Policy Implications of Stochastic Learning Using a Modified PAGE2002 Model

Stephan Alberth and Chris Hope

April 2006

CWPE 0633 *and* EPRG 0614

Policy Implications of Stochastic Learning using a modified PAGE2002 Model

Stephan Alberth[1] **& Chris Hope**[2]
The Judge Business School, University of Cambridge

Abstract. We consider the importance of Endogenous Technical Change (ETC) on the risk profiles for different abatement strategies using the PAGE2002 model with ETC. Three outcomes from this modelling research have significant impacts on the way we 'optimise' the greenhouse gas abatement path. Firstly, it was found that for most standard abatement paths there would be an initial "learning investment" required that would substantially reduce the unit costs of CO2 abatement as compared to a business as usual scenario. Secondly, optimising an abatement program where ETC has been included can lead to an increased risk profile during the time of widespread CO2 abatements due to the costs associated with learning. Finally, the inclusion of ETC leads to a slightly deferred optimised abatement path followed by a drastic abatement program that itself would seem highly impractical. Together, the results draw attention to the possibilities of uncovering uncertainty through proactive abatements.

Keywords: Endogenous Technical Change, Optimal Abatement, Climate Change

JEL Classification: O13 , Q55, Q56

---

[1] This article is based on research leading to an MPhil in Management Studies at the Judge Business School, University of Cambridge. Corresponding author, s.alberth@jbs.cam.ac.uk. Submitted to EPRG November 2005.
[2] Senior Lecturer at the Judge Business School, University of Cambridge. c.hope@jbs.cam.ac.uk

Table of Contents

Chapter 1. Introduction

Research into climate change is an important yet difficult undertaking. The importance of the general area of research lies in being able to understand and help to avoid the severe impacts that anthropogenic emissions may have on the climate and as a result on our world's environment and people. However the task at hand is extremely complex due to the unique characteristics of global warming. It is a problem that is "global, long-term (up to several centuries), and involves complex interactions between climatic, environmental, economic, political, institutional, social and technological processes" (IPCC, 2001c). To tackle these issues there are a large variety of models that have been designed with the aim of understanding and combining these processes. They are also sometimes referred to as E3 models representing three of the key areas affected, *Energy*, *Environment* and the *Economy* with high levels of uncertainties existing in all of these areas.

This paper examines the effect on abatement costs of adding endogenous technical change (ETC) to the PAGE2002 integrated assessment model. What defines the PAGE model as different to many other, often more complicated Integrated Assessment Models (IAM), is that it is a stochastic model that incorporates simplified environmental and economic equations to approximate complex phenomena such as radiative forcing and economic impacts. By using probability distributions for economic and environmental variables, the model calculates a range, and not a best guess outcome, for global mean temperatures, impacts and costs.

This work continues directly on from the work published by Alberth and Hope (2006) focussing on the addition of ETC into the PAGE2002 model, and the general effect that it had on the model's overall performance. The reader is requested to refer to this publication for a technical description of the *modified PAGE2002 model with Endogenous Technical Change*.

The aim of this current paper is to look more closely at the various changes in the risk structure of the model's appraisal of future abatement costs. When looking through the lens of a stochastic model, the importance of understanding how learning uncertainty changes through time and experience becomes paramount. A direct consequence is that, as with most situations of uncertainty, great value can be extracted through carefully managing risks and opportunities.

The following chapter will provide a brief summary to the modelling framework. Chapter 3 will look at a series of results in terms of standard abatement paths with chapter 4 looking at the results of optimising the abatement path to reduce discounted costs and impacts. Chapter 5 will draw conclusions from these results and suggest avenues for future work.

Chapter 2. Modelling Framework Summary

This research uses the PAGE2002 model, modified to include learning factors that relate directly to the cost of abatement. This section gives a brief introduction to the model, with details of learning parameters and their calibration in the following section. For a full discussion of the modified PAGE2002 model with endogenous technical change, please refer to Alberth and Hope (2006), and for a description of the PAGE2002 model without endogenous technical change, see Hope (2002, 2004).

2.1. The PAGE2002 model

The abatement cost calculation in the PAGE2002 model starts from the CO2 cutbacks needed to reach a certain stabilisation target. In this paper, the cutbacks are calculated as the difference between prescribed stabilisation scenario emission levels and the CPI or "Common POLES-IMAGE" (Den Elzen et Al. 2003) baseline 'Business As Usual' (BAU) scenario. The BAU scenario has essentially zero abatement costs.

The abatement costs incurred through reducing emissions can then be compared to the benefits gained through reduced global warming impacts. The model uses initial discount rates averaging around 5% per year (with a maximum of 11% for China), falling gradually to an average of about 4% per year over the 21st century. The model calculates annual abatement costs and the distribution of abatement NPVs for different CO2 abatement strategies and when using a genetic algorithm optimisation add-in, can also determine a (stochastically) optimised abatement path

2.2. Implicit learning and the treatment of technical change

In the standard PAGE2002 model, an abatement cost distribution is used to represent the future unit costs of cutbacks; this distribution does not vary with time or level of cutback, except that two types of cutbacks are modelled, namely *cheap cutbacks*, that relate to all levels of CO2 cutbacks and *added-cost cutbacks*, that are added to the cheap cutback costs for all cutbacks above a certain threshold. The cost of abatement within the original PAGE 2002 model is based on expert opinion, taking into account the availability of technologies, and expected future costs. Hence it can be said that the standard model does take learning into account, but only implicitly.

2.3. Alterations made to the PAGE2002 model

The two alterations made to the standard PAGE2002 model for this paper are:

a) the definition of cheap and expensive cutbacks, and

b) the inclusion of regional learning-by-doing for these two types of cutbacks based on cumulative CO2 reductions made with incomplete regional spill-over effects.

In the standard model, three uncertain parameters are used to model abatement costs, the unit cost of the cheapest control measures, the maximum cutback proportion that can be achieved by the cheap control measures, as a percentage of year 2000 emissions and the additional unit costs for reductions in excess of maximum proportion achieved by cheap control measures. However, this type of calculation leads to fixed levels of cheap cutbacks despite relatively large variations over time in energy requirements and emissions projected in the future, particularly in non-annex-1 countries. In this modification, the cut-off level between cheap and expensive cutbacks is defined as a percentage not of base year emissions, but of the region's business-as-usual emissions for the year in question.

The more significant change was to use variable costs that are calculated with respect to cumulative cutbacks made instead of using fixed costs for both the Cheap Cutbacks and the Added-Cost Cutbacks. Here the commonly used "Experience Curve", a method for calculating costs as a function of experience is gained, was modified for our purposes. A short explanation follows.

> *Empirical evidence for learning curves was first discovered in 1925 at the Wright-Patterson Air Force Base where it was discovered that plotting an aeroplane's manufacturing input against cumulative number of planes built on a log/log scale was found to result in a straight line (Wright 1936). The resulting benefits in efficiency Wright proclaimed to be a result of "Learning by Doing" in his 1936 publication. This "learning curve" was calculated for a manufacturing input such as time as shown in equation (1), where* N_t *is the labour requirements per unit output for period (t),* X_t *is the cumulative output in units by the end of the period where 'a' is the constant and 'b' the learning coefficient as determined by regression analysis.*

$$\log N_t = a - b \log X_t$$

Equation 1

> *The next major advancement in learning curves was made by Arrow in his 1962 publication (Arrow 1962, IEA 2000). He was able to generalise the learning concept and also put forward the idea that technical learning was a result of experience gained through engaging in the activity itself. Undertaking an activity, Arrow suggested, leads to a situation where "favourable responses are selected over time" (Arrow 1962, p156).*

> *During the 1960's the Boston Consulting Group (BCG) popularised the learning curve theory. They further developed the theory and published a number of articles on the subject (BCG 1968 in IEA 2000, Henderson 1973a, Henderson 1973b). They also coined the term "experience curve", as distinct from "learning curve" which related to 'unit total costs' as a function of 'cumulative output', rather than 'unit inputs' as a function of 'cumulative output' as shown in equation (2). In this equation the cost per unit '*C_t*' depends on the cumulative number of units produced '*X_t*' and the constant 'a' and coefficient 'b' that are found using a regression style analysis (Alberth and Hope, 2006).*

$$C_t = C_0 \left(\frac{X_t}{X_0} \right)^{-b}$$

Equation 2

Although the standard model allows the distribution of cheap cutback costs to include negative values, in the modified model all cutback costs are positive. This implies that at any point in time, all technologies leading to abatement are more expensive than the cost of traditional fossil fuel energy sources. Hence the modified model implicitly includes a floor cost (albeit variable) for energy production beneath which it is assumed low carbon alternatives cannot fall and towards which their costs gradually tend with learning. By associating a cost to carbon abatement and not the underlying energy production method itself, all learning takes place on this very aggregate level. Cost reductions occur according to the traditional learning by doing formula with a small autonomous learning element.

The resulting minimum abatement unit costs out of the 8 regions for the ‘450ppm’ stabilisation scenario are shown in Figures 1a and 1b. Through large-scale abatement the price per MT of CO2 abatement reduces significantly. By looking at the 5% to 95% error lines one can see that the model predicts quite a wide range of possible abatement costs.

