs are below the tax – this turns out to be a relatively small fraction of

¹⁴A direct optimization over all the policy parameters turned out to be troublesome due to the sample size. By using the optimized tax policy, the combined policy is guaranteed to both raise welfare and involve lower costs.

the sample and results in a large estimation error.¹⁵ Interestingly, the optimal initial permit level remains close to the 1990 emission level (8.5 GtC) for almost the entire forecast period. Setting the initial distribution to 8.5 GtC, in fact, has a negligible effect on the welfare (as was true in the static case). The gain in welfare is \$339 billion for the optimized permit-fee policy and \$340 billion for a policy with an initial permit distribution of 8.5 GtC in every period, both with a standard error of \$20 billion.

4.5 Costs Only Comparison

The purpose of this paper has been to compare alternative climate change policy instruments under a metric of *welfare maximization*. The main conclusion of this study is that the current view of climate damages leads to a relatively flat marginal benefit curve. In the presence of uncertainty, this generates a preference for tax rather than quantity instruments – or even better, a combined permit-fee system.

However, in order to perform this analysis, it has been necessary to lean on rather speculative assessments of climate damages. Even while allowing for considerable uncertainty, the underlying structure of damages – a quadratic function depending on temperature change – has not been questioned in this paper. This may tend to cast some doubt on the welfare calculations presented in support of the alternative policies.

For that reason, Table 2 presents a “cost-only” comparison of the policies derived in the previous sections. That is, the climate portion of the model (specifically, the damages due to global warming) is ignored and only the *cost* of the different reduction strategies is computed. As highlighted in Figure 10, however, it is important to keep in mind that the alternative policies have different emission consequences as well. The interpretation of this measure is that it reflects the expected cost of each policy in the event that there are, in fact, no consequences to global warming.

The permit policy, which has the lowest net benefits, also has the lowest gross cost. Intuitively, the policy simply fails to affect many states of nature since it is only binding 25% of the time. The

¹⁵This sampling error could be reduced with additional simulations. This increased uncertainty can also be observed in Figure 6 as a relatively flat ridge for a wide range of permit levels near the optimal tax.

Table 2: "Cost-only" Comparison of Alternative Policies
(measured in \$1989 billions; standard errors in parentheses)

	optimal ^a permit system	optimal ^a tax system	optimal ^a permit-fee system
benefit-cost ^b	69 (16)	338 (21)	339 (20)
cost-only ^c	-151 (5.5)	-187 (2.7)	-170 (2.7)

^aPolicies are optimal in that they maximize the benefit-cost measure. Large standard errors arise due to the variability of benefits (note that the variation is much smaller in the cost-only measures). Additional simulations can be used to reduce the sampling error further.

^bBenefit-cost measures the change in expected net present value of welfare when damages due to global warming are included in the simulation.

^cCost-only measures the change in expected net present value of welfare when damages are ignored (both in the baseline welfare simulations as well as for the given policies).

tax policy, which has much higher net benefits has 20% higher gross costs. These higher costs arise inevitably in the other 75% of the states of nature, relative to the permit policy, where the tax policy has some effect. Finally, the combined permit-fee system is only marginally better in benefit-cost terms than the pure tax policy, but cuts the gross cost difference in half. This cost savings arises because in those states of nature where permit level is either not binding or marginally binding, less aggressive reductions are undertaken. The fact that the net benefits remain virtually unchanged indicates that this slackness in the policy is occurring in states where marginal benefits are roughly equal to marginal costs.

Perhaps the most important observation from Table 2 is that while the benefit-cost comparison shows taxes and permit-fee systems to be much more efficient than a pure permit mechanism (a factor of five-to-one) the difference in costs is much smaller – at most 25%. This suggests that concern about a tax system turning out to be much more expensive than a quota or permit system is misplaced.¹⁶

¹⁶This statement, of course, is in reference to the cost difference among efficiently implemented policies in a first-best world. Issues related to cost-effectiveness and second-best distortions are not considered.

5 Conclusion

Discussions of alternative tax and permit mechanisms for combating climate change have generally ignored the fact that the costs and benefits of future reductions are highly uncertainty. Such uncertainty can lead to large efficiency differences between the two policies (Weitzman 1974). This paper has explored this question in the context of an integrated climate-economy model capable of simulating thousands of uncertain states of nature.

The resulting welfare analysis indicates that taxes perform much more efficiently than permits – by a factor of *five to one* (\$337 billion versus \$69 billion). This derives from the relatively flat marginal benefit curve associated with emission reductions. This flatness is partially a product of the assumption of quadratic damages, even allowing for the fact that the level of damages due to three degrees of warming is anywhere from zero to a three percent loss of GDP. It also arises naturally because greenhouse gases are stock pollutants so reductions in any period inevitably have little effect on the level of marginal benefits.

An important observation in the static analysis was the risk involved in setting the permit level *too low*. Not only does the optimal permit involve much lower welfare gains, but setting the permit level incorrectly can lead to massive losses. The tax instrument, in contrast, leads to welfare gains over a much wider range of values.

In addition to pure tax and permit systems, a combined permit-fee system was considered. Such a system involves an initial allocation of permits followed by the subsequent sale of additional permits at fixed price. The initial distribution function like a quota while the subsequent permit sales function like a price instrument. By making the initial distribution small or the subsequent sale price high, this combined system can be made to mimic either the pure tax or permit system.

The combined permit-fee system improves on the welfare outcome using taxes but only slightly (\$339 billion vs. \$337 billion). Interestingly, the optimal distribution of initial permits turns out to be insignificantly different than the proposed 1990 emission level target of 8.5 GtC. This mechanism thus combines the benefits of a tax system, namely efficiency, with the benefits of a permit system, e.g., flexible distribution of rents.

A Model Specification

A.1 Economic Behavior

Economic behavior within each state is derived from a representative agent model where consumption must be optimally allocated across time. In a typical model with constant exogenous productivity growth, agent preferences define a steady state to which the economy converges over time. In the presence of random shocks and slowly changing trends, the economy instead converges to a distribution of states (due to the random shocks) which is itself slowly evolving (due to the slowly changing trends). For the moment we ignore these changing trends and focus on a standard stochastic growth model.

The representative consumer in this model exhibits constant relative risk aversion τ with respect to consumption per capita. Utility is separable across time, discounted at rate ρ and weighted each period by population. With the further assumption that preferences satisfy the von Neuman-Morgenstern axioms, the consumer's optimization problem can be written as

$$\max_{\{C_t\}} E \left[\sum_{t=0}^{\infty} (1 + \rho)^{-t} N_t \frac{(C_t/N_t)^{1-\tau}}{1-\tau} \right] \quad (\text{A.1})$$

where C_t is consumption in period t and N_t is population. That is, the consumer maximizes expected discounted utility where each period's utility is population weighted. This consumption program $\{C_0, C_1, \dots\}$ is subject to the resource constraints describing production

$$Y_t = (A_t^* N_t)^{1-\theta} K_t^\theta \quad (\text{A.2})$$

and capital accumulation

$$K_{t+1} = K_t(1 - \delta_k) + Y_t - C_t \quad (\text{A.3})$$

where Y_t is aggregate output, $A_t^* N_t$ is effective labor input and K_t is the capital stock. A_t^* is a measure of productivity distinct from capital but not completely exogenous, as discussed later. The parameter θ summarizes the Cobb-Douglas production technology given in Equation (A.2)

and δ_k reflects the rate of capital depreciation in the capital accumulation equation (A.3). Finally, there is a transversality condition for a balanced growth steady state:

$$\rho > (1 - \tau) \times (\text{asymptotic growth rate of } A_t^*) \quad (\text{A.4})$$

This is always satisfied by assuming zero growth asymptotically.

Even with exogenous constant growth models for N_t and A_t^* , the dynamic optimization problem given by Equations (A.1–A.3) is difficult if not impossible to solve analytically.¹⁷ However, choosing

$$\Delta \ln(K_{t+1}/N_{t+1}) = \alpha_1 + \alpha_2(\ln(K_t/N_t) - \ln(A_t^*)) \quad (\text{A.5})$$

and

$$C_t = K_t(1 - \delta_k) + Y_t - K_{t+1} \quad (\text{A.6})$$

— where α_1 and α_2 are functions of the parameters $(\rho, \tau, \theta, \delta_k)$ — yields a close approximation of optimal consumer behavior around the balanced growth steady state. This technique of approximating optimal dynamic behavior has its origins in the real business cycle literature beginning with Kydland and Prescott (1982). It is also related to the technique of feature extraction discussed by Bertsekas (1995).¹⁸

Intuitively, Equation (A.5) approximates behavior around a balanced growth steady state. At such a steady state, $\ln(K_t/N_t) - \ln(A_t^*)$ is constant and $\Delta \ln(K_{t+1}/N_{t+1}) = (\text{growth rate of } A_t^*) = \alpha_1 + \alpha_2(\ln(K_t/N_t) - \ln(A_t^*)) = \text{constant}$. If some unforeseen shock moves the economy away from the equilibrium value of $K_t/(A_t^*N_t)$ and α_2 is negative, e.g., the steady state is stable, then the economy will move back toward the steady state. In particular, when $K_t/(A_t^*N_t)$ is too high, capital accumulation will slow. If $K_t/(A_t^*N_t)$ is too low, capital accumulation will increase. Importantly, even if the growth rate of A_t^* is not constant, this approximation performs well as long as expected productivity growth changes gradually.

¹⁷Long and Plosser (1983) derive an analytic solution for the case of $\delta_k = 1$ and $\lim \tau \rightarrow 1$ (log utility).

¹⁸See Appendix A of Pizer (1997) for a simple derivation of expressions for α_1 and α_2 .

A.2 Long-term Growth, Climate Behavior and Damages

This section explains the remainder of the state-contingent model – specifically the evolution of A_t^* and N_t . This includes exogenous growth projections, climate behavior and damages from global warming (based primarily on the DICE model, Nordhaus 1994b). Exogenous labor productivity A_t is modeled as a random walk in logarithms with an exponentially decaying drift. That is,

$$\log(A_t) = \log(A_{t-1}) + \gamma_a \exp(-\delta_a t) + \sigma_a \epsilon_t \quad (\text{A.7})$$

where γ_a is the initial growth rate, δ_a is the annual decline in the growth rate, σ_a is the standard deviation of the random growth shocks and ϵ_t is a standard NIID random shock. This means that productivity growth begins with a mean growth rate of γ_a (around 1.3%) in the first period and eventually declines to zero. In addition, random and permanent shocks change the level of productivity every period. The standard error of these shocks is σ_a .

Net labor productivity A_t^* is distinguished from this exogenous measure A_t by the fact that A_t^* describes the amount of output available for consumption and investment – after output has been reduced by control costs and climate damages. To that end, A_t^* is expressed as A_t multiplied by a factors describing these two phenomena:

$$A_t^* = \underbrace{\left(\frac{1 - b_1 \mu_t^{b_2}}{1 + (D_0/9) \cdot T_t^2} \right)^{\frac{1}{1-\theta}}}_{\text{damages}} A_t \quad (\text{A.8})$$

μ_t is the fractional reduction in greenhouse gas emissions at time t (the “control rate”) versus a business as usual/no government policy baseline, while b_1 and b_2 parameterize the cost of attaining these reductions. Since b_1 and b_2 are both positive, additional rates of control involve reductions in net productivity. T_t is the average surface temperature relative to pre-industrialization in degrees Celsius and D_0 is the fractional loss in aggregate GDP from a 3° temperature increase. For temperature changes less than 10°, this is essentially a quadratic damage function.¹⁹ Additional details about the control cost and damage functions can be found in Nordhaus (1993) and Nordhaus (1994a), respectively.

¹⁹Over larger ranges, the damage function becomes S-shaped.

Population is modeled in the same way as exogenous productivity but without the stochastic element:

$$\log(N_t) = \log(N_{t-1}) + \gamma_n \exp(-\delta_n t) \quad (\text{A.9})$$

where γ_n is the initial growth rate and δ_n is the annual decline in the growth rate. Note that these models predict zero growth asymptotically, though this may occur centuries in the future.²⁰

The remaining portion of the model explains the link between economic activity (measured as aggregate output Y_t) and warming (measured as the average surface temperature T_t). The first step is linking output to emissions:

$$E_t = \sigma_t(1 - \mu_t)Y_t \left(\frac{A_t}{A_1}\right)^{1-\theta} \quad (\text{A.10})$$

where E_t is emission of controllable greenhouse gases,²¹ σ_t is an exogenous trend in emissions/output, and μ_t is the rate of emissions reductions induced by the policymaker. The expression $Y_t \left(\frac{A_t}{A_1}\right)^{1-\theta}$ reflects raw output prior to the effects of climate damages and control costs. The model of σ_t is, as with labor productivity and population, based on exponentially decaying growth:

$$\log(\sigma_t) = \log(\sigma_{t-1}) + \gamma_\sigma \exp(-\delta_\sigma t) \quad (\text{A.11})$$

where γ_σ is the initial growth rate of emissions/output (a negative number) and δ_σ is the annual decline in the growth rate. Note that the annual decline in the emissions/output growth rate is the same as the annual decline in labor productivity growth (δ_a).

Emissions of greenhouse gases accumulate in the atmosphere according to:

$$M_t - 590 = \beta E_{t-1} + (1 - \delta_m)(M_{t-1} - 590) \quad (\text{A.12})$$

where M_t is the atmospheric concentration of greenhouse gases in billions of tons of carbon equivalent. β is a measure of the retention rate of emissions. Low values of β indicate that emissions do

²⁰For example, the range of parameters used in the simulations (with $\delta_{a/n} \in (0.25\%, 2.5\%)$) leads to a halving of the growth rates every 20 to 200 years.

²¹See discussion of controllable versus uncontrollable greenhouse gases in Nordhaus (1994b), page 74. For the most part, controllable greenhouse gases are CO₂ and CFCs and uncontrollable greenhouse gases are everything else.

not, in fact, accumulate while a value of unity would mean that every ton of emitted greenhouse gases remains in the atmosphere. The parameter δ_m plays the role of a depreciation rate: it is assumed that greenhouse gases in the atmosphere above the pre-industrialization level of 590 billion tons slowly decays. This decay reflects absorption of greenhouse gases into the oceans which are assumed to be an infinite sink.

Above average concentrations of greenhouse gases in the atmosphere lead to increased radiative forcings, a measure of the rate of transfer between solar energy produced by the sun and thermal energy stored in the atmosphere. This is modeled according to

$$F_t = 4.1 \times \log(M_t/590)/\log(2) + O_t \quad (\text{A.13})$$

where F_t measures radiative forcings in units of watts per meter squared. The specification is such that a doubling of greenhouse gas concentrations leads to a roughly four fold increase in forcings (since 590 is the concentration before industrialization). O_t in this relation represents radiative forcings due to other uncontrollable greenhouse gases and is assumed exogenous to the model:

$$O_t = \begin{cases} 0.2604 + 0.0125t - 0.000034t^2 & \text{if } t < 150 \\ 1.42 & \text{otherwise} \end{cases} \quad (\text{A.14})$$

Increased forcings lead to temperature changes according to

$$T_t = T_{t-1} + (1/R_1) [F_t - \lambda T_{t-1} - (R_2/\tau_{12})(T_{t-1} - T_{t-1}^*)] \quad (\text{A.15})$$

$$T_t^* = T_{t-1}^* + (1/R_2)(R_2/\tau_{12})(T_{t-1} - T_{t-1}^*) \quad (\text{A.16})$$

where T_t is the surface temperature and T_t^* is the deep ocean temperature, both expressed in changes relative to pre-industrialization levels in degrees Celsius. Note that if $M_t = 590$ and $O_t = 0$ (e.g., pre-industrialization), T_t and T_t^* will equilibrate to zero. The parameter λ describes the equilibrium change in surface temperature for a given change in radiative forcings. In particular, based on (A.13) and (A.15), a doubling of the concentration of greenhouse gases in the atmosphere will lead to a $4.1/\lambda$ rise in surface temperature in the long run. This parameter $4.1/\lambda$ is a measure of the temperature sensitivity of the atmosphere.²²

²² λ by itself is referred to as the climate feedback parameter.

The parameters R_1 , R_2 and τ_{12} describe the thermal capacity of the surface atmosphere and deep oceans and the rate of energy transfer between them, respectively.

A.3 Social Welfare

A distinguishing feature of this analysis is the use of an econometrically estimated parameter distribution describing uncertainty in the economic model. However, the consumer's objective function given by Equation (A.1) makes no allowance for uncertainty about the preference parameters ρ and τ which are fixed from his or her perspective. In order to encompass uncertainty about preferences, it is necessary to step back and imagine a social planner who would like to maximize the objective given in (A.1) but is unsure of the parameters. Since a policy change which raises the expected utility for one set of parameters may lower the expected utility for another set, the social planner will need to specify a social welfare function to compare gains and losses across states of nature. This social welfare function provides a single objective specifying how changes in utility measured with different preferences are aggregated.²³ It is important to recognize that although parameter values in the representative agent model can be inferred from observed consumer behavior, there is no information available to estimate parameters in a social welfare function. Such information would be revealed only by observing the behavior of an actual social planner. Instead, we must rely on social choice theory and our own sense of fairness to specify the relation.

It is useful to note that the common approach in the climate change literature skirts this issue of preference aggregation by reporting a range of policy prescriptions based on a range of possible preferences and states of nature. For example, Cline (1992) presents benefit-cost analyses for 92 different cases (Tables 7.3 and 7.4). Dowlatabadi and Morgan (1993) integrate out much of the uncertainty in their analysis, but still present results for 48 scenarios. Chapter 8 of Nordhaus (1994b) gives one of the few examples where even preference uncertainty is integrated out, yielding a single welfare metric and a single policy recommendation. In a similar analysis, however, Nordhaus and Popp (1997) choose to fix preferences because of the difficulties with preference aggregation.

²³E.g., providing a negative loss function across states of nature for the social planner.

