

**A Comparison of Cost-Containment Instruments
for US Carbon Reduction Policies**

by

Lisa M. Jakobovits

B.S., Science, Technology and Society
Stanford University, 2004

Submitted to the Engineering Systems Division
in Partial Fulfillment of the Requirements for the Degree of
Master of Science in Technology and Policy

at the

Massachusetts Institute of Technology

June 2008

© 2008 Massachusetts Institute of Technology
All rights reserved

Signature of Author.....

Engineering Systems Division
May 8, 2008

Certified by.....

Mort D. Webster
Visiting Professor of Earth Atmosphere and Planetary Sciences
Thesis Supervisor

Certified by.....

A. Denny Ellerman
Senior Lecturer Sloan School of Management
Thesis Supervisor

Accepted by.....

Dava J. Newman
Professor of Aeronautics and Engineering Systems
Director of Technology and Policy Program

**A COMPARISON OF COST-CONTAINMENT INSTRUMENTS
FOR US CARBON REDUCTION POLICIES**

by

LISA M. JAKOBOVITS

Submitted to the Engineering Systems Division

On May 8, 2008

In partial fulfillment of the requirements for the degree of
Master of Science in Technology and Policy

ABSTRACT

A cap-and-trade program, as is used in the European Trading Scheme, is currently the most widely discussed method in the US for reducing greenhouse gases. A basic cap-and-trade program operates by mandating a fixed level of emissions for a given period, issuing permits, and then allowing a market for those permits to develop. The resulting market price for emissions permits, and hence the economic impacts of the chosen policy, can only be estimated in advance with a high degree of uncertainty. Many of the current US cap-and-trade proposals contain provisions for cost-containment instruments which reduce the possible range of emissions prices. This paper analyzes the relative effectiveness of three such cost-containment instruments, including a safety valve, an intensity target, and banking and borrowing. The results presented rely on two computable general equilibrium models developed at the Massachusetts Institute of Technology, and show the predicted performance of these instruments under a simulated range of economic outcomes.

Thesis Supervisor: Mort D. Webster

Title: Visiting Professor of Earth Atmosphere and Planetary Sciences

Thesis Supervisor: A. Denny Ellerman

Title: Senior Lecturer Sloan School of Management

ACKNOWLEDGEMENTS

Special thanks to my advisor, Mort Webster, as well as to Denny Ellerman, Angelo Gurgel and Ian Sue Wing for their indispensable guidance and collaboration. Thanks also to the Doris Duke Charitable Foundation for providing funding for this work.

TABLE OF CONTENTS

1. INTRODUCTION.....	11
1.1 Background.....	13
1.1.1 Why Cap-and-Trade?.....	16
1.1.2 Cost-Containment Instruments	16
1.2 Prior Work.....	18
1.3 Research Questions	20
1.4 Summary of Conclusions	21
1.5 Outline of Chapters	21
2. SINGLE-PERIOD STUDY: METHODS AND SENSITIVITY ANALYSIS	23
2.1 Model Description	23
2.2 Parameterization of Uncertainty	25
2.3 Policy Design.....	26
2.4 Sensitivity Analysis.....	28
2.4.1 Marginal Abatement Curves under Uncertainty	28
2.4.2 Tax under Uncertainty	29
2.4.3 Cap under Uncertainty	31
3. SINGLE-PERIOD STUDY: RESULTS AND ANALYSIS	35
3.1 Results	35
3.1.1 Carbon Price	35
3.1.2 Abatement.....	37
3.1.3 Emissions.....	38
3.1.4 Welfare	39
3.1.5 Net Benefits	40
3.2 Safety Valve vs. Intensity Target	43
3.2.1 Safety Valve Design	43
3.2.2 Intensity Target Design: Correlation Effect.....	46
3.2.3 Combined Instrument: Intensity Target with Safety Valve	47
3.3 Uncertain Marginal Benefits	51
3.3.1 Constant Marginal Benefits	51
3.3.2 Non-Constant Marginal Benefits	55
3.4 Conclusions	64
4. MULTI-PERIOD STUDY: METHODS.....	67
4.1 Model Description	67
4.2 Parameterization of Uncertainty	70
4.3 Policy Design.....	72
4.3.1 Effect of Banking.....	73
4.3.2 Emissions Cap (No Cost Containment)	75
4.3.3 Safety Valve	75
4.3.4 Compensated Safety Valve.....	75

4.3.5	Borrowing	76
4.4	Policy Behavior under a Single Price Shock	76
5.	MULTI-PERIOD STUDY: RESULTS AND ANALYSIS	81
5.1	Results of Monte-Carlo Simulation	81
5.1.1	Carbon Price	81
5.1.2	Emissions	82
5.1.3	Abatement	84
5.1.4	Cumulative Emissions	85
5.1.5	Welfare	86
5.2	Borrowing vs. Compensated Safety Valve	87
5.2.1	Optimal Trigger Price	89
5.3	Borrowing vs. Safety Valve	91
5.3.1	Constant Net Benefits	91
5.3.2	Non-Constant Net Benefits	95
5.4	Conclusions	98
6.	CONCLUSIONS	101
6.1	Policy Implications	103
6.2	Suggestions for Further Work	114
7.	APPENDIX I – SENSITIVITY ANALYSIS OF SINGLE-PERIOD CGE MODEL	117
8.	APPENDIX II – DETERMINATION OF OPTIMAL R*	123

LIST OF TABLES

Table 1: Reference values of substitution elasticities across each aggregated sector	25
Table 2: Distribution of uncertain parameters used in CGE model, with values for the elasticities of substitution represented as a fraction of reference	25
Table 3: Relative Contribution of Uncertainties to the 5-95% range of carbon price and abatement under a tax and a cap	30
Table 4: Expected Abatement, Carbon Price, Welfare Loss, and Net Benefits under a Tax, Cap, Safety Valve and Intensity Target.....	35
Table 5: Net benefits under the intensity target when correlation between baseline emissions and GDP growth is varied.....	46
Table 6: Performance of policy instruments given varied assumptions about flat marginal benefits	55
Table 7: Performance of policy instruments given varied assumptions about non-constant marginal benefits	64
Table 8: Regions, Sectors and Factors in the MIT forward-looking EPPA model	69
Table 9: Mean, Min and Max values for carbon price, emissions, abatement, cumulative emissions and overall welfare for each policy for a range of cost shocks in 2005	81

LIST OF FIGURES

Figure 1: Emissions under current Senate proposals.....	12
Figure 2: Expected CO ₂ prices under current Senate proposals.....	12
Figure 3: Effect of Uncertainty on Price and Quantity Instruments.....	14
Figure 4: Effect of Relative MC and MB slopes on Choice of Instrument.....	15
Figure 5: Policy Design under Certainty.....	27
Figure 6: Marginal Cost of Abatement under Uncertainty.....	28
Figure 7: PDFs of abatement (top) and emissions (bottom) in 2015 under a tax, showing the relative contribution of each uncertain parameter.....	29
Figure 8: PDFs of carbon price (top) and abatement (middle) in 2015 under a cap, and baseline emissions (bottom), showing relative contribution of each uncertain parameter.....	32
Figure 9: Relationship between carbon price and % abatement under a tax for each uncertain parameter.....	34
Figure 10: CDF of carbon price (marginal cost) under uncertainty for a cap, tax, safety valve and intensity target.....	36
Figure 11: CDF of abatement under uncertainty for a cap, tax, safety valve and intensity target.....	37
Figure 12: CDF of total emissions under uncertainty.....	38
Figure 13: CDF of welfare loss under uncertainty for a cap, tax, safety valve, and intensity target.....	39
Figure 14: Trade-off between expected abatement and expected welfare loss.....	40
Figure 15: CDF of net benefits under uncertainty for a cap, tax, safety valve and intensity target.....	41
Figure 16: Relative advantage of intensity target or safety valve given deviation in marginal cost from optimal value of \$23/ton.....	43
Figure 17: Expected net benefits of safety valve after varying the quantity instrument and price instrument one standard deviation above and below reference.....	45
Figure 18: Relative performance of intensity target to safety valve when varying the correlation between emissions (regulated quantity) and GDP (indexed quantity).....	47
Figure 19: CDF of carbon Price under the “Intensity Target with Safety Valve” (I + SV) policy.....	48
Figure 20: CDF of abatement under the “Intensity Target with Safety Valve” (I + SV) policy.....	49
Figure 21: Abatement-welfare trade-off under the “Intensity Target with Safety Valve” (I + SV).....	49
Figure 22: Net benefits under each policy when all are equivalent in the “no-uncertainty” case.....	50
Figure 23: Flat marginal benefits defined as half and twice the reference tax.....	52
Figure 24: Net Benefits of Cap, SV, Intensity and I+SV as a fraction of Net Benefits under a Tax under constant but varied marginal benefits.....	53
Figure 25: Absolute Net Benefits (relative to tax) under constant but varied marginal benefits.....	54
Figure 26: Net Benefits relative to reference SV given variations in cap and trigger price under varied but constant marginal benefits.....	55
Figure 27: Sub-optimal performance of a cap when BAU emissions are uncertain and MBs are represented in terms of abatement.....	56
Figure 28: Marginal Costs in 2015 in terms of Emissions.....	57
Figure 29: Non-Constant Marginal Benefits (in terms of emissions).....	58

Figure 30: Net Benefits of Cap, SV, Intensity and I+SV as a fraction of Net Benefits under a Tax under non-constant marginal benefits	59
Figure 31: Net Benefits relative to reference SV given variations in cap and trigger price under non-constant marginal benefits.....	60
Figure 32: Optimal Design of Intensity Target when marginal benefits are sloped, assuming GDP and Emissions are perfectly correlated.....	62
Figure 33: Fully-indexed intensity target vs. optimal partially indexed target when the slopes of marginal costs and benefits are equivalent.....	63
Figure 34: Performance of optimal intensity target (intensity*) and fully indexed target (intensity) relative to other policy instruments.....	64
Figure 35: Distribution (normal) of shocks to fixed factor oil resources; only positive cost shocks (increases in oil supply) are considered.....	71
Figure 36: Effect of increase in oil resources on carbon price	72
Figure 37: Emissions Cap with (adjusted cap) and without (non-adjusted cap) banking.....	73
Figure 38: Effect of banking given a temporary shock in 2005	74
Figure 39: Effect of banking on discounted carbon price (in 2005 dollars).....	74
Figure 40: Emissions from 2005-2050 in response to a single temporary shock	77
Figure 41: Emissions in 2005 and 2010 under each policy in response to a single temporary shock	78
Figure 42: Carbon Price under each policy in response to a single temporary shock	79
Figure 43: Discounted carbon price (in 2005 dollars) in response to a single temporary shock.....	79
Figure 44: Welfare reduction from reference in response to a single temporary shock	80
Figure 45: CDF of 2005 carbon price under uncertainty for the no cost-containment, borrowing, safety valve, and compensated safety valve policies	82
Figure 46: CDF of 2005 BAU carbon emissions under uncertainty	83
Figure 47: CDF of 2005 emissions under uncertainty for the no cost-containment, borrowing, safety valve, and compensated safety valve policies	83
Figure 48: Marginal cost curve in terms of emissions in the no shock and median shock cases.....	84
Figure 49: CDF of 2005 carbon abatement under uncertainty for the no cost-containment, borrowing, safety valve, and compensated safety valve policies	85
Figure 50: Marginal cost curve in terms of abatement in the no shock and median shock cases.....	85
Figure 51: CDF of cumulative carbon emissions (from 2005-2050) under uncertainty for the no cost-containment, borrowing, safety valve, and compensated safety valve policies	86
Figure 52: CDF of overall welfare (from 2005-2050) under uncertainty for the no cost-containment, borrowing, safety valve, and compensated safety valve policies	87
Figure 53: Resulting 2005 carbon price given trigger prices above and below reference.....	89
Figure 54: Reduction in expected welfare (2005-2050) from no-policy case, given compensated safety valve trigger price values above and below reference	90
Figure 55: Welfare index (2005-2050) under uncertainty given compensated safety valve trigger price values above and below reference	91
Figure 56: Ratio of expected net benefits under safety valve to net benefits under borrowing, when marginal benefits are flat and equal to the reference carbon price	92
Figure 57: Flat marginal benefits (2005) defined as half and twice the no-shock carbon price	93
Figure 58: Expected net benefits of SV/Borrowing given a 5% discount rate	94

Figure 59: Expected net benefits of SV/Borrowing given constant MBs over a range of discount rates	95
Figure 60: Sloped marginal benefits (2005) defined as half and twice the slope of the marginal cost curve at the point of intersection.....	96
Figure 61: Expected net benefits of SV/Borrowing given sloped MBs over a range of discount rates	97
Figure 62: Efficient performance of borrowing (relative to safety valve) in 2005, and optimal performance of safety valve in subsequent years (2010-2050)	98
Figure 63: Sensitivity of no-policy emissions to AEEI, GDP growth, and individual elasticities of substitution (doubled and halved)	117
Figure 64: Sensitivity of no-policy emissions to individual elasticities of substitution (doubled and halved)	118
Figure 65: Sensitivity of emissions under a \$50 carbon tax to AEEI, GDP growth, and individual elasticities of substitution (doubled and halved)	118
Figure 66: Sensitivity of equivalent variation (welfare) under a \$50 carbon tax to AEEI, GDP growth, and individual elasticities of substitution (doubled and halved)	119
Figure 67: Sensitivity of equivalent variation (welfare) under a 5000mmt emissions cap to AEEI, GDP growth, and individual elasticities of substitution (doubled and halved)	119
Figure 68: Sensitivity of equivalent variation (welfare) under a 5000mmt emissions cap to individual elasticities of substitution (doubled and halved)	120
Figure 69: Sensitivity of emissions under a \$50 carbon tax to pairings of elasticities of substitution (both doubled and halved)	121
Figure 70: Sensitivity of equivalent variation (welfare) under a \$50 carbon tax to pairings of elasticities of substitution (both doubled and halved)	122

1. INTRODUCTION

A cap-and-trade program is currently the most widely discussed approach for a carbon reduction policy in the US (Paltsev et al., 2007), and is already in use in the European Emissions Trading Scheme (EU ETS). However, there is much agreement among economists that this policy is not the most efficient choice. Though much of the climate policy debate is centered on the potential consequences of undertaking too little (or no) abatement to address the problem, a cap-and-trade program may require more abatement than necessary – and at far greater cost than benefit. In the US specifically, policy design is crucial to ensuring the success of any carbon reduction policy. Unexpectedly high costs not only restrain economic growth, but may also inhibit the long-term viability of the policy. Cap-and-trade programs without built in cost-containment measures are particularly vulnerable to unexpectedly high costs (Pizer, 1997).

A cap-and-trade program operates by distributing or auctioning emissions permits to regulated industries, and then allowing a free market for those permits to develop. Currently, two of the most widely discussed cap-and-trade proposals are the Bingaman-Specter Low Carbon Economy Act of 2007 (S.1766) and the Lieberman-Warner Climate Security Act of 2008 (S.2191) (Environmental Protection Agency, 2008). The emissions levels stipulated by these bills, from 2010-2050, are shown in Figure 1. The Bingaman-Specter bill specifies emissions targets only through 2030 (solid line), however it recommends a 60% reduction below 1990 levels by 2050 (dashed line) if the five largest international trading partners of the US take similar action against climate change (Environmental Protection Agency, 2008).

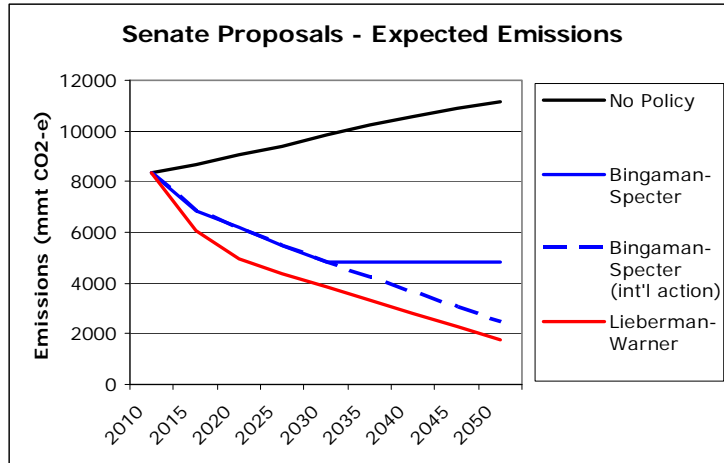


Figure 1: Emissions under current Senate proposals

(Source: EPA Analysis of the Lieberman-Warner Climate Security Act of 2008)

Although the quantity of emissions in each year is set by the cap under these policies, the resulting cost of emissions permits is determined by the market and cannot be known in advance. The expected price under the Lieberman-Warner bill is shown by the solid red line in Figure 2.

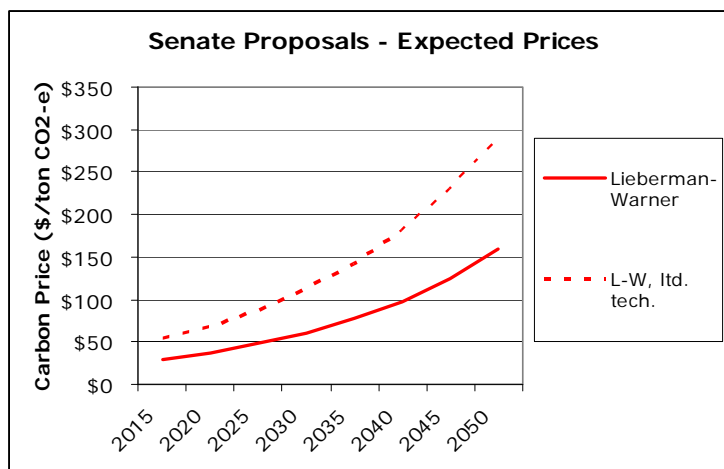


Figure 2: Expected CO₂ prices under current Senate proposals

(Source: EPA Analysis of the Lieberman-Warner Climate Security Act of 2008)

In terms of consumer effects, the \$159 carbon price in 2050 is predicted to add about \$68 to the cost of a barrel of oil, or about \$1.40 to the price of a gallon of gas (Environmental Protection Agency, 2008). This number may not sound unreasonable in

light of current gas prices, but is also far from certain. Constraining assumptions about the future availability of low-carbon technologies, specifically nuclear energy, biomass, and carbon capture and sequestration (CCS), increases the expected carbon price by 86% to the dashed red line shown in Figure 2. This begs the question: is the higher price too high? If not, what is? The question of how much we should be willing to pay is one of economics, specifically of the trade-off between the benefits of abatement and cost. To better understand this issue and how it relates to climate policy I first turn to a discussion of two policy instruments – a cap and a tax – and analyze their behavior under uncertainty.

1.1 BACKGROUND

Several studies have shown that a more efficient instrument than a cap-and-trade is a carbon tax (Weitzman, 1974, Pizer, 1997), which sets a fixed price on carbon emissions but yields an uncertain quantity of emissions. The reason has to do with the performance of these instruments under uncertainty. If the marginal costs of carbon reduction are not uncertain (i.e. they are known at the time the policy is designed), then a tax and cap will result an identical emissions level and carbon price (Figure 3, left). However, if marginal costs are uncertain and deviate from expectation, a tax and cap behave differently (Figure 3, right). A cap fixes the emissions level with the resulting price determined by the actual marginal cost curve. Conversely, a tax fixes the price with the resulting quantity is determined by the actual marginal cost curve. As seen in Figure 3 (right), neither of these instruments achieves the optimal price or quantity under uncertainty. However, as Weitzman (1974) demonstrated, the policy that minimizes

deadweight loss can be determined by the relative slopes of the marginal costs and benefits.

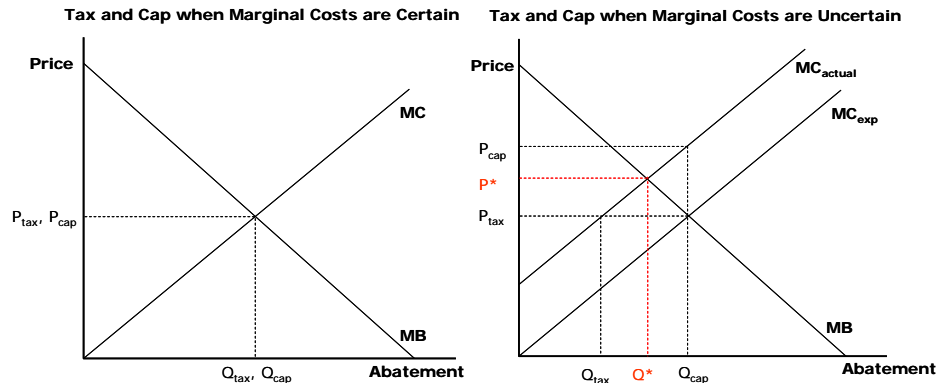


Figure 3: Effect of Uncertainty on Price and Quantity Instruments

Weitzman showed that when the marginal costs of abatement are relatively steeper than the marginal benefits, a tax is preferable to a cap. This result is demonstrated in Figure 4 (left). If the marginal cost of abatement is greater than expected, a tax (price instrument) comes much closer to the optimal target than does a cap (quantity instrument). The welfare loss under each policy is shown in Figure 4 (left), with the small triangle representing the loss under a tax, and the larger triangle representing the welfare loss under a cap. When marginal benefits are steeper than marginal costs, the opposite holds (Figure 4, right). Here, the cap is closer than a tax to the optimal target, and the welfare loss under the cap (small triangle) is less than the welfare loss under a tax (larger triangle). Uncertainty in marginal benefits has no effect on the optimal instrument choice (Weitzman, 1974), except in the case where marginal damages are correlated with marginal costs (Stavins, 1996).