Figure 1a and 1b

Figure 1a Unit CO2 abatement costs for cheap cutbacks – 450ppm

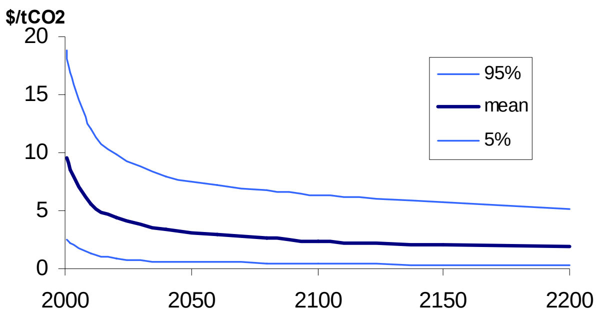

Figure 1b Unit CO2 abatement costs for added cost cutbacks – 450ppm

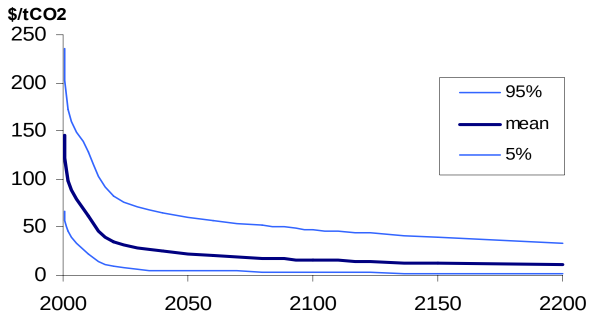

2.4. Learning rates for endogenous technical change

Due to the aggregate nature of CO2 abatement and the lack of previous work done in this area, we incorporate quite a wide distribution. The values are chosen to encompass the general results suggested by other authors, such as Zwaan (2004), Junginger et Al. (2005), Rubin (2004), WEA (2000) and Papineau (2004).

2.5. Autonomous learning rates

As well as endogenous technical learning, the PAGE2002 model with learning also includes a small amount of autonomous technical learning that occurs on top of ETC. These values are represented by a triangular distribution as shown in Table 1[3] and are based on judgement.

[3] The autonomous technical change level is smaller for the cheap cutbacks due to the fact that the CPI baseline from which emission reductions are calculated shows a reduction in total emissions. This we have assumed to be due to the autonomous technical development of technologies in the cheap cutbacks category.

Table 1 Key triangular distributions used for the PAGE2002 with learning model

	Key learning coefficient distributions	Mean	Units	Min	Mode	Max
	Uncertainty in baseline emission levels	-5	%	-30	-5	20
CL	CO2 abatement cost (cheap=1)	17	$/TCO2	0	10	40
CH	CO2 abatement cost (added cost=2)	250	$/TCO2	100	250	400
AL	Autonomous technical change 1	0.10	%/Year	0	0.1	0.2
AH	Autonomous technical change 2	0.25	%/Year	0	0.25	0.5
BL	Learning coefficient 1	0.2		0.04	0.2	0.36
BH	Learning coefficient 2	0.2		0.04	0.2	0.36
EL	Initial experience stock 1	1000	MT of CO2	100	1000	1900
EH	Initial experience stock 2	10	MT of CO2	1	10	19
RC	Regional crossover of experience	87	%	70	90	100
MAX	maximum cheap emissions as percentage of zero cost emissions	30	%	10	30	50

Chapter 3. Stabilisation Results Comparison

3.1. The three stabilisation scenarios

Part of the motivation for this research has been a request from the Innovation Modelling Comparison Project (IMCP) (Koehler et al 2006) who have requested that research groups document the effects of including endogenous technical change (ETC) into pre-existing E3 models. Hence the three stabilisation scenarios used have been the '450 ppm[4]', '500 ppm' and '550 ppm' scenarios requested to be modelled by the IMCP group, originally developed by Tom Wigley using the TAR version of the MAGICC model (Wigley, 2003).

Figure 2 Emission paths comparison of the BAU, 450 ppm, 500 ppm and 550 ppm scenarios.

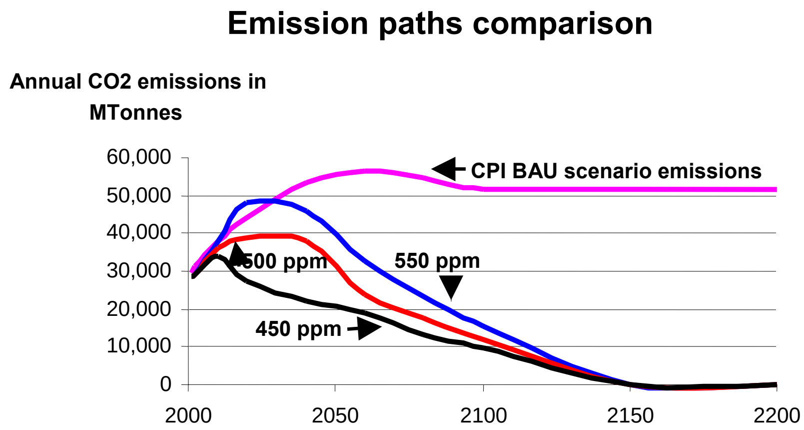

The "Common POLES-IMAGE" (CPI) scenario shown in Figure 2 is used as the BAU scenario for all of the stabilisation and optimisation scenarios tested as part of this research. It is also used by the IMCP as the common BAU scenario (Full details of the CPI scenario are included in Den Elzen et al., 2003) and is often referred to as the baseline scenario. CO2 abatements are calculated as the difference between baseline emissions and the scenario being tested.

Using both the modified and standard PAGE2002 models, the three stabilisation scenarios are modelled and the results collected and compared. The main emphasis in the presentation of results is on the lowest of the emission paths, the '450 ppm' scenario, however a full set of experiments have also been carried out and some of their more important results are presented in Annexe 1. Qualitatively, however the ETC results for the '500 ppm' and '550 ppm' are found to be very similar to the '450 ppm' scenario.

3.2. Learning Curve for Unit Abatement Costs

Abatements following the '450' ppm scenario can lead to dramatic cost reductions for both the cheap and added-cost cutbacks as modelled by the PAGE2002 model with learning. Figure 3 shows how the bulk of these benefits occur over the first 20 years with further important reductions still occurring over the following 40 or so years. However, due to the different cost structures, unit costs are not directly comparable to the standard model. With

[4] The term ppm stands for parts per million, and is used to define the amount of gas particles in the atmosphere.

the CPI scenario, the cost reductions of abatements are due mainly to the small autonomous learning coefficient, although, even with the baseline case, there may be some abatements due to the stochastic nature of the model. As expected, the initial important reductions in the cost of cheap and added-cost cutbacks slow down substantially. Furthermore there remains a great deal of uncertainty in unit costs where expensive cutbacks in the year 2050 may cost as much as \$50/tCO_2 at the 95% level or as little as \$3/tCO_2 at the 5% level depending on the parameters chosen at random from the given ranges. The mean value is \$15.7/tCO_2.

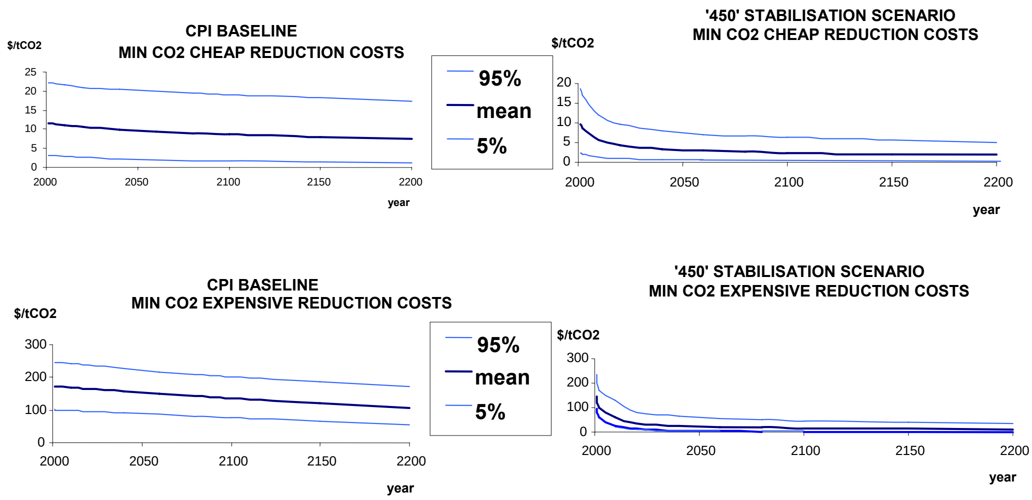

Figure 3   Cost of Cheap and Added-Cost Cutbacks as a function of time for the 450ppm

3.3. CO2 concentration and global mean temperature for the '450 ppm' scenario

Due to a feedback loop in the PAGE2002 model's carbon cycle that limits the natural ocean's carbon sequestration ability as the temperature rises (Hope, 2004), the PAGE2002 model's mean expected emissions are higher than the stated value for all three of the scenarios used by around 70 ppm. Furthermore, due to the stochastic nature of the model and the distributions for the input coefficients, it can be observed that the final concentration levels are highly uncertain as is clearly shown in Figure 4.