Regardless, these authors ubiquitously observe that uncertainty about time preference has large consequences for optimal policy choice.²⁴ Moreso, in fact, than uncertainty about climate sensitivity and damages. It therefore behooves us to seriously consider how to aggregate over uncertain preferences in the most reasonable way.

In this analysis social welfare is specified as an average of utility measured in each state of nature by Equation (A.1), *rescaled*. The rescaling serves to equate the marginal social welfare of one additional dollar of current consumption in all states of nature. While arbitrary, some adjustment is necessary to prevent the resulting policy prescription from being sensitive to the choice of units in the model.²⁵ Social welfare can then be written as

$$SW(x) = I^{-1} \sum_{i=1}^I u(x, i) \quad (A.17)$$

where $u(x, i)$ is rescaled utility in state i with outcome x and $SW(x)$ is the social welfare associated with x . The rescaling is such that

$$u(+\$1 \text{ in initial period}, i) - u(\emptyset, i) = u(+\$1 \text{ in initial period}, j) - u(\emptyset, j) \quad \forall i, j$$

That is, a policy corresponding to an extra dollar of consumption in the initial period is assumed to have the same utility gain in every state relative to a no policy ($\emptyset$) baseline.

This social welfare function has its origins in the literature on social choice. Harsanyi (1977) shows that in defining social welfare over lotteries, if individual preferences satisfy the von Neumann Morgenstern axioms then social welfare must have this weighted average form. Otherwise, social preferences will fail to mimic individual preferences over lotteries involving only that individual. This functional form can also be derived from the assumption of cardinal unit comparability, as discussed by Roberts (1980). More flexible forms require additional assumptions about level or scale comparability. Our choice of welfare functions is therefore less arbitrary than it might have originally appeared: a more flexible form requires both integrating out uncertainty from the representative agent's perspective (to satisfy Harsanyi's point) and more stringent assumptions about

²⁴See discussion in Arrow, Clive, Maler, Munasinghe, and Stiglitz (1996).

²⁵An explanation of this point is given in Appendix C of Pizer (1997).

the level of comparability (to satisfy Robert's point)²⁶

B Measuring Uncertainty

Estimates of uncertainty in the model come from two sources: econometric analysis and subjective assessment. The model involves nineteen different parameters. Six are parameters describing observable economic activity:

- pure time preference ρ ,
- risk aversion τ ,
- output-capital elasticity θ ,
- productivity growth γ_a ,
- variation in productivity growth σ_a^2 , and
- depreciation δ_k .

A joint distribution for these parameters is estimated with historical data. The remaining thirteen describe emissions:

- emissions rate growth γ_e ,

climate change:

- CO₂ retention rate β ,
- temperature sensitivity $\Delta T/\lambda$,
- CO₂ decay rate δ_m ,
- thermal capacities and conductivities R_1 , R_2 and τ_{12} ,

control costs and damages:

- cost function parameters b_1 and b_2 ,
- fractional loss of GDP for 3° temperature rise D_0 ,

and long-term growth trends

- population growth γ_n ,
- productivity slowdown and slowdown in the growth rate of emissions/output δ_a ,
- population slowdown δ_n .

Uncertainty about these parameters is based on subjective analysis.

The econometric analysis of the six economic parameters is based on post-war U.S. data.²⁷

²⁶Additional levels of comparability are especially difficult with the the constant coefficient of relative aversion (CRRA) form in Equation (A.1) where the parameter $\tau \geq 1$. Under these assumptions, utility is alternatively bounded from above or below.

²⁷See Chapter II of Pizer (1996).

Series describing aggregate investment, capital services, output and prices are fit to the model described by Equations (A.2),(A.3) and (A.5). The posterior parameter distribution which arises from this analysis is summarized in Table B.1.

Nordhaus (1994b)²⁸ develops a distribution for the remaining parameters based on a two-step subjective analysis. The first step involves testing his model's sensitivity to each parameter being changed, one at a time, to a more extreme value. Those parameters which produce the largest variance in model output are then further scrutinized. A discrete, five-value distribution is developed for seven of these thirteen variables. The other six are fixed at their best guess values. The distribution of the seven uncertain parameters is summarized in Table B.2. Values of the six fixed parameters as well as initial conditions for the model are given in Table B.3.

C Factors Which Complicate the Weitzman Analysis

In this section the factors which complicate Weitzman's (1974) original analysis are discussed. While these factors do not affect the basic intuition behind his result – that a flatter marginal benefit curve favors taxes – they do qualify it. In particular, the intersection of the expected marginal benefit and expected marginal cost curve can no longer be used to determine the optimal tax policy when there are non-linear marginal costs and non-additive shocks. As noted by Weitzman and others (Stavins 1996), correlation also changes the optimality result, with positive correlation among costs and benefits favoring permits and negative correlation favoring taxes. Other factors such as truncation and discounting further complicate the simple graphical analysis in Section 3.

C.1 Non-linearities and Non-additive Shocks

A simple graphical comparison of expected marginal benefits and expected marginal costs allows one to quickly derive the optimal quantity control. This works because these marginal measures are with respect to the quantity measure and expectations are taken holding the quantity fixed.

²⁸ Chapters 6 and 7.

Table B.1: Marginal distributions of uncertain economic parameters
 (narrow bars indicate values used in simulations without uncertainty)







description	symbol	equation	distribution
pure rate of time preference	ρ	(A.1)	
coefficient of risk aversion	τ	(A.1)	
output-capital elasticity	θ	(A.2)	
rate of capital depreciation	δ_k	(A.3)	
initial productivity growth rate	γ_a	(A.7)	
standard error of productivity shocks	σ_a	(A.7)	

Table B.2: Discrete distributions of uncertain climate/trend parameters
(narrow bars indicate values used in simulations without uncertainty)


description	symbol	equation	distribution
annual decline of population growth rate	δ_n	(A.9)	
annual decline of productivity growth rate	δ_a	(A.7),(A.11)	
initial growth rate of CO2 per unit output	γ_o	(A.10),(A.11)	
damage parameter (% loss of GDP for 3° temperature rise)	D_0	(A.8)	
cost function parameter	b_1	(A.8)	
retention rate for CO2 emissions	β	(A.12)	
temperature sensitivity to CO2 doubling (in °C)	$4.1/\lambda$	(A.13),(A.15)	

Table B.3: Description of fixed parameters^a

parameter	symbol	equation	units	value
cost function parameter	b_2	(A.8)		2.887
decay rate of atmospheric CO ₂	δ_m	(A.12)		0.00833
1/thermal capacity of atmosphere	$1/R_1$	(A.15)	°C-meter ² /watt-year	0.048
thermal conductivity b/w atmosphere and oceans	R_2/T_0	(A.15),(A.16)	watt/°C-meter ²	0.44
1/thermal capacity of deep oceans	$1/R_2$	(A.16)	°C-meter ² /watt-year	$\frac{0.002}{0.44}$
1995 global population	N_0	(A.9)	millions of people	5590
initial population growth	γ_n	(A.9)		0.0124
initial rate of CO ₂ emissions per unit of output	σ_0	(A.11)	billion tons CO ₂ per \$1989 trillions	0.385
1995 global capital stock	K_0	(A.2),(A.3)	\$1989 trillions	79.5
1995 global output	Y_0	(A.2),(A.3),(A.10)	\$1989 trillions	24.0
1995 atmospheric concentrations of CO ₂	M_0	(A.12)	billions of tons of C equivalent	763.6
1995 surface temperature ^b	T_0	(A.8),(A.15),(A.16)	°Celsius	0.763
1995 deep ocean temperature ^b	T_0^*	(A.15),(A.16)	°Celsius	0.117

^aAll fixed parameters are from Nordhaus (1994b). The parameters that do not depend on time are from Nordhaus' Table 2.4. Initial values for temperature, CO₂ concentrations, and output in 1995, as well as the initial annual growth rate for population, are based on the Nordhaus base case simulation. The 1995 capital stock is adjusted upward to reflect differences in the definition of capital as well as underlying parameter values. The decay rate of atmospheric CO₂ is divided by ten to convert from a decennial to annual rate. The annual thermal capacity of the ocean and atmosphere are from the second line of Nordhaus' Table 3.4b.

^bTemperatures are measured as deviations from the pre-industrialization level, circa 1900.

Figure C.1: Distribution of $\frac{dQ}{d(\text{tax})}$ 

In contrast, the optimal price control must be derived by taking the expectation of $\Delta \text{benefit} / \Delta \text{tax}$ and $\Delta \text{cost} / \Delta \text{tax}$, holding the tax level fixed, and finding the tax level where *these* two marginal measures are equal. This is not generally revealed by a simple diagram such as Figure 3. It is revealed, however, if the slope of the marginal cost schedule is constant and known. In that case, the condition for the optimal tax reduces to the intersection of the expected marginal cost and benefit curves:

$$\begin{aligned}
 E \left[\frac{\Delta \text{benefit}}{\Delta \text{tax}} \right] &= E \left[\frac{\Delta \text{cost}}{\Delta \text{tax}} \right] \\
 E \left[\frac{\Delta \text{benefit}}{\Delta \text{quantity}} \frac{\Delta \text{quantity}}{\Delta \text{tax}} \right] &= E \left[\frac{\Delta \text{cost}}{\Delta \text{quantity}} \frac{\Delta \text{quantity}}{\Delta \text{tax}} \right] \\
 E \left[\frac{\Delta \text{benefit}}{\Delta \text{quantity}} \frac{1}{\text{MC slope}} \right] &= E \left[\frac{\Delta \text{cost}}{\Delta \text{quantity}} \frac{1}{\text{MC slope}} \right] \\
 E \left[\frac{\Delta \text{benefit}}{\Delta \text{quantity}} \right] \frac{1}{\text{MC slope}} &= E \left[\frac{\Delta \text{cost}}{\Delta \text{quantity}} \right] \frac{1}{\text{MC slope}} \\
 E \left[\frac{\Delta \text{benefit}}{\Delta \text{quantity}} \right] &= E \left[\frac{\Delta \text{cost}}{\Delta \text{quantity}} \right] \\
 E[MB] &= E[MC]
 \end{aligned} \tag{C.18}$$

This condition, that the slope of the marginal cost curve is constant, will be violated if the cost curve is non-linear or there are non-additive shocks. Under these conditions, the slope of

the marginal cost curve cannot be factored outside the expectation, as shown in (C.18). To verify this condition, the slope of the marginal cost curve can be examined across states of nature and at different tax levels to see if it remains constant. Figure C.1 shows the distribution of slopes for a single tax level of \$5/tC.²⁹

Given these violations, it will not be possible to determine the optimal tax level from Figure 3 precisely, though it may still provide a rough approximation.


C.2 Correlation of Cost/Benefit Shocks

C.3 Truncation

An important point ignored by the Weitzman analysis is whether a particular quantity control is binding in every state of nature. Figure 4, for example, reveals that a quota of 12 GtC in 2010 would lie above the actual emission level roughly half the time. In these states of the world, it is inappropriate to count any benefits from the policy since the policy has no consequence. That is, in order to properly compare the expected marginal cost of a 12 GtC permit policy to its benefits, marginal benefits should only be counted when the policy is binding.

Figure C.2 shows the consequence of this calculation. The left panel (the same as Figure 2) shows the distribution of marginal benefits when emissions are fixed at – rather than limited to – different levels. This calculation ignores the actual level of emissions in 2010. In contrast, the right panel shows the marginal benefit associated with limiting emissions to the specified level (as would occur with a permit system). In those states of nature where uncontrolled emissions are below the value on the x-axis, they remain unchanged and there is no marginal benefit. Therefore at high permit levels (18 GtC) that lie above uncontrolled emission levels in every state, the marginal benefit is zero. Moving to the right, marginal benefits initially rise as additional states become contributing benefits. Eventually, the marginal benefit schedule must slope downwards as the marginal benefit in every state is declining.


²⁹From Figure 3 it is clear that the slope also changes at higher permit/tax levels.

Figure C.2: Distribution of Marginal Benefits in 2010^a

^aThese marginal benefits are based on the discounted (to 2010) value associated with emission reductions in the year 2010 only. Emissions in other periods are held at their uncontrolled levels. Values are expressed in \$2010 and, due to different discount rates across states of nature, will not be weighted equally when balancing costs and benefits.

^aThe left panel indicates the range of marginal benefits obtained by varying the level of 2010 emissions as shown, ignoring the forecast level of uncontrolled emissions. The right panel indicates the range of marginal benefits when permits are used to control the level of emissions at or below the permit level. That is, in cases where uncontrolled emissions are below the indicated level, there is no marginal benefit since emissions are not, in fact, being reduced. As shown in Figure 4, uncontrolled emissions in 2010 are unlikely to be more than 18 GtC. The marginal benefit when the permit level is set to 18 GtC is therefore zero.

Figure C.3: Distribution of Discount Factors in 2010



C.4 Discounting

Discounting is an important issue which is ignored in the simple static case. Even when considering policy in a single period (2010), it is necessary to wonder whether costs and benefits should be measured in 2010 or 1995. This is relevant because different states of nature will involve different discount rates. Specifically, those states with higher growth involve more discounting relative to states of nature with low growth. This follows the basic intuition that extra dollars are more valuable when one is poor than when one is rich.

The static analysis in Section 3 is based on dollar welfare measures in 2010. In contrast, the dynamic analysis in Section 4 is based on dollar welfare measures in 1995. Surprisingly, there is little difference in the initial policy outcome, suggesting the discrepancy is small. Figure C.3 shows the range of discount factors observed in 2010 relative to 1995.

References


- Arrow, K. J., W. R. Cline, K.-G. Maler, M. Munasinghe, and J. E. Stiglitz (1996). Intertemporal equity and discounting. In J. Bruce, H. Lee, and E. Haltes (Eds.), *Climate Change 1995: Economic and Social Dimensions of Climate Change*. Cambridge: Cambridge University Press.
- Bertsekas, D. P. (1995). *Dynamic Programming and Optimal Control*, Volume 1. Belmont, MA: Athena Scientific.
- Cline, W. (1992). *The Economics of Global Warming*. Washington: Institute of International Economics.
- Dowlatabadi, H. and M. Morgan (1993). A model framework for integrated studies of the climate problem. *Energy Policy* 21(3), 209-221.
- Goulder, L., I. W. Parry, and D. Burtraw (1996). Revenue-raising vs. other approaches to environmental protection: The critical significance of pre-existing tax distortions. NBER Working Paper No. 5641.
- Harsanyi, J. C. (1977). *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*. New York: Cambridge University Press.
- Intergovernmental Panel on Climate Change (1992). The 1992 IPCC supplement: Scientific assessment. In J. Houghton, B. Callander, and S. Varney (Eds.), *Climate Change 1992: The Supplemental Report to the IPCC Scientific Assessment*, pp. 1-22. Cambridge University Press.
- Kydland, F. E. and E. C. Prescott (1982). Time to build and aggregate fluctuations. *Econometrica* 50(6), 1345-1370.
- Lashof, D. A. and D. R. Ahuja (1990). Relative contributions of greenhouse gas emissions to global warming. *Nature* 344, 529-531.
- Long, J. B. and C. I. Plosser (1983). Real business cycles. *Journal of Political Economy* 91, 39-69.
- McKibbin, W. J. and P. J. Wilcoxon (1997). A better way to slow global climate change. Brookings Policy Brief No. 17, the Brookings Institution.
- Nordhaus, W. D. (1993). The cost of slowing climate change. *Energy Journal* 12(1), 37-65.
- Nordhaus, W. D. (1994a). Expert opinion on climate change. *American Scientist* 82(1), 45-51.
- Nordhaus, W. D. (1994b). *Managing the Global Commons*. Cambridge: MIT Press.
- Nordhaus, W. D. and D. Popp (1997). What is the value of scientific knowledge? an application to global warming using the PRICE model. *Energy Journal* 18(1), 1-45.
- Nordhaus, W. D. and G. W. Yohe (1983). Future paths of energy and carbon dioxide emissions. In N. R. Council (Ed.), *Changing Climate*, pp. 87-152. Washington: National Academy Press.
- Pepper, W., J. Leggett, R. Swart, J. Wasson, J. Edmonds, and I. Mintzer (1992). Emission scenarios for the IPCC: An update; assumptions, methodology and results. Detailed documentation in support of Chapter A3 of *Climate Change 1992: The Supplementary Report*

DISCUSSION PAPER

Optimal Choice of Policy Instrument and Stringency under Uncertainty: The Case of Climate Change

William A. Pizer

Discussion Paper 97-17



RESOURCES
FOR THE FUTURE

**Optimal Choice of Policy Instrument
and Stringency under Uncertainty:
The Case of Climate Change**

William A. Pizer

Discussion Paper 97-17

December 1996

Resources for the Future
1616 P Street, NW
Washington, DC 20036
Telephone 202-328-5000
Fax 202-939-3460

© 1996 Resources for the Future. All rights reserved.
No portion of this paper may be reproduced without permission
of the author.

Discussion papers are research materials circulated by their
authors for purposes of information and discussion. They have
not undergone formal peer review or the editorial treatment
accorded RFF books and other publications.

Optimal Choice of Policy Instrument and Stringency under Uncertainty: The Case of Climate Change

William A. Pizer

Abstract

The relevance of uncertainty in the climate change policy debate is without doubt. Surprisingly, there have been few attempts to examine the direct policy consequences of including uncertainty in an integrated climate-economy framework. This paper presents results concerning optimal policy stringency and instrument choice when economic and climate parameters assume distributions rather than single values. Uncertainty is found to raise the optimal level of emission reductions relative to an optimization based on one set of central parameter estimates. Much of this effect can be related to economic rather than climate uncertainty. Uncertainty also leads to a preference for taxes over quantity controls.

Previous studies of uncertainty in the climate change context have used a small number of states to measure the value of earlier information, learning and adaptation. This paper attempts to refocus attention on the more basic question of whether, in the absence of new information and learning, the inclusion of uncertainty yields significantly different policy conclusions. For policymakers confronting the problem of climate change today, this is the more relevant question.

