# Introduction

**Preface/Acknowledgements**

The idea for this book grew out of a conference held in January 2010 in Washington, DC at Resources for the Future (RFF). The conference was a joint effort of the Energy Initiative at the University of Chicago, the Center for Energy Economics and Policy at RFF, and the University of Illinois at Urbana-Champaign. It brought together a diverse group of senior policymakers and leading researchers from the fields of economics, physical sciences, computational sciences, and law. Ten original research papers were presented, along with comments and spirited discussion focused both on technical modeling strategy and direct policy implications.[1] All papers have been revised and updated into ten chapters for this book, and commentators have added their written thoughts as well.

We are most grateful to Heidi Levin, then Executive Director of Chicago's Energy Initiative, who helped to organize and superbly managed the conference logistics. We could not have managed the conference without Heidi. Thanks also to Steve Goldberg of Argonne National Laboratory who allowed us to draw on his vast knowledge of energy policy and politics, while providing wise counsel and encouragement at every stage. We also thank Aysha Ghadiali at RFF for her steadfast logistical support. Of course, we thank all of the authors in this volume for their hard work on their papers and commentaries. Finally, we thank our sponsors – the US Environmental Protection Agency Climate Economics Branch, Resources for the Future, the Energy Initiative at the University of Chicago, and three units at the University of Illinois at Urbana-Champaign: the Environmental Change Institute (ECI), the Center for Business and Public Policy (CBPP), and the Institute of Government and Public Affairs (IGPA).

Mark A. Cohen, Vanderbilt University and Resources for the Future
Don Fullerton, University of Illinois
Robert H. Topel, University of Chicago

**Introduction**

The chapters gathered here represent the first installment on a regular conference event, under the banner *Energy Policy Symposium*. The theme for the 2010 conference, held in Washington, DC at Resources for the Future (RFF) was *Distributional Aspects of Energy and Climate Policies*. In particular, the chapters examine policies that would 'price' carbon emissions or otherwise seek to mitigate anthropogenic climate change. Since energy and climate policies are so closely linked and are on the public agenda, our first conference focused on both. In soliciting contributions, we interpreted 'distributional' fairly broadly, to include impacts of pending or possible legislation on the living standards of households across the US income distribution, across generations, or across geographic areas, as well as international differences in the costs and benefits of climate policies that would affect countries' willingness to participate in harmonized international agreements.

Most observers are, by now, familiar with the broad outlines of global climate change. The consensus is growing among scientific and government institutions that the earth's atmosphere and oceans are warming faster than natural causes can explain, and that these changes are caused by human activity. Climate change includes not only warming, which threatens biodiversity and agricultural productivity, but also sea level rise and increased severity of droughts, floods, and other extreme weather events like hurricanes. The most significant source of this anthropogenic climate change is the emission of carbon dioxide, which comes primarily from the combustion of fossil fuels, as well as from other sources such as deforestation and changes in land use.

Many governments, either unilaterally or in the context of international agreements, either have taken action, or are considering taking action, to limit carbon emissions and so combat climate change. Serious proposals to reduce carbon emissions generally include some kind of carbon pricing system. Under a carbon pricing system, governments intervene in the market to internalize the costs of discharging carbon into the atmosphere – that is, to shift these costs from the planet as a whole to the individual emitter (for example, a 'polluter pays principle').

This 'polluter pays principle' actually has two important but very different interpretations. First, the efficiency interpretation is that the polluters should pay the right price in order to get them to abate emissions by all of the least expensive means. Second, the equity interpretation is that polluters should have to pay for the damages they cause. While most research in environmental economics is concerned with the efficiency interpretation, this conference and collection of chapters is more concerned with the equity interpretation. We note, however, that any legal requirement for the polluter to pay for damages is only the 'statutory incidence'. It does not dictate the final 'economic incidence', in terms of who really bears the burden after adjustments in the marketplace, such as changes in output prices, factor prices, and scarcity rents. The main purpose of this symposium is to focus on economic incidence.

Carbon pricing systems generally take one of two forms: a tax on emissions, or a cap-and-trade regime. A tax or price could be applied directly on the emission of carbon dioxide ($CO_2$) or other greenhouse gases (GHG). Emitters pay to the government a legally determined amount based on the volume of $CO_2$-equivalent gases that they release into the atmosphere. Under a cap-and-trade system, the government issues permits authorizing the holder to emit a certain volume of those gases into the atmosphere in a given time frame. The sum of the volumes authorized by all permits issued equals the cap – the total amount that can be emitted for the whole nation (or region or even world). Permit holders are then free to trade the permits amongst themselves, allowing the price to be set by the market.