Figure 4   CO2 concentration and Global mean temperature rise by 2100 from 2000 levels for the 450ppm scenario. CO2 Concentration have a mean of 521ppm and 5% and 95% confidence levels at 477ppm and 572ppm Respectively. Global Mean temperature change in shows a rise of 3.1 degrees centigrade and 5% and 95% confidence levels at 1.8 and 4.6 Respectively.

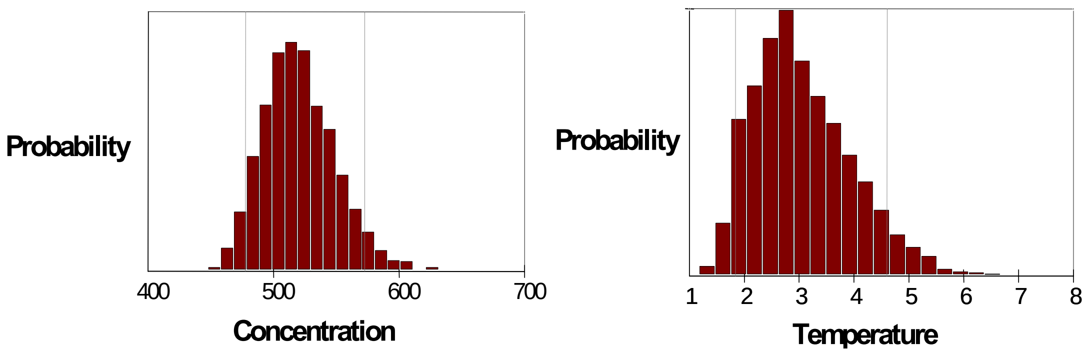

The temperature rise includes a 0.5 Celsius rise having already occurred by the year 2000, for both the standard and modified PAGE2002 models. This temperature rise is also a slightly conservative estimate of the IPCC's Working Group I report, which calculated the anthropogenic temperature rise by 2000 to be 0.6 +/- 0.2 Celsius (IPCC, 2001b). The resulting temperature rise despite following the more stringent 450 ppm scenario, is shown to still vary from about 1.5 to 6 degrees Celsius.

3.4. Discounted Total Abatement Costs for the '450 ppm' Scenario

As presented in Figure 5, the total abatement costs discounted to year 2000 dollars is between 2 and 13 trillion dollars, with a mean value of 6 trillion dollars.

Figure 5 Comparison of discounted abatement cost distributions between the standard PAGE2002 model and of the modified PAGE2002 model with a mean learning rate of 13% and with a zero learning rate.

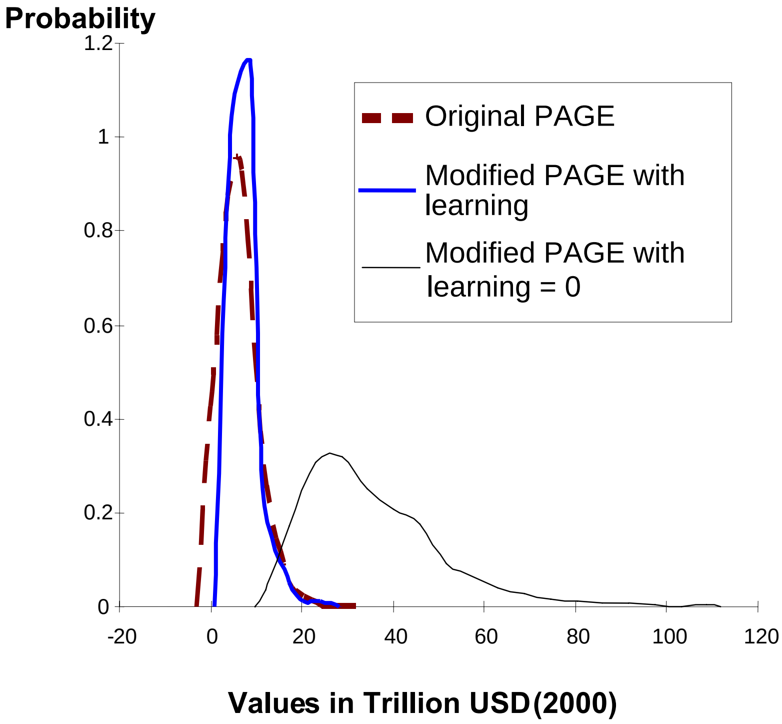

Despite the many differences in the way in which the standard PAGE2002 model and the modified PAGE model deal with abatement costs, the distribution of total abatement costs remain remarkably similar. Not only is this true with the '450 ppm' stabilisation scenario, but was also found for the '500 ppm' and '550 ppm' scenarios. The third line representing the modified model with no learning (neither autonomous learning nor learning-by-doing) represents the costs of abatements if renewable technologies were to remain fixed at 2000 prices. Unsurprisingly, these costs are well above those found by the model with endogenous technical change, and it's also interesting to note the increased level of uncertainty.

Figure 6 presents in sequence the input coefficients that have the most profound positive and negative influences on total abatement costs. The sensitivity analysis is carried out using "Rank Order Correlation" analysis that determines the level of correlation using ranks rather than the actual correlation of inputs to outputs and is based on Spearman rank correlation coefficients (Guide to using @risk p87). The values can vary between +1 and -1 signifying that the input and the output are completely correlated either positively or negatively. A value of 0 represents no correlation.

The importance of the learning coefficient (for added-cost cutbacks), the initial costs of cheap cutbacks and the elasticity of utility, which contributes to the discount rate, are evident. The negative correlation between the learning coefficient and total abatement costs is expected as an increase in a technology's learning coefficient leads to an increase in the learning rate which in turn reduces abatement costs. Similarly, a higher elasticity of utility and *pure-time-preference* (ptp) rate both reduce the importance of future costs as presented by the discount rate formula presented above. A larger initial knowledge stock, or initial cutback costs can only increase the total abatement costs and the sensitivity analysis in Figure 6 also shows this positive correlation clearly.

Figure 6 Correlation sensitivity for total abatement costs for the 450 stabilisation scenario.

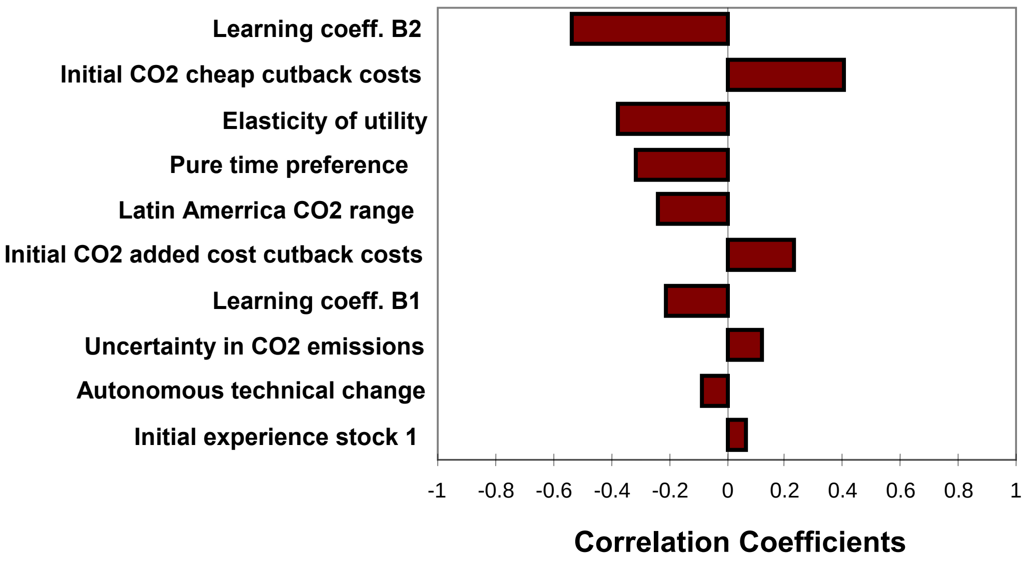

The sensitivity analysis in Figure 7 reveals the remaining importance of learning for the sum of total abatement and impact costs, with the learning coefficient for added-cost cutbacks, B2, being the fourth most important. Here a number of other important cost factors come into play such as the level of thermal forcing that occurs for a given concentration level, or the terms affecting the discount rate.

