

MIT Open Access Articles

Marginal Abatement Costs and Marginal Welfare Costs for Greenhouse Gas Emissions Reductions: Results from the EPPA Model

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Morris, Jennifer, Sergey Paltsev, and John Reilly. "Marginal Abatement Costs and Marginal Welfare Costs for Greenhouse Gas Emissions Reductions: Results from the EPPA Model." *Environmental Modeling & Assessment* 17.4 (2012): 325–336.

As Published: <http://dx.doi.org/10.1007/s10666-011-9298-7>

Publisher: Springer Netherlands

Persistent URL: <http://hdl.handle.net/1721.1/106558>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.



Marginal Abatement Costs and Marginal Welfare Costs for Greenhouse Gas Emissions Reductions: Results from the EPPA Model

Jennifer Morris · Sergey Paltsev · John Reilly

Received: 9 June 2010 / Accepted: 7 November 2011 / Published online: 27 November 2011
© Springer Science+Business Media B.V. 2011

Abstract Marginal abatement cost (MAC) curves, relationships between tonnes of emissions abated and the CO₂ (or greenhouse gas (GHG)) price, have been widely used as pedagogic devices to illustrate simple economic concepts such as the benefits of emissions trading. They have also been used to produce reduced-form models to examine situations where solving the more complex model underlying the MAC is difficult. Some important issues arise in such applications: (1) Are MAC relationships independent of what happens in other regions?, (2) are MACs stable through time regardless of what policies have been implemented in the past?, and (3) can one approximate welfare costs from MACs? This paper explores the basic characteristics of MAC and marginal welfare cost (MWC) curves, deriving them using the MIT Emissions Prediction and Policy Analysis model. We find that, depending on the method used to construct them, MACs are affected by policies abroad. They are also dependent on policies in place in the past and depend on whether they are CO₂-only or include all GHGs. Further, we find that MACs are, in general, not closely related to MWCs and therefore should not be used to derive estimates of welfare change. We also show that, as commonly constructed, MACs may be unreliable in replicating results of the parent model when used to simulate GHG policies. This is especially true if the policy simulations differ from the conditions under which the MACs were simulated.

Keywords Climate policy · Economic cost · Greenhouse gas emissions reductions · Marginal abatement cost

1 Introduction

Marginal abatement cost curves (MACs), relationships between tonnes of emissions abated and the carbon dioxide (CO₂) or greenhouse gas (GHG) price, have been the subject of many studies. In 1998, Ellerman and Decaux produced a much-used set of MACs from an early version of the MIT Emissions Prediction and Policy Analysis (EPPA) model. The EPPA model has since evolved a great deal [1], leading us to consider the re-estimation of a set of new MACs that better represent abatement costs as we now understand them given the advances and improvements we have made in modeling global greenhouse gas emissions. During this process, a number of issues have arisen: the stability of MACs to policies abroad, stability over time and dependency of MACs on previous policies, using MACs as a measure of welfare, and the inclusion of all GHGs. This paper explores these issues and offers sets of updated MACs from the EPPA model that analysts may find useful under some conditions as well as some cautions on their use.

Marginal abatement cost refers to the cost of eliminating an additional unit of emissions. Total abatement cost is simply the sum of the marginal costs, or the area under the MAC curve. A MAC curve for CO₂ emissions abatement can be constructed by plotting CO₂ prices (or equivalent CO₂ taxes) against a corresponding reduction amount for a specific time and region [2]. Construction of MACs involves multiple runs of a model to get different price–quantity pairs. MAC curves can be constructed for a single GHG or a combination GHGs if one has a weighting system for trading among them such as the Global Warming Potential (GWP) index. MACs have

Electronic supplementary material The online version of this article (doi:10.1007/s10666-011-9298-7) contains supplementary material, which is available to authorized users.

J. Morris · S. Paltsev · J. Reilly (✉)
Joint Program on the Science and Policy of Global Change,
Massachusetts Institute of Technology,
77 Massachusetts Avenue, E19-411,
Cambridge, MA 02139-4307, USA
e-mail: jreilly@mit.edu

been widely used as pedagogic devices to illustrate simple economic concepts such as the benefits of emissions trading. They have also been used to produce reduced-form models to examine situations where solving the more complex model underlying the MAC is difficult.

Ellerman and Decaux [2] and Klepper and Peterson [3] are two of the most commonly cited MAC studies. Ellerman and Decaux investigated the robustness of MACs with respect to different levels of abatement among regions and different scopes of emission trading. According to their definition, robustness refers to whether the MAC is virtually the same whatever the reductions of other countries. Klepper and Peterson also explored the robustness issue, arriving at somewhat different conclusions than Ellerman and Decaux.

Issues not explored by these previous authors include the stability of the MACs over time and closely related path dependency, whether measures of welfare can be derived from MACs, and the implications of expanding MACs to include all GHGs. MACs may change over time as a result of technological opportunities and resources and other conditions that may differ over time. By path dependency we mean specifically: Does a MAC constructed for a country in period $t = n$ depend on GHG policies implemented in periods $t = 0$ through $t = n - 1$? Many analyses, such as those that seek to demonstrate the potential benefits of emissions trading, must interpret MACs as equivalent to marginal welfare cost (MWC) curves. Since the EPPA model includes an explicit evaluation of welfare change, we are able to construct direct measures of MWC from the EPPA runs and compare them to MACs. Finally, since the early work of Ellerman and Decaux [2] and Klepper and Peterson [3], the importance of considering non-CO₂ GHGs in the design of policy has been realized and so we consider how that inclusion changes the basic shape of estimated MACs.