While economists long ago extolled the virtues of a market-based system of pricing externalities such as carbon emissions to achieve an efficient allocation of resources, the distributional consequences of any pricing policy will have important implications not only for those who stand to gain and lose but also for the political feasibility of any proposed solution. Ultimately, politicians will decide both the underlying framework (for example tax, cap-and-trade, or direct regulation) and the details of the proposed solution (for example who is regulated and by how much). Many factors impact the distribution of winners and losers, including the design of the tax or cap-and-trade system, the allocation of proceeds from the tax or permit sale, and the level at which the pricing system is adopted (state, national, or international levels). The interactions of these factors are enormously complex and subject to debate. Thus, serious analytical and empirical analysis can help identify the winners and losers as well as help craft policies that balance efficiency and political feasibility.

For example, revenue from a carbon tax or from the sale of emission permits under a cap-and-trade regime can be used in any number of ways. It can be distributed to the people using a variety of allocation formulas. It can be used as general tax revenue by the government. It can be directed to other programs designed to reduce carbon emissions, such as efficiency measures or research into clean technologies. Each of these alternatives implies different winners and losers, as well as different social costs and benefits.

The ten primary chapters (and ten commentary chapters) of this book address the distributional impact of efforts to mitigate global climate change by reducing carbon or other GHG emissions. That is, they seek to answer questions such as: to the extent that climate mitigation policies have costs, who will bear these costs? Will they fall on those of us living now, or on future generations? Will they fall on the rich, the poor, the middle class, or on everyone proportionally? Which countries will benefit, and which will suffer? If a global agreement is impossible, what will happen if countries act unilaterally, or in groups? What sort of nations would participate in such a pact? Would such a subglobal arrangement succeed in reducing global carbon emissions, or would it merely push carbon-intensive industries across borders?

Chapter 1, which is written by **Gary S. Becker, Kevin M. Murphy and Robert H. Topel**, begins by introducing the economic issues that underlie the problems of climate change and climate policy. The authors describe climate as a global public good, and they identify a number of complications that make climate change more than a simple externality problem. They stress the difficulty of valuing uncertain future damages and of setting an appropriate social discount rate. Notably, they suggest that a low discount rate may be appropriate in evaluating climate policies, because such policies may function as a kind of insurance – that is, they may pay off in times when other investments do not. This approach differs from a traditional economic approach that might suggest using a market rate of interest as a social discount rate. In addition, they note that different valuations of the marginal utility of future consumption across countries will make any global agreement that much more difficult to realize.

Chapter 2 is a comment by **Manasi Deshpande and Michael Greenstone** that applauds the paper's accomplishments, and then focuses on that last observation that differences in preferences across countries will make global agreement more difficult to accomplish. These commentators argue the opposite, namely, that differences between countries provide that much more 'gains from trade'. Developed countries might have more willingness and ability to pay for greenhouse gas emission reductions, but less developed countries might find less expensive ways of doing that abatement, so an appropriate agreement can provide gains to both.

In Chapter 3, **Louis Kaplow, Elisabeth Moyer and David A. Weisbach** delve more deeply into the issue of social discounting – which has a significant impact on the distribution of costs and benefits across generations. They draw a distinction between utility functions, in which the discount rate can be determined by empirical observation of actual human behavior, and social welfare functions, which require ethical judgments. The authors point out that the question of how people discount future consumption in practice is distinct from that of how we should discount the welfare of future generations. Small changes in a model's discount rate can result in large differences in the long-range projections and policy assessments produced using that model. The authors review two existing integrated assessment models, Nordhaus's

DICE model, and the PAGE2002 model used in the British government's Stern Review, which they believe fail to distinguish adequately between the empirical and ethical judgments involved. Ultimately, the authors challenge the concept that a single discount rate is truly applicable to social welfare problems, and they recommend drawing sharp distinctions between discounting in the context of utility functions versus discounting in the context of social welfare functions.

In a brief comment to Chapter 3, **Martin L. Weitzman** (in Chapter 4) is sympathetic to the call for clearly specifying whether the underlying discount rate is behavioral or ethically driven. However, he cautions that allowing for two distinct discount rates in integrated assessment models may turn out to be unwieldy and may raise more issues than it solves.