Figure 7 Correlation sensitivity analysis of total abatement costs and impacts combined

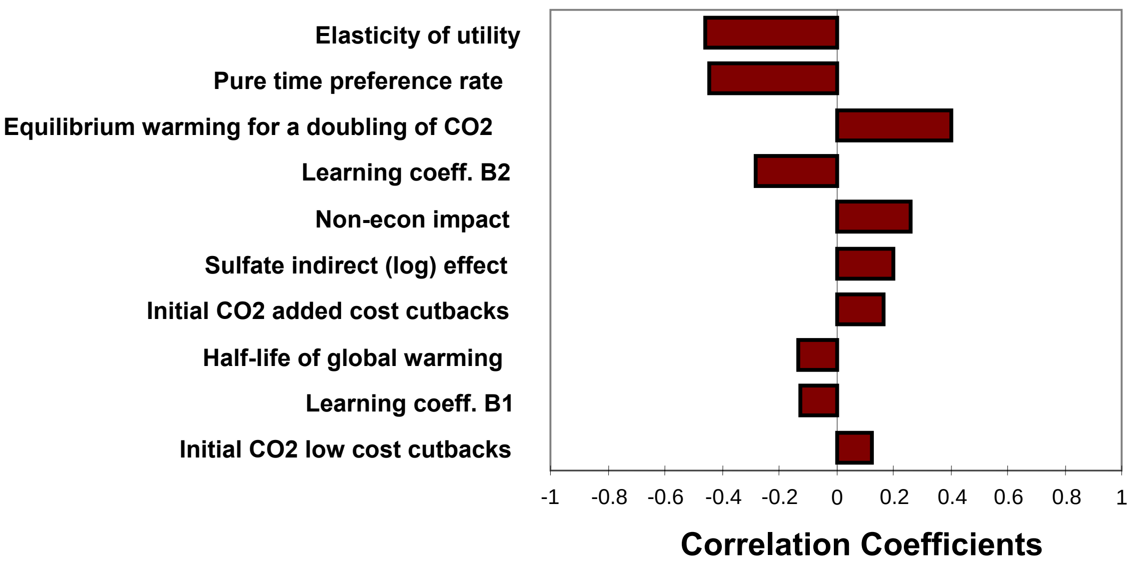

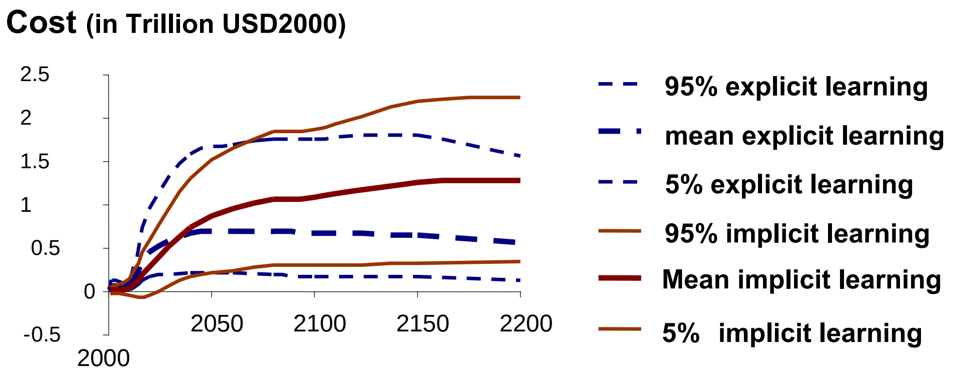

Figure 8 Comparison of total abatement costs of the PAGE2002 model with explicit learning to the PAGE2002 model with implicit learning.

Figure 8 provides a clear representation of how the implementation of explicit learning has affected the PAGE2002 model. Due to higher initial costs, cutbacks over the first part of the century actually cost more than in the standard model. These cost increases are gradually reduced to the point where the two models predict similar outcomes around the year 2040. After this point, costs in the modified model remain well below the levels in the standard model as technology continues to advance. Discounting means that the large cost benefits after the year 2050 are balanced by the minor cost increases during the first 50 years.

3.6. Abatement costs as a proportion of gross world product (GWP)

The modified PAGE2002 model envisages quite a large decrease in unit abatement costs as large scale abatements begin to take place for all of the scenarios considered. However, in order to reach these cost reductions, an investment which goes mainly into learning has to be made, as represented visually in Figure 9. Due to the increasing length of the time periods used, starting by using yearly calculations for the first 2 years, then larger time steps of 10, 20 and 50 years, it was impossible to obtain fine details over, for example, the entire first 10 or 20 years. Nevertheless, there is a clear increase in costs during the first two years representing the large costs brought about through large scale use of what remains a premature technology.

Figure 9 Abatement costs as Percentage of GWP for '450' ppm scenario

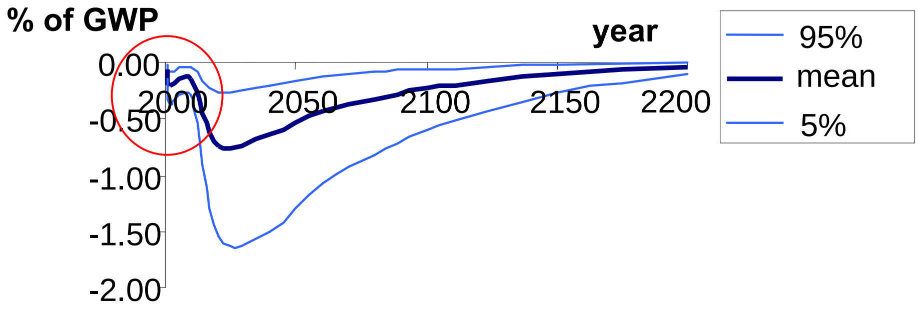

This double peak of costs as a proportion of GWP presents a particularly interesting feature that was also found for the '500 ppm' and '550 ppm' scenarios, and would most likely remain for all but the most gradual transitions from carbon intensive energy to their alternatives. In the case of the '450 ppm' scenario, this annual investment came to a mean expected value of 0.2% of GWP as circled in red in the diagram. This initial rise and fall in costs could also be referred to as a learning investment, since the main advantages earned from these initial abatements come not from the actual abatements made, but rather from the cost reductions that it would bring to future abatements. The possible policy implications of this feature are most interesting and confer with a popular call for an Apollo-style mission[5] to break the cost barriers of alternative energy technologies. The aim of such a mission would be to pass through the first of the two peaks, driving down the costs of renewables and generally increasing the knowledge base. However an important question to answer is: who would fund such a mission? As businesses noticed during the 70's and 80's, progress down the learning curve does not necessarily lead to sustained competitive advantage (Lieberman, 1987). Furthermore, research has also shown that justice or fairness can also play an important role in individual or policy decisions (Baron, 2000). Hence it is understandable that individual countries or regions are unwilling to support the cost of providing what essentially becomes a common good, unless the international sharing of such a cost is perceived as fair and just.

After running the model using the @risk[6] correlation sensitivity analysis, it was noted that the most important coefficient with respect to abatement costs is the learning rate of the added-cost cutbacks, closely followed by the discount rate coefficients. Similarly the original PAGE2002 model also finds that the discount rate coefficients are highly correlated to the value representing both impacts and abatement costs.

For a visual comparison of how the learning rates affect the NPV of different strategies, Figure 10 presents the mean NPV of the various abatement strategies using two different learning rates (LR), 13 % and 19%.

Figure 10 NPV of abatement for different learning rates and scenarios (LR)[7]

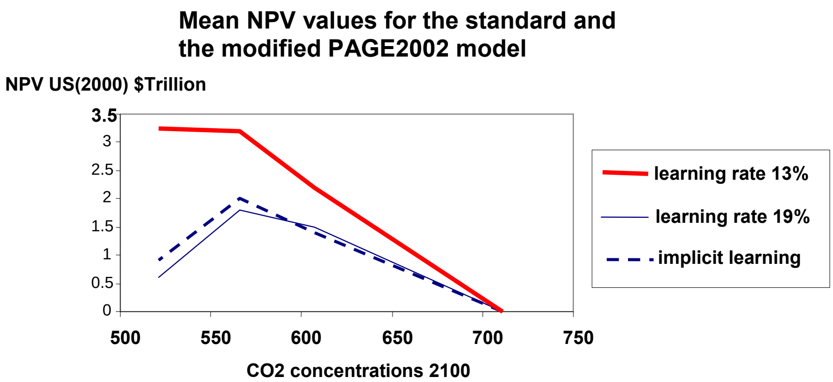

The 13% learning rate reflects a somewhat reserved view of how these technologies will perform in the future, whereas the higher learning rate of 19% assumes that technologies such as solar power will continue to see substantial cost reductions through to mass production on a global scale. The learning rate of 13% (or learning coefficient of 0.20) is the standard mean learning rate used for the modified PAGE2002 model. The dashed line

[5] An Apollo-style mission is exactly what the Apollo Alliance promotes. Although the focus is to "achieve energy independence in one generation" (http://www.apolloalliance.org/about_the_alliance/) their methods for achieving it are mainly through developing renewable energy solutions.
[6] @risk is a registered product part of the Palisade Decision Tools Suite
[7] See equation 11 and equation 12 for an explanation of the learning rate, progress ratio and learning coefficient.

represents the results from the standard PAGE2002 model before the addition of ETC. Table 2 shows this information such that it clearly presenting the mean NPV's found for each of the four scenarios. Overall what this information suggests is that the higher learning rate would encourage the adoption of a higher level of abatement (thus stabilising at lower greenhouse gas concentrations).

Table 2 NPV of abatement (Trillion US Dollars 2000) for different learning rates and scenarios in Table format

	CPI	550 ppm	500 ppm	450 ppm
implicit learning	0	1.4	2	0.9
learning rate 13%	0	1.5	1.8	0.6
learning rate 19%	0	2.2	3.2	3.25

The large differences in final NPV costs resulting from the simple comparison of two learning rates demonstrates their relevance in calculating abatement costs and NPV's, and hence the importance of accurately being able to determine their range of values. More importantly, given that we are now in a position of very low information with regards to what the future costs of abatement technologies may be, the key question here is "how much would it be worth for decision makers to reduce learning rate uncertainty"? This important point will be looked at again in chapter 5.