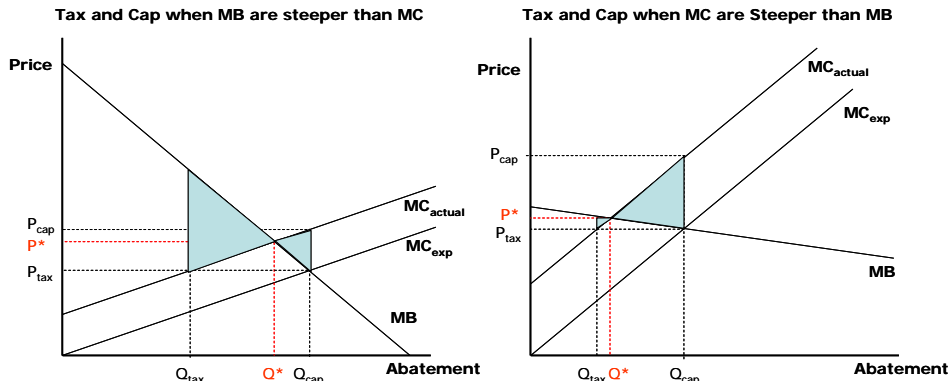


Figure 4: Effect of Relative MC and MB slopes on Choice of Instrument

In regard to carbon emissions, the current consensus holds that the marginal costs of abatement are steeper than the marginal damages (Pizer, 1999, Pizer, 2002), which leads to a preference for a tax over a cap. The intuition is that marginal costs relate to the flow (annual carbon emissions) while benefits from carbon abatement relate to the stock (total carbon in the atmosphere). That is, in a given year it is increasingly expensive to reduce emissions by an additional unit, but the benefit of those reductions is relatively constant because annual emissions have only a small effect on the total stock of carbon in the atmosphere. Work by Pizer (1997) shows the expected welfare gain under a tax to be at least five times that under a cap.

The Weitzman model is oversimplified, however, in that it concerns only a single period and a single cost shock. In reality, a carbon reduction policy spans multiple periods and is subject to both temporary shocks (uncorrelated across periods) and permanent shocks (correlated across periods). Temporary shocks might be caused by weather episodes or labor strikes, whereas the rates of economic growth and technology development are arguably more permanent (Parsons, 2007). The Weitzman model is analogous to considering only permanent cost shocks, thus consideration of temporary shocks complicates the problem (Parsons, 2007). Several authors, including Hoel and

Karp (2002), Newell and Pizer (2003), and Karp and Zhang (2005) have modeled this more complex structure of uncertainty, and all have found a tax to be preferable to a cap. However, Parsons (2007) has noted that these models do not allow banking and borrowing across time (a form of cost-containment), which add flexibility to the cap and may help to counteract temporary cost shocks.

1.1.1 Why Cap-and-Trade?

Given the more efficient performance of a price instrument for CO₂ reductions under uncertainty, one may wonder why there is a strong preference in the US for a cap-and-trade policy. The issue is largely one of political feasibility; tax hikes are not generally favored in the US political system (The Washington Post, 2007). A fixed or declining emissions cap has also been used in the US to control other types of pollution, including SO₂ (Ellerman and Montero, 2002) and NO_x (Pizer, 2005), and is also currently used in the EU Emissions Trading Scheme. Additionally, most US cap-and-trade proposals call for the majority of permits to be freely allocated, rather than auctioned, to regulated entities (Paltsev et al., 2007). A tax requires polluters to pay for every unit of emissions, whereas permits allocated for free entail no additional cost to firms up to the level of permits provided.

1.1.2 Cost-Containment Instruments

The political preference for a cap-and-trade program, along with the economic drawbacks associated with it, has led to a discussion of cost-containment instruments. Such instruments would be built in to the cap and trade program and would reduce the

potential for high costs. Recent carbon reduction proposals have specifically included a safety valve, intensity target, or banking and borrowing in conjunction with a cap-and-trade program (Pizer, 2005). A safety valve is essentially a cap-and-trade program with a ceiling price. If the market price of permits exceeds some set trigger price an unlimited number of additional permits are sold at that price, thus reverting to a carbon tax at the level of the trigger price (Pizer, 2005, Jacoby and Ellerman, 2004). An intensity target allows the cap in a given year to fluctuate relative to some other measurable quantity that is correlated with emissions, for example gross domestic product (GDP) (Newell and Pizer, 2006; Ellerman and Sue Wing, 2003; Sue Wing et al., 2006). Banking and borrowing allows regulated entities to shift their obligations across time periods (Paltsev et al., 2007). If a firm expects the cost of compliance to increase in the future at a rate greater than the discount rate, it may abate more in the current period, or purchase additional current period permits, and “bank” the unused permits for future use. Conversely, if a firm expects costs to decrease in the future it may abate less in the current period and use or purchase future period permits for current period compliance, thereby “borrowing” from future period allocations.

Prior work has shown the ability of a safety valve (Weitzman, 1978, Pizer, 2002, Jacoby and Ellerman, 2004, Roberts and Spence, 1976), intensity target (Ellerman and Sue Wing, 2003, Sue Wing et al., 2006, Newell and Pizer, 2006) and banking/borrowing (Pizer, 2005, Paltsev et al., 2007) to provide flexibility that reduces cost uncertainty relative to an absolute cap. For policymakers, this work implies that concerns about unexpectedly high costs can be effectively addressed by these cost-containment measures. However, the question of how these instruments perform relative to each other

remains unanswered, and is the focus of this thesis. Building on previous work, this study uses an empirical model to quantify the economic advantages of each of these instruments relative to a pure cap. It also poses a head-to-head comparison of each of these cost-containment instruments, and assesses whether one is best suited to control economic costs under a cap-and-trade program.

1.2 PRIOR WORK

Several studies have shown the safety valve to be preferable to an absolute cap and, in some cases, to a tax. Roberts and Spence(1976), and Weitzman (1978) showed analytically that under uncertain marginal costs and non-linear marginal damages, the optimal policy that combines price and quantity controls (such as a safety valve) performs better than the optimal tax or cap. Pizer (2002) used a single-period and a multi-period model to empirically determine the optimal cap, tax, and safety valve policies for the year 2010. He found the optimal tax to be substantially better than a pure cap, and the optimal safety valve to be only marginally better than a pure tax. Jacoby and Ellerman (2004) argue that the safety valve may be a useful short-term cost-containment measure, but that it is unlikely to be successful as a long-term instrument. They note that a main criticism of a safety valve, particularly from environmental groups, is that safety valve relaxes the emissions target when costs are unexpectedly high. Though there is not currently evidence that we approaching a “critical threshold” of atmospheric greenhouse gases (Jacoby and Ellerman, 2004), a safety valve does not fix the level of cumulative emissions allowed under the policy.

Other studies have examined an intensity target as an alternative to a pure cap. An intensity cap, like a safety valve, acts to control costs by relaxing the quantity target when costs are unexpectedly high, but does so by scaling the emissions cap in relation to GDP. If GDP growth is greater than expected, an intensity target requires less abatement and thus a lower cost than would an absolute cap. However, if GDP growth is lower than expected, an intensity target requires more abatement and higher cost than would an absolute cap (Ellerman and Sue Wing, 2003). Quirion (2005) studied the performance of an intensity target given two types of uncertainty: the slope of the marginal cost curve and business-as-usual (BAU) emissions. He found that an intensity target only yields higher expected welfare than a tax if the slope of marginal costs is roughly equal to the slope of marginal damages or if uncertainty about BAU emissions is extremely high. Newell and Pizer (2006) found that the performance of an intensity target is dependent on the correlation between the regulated quantity and the indexed quantity; in the case, the correlation between BAU emissions and GDP.

Banking and borrowing allow the cumulative abatement required by an absolute cap to be optimized over time. Paltsev et al (2007) found that under the recent cap-and-trade proposals, abatement efforts tend to be reallocated to the near term so that less stringent abatement is required later on. This analysis showed that banking results in a slight improvement in the net present value (NPV) of welfare in comparison to an absolute cap, but did not explicitly consider uncertainty or use a model that could perfectly optimize through time.

1.3 RESEARCH QUESTIONS

Many studies have examined each of these cost containment instruments and identified conditions under which each would be preferred to a cap, a tax, or both. However, the question more relevant to current climate policy is not whether to use a cost-containment instrument but rather which one to choose. The Lieberman-Warner and Bingaman-Specter bills differ significantly in their proposals for cost-containment instruments. The Bingaman-Specter bill specifies a Technology Accelerator Payment (TAP), which is equivalent to a safety valve in that places an upper limit on the market price of permits. This upper limit is to start at \$12/ton of CO₂ equivalent GHG reductions (CO₂-e) in 2012, and rises at a real rate of 5% each year (United States Senate Committee on Energy and Natural Resources, 2007). The Lieberman-Warner bill proposes to contain costs by allowing flexibility over time. Regulated entities can borrow permits from future years to account for up to 15% of their yearly emissions (S.2191).

A comparison of these bills thus requires a comparison of a safety valve to borrowing. In light of that, this thesis proposes to answer the following questions:

- (1) *How much do a safety valve, intensity target, and banking/borrowing improve expected welfare relative to an absolute cap?*
- (2) *Under what conditions is each of these instruments preferred to the others?*

It is important to note that if marginal costs were known with certainty, each of these instruments performs identically. Thus, to compare them head-to-head I must test their response to the distribution of possible future costs. I first use a single-period model of the year 2015, and test the relative performance of a safety valve and intensity target in

response to three combined types of uncertainty: GDP growth, autonomous energy efficiency improvements (AEEI) and elasticities of substitution. I then use a multi-period model of the years 2005-2050 which is able to fully optimize over time to test the performance of a safety valve relative to banking and borrowing provisions.

1.4 SUMMARY OF CONCLUSIONS

The single-period analysis finds that if both the safety valve and the intensity target are designed to be optimal in the reference (no uncertainty) case, the safety valve is preferred to the intensity target in all but a very narrow range of conditions. It also finds that combining the safety valve with the intensity target may be preferable to either instrument individually.

The multi-period analysis finds that the relative performance of a safety valve or borrowing provision depends, in part, on whether it is essential to achieve the proposed cumulative emissions cap. If the cumulative cap must be met, borrowing is the preferred instrument. The compensated safety valve performs almost as well as borrowing if the trigger price is set to the optimal level, but welfare is substantially reduced if the trigger price is set too low. If the cumulative cap is flexible, the traditional safety valve outperforms borrowing under most assumptions about net benefits.

1.5 OUTLINE OF CHAPTERS

Chapter 2 provides a description of the single-period model and explains the choice and parameterization of uncertainty. It also describes the design of the four tested

instruments – a tax, cap, safety valve and intensity target – which perform differently under uncertainty but identically in the no-uncertainty case.

Chapter 3 presents the results of the model and provides a direct comparison of the safety valve to the intensity target. It finds their relative performance to be affected by both the design of the safety valve and the correlation between no-policy emissions and GDP. It also tests a fifth policy instrument which is found to be superior in terms of net benefits: an intensity target with a safety valve.

Chapter 4 describes the multi-period model, which is a forward-looking version of the MIT Emissions Prediction and Policy Analysis (EPPA) model. The chapter also explains the parameterization of temporary (one year only) shock used in the model, and describes the four policies tested – a cap, safety valve, borrowing, and compensated safety valve (a safety valve with a cumulative cap over the policy horizon).

Chapter 5 presents the results of the multi-period model, and provides a comparison of the safety valve and borrowing. It finds that borrowing is always preferable to the compensated safety valve, but that the compensated safety valve can closely mimic borrowing if the trigger price is set correctly. This section also compares the traditional safety valve to borrowing on the basis on net benefits, and finds the safety valve to be preferable if marginal benefits are relatively flat.

Chapter 6 concludes, gives policy recommendations, and suggests areas for further study.

2. SINGLE-PERIOD STUDY: METHODS AND SENSITIVITY ANALYSIS

In this chapter, I describe the single-period static computable general equilibrium (CGE) model, which is used in Chapter 3 to analyze the relative performance of the intensity target and safety valve. The CGE model allows for a more detailed representation of the marginal costs of abatement than would an analytical model, and additionally allows for the treatment of several types of economic uncertainty. Here, economic uncertainty refers only to the marginal costs of reducing carbon emissions, as opposed to the marginal benefits.

This chapter begins with a description of the single-period model used in the analysis, and follows with a discussion of the how economic uncertainty is simulated within the model. It concludes by examining the sensitivity of emissions and carbon price under a cap and tax to each type of uncertainty.

2.1 MODEL DESCRIPTION

The model treats households as an aggregate representative agent with constant elasticity of substitution (CES) preferences. Industries are consolidated into 11 sectoral groupings, and are treated as representative firms with nested CES production technology. Bovenberg and Goulder's (1996) KLEM production technology and parameterization are adapted for this purpose.

The model's algebraic structure is numerically calibrated using U.S. data on inter-industry economic flows, primary factor demands, commodity uses and emissions in the

year 2004. Prices, economic quantities, and emissions of CO₂ in the year 2015 are simulated by scaling both the economy's aggregate factor endowment and the coefficients on energy within industries' cost functions and the representative agent's expenditure function. The probability distributions of these scaling factors are used as inputs into the model, and give rise to probability distributions for the future value of baseline national income, energy use and emissions.

The elasticities of substitution considered in this analysis are those for production, capital-labor, energy-materials, inter-fuels, fixed factors, and materials (denoted q , kl , em , e , f , and m , respectively). Probability distributions for these six parameters, when propagated through the model, generate probability distributions for the changes in income and emissions from their baseline levels in response to climate policy. The model also makes best-guess assumptions the output transformation elasticity between production for domestic and export markets (t), the armington substitution elasticity between domestic and imported goods (a), and the substitution elasticity between fossil and non-fossil resources ($f-nf$). The reference value for each of the substitution elasticities, across each sector classification, is shown in Figure 1. These values are derived and explained in detail in Webster, et al. (2007). Sensitivity of emissions, carbon price, and equivalent variation (welfare) to each of these nine elasticities is examined in Appendix I.

aggregate sector classification	q	f	kl	em	e	m	t	a	f-nf
agriculture	0.7	0.4	0.68	0.7	1.45	0.6	4	2.31	
crude oil and gas	0.7	0.4	0.68	0.7	1.45	0.6	0.6	5	
coal mining	0.7	0.4	0.8	0.7	1.08	0.6	4	1.14	
refineries	0.7	0.4	0.74	0.7	1.04	0.6	4	2.21	
gas works and distribution	0.7	0.4	0.96	0.7	1.04	0.6	4	1	
electric power	0.7	0.4	0.81	0.7	0.97	0.6	0.1	1	8
energy intensive industries	0.7	0.4	0.94	0.7	1.08	0.6	4	2.74	
transportation	0.7	0.4	0.8	0.7	1.04	0.6	0.1	1	
manufacturing	0.7	0.4	0.94	0.7	1.08	0.6	4	2.74	
services	0.7	0.4	0.8	0.7	1.81	0.6	1	1	
rest of economy	0.7	0.4	0.98	0.7	1.07	0.6	1	1	

Table 1: Reference values of substitution elasticities across each aggregated sector

2.2 PARAMETERIZATION OF UNCERTAINTY

For this single period analysis of carbon abatement policies I consider three types of uncertainty: the GDP growth rate of the economy between 2005 and 2015, the rate of autonomous energy efficiency improvement (AEEI), and the elasticities of substitution in the production functions. A brief summary of the probability distributions for the uncertainty parameters is given below in Table 2, and a more detailed description can be found in Webster et al (2007). For the elasticities of substitution, the distributions given below are normalized in relation to the reference value. That is, the distributions in Table 2 should be multiplied by the reference values given in Table 1 to determine the actual model inputs. The distribution below applies to, and is equivalent across, each of the 11 sectors.

	AEEI (%/yr)	GDP Growth (%/yr)	Elasticities of Substitution					
			f	q	kl	em	e	m
Mean	1.0	2.8	1.0	1.1	1.1	1.0	1.0	1.0
Standard Deviation	0.4	0.8	0.5	0.5	0.6	0.4	0.5	0.5
0.025	0.2	1.2	0.1	0.2	0.1	0.3	0.2	0.1
0.05	0.3	1.4	0.2	0.3	0.2	0.4	0.3	0.2
0.25	0.7	2.3	0.7	0.7	0.6	0.7	0.7	0.7
0.5	1.0	2.9	1.0	1.1	1.1	1.0	1.0	1.0
0.75	1.3	3.4	1.3	1.4	1.6	1.3	1.3	1.4
0.95	1.7	4.1	1.8	1.8	2.3	1.7	1.8	1.9
0.975	1.8	4.4	1.9	2.0	2.5	1.8	2.0	2.0

Table 2: Distribution of uncertain parameters used in CGE model, with values for the elasticities of substitution represented as a fraction of reference

Annual GDP growth rates are modeled as a random walk with drift (Stock and Watson, 1988, Schwartz and Smith, 2000). The volatility is estimated from GDP time series data for the U.S. economy from 1970-2000 (Bureau of Economic Analysis, 2007). The reference EIA forecast growth rate of 2.8% per annum (Energy Information Administration, 2007) is used to project from 2005 to 2015. The estimated volatility results in a distribution of future growth rates with +/- one standard deviation almost identical to the EIA high and low growth cases.

The AEEI parameter has a reference (mean) value of 1.0% p.a., consistent with many other energy economic models (Azar and Dowlatabadi, 1999). The uncertainty in AEEI is assumed to be normal with a standard deviation of 0.4% based on several analyses (Scott et al., 1999, Webster et al., 2002).

The uncertainties in the elasticities of substitution are based on literature survey of econometric estimates with published standard errors. The details of this survey and the synthesis of the standard errors into a probability distribution for each elasticity are documented fully in Webster et al (2007). The empirical probability distributions for each of these parameters are summarized above in Table 2, along with representative statistics.

2.3 POLICY DESIGN

These uncertain parameters are introduced into the CGE model through a Monte Carlo simulation, drawing 1000 random samples for each parameter value. For each of the 1000 samples I impose five alternative policy constraints: (1) reference (no policy), (2) emissions cap, (3) carbon tax, (4) safety valve, and (5) intensity target. The stringency of the emissions cap is defined as the expected CO₂ abatement under the 50% reduction

by 2050 emissions path in Paltsev et al. (2007) of 2050 Mt CO₂, leaving U.S. emissions in 2015 at 5000 Mt CO₂, at a marginal cost of \$23/ton CO₂. All other policy constraints are defined such that they are equivalent under certainty; the carbon tax is \$23/ ton CO₂, the trigger price of the safety valve is the expected carbon price of \$23/ton CO₂, and the intensity target constrains the emissions/GDP ratio to be that under the emissions cap (a 70% reduction from reference). See Figure 5.




Figure 5: Policy Design under Certainty

A critical assumption in the results shown here is that the marginal benefit of CO₂ abatement in 2015 is assumed to be constant at \$23/ton CO₂, implying that all policies are optimal in the no-uncertainty case. Additionally, the tax is assumed to be optimal under uncertainty, as it sets the carbon price equal to the marginal benefits in all cases.

2.4 SENSITIVITY ANALYSIS

The overall effect of uncertainty on the marginal cost of abatement is shown below. Additionally, the effect of individual parameter on emissions, abatement, and carbon price under a tax and cap is examined.

2.4.1 Marginal Abatement Curves under Uncertainty

When all three input parameters (GDP, AEEL, and elasticities of substitution) are known with certainty, the cost of a given quantity of abatement is also certain. This result is shown by the solid line in Figure 6. As reference emissions under certainty are 7050mmt, about 2050mmt of abatement is required to achieve the 5000mmt cap. This amount of abatement has a marginal cost of \$23/ton in the no-uncertainty case.