Key Words: uncertainty, climate change, taxes and quantity controls

JEL Classification No(s): C15, D81, Q28

Table of Contents

1. Introduction	1
2. Model	4
2.1 Overview	4
2.2 State-contingent Model	6
2.2.1 Economic Structure	6
2.2.2 Long-term Growth, Climate Behavior and Damages	7
2.2.3 Summary of State-contingent Climate-Economy Model	10
2.3 Quantifying Uncertainty	11
2.4 Social Welfare	12
3. Does Uncertainty Matter?	17
3.1 Policy Stringency	17
3.2 Welfare	19
4. Instrument Choice	20
4.1 Taxes versus rate controls	21
4.2 Value of emission rights	22
5. Conclusion	24
Appendix	25
A. Notes on Solution Algorithm	25
B. Notes on Econometrics	27
C. Notes on Utility Rescaling	29
References	30

List of Figures and Tables

Figure 1. State-contingent model	10
Figure 2. Optimal climate change policy with and without uncertainty	18
Figure 3. Decomposition of policy consequences	19
Figure 4. Intuition for taxes being preferred to rate instruments	23
Figure 5. Value of emission rights under alternative policies	24
Table 1. Marginal distributions of uncertain economic parameters	13
Table 2. Discrete distributions of uncertain climate / trend parameters	14
Table 3. Description of fixed parameters	15

Optimal Choice of Policy Instrument and Stringency Under Uncertainty: The Case of Climate Change

William A. Pizer¹

1 Introduction

Uncertainty is a pervasive feature of climate change analysis. The wide range of divergent opinions on the proper policy response to climate change is testament to the scale of the issue. However, there is a difference between recognizing uncertainty and using it to propose *one* policy response.

This paper acknowledges uncertainty and optimizes over it to recommend a single course of action. This permits a useful assessment of the importance of including uncertainty in the analysis: Does it make a difference versus optimizing with a single set of central parameter values? Do alternative policy instruments behave differently? Why?

Previous work in this vein has tended to miss these two basic questions. Frequently, the importance of uncertainty is framed in terms of learning: what happens if uncertainty is resolved now rather than later [Manne and Richels (1993,1995), Nordhaus (1994b), Manne (1996)]. This skirts the original question by failing to address how, without learning, uncertainty affects policy answers in the first place. The typical use of a highly stylized framework with only two or three states of nature also diminishes one's confidence that these models are capturing important features of the unknown parameter space. Simply considering high and low values for growth, climate sensitivity, damages and costs requires more than a dozen states of nature. Existing work which does consider a large space of uncertainty does so without a general equilibrium model of intertemporal prices [Dowlatabadi and Morgan (1993)]. This turns out to be inadequate for determining optimal policy and, by extension,

¹Fellow, Quality of the Environment Division, Resources for the Future. The author would like to thank Dale Jorgenson, Gary Chamberlain and seminar participants at Harvard University and Resources for the Future for helpful comments on earlier drafts of this paper.

for addressing uncertain preferences. Preferences, it turns out, are extremely important.⁴

In this paper a relatively simple integrated climate-economy model based on optimal intertemporal behavior is used to examine alternative policies under uncertainty. Several features distinguish this model from other climate-economy models. First, a fast simulation technique is used to permit thousands of different states of nature to be simulated simultaneously, each capturing optimal consumer behavior. Second, the distribution of the economic states of nature is based on econometrically estimated parameter distributions. Third, these states – each of which exhibits different consumer preferences – are consistently aggregated using an explicit social welfare function.

Each of these features is important given the question to be addressed. Without a general equilibrium framework, intertemporal prices are likely to be wrong. Higher productivity growth, for example, leads to more output and emissions but simultaneously higher interest rates. Ignoring the higher interest rates, or timing them incorrectly, will lead to incorrect intertemporal valuation and biased results.² Econometric estimates of economic parameter distributions are important to capture well-known parameter relations and their associated uncertainty. In order to be consistent with historic interest rates, for example, time preference and intertemporal substitutability must be positively correlated. By how much and with what degree of certainty is a question best addressed by data. Finally, given the extremely long horizons involved in climate change analysis, it is not surprising that the results are extremely sensitive to variation in preferences across time. Without a consistent method of aggregating over these preferences, optimal policy responses will depend quite heavily on which rate of pure time preference is assumed – without any assessment of how bad an incorrect guess will be.

The short answer to the question posed is that uncertainty does matter. First, ignoring it leads to policy recommendations which are too slack with welfare gains 30% below that of the optimal policy incorporating uncertainty. Second, a large part of the consequences – around

²An error in interest rates will bias the valuation since policy costs typically occur in earlier years than policy benefits.

fifty percent – arises solely from modeling uncertain preferences (versus uncertain technology or climate parameters). Third, taxes are significantly more efficient under uncertainty than rate controls. These three points could not be made without the kind of model developed in this paper.

Details of the integrated climate-economy model are laid out in Section 2. While the climate portion follows earlier work by Nordhaus (1994b), the economic model is based on a real business cycle (RBC) approach. Exogenous productivity changes are viewed as stochastic rather than deterministic phenomena and optimal intertemporal behavior is analytically approximated rather than numerically optimized. Also in this section the quantification of uncertainty is summarized and preference aggregation is discussed. Additional information about the RBC-style analytic approximation, econometric estimation, and the choice of social welfare functions is given in the appendix.

Section 3 examines the question of optimal control rates under uncertainty. In particular, do models with and without uncertainty recommend the same reductions? To answer this question, two policies are computed: one which is optimal with uncertainty as described in Section 2 and another which is optimal using median parameter values throughout the model. These policies are compared both in terms of stringency and social welfare relative to the policy ignoring uncertainty. The optimal policy under uncertainty involves an additional 6% reduction in CO₂ emissions initially (an 8.4% versus 7.9% control rate), rising to an additional 20% in 30 years and over 50% in 80 years. Perhaps surprisingly, around half of this increase can be reproduced by modeling preferences as the sole source of uncertainty. The certainty equivalent welfare gain is \$73 in current consumption per capita for the optimal policy which incorporates uncertainty versus only \$58 for the optimal policy ignoring uncertainty.

Section 4 explores the difference in efficiency between price and quantity instruments as a mechanism for regulating carbon dioxide emissions. From the early work of Weitzman (1974), we know that uncertainty about marginal costs leads to unequal consequences for price and quantity controls which are otherwise equivalent. However, it is unclear from

casual observation which instrument will be more efficient and whether the difference will be of practical significance. This section shows that taxes are preferable to quantity controls and that in terms of social welfare, the difference is significant. The optimal tax generates a \$86 certainty equivalent gain while the optimal rate generates a \$73 gain.

2 Model

2.1 Overview

The purpose of the model is to translate a climate change policy – either a tax or control rate – into a meaningful welfare measure of its consequences. This welfare measure, in turn, has two purposes: 1) to facilitate the calculation of optimal policies by maximizing the measure over a well-defined class of policies, and 2) to provide a metric for comparing policies drawn from different classes. Such classes might include taxes versus quantity controls, or policies which consider uncertainty versus those which ignore it.

The process of computing the welfare associated with a given policy can be viewed as a three-step process. First, for any given state of nature we need to determine how the economy and climate will evolve in the presence of the policy. This requires a model of consumer and producer behavior, a climate model, and links between the two. Second, a sample of states of nature (sets of values for all the model parameters) must be generated representing our view of uncertainty. Third, using the model to generate a set of outcomes for each of the sampled states of nature based on a particular policy, we need to assign a single measure of the social value associated with that policy. When preferences vary from state to state, this is a non-trivial task even when a representative agent model is used within each state.

In the first step, a modified stochastic growth model is used to capture consumer behavior in each state. Standard stochastic growth models choose the level of aggregate consumption each period based on the optimal behavior of a representative consumer in the face of

random IID exogenous shocks to productivity.³ In contrast, this paper models the trend in productivity as both declining over time and endogenized to depend on climate change. However, these distinctions are ignored by the representative consumer.⁴ This strategy has two important consequences. Consumption remains optimally allocated across time in every state. Therefore, the intertemporal consequences of any climate change policy will be correctly valued. Equally important, for any given state the evolution of the economy and climate is easily determined: the 400 period path for each of 1000 states can be computed in less than 5 seconds on a desktop computer.

For the second step, uncertainty is quantified two ways. Economic uncertainty (describing preferences, technology and current exogenous productivity trends) is based on an econometric analysis of historical data. Climate and long-term trend uncertainty (describing the climate consequences of economic activity, damages due to climate change, and eventual slowdowns in productivity and population growth) is based on discrete distributions developed in previous work.

The third step is to aggregate the results over many states into a single welfare measure. Based on the idea of cardinal unit comparability among suitably scaled utility measures in each state, social welfare is computed as an arithmetic average. This approach requires a somewhat arbitrary choice of utility scaling and eliminates the possibility of distributional considerations. While subject to criticism on both issues, it highlights a point which has been skipped in much of the discussion of climate change. Uncertainty about consumer/social preferences cannot be overcome without some aggregation scheme.

³For a recent summary of the solution techniques for stochastic growth models see Taylor and Uhlig (1990).

⁴The productivity slowdown and economy-climate links are based on William Nordhaus' (1994) DICE model. The fact that the consumer ignores these features has a negligible impact on the resulting path of consumption because the productivity slowdown and climate consequences (e.g., loss of output) occur quite gradually in the DICE model.

2.2 State-contingent Model

2.2.1 Economic Structure

Economic behavior within each state is derived from a representative agent model where consumption must be optimally allocated across time. In a model with constant exogenous productivity growth, preferences lead the economy to converge to a steady state. In the presence of random shocks and slowly changing trends, the economy instead converges to an ergodic distribution of states (due to the random shocks) which is itself slowly evolving (due to the slowly changing trends).

The representative consumer in this model exhibits constant relative risk aversion τ with respect to consumption per capita. Utility is separable across time, discounted at rate ρ and weighted each period by population. With the further assumption that preferences satisfy the von Neuman-Morgenstern axioms, the consumer's optimization problem can be written as

$$\max_{\{C_t\}_{t \in \{0,1,\dots\}}} E \left[\sum_{t=0}^{\infty} (1 + \rho)^{-t} N_t \frac{(C_t/N_t)^{1-\tau}}{1-\tau} \right] \quad (1)$$

where C_t is consumption in period t and N_t is population. This consumption program $\{C_0, C_1, \dots\}$ is subject to the resource constraints

$$Y_t = (A_t^* N_t)^{1-\theta} K_t^\theta \quad (2)$$

and

$$K_{t+1} = K_t(1 - \delta_k) + Y_t - C_t \quad (3)$$

where Y_t is aggregate output, A_t^* is net labor productivity and K_t is the capital stock. The parameter θ summarizes the Cobb-Douglas production technology given in Equation (2) and δ_k reflects the rate of capital depreciation in the capital accumulation equation (3). Finally, there is a transversality condition for a balanced growth steady state,

$$\rho > (1 - \tau) \times (\text{asymptotic growth rate of } A_t^*) \quad (4)$$

This is always satisfied by assuming zero growth asymptotically. 5

Even with exogenous constant growth models for N_t and A_t^* , the dynamic optimization problem given by Equations (1–3) is difficult if not impossible to solve analytically.⁵ However, choosing

$$\Delta \ln(K_{t+1}/N_{t+1}) = \alpha_1 + \alpha_2(\ln(K_t/N_t) - \ln(A_t^*)) \quad (5)$$

and

$$C_t = K_t(1 - \delta_k) + Y_t - K_{t+1} \quad (6)$$

– where α_1 and α_2 are functions of the parameters $(\rho, \tau, \theta, \delta_k)$ – yields a close approximation of optimal consumer behavior. This technique of approximating optimal dynamic behavior has its origins in the real business cycle literature beginning with Kydland and Prescott (1982). A simple derivation of expressions for α_1 and α_2 is given in Appendix A.

2.2.2 Long-term Growth, Climate Behavior and Damages

This section explains the remainder of the state-contingent model, including growth projections, climate behavior and damages from global warming [based primarily on the DICE model, Nordhaus (1994b)]. Net labor productivity, A_t^* , is modeled as the product of exogenous, gross labor productivity, A_t , a factor describing control costs, and a factor describing climate damages:

$$A_t^* = \underbrace{\left(\frac{1 - b_1 \mu_t^{b_2}}{1 + (D_0/9) \cdot T_t^2} \right)}_{\text{damages}}^{\text{control costs}}^{-\frac{1}{1-\theta}} A_t \quad (7)$$

μ_t is the fractional reduction in CO₂ emissions at time t (the “control rate”) versus a business as usual/no government policy baseline, while b_1 and b_2 parameterize the cost of attaining these reductions. T_t is the average surface temperature relative to pre-industrialization and

⁵Long and Plosser (1983) derive an analytic solution for the case of $\delta_k = 1$ and $\lim \tau \rightarrow 1$ (log utility).

D_0 is the fractional loss in aggregate GDP from a 3° temperature increase. The form of the control cost function is derived from an econometric analysis of a number of previous studies. The damage function is considerably more ad hoc. See Nordhaus (1993) and Nordhaus (1994a) for additional details on costs and damages, respectively.

Exogenous labor productivity A_t is modeled as a random walk in logarithms with an exponentially decaying drift. That is,

$$\log(A_t) = \log(A_{t-1}) + \gamma_a \exp(-\delta_a t) + \sigma_a \epsilon_t \quad (8)$$

where γ_a is the initial growth rate, δ_a is the annual decline in the growth rate, σ_a is the standard deviation of the random growth shocks and ϵ_t is a standard NIID random shock. Population is modeled similarly but without the stochastic element:

$$\log(N_t) = \log(N_{t-1}) + \gamma_n \exp(-\delta_n t) \quad (9)$$

where γ_n is the initial growth rate and δ_n is the annual decline in the growth rate. Note that these models predict zero growth asymptotically, though this may occur centuries in the future.⁶

The remaining portion of the model explains the link between economic activity (measured as aggregate output Y_t) and warming (measured as the average surface temperature T_t). The first step is linking output to emissions:

$$E_t = \sigma_t(1 - \mu_t)Y_t \left(\frac{A_t}{A_t^*} \right)^{1-\theta} \quad (10)$$

where E_t is emission of CO₂, σ_t is an exogenous trend in emissions/output, and μ_t is the rate of emissions reductions induced by the policymaker. The expression $Y_t \left(\frac{A_t}{A_t^*} \right)^{1-\theta}$ reflects “raw” output prior to reductions due to damages and control costs. The model of σ_t is, as with labor productivity and population, based on exponentially decaying growth:

$$\log(\sigma_t) = \log(\sigma_{t-1}) + \gamma_\sigma \exp(-\delta_\sigma t) \quad (11)$$

⁶For example, the range of parameters used in the simulations (with $\delta_{a/n} \in (0.25\%, 2.5\%)$) leads to a halving of the growth rates every 20 to 200 years.

where γ_o is the initial growth rate of emissions/output (a negative number) and δ_a is the annual decline in the growth rate. Note that the annual decline in the emissions/output rate is the same as the annual decline in labor productivity growth (δ_a).

Emissions of CO₂ accumulate in the atmosphere according to:

$$M_t - 590 = \beta E_{t-1} + (1 - \delta_m)(M_{t-1} - 590) \quad (12)$$

where M_t is the atmospheric concentration of CO₂ in billions of tons of carbon equivalent. β is a measure of the retention rate of emissions. Low values of β indicate that emissions do not, in fact, accumulate while a value of unity would mean that every ton of emitted CO₂ becomes a ton of atmospheric CO₂. The parameter δ_m plays the role of a depreciation rate: it is assumed that CO₂ in the atmosphere above the pre-industrialization level of 590 billion tons slowly decays. This decay reflects absorption of CO₂ into the oceans which are assumed to be an infinite sink.

Above average concentrations of CO₂ in the atmosphere lead to increased radiative forcings, a measure of the rate of transfer between solar energy produced by the sun and thermal energy stored in the atmosphere. This is modeled according to

$$F_t = 4.1 \times \log(M_t/590)/\log(2) + O_t \quad (13)$$

where F_t measures radiative forcings in units of watts per meter squared. The specification is such that a doubling of CO₂ concentrations leads to a roughly four fold increase in forcings. O_t in this relation represents radiative forcings due to other greenhouse gases and is assumed exogenous to the model:


$$O_t = \begin{cases} 0.2604 + 0.0125t - 0.000034t^2 & \text{if } t < 150 \\ 1.42 & \text{otherwise} \end{cases} \quad (14)$$

Increased forcings lead to temperature changes according to

$$T_t = T_{t-1} + (1/R_1) [F_t - \lambda T_{t-1} - (R_2/\tau_{12})(T_{t-1} - T_{t-1}^*)] \quad (15)$$

$$T_t^* = T_{t-1}^* + (1/R_2)(R_2/\tau_{12})(T_{t-1} - T_{t-1}^*) \quad (16)$$

Figure 1: State-contingent model



where T_t is the surface temperature and T_t^* is the deep ocean temperature, both expressed in changes relative to pre-industrialization levels in degrees Celsius. Note that if $M_t = 590$ and $O_t = 0$ (e.g., pre-industrialization), T_t and T_t^* will equilibrate to zero. The parameter λ describes the equilibrium change in surface temperature for a given change in radiative forcings. In particular, based on (13) and (15), a doubling of the concentration of CO_2 in the atmosphere will lead to a $4.1/\lambda$ rise in surface temperature in the long run. This parameter $4.1/\lambda$ is a measure of the temperature sensitivity of the atmosphere.⁷

The parameters R_1 , R_2 and τ_{12} describe the thermal capacity of the surface atmosphere and deep oceans and the rate of energy transfer between them, respectively.

2.2.3 Summary of State-contingent Climate-Economy Model

The state-contingent model is now complete. Figure 1 shows schematically how the model evolves over time and how policy (a path of reductions $\{\mu_t\}$) is translated into a state-contingent utility measure (as measured in Equation (1)).