Chapter 4. Optimisation Results

4.1. Comparison of optimisation scenarios

Using a stochastic optimisation feature, 'Risk Optimiser', we are able to choose an emission scenario that minimises the total value of impacts and costs. 'Risk Optimiser' works by running a large number of @risk simulations and by varying assigned inputs between each simulation using genetic algorithms. The aim of this optimisation is to minimise the total amount of discounted abatement costs and global warming impacts through choosing different emissions paths and associated abatements when compared to the BAU scenario. Risk Optimiser uses the initial emission paths as a base and then, through a large number of mutations, it narrows in on the desired emission path. Each time a new best 'recipe' is found, the optimisation feature uses these values as the starting point from which mutations are made. Initially, 100 iterations per simulation was used followed by an optimisation with reduced variable ranges and 1000 simulations to fine tune the results. This method uses a single optimising agent deciding all global CO2 emission paths for the entire 200 year time frame of the PAGE2002 model. The agent's aim is simply to minimise abatement costs and impacts to climate change discounted to year 2000 dollars.

Optimisation scenarios starting from varied initial values should all lead to the same optimisation if a single global minimum value does exist. On the other hand, in cases where there may be many local minima, the optimisations may arrive at different results. The inclusion of learning has often been shown to lead to non-linear style outputs with multiple 'optimised' solutions (Barreto et al. 2004, Löschel 2002). Authors such as Manne & Barreto (2001) further suggest that one should use a variety of different commercial applications and for each, try various starting positions to overcome the problem of finding the global minimum. Although the adoption of various applications was not possible as part of this exercise due to time constraints, various starting positions were used to try and detect other local minima. Nevertheless for all of the starting points that were tested, a single set of final emission paths was found, a result that is not unsurprising considering the level of aggregation of technologies.

The following results look at a number of optimisations made with some changes to the basic parameters however we also compare the modified PAGE2002 optimisation results with the standard PAGE2002 optimisation results to see how the inclusion of an explicit cost function affects the way abatements are optimally made.

4.2. Comparison of standard to modified optimisation

Figure 11 Distribution of discounted preventative costs for the optimisation scenarios with the standard and modified PAGE2002 models.

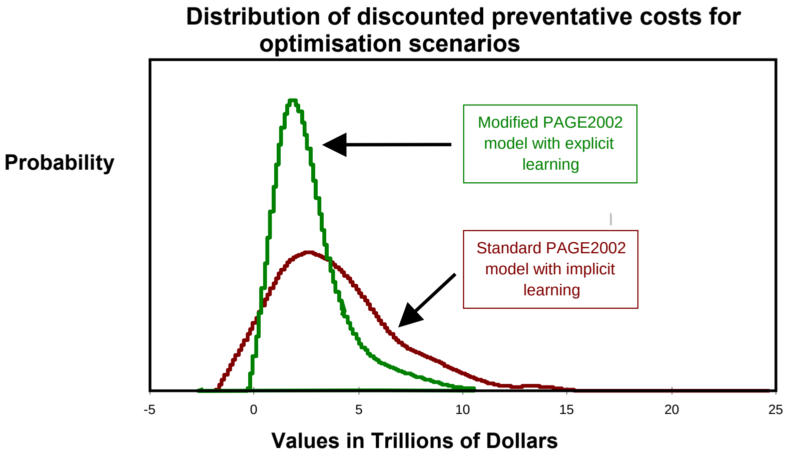

Figure 11 compares the distribution of discounted total abatement costs between the standard model's optimised emission path using fixed costs, and the modified model's emission path that explicitly models costs. The modified PAGE2002 model's results show less dispersion than that of the standard PAGE2002 model. The total discounted costs are also reduced from 3.6 Trillion to 2.8 Trillion due to the effects of learning, however these differences are highly dependant on the learning parameters used. This suggests that not only could the total abatement costs be reduced due to the experience curve effect, but also that the uncertainty of abatement costs might also be reduced. The reduction in risk, however, may owe itself to the delayed abatement path chosen. Perversely, although the uncertainty of discounted costs has reduced, it would seem that an "optimised" abatement path actually increases future abatement cost uncertainties when comparing the model with ETC to the one without. This is shown below.

From Figure 12, one can note that the total amount of annual abatements made in weight are roughly the same for both the modified and the standard model with mean concentration levels for the year 2100 in the vicinity of 500ppm. What changes, however, is the way in which these reductions are made and the apparent annual abatement cost uncertainty that learning brings in to the model. The optimisation path for the modified version of PAGE2002 abates less during the first time periods, but then accelerates its abatement program towards the middle of the century, overtaking for a time those made by the standard model due to a sharp jump in abatements.

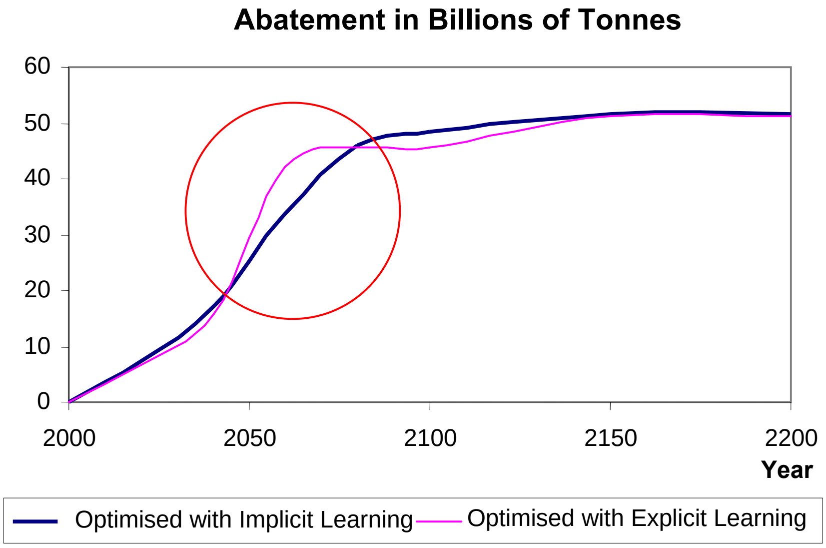

Figure 12 Comparison of optimised abatement with and without learning

The importance of "path dependence", as noted by Grubb & Köhler (2000) comes to mind in this example. Path dependence is the term describing the difficulty encountered by new technologies requiring a different set of infrastructures to become competitive and experience large market growth. Furthermore, the lifetime of existing capital tends to slow down even the most economically viable solution.

Figure 13 The 'dash for gas' in the UK and OECD Europe[8]

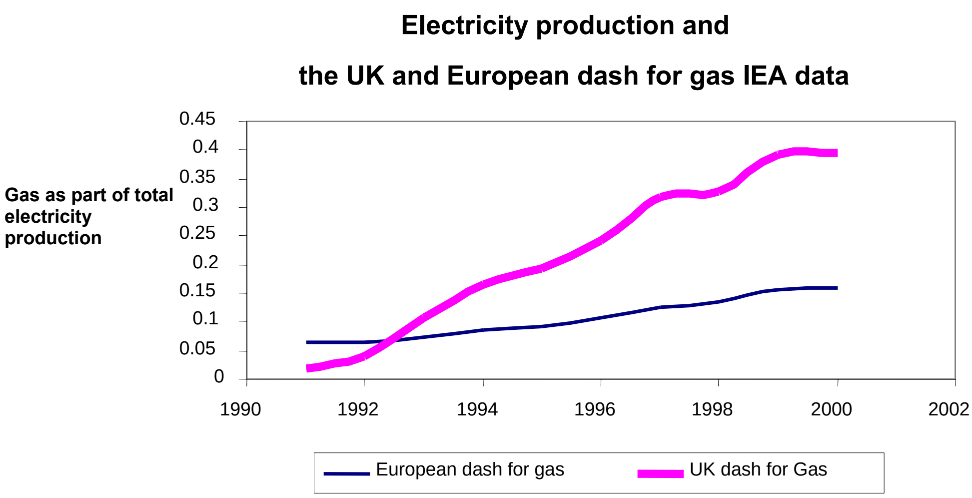

A parallel example of the uptake of an economically viable energy technology in the UK is the 'dash for gas'. Due to superior performance both in terms of green house gases and economics, huge investments have been made into Combined Cycle Gas Turbine (CCGT) technology and its dispersion as a fraction of total electricity produced rose almost 40% over a ten year period. Nevertheless on a Europe wide level, its market share increased by only 10% during this same period, as shown in Figure 13. The optimised scenario presented in Figure 12 when compared to the original CPI baseline emissions in Figure 2 require global

[8] Using IEA data from 1991 to 2000 inclusive

abatements from all industries to jump from about 20% to 80% of the CPI scenario, and this on a global scale in the period of 20 years (between 2040 and 2060). Such a sharp increase in abatements on a global scale may prove to be unrealistic and hence the limitation in any model showing such jumps should be looked into further.

As can be seen in Figure 14, the added-cost abatements for the optimised scenario show far slower cost reductions, than did the standard ‘450 ppm’ (as was found for the ‘500 ppm’ and ‘550 ppm’) abatement scenarios. Almost no added-cost cutbacks are made for the first quarter of the century, though, even the small amount of cutbacks made are able to substantially reduce costs, before the intensive cutbacks are made around 2050.