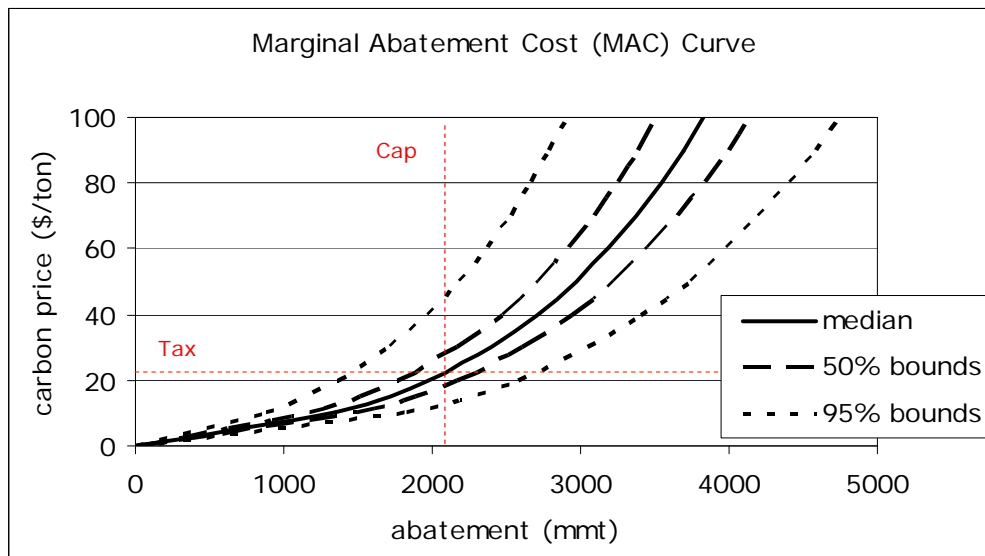


Figure 6: Marginal Cost of Abatement under Uncertainty

The dashed lines in Figure 6 show the 50% and 95% bounds of the marginal cost of abatement when all parameters are uncertain. Under uncertainty, a \$23/ton tax does not always result in 2050mmt of abatement; rather, abatement ranges from 1447mmt and

2724mmt in 95% of cases. Similarly, 2050mmt of abatement does not always require a carbon price of \$23/ton, but costs between \$7 and \$50 per ton in 95% of cases. Note that the cap fixes total emissions rather than abatement, thus abatement fluctuates under a cap as well as under a tax: 966-3207mmt and 1200-2724mmt, respectively, in the 95% range.

2.4.2 Tax under Uncertainty

Under uncertainty, a tax has a fixed marginal cost of abatement (carbon price), but uncertain abatement and total emissions.

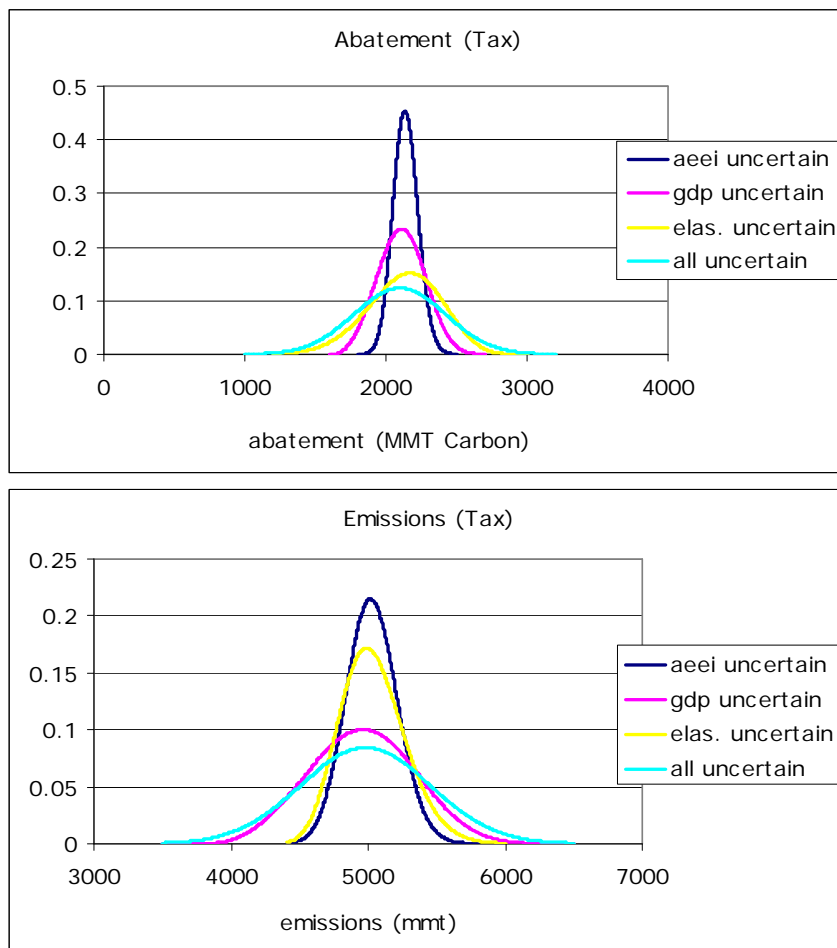


Figure 7: PDFs of abatement (top) and emissions (bottom) in 2015 under a tax, showing the relative contribution of each uncertain parameter

The carbon price under a tax is by definition \$23/ton. Possible abatement varies from roughly 1200mmt carbon to 3090mmt CO₂ (Figure 7). Uncertainty in the elasticities of substitution is the largest contributor to uncertainty in abatement, as the effectiveness of the tax is largely dependent on the cost of using less energy or switching to less carbon-intensive energy sources. Given the 5-95% range of possible abatement when all parameters are uncertain, uncertainty in only elasticities of substitution will vary abatement by 81% of this range (Table 3). GDP is the next largest contributor to uncertainty in abatement (50% of the all-uncertain range), and varies positively with the amount of abatement. AEEI is the smallest contributor and varies inversely with abatement, as more rapid autonomous efficiency improvements reduce baseline emissions.

<i>carbon price (\$/ton)</i>					
	variance	5% bound	95 % bound	range	%range
tax_all	0	22.7	22.7	0	100%
tax_aeei	0	22.7	22.7	0	100%
tax_elas	0	22.7	22.7	0	100%
tax_gdp	0	22.7	22.7	0	100%
tax_elasgdp	0	22.7	22.7	0	100%
tax_elasaeei	0	22.7	22.7	0	100%
tax_gdpaei	0	22.7	22.7	0	100%
<hr/>					
cap_all	117.1	8.4	42.4	34.0	100%
cap_aeei	16.7	16.5	30.0	13.4	39%
cap_elas	36.4	16.2	35.3	19.0	56%
cap_gdp	66.4	9.7	36.3	26.6	78%
cap_elasgdp	105.6	9.0	40.7	31.8	93%
cap_elasaeei	56.0	14.3	37.3	23.0	67%
cap_gdpaei	77.9	9.3	37.7	28.4	84%
<hr/>					
<i>abatement (\$/ton)</i>					
	variance	5% bound	95 % bound	range	%range
tax_all	103800	1546	2632	1086	100%
tax_aeei	7777	1999	2289	290	27%
tax_elas	69829	1653	2538	885	81%
tax_gdp	26853	1843	2382	539	50%
tax_elasgdp	95426	1580	2623	1043	96%
tax_elasaeei	77274	1637	2561	924	85%
tax_gdpaei	32533	1827	2417	590	54%
<hr/>					
cap_all	351737	1112	3046	1934	100%
cap_aeei	75306	1702	2604	902	47%
cap_elas	1246	2079	2196	117	6%
cap_gdp	297573	1153	2948	1795	93%
cap_elasgdp	298404	1164	2965	1800	93%
cap_elasaeei	76088	1711	2606	895	46%
cap_gdpaei	351231	1112	3036	1924	99%

Table 3: Relative Contribution of Uncertainties to the 5-95% range of carbon price and abatement under a tax and a cap

The contribution of each uncertain parameter to variation in total emissions is shown at the bottom of Figure 7. Although elasticities of substitution are the largest contributors to variation in abatement under a tax, GDP is the largest contributor to variation in total emissions. As will be discussed in more detail below, GDP is responsible for most of the variation in baseline emissions, which overshadows the variation in abatement under a tax. Uncertainty in elasticities of substitution and AEEI contribute roughly equally to uncertainty in total emissions. Though elasticities of substitution contribute more significantly to variation in abatement, AEEI contributes more to variation in baseline emissions (Figure 6), causing the difference of the two (total emissions under a tax) to be approximately equivalent.

2.4.3 Cap under Uncertainty

Carbon price, abatement and total emissions under a cap are shown in Figure 8. Under a cap, total emissions are fixed at 5000mmt, while both carbon price and abatement vary under uncertainty. The amount of abatement under a cap is determined by no-policy emissions less the emissions allowed under the cap. Uncertain parameters thus influence the distribution of possible abatement to the same extent that they influence the distribution of possible no-policy emissions (Figure 8, middle and bottom).

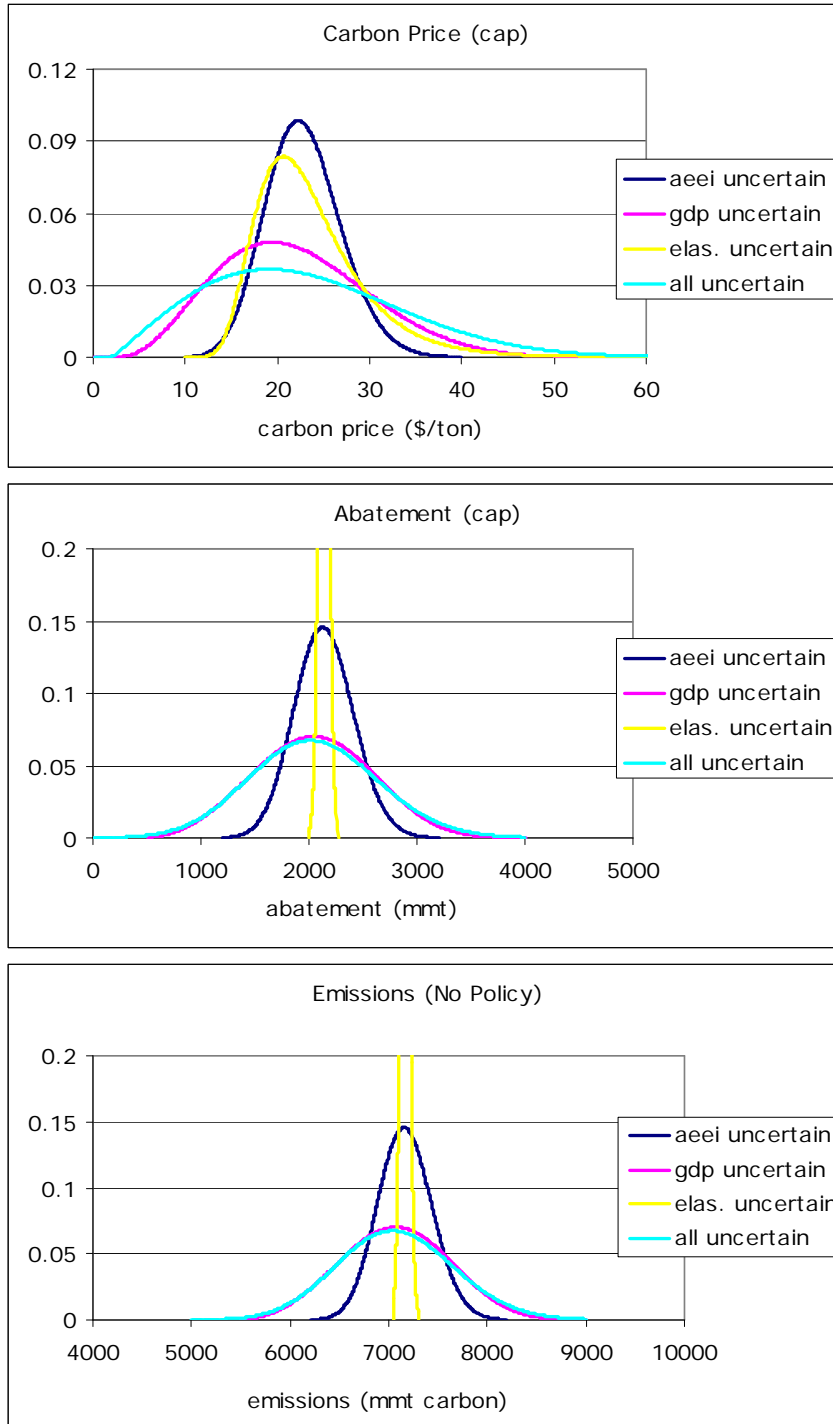


Figure 8: PDFs of carbon price (top) and abatement (middle) in 2015 under a cap, and baseline emissions (bottom), showing relative contribution of each uncertain parameter

GDP is the largest driver behind uncertainty in emissions (93% of the all-uncertain range, see Table 3), and varies positively with baseline emissions. AEEI

uncertainty is the next-largest contributor (47% of the all-uncertain range), and varies negatively with no-policy emissions. Uncertainty in elasticities of substitution is negligible (6% of the all-uncertain range), as these parameters generally matter only after substitution between energy sources is forced by policy.

Uncertainty in GDP is also the largest contributor to uncertainty in carbon price (78% of all-uncertain range, see Table 3), as its effect on no-policy emissions determines the required level and hence cost of abatement. Uncertainty in elasticities of substitution contribute significantly to carbon price (56% of all-uncertain range), as these elasticities determine how expensive it is to switch between energy sources to meet the emissions constraint.

The relationship between these uncertain parameters can be understood more intuitively by examining the cost of any given percent reduction in baseline emissions (Figure 9). GDP and AEEI affect baseline emissions, meaning they vary the fraction of abatement that is required by a policy, but not the cost associated with a given fractional reduction. Elasticities of substitution affect the slope of the marginal cost curve, meaning they determine the cost of a given fractional reduction in emissions but do not change the amount of abatement that is required. When all three parameters are uncertain both baseline emissions (horizontal intercept of the marginal cost curve) and cost (slope of the marginal cost curve) are unknown, yielding uncertainty along both dimensions.




Figure 9: Relationship between carbon price and % abatement under a tax for each uncertain parameter

3. SINGLE-PERIOD STUDY: RESULTS AND ANALYSIS

Having examined how a pure tax and cap respond to uncertain input parameters in the model, I now introduce two more politically-feasible policies: the safety valve and the intensity target. In this section, all policies are designed based on the reference (no-uncertainty) case, as explained above. This chapter first describes the relative performance of these four policies in terms of carbon price, abatement, total emissions, welfare loss, and net benefits. It then provides a direction comparison of the safety valve and intensity target in relation to three factors: the design of the safety valve, the correlation between BAU emission and GDP growth, and assumptions made about net benefits.

3.1 RESULTS

The results of the Monte Carlo simulation are summarized below in Table 4.

	Abatement (<i>mmt Carbon</i>)			Carbon Price (\$/ton)			Emissions (<i>mmt Carbon</i>)			Welfare Loss (%)			Net Benefits (<i>billion \$</i>)		
	<i>mean</i>	<i>min</i>	<i>max</i>	<i>mean</i>	<i>min</i>	<i>max</i>	<i>mean</i>	<i>min</i>	<i>max</i>	<i>mean</i>	<i>min</i>	<i>max</i>	<i>mean</i>	<i>min</i>	<i>max</i>
tax	2099	1200	3090	23	23	23	4982	3710	6410	-0.15	-0.20	-0.09	25341	14467	36491
cap	2052	344	3958	23	2	84	5027	5027	5027	-0.15	-0.62	-0.01	23231	-16678	35025
safety	1887	344	3090	19	2	23	5191	5027	6410	-0.12	-0.20	-0.01	24273	6932	36491
intensity	2108	1311	3318	24	9	77	4971	3913	6206	-0.16	-0.57	-0.04	24363	-10947	36386

Table 4: Expected Abatement, Carbon Price, Welfare Loss, and Net Benefits under a Tax, Cap, Safety Valve and Intensity Target

3.1.1 Carbon Price

A cumulative distribution function (CDF) of possible carbon prices under each of the four policies is shown in Figure 10. The carbon price is a constant \$23/ton under the

tax, while the cap results in a carbon price with a 95% range of \$7-\$50/ton. The carbon price under a safety valve is the carbon price under the cap if the market price under the cap is less than \$23/ton, otherwise the carbon price is a constant \$23/ton. Since the cap and tax used to design the safety valve are equivalent in the median case, the safety valve is expected to behave like a tax with approximately 50% probability. Lowering the trigger price or increasing the stringency of the cap will increase the probability that the safety valve will behave like a tax in a larger fraction of the cases.

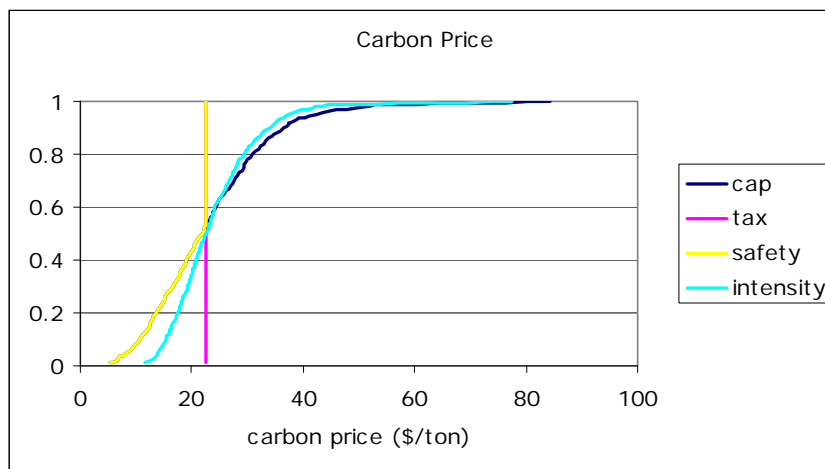


Figure 10: CDF of carbon price (marginal cost) under uncertainty for a cap, tax, safety valve and intensity target

The carbon price under an intensity target behaves similarly to that of a cap, but has lower variance (i.e., less uncertainty in marginal cost). In the bounded 95% range, the marginal cost is \$13-\$42/ton. If baseline emissions are low, the intensity target requires more abatement than a cap, and thus imposes a higher carbon price. If baseline emissions are high, the intensity target allows a more relaxed target than a cap, and thus a relatively lower carbon price. As will be shown in the next section, the intensity target behaves more like a tax in terms of abatement, but more like a cap in terms of carbon price. The intensity target “chooses” the quantity to abate by maintaining a constant emissions/GDP

ratio, which substantially narrows the range of abatement to more closely mirror a tax. However, as under a cap, elasticities of substitution have little influence in determining the amount of abatement required, but are an important factor in determining how much that abatement will cost (Figure 8). This effect causes carbon prices under both the cap and intensity target to be highly variable.

3.1.2 Abatement

A CDF of possible abatement under uncertainty is shown for each policy in Figure 11.




Figure 11: CDF of abatement under uncertainty for a cap, tax, safety valve and intensity target

As shown in Chapter 2, the cap has a larger variance in abatement than the carbon tax. If baseline emissions are higher than expected in the no-uncertainty case, a cap will require more abatement than a tax; if baseline emissions are lower than expected, a cap will require less abatement than a tax. Abatement under the intensity target instrument is very similar to that of the tax. The intensity target, like the tax, requires less abatement in the high emissions (high cost) cases, but requires more abatement in the low emissions cases. Thus, an intensity target requires additional effort in a low-growth (low-cost)

situation that under a cap would require little or no emissions reductions. In contrast, the safety valve is equivalent to the quantity instrument in the low-cost cases. For a weak enough quantity target, the cap and the safety valve allow a non-binding cap with a zero carbon price. For the high-cost cases, the safety valve reverts to a price instrument. As a result, the safety valve has the lowest expected (median) abatement of the four policy instruments (Table 4), as it combines the “easiest” cases of the cap and tax.

3.1.3 Emissions

Total emissions for each of the four policies are shown in Figure 12. The cap, by definition, has fixed total emissions at 5000mmt. Total emissions under a tax and intensity target are roughly equivalent since abatement under these policies is also similar, and vary respectively from 4036-5935mmt and 4228-5752mmt in the bounded 95% range. A safety valve maintains total emissions at a level greater than or equal to 5000mmt. If the cap allows for more total emissions than a tax, the safety valve acts as a tax; if the tax allows for more total emissions than a cap, the safety valve acts as a cap.