At any point in time, the vector $\{K_t, A_t, N_t, \sigma_t, M_t, T_t, T_t^*\}$ completely describes the state of the economy and climate. The three somewhat distinct portions of the model – economy,

⁷ λ by itself is referred to as the climate feedback parameter.

climate and trends – are then used to evolve the current state into the next period state. The economy and climate models make use of the policy variable μ_t describing emission reductions in order to compute control costs and realized emissions, respectively. The realized path of consumption is an output of the economic model which is then used to evaluate the utility associated with implementing the policy in this state.

A key innovation relative to similar models is the use of Equation (5) to capture optimal intertemporal consumer behavior without performing a numerical optimization. This reduces by an order of magnitude the time required to evaluate policy while preserving the proper intertemporal prices.

2.3 Quantifying Uncertainty

Estimates of uncertainty in the model come from two sources: econometric analysis and subjective assessment. The model involves nineteen different parameters. Six are parameters describing observable economic activity:

- pure time preference ρ ,
- risk aversion τ ,
- output-capital elasticity θ ,
- productivity growth γ_a ,
- variation in productivity growth σ_a^2 , and
- depreciation δ_k .

A joint distribution for these parameters is estimated with historical data. The remaining thirteen describe emissions:

- emissions rate growth γ_σ ,

climate change:

- CO₂ retention rate β ,
- temperature sensitivity $4.1/\lambda$,
- CO₂ decay rate δ_m ,
- thermal capacities and conductivities R_1 , R_2 and τ_{12} ,

control costs and damages:

- cost function parameters b_1 and b_2 ,

- fractional loss of GDP for 3° temperature rise D_0 ,

and long-term growth trends

- population growth γ_n ,
- productivity slowdown δ_a ,
- population slowdown δ_n .

Uncertainty about these parameters is based on subjective analysis.

The econometric analysis of the six economic parameters is based on post-war U.S. data as described Appendix B. Series describing aggregate investment, capital services, output and prices are fit to the model described by Equations (2),(3) and (5). The posterior parameter distribution which arises from this analysis is summarized in Table 1.

Nordhaus (1994b)⁸ develops a distribution for the remaining parameters based on a two-step subjective analysis. The first step involves testing his model's sensitivity to each parameter being changed, one at a time, to a more extreme value. Those parameters which produce the largest variance in model output are then further scrutinized. A discrete, five-value distribution is developed for seven of these thirteen variables. The remaining six are fixed at their best guess values. The distribution of the seven uncertain parameters is summarized in Table 2. Values of the six fixed parameters as well as initial conditions for the model are given in Table 3.

2.4 Social Welfare

Once we concede that preferences are imprecisely known, the next logical step is to formulate a social welfare function to aggregate over these preferences. This social welfare function then allows the policymaker to compare alternatives in the face of such uncertainty. While it is more common to envision aggregation over different individuals, here different states of nature are being aggregated, each with a single but unique representative agent. It is important to recognize that although the parameters of a representative agent model can be inferred from observed consumer behavior, there is no information available to infer

⁸Chapters 6 and 7.

Table 1: Marginal distributions of uncertain economic parameters
(narrow bars indicate values used in simulations without uncertainty)







description	symbol	equation	distribution
pure rate of time preference	ρ	(1)	
coefficient of risk aversion	τ	(1)	
output-capital elasticity	θ	(2)	
rate of capital depreciation	δ_k	(3)	
initial productivity growth rate	γ_a	(8)	
standard error of productivity shocks	σ_a	(8)	

Table 2: Discrete distributions of uncertain climate/trend parameters
(narrow bars indicate values used in simulations without uncertainty)








description	symbol	equation	distribution
annual decline of population growth rate	δ_n	(9)	
annual decline of productivity growth rate	δ_a	(8),(11)	
initial growth rate of CO ₂ per unit output	γ_σ	(10),(11)	
damage parameter (% loss of GDP for 3° temperature rise)	D_0	(7)	
cost function parameter	b_1	(7)	
retention rate for CO ₂ emissions	β	(12)	
temperature sensitivity to CO ₂ doubling (in °C)	$4.1/\lambda$	(13),(15)	

Table 3: Description of fixed parameters^a

parameter	symbol	equation	units	value
cost function parameter	b_2	(7)		2.887
decay rate of atmospheric CO ₂	δ_m	(12)		0.00833
1/thermal capacity of atmosphere	$1/R_1$	(15)	°C-meter ² /watt-year	0.048
thermal conductivity b/w atmosphere and oceans	R_2/τ_{12}	(15),(16)	watt/°C-meter ²	0.44
1/thermal capacity of deep oceans	$1/R_2$	(16)	°C-meter ² /watt-year	$\frac{0.002}{0.44}$
1995 global population	N_0	(9)	millions of people	5590
initial population growth	γ_n	(9)		0.0124
initial rate of CO ₂ emissions per unit of output	σ_0	(11)	billion tons CO ₂ per \$1989 trillions	0.385
1995 global capital stock	K_0	(2),(3)	\$1989 trillions	79.5
1995 global output	Y_0	(2),(3),(10)	\$1989 trillions	24.0
1995 atmospheric concentrations of CO ₂	M_0	(12)	billions of tons of C equivalent	763.6
1995 surface temperature ^b	T_0	(7),(15),(16)	°Celsius	0.763
1995 deep ocean temperature ^b	T_0^*	(15),(16)	°Celsius	0.117

^aAll fixed parameters are from Nordhaus (1994b). The parameters that do not depend on time are from Nordhaus' Table 2.4. Initial values for temperature, CO₂ concentrations, and output in 1995, as well as the initial annual growth rate for population, are based on the Nordhaus base case simulation. The 1995 capital stock is adjusted upward to reflect differences in the definition of capital as well as underlying parameter values. The decay rate of atmospheric CO₂ is divided by ten to convert from a decennial to annual rate. The annual thermal capacity of the ocean and atmosphere are from the second line of Nordhaus' Table 3.4b.

^bTemperatures are measured as deviations from the pre-industrialization level, circa 1900.

parameters of the social welfare function. Instead, we must rely on social choice theory and our own sense of fairness.

The common approach in the climate change literature gets around the issue of uncertain preferences by reporting a range of policy prescriptions based on a range of possible preferences and states of nature. For example, Cline (1992) presents benefit-cost analyses for 92 different cases (Tables 7.3 and 7.4). Dowlatabadi and Morgan (1993) integrate out much of the uncertainty in their analysis, but still present results for 48 scenarios. Nordhaus (1994b) gives one of the few examples where *all* uncertainty including preferences is integrated out to yield a single policy recommendation (Chapter 8). However, the social welfare function implicit in the analysis is not revealed.

As most authors observe, uncertainty about time preference tends to dominate many policy considerations. Moreso, in fact, than uncertainty about climate sensitivity and damages. It is therefore important to consider how to aggregate over such preferences rather than guessing.

In this analysis it is assumed that social welfare is an average of rescaled utility measured in each state of nature (given by Equation (1)). The rescaling serves to equate the marginal social welfare of one additional dollar of current consumption in all states of nature. While arbitrary, some kind of adjustment is necessary to prevent the resulting policy prescription from being sensitive to the choice of units in the model (e.g., using dollars rather than thousands of dollars as the unit of account; see explanation in Appendix C). Social welfare can then be written as

$$SW(x) = I^{-1} \sum_{i=1}^I u(x, i) \quad (17)$$

where $u(x, i)$ is rescaled utility in state i with outcome x and $SW(x)$ is the social welfare associated with x . The rescaling is such that

$$u(+\$1 \text{ p.c. in initial period}, i) - u(\emptyset, i) = u(+\$1 \text{ p.c. in initial period}, j) - u(\emptyset, j) \quad \forall i, j$$

That is, a policy corresponding to an extra dollar of per capita in the initial period is assumed

to have the same utility in every state relative to a no policy ($\emptyset$) baseline.

This social welfare function has its origins in the literature on social choice. Harsanyi (1977) shows that in defining social welfare over lotteries, if individual preferences satisfy the von Neumann Morgenstern axioms then social welfare must have this weighted average form. Otherwise, social preferences will fail to mimic individual preferences over lotteries involving only that individual. This functional form can also be derived from the assumption of cardinal unit comparability, as discussed by Roberts (1980). A more general functional form therefore requires both integrating out uncertainty from the representative agent's perspective (to satisfy Harsanyi's point) and more stringent assumptions about the level of utility comparability (to satisfy Robert's point).


3 Does Uncertainty Matter?

To understand whether uncertainty significantly affects decision making and welfare consequences, two optimizations are performed. The first looks for the optimal CO₂ reductions based on the model developed in the previous section and the parameter distributions given in Tables 1 and 2. The second assumes there is no uncertainty. All uncertain parameters are fixed at their median values and the variance of the random exogenous growth shocks is set to zero (as indicated by the narrow bars in Tables 1 and 2). There is only one state of nature.

3.1 Policy Stringency


The resulting policies are shown in Figure 2. The initial difference between policies is small, 7.9% ignoring uncertainty versus 8.6% with uncertainty included. In thirty years these differences are higher, 10.4% versus 12.8% in 2025. By the end of the next century, however, the rate is almost doubled when uncertainty is considered. Uncertainty clearly has a dramatic effect on the future path of the optimal control rate, but with a relatively small consequence for immediate decisions.

Figure 2: Optimal climate change policy with and without uncertainty



What explains this pattern of small policy consequences now and large ones later? Surprisingly, much of the consequences (roughly 50%) can be traced to uncertainty about future interest rates, rather than climate or other economic uncertainty. This can be seen by reoptimizing the policy for a model with only pure time preference ρ and risk aversion τ uncertain, as illustrated in Figure 3. In contrast, a reoptimized policy based on modeling only the climate parameters (γ_σ , D_0 , b_1 , β , and $4.1/\lambda$) as uncertain yields essentially the same results as completely ignoring uncertainty.

Intuitively, Figure 3 shows the consequences of non-linearities in the first order conditions generating the optimal policy. If these first order conditions were linear in the uncertain parameters, replacing uncertain parameters with their means would have little effect. As it turns out, these conditions are highly non-linear functions of the interest rate but fairly linear in terms of the climate variables. Recall, for example, that the value of an annuity with coupon P and interest rate i is P/i – a highly non-linear function near $i = 0$. Furthermore, the interest rate is relatively well known during the initial periods of the simulation but less precisely known – and possibly near zero – in the future. Thus the policy effects of uncertainty are higher in the future.



Why does uncertainty about the interest rate rise in the future? The historic interest rate is an observed quantity and therefore easily predicted over short horizons. More generally, it depends in a simple way on future growth, risk aversion and time preference via the Euler equation: $i = \rho + \tau \cdot g_{a,t}$. Here i is the interest rate, ρ the pure rate of time preference, τ the coefficient of relative risk aversion and $g_{a,t}$ the rate of future productivity growth at time t . Future uncertainty arises because we assume productivity growth $g_{a,t}$ eventually declines to zero (according to Equation (8)) leaving only pure time preference ρ to motivate non-zero interest rates. While historical data tells us that with growth of 1.3%, interest rates will be 5%, we have much less information about how to dissect the observed 5% into its component pieces ρ and $\tau \cdot 1.3\%$. Therefore, we know less about what the interest rate will be once growth declines.

3.2 Welfare

In spite of the differences depicted in Figure 2, the welfare consequences of the two policies remain unclear. Suppose we analyze both policies using the model in Section 2, reflecting the uncertainty summarized in Tables 1 and 2. Does the policy which ignores uncertainty leave

us substantially worse off, given that uncertainty is in fact present? To answer this question, it is necessary to convert the social welfare measure used to determine these optimal policies in Equation (17) into a more meaningful metric. A natural vehicle for comparison is the certainty equivalent change in first period consumption yielding the same welfare consequences as the given policy. For example, implementing the policy which ignores uncertainty improves social welfare by 0.32 units based on Equation (17).⁹ A similar improvement in social welfare can be achieved by exogenously raising per capita consumption by \$58 in the first period in every state of nature.

The alternative policy which optimizes over uncertainty leads to a certainty equivalent gain of \$73 – a 25% improvement. The source of this improvement tends to be those states with low costs ($b_1 \downarrow$), high damages ($D_0 \uparrow$), high temperature sensitivity ($4.1/\lambda \uparrow$), low pure rates of time preference ($\rho \downarrow$), rapid productivity slowdown ($\delta_a \uparrow$), but high population growth ($\delta_n \downarrow$). In these states, the marginal costs to reduce emissions are low and the marginal benefits are high. Therefore, a more stringent policy will be preferred. While opposing cases (e.g., high costs, low damages) favor less stringent policy, such cases tend not to be as extreme. This is the non-linearity discussed earlier. In other words, once uncertainty is considered bad states of the world tend to be more extreme than good states and motivate additional stringency.

4 Instrument Choice

Uncertainty about the cost of reducing emissions leads to alternative consequences for quantity and price instruments. This suggests that welfare gains might be achieved by using taxes rather than rate controls to reduce emissions. The obvious experiment is to optimize

⁹It turns out that computing the welfare gain is complicated by extremely fat tails in the distribution of state-contingent welfare gains. Even with 50,000 draws, the mean welfare gain had not yet converged to its asymptotic normal distribution (a sample of 38 such means rejected normality at the 10^{-5} level). The somewhat ad hoc solution used here censors all gains higher than +150 and lower than -12 utility units. This corresponds to censoring at -\$1,200 and +\$11 million in terms of certainty equivalent gains and affects 0.1% of the states. This approach was chosen so that the reported results, if anything, are understated.

over a tax policy and then compare the resulting welfare gains.

Weitzman (1974) first demonstrated the dichotomy between price and quantity controls with uncertain costs in a simple partial equilibrium model. This model incorporates dynamic general equilibrium, leading to a more complex expression for the marginal cost of reduced emissions in period t :

$$MC_t = \frac{b_1 \cdot b_2}{\sigma_t \cdot (1 + D_0/9 \cdot T_t^2)} \mu_t^{b_2-1} \quad (18)$$

where b_1 and b_2 are cost function parameters, D_0 is the damage from a 3° temperature rise, T_t is the temperature rise since industrialization, μ_t is the control rate, and σ_t is the rate of emissions per unit of output. Since b_1 , σ_t , D_0 , and T_t vary among states of nature, the marginal cost will as well. Optimizing agents who set marginal cost equal to an imposed tax rate in each period will choose different control rates depending on the parameter values in their particular state of nature. Since a rate control would fix the emission reduction rate across all states, there is no rate control which can reproduce the consequences of a tax instrument in all states.

4.1 Taxes versus rate controls

In order to implement a tax instrument in the model, the control rate in each state of nature and time period is replaced with an expression for μ_t derived from (18).¹⁰ MC in this expression is replaced with the chosen tax rate. The optimal tax path is then determined by maximizing the social welfare expression in (17).

The tax instrument does in fact lead to a higher welfare gain compared to the rate instrument. When converted to a certainty equivalent change in first period consumption, the gain is \$86 per person for the tax. This is a 20% improvement over the rate instrument, which yields only \$73. Note that compared to the rate instrument computed in the absence of uncertainty, this represents a 50% overall gain. That is, a policymaker who ignores

¹⁰Derived rates higher than 100% are set to 100%, indicating that the tax induces complete elimination of emissions


uncertainty and implements the optimal rate policy would improve social welfare by \$58 per person (where the measured gain is based on the assumption that uncertainty does, in fact, exist). He or she would not have any reason to consider tax instruments since they are equivalent in the absence of uncertainty. Compare this outcome to that of a second policymaker who considers uncertainty and computes both optimal tax and control rate policies: This policymaker would see that taxes perform better and, in implementing the optimal tax policy, would improve social welfare by \$86 per person. This represents a net improvement of nearly 50% over the first policymaker's outcome.

Why are taxes preferred to rate controls? Weitzman's original insight was that this would be true when the elasticity of reductions demand (e.g., benefits) exceeded the elasticity of reductions supply (e.g., costs). The left panel of Figure 4 illustrates this intuition: when the cost curve shifts (dashed lines), the optimal shadow price (intersection of costs and expected benefits) is relatively constant compared to the control rate. This is only part of the reason, however. Weitzman's analysis was based on constant elasticities and, importantly, independent shocks to supply and demand. Here, high benefits occur when productivity growth declines rapidly ($\delta_a \uparrow$), in turn permitting low rates of pure time preference to generate high expected benefits more quickly, as discussed in Section 3.1. But a high value of δ_a , by the definition of the emissions rate (11), also leads to higher values of σ_t and low marginal control costs by (18) (recall that growth in σ_t is initially negative). Hence upward shifts in the benefits curve are correlated with downward shifts in the cost curve. Examining the right panel of Figure 4, we see that this also leads to a preference for taxes since the optimal shadow price is again less variable than the optimal rate.

4.2 Value of emission rights

Another way to view the dichotomy between taxes and rate controls is in terms of the value of emission rights. Based on a tax instrument, the value of emission rights is held constant across states at the given tax rate. Under a rate or quantity control, however, emission

Figure 4: Intuition for taxes being preferred to rate instruments
(dashed lines indicate particular realizations of uncertain costs and benefits)




Original Weitzman (1974) result that relatively flatter (higher elasticity) marginal benefits favor taxes. Note that all three intersections correspond to similar shadow prices but different control rates. Hence the deadweight loss will be lower with a fixed shadow price rather than a fixed control rate.

Correlation of upward shifts in benefits with downward shifts in costs favors taxes. Note that both intersections lead to similar shadow prices. As in the left panel, the deadweight loss is lower with the shadow price fixed (via a tax) rather than the control rate fixed.

rights will have different values in different states. Using the results of many simulations, it is possible to plot the distribution of resulting values and obtain a visual sense of the difference between instruments.

Figure 5 shows the median, 50% and 95% forecast intervals for the value of emission rights under the optimal rate control, along with the fixed values under the optimal tax policy (the optimal tax policy which ignores uncertainty is presented for comparison). While the median case under the optimal rate control closely tracks the optimal tax policy, the 95% forecast interval includes widely divergent values. For example, the median value of a right to emit one ton of carbon dioxide rises to around \$20 after 50 years. It is entirely possible, however, that under an optimal rate control the value will be over \$100. With these magnitudes of difference between the two instruments in terms of marginal costs, the distinction in welfare consequences discussed in the previous section is no longer surprising.