While the first four chapters focus primarily on the global public good nature of climate change and concern themselves with big picture issues such as the tradeoff between current and future generations, the next few chapters examine the impact of climate policy on the distribution of income for households and regions in the US.

In Chapter 5, **Don Fullerton and Garth Heutel** examine the distributional impacts of a carbon tax, using analytical general equilibrium results. They present a simple model dividing the economy into a 'clean' sector that employs both labor and capital and a 'dirty' sector that also creates pollution. The analytical model is used to show effects of a pollution tax on the returns to labor and capital (the 'sources side') as well as on the prices of clean and dirty outputs (the 'uses side'). Then they use household survey data to split the population into income deciles and calculate the burden on each income group. They find a regressive burden on the uses side and a burden on the sources side that can be progressive, regressive, or U-shaped. Notably, they find that once life-cycle effects are considered, by using expenditure deciles, the net distributional impact of a carbon tax is less clear.

In his comment, Chapter 6, **Samuel Kortum** provides additional intuition for their model by use of a diagram and analogy to a simple model of international trade. He also notes a 'fundamental inconsistency' in the Fullerton–Heutel paper, which employs the simplification of a representative agent model to solve for general equilibrium outcomes analytically, but then applies those output and factor price results to a set of ten different groups' expenditure patterns.

While many analysts prefer use of a carbon tax to reduce emissions, recent policy discussion in the United States has centered on three different cap-and-trade proposals that were proposed in the 111th Congress. In 2009, the House of Representatives passed the American Clean Energy and Security Act, more commonly known as the Waxman–Markey Bill after its authors, Congressmen Henry Waxman and Edward Markey. Two competing proposals were considered by the Senate: a bill introduced by Senators John Kerry and Barbara Boxer, which was very similar to Waxman and Markey's House Bill; and an alternative plan, put forward by Congresswomen Maria Cantwell and Susan Collins, which proposed a very different formula for the initial allocation of emission permits, and for the distribution of revenues from the auction of permits. In particular, the Cantwell–Collins Bill would have rebated a significant portion of auction revenues in a lump sum form to households to offset rising energy prices. While the 111th Congress ultimately did not pass cap-and-trade legislation, the outline of any future bill – whether it involves some form of emissions fee or trading scheme – will undoubtedly consider similar provisions. The next chapters model the impacts of these alternative cap-and-trade proposals, with special attention to distributional outcomes.

Significant points of disagreement among the different cap-and-trade proposals included the methods for allocating the initial carbon emissions permits (or equivalently, for distributing the revenue from those permits that were to be auctioned by government). In Chapter 7, **Joshua Blonz, Dallas Burtraw and Margaret A. Walls** examine the permit allocation framework included in the Waxman–Markey Bill. Specifically, they highlight uncertainty in the impacts of three provisions: (1) the free allocation of some emissions permits to local electricity distributors; (2) the dedication of a portion of carbon permit auction revenues to efficiency programs; and (3) the dedication of another portion of those revenues to clean technology development. Alternatively modeling optimistic and pessimistic assumptions regarding the effects of these provisions, they identify what they consider 'substantial' uncertainty regarding the cost of the Waxman–Markey proposal on an average household. For example, they estimate the annual cost to the average household to range between $133 and $418 – with this range being caused by uncertainty in the ultimate allowance price. In contrast, they model an alternative proposal in which 75 percent of allowances are auctioned, with funds returned to households in the form of lump-sum rebates (oftentimes called a 'cap-and-dividend' program). The authors find such a policy could both reduce uncertainty and produce a progressive distribution of burdens.

In a brief comment to the Blonz, Burtraw and Walls chapter, **Arik M. Levinson**, in Chapter 8, attempts to place these uncertainties into context. In particular, he notes that households face higher expected expenditure fluctuations due to uncertainties in the treatment of offsets and banking provisions of the law, as well as in the ultimate cost of energy. For example, annual fluctuations in the price of gasoline alone could affect the average household by plus-or-minus $2,000 – far in excess of any fluctuations due to permit prices.