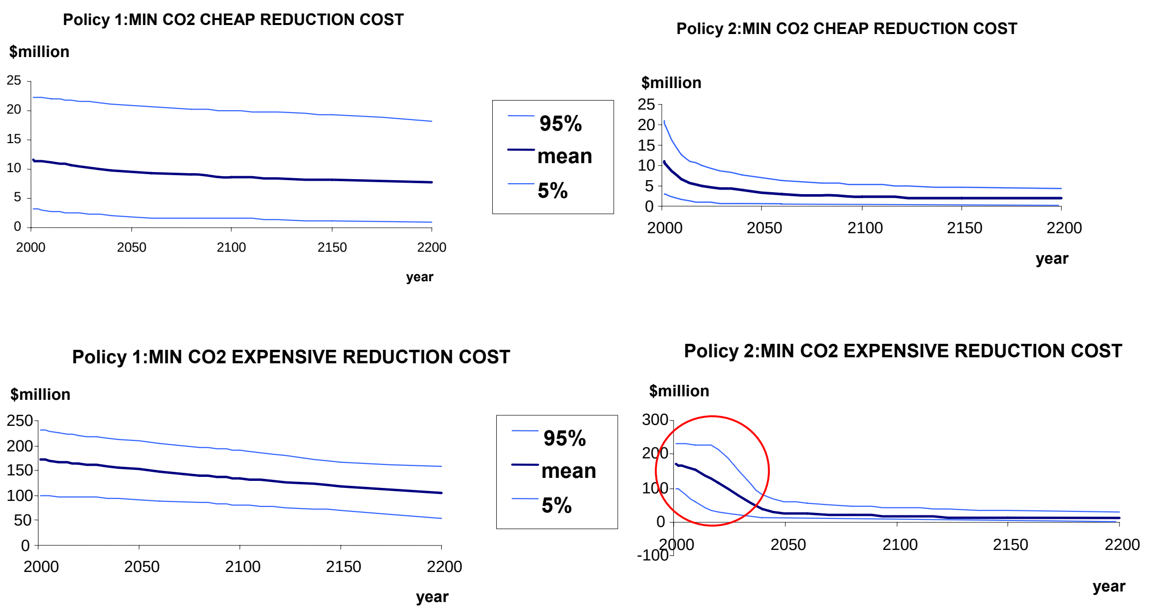

Figure 14 Cheap and added-cost cutbacks.

Once again from Figure 15 we see how the modified PAGE2002 model with explicit learning shows a later start, and then calls for greater action than the standard PAGE2002 model. It is at this point that we see how the long term optimisation decision actually leads to an increased level of cost uncertainty in the future. As circled in red, the compressed structure that learning brings to the optimised abatement plan greatly increases the associated risk. Whereas with the standard model the mean expected costs of abatements rise to around 0.5% (the thick dotted line), with only a 5% chance that they rise above 0.9% of GWP (the lower thin dotted line). With the addition of an explicit abatement cost function with endogenous learning, the possibility of substantial costs rises dramatically, and even though the highest mean abatement cost remains at around 0.5% (the thick solid line), there is a 5% chance that more than 1.5% of GWP would be needed to cover the costs of abatement. This is due to the uncertainty in abatement costs being magnified during the initial period of large-scale abatements as this is the time where the ‘learning costs’, the costs associated with driving technologies down the learning curve, would be most important.

Figure 15 Comparison of optimised scenarios with explicit and implicit learning as percentage of GWP

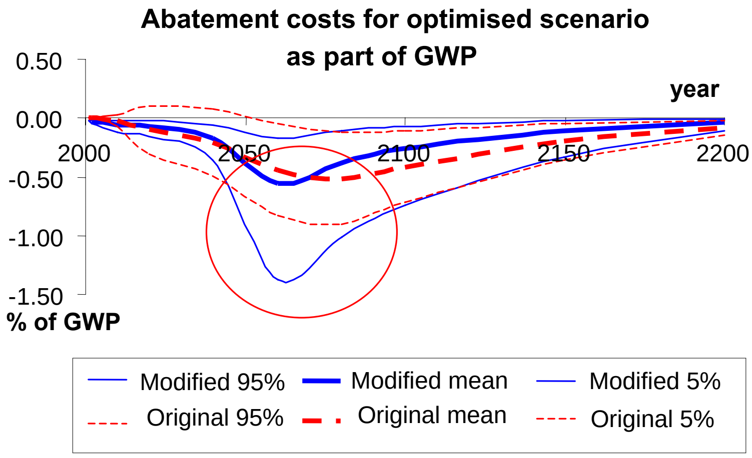

The diagram also reveals that the optimisation has completely smoothed out the initial peak of the learning investment for the modified version of the model that was previously demonstrated in the non-optimised stabilisation scenario (see Figure 9).

4.3. Spill-over and the costs of climate change

Figure 16, along with the sensitivity analysis in Chapter 4 of this paper[9], demonstrates that the level of sharing of experience between regions has minor importance. Normally a mean value of 0.87 is used, with upper and lower limits of 1 and 0.8 respectively. This is presented by the solid line in the Figure. Changing the mean value to 0.5 with upper and lower limits of 0.7 and 0.3 respectively naturally lead to increased unit costs since each region has to pay for their own learning, though, as seen from the diagram, the overall effect is found to be only small, in agreement with the above mentioned sensitivity analysis.

Figure 16 Abatement costs as part of GWP with only 50% cross-regional spill-over

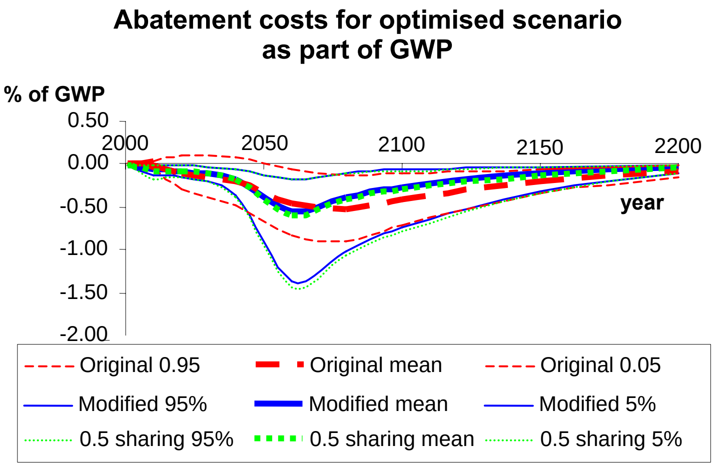

[9] The coefficient is found to be of little importance in the sensitivity analysis, not registering in the Figure.

4.4. Effect of increasing the learning rate

Increasing learning rates from a mean value of 13% to a mean value of 19% (that is, to change the learning coefficient from a distribution centred on 0.2 to a distribution centred on 0.3) tend to lead towards higher abatement levels that occur earlier, as shown in Figure 17. The burden is also substantially reduced.

Figure 17 Comparison of abatement costs as percentage of GWP for LR= 13% and 19%

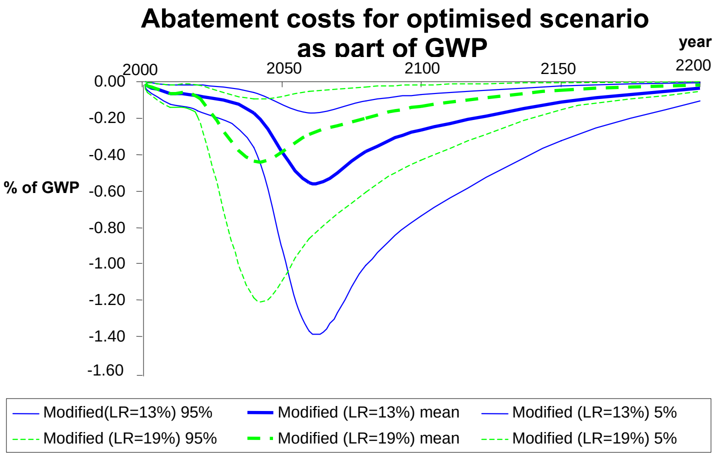

The increased learning rate used in this example has improved the cost effectiveness of abating CO2 emissions which in turn has lead to higher abatement levels. Mean global temperatures are around 0.4 degrees Celsius lower than previously calculated with both the modified version (using a learning coefficient of 0.2) and the standard version of PAGE2002. Since learning rates of different technologies have been shown to vary from 5 to 35% (examples are given in Alberth & Hope 2006), the choice of the learning rate distribution could certainly be an important source of error, and hence should be a priority of future research. This is in agreement with the sensitivity analysis of the previous chapter.

4.5. A short note on the effectiveness of optimisation

To give a representation of the accuracy with which Risk Optimiser is able to choose an abatement path in the PAGE2002 model, a comparison between the initial 'best optimised values' to the final 'best optimised values' is made. Each time a simulation's 'recipe' describing its emission path leads to a 'best optimised value' of the target variable, the model records these values creating a list of all new best emission paths. Figure 18 compares the initial values found to the final values of a total of 100 'best optimised values' found for the 7000 simulations carried out. The top two diagrams show the first half and final half of the values, while the bottom two diagrams compare just the first 10 and final 10 best values. In this example, the starting point used was already close to optimised with a mean discounted cost of US$37.5 trillion. After 50 scenarios this was reduced to US$34.8 trillion, and after 100 scenarios the value further decrease by only US$0.1 trillion to US$34.7 trillion.