In Section 2, we describe the version of the EPPA model used here and construct MACs for a set of core cases. In Section 3, we also test the stability of MACs to policies abroad following protocols developed by Klepper and Peterson [3]. In Section 4, we explore issues that were not investigated by Ellerman and Decaux or Klepper and Peterson, including stability over time and path dependency, the relationship of MACs to MWCs, and the inclusion of all GHGs in a policy. Section 5 then takes into account all the issues previously discussed to develop a set of MACs that, if one must rely on them, are derived under conditions that are relevant to current policy discussions. Section 6 offers conclusions and cautions on the use of MACs.

2 The EPPA Model

To construct the MAC curve, we use version 4 of the MIT EPPA model. The standard version of the EPPA model is a

multi-region, multi-sector recursive–dynamic representation of the global economy [1]. In a recursive–dynamic solution, economic actors are modeled as having “myopic” expectations.¹ This assumption means that current period investment, savings, and consumption decisions are made on the basis of current period prices. This version of the model is applied below.

The level of aggregation of the model is presented in Table 1. Each country/region includes detail on economic sectors (agriculture, services, industrial and household transportation, energy-intensive industry) and a more elaborated representation of energy sector technologies. The model includes representation of abatement of both CO₂ and non-CO₂ greenhouse gas emissions (CH₄, N₂O, hydrofluorocarbons, perfluorocarbons, and SF₆). For non-CO₂ gases, calculations consider both the emissions mitigation that occurs as a by-product of actions directed at CO₂ and reductions resulting from gas-specific control measures.

When emissions constraints on certain countries, gases, or sectors are imposed in a computable general equilibrium (CGE) model such as EPPA, the model calculates a shadow value of the constraint which is interpretable as a price that would be obtained under an allowance market that developed under a cap and trade system. Those prices are the marginal costs used in the construction of MAC curves. They are plotted against a corresponding amount of abatement, which is the difference in emissions levels between an unconstrained business-as-usual reference case and a policy-constrained case. The solution algorithm of the EPPA model finds least-cost reductions for each gas in each sector, and if emissions trading is allowed, it equilibrates the prices among sectors and gases (using GWP weights). This set of conditions, often referred to as “what” and “where” flexibility, will tend to lead to least-cost abatement.

The results depend on a number of aspects of model structure and particular input assumptions that greatly simplify the representation of economic structure and decision making. For example, the difficulty of achieving any emissions path is influenced by assumptions about population and productivity growth that underlie the no-policy reference case. The simulations also embody a particular representation of the structure of the economy, including the relative ease of substitution among the inputs to production and the behavior of consumers in the face of changing prices of fuels, electricity, and other goods and services. Further critical assumptions must be made about

¹ The EPPA model can also be solved as a forward looking model [4]. Solved in that manner, the behavior is very similar in terms of abatement and CO₂-e prices compared to a recursive solution with the same model features. However, the solution requires elimination of some of the technological alternatives.

Table 1 EPPA model details

Country or region ^a	Sectors	Factors
USA	Non-energy	Capital
Canada	Agriculture	Labor
Japan	Services	Crude oil resources
European Union+	Energy-intensive products	Natural gas resources
Australia and New Zealand	Other industries products	Coal resources
Eastern Europe	Transportation	Shale oil resources
Former Soviet Union	Household transportation	Nuclear resources
India	Energy	Hydro resources
China	Coal	Wind/solar resources
Indonesia	Crude oil	Land
Higher-income East Asia	Refined oil	
Mexico	Natural gas	
Central and South America	Electric: fossil	
Middle East	Electric: hydro	
Africa	Electric: nuclear	
Rest of World (ROW)	Electric: solar and wind	
	Electric: biomass	
	Electric: gas combined cycle	
	Electric: gas with CCS	
	Electric: coal with CCS	
	Oil from shale	
	Synthetic gas	
	Liquids from biomass	

^aSpecific detail on regional groupings is provided in Paltsev et al. [1]

the cost and performance of new technologies and what might limit their market penetration. Alternatives to conventional technologies in the electric sector and in transportation are particularly significant.

We construct a set of “core” cases that assume all countries and regions pursue the same policy, which reduces CO₂ emissions by 1%, 5%, 10%, 20%, 30%, 40%, or 50% below the reference no-policy level in all time periods. The seven points on each of the MAC curves represent these seven reduction levels, and the absolute reductions resulting from each level can be read off of the *x*-axis of the graphs. The policy starts in 2010 and remains the same through 2050. So a 50% policy means that beginning in 2010 and each year thereafter, emissions must be reduced by 50% below reference. To focus on the domestic costs of abatement in each region, there is no emissions trading among regions/countries, but implicitly a trading system operates *within* each region/country. All marginal costs are expressed as 2005 dollars per tonne of CO₂, and the quantities of CO₂ emissions reduced are in million metric tonnes of CO₂. MACs are snapshots of costs at a particular point in time and change over time. To reflect this, Fig. 1 shows the MACs that result from these “core” assumptions for three different time periods (2010, 2020, and 2050) for USA, European Union, and Japan.