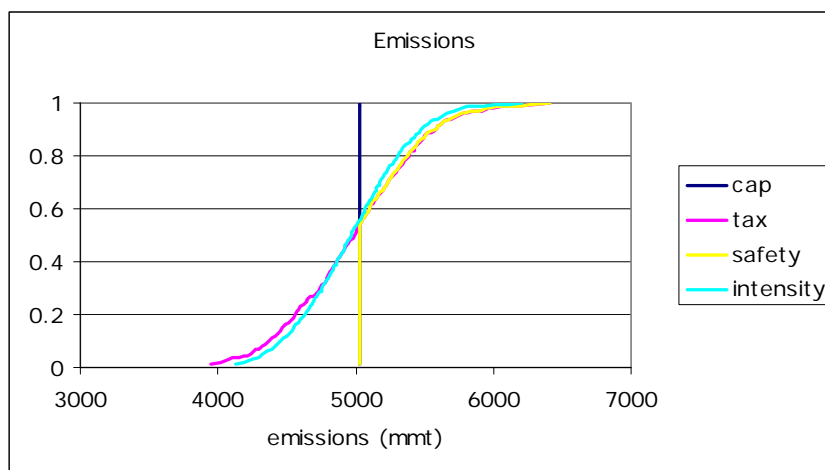


Figure 12: CDF of total emissions under uncertainty

3.1.4 Welfare

CDFs of welfare loss, measured as the percent change in equivalent variation, are shown for each policy in Figure 13. The tax has the least variability in welfare loss, and the cap has the most. The intensity target has only slightly less variability than a cap (Figure 13). The variability in welfare loss under a cap and intensity target is related to the substantial variability in carbon price under those instruments. Uncertainty in the elasticities of substitution are the largest determinant of cost (and hence welfare loss) under the cap and intensity target, but have little bearing on required amount of abatement. Thus a high level of abatement may be required when marginal costs are particularly high, or vice-versa, leading to potential extremes in welfare loss.




Figure 13: CDF of welfare loss under uncertainty for a cap, tax, safety valve, and intensity target

The safety valve is stochastically dominant (CDF is to the right of all others), meaning that the policy cost under the safety valve is less than or equal to the cost under any other instrument in all possible cases. Thus for the purposes of minimizing mitigation costs, the safety valve is preferred. However, as shown above, the safety valve also has the lowest expected abatement. In part, the safety valve has lower costs because it abates

less than other instruments. A graph depicting the relationship between welfare and abatement is shown in Figure 14, with the predictable result that increased abatement also leads to increased welfare impacts.

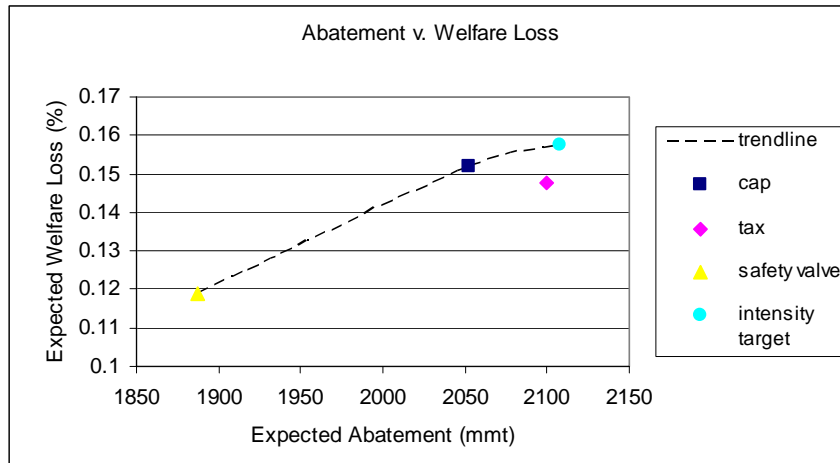


Figure 14: Trade-off between expected abatement and expected welfare loss

The only clear result from this graph is that the tax is optimal to the cap in the expected case, as the tax allows for more abatement than a tax with lower welfare cost. In order to compare these instruments on an equivalent basis, it is necessary to account for not only the costs but the also benefits of each policy instrument, to which I turn next.

3.1.5 Net Benefits

I first calculate the marginal cost curve for each of the 1000 scenarios. This is done by determining level of abatement in each scenario resulting from a series of carbon taxes (\$0 to \$200, in \$10 increments). I then fit a cubic curve to each of the 1000 scenarios, resulting in an average R^2 value of .9991. The total cost of abatement in each scenario is calculated by integrating under the curve. I assume that the marginal benefits of abatement are constant at the value of the reference tax, \$23/ton. Net benefits are calculated by subtracting total costs from total benefits. Because the tax always abates to

\$23/ton, i.e. the point where marginal cost equals marginal benefit, the tax is by definition optimal in terms of net benefits. This result is shown in Figure 15, where the CDF of net benefits under the tax stochastically dominates the CDFs for all other instruments (highest net benefits in all cases). Similarly, the cap is stochastically dominated by all other instruments. This is another way to demonstrate Weitzman's (1974) result that in the case of a constant marginal damage function, a price instrument is preferred to a quantity instrument.

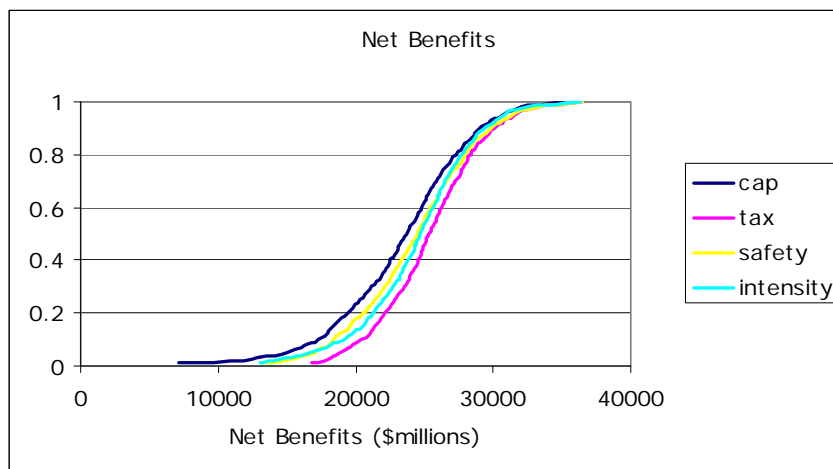


Figure 15: CDF of net benefits under uncertainty for a cap, tax, safety valve and intensity target

The CDFs of net benefits for the safety valve and the intensity target are between those of the pure price and pure quantity instrument, indicating that both are preferable to a pure cap in terms of net benefits and neither is preferred to a tax. The relative preference between the safety valve and intensity target cannot be determined from these results; the curves for these “second-best” policies cross each other, indicating that neither stochastically dominates the other.

Whether a safety valve or intensity target yields higher net benefits is determined by the relative carbon prices under the policy, as is seen in Figure 16. These carbon prices, by definition, represent the marginal cost of abatement in each scenario. When the

safety valve behaves as a price instrument, and thus imposes a tax equal to the marginal benefits, it is always optimal to the intensity target. When the safety valve behaves as a quantity instrument, and thus imposes a carbon price below the optimal level, its relative advantage depends on the corresponding price under the intensity target. If the intensity target imposes a carbon price less than \$23 (points below the dashed line), then the policy with the higher carbon price (and hence higher abatement) is preferred. If the intensity target imposes a carbon price more than \$23 (points above the dashed line), then the preferred instrument depends on the relative distance of each carbon price from the optimal tax. The situation is complicated by the non-linear slope of the marginal cost curve, but roughly, the instrument which enforces a carbon price closest to the optimal \$23/ton is preferred. For example, if the carbon price under the intensity target is \$40/ton and under the safety valve the price is \$10/ton, the safety valve is preferred because $\|23 - 10\| < \|23 - 40\|$. The preferred policy instrument in terms of expected net benefits is the one for which there is a greater density of samples in the area where it is preferred. In the results shown here, the intensity target performs slightly better than the safety valve (Table 4), but they are close enough to be statistically indistinguishable.

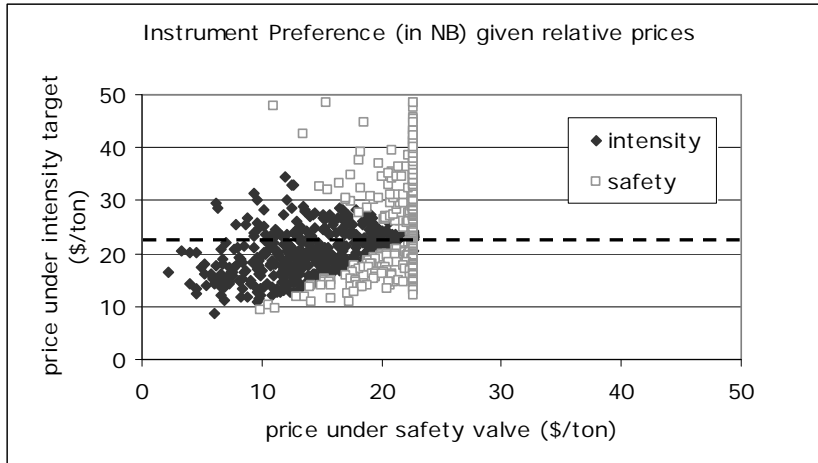


Figure 16: Relative advantage of intensity target or safety valve given deviation in marginal cost from optimal value of \$23/ton

3.2 SAFETY VALVE VS. INTENSITY TARGET

The above findings apply only when these “second-best” policies are designed to be equivalent in the no-uncertainty case, and in particular when the safety valve has a trigger price equal to level of marginal benefits. In practice, the design of these policies is likely to be somewhat arbitrary. This section explores the relative performance of the safety valve relative to the intensity target when (1) the design of the safety valve is altered, (2) the correlation between GDP and baseline emissions is varied, and (3) the safety valve and intensity target are used in combination.

3.2.1 Safety Valve Design

A safety valve has two design variables that can be changed: the cap and the trigger price. A tighter cap raises the expected market price of permits and increases the likelihood that the safety valve will act as a tax. A looser cap lowers the expected market price of permits and increases the likelihood that the safety valve will act as a cap. If the

cap is so loose that it is non-binding (i.e. allowed emissions exceed no-policy emissions), the situation is equivalent to having no policy. Increasing the other design variable, the trigger price, increases the likelihood that the safety valve will operate as a cap. If the trigger price is decreased, the safety valve is more likely to operate as a tax.

Figure 17 shows the effect on net benefits when these design variables are altered in combination. In this experiment, the “loose cap” (5500mmt) represents a one standard deviation increase in the reference cap, as determined by uncertainty in emissions under a \$23/ton tax. The “tight cap” (4500mmt) is a one standard deviation decrease in the reference cap. Similarly, the “high trigger price” (\$34/ton) is a one standard deviation increase in the reference tax, as determined by uncertainty in carbon price under a 5000mmt cap. The “low trigger price” (\$12/ton) represents a one standard deviation decrease from the reference. As seen in Figure 17, increasing the stringency of the cap yields higher net benefits. This is not surprising given that the optimal safety valve would have a quantity limit of zero, and would thus always behave as a tax. Figure 17 also shows that it is preferable in terms of net benefits to impose an overly stringent policy (tight cap or high trigger price), than to impose an overly relaxed policy (loose cap, low trigger price). The policies using either a loose cap or low trigger price are shown to be sub-optimal to a pure cap.




Figure 17: Expected net benefits of safety valve after varying the quantity instrument and price instrument one standard deviation above and below reference

Only the best-performing safety valve (tight cap, reference trigger price) attains higher expected net benefits than the reference intensity target. As noted above, the reason this design outperforms others is because the safety valve acts as a tax in the highest fraction of cases. The standard safety valve (ref cap, ref tp) does slightly worse in terms of net benefits than the intensity target, though in relation to other possible safety valve designs this difference is negligible.

Thus, the only safety valve better than an intensity target is one that is designed to act as the optimal tax in more than half of the possible scenarios. However, assumptions regarding factors affecting the performance of the intensity target may also be altered. I now turn to the question of how these factors, specifically the correlation of emissions with GDP, affect the performance of the intensity target relative to the safety valve.

3.2.2 Intensity Target Design: Correlation Effect

The relative effectiveness of an intensity target is determined by the correlation between the regulated quantity (emissions) and the indexed quantity (GDP) that determines the policy for a given period. This relationship has previously been demonstrated analytically by Newell and Pizer (2006) and Webster et al (2007). I test this relationship empirically by changing the variance in the GDP growth rate in order to alter the correlation of GDP and emissions. Increasing the variance in GDP uncertainty increases the variance in emissions, thereby reducing the relative influence of other uncertain parameters (AEEI and elasticities of substitution). This effect causes emissions to become more closely correlated with GDP. GDP becomes less closely correlated with emissions if its variance is reduced, as the effect of other uncertainties gains relative importance. The reference intensity target used here has a correlation of about 0.87; I alter GDP variance to increase the correlation to 0.93 and to reduce it to 0.65. The results, in terms of net benefits under the intensity target, are shown in Table 5.

correlation	Intensity NB
0.93	24846
0.87	24363
0.85	24241
0.81	24142
0.76	24079
0.65	23952

Table 5: Net benefits under the intensity target when correlation between baseline emissions and GDP growth is varied

The reference intensity target, with an emissions/GDP correlation of 0.87, has been shown to be slightly preferable to the reference safety valve (Table 4). Decreasing this correlation, however, decreases the net benefits of the intensity target to the extent that the safety valve is preferred. This result is shown in Figure 18. The “indifference

point” between the safety valve and intensity target occurs at a correlation of 0.86, or just slightly less than correlation in the reference intensity target. Thus, if GDP and emissions are expected to exhibit high correlation with each other in a given period, the intensity target is preferable to a safety valve. If GDP and emissions do not closely correlate, the safety valve is preferable to the intensity target.




Figure 18: Relative performance of intensity target to safety valve when varying the correlation between emissions (regulated quantity) and GDP (indexed quantity)

3.2.3 Combined Instrument: Intensity Target with Safety Valve

A remaining policy option involves the use of an intensity-based quantity target with a safety valve. Under such a policy, the market price of permits is that given by the intensity target, but unlimited permits are available at the trigger price (\$23/ton).

Perhaps surprisingly, an intensity target combined with a safety valve entails higher expected costs than a pure safety valve. Both the intensity target and safety valve may be thought of as cost-controlling measures when output or costs are unexpectedly high, but the intensity target increases the stringency of the cap when output is unexpectedly low. A cap may require little or no abatement if baseline emissions are low,

but an intensity target applies a target relative to baseline emissions and thus requires that reductions are not negligible.




Figure 19: CDF of carbon Price under the “Intensity Target with Safety Valve” (I + SV) policy

An intensity target thus requires higher costs than a cap in the cases where it is “easiest” to meet the emissions target. Under a pure cap, the market price may be as little as \$2/ton, whereas under an intensity target the market price will be at least \$9/ton (Table 3, Figure 19). Since the safety valve ensures that the market price of permits may not exceed \$23/ton, there is much less uncertainty in marginal cost under the combined instrument than under an intensity target used alone. Possible abatement under the combined instrument is also less uncertain (has a smaller variance) than under the pure safety valve. (Figure 20).




Figure 20: CDF of abatement under the “Intensity Target with Safety Valve” (I + SV) policy

In comparing expected abatement and welfare loss it is unclear whether the combined instrument should be preferred to a pure safety valve and intensity target. The combined instrument has higher expected abatement and welfare loss than the safety valve, and lower expected abatement and welfare loss than the intensity target (Figure 21).

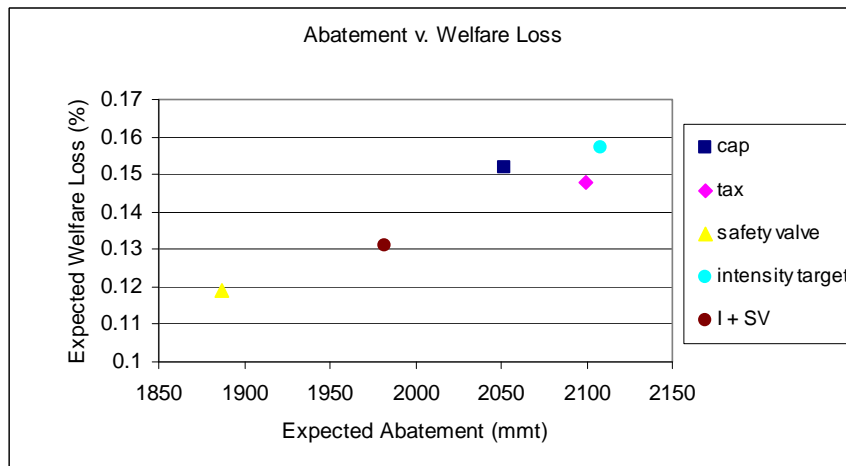


Figure 21: Abatement-welfare trade-off under the “Intensity Target with Safety Valve” (I + SV) Policy

In terms of expected net benefits, however, the combined instrument performs better than both the pure safety valve and intensity target (Figure 22). The reason may be

seen in Figure 19, which shows the CDFs of carbon price under each of the five policies. Given certain marginal damages, the optimal policy sets a tax equal to damages, in this case \$23/ton. Of the four sub-optimal policies, the combined instrument behaves most closely to the tax. A pure intensity target requires more abatement than the combined instrument policy, but in some cases imposes a very high carbon price which may yield small or even negative net benefits. A traditional safety valve allows for very low carbon prices in some cases which also yield very small net benefits. The combined instrument eliminates these extremes and in a sense induces a price ceiling (through the safety valve) and price floor (through the indexed target). These instruments work together to keep the market price of carbon in the near-optimal range, which in turn yields higher expected net benefits than either instrument operating alone. This is reminiscent of work by Roberts and Spence (1976), who showed a cap combined with a price ceiling and price floor to be the optimal hybrid instrument.




Figure 22: Net benefits under each policy when all are equivalent in the “no-uncertainty” case

3.3 UNCERTAIN MARGINAL BENEFITS

The above results rely on the assumptions that marginal benefits are constant with respect to the level of abatement and are equivalent to the level of the reference tax (\$23/ton). I first relax the latter assumption, and test the performance of the policies given constant marginal benefits equal to half (\$11.5/ton) and twice (\$46/ton) the reference level. I then relax the assumption that the marginal benefits are constant, and evaluate the policies under the conditions that the slope of the marginal benefits is equal to, twice, and half the slope of the expected marginal costs at the point of intersection.

3.3.1 Constant Marginal Benefits

If marginal benefits are constant and equal to the level of the tax (as in the results above), the tax by definition is optimal in all cases. If, however, marginal benefits differ from this level the tax does not necessarily perform better than the safety valve or intensity target. If marginal benefits are twice the reference tax (\$46/ton), the optimal level of abatement in the expected case is 2800mmt, substantially more than the 2050mmt required by the tax. If marginal benefits are half the reference tax (\$11.5/ton), the optimal level of abatement is only 1100mmt. The optimal level of abatement in the expected case is determined by the intersection of the marginal benefit curve with the expected marginal abatement cost curve (see Figure 23).




Figure 23: Flat marginal benefits defined as half and twice the reference tax

The performance of each policy relative to a tax is shown below in Figure 24. If marginal benefits are equal to the reference tax (\$23/ton), the ordering of the policy instruments is that shown in the results above (Figure 22). If marginal benefits are twice the reference tax (\$46/ton), the tax still has the highest net benefits of all instruments. The ordering of the next-best policies changes, with the intensity target yielding the second-highest net benefits (98.9% of tax), followed by the intensity target with safety valve (95.8% of tax), the cap (95.6% of tax) and then safety valve (91.9% of tax). The safety valve thus performs worse than all other instruments when marginal benefits are underestimated. The relative ordering of policies is shown most clearly in Figure 25, which highlights the difference between instruments in absolute, rather than relative, terms. In relative terms, however, the difference between policies when marginal benefits are underestimated is very small compared to the difference when they are overestimated (Figure 24).

That is, when marginal benefits are only half the reference (\$11.5/ton), the safety valve and intensity target with safety valve perform significantly better than the \$23/ton tax. The intensity target and cap perform significantly worse. The cap in this scenario

actually has negative net benefits, meaning the costs of abatement outweigh the avoided damages.

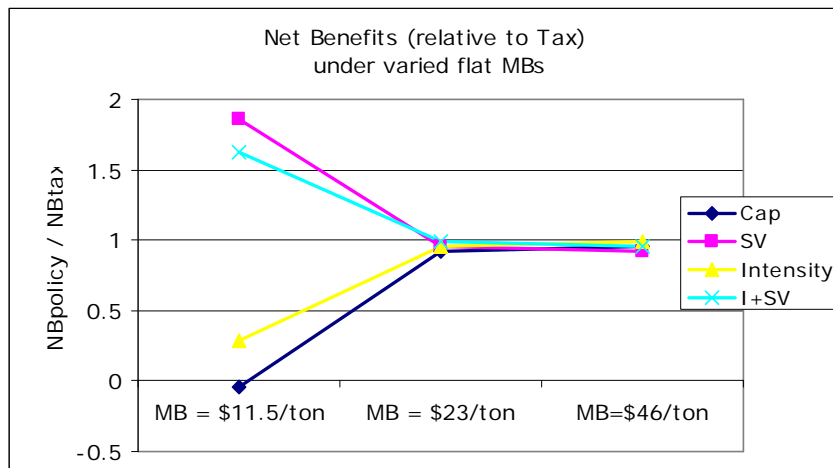


Figure 24: Net Benefits of Cap, SV, Intensity and I+SV as a fraction of Net Benefits under a Tax under constant but varied marginal benefits

From this analysis we can infer that policy instrument choice, in relative terms, matters most when marginal benefits are overestimated. In this case a safety valve, which has lower expected abatement, has net benefits that are 6.5 times higher than the intensity target. If net benefits are underestimated the intensity target performs less than 8% better than the safety valve. These results are summarized in Table 6.