Figure 5: Value of emission rights under alternative policies



5 Conclusion

This paper has sought to inject discussions of climate change policy with a dose of skepticism regarding the limitations of models which ignore uncertainty. Excluding uncertainty tends to reduce expected benefits due to highly non-linear relations leading to policy recommendations which are too lax. Much of this effect can be related to agent preferences; over half the noted consequences can be replicated with preferences as the sole source of uncertainty. This raises the issue of how to aggregate over preferences, a question which is at once both controversial but necessary. Finally, uncertainty about costs introduces a dichotomy between price and quantity controls. In the case of carbon dioxide emissions, taxes perform better. The welfare gain associated with the optimal tax instrument is \$86 per person whereas the gain associated with the optimal rate instrument is only \$73. An optimal rate policy computed in the absence of uncertainty has a gain of only \$58.

While the results of this analysis should be taken with many of the same caveats accorded other integrated assessment models (the dependence on stylized and simplified models, parameterization, etc.), there are important distinctions to be made. This is the first dynamic general equilibrium model to incorporate large numbers of uncertain states of nature. It is

also takes advantage of econometrically quantified uncertainty. Lastly, the model explicitly addresses the issue of aggregation over uncertain preferences.

The most important consequence of this paper may be to encourage a reassessment of the current agenda for climate change research. Considerable effort has been directed towards the importance of learning; perhaps more effort should be directed at the consequences of uncertainty itself – without learning. A great deal of attention has focused on climate uncertainty; this paper suggests that considerable uncertainty surrounds the economy. In particular, the course of productivity growth has an important bearing on both future emissions and the valuation of future benefits via the interest rate. Finally, much of the policy discussion seems to have shifted towards quantity controls as the preferred instrument; this paper suggests that efficiency arguments may point the other way, towards taxes.

Appendix

A Notes on Solution Algorithm

This section explains the solution to the consumer optimization problem (1) given by Equation (5). In words, (5) is a log-linear approximation of the optimal decision rule based on the deterministic steady-state level of capital per efficiency unit of labor, e.g., $K_t/(A_t^* N_t)$. This approach turns out to work rather well for three reasons: First, the economy starts off relatively close to the steady state defined by the initial growth rates. Second, the decline in deterministic growth rates, described in Equations (8–9), has little effect on the approximation parameters. Third, the shocks away from the steady state – from both exogenous stochastic productivity shocks as well as endogenous climate change effects – are relatively small. While there are many ways of deriving expressions for α_1 and α_2 in terms of the underlying parameters, the simplest is based on a logarithmic approximation to the saddle path around the steady state in $C_t/(A_t^* N_t)$ and $K_t/(A_t^* N_t)$ space.

We begin with the Euler equation arising in a balanced growth equilibrium (where $\gamma_a =$

$\log(A_t^*) - \log(A_{t-1}^*)$ and $\gamma_n = \log(N_t) - \log(N_{t-1})$ are constant $\forall t$:¹¹

$$\left(\frac{C_{t+1}/N_{t+1}}{C_t/N_t}\right)^\tau = (1 + \rho)^{-1} (\theta(K_{t+1}/A_{t+1}^* N_{t+1})^{\theta-1} + 1 - \delta_k) \quad (\text{A.1})$$

At the balanced growth equilibrium $\frac{C_{t+1}/N_{t+1}}{C_t/N_t} = \exp(\gamma_a)$ so (A.1) can be solved for steady state capital stock per efficiency unit of labor, $K_{t+1}/(A_{t+1}^* N_{t+1}) = \tilde{K}^\dagger = \left(\frac{(1+\rho)\exp(\gamma_a\tau) - (1-\delta_k)}{\theta}\right)^{\frac{1}{\theta-1}}$.

The capital accumulation equation (3),

$$K_{t+1} = K_t(1 - \delta_k) + (A^* N)^{1-\theta} K_t^\theta - C_t \quad (\text{A.2})$$

then yields steady state consumption per efficiency unit of labor $C_t/(A_t^* N_t) = \tilde{C}^\dagger = \tilde{K}^{\dagger\theta} + (1 - \delta_k - \exp(\gamma_a + \gamma_n))\tilde{K}^\dagger$.

The approximation (5) is made by taking total derivatives of both (A.1–A.2) with respect to $\log(\tilde{K}_{t+1})$, $\log(\tilde{K}_t)$, $\log(\tilde{C}_{t+1})$ and $\log(\tilde{C}_t)$ around the steady-state values $\tilde{K}_t = \tilde{K}_{t+1} = \tilde{K}^\dagger$ and $\tilde{C}_t = \tilde{C}_{t+1} = \tilde{C}^\dagger$. This leads to the expressions

$$\tau(dc_{t+1} - dc_t) = \kappa_{ck} dk_{t+1} \quad (\text{A.3})$$

$$\exp(\gamma_a + \gamma_n)dk_{t+1} = \kappa_{kk} dk_t + \kappa_{kc} dc_t \quad (\text{A.4})$$

where $dc_t = \log(\tilde{C}_t) - \log(\tilde{C}^\dagger)$ is the log deviation of $\tilde{C}_t$ and $dk_t = \log(\tilde{K}_t) - \log(\tilde{K}^\dagger)$ is the log deviation of $\tilde{K}_t$. κ_{ck} , κ_{kk} and κ_{kc} are constants, the derivatives of the log of the right hand side of (A.1) with respect to $\log(\tilde{K}_{t+1})$ and of the log of the right hand side of (A.2) with respect to $\log(\tilde{K}_t)$ and $\log(\tilde{C}_t)$.¹² Subtracting (A.4) from the same expression evaluated one period forward yields

$$\exp(\gamma_a + \gamma_n)(dk_{t+2} - dk_{t+1}) = \kappa_{kk}(dk_{t+1} - dk_t) + \kappa_{kc}(dc_{t+1} - dc_t) \quad (\text{A.5})$$

Equation (A.3) can then be used to substitute for $dc_{t+1} - dc_t$, leaving a second order difference equation in dk_t :

$$\exp(\gamma_a + \gamma_n)(dk_{t+2} - dk_{t+1}) = \kappa_{kk}(dk_{t+1} - dk_t) + \kappa_{kc}(\kappa_{ck}/\tau) dk_{t+1} \quad (\text{A.6})$$

¹¹The Euler equation can be derived from Equation (1) by ignoring random productivity shocks and examining a perturbation which decreases consumption in period t , invests the extra output, then consumes the gross return (interest and principal) in period $t + 1$. If the path of consumption is indeed optimal, this should have no effect on utility and leads to the condition given in (A.1).

¹²Note that $\kappa_{ck} = (1 + \rho)^{-1}(\theta - 1)\theta\tilde{K}^{\dagger\theta-1}\exp(-\gamma_a\tau)$, $\kappa_{kk} = (1 + \rho)\exp(\gamma_a\tau)$, and $\kappa_{kc} = -\tilde{C}^\dagger/\tilde{K}^\dagger$.

From (5) and the observation that when $\tilde{K}_t = \tilde{K}^\dagger$ the growth of capital stock per capita is the steady state rate of γ_a , so $\Delta k_t = \gamma_a$, we know that $\gamma_a = \alpha_1 + \alpha_2 \log(\tilde{K}^\dagger)$. This leads to

$$\alpha_1 = \gamma_a - \alpha_2 \log(\tilde{K}^\dagger) \quad (\text{A.7})$$

and suggests a guess for solution to (A.6):

$$dk_{t+1} = (\alpha_2 + 1)dk_t \quad (\text{A.8})$$

(since $dk_{t+1} - dk_t = \Delta k_{t+1} - \gamma_a = \alpha_2 dk_t$). Then (A.6) yields the characteristic equation

$$\exp(\gamma_a + \gamma_n) [(\alpha_2 + 1)^2 - (\alpha_2 + 1)] = \kappa_{kk}((\alpha_2 + 1) - 1) + \kappa_{kc}(\kappa_{ck}/\tau)(\alpha_2 + 1) \quad (\text{A.9})$$

which can be used to find α_2 in terms of κ_{kk} , κ_{kc} and κ_{ck} :

$$(\alpha_2 + 1) = \frac{e^{\gamma_a + \gamma_n} + \kappa_{kk} + \kappa_{kc}(\kappa_{ck}/\tau) - \sqrt{(e^{\gamma_a + \gamma_n} + \kappa_{kk} + \kappa_{kc}(\kappa_{ck}/\tau))^2 - 4e^{\gamma_a + \gamma_n}\kappa_{kk}}}{2e^{\gamma_a + \gamma_n}} \quad (\text{A.10})$$

For $\rho = 0.04$, $\tau = 1.2$, $\theta = 0.38$, $\delta_k = 0.05$, $\gamma_a = 0.013$ and $\gamma_n = 0.012$ we have $\tilde{K}^\dagger = 7.8$, $C^\dagger = 1.6$, $\kappa_{ck} = -0.062$, $\kappa_{kk} = 1.06$ and $\kappa_{kc} = -0.20$. The yields $\alpha_2 = -0.084$ and $\alpha_1 = 0.185$.

The fact that we can find a value of α_2 which satisfies (A.6) validates our initial guess. Namely, a linear rule for Δk_{t+1} in terms of $k_t - a_t^*$ as in (5) or equivalently (A.8).

B Notes on Econometrics

The model given by Equations (2), (3) and (5) can be coupled with the price implications of market equilibrium (Sheperd's Lemma) and a simplified model of productivity shocks (ignoring the future slowdown and climate consequences) to yield the following econometric

model:

$$\Delta a_t^* = \gamma_a + \epsilon_t \quad (\text{B.11})$$

$$\Delta k_t = (\alpha_1 - \alpha_2 a_0^*) + \alpha_2(k_{t-1} - (a_{t-1}^* - a_0^*)) + \eta_{2t} \quad (\text{B.12})$$

$$y_t = (1 - \theta)a_0^* + \theta k_t + (1 - \theta)(a_t^* - a_0^*) + \eta_{1t} \quad (\text{B.13})$$

$$\frac{K_{t+1} - (Y_t - C_t)}{K_t} = 1 - \delta_k + \eta_{3t} \quad (\text{B.14})$$

$$\frac{P_{K,t} K_t}{P_{Y,t} Y_t} = \theta + \nu_t \quad (\text{B.15})$$

$$\nu_t = r_v \nu_{t-1} + \eta_{4t} \quad (\text{B.16})$$

Here $a_t^* = \log(A_t^*)$, $k_t = \log(K_t/N_t)$, $y_t = \log(Y_t/N_t)$, $P_{K,t}$ is the price of capital services and $P_{Y,t}$ is the price of output; the remaining variables are as defined in Section 2.¹³ The disturbances η_t and ϵ_t are assumed to be jointly NIID with zero covariance (note that the capital share is defined with an autoregressive error). This specification defines a likelihood function describing the distribution of data and unobserved productivity shocks conditional on the model parameters.

This likelihood function coupled with a prior distribution over the parameters can be used to obtain a function describing the posterior density of the parameters and productivity shocks conditional on the observed data.¹⁴ This is Bayes' rule. Obtaining a sample of draws from the posterior density is difficult, however, as the distribution is non-standard. To generate the distributions shown in Table 1, the Gibbs sampler was used.¹⁵ Rather than drawing combinations of all the parameters at once, the Gibbs sampler draws each unobserved parameter sequentially based on its distribution conditional on the previous draw of all other parameters. A chain of draws is formed which converges to the joint posterior

¹³The data, for the years 1952-1992, is from the U.S. Worksheets [Jorgenson (1980); see Ho (1989) and Wilcoxon (1988) for further details]. Output is a divisia index of output of consumption goods (c790+c781), investment goods (c789) and government services (c943). Capital is aggregate capital services (c667). Changes in relative prices among output goods (c792, c761, c791, c395) are subsumed into quantity changes for simplicity. Population is 16+ (c941).

¹⁴The Jeffrey's prior (see Section 2.8 of Gelman, Carlin, Stern, and Rubin (1995)) is used for each parameter, restricted to the economically relevant parameter space (e.g., $0 < \theta < 1$, $\delta_k > 0$, $\alpha_2 < 0$).

¹⁵Geman and Geman (1984), Tanner and Wong (1987), Jacquier, Polson, and Rossi (1994).

distribution.

The distribution used in this paper is based on ten chains of 1500 draws with the first 500 draws dropped to reduce start-up effects. The sample was further reduced by about 10% by requiring that the reduced form parameters α_1 and α_2 map back into meaningful structural parameters ρ and τ .

C Notes on Utility Rescaling

Policy consequences in a given state are always expressible in terms of a first period consumption equivalent. Letting $C_{i,1}(x)$ be this consumption equivalent for state i and policy x , a social welfare function which simply averages utility yields

$$SW(x) = I^{-1} \sum_i \frac{(C_{i,1}(x)/N_1)^{1-\tau_i}}{1-\tau_i}$$

where N_1 is the first period population (known with certainty) and τ_i is the coefficient of relative risk aversion in state i . But scaling the units of consumption by a factor of ψ changes social the social welfare associated with policy x to:

$$SW(x) = I^{-1} \sum_i \psi^{1-\tau_i} \frac{(C_{i,1}(x)/N_1)^{1-\tau_i}}{1-\tau_i}$$

Unless the consequences are the same for each state, this simple change in units – which leads to an unintentional reweighting among states – will change the ordering among policies. The proposed solution takes the form

$$SW(x) = I^{-1} \sum_i (C_{i,1}(\emptyset)/N_1)^{\tau_i} \frac{(C_{i,1}(x)/N_1)^{1-\tau_i}}{1-\tau_i}$$

where $C_{i,1}(\emptyset)$ is the consumption level in the absense of policy and

$$u(x, i) = (C_{i,1}(\emptyset)/N_1)^{\tau_i} \frac{(C_{i,1}(x)/N_1)^{1-\tau_i}}{1-\tau_i}$$

is the rescaled utility measure in (17). Note that renormalizing the units of consumption no longer affects the policy ordering.

References

- Cline, W. (1992). *The Economics of Global Warming*. Washington: Institute of International Economics.
- Dowlatabadi, H. and M. Morgan (1993). A model framework for integrated studies of the climate problem. *Energy Policy* 21(3), 209–221.
- Gelman, A., J. B. Carlin, H. S. Stern, and D. B. Rubin (1995). *Bayesian Data Analysis*. New York: Chapman & Hall.
- Geman, S. and D. Geman (1984). Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 6, 721–741.
- Harsanyi, J. C. (1977). *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*. New York: Cambridge University Press.
- Ho, M. S. (1989). *The effects of external linkages on U.S. economic growth: A dynamic general equilibrium analysis*. Ph. D. thesis, Harvard University.
- Jacquier, E., N. G. Polson, and P. E. Rossi (1994). Bayesian analysis of stochastic volatility models. *Journal of Business and Economic Statistics* 12(4), 371–389.
- Jorgenson, D. W. (1980). Accounting for capital. In G. M. von Furstenberg (Ed.), *Capital, Efficiency, and Growth*, pp. 251–320. Cambridge: Ballinger.
- Kydland, F. E. and E. C. Prescott (1982). Time to build and aggregate fluctuations. *Econometrica* 50(6), 1345–1370.
- Long, J. B. and C. I. Plosser (1983). Real business cycles. *Journal of Political Economy* 91, 39–69.
- Manne, A. and R. Richels (1995). The greenhouse debate: Economic efficiency, burden sharing and hedging strategies. *The Energy Journal* 16(4), 1–37.
- Manne, A. S. (1996, April). Hedging strategies for global carbon dioxide abatement: A summary of poll results. Energy Modeling Forum 14 Subgroup – Analysis for Decisions Under Uncertainty.
- Manne, A. S. and R. G. Richels (1992). *Buying Greenhouse Insurance*. Cambridge: MIT Press.
- Nordhaus, W. D. (1993). The cost of slowing climate change. *Energy Journal* 12(1), 37–65.
- Nordhaus, W. D. (1994a). Expert opinion on climate change. *American Scientist* 82(1), 45–51.
- Nordhaus, W. D. (1994b). *Managing the Global Commons*. Cambridge: MIT Press.
- Roberts, K. (1980). Interpersonal comparability and social choice theory. *Review of Economic Studies* 47, 421–439.
- Tanner, M. and W. Wong (1987). The calculation of posterior distributions by data augmentation. *Journal of the American Statistical Association* 82, 528–550.


- Taylor, J. B. and H. Uhlig (1990). Solving nonlinear stochastic growth models: A comparison of alternative solution methods. *Journal of Business and Economic Statistics* 8(1), 1-17.
- Weitzman, M. L. (1974). Prices vs. quantities. *Review of Economic Studies* 41(4), 477-491.
- Wilcoxon, P. J. (1988). *The effects of environmental regulation and energy prices on U.S. economic performance*. Ph. D. thesis, Harvard University.

DISCUSSION PAPER

Optimal Choice of Policy Instrument and Stringency under Uncertainty: The Case of Climate Change

William A. Pizer

Discussion Paper 97-17



RESOURCES
FOR THE FUTURE

**Optimal Choice of Policy Instrument
and Stringency under Uncertainty:
The Case of Climate Change**

William A. Pizer

Discussion Paper 97-17

December 1996

Resources for the Future
1616 P Street, NW
Washington, DC 20036
Telephone 202-328-5000
Fax 202-939-3460

© 1996 Resources for the Future. All rights reserved.
No portion of this paper may be reproduced without permission
of the author.

Discussion papers are research materials circulated by their
authors for purposes of information and discussion. They have
not undergone formal peer review or the editorial treatment
accorded RFF books and other publications.

Optimal Choice of Policy Instrument and Stringency under Uncertainty: The Case of Climate Change

William A. Pizer

Abstract

The relevance of uncertainty in the climate change policy debate is without doubt. Surprisingly, there have been few attempts to examine the direct policy consequences of including uncertainty in an integrated climate-economy framework. This paper presents results concerning optimal policy stringency and instrument choice when economic and climate parameters assume distributions rather than single values. Uncertainty is found to raise the optimal level of emission reductions relative to an optimization based on one set of central parameter estimates. Much of this effect can be related to economic rather than climate uncertainty. Uncertainty also leads to a preference for taxes over quantity controls.