A striking result in Fig. 1 is that the later-year MACs are lower than earlier-year MACs, and this is especially pronounced in the 2050 MAC. In general, later-year MACs have a flatter shape, sometimes even an S-shape or step function. The result is due to (1) the availability of more technological options after 2020 and (2) path dependency which will be demonstrated more directly in Section 4.1. Several of the advanced electric generation options—those that feature CCS—are promising but are at a stage of development where most believe it would take something like a decade to get even a large-scale demonstration project completed. Thus, while EPPA has these technologies in the model, they are simply prohibited from entry until 2025 at the earliest. Thus, they have no effect on the MACs in 2010 and 2020.

As we will see in later sections, path dependency is playing a large role in the shape of the MAC in later years, and so this simulation design where high levels of abatement are simulated from 2010 onward, while useful in demonstrating behavior of the model, is generally not very realistic. Most policy proposals envision a gradual tightening of the reduction over time to avoid unnecessarily high near-term costs. Or policies envision banking and borrowing—allowing agents to reallocate reductions through time in an economically rational way to again avoid excessive near-term costs.

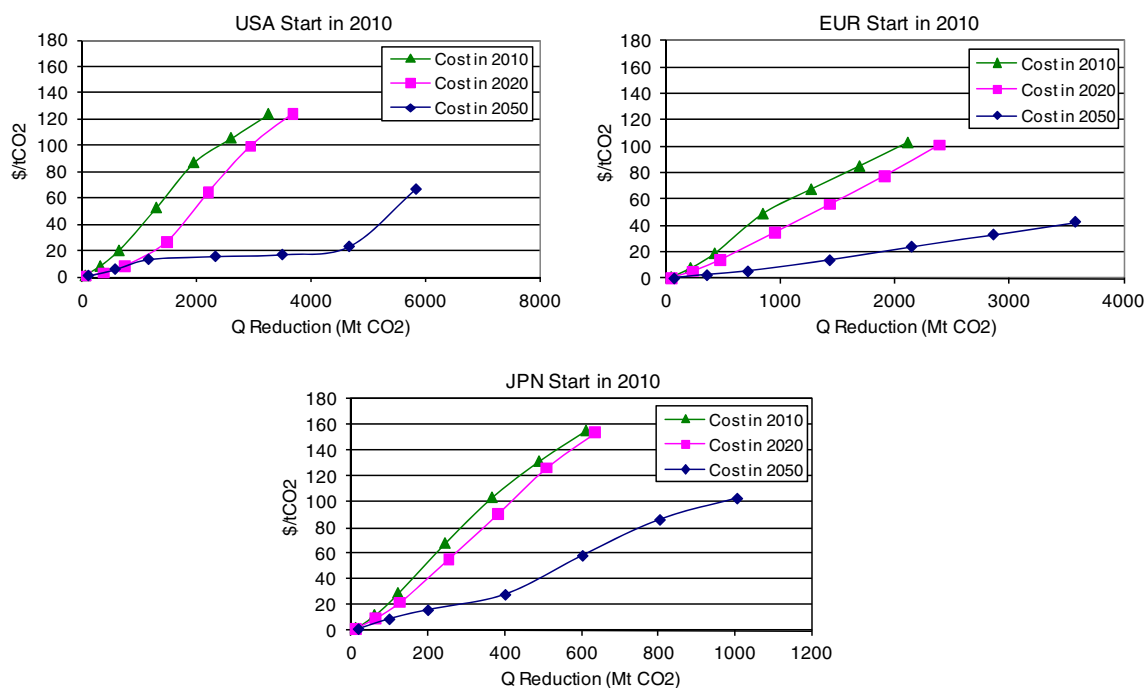


Fig. 1 MACs in 2010, 2020, and 2050 for the “core” scenario when the reduction policy is started in 2010 (points represent 1%, 5%, 10%, 20%, 30%, 40%, and 50% reductions)

3 Stability to Policies Abroad

One issue explored by Ellerman and Decaux [2] was the stability of MACs to policies abroad. They found that the EPPA-based curves were very stable and therefore robust to other countries’ behaviors. A later paper by Klepper and Peterson [3] challenged this stability using a different CGE model (DART). They attributed the shifts in national MAC curves to the changes in energy prices resulting from different global abatement levels. The difference in the approaches is a matter of defining the baseline. Emissions in a particular country, the USA for example, change depending on the policies abroad as those policies affect energy prices and have trade effects. The baseline from which to calculate the cost of reductions within the USA can either remain unchanged, always using the original emissions level resulting from original assumptions about policies abroad, or can change to represent the new emissions levels that result from changing policies abroad. The Klepper and Peterson approach always uses the original baseline, and as a result, a certain reduction in the USA would require a higher price (a shift of MAC) if the rest of the world had already acted and changed energy prices. The Ellerman–Decaux approach uses a baseline that changes depending on the policy abroad, and as a result, the cost of a particular reduction is stable to policies abroad. Klepper and Peterson found a shifting MAC when they held the baseline unchanged but obtained a similar result to

Ellerman and Decaux when they changed the baseline. There are actually additional ways to design the MAC construction. For example, one approach is to estimate each country’s abatement curve when all other countries are doing nothing, and then any terms of trade effects would be due only to actions within the country.

In our view, which approach to use depends on how one sees policy developing and how results from such an exercise might be used. If the country of interest, for example the USA, has remained out of an international agreement while other countries have committed to a clear policy, then the USA baseline should be the emissions that result in the USA given that specific international policy (i.e., USA does not participate in the policy but is affected by the actions of other countries). The MAC should then refer to reductions below that baseline, and the international policy should not be changed while constructing the MAC. This first case is that of a country considering unilateral policy action given that they know what policy the rest of the world is pursuing. The country’s reference should include the fact that other countries are committed to pursuing policies and take into account whatever impacts those policies have on energy markets. One would not want to have additional trade effects represented in either the position or the slope of the MAC.