In Chapter 9, **Ian W.H. Parry and Roberton C. Williams III** also examine the allocation of proceeds from a carbon permit auction. Specifically, they analyze four alternatives for distributing auction revenue: (1) granting all permits to existing emitters for free; (2) a full cap-and-dividend plan in which proceeds are distributed to households in equal lump sums; (3) using auction revenues for proportional offset of income taxes; and (4) a distribution-neutral income tax offset in which benefits are steered towards lower-income households to offset the regressive nature of the carbon price itself. They find that the overall welfare cost, as well as the distributional effects of cap-and-trade, depends greatly on how auction revenues are allocated, and they identify a clear tradeoff between efficiency and progressivity. The authors also emphasize the significance of the revenue-recycling effect, realized from using auction revenues to offset income or other distortionary taxes.

In a brief comment, Chapter 10, **William Randolph** notes that an important result of the Parry–Williams Model is that the fourth method – a distribution-neutral income tax offset – can be used to make the total efficiency cost of a cap-and-trade policy almost identical to the direct cost of emission reductions. That is, virtually all distortions can be eliminated. He also reminds us that like all of the chapters in this book, their model only considers the costs of climate policies – not the offsetting benefits.

In Chapter 11, **Sebastian Rausch, Gilbert E. Metcalf, John M. Reilly and Sergey Paltsev** employ the US Regional Energy Policy Model (USREP), a computable general equilibrium (CGE) model that has been used to examine energy and greenhouse gas policies on regions, sectors, industries and household income classes. The authors extended USREP using a recursive dynamic formulation to compare the distributional impacts of the three competing proposals in the 111th Congress. The authors find that carbon pricing itself is either distribution-neutral or progressive. The seeming conflict between this conclusion and that reached by Parry and Williams, who found pricing to be regressive when considered in isolation, stems from the inclusion of government transfers in the analysis. Rausch et al. assume that government transfers, that are received mostly by poorer households, are indexed to the price level in a way that protects those households from the costs of carbon pricing, which makes a carbon price more progressive than when transfers are excluded from the model. The authors find that the Cantwell–Collins 'cap-and-dividend' plan (which called for most auction revenues to be distributed to households via lump sum rebates), is less progressive than the more complex Waxman–Markey or Kerry–Boxer proposals. Like Parry and Williams, the authors identify a tradeoff between equity and efficiency, with the more progressive plans imposing a greater welfare cost than the less progressive Cantwell–Collins proposal. This paper also devotes some attention to the regional distribution of costs and benefits; they find that the Waxman–Markey and Kerry–Boxer proposals overcompensate some affected regions in the years immediately following adoption.

In a brief comment to Chapter 11, **Shanta Devarajan** points out in Chapter 12 some important limitations of this and other CGE models used to simulate the impact of climate policy. First, he notes that all general equilibrium models assume full-employment – something that hardly seemed realistic in 2010 with US unemployment hovering at 10 percent. Similarly, he notes that their recursive model is not truly dynamic and thus ignores many important consumer and producer choices. Thus, Devarajan points out the difficulty of any such modeling exercise – the tradeoff between realism and complexity. Finally, he points out another subtle difficulty in using the recursive model to predict the political feasibility of differing policy alternatives: since any progressive aspects of these policies are short-lived, it is not clear that the winners will lobby for them; nor will losers necessarily fight them.

In Chapter 13, **Dale W. Jorgenson, Richard Goettle, Mun S. Ho, Daniel T. Slesnick and Peter J. Wilcoxen** present a new method for evaluating the effects of carbon pricing regimes. The authors have revised the Intertemporal General Equilibrium Model (IGEM), originally introduced by Jorgenson and Wilcoxen in 1990 and now employed by the Environmental Protection Agency, to include a model of household behavior developed by Jorgenson and Slesnick. Using this updated model to incorporate labor–leisure choices into their analysis, they evaluate the impact of the Waxman–Markey Bill on 244 different types of households. Their conclusion reinforces the findings of some previous studies that such a regime, if adopted, would impact different regions and demographics in very different ways, but would, on balance, be regressive.

In his comment on Jorgenson et al. in Chapter 14, **Thomas Hertel** first articulates four important contributions: intertemporal decision-making, endogenous technical progress, substantial disaggregation among households, and econometrically estimated consumer and producer behaviors. He then notes that while the model has regional households, it does not have regional production of electricity with changes in relative costs that depend on differential coal intensity of particular regions. Also, many estimated parameters are used in this complicated CGE model, so Hertel recommends that the authors provide summary tables showing overall consumer responsiveness to energy prices – and then vary those parameters in sensitivity analysis.