It is also useful to look at net benefits in absolute terms, as more is at stake when net benefits are higher than expected. The absolute difference (in million \$) between net benefits under a tax and all other policy instruments is shown in Figure 25. It is foremost apparent that all cost-containment instruments (SV, Intensity, I+SV) are very effective at dampening uncertainty about cost relative to a cap. An intensity target has the least variation in net benefits, but is also inferior to a tax under all assumptions. A safety valve has the most variation of the three cost-containment instruments, but performs better than the tax when marginal benefits are overestimated.




Figure 25: Absolute Net Benefits (relative to tax) under constant but varied marginal benefits

Safety Valve Design

Lastly, I turn to the question of how the design of the safety valve affects its performance when marginal benefits are constant but varied. As above, I vary the cap one standard deviation above and below the reference (loose cap, tight cap) as well as the trigger price (high tp, low tp). A loose cap or high trigger price yields expected abatement lower than the reference SV, and a tight cap or low trigger price yields expected abatement higher than the reference safety valve. As seen in Figure 26, these design factors matter more, in relative terms, when marginal benefits are overestimated. Also, varying the trigger price, as opposed to the cap, yields higher uncertainty about net benefits. This implies that a safety valve with the trigger price set at the expected carbon price is a more robust cost-containment instrument than a safety valve with the cap set at expected emissions.




Figure 26: Net Benefits relative to reference SV given variations in cap and trigger price under varied but constant marginal benefits

		Constant but Varied Marginal Benefits					
		MB = ref		MB = ref/2		MB = ref*2	
		NB (\$ millions)	Fraction of Tax	NB (\$ millions)	Fraction of Tax	NB (\$ millions)	Fraction of Tax
Policies	Tax	25362	1.0000	1533	1.0000	73019	1.0000
	Cap	23231	0.9160	-58	-0.0379	69808	0.9560
	SV	24262	0.9566	2843	1.8540	67100	0.9189
	Intensity	24363	0.9606	436	0.2844	72218	0.9890
	I+SV	24985	0.9851	2490	1.6242	69975	0.9583

		NB (\$ millions)	Fraction of SV	NB (\$ millions)	Fraction of SV	NB (\$ millions)	Fraction of SV
SV design	SV (ref cap, ref tp)	24262	1.0000	2843	1.0000	67100	1.0000
	ref cap, high tp	23837	0.9825	953	0.3354	69605	1.0373
	ref cap, low tp	21984	0.9061	5438	1.9129	55078	0.8208
	tight cap, ref tp	25181	1.0379	1887	0.6638	71770	1.0696
	loose cap, ref tp	21816	0.8992	4242	1.4923	56963	0.8489

Table 6: Performance of policy instruments given varied assumptions about flat marginal benefits

3.3.2 Non-Constant Marginal Benefits

In the above results marginal benefits are varied but are constant with respect to emissions, meaning that each metric ton of carbon reduction is of equivalent benefit. Since the marginal cost curve is much steeper than the marginal benefit curve, a tax is always optimal to a cap under uncertainty (Weitzman, 1974).

If, in the other extreme, marginal benefits were to have an infinite (as opposed to zero) slope, a cap would always be preferable to a tax under uncertainty. A perfectly vertical marginal benefits curve indicates that any emissions exceeding the point at which

it crosses the marginal cost curve would lead to catastrophic damages. A cap, which guarantees that the appropriate level of emissions, is always optimal in this case.

To show this result empirically it is necessary to represent marginal costs and benefits in terms of emissions rather than abatement. This is because uncertainty in GDP and AEEI lead to uncertainty in baseline emissions (Figure 8), meaning that attaining an emissions cap of 5000mmt results in an uncertain quantity of abatement. If marginal benefits were set to be infinite at the level of expected abatement, a cap would not be optimal in all cases (as it should by definition) since abatement under the cap deviates widely from the expected case (Figure 27).

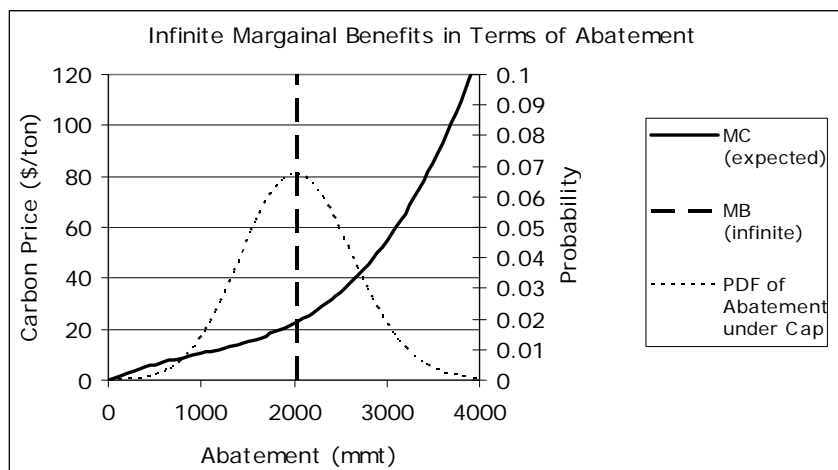


Figure 27: Sub-optimal performance of a cap when BAU emissions are uncertain and MBs are represented in terms of abatement

I resolve this issue by fitting a cubic marginal cost curve to each of the 1000 cases in terms of emissions (Figure 28). I then test the performance of the policy instruments not under infinite marginal damages (where we know the cap will outperform the other instruments), but under three more interesting cases: (1) the slope of the marginal benefits equals the slope of the expected marginal costs, (2) the slope of the marginal benefits is

half the slope of the expected marginal costs, and (3) the slope of the marginal benefits is twice the slope of the expected marginal costs (Figure 29).

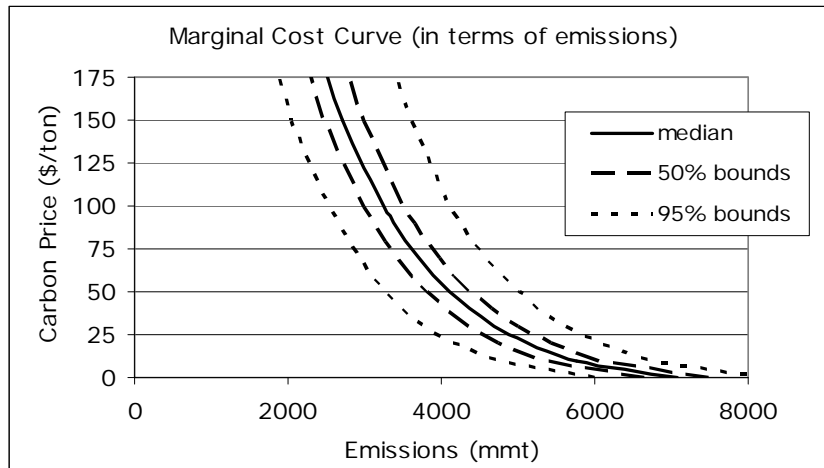


Figure 28: Marginal Costs in 2015 in terms of Emissions

I represent the marginal benefits as linear functions. In the case where the slope of the marginal benefits equals the slope of expected marginal costs, I define the slope of the marginal benefits such that expected benefits under the tax and cap are equivalent. This is consistent with Weitzman's (1974) result that if the slopes of the marginal costs and benefits are equal a cap and tax will yield the same net benefits. The determined slope of the marginal benefits is roughly equivalent to slope of marginal cost curve at the level of emissions under the cap (i.e. the derivative of the marginal cost curve at $x = 5000$).




Figure 29: Non-Constant Marginal Benefits (in terms of emissions)

The performance of a tax, cap, safety valve, intensity target, and intensity target with safety valve under these assumptions are shown in Figure 30. As above, I represent the net benefits of the latter four policies as a fraction of net benefits under the tax. In comparison to varying the level of constant marginal benefits, the difference between the policies given non-constant marginal benefits is small. The largest deviation between any two policies under any of the three assumptions is less than 2.5% (Table 7). Though the difference is small, the safety valve performs better than the intensity target in each of the three scenarios. In addition, the safety valve performs better than all other instruments when the slope of marginal benefits equals the slope of marginal costs. This is consistent with earlier findings (Weitzman, 1978, Roberts and Spence, 1976) that an instrument combining both price and quantity controls can be optimal to either a tax or cap.




Figure 30: Net Benefits of Cap, SV, Intensity and I+SV as a fraction of Net Benefits under a Tax under non-constant marginal benefits

Safety Valve Design

By definition, the optimal safety valve should perform at least as well as a pure tax or pure cap since it combines both of those instruments. Since the safety valve does not perform better than a tax when $MB = MC/2$ or better than a cap when $MB=MC*2$, this implies that the design of this instrument is suboptimal in these cases. I thus test the performance of my standard four safety valve cases (ref cap, high tp; ref cap, low tp; tight cap, ref tp; loose cap, ref tp) given non-constant marginal benefits. The results (Figure 31) show that in the $MB=MC/2$ case, the best safety valve is a tight cap with the reference (\$23/ton) trigger price. This is intuitive since the pure tax performs better than the pure cap in this scenario; combining a tight cap with the reference trigger price ensures that in all but very low-cost cases the safety valve will act as a tax. Similarly, in the $MB=2*MC$ case, the best safety valve is the reference cap with a high trigger price. The cap outperforms the tax when marginal benefits have a steeper slope than marginal costs. The “ref cap, high tp” safety valve acts as a cap in all but very high-cost cases, and thus is the preferred instrument (Table 7).




Figure 31: Net Benefits relative to reference SV given variations in cap and trigger price under non-constant marginal benefits

In summary, if the slope marginal benefits is expected to exceed the slope of marginal costs, the most important aspect of safety valve design is to set the cap at the expected level of emissions. Conversely, if marginal benefits are thought to have a slope less than the marginal costs, the safety valve should be designed such that the trigger price is set at the expected carbon price.

Intensity Target Design

When marginal benefits are assumed to have a slope greater than zero, the intensity target does not perform well relative to the reference safety valve (Figure 30). Much of the reason is that the intensity target is not optimally designed for these conditions. To understand this, consider the case in which GDP growth is the only variable which determines BAU emissions (i.e. baseline emissions and GDP have a correlation of 1). BAU emissions (Q_{bau}) can be represented as a multiple of GDP (x) such that:

$$(1) \quad Q_{bau} = \alpha \cdot x$$

The intensity target requires the ratio of emissions (Q) to GDP to be constant at r , such that:

$$(2) \quad \frac{Q}{x} = r$$

Substituting equation (1) into equation (2), we can see that the relationship between allowed emissions and baseline emissions must be constant when GDP is the only uncertain parameter.

$$(3) \quad \frac{r}{\alpha} = \frac{Q}{Q_{bau}}$$

Also, when GDP is the only uncertain parameter a given percent reduction always has the same cost (Figure 9). In the reference case, r/α is about 70%, meaning that the intensity target always requires a 30% reduction in BAU emissions at a cost of \$23/ton. The intensity target thus acts as a \$23/ton tax when GDP is the only uncertain parameter, meaning that it performs optimally when marginal benefits are flat. When marginal benefits are sloped, however, the \$23/ton tax is no longer optimal. As seen in Figure 32, a tax yields emissions Q^P which are higher than the optimal level. An intensity instrument that is less responsive to deviations in baseline emissions, however, would yield the optimal emissions Q^I .

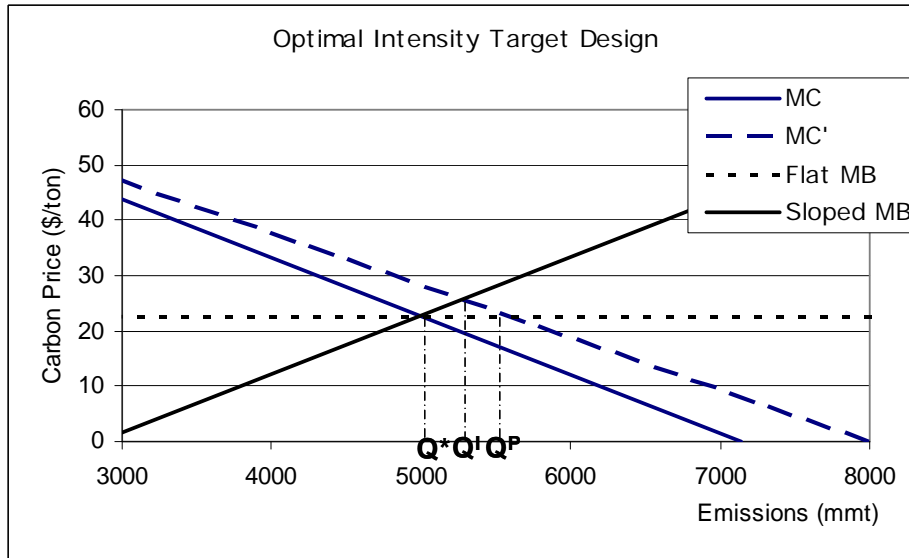


Figure 32: Optimal Design of Intensity Target when marginal benefits are sloped, assuming GDP and Emissions are perfectly correlated

The type of intensity target discussed up until this point may be described as a fully indexed target, meaning that the emissions target scales completely with gdp such that:

$$(4) \quad Q = rx$$

However, a partially indexed intensity target is one that is less widely discussed in policy (Newell and Pizer, 2006), but has been previously been shown to be preferable from a cost perspective (Ellerman and Sue Wing, 2003). Under this target, deviations from expected GDP would define the level of deviation from the expected cap, meaning:

$$(5) \quad Q = q^* + r^*(x - \bar{x})$$

Here, q^* is the level of emissions required by the reference cap (5000mmt in 2015) and $(x - \bar{x})$ is the deviation from expected GDP. The degree to which the intensity emissions target (Q) adjusts in relation to changes in expected GDP is determined by r^* .

A more detailed description of r^* may be found in Appendix II and Newell and Pizer

(2006), but in essence the optimal r^* varies inversely with the slope of marginal benefits. When marginal benefits are flat the optimal $r^* = r$, which causes the intensity target to adjust fully to changes in GDP. When marginal benefits are vertical the optimal $r^* = 0$, since a pure cap with no adjustment always yields optimal emissions. Similarly, when the slope of the marginal benefits equals the slope of the marginal costs, $r^* \approx r/2$, meaning that the cap adjusts to changes in GDP half as much as it would under the fully indexed target. This is shown in Figure 33, with the black link depicting the relationship between the intensity cap and GDP under a fully indexed target, and the orange line depicting the relationship under the optimal (partially-indexed) target assuming the slopes of marginal costs and benefits to be equivalent.




Figure 33: Fully-indexed intensity target vs. optimal partially indexed target when the slopes of marginal costs and benefits are equivalent

The partially indexed (optimal) intensity target has the highest net benefits of all policy instruments, while the fully indexed intensity target has the lowest net benefits (Figure 34). However, designing the intensity target optimally requires prior knowledge of the relative slopes between marginal costs and benefits, which cannot be known in advance with certainty. Currently, marginal benefits are thought to be relatively flat over

a single period (Jacoby and Ellerman, 2004), implying that the fully indexed intensity target should be the design of choice. However, as seen below, this instrument is not very robust if marginal benefits prove to be steeper than the expectation.

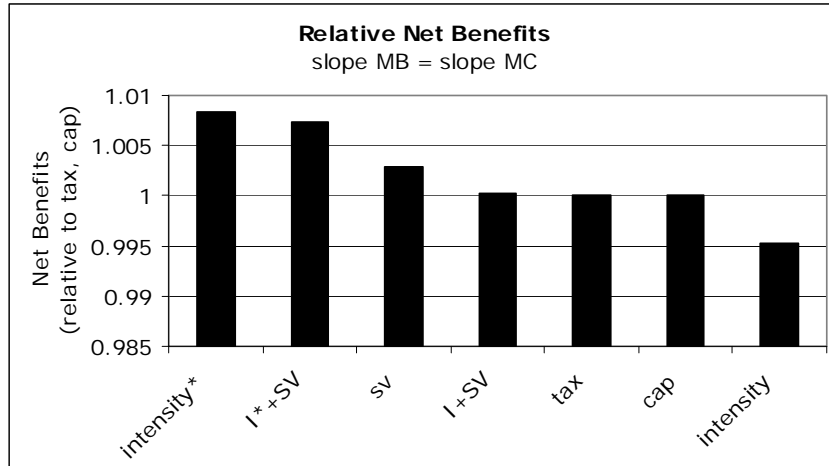


Figure 34: Performance of optimal intensity target (intensity*) and fully indexed target (intensity) relative to other policy instruments

		Linear (non-flat) Marginal Benefits					
		MB = MC		MB = MC/2		MB = MC*2	
		NB (\$ millions)	Fraction of Tax	NB (\$ millions)	Fraction of Tax	NB (\$ millions)	Fraction of Tax
Policies	Tax	65222	1.0000	45264	1.0000	105136	1.0000
	Cap	65222	1.0000	44226	0.9771	107213	1.0198
	SV	65412	1.0029	44845	0.9907	106547	1.0134
	Intensity	64908	0.9952	44636	0.9861	105453	1.0030
	I+SV	65240	1.0003	45089	0.9961	105543	1.0039
	Intensity*	65771	1.0084				
	I*+SV	65696	1.0073				
SV design	SV (ref cap, ref tp)	65412	1.0000	44845	1.0000	106547	1.0000
	ref cap, high tp	65684	1.0042	44761	0.9981	107531	1.0092
	ref cap, low tp	58869	0.9000	40427	0.9015	95754	0.8987
	tight cap, ref tp	65623	1.0032	45375	1.0118	106120	0.9960
	loose cap, ref tp	61358	0.9380	41587	0.9273	100900	0.9470

Table 7: Performance of policy instruments given varied assumptions about non-constant marginal benefits

3.4 CONCLUSIONS

The single-period analysis has shown that if marginal benefits are constant and both instruments are optimally designed, a high level of correlation between the cost

uncertainty and the index uncertainty is required to justify the choice of an intensity target as a regulatory instrument over a safety valve. In actual policy discussions, the design details of each instrument are critical to the choice between instruments. For example, a safety valve with a trigger price much lower than the marginal benefits will be much less efficient, and an intensity target may be superior. In terms of cost effectiveness, a safety valve results in a smaller amount of welfare loss than an intensity target, but also achieves less carbon abatement. An instrument which combines both the safety valve and intensity target falls between the two in terms of abatement and welfare loss, but may be preferable (in terms of net benefits) to either used individually.

Since the safety valve and intensity target differ in expected abatement (and cost) they must be compared in terms of net benefits. Though the marginal benefits of abatement are highly uncertain, the safety valve proves to be more robust than the intensity target under a varied range of assumptions about the benefits of abatement. The relative difference between the safety valve and intensity target is greatest when the benefits of abatement are smaller than expected. In this case, the safety valve performs better than all other instruments by allowing higher expected emissions. Similarly, the safety valve performs better than the intensity target when marginal benefits have a greater slope than expected, as the fully-indexed design of the intensity target is sub-optimal in this case.

The analysis presented up to this point has focused exclusively on a single period of relatively few years. Over a longer time horizon of multiple periods, an additional question addressed in this study is how banking and borrowing of emissions permits performs relative to either a safety valve or an intensity target. As the single-period

analysis showed the safety valve to be the more robust policy instrument, the multi-period analysis has addressed how the safety valve performs relative to borrowing emissions from future periods.

4. MULTI-PERIOD STUDY: METHODS

In this chapter, I describe the multi-period CGE model, which is used in Chapter 5 to analyze the relative performance of the safety valve and borrowing. A multi-period model is necessary to analyze borrowing, as emissions “borrowed” in the first period must be compensated for in later years. Additionally, the multi-period model allows the economic effect of enforcing a cumulative target under the safety valve to be analyzed. The multi-period analysis aims to discover the assumptions and economic conditions under which the safety valve is preferable to borrowing, and vice-versa.

This chapter begins with a description of the multi-period model used in the analysis, and follows with a discussion of the how economic uncertainty is simulated within the model. It then describes the formulation of the policies examined in this analysis. The chapter ends with a “sample” case, where a single shock is evaluated in terms of its effect on emissions, carbon price, and welfare.

4.1 MODEL DESCRIPTION

The model used in this analysis is a forward looking two-region general equilibrium model developed from the recursive-dynamic multi-region MIT Emissions Prediction and Policy Analysis (EPPA) model. The latest version of the recursive-dynamic EPPA model is documented fully in (Paltsev et al., 2005), and the forward-looking model is documented in (Gurgel et al., 2007) and (Babiker et al., 2008). The recursive-dynamic version of EPPA optimizes decisions (across sectors, regions, goods, and inputs) within a time period, but is considered “myopic” in that investments are made

as if input costs and output prices will remain unchanged in the future. In contrast, the forward-looking version of EPPA optimizes decisions both within a time period and across time periods. That is, economic actors in the forward-looking model have perfect foresight – they precisely know the economic conditions in all periods covered by the model, and are able to adjust savings and consumption accordingly. In reality, of course, the situation is somewhere in between (actors have expectations about the future that may or may not be realized). Results from this model are thus idealized in the sense that regulated entities have perfect information about the future.