Different situations may arise when a country is involved in multilateral negotiations. In such a situation, the baseline a country is working from likely reflects no additional action by others, as all countries would impose a policy

simultaneously. However, if the country has proposed additional cuts for itself and one would like to evaluate the costs of different levels of abatement in that country, then one would like a MAC that was shifted to represent other countries' proposed policies.²

How much do these different experimental designs affect the MAC? We demonstrate these differences with the EPPA model by creating three cases in which the USA does either a 0%, 1%, 5%, or 10% reduction in each of the years 2010, 2015, and 2020 while (1) the rest of the world does nothing (case denoted as "US ONLY"), (2) other Annex 1 countries reduce by 10% in 2010 and 2015 and by 20% in 2020 (case denoted as "ANNX1"), or (3) the rest of the world reduces by 10% in 2010 and 2015 and by 20% in 2020 (case denoted as "ROW"). The two constructions of the MACs of these policies are shown in Fig. 2 for the USA in 2020. Using the Klepper and Peterson experimental design, the abatement target is based on the baseline emissions in the US-only case, and thus, the target is the same for all three cases. When policies in other countries are considered, energy prices fall which leads to higher emissions in the USA (absent US policy). Therefore, the marginal cost of achieving the same absolute abatement target in the US increases as more countries participate in a policy. Similar to their result, the percentage differences in the MACs are quite large at low levels of reductions. Under the Ellerman and Decaux design, each case uses different baseline emissions so different absolute abatement targets result. The more policies abroad, the higher US emissions (absent US policy) so the greater amount of abatement required. The curves for the three cases lie nearly on top of one another, showing the stability Ellerman and Decaux found,³ but note that the cost of achieving a given percentage reduction (e.g., 10%) in the USA is higher when more countries have policies.

Like Klepper and Peterson, we find the difference in the MACs produced using their method to be the result of changes in energy prices caused by global abatement that affect countries even though the emission trading systems of countries are not linked. Mitigation policy abroad reduces the world oil price as countries demand less oil in order to meet their reduction targets. At lower oil prices, countries would be inclined to use more oil which would in turn create more emissions. Meeting a reduction target in the face of this situation would therefore require a higher CO₂ price to make alternatives economically attractive. In effect, the CO₂ price needs to be higher to make up for the

drop in the world oil price.⁴ Another energy price playing an important role is that of biofuels. More stringent mitigation policy abroad also leads to greater global biofuel use, and the resulting higher biofuel prices make reductions more expensive.

A few general observations: (1) The different experimental designs for constructing MACs can lead to fairly large differences. (2) There is no universally correct approach as it depends on how the MACs are being used to inform decisions and what other ancillary information is being used—whether shifts in the baseline are being considered separately or not. (3) There are an unlimited number of variants with more or less participation of other countries at different levels of abatement, responding or not to changes in the abatement level of other countries. (4) In principle, one would want to produce a set of MACs for the exact conditions one wished to examine but that entails running the parent model many more times to produce the MAC than would be necessary to simply examine the policy with the parent model. These considerations thus lead to the conclusion that any particular set of MACs can at best only provide a rough approximation of the marginal abatement cost in a particular country, and using them as a basis for a reduced-form model has limits in that they will not be completely consistent with different policies simulated with the parent model.

4 Additional Issues

In this section, we consider other important issues that were not investigated by Ellerman and Decaux or Klepper and Peterson. We look at potential path dependency, the relationship between MACs and welfare, and the inclusion of all GHGS in a reduction policy. We construct MACs for USA, Japan, European Union, China, India, and Middle East to illustrate the effects of these issues.

4.1 Path Dependency

Path dependency refers to whether a MAC constructed for a country in period $t=n$ depends on policies in periods $t=0$ through $t=n-1$. In order to explore this issue, we constructed MACs for 2050 for three cases that have different time frames of policy implementation. The first case is that used earlier where all countries are doing the same policy (1%, 5%, 10%, 20%, 30%, 40%, or 50% reductions) each period starting in 2010. The second case has all countries doing the same policy starting in 2050 and doing nothing before then. The third case develops a more

² The right set of MACs for a country is somewhat less clear when international emissions trading is allowed. See Morris et al. [5] for further discussion.

³ For additional comparison with the Ellerman and Decaux analysis, see Morris et al. [5].

⁴ For additional examples demonstrating this result, see Morris et al. [5].




Fig. 2 Robustness using different constructions of MACs for 2020: **a** Klepper–Peterson construction and **b** Ellerman–Decaux construction

realistic path of emissions reductions from 2010 to 2050 that gradually increases and incorporates a delay for developing countries. This path for developed regions (USA, European Union, Japan, Canada, Australia and New Zealand, Eastern Europe, and Former Soviet Union) and developing regions (the rest of the regions listed in Table 1) is detailed in Table 2.

Figure 3 shows the MACs for the three different cases for the USA, European Union, Japan, China, India, and Middle East in 2050. The differences in marginal abatement costs between the three cases are substantial. The figures clearly show that MAC curves for the same time period and region with the same constraint have different shapes depending on what policies were enacted in the past. The stronger and longer the policy in the past, the lower the marginal abatement costs for a given reduction in a given year. This path dependency is a major explanation for the flat and S-shaped 2050 MACs we saw in Section 2 and repeated here.