The remaining chapters move beyond the confines of the US and consider global aspects of climate policy. In Chapter 15, **Joshua Elliott, Ian Foster, Kenneth Judd, Elisabeth Moyer and Todd Munson** present a new computable general equilibrium model called CIM-EARTH. The authors claim that many of the limitations of previous models are due to computational and methodological constraints that can be overcome using advances in computer architecture and numerical methods. CIM-EARTH allows for a detailed representation of individual industries, income cohorts and generations. The model also has the important feature that it is open and accessible to other researchers who may modify assumptions, test different data, develop more sophisticated uncertainty scenarios and see how they affect results. They demonstrate CIM-EARTH by analyzing the hypothetical adoption of emissions reduction by the nations of the Kyoto Protocol's Annex B. Specifically they predict the 'leakage' of emissions that would result from such an agreement, that is, the extent to which carbon emissions simply shift from participating countries to nonparticipating countries.

In a short comment on Elliott et al., **Don Fullerton** notes in Chapter 16 that it is really a progress report on a newer, larger, and even more complicated model. It will therefore run smack up against the main problem facing CGE modelers: the workings of the model are opaque to readers. They cannot see the effects of alternative assumptions, only effects of assumptions the authors choose to vary. In this case, Elliot et al. vary two assumptions about the baseline and thereby generate 25 different baselines. Each is used to calculate results of the same policy change. Readers do not see the effects of thousands of other assumptions, such as alternative aggregations, elasticity parameters, functional forms, abatement technologies, imperfect competition, factor mobility and a host of possible market imperfections not discussed in the paper.

In Chapter 17, **Christoph Boehringer, Carolyn Fischer and Knut Einar Rosendahl** consider subglobal climate agreements – that is, climate reduction policies adopted by countries or groups of countries without a fully global agreement. They examine the welfare effects of such unilateral policies in the US and the EU, both on the participating countries and on their trading partners. They also evaluate the likely effects of output-based rebating and border carbon adjustments, policies aimed at reducing 'leakage'. The authors find that the situation where countries participate in subglobal climate agreements would likely result in welfare losses both for the participating countries and for petroleum exporting nations, but could result in welfare gains for other trading partners. They predict that efforts to reduce leakage would likely be ineffective – reducing such leakage by at most 22 percent – and that such policies would have little effect, either positive or negative, on overall welfare.

In his comment (Chapter 18), **Rodney D. Ludema** points out that international trade is usually thought to impose extra costs for participating countries, whereas Boehringer et al. find the opposite. Unilateral climate policy helps them, as it reduces the worldwide price of oil and improves the terms of trade for participating countries. Because the border carbon adjustments effectively turn a tax on production into a tax on consumption, Ludema points out that they also change the burdens of the policy from producers to consumers.

In Chapter 19, **Charles D. Kolstad** examines the decision countries face in choosing to join or not join an international environmental agreement (IEA). Kolstad, expanding on previous modeling efforts that assumed identical countries, models the behavior of countries that differ in size, in the damage they suffer from warming, and in the costs they assume by adopting abatement measures. He finds that this heterogeneity has implications for both the composition and welfare impacts of IEAs. That is, a more heterogeneous pool of countries could result in an IEA that includes more countries but that imposes greater welfare costs on its members. Kolstad identifies two possible equilibria, maximizing either efficiency or equity. Then, under each such goal, he explores the likely form of the predicted international agreements.

In a brief comment in Chapter 20, **Scott Barrett** refers to an alternative modeling strategy that would yield an equilibrium in which *ex ante* both fairness and efficiency are achieved. He suggests that this model could be extended to the case of country heterogeneity outlined by Kolstad. However, he also notes that while this outcome appears consistent with the Kyoto Protocol, the fact that this treaty was unable to sustain a consensus reminds us of a major shortcoming of research that relies upon such models to predict behavior – in particular, the global community has thus far lacked an effective, credible punishment mechanism to enforce any such equilibrium.