The countries, sectors and factors in the forward-looking EPPA model are detailed below in Table 8. Instead of the sixteen regions included in recursive dynamic EPPA the model is collapsed to two regions: the US and the Rest of World (ROW). For the purpose of this US-focused study, the ROW is assumed to have no policy for the time horizon of the model. The forward-looking model solves for the time period 2005-2050 in 5-year steps, and the first policy period is assumed to be 2005. Emissions in that period are thus predicted by the model rather than based on from historical data.

Country or Region [†]	Sectors	Factors
United States (USA)	Non-Energy	Capital
ROW Aggregation from standard EPPA:	Agriculture (AGRI)	Labor
Canada (CAN)	Services (SERV)	Crude Oil Resources
Japan (JPN)	Energy-Intensive Products (EINT)	Natural Gas Resources
European Union+ (EUR)	Other Industries Products (OTHR)	Coal Resources
Australia & New Zealand (ANZ)	Energy	Shale Oil Resources
Former Soviet Union (FSU)	Coal (COAL)	Nuclear Resources
Eastern Europe (EET)	Crude Oil (OIL)	Hydro Resources
India (IND)	Refined Oil (ROIL)	Wind/Solar Resources
China (CHN)	Natural Gas (GAS)	Land
Higher Income East Asia (ASI)	Electric: Fossil (ELEC)	
Mexico (MEX)	Electric: Hydro (HYDR)	
Central & South America (LAM)	Electric: Nuclear (NUCL)	
Middle East (MES)	Electric: Solar and Wind (SOLW)	
Africa (AFR)	Electric: Biomass (BIOM)	
Rest of World (ROW)	Electric: NGCC	
	Oil from Shale (SYNO)	
	Synthetic Gas (SYNG)	

[†] Specific detail on regional groupings is provided in Paltsev *et al.* (2005)

Table 8: Regions, Sectors and Factors in the MIT forward-looking EPPA model

The utility function used in forward-looking EPPA is a constant inter-temporal elasticity of substitution function. The representative agent in each region maximizes this utility function, subject to a budget constraint, technology and the evolution of capital stock in the economy. Each region is endowed with an initial stock of capital, labor and energy resources. As described below, the supply of oil resources in 2005 is increased in the model to simulate a cost shock in that year.

The above optimization problem is converted into a market equilibrium formulation using the mixed complementarity problem (MCP) algorithm (Mathiesen, 1985, Rutherford, 1995), and solved numerically using the General Algebraic Modeling System (GAMS) software (Brooke *et al.*, 1998).

4.2 PARAMETERIZATION OF UNCERTAINTY

The cost shock considered in this analysis is “temporary” in that it only affects the first policy period. The marginal cost curves in all years following (2010-2050) are assumed to revert to the expected marginal cost curves for each year. Such a shock could be caused by a temporary change in the supply or demand of energy, or alternatively as a result of market fluctuations in adjusting to a new policy. This type of shock is distinct from a “permanent” shock, in which the cost in every policy period (2005-2050) would be correlated across time. A GDP shock is generally an example of a permanent shock, as GDP in each year is affected by the previous year’s GDP. Assume for example that GDP is expected to grow at 3% each year between 2005 and 2050, but that the GDP growth rate in 2005 is instead 4%. Even if GDP growth in every subsequent year reverts to expectation (3%), total GDP in each of these years will be higher than expected since all growth is compounded on the 2005 level.

In contrast, the temporary shock used here only raises the 2005 carbon price above expectation. This is similar to the single-period analysis as the cost shock occurs in only one period, but the multi-period model allows emissions to be traded between periods in response to this shock. In the model, the shock is simulated by increasing the fixed factor supply of oil. Increasing the supply of this carbon-intensive energy source lowers oil prices, thus increasing BAU emissions in 2005. Since BAU emissions are higher than expected, the carbon price required to meet the carbon emissions constraint is also higher than expected.

The distribution of shocks to the fixed factor oil resource is shown below in Figure 35. The distribution denotes the factor by which oil resources are increased, and

is normal with a mean of 1. A shock value of 3 indicates that oil resources in 2005 are considered to be 3 times the reference value. The expected shock of 1 thus represents no change to reference oil resources.

As shown in Figure 35, only cost shocks greater than one are considered in this analysis. Cost shocks less than one would imply that the cost of meeting the emissions target is cheaper than expected. The cost-containment instruments considered here are intended to restrain costs from becoming too high; thus only higher-than-expected cost shocks are relevant.




Figure 35: Distribution (normal) of shocks to fixed factor oil resources; only positive cost shocks (increases in oil supply) are considered

Figure 36 shows the effect of an increase in oil resources on 2005 carbon price. The emissions target used here is extrapolated from the 50% reduction path in Paltsev et al. (2007), and is described as the “non-adjusted path” in later sections of this analysis. The factor shock to carbon price denotes the factor increase in carbon price relative to the reference (no-shock) value under the target. Thus, it is about twice as expensive to meet the emissions target when oil resources are increased by a factor of 6, as compared to when oil resources are unchanged (factor shock of 1).




Figure 36: Effect of increase in oil resources on carbon price

4.3 POLICY DESIGN

As in the single period model (chapters 2 and 3), the shock to oil resources is introduced into the CGE model through a Monte Carlo simulation. I draw 2000 random samples from the normal distribution of shock to the fixed factor oil supply, and run only those greater than the reference (roughly half) through the model. For each shock value I impose five types of policy constraints: (1) reference (no policy), (2) emissions cap, (3) safety valve, (4) compensated safety valve, and (5) borrowing, which are detailed in this section.

The stringency of the emissions cap is again defined as the expected CO₂ abatement under the 50% reduction by 2050 emissions path in Paltsev et al. (2007). Emissions reductions are assumed to begin in 2005, with the amount required in that year determined by the linear path between 1997 historical emissions and the required target in 2030 (20% below 1997 emissions). See “Non-Adjusted Cap” in Figure 37.

To obtain the emissions quotas in each year, I allow intertemporal trading (banking and borrowing) in the CGE model. This finds the optimal path through time that

achieves the same cumulative emissions from 2005-2050, as seen by the “Adjusted Cap” in Figure 37. As is clear from the figure, the optimal temporal path results in net banking relative to the initial allocation. This means that abatement in earlier years (before 2030) exceeds the required level and the saved permits are then used in later years (after 2030).




Figure 37: Emissions Cap with (adjusted cap) and without (non-adjusted cap) banking

4.3.1 Effect of Banking

Both the Lieberman-Warner and Bingaman-Specter Bills allow net banking (United States Senate Committee on Energy and Natural Resources, 2007, Environmental Protection Agency, 2008) and, as depicted in Figure 37, this is expected to lead to more stringent action in early periods. If an unexpected cost shock occurs in the first year of the policy period, firms will simply bank less permits than in the no-shock case. This is shown in Figure 38, where a shock simulated by four times the fixed factor resources still results in 2005 emissions that are well below the required cap. Under the shock, 2005 emissions are slightly elevated above the no-shock banking path, and emissions in all subsequent years are slightly reduced.




Figure 38: Effect of banking given a temporary shock in 2005

In this sense, banking alone acts as a cost-containment measure against first-period shocks. This result can also be seen in the expected carbon prices shown in Figure 39. Banking smooths the cost of abatement such that the resulting discounted carbon price (in 2005 dollars) is constant through time. Under no-shock expectations, the carbon price each year is about \$61/ton CO₂; a cost shock in 2005 raises the yearly price slightly to about \$62/ton CO₂. If banking is not allowed, the entire cost of the shock must be endured in 2005, nearly doubling the expected carbon price in that year.




Figure 39: Effect of banking on discounted carbon price (in 2005 dollars)

4.3.2 Emissions Cap (No Cost Containment)

Because net banking behavior is observed in the presence of a substantial cost shock in 2005, neither borrowing nor a safety valve are likely to be activated. Since the goal of this section is to compare the safety valve (specified in the Bingaman-Specter bill) to borrowing (specified in the Lieberman-Warner bill) I assume that the initial permit allocation is given by the “Adjusted Cap” in Figure 37. This is equivalent to assuming that the pure cap is designed to be optimal in the expected (no-shock) case. The pure cap is defined as “no cost-containment” since banking and borrowing are not allowed in response to the shock.

4.3.3 Safety Valve

As with the single-period analysis, the safety valve is designed to perform identically to the pure cap in the no-shock case. The safety valve trigger price is thus set to the expected 2005 price under the cap, about \$61/ton CO₂. The safety valve may only be triggered in the first period, and no banking or borrowing is allowed. That is, the safety valve is only intended to protect against cost shocks occurring in the first period of the policy. Permits cannot be purchased at the safety valve price, banked, and used in future years.

4.3.4 Compensated Safety Valve

One of the prime oppositions to the safety valve is that it allows emissions to exceed the cumulative level specified by the cap, potentially resulting in irreversible climate damages if that cumulative cap is important (Jacoby and Ellerman, 2004). The

compensated safety valve addresses this issue by requiring abatement avoided in the first period to be compensated for in later years. As under the traditional safety valve, unlimited permits may be purchased at the trigger price in 2005. However, emissions allowances in later years are evenly reduced to compensate for the extra emissions in 2005.

4.3.5 Borrowing

This policy allows permits to be optimally traded across time in response to the shock, but requires that the cumulative emissions cap be met. Since only positive cost shocks are considered, net borrowing always results. Permits are borrowed from future years to be used in 2005 such that the burden of the cost shock is shared throughout the length of the policy (see Figure 39).

4.4 POLICY BEHAVIOR UNDER A SINGLE PRICE SHOCK

Relative effects of these policies may be illustrated by comparing their behavior under a single temporary shock of four times the reference oil resources. Figure 40 shows emissions from 2005-2050 under each of the four policies in response to the shock.




Figure 40: Emissions from 2005-2050 in response to a single temporary shock

The differences in 2005 and 2010 emissions are highlighted in the bar graphs in Figure 41. “No Cost Containment” represents the pure cap, which requires the lowest emissions in 2005 since the cap in that year is not flexible. Borrowing allows for more emissions than under the cap but less than under the safety valve. Both the Safety Valve and Compensated Safety Valve allow the most emissions in 2005, as additional permits are purchased at the trigger price. In 2010, and all subsequent periods, the situation changes. Emissions are most constrained (lowest) under the compensated safety valve, because the additional emissions in 2005 must be made up for. Borrowing has slightly higher emissions than the compensated safety valve in 2010. This is because emissions under borrowing were lower in 2005 than under the compensated safety valve, and therefore fewer emissions need to be reduced in future periods. The safety valve and no cost-containment policies have the highest emissions in 2010; the safety valve does require the extra 2005 emissions to be compensated for and thus allows emissions up to the level of the cap.




Figure 41: Emissions in 2005 and 2010 under each policy in response to a single temporary shock

The resulting carbon prices can be seen in Figures 42 and 43. The differences in carbon price under each policy can be seen most clearly in Figure 43, which shows each year's carbon price in discounted 2005 dollars. As expected, the no cost-containment policy (cap) has the highest 2005 carbon price, as this policy requires the 2005 emissions target to be met at any cost. The borrowing policy smooths the discounted cost over time such that it is roughly equal in each year. The safety valve results in the lowest carbon price in 2005 (which is pre-determined by the specified trigger price), and has a carbon price equal to the no cost-containment policy in each period following. The price under the compensated safety valve in 2005 is also pre-determined by the trigger price, but the price in each subsequent year is slightly higher than the price under borrowing.




Figure 42: Carbon Price under each policy in response to a single temporary shock




Figure 43: Discounted carbon price (in 2005 dollars) in response to a single temporary shock

The performance of the policies in terms of expected overall welfare (2005-2050) is shown in Figure 44, in terms of the reduction from no-policy welfare. Welfare is measured as loss in consumption, and here refers only to the cost of the policy (not to benefits associated with carbon abatement). The safety valve has the highest welfare (lowest reduction from reference), but also requires the least abatement as described

above. The welfare loss under the compensated safety valve and borrowing is approximately equal, but significantly less than under the no cost-containment policy. Since the no cost-containment policy enforces the same cumulative cap as the compensated safety valve and borrowing policies, but at greater cost, it can be considered inferior to both.

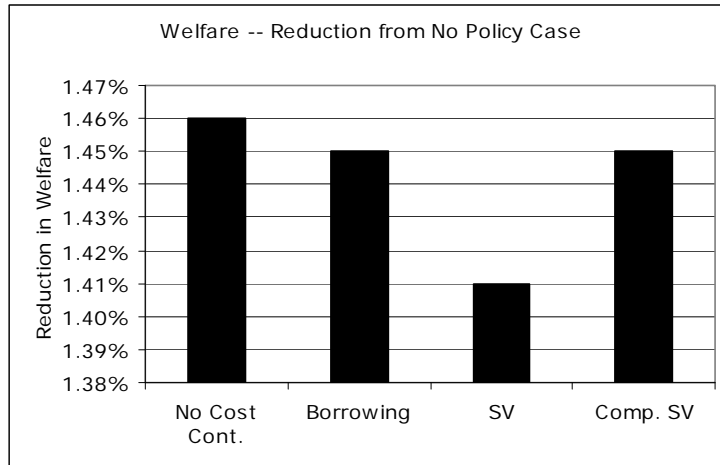


Figure 44: Welfare reduction from reference in response to a single temporary shock

As it is unclear from this one-shock example how borrowing compares to either the compensated safety valve or traditional safety valve, I now turn to the results of the Monte Carlo simulation. I first compare borrowing to the compensated safety valve in terms of overall welfare, as both attain the same level of cumulative emissions. I then compare borrowing to the traditional safety valve in terms of net benefits, given a variety of assumptions about the marginal benefits of abatement.

5. MULTI-PERIOD STUDY: RESULTS AND ANALYSIS

Having examined the performance of the safety valve and borrowing under a single temporary shock, this chapter now turns to the results of the Monte Carlo simulation. This chapter first describes the range of carbon price, emissions, abatement, cumulative emissions and welfare possible under the given distribution of cost shocks. It then directly compares the compensated safety valve and borrowing on the basis of overall welfare, specifically examining the effect of varying the trigger price used in the compensated safety valve. Finally, this chapter directly compares the traditional safety valve to borrowing under a variety of assumptions about the net benefits of abatement.

5.1 RESULTS OF MONTE-CARLO SIMULATION

Introducing a series of random shocks to oil resources through a Monte Carlo simulation yields the following results, as summarized below (Table 9) and presented in detail in this section.

	Carbon Price (2005) (\$/ton)			Emissions (2005) (mmt Carbon)			Abatement (2005) (mmt Carbon)			Emissions (2005-2050) (mmt)			Welfare (2005-2050) (fraction no policy wel.)		
	mean	min	max	mean	min	max	mean	min	max	mean	min	max	mean	min	max
No Cost-Cont	70.9	60.6	79.4	4060	4060	4060	2147	1457	2374	47717	47717	47717	0.98521	0.98496	0.98555
Borrowing	61.2	60.6	62.0	4254	4060	4428	1953	1457	2042	47717	47717	47717	0.98526	0.98496	0.98567
SV	60.8	60.6	60.6	4264	4060	4454	1943	1456	2034	47921	47717	48111	0.98548	0.98506	0.98608
Comp. SV	60.8	60.6	60.6	4264	4060	4454	1943	1456	2034	47717	47717	47717	0.98524	0.98496	0.98563

Table 9: Mean, Min and Max values for carbon price, emissions, abatement, cumulative emissions and overall welfare for each policy for a range of cost shocks in 2005

5.1.1 Carbon Price

This analysis only considers a range of positive cost shocks, meaning the carbon price under the cap must always be greater than the no-shock carbon price. The safety

valve trigger price is set at the level of the no-shock price (\$60.6/ton). Because the price under the cap always exceeds this level, the safety valve is always triggered. The 2005 carbon price under the safety valve and compensated safety valve is thus a constant \$60.6/ton (Figure 45).

When borrowing is allowed, the carbon price deviates slightly, ranging from \$60.6/ton to \$62/ton (Figure 45). Given a very large cost shock in 2005, it is optimal to raise the carbon price (and lower emissions) slightly relative to the no-shock price so that fewer emissions have to be made up in the future. With no cost-containment, the carbon price experiences a 31% increase, ranging from \$60.6/ton to \$79.4/ton (Figure 45), as it becomes increasingly more expensive to meet the required emissions target.




Figure 45: CDF of 2005 carbon price under uncertainty for the no cost-containment, borrowing, safety valve, and compensated safety valve policies

5.1.2 Emissions

The cost shock simulated in this model involves an increase in oil resources such that emissions in the BAU case increase above the projected 2005 emissions level (Figure 46). This figure also shows that the marginal increase in emissions declines in response to each marginal increase in the oil supply.




Figure 46: CDF of 2005 BAU carbon emissions under uncertainty

In terms of emissions, the no cost-containment policy requires that the 2005 cap be met and thus emissions are held constant at 4060mmt (Figure 47). Under the safety valve and compensated safety valve the carbon price is fixed, and emissions are a function of the marginal cost curve. Larger cost shocks yield higher emissions, which increase approximately linearly. Under borrowing 2005 emissions are slightly lower than under the safety valve, a result of the slightly higher carbon price (Figure 47).




Figure 47: CDF of 2005 emissions under uncertainty for the no cost-containment, borrowing, safety valve, and compensated safety valve policies

The marginal cost curve, in terms of emissions, is shown in Figure 48. As expected, the cost of achieving any given level of emissions is greater in the shock than the no shock case. The difference in price between the no shock and shock case declines, however, as the emission target becomes more stringent.




Figure 48: Marginal cost curve in terms of emissions in the no shock and median shock cases

5.1.3 Abatement

Abatement in each year is defined as BAU emissions less policy emissions. Policy emissions under the no cost-containment policy are constant, and thus abatement under this policy follows the same trend as BAU emissions in that it increases in response to larger cost shocks. Abatement under the safety valve and borrowing policies also increases as the size of the cost shock increases (Figure 49). Figure 50 shows the marginal cost curve under the no shock and median shock cases in terms of abatement. This figure indicates that a given carbon tax (ex. \$50/ton) results more abatement under the shock as compared to under no shock. As shown in Figure 48, emissions under a given carbon tax also increase in response to the oil shock. This effect is explained by the increase in BAU emissions under the oil shock. The level of BAU emissions must,

by definition, be equal to the sum of emissions and abatement, and thus is shared between both variables.




Figure 49: CDF of 2005 carbon abatement under uncertainty for the no cost-containment, borrowing, safety valve, and compensated safety valve policies




Figure 50: Marginal cost curve in terms of abatement in the no shock and median shock cases

5.1.4 Cumulative Emissions

In terms of cumulative emissions (2005-2050), only the safety valve differs from the other three policies. The no cost-containment, borrowing, and compensated safety valve policies all enforce a cumulative emissions target, and thus have the same total emissions although emissions levels in individual years may differ. The safety valve, in

contrast, is not required to make up any 2005 emissions which exceed the original cap, and thus has cumulative emissions greater than or equal to the other policies (Figure 51). Cumulative emissions under the safety valve are determined by the level of the cost shock in 2005; a larger shock leads to higher 2005 emissions and thus higher cumulative emissions.




Figure 51: CDF of cumulative carbon emissions (from 2005-2050) under uncertainty for the no cost-containment, borrowing, safety valve, and compensated safety valve policies

5.1.5 Welfare

Total welfare, representing the loss in consumption over the policy period 2005-2050, is shown for all policies in Figure 52. Welfare is represented as a fraction of no-policy welfare, and again refers only to the economic cost of adhering to the policy. The safety valve has the highest welfare of all a policies (CDF to the right of all others) but, as seen in the previous section, this policy also allows the highest cumulative emissions. Comparing the trade-off between welfare and emissions requires the benefits of abatement to be quantified, which is explored in section 5.3.

The no cost-containment, borrowing and compensated safety valve policies have the same cumulative emissions, and therefore the same benefits, so a comparison of cost

is more meaningful. As shown in Figure 52 the no cost-containment policy has the lowest welfare (CDF to the left of all others), followed by the compensated safety valve. The borrowing policy stochastically dominates both the no cost-containment and compensated safety valve policies, meaning that it is the optimal policy if the cumulative emissions cap must be met. This result is expected since the model used here has perfect foresight, and is able to choose the 2005 carbon price under each cost shock that maximizes total welfare.