What can be seen from these diagrams is that, although there remains a level of uncertainty after the 7000 simulations, the emission paths remain quite stable. This suggests that the @risk optimisation tool is quite effective at finding the optimum emission paths within the IAM framework used here.

Figure 18   Effectiveness of optimisation over 7000 simulations

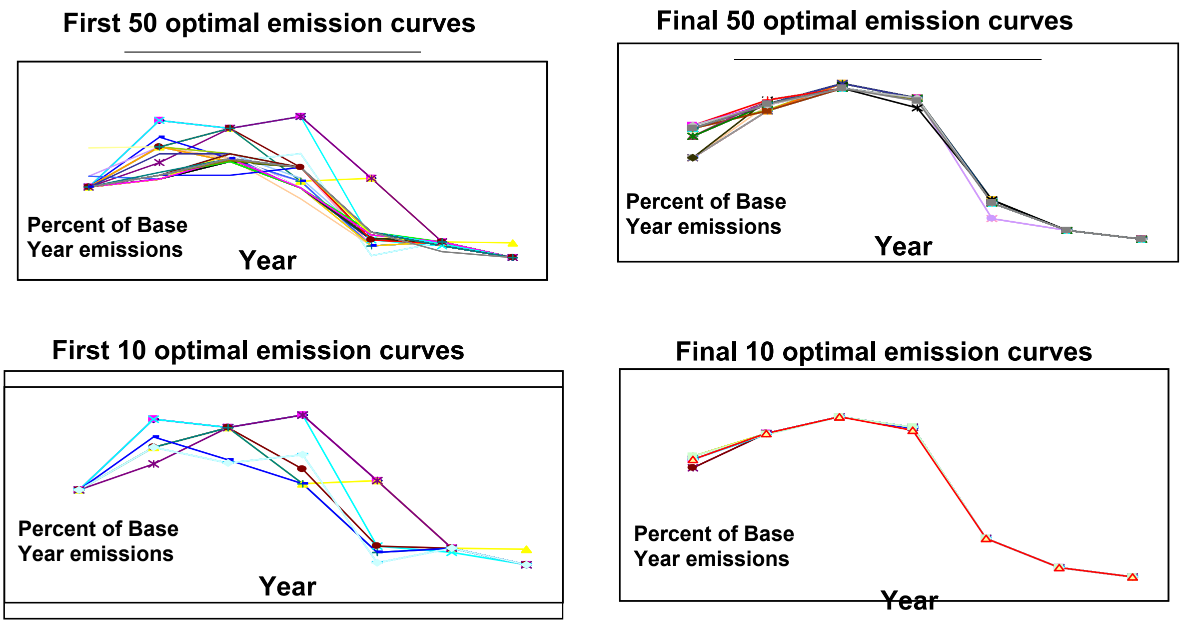

4.6. Optimisation overview

The use of a single optimising agent to reduce discounted future costs and impacts leads to emission paths that generate high levels of CO2 abatement and thereby reduce mean expected CO2 concentration to levels generally below 500 ppm. The mean annual non-discounted costs of abatement are also reasonably low, with at most a mean expected value of 0.5% of GWP, with a 5% chance of costs rising above 1.5%[10]. The increased abatement levels lead to a large reduction in CO2 concentrations that fall to around 450 ppm. A weakness in the optimisation method used in the PAGE2002 model is that it does not take into account path dependence, whose path restricting presence could only serve to increase the level of abatement costs. The major drawback of using a single optimisation agent calculating from base year information is that the model does not allow for learning in the broader sense. Future decisions will be made using relevant knowledge that exists today, but will also be based on knowledge gained between now and that time. In other words, future decision makers will be able to make far better decisions than what this type of optimisation can allow for. Included in that will be an increase in the level of learning about climate sensitivities but there is also sure to be learning about the future state of technologies, and better projections about future costs. In this sense we foresee among other things an important level of "learning about learning".

[10] Please refer to Figure 15

Chapter 5. Concluding Remarks

5.1. The PAGE2002 model with learning

How the modified PAGE2002 compares to other models

Despite a number of important changes in the way abatement costs are modelled and with the addition of ETC into the model, the **discounted** abatement costs for the three stabilisation scenarios modelled remain very similar to those of the standard PAGE2002 model (as presented in chapter 4 and also shown in annexe 1). This is found to be the result of a relatively small increase in immediate abatement costs traded off against a much larger reduction in heavily discounted future abatement costs. The similarity of the two models, however, is heavily dependant on the coefficients used, and the sensitivity analysis further demonstrates that the learning coefficient has a strong impact on the calculation of total abatement costs. When looking closer at the unit abatement costs for the '450 ppm', '500 ppm' and '550 ppm' scenarios, it can be seen that the large amounts of abatements necessary to reach any of the three stabilisation scenarios are able to drastically reduce prices from 2000 levels. The difference between the 450 ppm and 550 ppm, on a log scale, is only minor compared to the difference between the historical abatements made before the year 2000. Hence, as the vast majority of learning takes place very early on in the learning curve, all three stabilisation scenarios make the majority of their abatements at similar costs.

Figure 19 compares the modified PAGE2002 model's total discounted abatement costs to a number of other models as used in the IPCC report. As one would hope from a stochastic model based on many of the conclusions drawn from the IPCC, our stochastic model, including the 5% and 95% error bars, is able to encompass the majority of the modelling results that formed part of their report. However, the costs for low stabilisation levels do seem to veer towards the higher values, possibly due to the feedback loop in the PAGE2002 carbon cycle (Hope, 2004).

Figure 19 Abatement costs for a number of different models[11].

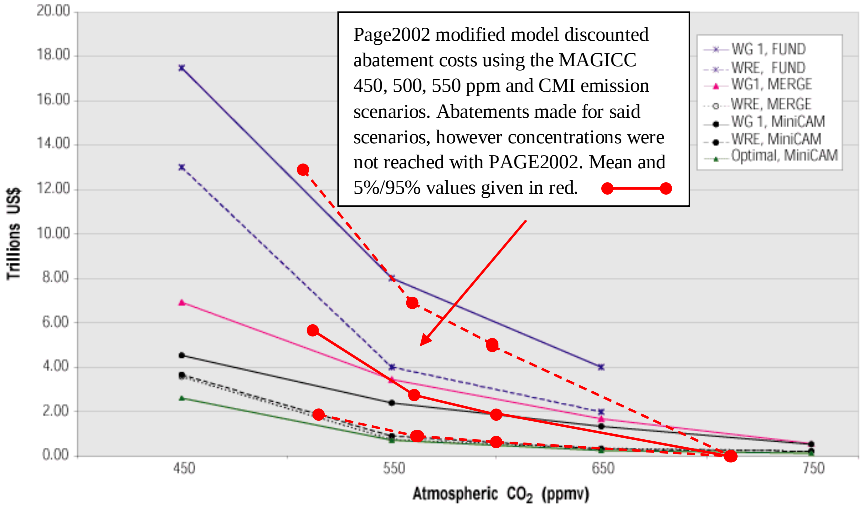

[11] Here the Page abatement costs has been superimposed on the IPCC Climate Change 2001: Working Group III: Mitigation diagram from chapter 8.4.2 (IPCC, 2001c). When using the CPI as the BAU scenario, the zero cost concentration level in the year 2100 is found to have a mean of 711 ppm

Learning Costs

It is found that for all but the most gradual transition from carbon intensive energy to their alternatives, there is a sizeable investment required which is represented by a relatively small initial peak in the abatement cost curve[12]. The associated costs would most likely relate to strategic deployment of technologies where costs are above those of the traditionally used technologies and could either be funded by government initiatives or by consumer preference. Most importantly, the deployment of these technologies would lead to increased learning-by-doing and, based on historical data, this should lead to cost reductions. Despite the fact that the initial learning cost, relative to the long term costs of stabilisation, is quite small, it is possible that it is this type of cost and not the more substantial later abatement costs that pose the greatest problem from a political perspective.

In order for such an injection of learning and funds to take place, one possible strategy would be the development of an Apollo style mission to promote greenhouse gas abating technologies, thereby pushing their costs down. Not only would this have an impact on the cost of future abatements, but it would also increase our knowledge of how these technologies are likely to progress, including greatly improved data of the associated learning rates.

5.2. Optimisation under great uncertainty

As shown in all of the scenarios, though particularly in the optimisation scenarios, the addition of ETC has brought an increased level of uncertainty to the future annual abatement costs. Specifically, the modified PAGE2002 found that strategically delaying abatements to a specific point in time carries the entire risk associated with the learning curve up to that point. This brings to mind the IEA (2000) caution discussed in chapter 2 about relying on cost reductions that may not materialise, with the inverse being equally true in that we may be assuming future abatement costs that are highly overrated.

The only sure method to reduce the risk of excessive abatement costs would be to actually make substantial headway in reaching the required cost reductions through early action, despite the fact that this would not lead to an “optimal” abatement path according to the criteria used here. Perhaps this points to a deficiency in the way the optimisation has been carried out, where costs and not utility incorporating risk aversion has been measured. Having taken the first steps toward decarbonisation, and armed with greater knowledge of the true costs of CO2 abatement after substantial learning has taken place, researchers and decision makers would be in a far better position to make decisions about the future thereby reducing the economic shock bestowed upon those living during the time of large-scale abatements. Furthermore the true learning rates of various technologies would become more apparent. Thus a substantial part of the value of strategic deployment of abatement technologies, beyond allowing for the technology to reach cost-effectiveness, may lie in allowing decision makers to choose the best plan of action for the future. To measure the potential value of immediate action, as opposed to deferred abatement, a real options analysis across a number of different models should be carried out, with the cost of the option being the learning cost, and the potential value of the option calculated as the potential to re-optimise future abatement paths more accurately allowing for overall reductions in costs and damages.