## Concluding Remarks

Large-scale economic policy evaluation and calibration is always a daunting exercise, because we economists know so little about the underlying social and economic processes on which key outcomes depend. In describing the similar uncertainties that arise in military policy, former Defense Secretary Donald Rumsfeld usefully divided the space of relevant information into three parts – (1) the things we know; (2), the things we know we do not know; and (3) the things we do not know we do not know. Applied to evaluation of the economic impacts of energy and climate policies, the sad facts are that (1) is small, (2) is very large and – one suspects – (3) is biggest of all. Burdened with such ignorance, it is tempting to emulate the modesty of Mark Twain's famous response to a similarly difficult question: 'I was gratified to be able to answer promptly, and I did. I said I didn't know.'[2]

That will not occur, and not merely because economists are immodest. The development and use of abundant energy is essential to the maintenance and spread of economic well-being – rich countries use more energy per capita, and poor countries will use more if they are ever to grow rich. The fact is that new and far-reaching climate/energy policies will almost certainly be enacted – either unilaterally by certain developed countries or cooperatively through international agreements. Given the importance of energy to economic welfare, those policies have the potential to do great harm or good – with anticipated rent-seeking on a global scale, the cure could be worse than the disease. Thus, informed policy options are essential, especially with regard to economic impacts. But as Rumsfeld's partition demonstrates, part of being economically 'informed' is to acknowledge and incorporate the huge uncertainties underlying all policy choices in this area. Recognizing that we do not know much may be the most important aspect of wise energy policy.

Here is an example in the context of the distributional impacts discussed in several chapters of this book. Among the things we think we 'know' as economists are the directional 'signs' of most of the distributional effects of embedding a carbon price into the private cost of using traditional fuel sources. Economists are pretty good at signing those types of things, so we are confident that living standards of low-income households would be affected proportionally more because they spend a larger share of income on energy. And were policy designed to offset those effects in the name of equity, we will surely sacrifice some efficiency. With some nuances related to the details of various policies, that is what these papers found. The additional insights from our authors that would usefully inform policymakers come from putting numbers on the effects via tightly specified models – how large are the effects for various groups or countries, and how sensitive are they to various policy adjustments? What are the efficiency costs of redistribution? The danger is that such seeming precision gives a false sense of confidence, putting numbers on magnitudes we actually know very little about simply because our models are capable of generating them.

The point is stronger when we advance to Rumsfeld's second category – the large set of 'known unknowns'. Among these unknowns, to name just a few, are: the nature and likelihoods of catastrophic climate outcomes in the near or distant future; the pace of market-driven innovation of alternative energy sources; which countries will choose to participate in more-or-less harmonized climate policies and when; the pace of future economic growth; and the long- and short-run substitution opportunities that arise in myriad industries world-wide. As demonstrated in chapters 11, 12 and 13, the (dynamic) computable general equilibrium (CGE) models that very likely will be the workhorses of policy evaluation must assume particular structures and magnitudes for all of these possibilities, over very long horizons. And while advances in computing power allow us to build and analyze ever more sophisticated CGE models, the fact is that these models require vast amounts of empirical knowledge that is, by comparison, in the Stone Age – as if we had invented air conditioning but not electricity. After decades of empirical research, for example, economists have reached little consensus on such basics as the elasticity of labor supply or the willingness of consumers to substitute consumption over time, to say nothing of substitution over generations. Do not even ask what we know about affecting the pace of innovation or the design of research incentives. Now add in the 'unknown unknowns' and we rest our case.

We are not suggesting that the work not be done – what is the alternative? But rather that we acknowledge and even emphasize the fragility and uncertainty of what we can tell policymakers. Twain's prompt 'I don't know' is not so far from where we are, and it can be said with great confidence because it is true.

We close with two things that we do know. First, energy and climate policies are being implemented as we speak. Even in the absence of proactive legislative or regulatory action, existing taxes, subsidies, and public and private investment decisions are being made in anticipation of future policies and climate realities. Thus, any rigorous analysis based on the best available current data and modeling capability is better than none, if appropriately caveated and undertaken with serious sensitivity analysis.

Second, people live in both Sarasota and Saskatoon, and living standards in those locations are not much different. In other words, people are resourceful and, within some bounds, are able to accommodate their lives to significant differences in climate. With what seem to be fairly dim prospects for internationally harmonized policies that would do much good, adaptation is likely to be one method, or perhaps even the most important method, for dealing with climate change.

## Notes

1. All ten of these papers have been published in a special refereed issue of *The B.E. Journal of Economic Analysis & Policy*, along with a 'comment' on each of them.
2. Mark Twain, *Life on the Mississippi*.