Figure 52: CDF of overall welfare (from 2005-2050) under uncertainty for the no cost-containment, borrowing, safety valve, and compensated safety valve policies

5.2 BORROWING VS. COMPENSATED SAFETY VALVE

The compensated safety valve is inherently less flexible than borrowing, because it fixes the 2005 carbon price before the cost shock is known. Borrowing allows the 2005 carbon price to deviate along with the shock, such that different shock results in a different carbon price. As a result, the compensated safety valve can never perform better than this idealized version of borrowing. However, it is still an important instrument to consider in that it addresses the main opposition to the traditional safety valve: the lack of

a cumulative emissions target. Information about the benefits of carbon abatement will change over time, and it may be realized in later years that the excess emissions allowed under a safety valve must be compensated for. Additionally, borrowing may not function as smoothly as the economic model predicts. Regulated entities do not have perfect foresight, borrowing may have penalties associated with it, and the length of the policy period may be uncertain. The performance of the compensated safety valve is less dependent on these factors, because emissions in the first policy period is determined by the marginal cost curve rather than the agent's expectations about the future. For these reasons the political process may favor the compensated safety valve over borrowing. The remainder of this section is thus devoted to the effect of the chosen trigger price on welfare under the compensated safety valve.

Trigger price values above and below the reference (no-shock) value of \$61/ton CO₂ are examined. The trigger price is raised to \$70/ton (15% above reference) and \$79/ton (30% above reference), and lowered to \$42/ton (30% below reference) and \$18/ton (70% below reference). The resulting carbon price in 2005 under the compensated safety valve policy is shown in Figure 53. In the reference (\$61/ton), \$42/ton and \$18/ton cases the trigger price is always binding, resulting in a constant marginal cost (or carbon price) across all cost shocks. In the \$70/ton case the trigger price is binding for approximately half of the shock values. The remaining shock values are not large enough to raise the market price of carbon above the trigger price value, and so the carbon price is equal to the carbon price under the cap. In the \$79/ton case the trigger price is only binding in 5% of cases, making it only slightly different from the no cost-containment policy.




Figure 53: Resulting 2005 carbon price given trigger prices above and below reference

5.2.1 Optimal Trigger Price

Expected welfare under the compensated safety valve given, variable trigger prices is shown in Figure 54. Welfare is depicted in terms of the reduction from no-policy welfare, the policies represented are sorted from lowest to highest reduction. Borrowing entails the least reduction in welfare (as expected) and is closely followed by the reference trigger price case, indicating that the reference trigger price has the highest welfare of all prices tested. Raising the trigger price to \$70/ton, so that it is binding in only 50% of cases, decreases welfare slightly further. Raising the trigger price to \$79/ton, so that it is binding in only 5% of cases, causes expected welfare to be roughly equivalent to the no cost-containment case. Lowering the trigger price to \$42/ton leads to even lower expected welfare than having no cost-containment at all. Lowering the trigger price to only \$18/ton leads to an even larger reduction in welfare.




Figure 54: Reduction in expected welfare (2005-2050) from no-policy case, given compensated safety valve trigger price values above and below reference

The clear insight from the above graph is that setting the trigger price below the reference level is far worse in terms of welfare than setting the trigger price above the reference level. A low trigger price relaxes the 2005 emissions constraint, and the cost of making those emissions up in later years is greater than the cost saved in 2005. As shown in Figure 55, this result is true under all possible cost shock values, as the \$42/ton and \$18/ton cases are stochastically dominated by all other trigger price values.

The \$18/ton trigger price (-70%) is roughly equivalent to the expected carbon price under the non-adjusted emissions path (see Figure 37), i.e. the original emissions path without accounting for banking. A trigger price set at this level has the lowest expected welfare. Thus, a policy with banking and safety valve (such as Bingaman-Specter) that enforces a cumulative emissions target should take banking into consideration when choosing the initial trigger price. The optimal trigger price is close to the expected carbon price under the inter-temporal least-cost emissions path.




Figure 55: Welfare index (2005-2050) under uncertainty given compensated safety valve trigger price values above and below reference

5.3 BORROWING VS. SAFETY VALVE

Given that the safety valve and borrowing policies result in different cumulative emissions over the horizon of the policy, a comparison in terms of welfare is not very meaningful. As in the single-period analysis, I compare these policies based on varied assumptions about the net benefits of abatement. One complication of calculating net benefits over multiple periods, as opposed to a single period, is that the chosen discount rate can affect the optimal instrument choice. Because the appropriate discount rate is a controversial aspect of climate policy (Nordhaus, 2007), the results shown below are given for a range of possible discount rates.

5.3.1 Constant Net Benefits

Figure 56 shows the ratio of safety valve net benefits to borrowing net benefits when the marginal benefits of abatement are assumed to be constant in each period and equal to the (adjusted) no-shock carbon price. In this figure, and in others following, the

safety valve is preferred in the region above the orange line, whereas borrowing is preferred in the area below the orange line. As can be seen in the graph, the safety valve has a slight advantage over borrowing for all discount rates considered (1-15%), though its advantage narrows as the discount rate rises.




Figure 56: Ratio of expected net benefits under safety valve to net benefits under borrowing, when marginal benefits are flat and equal to the reference carbon price

The reason the safety valve has higher net benefits than borrowing for all discount rates has to do with the safety valve design. The trigger price in all years is set to the no-shock carbon price, meaning that the trigger price is equal to the marginal benefit of abatement. Since only positive cost shocks are considered the safety valve is always triggered, and thus the safety valve acts as a tax equal to the level of the marginal benefits. As demonstrated in the single-period model and as shown in Figure 56, a tax set to the level of the marginal benefits of abatement always results in the optimal quantity of abatement. Under a positive cost shock, the borrowing policy always abates a bit more than optimal in 2005 and in every year thereafter (resulting in greater cumulative abatement), which results in some loss of efficiency.




Figure 56: Safety valve design is optimal when marginal benefits of abatement are flat and equal to the trigger price

I next test the relative performance of the safety valve to borrowing when net benefits are constant and equal to twice (\$121/ton) and half (\$30/ton) the no-shock carbon price in each period. These assumptions are shown for 2005 in relation to the expected marginal cost curve in Figure 57.

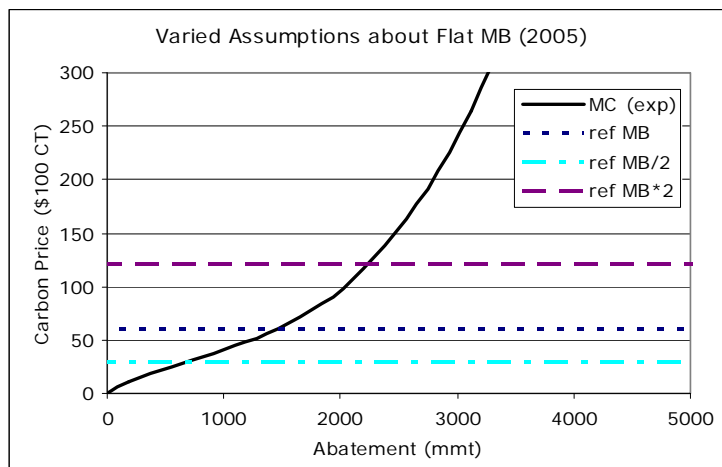


Figure 57: Flat marginal benefits (2005) defined as half and twice the no-shock carbon price

The performance of the safety valve relative to borrowing under the varied assumptions about net benefits is shown in Figure 58, assuming a 5% discount rate. The safety valve performs slightly better than borrowing in the reference case (see Figure 56).

However, this difference is negligible when compared to the safety valve's advantage when marginal benefits are only half the no-shock carbon price in all years (ref MB/2). In this case, the safety valve performs about 3.5% better than borrowing. If reference benefits are twice the no-shock carbon price (ref MB*2), then borrowing outperforms the safety valve, but only by about 0.5%.




Figure 58: Expected net benefits of SV/Borrowing given a 5% discount rate

The relationship shown in Figure 58 holds for all discount rates (see Figure 59), though the advantage of the optimal policy increases as the discount rate rises. The explanation for this is that a higher discount rate places greater weight on the first policy period (2005) in the calculation of overall net benefits. The difference in expected abatement between borrowing and safety valve is greatest in the first period (due to the cost shock), causing the differential in expected net benefits to be greatest in that period as well. Thus, the policy which provides the highest net benefits in 2005 (the safety valve in the MB/2 case and borrowing in the MB*2 case) becomes increasingly favored as a higher discount rate is used.




Figure 59: Expected net benefits of SV/Borrowing given constant MBs over a range of discount rates

One can also conclude from Figure 59 that the safety valve is the more robust than borrowing under constant but uncertain marginal benefits (assuming uncertainty is symmetric). When marginal benefits are twice the expected amount, net benefits under borrowing are in all cases less than 1% better than net benefits under the safety valve. However, when net benefits are half the expected amount the safety valve performs up to 17% better than borrowing (assuming a 10% discount rate).

5.3.2 Non-Constant Net Benefits

This section addresses the question of how the safety valve performs relative to borrowing when marginal benefits are sloped. As in the single-period analysis, it is necessary to represent marginal costs and benefits in terms of emissions rather than abatement. This is again because BAU emissions shift along with the cost shock, and thus there is no level of abatement for which a cap would always be optimal if marginal benefits were infinite (perfectly vertical at the specified level). (see Figure 27)

For this analysis I again assume that marginal benefits are (1) equal to the slope of marginal costs (at the point of intersection), (2) twice the slope of marginal costs, and (3) half the slope of marginal costs. These are shown graphically in terms of expected 2005 emissions in Figure 60. For years 2010-2050 the same assumptions are made about marginal benefits in relation to the (certain) marginal costs at each 5-year interval.




Figure 60: Sloped marginal benefits (2005) defined as half and twice the slope of the marginal cost curve at the point of intersection

The results over a 1-15% range of discount rates are shown in Figure 61. Net benefits under the reference constant marginal benefits (from Figure 56) are also shown. The trend seen in the graph is that the greater the slope of the marginal benefits, and the higher the appropriate discount rate, the better borrowing performs in relation to the safety valve. When the slope of marginal benefits equals the slope of marginal costs, borrowing is preferred over the safety valve when a discount rate greater than 13% is used. When the slope of marginal benefits is twice the slope of marginal costs, borrowing is preferred over the safety valve when a discount rate greater than 5% is used. If the

slope of marginal benefits is only half the slope of marginal costs, a discount rate greater than 15% is required for borrowing to be preferred over the safety valve.




Figure 61: Expected net benefits of SV/Borrowing given sloped MBs over a range of discount rates

The economic theory behind this result is shown in Figure 62. Since the only cost shocks considered here are positive, the actual marginal cost curve (in terms of emissions) is an upward vertical shift from the expected marginal cost curve (see Figure 48). Given an upward shift, the optimal carbon price in 2005 (P^*) is always higher than the safety valve trigger price (P_{SV}). The carbon price under borrowing (P_B) is always slightly above the safety valve level, and thus there are higher net benefits (less deadweight loss) under this policy than under the safety valve (see Figure 62, left). In the following years (2010-2050), the situation is reversed. The safety valve trigger price always results in the optimal level of emissions (since marginal costs are as expected), whereas borrowing results than a higher than optimal carbon price.

Borrowing thus performs better than the safety valve in 2005, but worse in years 2010-2050. If a high discount rate is used, net benefits in 2005 have more weight relative to net benefits in later years, and thus borrowing does better under those assumptions.




Figure 62: Efficient performance of borrowing (relative to safety valve) in 2005, and optimal performance of safety valve in subsequent years (2010-2050)

5.4 CONCLUSIONS

An important observation made in the multi-period analysis is that banking, though not explicitly intended as a cost-containment instrument, reduces the need for either a safety valve or borrowing in the first period. Both the Bingaman-Specter and Lieberman-Warner bills set relatively relaxed targets in the initial policy periods (until 2025), but more stringent action in later periods (2030-2050). The result under banking is that more emission reductions than required are undertaken in the early periods to reduce the burden in later periods. This provides an inherent cushion for cost shocks in the first policy period, as the intended emissions target in the first period is flexible (agents may choose to bank less).

If, however, the specified emissions path were to represent the inter-temporally optimal policy in the absence of shocks (the “adjusted path” in this analysis), either borrowing or a safety valve would be effective as a cost-containment instrument. Even for a non-optimal emissions path, these instruments could also be useful in later policy periods (when a more stringent emissions target is required) or could also be important if the first-period cost shock were extremely large.

Under the “adjusted” (post-banking) path, this analysis finds that a safety valve, compensated safety valve or borrowing are all effective in reducing costs below the level required by a pure cap with no cost-containment. The analysis first compares the compensated safety valve to borrowing, which may be compared on the basis of abatement costs since both policies enforce an equivalent cumulative target. Borrowing is found to be stochastically dominant to the compensated safety valve in terms of welfare (it dominates regardless of the magnitude of the cost shock), but its performance is reliant on the assumption made by the model that all agents have perfect information about the future. In reality, agents may speculate about the future, but cannot know the economic conditions with certainty. Thus, borrowing is unlikely to perform as well as the model predicts.

The performance of the compensated safety valve is not dependent on actors’ knowledge about the future; the policy sets a maximum price before the cost shock is known, and the resulting emissions are determined by the marginal cost curve. If the compensated safety valve’s trigger price is set to the no-shock carbon price under the adjusted path, it performs nearly as well as the idealized version of banking. If the trigger price is set too high, abatement costs will be no higher than under a pure cap (no cost-

containment). If the trigger price is set too low, meeting the cumulative cap may become extremely costly – more so than having no cost-containment at all.

The traditional safety valve must be compared to borrowing on the basis of net benefits, because the policies have different cumulative caps over the policy horizon. The safety valve, unlike borrowing, does not require excess emissions allowed in the first period to be made up for in later years. The multi-period analysis showed the safety valve to be preferable to borrowing (for any discount rate) if the marginal benefits are constant and equal to the expected (no-shock) carbon price. If marginal benefits are constant and twice the expected carbon price, borrowing is slightly preferred; if they are half the expected carbon price, the safety valve is preferred by a substantial margin. If marginal benefits are sloped in all periods, borrowing is only preferred to the safety valve if the slope of the marginal benefits are steep (slope greater than the marginal costs) and the appropriate discount rate is high (at least 13%).

6. CONCLUSIONS

A cap-and-trade program is currently the most politically palatable carbon reduction policy in the US, but because of uncertainty has the potential for extremely high policy costs. If costs are much higher than expected under a proposed cap-and-trade program, the economic impacts of the policy may be severe and any future carbon reduction policy may lose credibility. Although a carbon tax has been shown to be far less vulnerable to cost uncertainty, its unpopularity has led the political process to consider alternative cost-containment measures to augment a cap-and-trade approach. Recently, three regulatory instruments that address cost uncertainty under a cap-and-trade program have been proposed, including a safety valve, an intensity target and borrowing. While previous studies have demonstrated the advantage of each of these instruments in relation to a pure cap-and-trade program, they have not examined how these instruments perform in a relative sense. This thesis has used both a single-period and multi-period economic model to compare these cost-containment instruments directly to each other.

First, a single-period economic model of 2015 was used to test the performance of a pure cap, tax, safety valve and intensity target. The types of uncertainty considered included the GDP growth rate, AEEI, and elasticities of substitution. As expected, the pure cap was shown to have the greatest potential for welfare loss, the tax the least, and the safety valve and intensity target fell in between. Factors affecting the relative performance of the safety valve to the intensity target were found to include, 1) safety valve design (especially trigger price), 2) correlation between emissions and GDP, and 3) assumptions about net benefits.

With regard to safety valve design, setting the trigger price at the expected optimal price was found to be more important than setting the cap at the expected optimal quantity. The optimal safety valve is one which will act as the optimal tax in most cases, and is achieved by pairing a very stringent cap with the expected optimal trigger price. With regard to correlation, a very high (.86) correlation was required between emissions and GDP for the intensity target to be preferable to the reference safety valve. This correlation is higher than that predicted by a linear analytical model, as the cubic slope of the marginal cost curve tends to favor the safety valve. Finally, with regard to assumptions about net benefits, the safety valve out-performed the intensity target when constant marginal benefits were lower than expected, or when marginal benefits were non-constant. If constant marginal benefits are greater than expected, the safety valve performs worse than the intensity target, but the difference between the two is relatively small in this circumstance. Finally, the combined intensity target with safety valve instrument was found to be preferable to either individual instrument under most assumptions about net benefits, because the safety valve provides a price ceiling while the intensity target essentially provides a price floor.

The safety valve was thus found in the single-period analysis to be preferable to the intensity target under most assumptions. A multi-period model was then used to compare the performance of the safety valve relative to borrowing. The optimal instrument choice was found to hinge on whether meeting the cumulative emissions cap over the specified period (2005-2050) was considered essential. Borrowing was found to be the optimal instrument for meeting the cumulative cap, though its performance is dependent on the foresight of regulated agents. The compensated safety valve's

performance is dependant on its design, and was found to perform almost as well as borrowing if the first period trigger price is set to the expected carbon price.

If the cumulative emissions cap is flexible, borrowing and the traditional safety valve are best compared in terms of net benefits. The safety valve was found to be preferred to borrowing if marginal benefits were constant at less than or equal to the expected level. Borrowing was preferred if marginal benefits were constant and greater than expected, though to a much lesser extent. If marginal benefits are non-constant, the safety valve is favored if the slope of marginal benefits is less than marginal costs, and borrowing is favored if the slope of the marginal benefits is greater than the marginal costs. The discount rate also affects policy choice; the higher the discount rate, the better borrowing performs relative to the safety valve.

6.1 POLICY IMPLICATIONS

From this analysis, several clear implications for policymakers emerge. These findings, which are intended to guide the choice and design of a cost-containment instrument, are detailed below.

1) Each of the cost-containment instruments examined are preferable to the pure cap

In terms of economic efficiency, the intensity target, safety valve and borrowing were all found to outperform the pure cap. Economic efficiency was measured both in terms of welfare loss (the loss in consumption resulting from the cost of the policy), and in terms of net benefits.

Borrowing and the compensated safety valve were compared to the pure cap in terms of welfare, as all policies resulted in an equivalent level of cumulative emissions. Both borrowing and the compensated safety valve had smaller welfare loss than the cap for any given cost shock. In addition, banking also improved welfare relative to the pure cap.

The intensity target resulted in higher expected abatement than the pure cap, but also higher abatement costs. The safety valve resulted in lower expected abatement than the pure cap, but had lower abatement costs. A more meaningful comparison of these instruments to the pure cap was in terms of net benefits. The intensity target had higher net benefits than the pure cap for each level of correlation between emissions and GDP examined. The safety valve was also preferable in terms of net benefits, unless the safety valve was designed as a very relaxed policy (low trigger price or loose cap).

In general, each of the above instruments was found to be more flexible than a pure cap, and this flexibility allowed cost savings. A pure cap requires an emissions target to be met, no matter how high the resulting cost is. Each cost-containment instrument provides an alternative option. In high-cost years, borrowing allows reductions to be reallocated to future periods, the intensity target relaxes the emissions target, and the safety valve imposes a tax rather than a cap. The resulting effect of each of these instruments is to require relatively more abatement in low-cost periods, and less abatement in high-cost periods. Since emissions reductions are required at the time when they are most cost-effective, net benefits under each of the instruments is higher than under a pure cap.

2) The preferred cost-containment instrument is dependant on several economic assumptions

The instrument which performs best in terms of economic efficiency is determined by several assumptions that go into this calculation. These assumptions, and their effects on the relative performance of cost-containment measures, are detailed below.

Correlation between BAU emissions and GDP growth

The single-period analysis showed the intensity target to be preferable to the reference safety valve only if BAU emissions and GDP had a correlation of 0.86 or greater. Though GDP growth is arguably the most influential factor BAU emissions (if it were not, regulators would do better to index the cap with the more influential factor), many other factors determine BAU emissions in a given year. Deviations from the expected demand for energy, for example through extreme weather or efficiency improvements in technology, will effect no-policy emissions, as will deviations from the expected supply of energy, for example through changing political regimes. For the intensity target to be the optimal cost-containment instrument, this other “noise” must be expected to constitute only small changes in no-policy emissions.

Additionally, lag-time effects may reduce the correlation between BAU emissions and GDP. These effects are not captured in the single-period model, as the quantity target in 2015 is simultaneously updated with the GDP growth rate. In reality, this is impossible, meaning that the emissions target would likely be based on

the GDP growth rate in the prior period. The growth rate will likely have changed by the time the emissions target has been implemented, further reducing the expected level of correlation. Given these compounding effects, it is unlikely that the correlation between BAU emissions and GDP is high enough for the intensity target to be the preferred instrument.

Importance of cumulative cap

A potential disadvantage of some cost-containment instruments is that they allow cumulative emissions over the policy period to vary from that required under the pure cap. This is the case for both the intensity target and safety valve. The intensity target has been shown to increase expected abatement relative to the pure cap, whereas the safety valve has been shown to decrease expected abatement. Though we may not be approaching a critical threshold of CO₂ emissions in the near future, there is strong belief by many that long-term abatement must be stabilized below some threshold. For example, an intended goal of the EU ETS is to ensure that global average temperature increases do not exceed pre-industrial levels by more than 2 degrees Celsius (Commission of the European Communities, 2007).