Previous studies of uncertainty in the climate change context have used a small number of states to measure the value of earlier information, learning and adaptation. This paper attempts to refocus attention on the more basic question of whether, in the absence of new information and learning, the inclusion of uncertainty yields significantly different policy conclusions. For policymakers confronting the problem of climate change today, this is the more relevant question.

Key Words: uncertainty, climate change, taxes and quantity controls

JEL Classification No(s): C15, D81, Q28

Table of Contents

1. Introduction	1
2. Model	4
2.1 Overview	4
2.2 State-contingent Model	6
2.2.1 Economic Structure	6
2.2.2 Long-term Growth, Climate Behavior and Damages	7
2.2.3 Summary of State-contingent Climate-Economy Model	10
2.3 Quantifying Uncertainty	11
2.4 Social Welfare	12
3. Does Uncertainty Matter?	17
3.1 Policy Stringency	17
3.2 Welfare	19
4. Instrument Choice	20
4.1 Taxes versus rate controls	21
4.2 Value of emission rights	22
5. Conclusion	24
Appendix	25
A. Notes on Solution Algorithm	25
B. Notes on Econometrics	27
C. Notes on Utility Rescaling	29
References	30

List of Figures and Tables

Figure 1. State-contingent model	10
Figure 2. Optimal climate change policy with and without uncertainty	18
Figure 3. Decomposition of policy consequences	19
Figure 4. Intuition for taxes being preferred to rate instruments	23
Figure 5. Value of emission rights under alternative policies	24
Table 1. Marginal distributions of uncertain economic parameters	13
Table 2. Discrete distributions of uncertain climate / trend parameters	14
Table 3. Description of fixed parameters	15

Optimal Choice of Policy Instrument and Stringency Under ² Uncertainty: The Case of Climate Change

William A. Pizer¹

1 Introduction

Uncertainty is a pervasive feature of climate change analysis. The wide range of divergent opinions on the proper policy response to climate change is testament to the scale of the issue. However, there is a difference between recognizing uncertainty and using it to propose *one* policy response.

This paper acknowledges uncertainty and optimizes over it to recommend a single course of action. This permits a useful assessment of the importance of including uncertainty in the analysis: Does it make a difference versus optimizing with a single set of central parameter values? Do alternative policy instruments behave differently? Why?

Previous work in this vein has tended to miss these two basic questions. Frequently, the importance of uncertainty is framed in terms of learning: what happens if uncertainty is resolved now rather than later [Manne and Richels (1993,1995), Nordhaus (1994b), Manne (1996)]. This skirts the original question by failing to address how, without learning, uncertainty affects policy answers in the first place. The typical use of a highly stylized framework with only two or three states of nature also diminishes one's confidence that these models are capturing important features of the unknown parameter space. Simply considering high and low values for growth, climate sensitivity, damages and costs requires more than a dozen states of nature. Existing work which does consider a large space of uncertainty does so without a general equilibrium model of intertemporal prices [Dowlatabadi and Morgan (1993)]. This turns out to be inadequate for determining optimal policy and, by extension,

¹Fellow, Quality of the Environment Division, Resources for the Future. The author would like to thank Dale Jorgenson, Gary Chamberlain and seminar participants at Harvard University and Resources for the Future for helpful comments on earlier drafts of this paper.

for addressing uncertain preferences. Preferences, it turns out, are extremely important.⁴

In this paper a relatively simple integrated climate-economy model based on optimal intertemporal behavior is used to examine alternative policies under uncertainty. Several features distinguish this model from other climate-economy models. First, a fast simulation technique is used to permit thousands of different states of nature to be simulated simultaneously, each capturing optimal consumer behavior. Second, the distribution of the economic states of nature is based on econometrically estimated parameter distributions. Third, these states – each of which exhibits different consumer preferences – are consistently aggregated using an explicit social welfare function.

Each of these features is important given the question to be addressed. Without a general equilibrium framework, intertemporal prices are likely to be wrong. Higher productivity growth, for example, leads to more output and emissions but simultaneously higher interest rates. Ignoring the higher interest rates, or timing them incorrectly, will lead to incorrect intertemporal valuation and biased results.² Econometric estimates of economic parameter distributions are important to capture well-known parameter relations and their associated uncertainty. In order to be consistent with historic interest rates, for example, time preference and intertemporal substitutability must be positively correlated. By how much and with what degree of certainty is a question best addressed by data. Finally, given the extremely long horizons involved in climate change analysis, it is not surprising that the results are extremely sensitive to variation in preferences across time. Without a consistent method of aggregating over these preferences, optimal policy responses will depend quite heavily on which rate of pure time preference is assumed – without any assessment of how bad an incorrect guess will be.

The short answer to the question posed is that uncertainty does matter. First, ignoring it leads to policy recommendations which are too slack with welfare gains 30% below that of the optimal policy incorporating uncertainty. Second, a large part of the consequences – around

²An error in interest rates will bias the valuation since policy costs typically occur in earlier years than policy benefits.

fifty percent – arises solely from modeling uncertain preferences (versus uncertain technology⁵ or climate parameters). Third, taxes are significantly more efficient under uncertainty than rate controls. These three points could not be made without the kind of model developed in this paper.

Details of the integrated climate-economy model are laid out in Section 2. While the climate portion follows earlier work by Nordhaus (1994b), the economic model is based on a real business cycle (RBC) approach. Exogenous productivity changes are viewed as stochastic rather than deterministic phenomena and optimal intertemporal behavior is analytically approximated rather than numerically optimized. Also in this section the quantification of uncertainty is summarized and preference aggregation is discussed. Additional information about the RBC-style analytic approximation, econometric estimation, and the choice of social welfare functions is given in the appendix.

Section 3 examines the question of optimal control rates under uncertainty. In particular, do models with and without uncertainty recommend the same reductions? To answer this question, two policies are computed: one which is optimal with uncertainty as described in Section 2 and another which is optimal using median parameter values throughout the model. These policies are compared both in terms of stringency and social welfare relative to the policy ignoring uncertainty. The optimal policy under uncertainty involves an additional 6% reduction in CO₂ emissions initially (an 8.4% versus 7.9% control rate), rising to an additional 20% in 30 years and over 50% in 80 years. Perhaps surprisingly, around half of this increase can be reproduced by modeling preferences as the sole source of uncertainty. The certainty equivalent welfare gain is \$73 in current consumption per capita for the optimal policy which incorporates uncertainty versus only \$58 for the optimal policy ignoring uncertainty.

Section 4 explores the difference in efficiency between price and quantity instruments as a mechanism for regulating carbon dioxide emissions. From the early work of Weitzman (1974), we know that uncertainty about marginal costs leads to unequal consequences for price and quantity controls which are otherwise equivalent. However, it is unclear from

casual observation which instrument will be more efficient and whether the difference will be of practical significance. This section shows that taxes are preferable to quantity controls and that in terms of social welfare, the difference is significant. The optimal tax generates a \$86 certainty equivalent gain while the optimal rate generates a \$73 gain.

2 Model

2.1 Overview

The purpose of the model is to translate a climate change policy – either a tax or control rate – into a meaningful welfare measure of its consequences. This welfare measure, in turn, has two purposes: 1) to facilitate the calculation of optimal policies by maximizing the measure over a well-defined class of policies, and 2) to provide a metric for comparing policies drawn from different classes. Such classes might include taxes versus quantity controls, or policies which consider uncertainty versus those which ignore it.

The process of computing the welfare associated with a given policy can be viewed as a three-step process. First, for any given state of nature we need to determine how the economy and climate will evolve in the presence of the policy. This requires a model of consumer and producer behavior, a climate model, and links between the two. Second, a sample of states of nature (sets of values for all the model parameters) must be generated representing our view of uncertainty. Third, using the model to generate a set of outcomes for each of the sampled states of nature based on a particular policy, we need to assign a single measure of the social value associated with that policy. When preferences vary from state to state, this is a non-trivial task even when a representative agent model is used within each state.

In the first step, a modified stochastic growth model is used to capture consumer behavior in each state. Standard stochastic growth models choose the level of aggregate consumption each period based on the optimal behavior of a representative consumer in the face of

random IID exogenous shocks to productivity.³ In contrast, this paper models the trend⁴ in productivity as both declining over time and endogenized to depend on climate change. However, these distinctions are ignored by the representative consumer.⁴ This strategy has two important consequences. Consumption remains optimally allocated across time in every state. Therefore, the intertemporal consequences of any climate change policy will be correctly valued. Equally important, for any given state the evolution of the economy and climate is easily determined: the 400 period path for each of 1000 states can be computed in less than 5 seconds on a desktop computer.

For the second step, uncertainty is quantified two ways. Economic uncertainty (describing preferences, technology and current exogenous productivity trends) is based on an econometric analysis of historical data. Climate and long-term trend uncertainty (describing the climate consequences of economic activity, damages due to climate change, and eventual slowdowns in productivity and population growth) is based on discrete distributions developed in previous work.

The third step is to aggregate the results over many states into a single welfare measure. Based on the idea of cardinal unit comparability among suitably scaled utility measures in each state, social welfare is computed as an arithmetic average. This approach requires a somewhat arbitrary choice of utility scaling and eliminates the possibility of distributional considerations. While subject to criticism on both issues, it highlights a point which has been skipped in much of the discussion of climate change. Uncertainty about consumer/social preferences cannot be overcome without some aggregation scheme.

³For a recent summary of the solution techniques for stochastic growth models see Taylor and Uhlig (1990).

⁴The productivity slowdown and economy-climate links are based on William Nordhaus' (1994) DICE model. The fact that the consumer ignores these features has a negligible impact on the resulting path of consumption because the productivity slowdown and climate consequences (e.g., loss of output) occur quite gradually in the DICE model.

2.2 State-contingent Model

2.2.1 Economic Structure

Economic behavior within each state is derived from a representative agent model where consumption must be optimally allocated across time. In a model with constant exogenous productivity growth, preferences lead the economy to converge to a steady state. In the presence of random shocks and slowly changing trends, the economy instead converges to an ergodic distribution of states (due to the random shocks) which is itself slowly evolving (due to the slowly changing trends).

The representative consumer in this model exhibits constant relative risk aversion τ with respect to consumption per capita. Utility is separable across time, discounted at rate ρ and weighted each period by population. With the further assumption that preferences satisfy the von Neuman-Morgenstern axioms, the consumer's optimization problem can be written as

$$\max_{C_t} E \left[\sum_{t=0}^{\infty} (1 + \rho)^{-t} N_t \frac{(C_t/N_t)^{1-\tau}}{1-\tau} \right] \quad (1)$$

where C_t is consumption in period t and N_t is population. This consumption program $\{C_0, C_1, \dots\}$ is subject to the resource constraints

$$Y_t = (A_t^* N_t)^{1-\theta} K_t^\theta \quad (2)$$

and

$$K_{t+1} = K_t(1 - \delta_k) + Y_t - C_t \quad (3)$$

where Y_t is aggregate output, A_t^* is net labor productivity and K_t is the capital stock. The parameter θ summarizes the Cobb-Douglas production technology given in Equation (2) and δ_k reflects the rate of capital depreciation in the capital accumulation equation (3). Finally, there is a transversality condition for a balanced growth steady state,

$$\rho > (1 - \tau) \times (\text{asymptotic growth rate of } A_t^*) \quad (4)$$

This is always satisfied by assuming zero growth asymptotically. ✱

Even with exogenous constant growth models for N_t and A_t^* , the dynamic optimization problem given by Equations (1–3) is difficult if not impossible to solve analytically.⁵ However, choosing

$$\Delta \ln(K_{t+1}/N_{t+1}) = \alpha_1 + \alpha_2(\ln(K_t/N_t) - \ln(A_t^*)) \quad (5)$$

and

$$C_t = K_t(1 - \delta_k) + Y_t - K_{t+1} \quad (6)$$

– where α_1 and α_2 are functions of the parameters $(\rho, \tau, \theta, \delta_k)$ – yields a close approximation of optimal consumer behavior. This technique of approximating optimal dynamic behavior has its origins in the real business cycle literature beginning with Kydland and Prescott (1982). A simple derivation of expressions for α_1 and α_2 is given in Appendix A.

2.2.2 Long-term Growth, Climate Behavior and Damages

This section explains the remainder of the state-contingent model, including growth projections, climate behavior and damages from global warming [based primarily on the DICE model, Nordhaus (1994b)]. Net labor productivity, A_t^* , is modeled as the product of exogenous, gross labor productivity, A_t , a factor describing control costs, and a factor describing climate damages:

$$A_t^* = \underbrace{\left(\frac{1 - b_1 \mu_t^{b_2}}{1 + (D_0/9) \cdot T_t^2} \right)}_{\text{damages}}^{-\frac{1}{1-\theta}} A_t \quad (7)$$

μ_t is the fractional reduction in CO₂ emissions at time t (the “control rate”) versus a business as usual/no government policy baseline, while b_1 and b_2 parameterize the cost of attaining these reductions. T_t is the average surface temperature relative to pre-industrialization and

⁵Long and Plosser (1983) derive an analytic solution for the case of $\delta_k = 1$ and $\lim \tau \rightarrow 1$ (log utility). ✱

D_0 is the fractional loss in aggregate GDP from a 3° temperature increase. The form of the control cost function is derived from an econometric analysis of a number of previous studies. The damage function is considerably more ad hoc. See Nordhaus (1993) and Nordhaus (1994a) for additional details on costs and damages, respectively.

Exogenous labor productivity A_t is modeled as a random walk in logarithms with an exponentially decaying drift. That is,

$$\log(A_t) = \log(A_{t-1}) + \gamma_a \exp(-\delta_a t) + \sigma_a \epsilon_t \quad (8)$$

where γ_a is the initial growth rate, δ_a is the annual decline in the growth rate, σ_a is the standard deviation of the random growth shocks and ϵ_t is a standard NIID random shock. Population is modeled similarly but without the stochastic element:

$$\log(N_t) = \log(N_{t-1}) + \gamma_n \exp(-\delta_n t) \quad (9)$$

where γ_n is the initial growth rate and δ_n is the annual decline in the growth rate. Note that these models predict zero growth asymptotically, though this may occur centuries in the future.⁶

The remaining portion of the model explains the link between economic activity (measured as aggregate output Y_t) and warming (measured as the average surface temperature T_t). The first step is linking output to emissions:

$$E_t = \sigma_t(1 - \mu_t)Y_t \left(\frac{A_t}{A_t^*} \right)^{1-\theta} \quad (10)$$

where E_t is emission of CO₂, σ_t is an exogenous trend in emissions/output, and μ_t is the rate of emissions reductions induced by the policymaker. The expression $Y_t \left(\frac{A_t}{A_t^*} \right)^{1-\theta}$ reflects “raw” output prior to reductions due to damages and control costs. The model of σ_t is, as with labor productivity and population, based on exponentially decaying growth:

$$\log(\sigma_t) = \log(\sigma_{t-1}) + \gamma_\sigma \exp(-\delta_\sigma t) \quad (11)$$

⁶For example, the range of parameters used in the simulations (with $\delta_{a/n} \in (0.25\%, 2.5\%)$) leads to a halving of the growth rates every 20 to 200 years.

where γ_σ is the initial growth rate of emissions/output (a negative number) and δ_a is the annual decline in the growth rate. Note that the annual decline in the emissions/output rate is the same as the annual decline in labor productivity growth (δ_a).

Emissions of CO₂ accumulate in the atmosphere according to:

$$M_t - 590 = \beta E_{t-1} + (1 - \delta_m)(M_{t-1} - 590) \quad (12)$$

where M_t is the atmospheric concentration of CO₂ in billions of tons of carbon equivalent. β is a measure of the retention rate of emissions. Low values of β indicate that emissions do not, in fact, accumulate while a value of unity would mean that every ton of emitted CO₂ becomes a ton of atmospheric CO₂. The parameter δ_m plays the role of a depreciation rate: it is assumed that CO₂ in the atmosphere above the pre-industrialization level of 590 billion tons slowly decays. This decay reflects absorption of CO₂ into the oceans which are assumed to be an infinite sink.

Above average concentrations of CO₂ in the atmosphere lead to increased radiative forcings, a measure of the rate of transfer between solar energy produced by the sun and thermal energy stored in the atmosphere. This is modeled according to

$$F_t = 4.1 \times \log(M_t/590)/\log(2) + O_t \quad (13)$$

where F_t measures radiative forcings in units of watts per meter squared. The specification is such that a doubling of CO₂ concentrations leads to a roughly four fold increase in forcings. O_t in this relation represents radiative forcings due to other greenhouse gases and is assumed exogenous to the model:


$$O_t = \begin{cases} 0.2604 + 0.0125t - 0.000034t^2 & \text{if } t < 150 \\ 1.42 & \text{otherwise} \end{cases} \quad (14)$$

Increased forcings lead to temperature changes according to

$$T_t = T_{t-1} + (1/R_1) [F_t - \lambda T_{t-1} - (R_2/\tau_{12})(T_{t-1} - T_{t-1}^*)] \quad (15)$$

$$T_t^* = T_{t-1}^* + (1/R_2)(R_2/\tau_{12})(T_{t-1} - T_{t-1}^*) \quad (16)$$

Figure 1: State-contingent model



where T_t is the surface temperature and T_t^* is the deep ocean temperature, both expressed in changes relative to pre-industrialization levels in degrees Celsius. Note that if $M_t = 590$ and $O_t = 0$ (e.g., pre-industrialization), T_t and T_t^* will equilibrate to zero. The parameter λ describes the equilibrium change in surface temperature for a given change in radiative forcings. In particular, based on (13) and (15), a doubling of the concentration of CO_2 in the atmosphere will lead to a $4.1/\lambda$ rise in surface temperature in the long run. This parameter $4.1/\lambda$ is a measure of the temperature sensitivity of the atmosphere.⁷

The parameters R_1 , R_2 and τ_{12} describe the thermal capacity of the surface atmosphere and deep oceans and the rate of energy transfer between them, respectively.