Interestingly, the 2050 MAC derived from the Path policy in Middle East does not begin near zero. Investigating further to find the source of the higher starting cost in the Middle East, we found that leakage from the rest of world was responsible. In particular, when we simulate the Middle East with no policy while the rest of the world follows the Path policy, emissions in Middle East are more than double those in the original reference case (no policy anywhere). This happens because energy activities are relocated from other countries into Middle East, where energy sources are located. This increased activity results in significantly higher emissions in Middle East than in the

original reference case. Thus, even returning to the original reference level emissions from this new higher level of emissions involves significant cost. This is a fairly extreme example of the Klepper and Peterson result. For the other regions, a 1% reduction in 2050 costs nearly nothing, regardless of what was done in previous years.

4.2 Measure of Welfare

Neither Ellerman and Decaux nor Klepper and Peterson investigated the relationship of MACs to welfare, even though many users of MACs integrate them to measure “gains from trade” which is a welfare concept. In an idealized neoclassical economic setting, first-order conditions from consumer welfare maximization involve consumers equating marginal welfare to the price of all goods and similarly producers setting marginal cost to the price of all goods. On that basis, the MAC and MWC curves should be identical. However, actual economies represented in computable general equilibrium models can diverge from this partial equilibrium, first-best neoclassical world. Goulder [6] showed generally that the CO₂ price can be a poor indicator of welfare when there are other distorting policies. Metcalf et al. [7] demonstrate how CO₂ pricing can interact with pre-existing energy taxes to exacerbate deadweight loss and raise the welfare cost of mitigation policy. Paltsev et al. [8] illustrate the effects of tax interactions and terms of trade effects diagrammatically and show how it can result in emissions trading being welfare worsening. Webster et al. [9] estimate welfare benefits for emissions trading in a stochastic setting using a

Table 2 Path policy

Year	2010 (%)	2015 (%)	2020 (%)	2025 (%)	2030 (%)	2035 (%)	2040 (%)	2045 (%)	2050 (%)
Developed	10	10	20	20	20	35	35	35	50
Developing	0	0	0	10	10	20	20	20	35