[12] As seen in Figure 9

Reference List

Alberth, S. & Hope, C., 2006, Developing the PAGE2002 model with Endogenous Technical Change, *EPRG working paper series, University of Cambrige.*

Apollo alliance, 2003. Retrieved: May 23, 2005 from http://www.apolloalliance.org/about_the_alliance/

Arrow, K, 1962, 'The economic implications of Learning by Doing', *Review of Economic Studies*, vol. 29, pp155-173.

Baron, 2000, 'Can we use human judgements to determine the discount rate?' *Risk analysis*, vol. 20, no. 6. Retrieved: 31 January, 2005 from Ebsco.

Barreto, L. & Kypreos, S., 2004, 'Endogenizing R&D and market experience in the "bottom-up" energy-system ERIS model', *Technovation*, vol 24 pp615-629. Retrieved: 7 May 2005 from Science Direct.

Den Elzen, Berk, Lucas, Eikhout & Vuuren, 2003, 'Exploring climate regimes for differentiation of commitments to achieve the EU climate target', *RIVM report* 728001023/2003. Retrieved: 17 May, 2005 from http://www.rivm.nl/bibliotheek/rapporten/728001023.pdf

Grubb, M. & Köhler, J., 2000, 'Induced Technical Change: Evidence and implications for energy-environmental modelling and policy', *Department of applied economics Working Paper Series*. Retrieved: 16 May 2005 from http://www.econ.cam.ac.uk/dae/repec/cam/pdf/wp0031.pdf.

Guide to using @risk, 2002, Palisade, Retrieved 24 Feb 2005 www.palisade.com.

Henderson, Bruce, 1973a, 'The experience curve reviewed – II. History', *BCG*, No 125. Retrieved: 3 May, 2005 from www.bcg.com.

Henderson, Bruce, 1973b, 'The experience curve reviewed – IV'. The growth share matrix, *BCG*, No 149. Retrieved: 3 May, 2005 from www.bcg.com.

Hope, C., 2003, The marginal impacts of CO2, CH4 and SF6 emissions, *Judge Business School Cambridge University*, Working Paper Series WP 10/2003

Hope, C., 2006, The marginal impact of CO2 from PAGE2002: An integrated assessment model incorporating the IPCC's five reasons for concern, *Integrated Assessment* Journal, Vol 6, issue 1

IEA, 2000, 'Experience Curves for Energy Technology Policy', *IEA Paris* , France.

IPCC, 2001a, *Climate change 200: Impacts, Adaptation, and Vulnerability*, Contribution of Working Group II to the Third Assessment Report of the Intergovernmental Panel on Climate Change.

IPCC, 2001b, *Climate change 200: The Scientific Basis*,, Contribution of Working Group I *to the Third Assessment Report of the Intergovernmental Panel on Climate Change.*

IPCC, 2001c, *Climate Change 2001: Mitigation*, Contributions from Working Group III *to the Third Assessment Report of the Intergovernmental Panel on Climate Chang.*

Junginger, M et al., 2005, 'Global experience curves for wind farms', *Energy Policy*, vol. 33, pp 133–150. Retrieved: 5 May, 2005 from Science direct.

Lieberman, M., 1987, 'The learning curve, diffusion, and competitive strategy', *Strategic Management Journal*, vol. 8, pp 441-452

Köhler, J., M. Grubb, D. Popp, O. Edenhofer and M. Galeotti (Forthcoming 2006) The Transition to Endogenous Technical Change in Climate-Economy Models. *The Energy Journal Special Issue, Endogenous Technological Change and the Economics of Atmosperic Stabilization*

Löschel, A., 2002, 'Technological change in economic models of environmental policy: a survey', *Ecological Economics, vol.* 43, pp 105-126. Retrieved: 1 May, 2005 from Science Direct.

Manne, A.S. & Barreto, L., 2001, 'Learn-by-doing and carbon dioxide abatement. *Interim report', IR-01-057, IIASA*, Luxemburg, Austria. Retrieved: 27 April, 2005 from http://www.iiasa.ac.at/Publications/Documents/IR-01-057.pdf.

Papineau, 2004, 'An economic perspective on experience curves and dynamic economies in renewable energy technologies', *Energy Policy, Article in press*. Retrieved: 7 February, 2005 from Science Direct.

Rubin, Yeh & Hounshell, 2004, 'Learning curves for environmental technology and their importance for climate policy analysis', *Energy*, 29 pp1551-1559. Retrieved 4 March, 2005 from Science Direct.

WEA, 2000, 'World Energy Assessment'. In: Turkenburg, W.C. (Ed.), *Renewable Energy Technologies*. New York/London.

Wigley, T., 2003, 'Modelling climate change under no-policy and policy emission pathways', *OECD, ENV/EPOC/GSP(2003)7/FINAL*. Retrieved: 25 May, 2005 from http://www.oecd.org/dataoecd/33/13/2489543.pdf.

Wright, T., 1936, 'Factors Affecting the Cost of Airplanes', *Journal of Aeronautical Sciences*, Vol. 3 p122.

Zwaan, B., 2004, "The learning potential of photovoltaics: implications for energy policy", *Energy Policy*, Volume 32, Issue 13, accessed via Science direct, 24/10/2004

Annexe 1. Summary Table of CMI baseline and IMCP scenarios

Table 3 Table of comparative results for the Baseline, 450, 500 and 550 ppm scenarios using both the standard PAGE2002 model with implicit learning as well as the modified PAGE2002 model with explicit learning with both a learning rate of 13% and 19%.

		CPI Baseline		450ppm scenario			500ppm scenario			550ppm scenario			
	Results (5%, **Mean**, 95%)	mean learning rate =13%	**implicit** learning	mean learning rate =13%	mean learning rate =19%	**implicit** learning	mean learning rate =13%	mean learning rate =19%	**implicit** learning	mean learning rate =13%	mean learning rate =19%	**implicit** learning	Units
Environmental Variables year 2100	year 2100 CO2 concentrations	640, **711**, 790		478, **521**, 573			515, **566,** 625			550, **607**, 673			ppm
	Total year 2100 Temperature Rise	2.5, **3.9**, 6.0		1.8, **3.1**, 4.6			2.0, **3.3**, 5.2			2.2, **3.5**, 5.3			°c
Unit Abatement Costs year 2100	Cheap Cutbacks	2, **9**, 20	-10, **10**,30	0.4, **2.4**, 6.2	.15, **1.3,** 4.1	-10, **10**,30	0.4, **2.6**, 6.5	.17, **1.4,** 4.3	-10, **10**,30	0.5, **2.6**, 6.5	0.18, **1.45,** 4.2	-10, **10**,30	Dollars/ Tonne CO2 (2000)
	Added-cost Cutbacks	90, **180**, 210	25, **35**, 45	2.8, **16.4**, 47	0.45, **6.92,** 25	25, **35**, 45	4.0, **17.2**, 44	0.6, **7.9,** 29.3	25, **35**, 45	4.3, **18.3**, 47	0.7, **8.45,** 30.7	25, **35**, 45	Dollars/ Tonne CO2 (2000)
Total Costs	Discounted Abatement Costs	0, **0.04**, 0.15	-0.3, **0.06**, 0.3	2.1, **6.0**, 13.2	0.8, **3.2,** 8.4	0.17, **5.8,** 13.4	1.1, **3.2,** 7.0	0.5, **1.8,** 4.4	-0.8, **3.1,** 8.1	0.8, **2.2**, 5.2	0.4, **1.3,** 3	0.7, **2.3,** 6.0	Trillion dollars (2000)
	Total Costs	4.7, **18,** 51	4.5, **19,** 51	7.3, **17.8,** 37.5	5.4, **15,** 34	6.1, **17.8,** 38.5	5.9, **16.3,** 37.9	5.2, **15.0,** 35.5	5.0, **16.2,** 39.0	5.6, **17**, 41	5, **15.7,** 40.5	5.1, **17**, 41	Trillion dollars (2000)
	NPV of Abatement	0	0	-8.5, **0.6,** 15.1	-4.4, **3.25,** 19.0	-8.1, **0.9,** 16.3	-3.7, **1.8,** 13	-1.0, **3.2,** 14.3	-4.2, **2.0,** 14.2	-2.4, **1.5**, 9.7	-1, **2.2,** 10.9	-3.2, **1.4,** 10	Trillion dollars (2000)

Boxes shaded in grey represent the standard model's cost factors. Due to changes in the way costs are calculated in the modified PAGE2002 model, these two can not be directly compared. The three numbers in each cell represent the 5%, Mean, and 95% levels of the model's output. The 13% learning rate and 19% learning rate correspond to a learning coefficient of 0.2 and 0.3 consecutively. As with the other parts of the model, the coefficient was specified as a distribution and not as a best guess, with the higher of the two having a triangular distribution of mean 0.3 with a minimum value of 0.06 and maximum value of 0.54.