2.2.3 Summary of State-contingent Climate-Economy Model

The state-contingent model is now complete. Figure 1 shows schematically how the model evolves over time and how policy (a path of reductions $\{\mu_t\}$) is translated into a state-contingent utility measure (as measured in Equation (1)).

At any point in time, the vector $\{K_t, A_t, N_t, \sigma_t, M_t, T_t, T_t^*\}$ completely describes the state of the economy and climate. The three somewhat distinct portions of the model – economy,

⁷ λ by itself is referred to as the climate feedback parameter.

climate and trends – are then used to evolve the current state into the next period state. The economy and climate models make use of the policy variable μ_t describing emission reductions in order to compute control costs and realized emissions, respectively. The realized path of consumption is an output of the economic model which is then used to evaluate the utility associated with implementing the policy in this state.

A key innovation relative to similar models is the use of Equation (5) to capture optimal intertemporal consumer behavior without performing a numerical optimization. This reduces by an order of magnitude the time required to evaluate policy while preserving the proper intertemporal prices.

2.3 Quantifying Uncertainty

Estimates of uncertainty in the model come from two sources: econometric analysis and subjective assessment. The model involves nineteen different parameters. Six are parameters describing observable economic activity:

- pure time preference ρ ,
- risk aversion τ ,
- output-capital elasticity θ ,
- productivity growth γ_a ,
- variation in productivity growth σ_a^2 , and
- depreciation δ_k .

A joint distribution for these parameters is estimated with historical data. The remaining thirteen describe emissions:

- emissions rate growth γ_σ ,

climate change:

- CO₂ retention rate β ,
- temperature sensitivity $4.1/\lambda$,
- CO₂ decay rate δ_m ,
- thermal capacities and conductivities R_1 , R_2 and τ_{12} ,

control costs and damages:

- cost function parameters b_1 and b_2 ,

- fractional loss of GDP for 3° temperature rise D_0 ,

and long-term growth trends

- population growth γ_n ,
- productivity slowdown δ_a ,
- population slowdown δ_n .

Uncertainty about these parameters is based on subjective analysis.

The econometric analysis of the six economic parameters is based on post-war U.S. data as described Appendix B. Series describing aggregate investment, capital services, output and prices are fit to the model described by Equations (2),(3) and (5). The posterior parameter distribution which arises from this analysis is summarized in Table 1.

Nordhaus (1994b)⁸ develops a distribution for the remaining parameters based on a two-step subjective analysis. The first step involves testing his model's sensitivity to each parameter being changed, one at a time, to a more extreme value. Those parameters which produce the largest variance in model output are then further scrutinized. A discrete, five-value distribution is developed for seven of these thirteen variables. The remaining six are fixed at their best guess values. The distribution of the seven uncertain parameters is summarized in Table 2. Values of the six fixed parameters as well as initial conditions for the model are given in Table 3.

2.4 Social Welfare

Once we concede that preferences are imprecisely known, the next logical step is to formulate a social welfare function to aggregate over these preferences. This social welfare function then allows the policymaker to compare alternatives in the face of such uncertainty. While it is more common to envision aggregation over different individuals, here different states of nature are being aggregated, each with a single but unique representative agent. It is important to recognize that although the parameters of a representative agent model can be inferred from observed consumer behavior, there is no information available to infer

⁸Chapters 6 and 7.

Table 1: Marginal distributions of uncertain economic parameters
(narrow bars indicate values used in simulations without uncertainty)







description	symbol	equation	distribution
pure rate of time preference	ρ	(1)	
coefficient of risk aversion	τ	(1)	
output-capital elasticity	θ	(2)	
rate of capital depreciation	δ_k	(3)	
initial productivity growth rate	γ_a	(8)	
standard error of productivity shocks	σ_a	(8)	

Table 2: Discrete distributions of uncertain climate/trend parameters
(narrow bars indicate values used in simulations without uncertainty)








description	symbol	equation	distribution
annual decline of population growth rate	δ_n	(9)	
annual decline of productivity growth rate	δ_a	(8),(11)	
initial growth rate of CO2 per unit output	γ_σ	(10),(11)	
damage parameter (% loss of GDP for 3° temperature rise)	D_0	(7)	
cost function parameter	b_1	(7)	
retention rate for CO2 emissions	β	(12)	
temperature sensitivity to CO2 doubling (in °C)	$4.1/\lambda$	(13),(15)	

Table 3: Description of fixed parameters^a

parameter	symbol	equation	units	value
cost function parameter	b_2	(7)		2.887
decay rate of atmospheric CO ₂	δ_m	(12)		0.00833
1/thermal capacity of atmosphere	$1/R_1$	(15)	°C-meter ² /watt-year	0.048
thermal conductivity b/w atmosphere and oceans	R_2/τ_{12}	(15),(16)	watt/°C-meter ²	0.44
1/thermal capacity of deep oceans	$1/R_2$	(16)	°C-meter ² /watt-year	$\frac{0.002}{0.44}$
1995 global population	N_0	(9)	millions of people	5590
initial population growth	γ_n	(9)		0.0124
initial rate of CO ₂ emissions per unit of output	σ_0	(11)	billion tons CO ₂ per \$1989 trillions	0.385
1995 global capital stock	K_0	(2),(3)	\$1989 trillions	79.5
1995 global output	Y_0	(2),(3),(10)	\$1989 trillions	24.0
1995 atmospheric concentrations of CO ₂	M_0	(12)	billions of tons of C equivalent	763.6
1995 surface temperature ^b	T_0	(7),(15),(16)	°Celsius	0.763
1995 deep ocean temperature ^b	T_0^*	(15),(16)	°Celsius	0.117

^aAll fixed parameters are from Nordhaus (1994b). The parameters that do not depend on time are from Nordhaus' Table 2.4. Initial values for temperature, CO₂ concentrations, and output in 1995, as well as the initial annual growth rate for population, are based on the Nordhaus base case simulation. The 1995 capital stock is adjusted upward to reflect differences in the definition of capital as well as underlying parameter values. The decay rate of atmospheric CO₂ is divided by ten to convert from a decennial to annual rate. The annual thermal capacity of the ocean and atmosphere are from the second line of Nordhaus' Table 3.4b.

^bTemperatures are measured as deviations from the pre-industrialization level, circa 1900.

parameters of the social welfare function. Instead, we must rely on social choice theory and our own sense of fairness.

The common approach in the climate change literature gets around the issue of uncertain preferences by reporting a range of policy prescriptions based on a range of possible preferences and states of nature. For example, Cline (1992) presents benefit-cost analyses for 92 different cases (Tables 7.3 and 7.4). Dowlatabadi and Morgan (1993) integrate out much of the uncertainty in their analysis, but still present results for 48 scenarios. Nordhaus (1994b) gives one of the few examples where *all* uncertainty including preferences is integrated out to yield a single policy recommendation (Chapter 8). However, the social welfare function implicit in the analysis is not revealed.

As most authors observe, uncertainty about time preference tends to dominate many policy considerations. Moreso, in fact, than uncertainty about climate sensitivity and damages. It is therefore important to consider how to aggregate over such preferences rather than guessing.

In this analysis it is assumed that social welfare is an average of rescaled utility measured in each state of nature (given by Equation (1)). The rescaling serves to equate the marginal social welfare of one additional dollar of current consumption in all states of nature. While arbitrary, some kind of adjustment is necessary to prevent the resulting policy prescription from being sensitive to the choice of units in the model (e.g., using dollars rather than thousands of dollars as the unit of account; see explanation in Appendix C). Social welfare can then be written as

$$SW(x) = I^{-1} \sum_{i=1}^I u(x, i) \quad (17)$$

where $u(x, i)$ is rescaled utility in state i with outcome x and $SW(x)$ is the social welfare associated with x . The rescaling is such that

$$u(+\$1 \text{ p.c. in initial period}, i) - u(\emptyset, i) = u(+\$1 \text{ p.c. in initial period}, j) - u(\emptyset, j) \quad \forall i, j$$

That is, a policy corresponding to an extra dollar of per capita in the initial period is assumed

to have the same utility in every state relative to a no policy ($\emptyset$) baseline.

This social welfare function has its origins in the literature on social choice. Harsanyi (1977) shows that in defining social welfare over lotteries, if individual preferences satisfy the von Neumann Morgenstern axioms then social welfare must have this weighted average form. Otherwise, social preferences will fail to mimic individual preferences over lotteries involving only that individual. This functional form can also be derived from the assumption of cardinal unit comparability, as discussed by Roberts (1980). A more general functional form therefore requires both integrating out uncertainty from the representative agent's perspective (to satisfy Harsanyi's point) and more stringent assumptions about the level of utility comparability (to satisfy Robert's point).


3 Does Uncertainty Matter?

To understand whether uncertainty significantly affects decision making and welfare consequences, two optimizations are performed. The first looks for the optimal CO₂ reductions based on the model developed in the previous section and the parameter distributions given in Tables 1 and 2. The second assumes there is no uncertainty. All uncertain parameters are fixed at their median values and the variance of the random exogenous growth shocks is set to zero (as indicated by the narrow bars in Tables 1 and 2). There is only one state of nature.

3.1 Policy Stringency


The resulting policies are shown in Figure 2. The initial difference between policies is small, 7.9% ignoring uncertainty versus 8.6% with uncertainty included. In thirty years these differences are higher, 10.4% versus 12.8% in 2025. By the end of the next century, however, the rate is almost doubled when uncertainty is considered. Uncertainty clearly has a dramatic effect on the future path of the optimal control rate, but with a relatively small consequence for immediate decisions.

Figure 2: Optimal climate change policy with and without uncertainty



What explains this pattern of small policy consequences now and large ones later? Surprisingly, much of the consequences (roughly 50%) can be traced to uncertainty about future interest rates, rather than climate or other economic uncertainty. This can be seen by reoptimizing the policy for a model with only pure time preference ρ and risk aversion τ uncertain, as illustrated in Figure 3. In contrast, a reoptimized policy based on modeling only the climate parameters (γ_σ , D_0 , b_1 , β , and $4.1/\lambda$) as uncertain yields essentially the same results as completely ignoring uncertainty.

Intuitively, Figure 3 shows the consequences of non-linearities in the first order conditions generating the optimal policy. If these first order conditions were linear in the uncertain parameters, replacing uncertain parameters with their means would have little effect. As it turns out, these conditions are highly non-linear functions of the interest rate but fairly linear in terms of the climate variables. Recall, for example, that the value of an annuity with coupon P and interest rate i is P/i – a highly non-linear function near $i = 0$. Furthermore, the interest rate is relatively well known during the initial periods of the simulation but less precisely known – and possibly near zero – in the future. Thus the policy effects of uncertainty are higher in the future.



Why does uncertainty about the interest rate rise in the future? The historic interest rate is an observed quantity and therefore easily predicted over short horizons. More generally, it depends in a simple way on future growth, risk aversion and time preference via the Euler equation: $i = \rho + \tau \cdot g_{a,t}$. Here i is the interest rate, ρ the pure rate of time preference, τ the coefficient of relative risk aversion and $g_{a,t}$ the rate of future productivity growth at time t . Future uncertainty arises because we assume productivity growth $g_{a,t}$ eventually declines to zero (according to Equation (8)) leaving only pure time preference ρ to motivate non-zero interest rates. While historical data tells us that with growth of 1.3%, interest rates will be 5%, we have much less information about how to dissect the observed 5% into its component pieces ρ and $\tau \cdot 1.3\%$. Therefore, we know less about what the interest rate will be once growth declines.

3.2 Welfare

In spite of the differences depicted in Figure 2, the welfare consequences of the two policies remain unclear. Suppose we analyze both policies using the model in Section 2, reflecting the uncertainty summarized in Tables 1 and 2. Does the policy which ignores uncertainty leave

us substantially worse off, given that uncertainty is in fact present? To answer this question, it is necessary to convert the social welfare measure used to determine these optimal policies in Equation (17) into a more meaningful metric. A natural vehicle for comparison is the certainty equivalent change in first period consumption yielding the same welfare consequences as the given policy. For example, implementing the policy which ignores uncertainty improves social welfare by 0.32 units based on Equation (17).⁹ A similar improvement in social welfare can be achieved by exogenously raising per capita consumption by \$58 in the first period in every state of nature.

The alternative policy which optimizes over uncertainty leads to a certainty equivalent gain of \$73 – a 25% improvement. The source of this improvement tends to be those states with low costs ($b_1 \downarrow$), high damages ($D_0 \uparrow$), high temperature sensitivity ($4.1/\lambda \uparrow$), low pure rates of time preference ($\rho \downarrow$), rapid productivity slowdown ($\delta_a \uparrow$), but high population growth ($\delta_n \downarrow$). In these states, the marginal costs to reduce emissions are low and the marginal benefits are high. Therefore, a more stringent policy will be preferred. While opposing cases (e.g., high costs, low damages) favor less stringent policy, such cases tend not to be as extreme. This is the non-linearity discussed earlier. In other words, once uncertainty is considered bad states of the world tend to be more extreme than good states and motivate additional stringency.

4 Instrument Choice

Uncertainty about the cost of reducing emissions leads to alternative consequences for quantity and price instruments. This suggests that welfare gains might be achieved by using taxes rather than rate controls to reduce emissions. The obvious experiment is to optimize

⁹It turns out that computing the welfare gain is complicated by extremely fat tails in the distribution of state-contingent welfare gains. Even with 50,000 draws, the mean welfare gain had not yet converged to its asymptotic normal distribution (a sample of 38 such means rejected normality at the 10^{-5} level). The somewhat ad hoc solution used here censors all gains higher than +150 and lower than -12 utility units. This corresponds to censoring at -\$1,200 and +\$11 million in terms of certainty equivalent gains and affects 0.1% of the states. This approach was chosen so that the reported results, if anything, are understated.

over a tax policy and then compare the resulting welfare gains.

Weitzman (1974) first demonstrated the dichotomy between price and quantity controls with uncertain costs in a simple partial equilibrium model. This model incorporates dynamic general equilibrium, leading to a more complex expression for the marginal cost of reduced emissions in period t :

$$MC_t = \frac{b_1 \cdot b_2}{\sigma_t \cdot (1 + D_0/9 \cdot T_t^2)} \mu_t^{b_2-1} \quad (18)$$

where b_1 and b_2 are cost function parameters, D_0 is the damage from a 3° temperature rise, T_t is the temperature rise since industrialization, μ_t is the control rate, and σ_t is the rate of emissions per unit of output. Since b_1 , σ_t , D_0 , and T_t vary among states of nature, the marginal cost will as well. Optimizing agents who set marginal cost equal to an imposed tax rate in each period will choose different control rates depending on the parameter values in their particular state of nature. Since a rate control would fix the emission reduction rate across all states, there is no rate control which can reproduce the consequences of a tax instrument in all states.

4.1 Taxes versus rate controls

In order to implement a tax instrument in the model, the control rate in each state of nature and time period is replaced with an expression for μ_t derived from (18).¹⁰ MC in this expression is replaced with the chosen tax rate. The optimal tax path is then determined by maximizing the social welfare expression in (17).

The tax instrument does in fact lead to a higher welfare gain compared to the rate instrument. When converted to a certainty equivalent change in first period consumption, the gain is \$86 per person for the tax. This is a 20% improvement over the rate instrument, which yields only \$73. Note that compared to the rate instrument computed in the absence of uncertainty, this represents a 50% overall gain. That is, a policymaker who ignores

¹⁰Derived rates higher than 100% are set to 100%, indicating that the tax induces complete elimination of emissions


uncertainty and implements the optimal rate policy would improve social welfare by \$58 per person (where the measured gain is based on the assumption that uncertainty does, in fact, exist). He or she would not have any reason to consider tax instruments since they are equivalent in the absence of uncertainty. Compare this outcome to that of a second policymaker who considers uncertainty and computes both optimal tax and control rate policies: This policymaker would see that taxes perform better and, in implementing the optimal tax policy, would improve social welfare by \$86 per person. This represents a net improvement of nearly 50% over the first policymaker's outcome.

Why are taxes preferred to rate controls? Weitzman's original insight was that this would be true when the elasticity of reductions demand (e.g., benefits) exceeded the elasticity of reductions supply (e.g., costs). The left panel of Figure 4 illustrates this intuition: when the cost curve shifts (dashed lines), the optimal shadow price (intersection of costs and expected benefits) is relatively constant compared to the control rate. This is only part of the reason, however. Weitzman's analysis was based on constant elasticities and, importantly, independent shocks to supply and demand. Here, high benefits occur when productivity growth declines rapidly ($\delta_a \uparrow$), in turn permitting low rates of pure time preference to generate high expected benefits more quickly, as discussed in Section 3.1. But a high value of δ_a , by the definition of the emissions rate (11), also leads to higher values of σ_t and low marginal control costs by (18) (recall that growth in σ_t is initially negative). Hence upward shifts in the benefits curve are correlated with downward shifts in the cost curve. Examining the right panel of Figure 4, we see that this also leads to a preference for taxes since the optimal shadow price is again less variable than the optimal rate.

4.2 Value of emission rights

Another way to view the dichotomy between taxes and rate controls is in terms of the value of emission rights. Based on a tax instrument, the value of emission rights is held constant across states at the given tax rate. Under a rate or quantity control, however, emission

Figure 4: Intuition for taxes being preferred to rate instruments
(dashed lines indicate particular realizations of uncertain costs and benefits)




Original Weitzman (1974) result that relatively flatter (higher elasticity) marginal benefits favor taxes. Note that all three intersections correspond to similar shadow prices but different control rates. Hence the deadweight loss will be lower with a fixed shadow price rather than a fixed control rate.

Correlation of upward shifts in benefits with downward shifts in costs favors taxes. Note that both intersections lead to similar shadow prices. As in the left panel, the deadweight loss is lower with the shadow price fixed (via a tax) rather than the control rate fixed.

rights will have different values in different states. Using the results of many simulations, it is possible to plot the distribution of resulting values and obtain a visual sense of the difference between instruments.