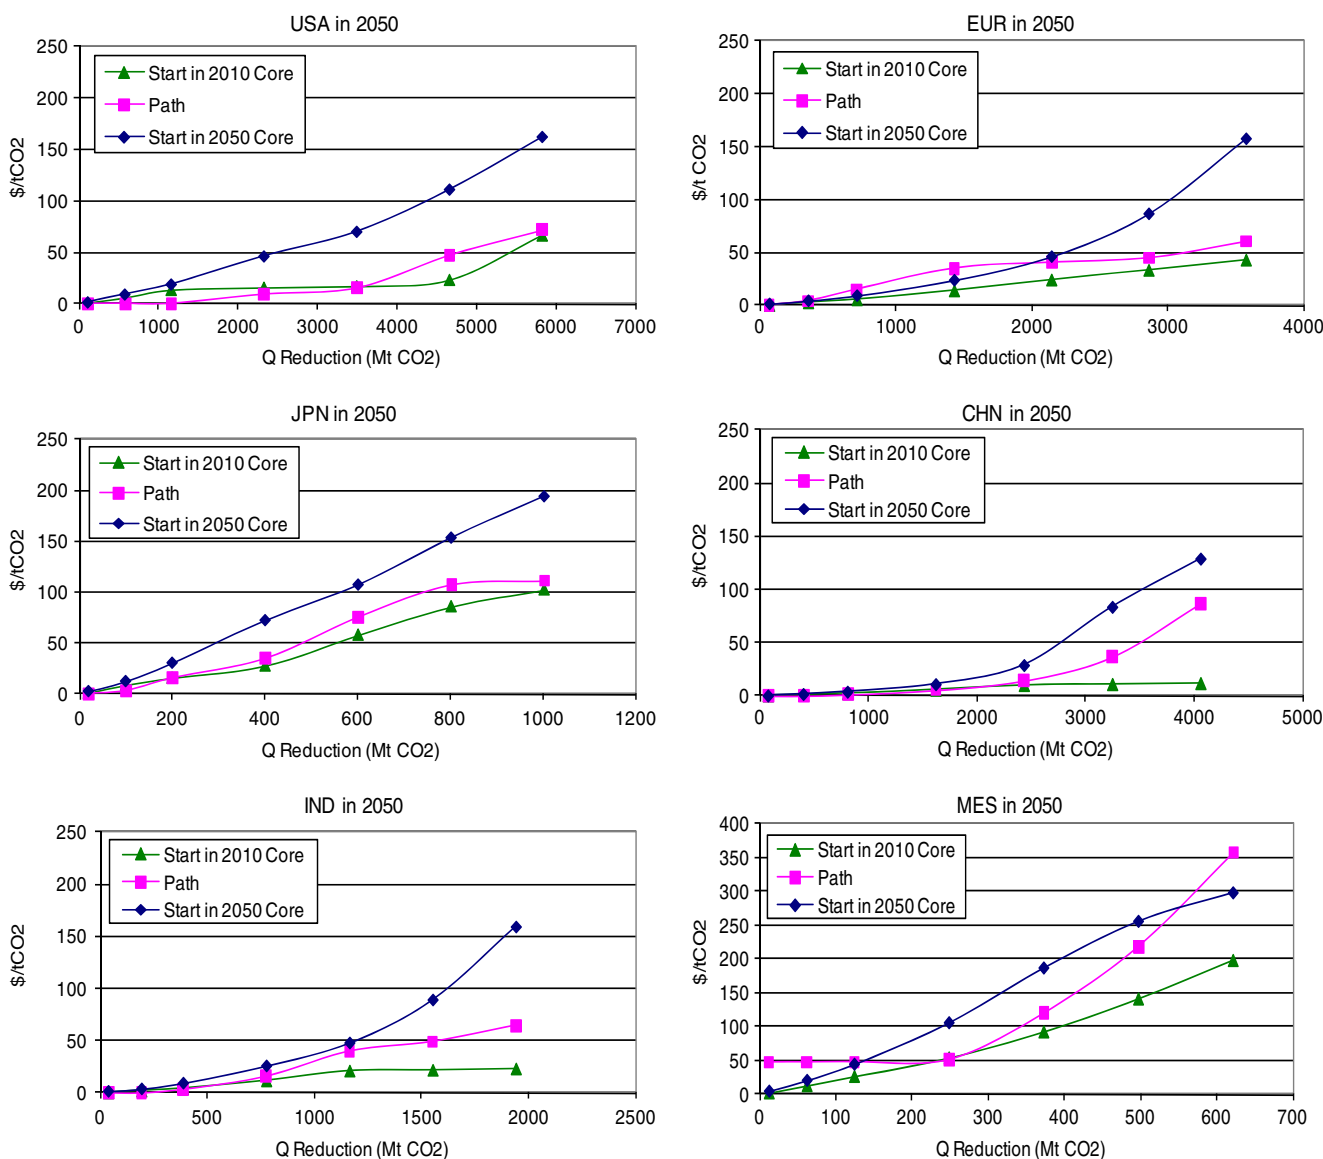


Fig. 3 MACs in 2050 when the reduction policy starts in 2010, 2050, or follows a reduction path

reduced-form MAC model and the parent model and show large differences. Thus, the fact that tax distortions and terms of trade effects, the source of instability in Klepper and Peterson’s [3] MACs, can also affect welfare estimates is not a new result. In this regard, it is useful to consider the welfare results from a CGE model to be driven by two components: (1) the direct welfare costs of abatement that can be measured as the integral under the MAC and (2) indirect welfare effects that involve terms of trade effects, interactions with other distortions, and saving and growth effects from policies in earlier years. Paltsev et al. [8] derive from a CGE model a method to estimate the direct costs and show them to be nearly identical to the MAC integration, but leaving a substantial residual difference when compared to the total welfare cost. So while the MAC

may be a good estimate of direct costs, it does not capture the indirect costs which are captured in the MWC.

Here we show divergence in welfare and marginal abatement costs by estimating MWC curves that can then be compared to MAC curves. To derive MWC curves, we note that marginal welfare cost refers to the welfare loss associated with abating an additional unit of emissions. For our welfare measure, we use equivalent variation and we monetize it as a change in aggregate market consumption for a representative agent in a region. The CO₂ price simulated by the model is a marginal concept, directly related to the shadow value of a Lagrangian maximization problem. However, the welfare index monetized as equivalent variation is a total cost concept—simply dividing the monetized welfare loss for a particular policy level compared to no policy gives an

average loss rather than a marginal loss. We therefore numerically approximate marginal welfare change by calculating the welfare change over a discrete but small change in the abatement level and then use the average welfare change over that small discrete change as an approximation of the marginal welfare change.

To calculate the marginal welfare cost, we ran the seven reduction policies (1%, 5%, 10%, 20%, 30%, 40%, and 50% reductions) plus another seven policies requiring an additional 1% reduction (i.e., 2%, 6%, 11%, 21%, 31%, 41%, and 51% reductions). We then calculated the change in welfare resulting from the additional 1% of reductions (for example, the change in welfare when comparing a 40% reduction to a 41% reduction). To make this cost measure comparable to the marginal abatement cost, we divided the monetized welfare change resulting from an additional 1% of reductions by the number of tonnes of CO₂ comprising that additional 1%. As a result, both the MWC and MAC are estimated in dollars per tonne of CO₂.

Figure 4 shows the MWCs and the MACs for the USA, European Union, Japan, China, India, and Middle East in 2010 for the case in which all countries do the same policy which starts in 2010. The basic result is that MWCs are not the same as the MACs, and they differ in some not unexpected ways given what has been learned in previous work. For example, high fuel taxes in Europe and Japan are tax distortions that are exacerbated by CO₂ policy and so the MWCs do not align very well with the MACs. The marginal deadweight loss can be many times larger than the direct cost when the CO₂ price is low because of these tax distortions and so we see especially at lower levels of abatement that the MWC is quite high compared with the MAC in these regions.

The USA has few such taxes and the MWC matches the MAC more closely. For the USA, we often see terms of trade benefits through the oil market, and so it is not surprising that we see the MWC to be somewhat below the MAC. Terms of trade benefits through energy markets also likely contribute to lower MWC for other regions that are net importers. For Middle East, we see the opposite—as a large energy exporter Middle East faces terms of trade losses that lead to MWC being far above the MAC.

We find that the difference between MWCs and MACs for countries in 2050 is significantly greater than in 2010. Welfare levels in 2050 are affected by the policy in 2050 directly (through marginal abatement costs), indirectly through terms of trade and interaction with distortions and, in addition, by previous year policies through effects on GDP, savings, and investment. Thus, it is not surprising that the marginal welfare cost in 2050 bears little resemblance to the MAC. Starting the policy in 2050 eliminates the GDP, savings, and investment effects from previous years, and the difference between the MWC and

MAC curves decreases significantly. Much of the very different MWC behavior in 2050 can thus be explained as the residual welfare effects of policies in prior years. After adjusting for those residual effects by beginning the policy in the year examined, the MWC and MAC are more similar, but are still not equal.