Borrowing and the compensated safety valve are both cost-containment instruments which enforce a cumulative cap, and thus may be preferred over either the safety valve or an intensity target. The reference safety valve enforces the cumulative emissions target only if costs are equal to or less than expected. The intensity target, though it leads to less cumulative emissions than the cap in the

expected case, allows emissions to exceed the cumulative cap if GDP growth is higher than expected.

Foresight and Rationality of Agents

This analysis has assumed that regulated agents are perfectly rational and have complete foresight. The assumptions cause borrowing to perform better than it is likely to in reality. In regard to rationality, agents may place a higher emphasis on short-term versus long-term considerations. Doing so could lead to excessive borrowing in early periods, with the costs of steep reductions delayed to a future date. Additionally, agents will not have perfect information about what the future holds. If the temporary shock is thought to be permanent, and agents expect that costs will remain high, net borrowing will be less than optimal. Similarly, if agents believe that the cap will be relaxed in the future (and it is not), net borrowing will be greater than optimal. If no borrowing occurs, the policy is analogous to having no cost containment. However, if too much borrowing occurs, the system is analogous to a compensated safety valve with a low trigger price, as the cost of making up future emissions will be very high. Such a policy could have greater welfare loss than a pure cap.

Slope and magnitude of marginal benefits

Given that the marginal benefits of abatement are highly uncertain, this analysis has tested the relative performance of these cost-containment instruments under a variety of assumptions. The safety valve requires the least abatement of all

instruments, and is the most efficient if marginal benefits of abatement have been overestimated. If the marginal benefits of abatement have been underestimated, the safety valve performs the worst of all instruments. The intensity target, which has the highest expected abatement, performs the best in this circumstance, followed by the instruments which enforce the cumulative cap (borrowing, compensated safety valve, and pure cap).

However, the relative gain in net benefits when the safety valve is preferred is far greater than the loss in net benefits than when any of the other policies is most preferred. Additionally, when net benefits are non-constant, but have a slope less than the marginal costs (likely the case for a stock pollutant such as carbon), the safety valve outperforms the other two instruments. The safety valve thus appears to be the most robust cost-containment instrument over the range of possibilities for the marginal benefits of carbon abatement.

3) Each cost-containment instrument's economic performance is sensitive to design

Regardless of which cost-containment instrument is selected, its efficiency is dependent on not only the factors mentioned above, but also on specific design choices. Furthermore, whether or not the intended cost-containment instrument is invoked at all depends on the design level of the cap. Under the "non-adjusted" path considered in this analysis, which is typical of most legislative proposals, the safety valve or borrowing would only have come into play under very large cost shocks. Rather than trigger the (reference) safety valve or borrow from future periods, agents

would simply have banked less. The relative quantity targets in each year, rather than the overall cumulative cap, determines the extent to which banking is able to dampen cost shocks. Relaxing the target in earlier years and tightening the target in later years both work to stimulate net banking in early periods.

Banking

Banking, like borrowing, is a form of inter-temporal borrowing, and its efficiency thus depends on the accuracy of agents' expectations. Uncertainty about the future price of carbon, the length of the policy period, or the penalties of future non-compliance may cause the amount of banking that occurs to differ from the optimum. Thus, in designing a policy with banking, every effort should be made to make these aspects of the policy transparent to regulated agents.

Intensity Target

The intensity target's performance is primarily related to the correlation between baseline emissions and GDP, but there is still some flexibility in design. The indexed cap can vary partially, rather than fully, with the GDP growth rate as was demonstrated in the single period analysis. Under a fully indexed cap, a percent change in the level of GDP translates to an equivalent percent change in the level of the cap. Under a partially indexed cap, a percent change in GDP translates to a smaller percent change in the level of the cap.

If marginal benefits are constant, a fully indexed intensity target is optimal. However, if marginal benefits are non-constant, a partially indexed target is optimal.

The steepness of the marginal benefits determines the optimal level of indexing for an intensity target. Very steep marginal benefits indicate that the indexed cap should vary only slightly with GDP, while flatter marginal benefits indicate that the indexed cap should vary more fully with GDP.

Safety Valve

The safety valve has two design variables that can be altered: the cap and the trigger price. The trigger price has been shown to be the more important of the design variables, but the level of the cap is far from inconsequential. A safety valve essentially specifies a tax and a cap and allows agents to choose the easier of the two. Thus, the only way to increase the stringency of a safety valve policy is to change the “easiest” of its two design variables.

The stringency of the safety valve is particularly important when a cumulative target is enforced. If the reference cap is combined with a low trigger price, the cost of making up emissions in later years exceeds the cost savings in the first period. This causes welfare loss under the safety valve to be substantially more than under the pure cap. However, if the trigger price is set too high, the policy cannot perform any worse than having no cost containment. Even if the trigger price is set so high that it is never binding, the policy will simply operate as a pure cap. Though no design of the compensated safety valve performs as well as borrowing, the reference cap combined with the reference trigger price is able to come close.

If a traditional safety valve (no cumulative target) is chosen, an overly stringent policy still performs better than an overly relaxed policy. The optimal

safety valve (assuming net benefits are constant) is one which acts like the optimal tax in most cases. This is achieved by combining a stringent cap with the reference tax. A safety valve which combines the reference cap and reference tax is the next-best in terms of net benefits.

Additionally, the extent to which banking may occur is essential to consider if banking is used in conjunction with a safety valve. The expected carbon price with net banking will generally be higher than the expected price without net banking. Thus, setting the trigger price without taking banking into account will result in a trigger price that is too low (and potentially worse than having no cost containment at all).

Borrowing

Borrowing, as modeled in the multi-period analysis, had no constraints or penalties associated with it. This design would allow maximum flexibility, though in practice its efficiency would be constrained by the accuracy of agents' expectations. The more constraints placed on borrowing, the smaller its advantage in relation to pure cap. However, constraints may have positive effect of discouraging agents from over-borrowing. Excessive borrowing in early periods leaves large reductions to be made later, which can lead to greater welfare loss than having no cost-containment at all.

Combined Instrument

The combined instrument considered in this analysis was an intensity target with a safety valve. This instrument was found to perform better, in terms of net benefits, than either an intensity target or safety valve used alone. The advantage of this instrument was that worked to not only restrain very high costs, but also to prevent very low costs. The intensity target essentially acts as a price floor (by making the cap more stringent in low-growth scenarios), whereas the safety valve acts as a price ceiling. Presumably, this type of policy could be replicated by having a safety valve with both a price ceiling (trigger price) and a price floor. Although the optimal level of these prices was not examined in this thesis, the above analysis suggests that cost containment may be important under low-cost as well as high-cost outcomes.

4) The choice of cost-containment instrument is necessarily linked with the choice of decision-making authority

A final consideration in the choice of cost containment instruments is the preference for the implied decision making authority. Borrowing and the safety valve differ widely in regard to the decision-making agent that they empower.

Borrowing, in its purest sense, is controlled by the market. In this case, its performance is determined by individual emitters' expectations about the future. These expectations may include the price of carbon in future years, the duration of the policy period, and the likelihood that the policy will be revoked or amended.

Emitters may also value short-term considerations differently than long-term ones, which could influence the amount of borrowing in early periods.

The Lieberman-Warner bill proposes to allow borrowing, and but differs from “pure” borrowing in that it stipulates the creation of a Carbon Market Efficiency Board. This board would be able to intervene in the market if carbon prices get too high, either by changing the limits or penalties on borrowing or temporarily adjusting the emissions cap. Thus, much of the decision making under this policy would fall not only to the agents but also to a specialized board.

The design of the safety valve, in contrast, would likely be decided by Congress and specified within the policy itself. The performance of the safety valve is thus dependant on policy makers’ foresight, as their estimate of the optimal cap and tax will determine the efficiency of the policy. A safety valve is less flexible than borrowing, as the carbon price specified before the cost of carbon abatement is known. The safety valve is also more transparent, however, since the trigger price is specified at the outset of the policy. Whether the cap or tax is the binding instrument is determined by the market, but is only a function of marginal costs rather than agents’ expectations about the future.

In summary, borrowing is the most flexible instrument and also the one which requires the least policy intervention. However, borrowing is also less transparent, especially if the constraints on borrowing and the level of the cap may be adjusted at any time. The safety valve provides the opposite; less flexibility with more transparency. The performance of both instruments, however, is dependant on the

foresight of the relevant decision maker – individual emitters or the Carbon Market Efficiency Board for borrowing, and Congress for the safety valve.

6.2 SUGGESTIONS FOR FURTHER WORK

The above analysis has led to many additional questions, and these questions suggest avenues for further research. This thesis concludes with several proposals for future studies.

1) Examine the effect of other types of shocks across multiple periods

The analysis above has examined only a temporary shock which occurs in the first period of the policy. This is a very specialized case, and may not be generalizable across other types of cost shocks. For example, the relative performance of cost-containment instruments may differ if the shock is permanent (correlated shocks across all periods). Additionally, temporary shocks which occur towards the middle or end of the policy may have a different effect than shock which occurs in the beginning. Testing the performance of these cost-containment instruments across a range of possible types of cost shocks could provide a more complete comparison.

2) Simulate imperfect decision making by assigning probabilities to potential outcomes

The results above are sensitive to the accuracy of agent decision making. The forward-looking model assumes that agents have perfect foresight, which is not realistic. Agents may have some expectations about the future, which may or may

not be realized, and are likely to hedge their position across a range of possible outcomes. Such a situation could be modeled by identifying possible outcomes (shock is temporary, shock is permanent, shocks occurs in multiple periods, etc), and assigning expected probabilities to each of them. This is likely to give a more realistic model of borrowing, given uncertainty about future cost shocks.

3) Test the sensitivity of the results to the expected length of the policy period

An additional uncertainty is the length of the policy period. Even though the policy period may be stated under the initial policy, agents are likely to be uncertain whether this policy will be extended, shortened, or revised before that date is reached. Agents' expectations of how future amendments to the policy will affect the current borrowing scheme are likely to guide their behavior. Additionally, the net benefits under a safety valve or compensated safety valve are likely sensitive to policy adjustments and the policy horizon. To incorporate this uncertainty, the sensitivity of results to longer and shorter periods, as well as to mid-period policy revisions, could be determined.

4) Further examine the advantage of an instrument with both a price ceiling and price floor

Finally, the performance of an instrument with both a price ceiling and floor could be tested more extensively. This thesis determined that an intensity target combined with a safety valve (which essentially acts as a floor and ceiling price) may be preferable to either individually. However, the analysis used only a single-period model and did not vary the stringency of either instrument. A more comprehensive

analysis could test a variety of floor and ceiling prices in response to shocks over a longer time horizon. Such analysis would require that both positive and negative cost uncertainty be considered.

In conclusion, the choice of cost-containment instrument is not a decision that is limited to economic considerations. If it were, the safety valve would prevail under most economic assumptions, so long as the cap and trigger price are well designed. However, other instruments may be still preferable (a cap with floor and ceiling price), and many additional factors complicate the best choice of instrument. This analysis has highlighted a few of those factors, and has suggested further research that will lead to more informed decision making.

7. APPENDIX I – SENSITIVITY ANALYSIS OF SINGLE-PERIOD CGE MODEL

The graphs below test the sensitivity of the single period model to each individual type of uncertainty which is changed within the model. The rate of aeei, gdp growth, and the individual elasticities of substitution are doubled and halved, and analyzed under the no policy, 5000mmt cap, and \$50/ton CO₂ carbon tax. baseline national income, energy use and emissions. As noted in Chapter 2, the elasticities of substitution considered are those for production (*q*), capital-labor (*kl*), energy-materials (*em*), inter-fuels (*e*), fixed factors (*f*), materials (*m*), domestic and export production markets (*t*), domestic and imported goods (*a*), and fossil and non-fossil resources (*f-nf*).




Figure 63: Sensitivity of no-policy emissions to AEEI, GDP growth, and individual elasticities of substitution (doubled and halved)




Figure 64: Sensitivity of no-policy emissions to individual elasticities of substitution (doubled and halved)




Figure 65: Sensitivity of emissions under a \$50 carbon tax to AEEL, GDP growth, and individual elasticities of substitution (doubled and halved)




Figure 66: Sensitivity of equivalent variation (welfare) under a \$50 carbon tax to AEEI, GDP growth, and individual elasticities of substitution (doubled and halved)




Figure 67: Sensitivity of equivalent variation (welfare) under a 5000mmt emissions cap to AEEI, GDP growth, and individual elasticities of substitution (doubled and halved)




Figure 68: Sensitivity of equivalent variation (welfare) under a 5000mmt emissions cap to individual elasticities of substitution (doubled and halved)




Figure 69: Sensitivity of emissions under a \$50 carbon tax to pairings of elasticities of substitution (both doubled and halved)




Figure 70: Sensitivity of equivalent variation (welfare) under a \$50 carbon tax to pairings of elasticities of substitution (both doubled and halved)

8. APPENDIX II – DETERMINATION OF OPTIMAL R*

Given quadratic costs and benefits measured in relation to the expected optimal quantity q^* , costs are given by

$$(1) \quad C(q) = c_0 + (c_1 - \theta_c)(q - q^*) + \frac{c_2}{2}(q - q^*)^2$$

where θ_c is a mean-zero random shock to marginal costs with variance σ_c^2 (Newell and Pizer, 2006). Benefits are given by

$$(2) \quad B(q) = b_0 + b_1(q - q^*) - \frac{b_2}{2}(q - q^*)^2$$

The optimal r^* (r^{**} in (Newell and Pizer, 2006)) is calculated by

$$(3) \quad r^* = (\sigma_{cx} / \sigma_x^2) / (b_2 + c_2)$$

where σ_x^2 is the variance of x , and σ_{cx} is the covariance of x and θ_c .

For the purposes of my data, I assumed θ_c to be the deviation from the expected carbon price under the cap. Since marginal costs were cubic rather than quadratic, I assumed c_2 to be the slope of the expected marginal cost curve at q^* .

REFERENCES

- AZAR, C. & DOWLATABADI, H. (1999) A Review of Technical Change Assessment of Climate Policy. *Annual Review of Energy and the Environment*, 24, 513-544.
- BABIKER, M., GURGEL, A., PALTSEV, S. & REILLY, J. (2008) A Forward-Looking Version of the MIT Emissions Prediction and Policy Analysis (EPPA) Model. Cambridge, MA, MIT.
- BOVENBERG, A. L. & GOULDER, L. H. (1996) Optimal Environmental Taxation in the Presence of Other Taxes: General-Equilibrium Analyses. *The American Economic Review*, 86, 985-1000.
- BROOKE, A. D., KENDRICK, A. M. & RAMAN, R. (1998) GAMS: a user's guide. GAMS Development Corporation.
- BUREAU OF ECONOMIC ANALYSIS (2007) Current-dollar and 'real' GDP.
- COMMISSION OF THE EUROPEAN COMMUNITIES (2007) Limiting Global Climate Change to 2 degrees Celsius. The way ahead for 2020 and beyond *Commission of the European Communities*. Brussels.
- ELLERMAN, A. D. & MONTERO, J.-P. (2002) The Temporal Efficiency of SO₂ Emissions Trading. *Cambridge Working Papers in Economics*.
- ELLERMAN, A. D. & SUE WING, I. (2003) Absolute vs. Intensity-Based Emission Caps. *Climatic Policy*, 2, S7-S20.
- ENERGY INFORMATION ADMINISTRATION (2007) Annual Energy Outlook 2007 with Projections to 2030.
- ENVIRONMENTAL PROTECTION AGENCY (2008) The United States Environmental Protection Agency's Analysis of Senate Bill S.2191 in the 110th Congress, the Lieberman-Warner Climate Security Act of 2008
- GURGEL, A., PALTSEV, S., REILLY, J. & METCALF, G. (2007) US Greenhouse Gas Cap-and-Trade Proposals: Application of a Forward-Looking Computable General Equilibrium Model. *Report 150*. Cambridge, MA, MIT Joint Program on the Science and Policy of Global Change.
- HOEL, M. & KARP, L. (2002) Taxes Versus Quotas for a Stock Pollutant. *Resource and Energy Economics*, 367-384.
- JACOBY, H. D. & ELLERMAN, A. D. (2004) The safety valve and climate policy. *Energy Policy*, 32, 481-491.
- KARP, L. & ZHANG, J. (2005) Regulation of Stock Externalities with Correlated Abatement Costs. *Environmental and Resource Economics*, 273-299.
- MATHIESEN, L. (1985) Computation of Economic Equilibrium by a Sequence of Linear Complementarity Problems. *Mathematical Programming Study*, 23, 144-162.
- NEWELL, R. G. & PIZER, W. A. (2003) Regulating Stock Externalities Under Uncertainty. *Journal of Environmental Economics and Management*, 45, 416-432.
- NEWELL, R. G. & PIZER, W. A. (2006) Indexed Regulation. *Discussion Paper No. 06-32*. Washington, D.C., Resources for the Future.
- NORDHAUS, W. D. (2007) A Review of the *Stern Review on the Economics of Climate Change*. *Journal of Economic Literature*, 45, 686-702.
- PALTSEV, S., REILLY, J., JACOBY, H., ECKAUS, R. S., MCFARLAND, J., SAROFIM, M. C., ASADOORIAN, M. & BABIKER, M. (2005) The MIT

- Emissions Prediction and Policy Analysis (EPPA) Model: Version 4. *Report 146*. Cambridge, MA, MIT Joint Program on the Science and Policy of Global Change.
- PALTSEV, S., REILLY, J., JACOBY, H., GURGEL, A., METCALF, G., SOKOLOV, A. & HOLAK, J. (2007) Assessment of U.S. Cap-and-Trade Proposals. *Report 146*. Cambridge, MA, Joint Program on the Science and Policy of Global Change, MIT.
- PARSONS, J. E. (2007) Stocks & Shocks: A Clarification in the Debate Over Price vs. Quantity Controls for Greenhouse Gases. Cambridge, MA.
- PIZER, W. A. (1997) Prices vs. Quantities Revisited: The Case of Climate Change. *RFF Discussion Paper 98-02*. Washington, D.C., Resources for the Future.
- PIZER, W. A. (1999) Optimal Choice of Policy Instrument and Stringency under Uncertainty: The Case of Climate Change. *Resource and Energy Economics*, 21, 255-287.
- PIZER, W. A. (2002) Combining Price and Quantity Controls to Mitigate Global Climate Change. *Journal of Public Economics*, 85, 409-434.
- PIZER, W. A. (2005) Climate Policy Design Under Uncertainty. *Discussion Paper No. 05-44*. Washington, D.C., Resources for the Future.
- QUIRION, P. (2005) Does Uncertainty Justify Intensity Emission Caps? *Resource and Energy Economics*, 27, 343-353.
- ROBERTS, M. J. & SPENCE, M. (1976) Effluent Charges and Licenses Under Uncertainty. *Journal of Public Economics*, 5, 193-208.
- RUTHERFORD, T. F. (1995) Extension of GAMS for complementarity problems arising in applied economic analysis. *Journal of Economic Dynamics and Control*, 19, 1299-1324.
- SCHWARTZ, E. & SMITH, J. E. (2000) Short-Term Variations and Long-Term Dynamics in Commodity Prices. *Management Science*, 46, 893-911.
- SCOTT, M. J., SANDS, R. D., EDMONDS, J., LIEBETRAU, A. M. & ENGEL, D. W. (1999) Uncertainty in Integrated Assessment Models: Modeling with MiniCAM 1.0. *Energy Policy*, 27, 855-879.
- STAVINS, R. N. (1996) Correlated Uncertainty and Policy Instrument Choice. *Journal of Environmental Economics and Management*, 30, 218-232.
- STOCK, J. H. & WATSON, M. W. (1988) Variable Trends in Economic Time Series. *The Journal of Economic Perspectives*, 2, 147-174.
- SUE WING, I., ELLERMAN, A. D. & SONG, J. (2006) Absolute vs. Intensity Limits for CO₂ Emission Control: Performance Under Uncertainty. *Report No. 130*. Cambridge, MA, Joint Program on the Science and Policy of Global Change, MIT.
- THE WASHINGTON POST (2007) Warming Proposals: The Presidential Candidates and Climate Change. *The Washington Post*.
- UNITED STATES SENATE COMMITTEE ON ENERGY AND NATURAL RESOURCES (2007) Low Carbon Economy Act Two-Pager.
- WEBSTER, M. D., BABIKER, M., MAYER, M., REILLY, J. M., HARNISCH, J., HYMAN, R., SAROFIM, M. C. & WANG, C. (2002) Uncertainty in Emissions Projections for Climate Models. *Atmospheric Environment*, 36, 3659-3670.

- WEBSTER, M. D., WING, I. S., JAKOBOVITS, L. & FELGENHAUER, T. (2007)
Uncertainty in Costs and Abatement from Near-Term Carbon Reduction Policies
in the U.S. *working paper*.
- WEITZMAN, M. L. (1974) Prices vs. Quantities. *The Review of Economic Studies*, 41,
477-491.
- WEITZMAN, M. L. (1978) Optimal Rewards for Economic Regulation. *American
Economic Review*, 68, 683-691.