Figure 5 shows the median, 50% and 95% forecast intervals for the value of emission rights under the optimal rate control, along with the fixed values under the optimal tax policy (the optimal tax policy which ignores uncertainty is presented for comparison). While the median case under the optimal rate control closely tracks the optimal tax policy, the 95% forecast interval includes widely divergent values. For example, the median value of a right to emit one ton of carbon dioxide rises to around \$20 after 50 years. It is entirely possible, however, that under an optimal rate control the value will be over \$100. With these magnitudes of difference between the two instruments in terms of marginal costs, the distinction in welfare consequences discussed in the previous section is no longer surprising.

Figure 5: Value of emission rights under alternative policies



5 Conclusion

This paper has sought to inject discussions of climate change policy with a dose of skepticism regarding the limitations of models which ignore uncertainty. Excluding uncertainty tends to reduce expected benefits due to highly non-linear relations leading to policy recommendations which are too lax. Much of this effect can be related to agent preferences; over half the noted consequences can be replicated with preferences as the sole source of uncertainty. This raises the issue of how to aggregate over preferences, a question which is at once both controversial but necessary. Finally, uncertainty about costs introduces a dichotomy between price and quantity controls. In the case of carbon dioxide emissions, taxes perform better. The welfare gain associated with the optimal tax instrument is \$86 per person whereas the gain associated with the optimal rate instrument is only \$73. An optimal rate policy computed in the absence of uncertainty has a gain of only \$58.

While the results of this analysis should be taken with many of the same caveats accorded other integrated assessment models (the dependence on stylized and simplified models, parameterization, etc.), there are important distinctions to be made. This is the first dynamic general equilibrium model to incorporate large numbers of uncertain states of nature. It is

also takes advantage of econometrically quantified uncertainty. Lastly, the model explicitly addresses the issue of aggregation over uncertain preferences.

The most important consequence of this paper may be to encourage a reassessment of the current agenda for climate change research. Considerable effort has been directed towards the importance of learning; perhaps more effort should be directed at the consequences of uncertainty itself – without learning. A great deal of attention has focused on climate uncertainty; this paper suggests that considerable uncertainty surrounds the economy. In particular, the course of productivity growth has an important bearing on both future emissions and the valuation of future benefits via the interest rate. Finally, much of the policy discussion seems to have shifted towards quantity controls as the preferred instrument; this paper suggests that efficiency arguments may point the other way, towards taxes.

Appendix

A Notes on Solution Algorithm

This section explains the solution to the consumer optimization problem (1) given by Equation (5). In words, (5) is a log-linear approximation of the optimal decision rule based on the deterministic steady-state level of capital per efficiency unit of labor, e.g., $K_t/(A_t^* N_t)$. This approach turns out to work rather well for three reasons: First, the economy starts off relatively close to the steady state defined by the initial growth rates. Second, the decline in deterministic growth rates, described in Equations (8–9), has little effect on the approximation parameters. Third, the shocks away from the steady state – from both exogenous stochastic productivity shocks as well as endogenous climate change effects – are relatively small. While there are many ways of deriving expressions for α_1 and α_2 in terms of the underlying parameters, the simplest is based on a logarithmic approximation to the saddle path around the steady state in $C_t/(A_t^* N_t)$ and $K_t/(A_t^* N_t)$ space.

We begin with the Euler equation arising in a balanced growth equilibrium (where $\gamma_a =$

$\log(A_t^*) - \log(A_{t-1}^*)$ and $\gamma_n = \log(N_t) - \log(N_{t-1})$ are constant $\forall t$:¹¹

$$\left(\frac{C_{t+1}/N_{t+1}}{C_t/N_t} \right)^\tau = (1 + \rho)^{-1} (\theta(K_{t+1}/A_{t+1}^* N_{t+1})^{\theta-1} + 1 - \delta_k) \quad (\text{A.1})$$

At the balanced growth equilibrium $\frac{C_{t+1}/N_{t+1}}{C_t/N_t} = \exp(\gamma_a)$ so (A.1) can be solved for steady state capital stock per efficiency unit of labor, $K_{t+1}/(A_{t+1}^* N_{t+1}) = \tilde{K}^\dagger = \left(\frac{(1+\rho)\exp(\gamma_a\tau) - (1-\delta_k)}{\theta} \right)^{\frac{1}{\theta-1}}$.

The capital accumulation equation (3),

$$K_{t+1} = K_t(1 - \delta_k) + (A^* N)^{1-\theta} K_t^\theta - C_t \quad (\text{A.2})$$

then yields steady state consumption per efficiency unit of labor $C_t/(A_t^* N_t) = \tilde{C}^\dagger = \tilde{K}^{\dagger\theta} + (1 - \delta_k - \exp(\gamma_a + \gamma_n))\tilde{K}^\dagger$.

The approximation (5) is made by taking total derivatives of both (A.1–A.2) with respect to $\log(\tilde{K}_{t+1})$, $\log(\tilde{K}_t)$, $\log(\tilde{C}_{t+1})$ and $\log(\tilde{C}_t)$ around the steady-state values $\tilde{K}_t = \tilde{K}_{t+1} = \tilde{K}^\dagger$ and $\tilde{C}_t = \tilde{C}_{t+1} = \tilde{C}^\dagger$. This leads to the expressions

$$\tau(dc_{t+1} - dc_t) = \kappa_{ck} dk_{t+1} \quad (\text{A.3})$$

$$\exp(\gamma_a + \gamma_n)dk_{t+1} = \kappa_{kk} dk_t + \kappa_{kc} dc_t \quad (\text{A.4})$$

where $dc_t = \log(\tilde{C}_t) - \log(\tilde{C}^\dagger)$ is the log deviation of $\tilde{C}_t$ and $dk_t = \log(\tilde{K}_t) - \log(\tilde{K}^\dagger)$ is the log deviation of $\tilde{K}_t$. κ_{ck} , κ_{kk} and κ_{kc} are constants, the derivatives of the log of the right hand side of (A.1) with respect to $\log(\tilde{K}_{t+1})$ and of the log of the right hand side of (A.2) with respect to $\log(\tilde{K}_t)$ and $\log(\tilde{C}_t)$.¹² Subtracting (A.4) from the same expression evaluated one period forward yields

$$\exp(\gamma_a + \gamma_n)(dk_{t+2} - dk_{t+1}) = \kappa_{kk}(dk_{t+1} - dk_t) + \kappa_{kc}(dc_{t+1} - dc_t) \quad (\text{A.5})$$

Equation (A.3) can then be used to substitute for $dc_{t+1} - dc_t$, leaving a second order difference equation in dk_t :

$$\exp(\gamma_a + \gamma_n)(dk_{t+2} - dk_{t+1}) = \kappa_{kk}(dk_{t+1} - dk_t) + \kappa_{kc}(\kappa_{ck}/\tau) dk_{t+1} \quad (\text{A.6})$$

¹¹The Euler equation can be derived from Equation (1) by ignoring random productivity shocks and examining a perturbation which decreases consumption in period t , invests the extra output, then consumes the gross return (interest and principal) in period $t + 1$. If the path of consumption is indeed optimal, this should have no effect on utility and leads to the condition given in (A.1).

¹²Note that $\kappa_{ck} = (1 + \rho)^{-1}(\theta - 1)\theta\tilde{K}^{\dagger\theta-1}\exp(-\gamma_a\tau)$, $\kappa_{kk} = (1 + \rho)\exp(\gamma_a\tau)$, and $\kappa_{kc} = -\tilde{C}^\dagger/\tilde{K}^\dagger$.

model:

$$\Delta a_t^* = \gamma_a + \epsilon_t \quad (\text{B.11})$$

$$\Delta k_t = (\alpha_1 - \alpha_2 a_0^*) + \alpha_2(k_{t-1} - (a_{t-1}^* - a_0^*)) + \eta_{2t} \quad (\text{B.12})$$

$$y_t = (1 - \theta)a_0^* + \theta k_t + (1 - \theta)(a_t^* - a_0^*) + \eta_{1t} \quad (\text{B.13})$$

$$\frac{K_{t+1} - (Y_t - C_t)}{K_t} = 1 - \delta_k + \eta_{3t} \quad (\text{B.14})$$

$$\frac{P_{K,t} K_t}{P_{Y,t} Y_t} = \theta + \nu_t \quad (\text{B.15})$$

$$\nu_t = r_\nu \nu_{t-1} + \eta_{4t} \quad (\text{B.16})$$

Here $a_t^* = \log(A_t^*)$, $k_t = \log(K_t/N_t)$, $y_t = \log(Y_t/N_t)$, $P_{K,t}$ is the price of capital services and $P_{Y,t}$ is the price of output; the remaining variables are as defined in Section 2.¹³ The disturbances η_t and ϵ_t are assumed to be jointly NIID with zero covariance (note that the capital share is defined with an autoregressive error). This specification defines a likelihood function describing the distribution of data and unobserved productivity shocks conditional on the model parameters.

This likelihood function coupled with a prior distribution over the parameters can be used to obtain a function describing the posterior density of the parameters and productivity shocks conditional on the observed data.¹⁴ This is Bayes' rule. Obtaining a sample of draws from the posterior density is difficult, however, as the distribution is non-standard. To generate the distributions shown in Table 1, the Gibbs sampler was used.¹⁵ Rather than drawing combinations of all the parameters at once, the Gibbs sampler draws each unobserved parameter sequentially based on its distribution conditional on the previous draw of all other parameters. A chain of draws is formed which converges to the joint posterior

¹³The data, for the years 1952-1992, is from the U.S. Worksheets [Jorgenson (1980); see Ho (1989) and Wilcoxon (1988) for further details]. Output is a divisia index of output of consumption goods (c790+c781), investment goods (c789) and government services (c943). Capital is aggregate capital services (c667). Changes in relative prices among output goods (c792, c761, c791, c395) are subsumed into quantity changes for simplicity. Population is 16+ (c941).

¹⁴The Jeffrey's prior (see Section 2.8 of Gelman, Carlin, Stern, and Rubin (1995)) is used for each parameter, restricted to the economically relevant parameter space (e.g., $0 < \theta < 1$, $\delta_k > 0$, $\alpha_2 < 0$).

¹⁵Geman and Geman (1984), Tanner and Wong (1987), Jacquier, Polson, and Rossi (1994).

From (5) and the observation that when $\tilde{K}_t = \tilde{K}^\dagger$ the growth of capital stock per capita is the steady state rate of γ_a , so $\Delta k_t = \gamma_a$, we know that $\gamma_a = \alpha_1 + \alpha_2 \log(\tilde{K}^\dagger)$. This leads to

$$\alpha_1 = \gamma_a - \alpha_2 \log(\tilde{K}^\dagger) \quad (\text{A.7})$$

and suggests a guess for solution to (A.6):

$$dk_{t+1} = (\alpha_2 + 1)dk_t \quad (\text{A.8})$$

(since $dk_{t+1} - dk_t = \Delta k_{t+1} - \gamma_a = \alpha_2 dk_t$). Then (A.6) yields the characteristic equation

$$\exp(\gamma_a + \gamma_n) [(\alpha_2 + 1)^2 - (\alpha_2 + 1)] = \kappa_{kk}((\alpha_2 + 1) - 1) + \kappa_{kc}(\kappa_{ck}/\tau)(\alpha_2 + 1) \quad (\text{A.9})$$

which can be used to find α_2 in terms of κ_{kk} , κ_{kc} and κ_{ck} :

$$(\alpha_2 + 1) = \frac{e^{\gamma_a + \gamma_n} + \kappa_{kk} + \kappa_{kc}(\kappa_{ck}/\tau) - \sqrt{(e^{\gamma_a + \gamma_n} + \kappa_{kk} + \kappa_{kc}(\kappa_{ck}/\tau))^2 - 4e^{\gamma_a + \gamma_n}\kappa_{kk}}}{2e^{\gamma_a + \gamma_n}} \quad (\text{A.10})$$

For $\rho = 0.04$, $\tau = 1.2$, $\theta = 0.38$, $\delta_k = 0.05$, $\gamma_a = 0.013$ and $\gamma_n = 0.012$ we have $\tilde{K}^\dagger = 7.8$, $C^\dagger = 1.6$, $\kappa_{ck} = -0.062$, $\kappa_{kk} = 1.06$ and $\kappa_{kc} = -0.20$. The yields $\alpha_2 = -0.084$ and $\alpha_1 = 0.185$.

The fact that we can find a value of α_2 which satisfies (A.6) validates our initial guess. Namely, a linear rule for Δk_{t+1} in terms of $k_t - a_t^*$ as in (5) or equivalently (A.8).

B Notes on Econometrics

The model given by Equations (2), (3) and (5) can be coupled with the price implications of market equilibrium (Sheperd's Lemma) and a simplified model of productivity shocks (ignoring the future slowdown and climate consequences) to yield the following econometric

distribution.

The distribution used in this paper is based on ten chains of 1500 draws with the first 500 draws dropped to reduce start-up effects. The sample was further reduced by about 10% by requiring that the reduced form parameters α_1 and α_2 map back into meaningful structural parameters ρ and τ .

C Notes on Utility Rescaling

Policy consequences in a given state are always expressible in terms of a first period consumption equivalent. Letting $C_{i,1}(x)$ be this consumption equivalent for state i and policy x , a social welfare function which simply averages utility yields

$$SW(x) = I^{-1} \sum_i \frac{(C_{i,1}(x)/N_1)^{1-\tau_i}}{1-\tau_i}$$

where N_1 is the first period population (known with certainty) and τ_i is the coefficient of relative risk aversion in state i . But scaling the units of consumption by a factor of ψ changes social the social welfare associated with policy x to:

$$SW(x) = I^{-1} \sum_i \psi^{1-\tau_i} \frac{(C_{i,1}(x)/N_1)^{1-\tau_i}}{1-\tau_i}$$

Unless the consequences are the same for each state, this simple change in units – which leads to an unintentional reweighting among states – will change the ordering among policies. The proposed solution takes the form

$$SW(x) = I^{-1} \sum_i (C_{i,1}(\emptyset)/N_1)^{\tau_i} \frac{(C_{i,1}(x)/N_1)^{1-\tau_i}}{1-\tau_i}$$

where $C_{i,1}(\emptyset)$ is the consumption level in the absense of policy and

$$u(x, i) = (C_{i,1}(\emptyset)/N_1)^{\tau_i} \frac{(C_{i,1}(x)/N_1)^{1-\tau_i}}{1-\tau_i}$$

is the rescaled utility measure in (17). Note that renormalizing the units of consumption no longer affects the policy ordering.

References

- Cline, W. (1992). *The Economics of Global Warming*. Washington: Institute of International Economics.
- Dowlatabadi, H. and M. Morgan (1993). A model framework for integrated studies of the climate problem. *Energy Policy* 21(3), 209–221.
- Gelman, A., J. B. Carlin, H. S. Stern, and D. B. Rubin (1995). *Bayesian Data Analysis*. New York: Chapman & Hall.
- Geman, S. and D. Geman (1984). Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 6, 721–741.
- Harsanyi, J. C. (1977). *Rational Behavior and Bargaining Equilibrium in Games and Social Situations*. New York: Cambridge University Press.
- Ho, M. S. (1989). *The effects of external linkages on U.S. economic growth: A dynamic general equilibrium analysis*. Ph. D. thesis, Harvard University.
- Jacquier, E., N. G. Polson, and P. E. Rossi (1994). Bayesian analysis of stochastic volatility models. *Journal of Business and Economic Statistics* 12(4), 371–389.
- Jorgenson, D. W. (1980). Accounting for capital. In G. M. von Furstenberg (Ed.), *Capital, Efficiency, and Growth*, pp. 251–320. Cambridge: Ballinger.
- Kydland, F. E. and E. C. Prescott (1982). Time to build and aggregate fluctuations. *Econometrica* 50(6), 1345–1370.
- Long, J. B. and C. I. Plosser (1983). Real business cycles. *Journal of Political Economy* 91, 39–69.
- Manne, A. and R. Richels (1995). The greenhouse debate: Economic efficiency, burden sharing and hedging strategies. *The Energy Journal* 16(4), 1–37.
- Manne, A. S. (1996, April). Hedging strategies for global carbon dioxide abatement: A summary of poll results. Energy Modeling Forum 14 Subgroup – Analysis for Decisions Under Uncertainty.
- Manne, A. S. and R. G. Richels (1992). *Buying Greenhouse Insurance*. Cambridge: MIT Press.
- Nordhaus, W. D. (1993). The cost of slowing climate change. *Energy Journal* 12(1), 37–65.
- Nordhaus, W. D. (1994a). Expert opinion on climate change. *American Scientist* 82(1), 45–51.
- Nordhaus, W. D. (1994b). *Managing the Global Commons*. Cambridge: MIT Press.
- Roberts, K. (1980). Interpersonal comparability and social choice theory. *Review of Economic Studies* 47, 421–439.
- Tanner, M. and W. Wong (1987). The calculation of posterior distributions by data augmentation. *Journal of the American Statistical Association* 82, 528–550.

- Taylor, J. B. and H. Uhlig (1990). Solving nonlinear stochastic growth models: A comparison of alternative solution methods. *Journal of Business and Economic Statistics* 8(1), 1-17.
- Weitzman, M. L. (1974). Prices vs. quantities. *Review of Economic Studies* 41(4), 477-491.
- Wilcoxon, P. J. (1988). *The effects of environmental regulation and energy prices on U.S. economic performance*. Ph. D. thesis, Harvard University.

What levels in each year, 2010 to 2050? (Fig 2 is reduction) P17 diff

EXECUTIVE OFFICE OF THE PRESIDENT
COUNCIL OF ECONOMIC ADVISERS

(esp. rel. to 1990)
[+ convert. index of carbon, 1990 in Fig 5]
Jeff
P17
"1990 in 2010
2020
peak in 2015 2040
peak in 2050
optimal path

This is a paper by the person that I invited to come in for lunch or breakfast sometime labor day week. The presentation would actually be an extension/update of this which is described in this RFF press release. Let me know when you would like to schedule

this or if you have
any questions. You can
also e-mail him (pizer@
rff.org) or call .

Sarah