In general, the existence of strong shifts in the magnitude and sign of the indirect welfare effects—such as strong tax interaction effects at low levels of abatement and then possibly large terms of trade benefits at one point and large terms of trade losses at another—likely explains the roller coaster shape of the MWC curves that persist for some regions.⁵ The goal here is to visually demonstrate that, however one might fill in between the points we have simulated, it is clear that there are large differences between the MAC and the MWC.

Our general conclusion is that MWC and MAC curve comparison confirms a substantial body of literature that has shown welfare results from CGE models that cannot easily be explained by the CO₂ price. With a relatively simple CGE model—a static one period setting, no tax distortions, a small open economy, and/or no consideration of policies abroad—one might expect to see a close relationship between the MWC and MAC. But once in a dynamic setting with changing policies abroad, trade effects, and existing tax distortions, it is not surprising that there is little correspondence. While this result is disconcerting for MAC-based analysis, at the same time it offers little comfort for CGE analysis. It is unlikely that we could ever hope to accurately represent all of the various tax and other distortions in an economy, yet these results show that interactions with such distortions can dominate estimates of welfare changes. The analysis thus is a general caution about over-interpreting welfare results in a world that is obviously not the idealized one of neoclassical economics.

4.3 Other GHGs

When Ellerman and Decaux and Klepper and Peterson completed their work, many modeling exercises had not formally introduced the non-CO₂ GHGs. Modeling of the non-CO₂ GHGs has advanced, and policy discussions have also recognized the importance of including them. We therefore constructed MACs for policies aimed at all GHGs, allowing trading among them at their GWP weights. We apply the policy in which all countries pursue the same reductions in all years starting in 2010 to all GHGs to create the MACs in Fig. 5, which are compared with the MACs from the CO₂-only policy from Fig. 1.

⁵ For discussion of terms of trade effects in specific policy settings in the US, see Paltsev et al. [10].

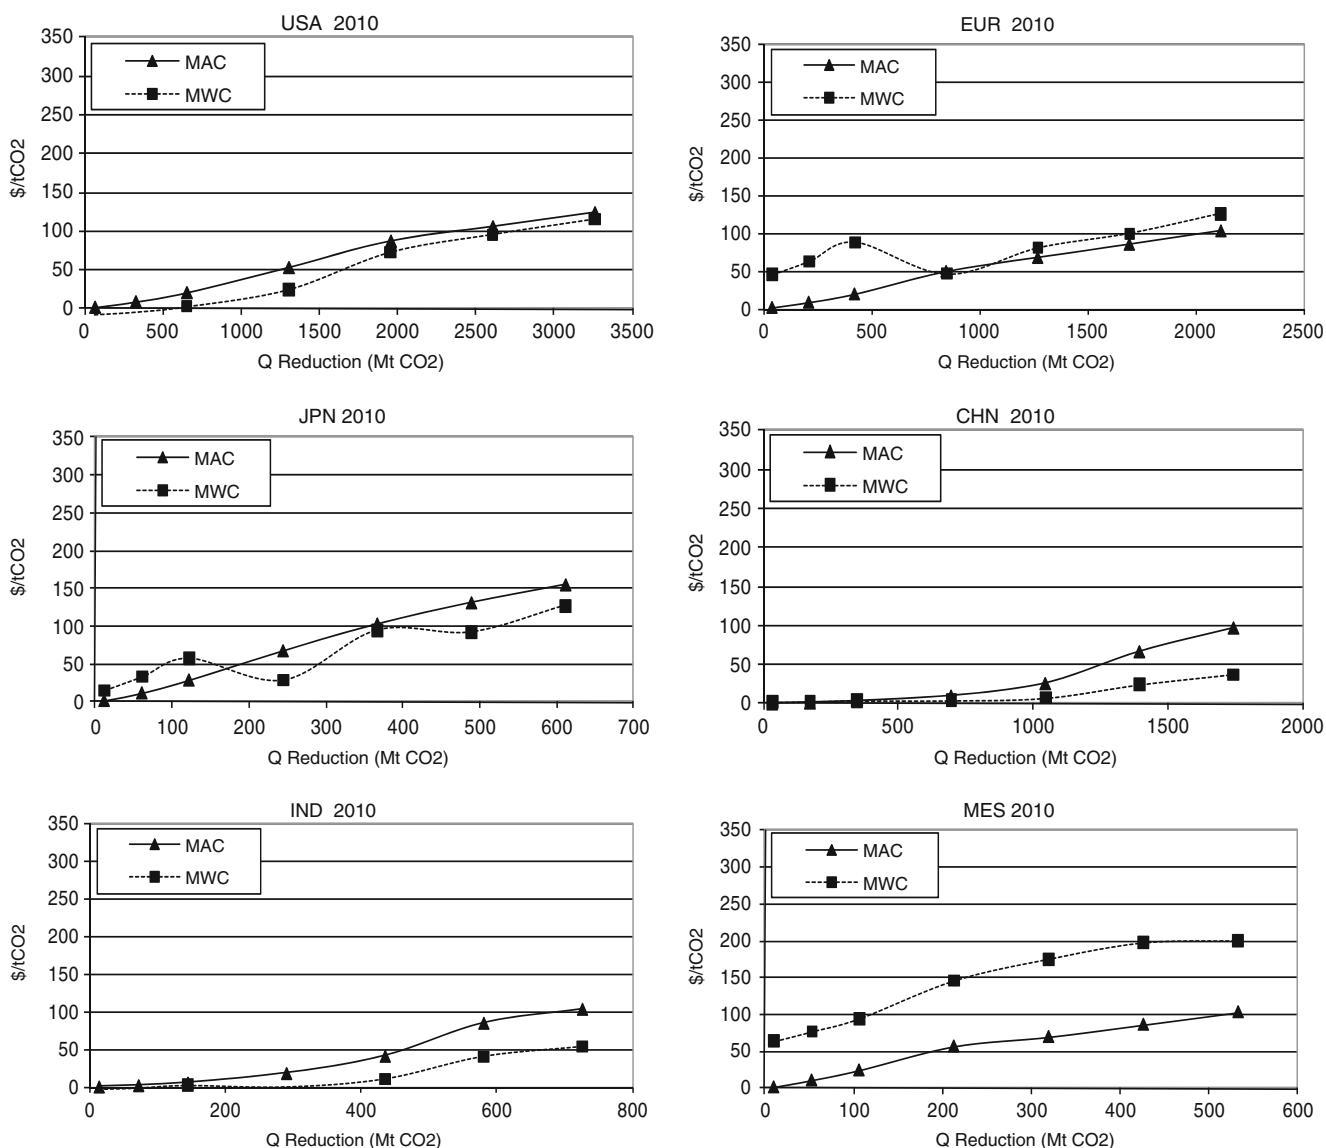


Fig. 4 Marginal abatement cost curves and marginal welfare cost curves in 2010 when the reduction policy starts in 2010

The inclusion of all GHGs expands abatement opportunities especially at low marginal abatement costs. Many non-CO₂ gases offer relatively inexpensive abatement opportunities, especially when one considers their high GWPs. These opportunities create a low shallow slope in the initial part of the MAC, essentially shifting the MAC outward. Once the non-CO₂ gases are mostly controlled, no more abatement opportunities exist for them and the remainder of the curve involves mostly CO₂ reductions.

5 If You Must Have MAC Curves

In this section, we make a best attempt to estimate MACs under the type of conditions that are relevant to existing policy discussions. We present graphs for a set of regions

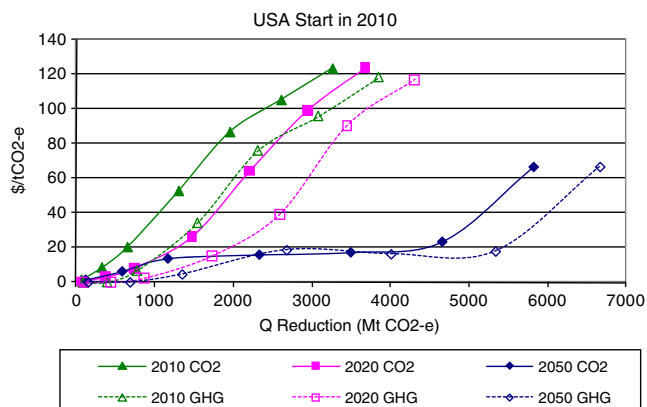


Fig. 5 MACs in 2010, 2020, and 2050 when the policy started in 2010 applies to just CO₂ and when it applies to all GHGs

for 2020. In Online Resource 1, we provide these more realistic MACs (in Excel format) for all regions for 2010 and 2050 in addition to 2020. The MACs were simulated to include all GHGs trading at GWP weights. The path dependency issue means we must consider carefully the policy environment over the full time horizon. The international policy environment is also of importance. We thus follow the reduction path for the developed and developing world previously detailed in Table 2. This path involves gradual tightening of the policy over time with reductions delayed in the developing countries. Each country's MAC is estimated separately, with other countries at the Table 2 level in that year and prior years. Constructing the 2020 MACs holds in place the 2010 and 2015 policy as described in Table 2 and simulates the model for each MAC point in 2020 for each country.

The graphed MAC data are shown for several example countries in Fig. 6. We illustrate them in two sections, one section with a heavy line and one with a lighter line. The different line weights are used to convey the idea that, given the construction approach, some parts of the MAC are more relevant than others for the year being considered. The lower parts of the MACs are shaded more heavily for 2020. It seems less likely that a country would switch from a mild or no policy in 2015 to a 40% or 50% cut in 2020. A more realistic estimate of the cost of large cuts in 2020 should probably be simulated assuming deeper cuts in 2010 and 2015.

6 Conclusions

Many analysts have found marginal abatement cost curves to be useful devices for illustrating economic issues associated with greenhouse gas abatement. As pedagogic tools, they follow in a long tradition in economics of using graphical analysis of supply and demand curves to represent markets for normal goods. In many applications, the use of MACs has gone well beyond pedagogy. They have been used as the basis for reduced-form models to help illustrate likely CO₂ prices, emissions trading, and welfare costs of different abatement levels. For such purposes, one would like to have some confidence that a MAC-based analysis would provide the same result as the parent model from which it was derived. Early analyses of MACs under relatively limited conditions suggested a somewhat surprising robustness. The specific test of robustness was whether a MAC in one country was affected by the level of the policy in another country. If the MAC is affected by the level of policy elsewhere and the intention is to examine emissions trading when the level of abatement in other countries is changing in the analysis, then an unstable MAC would clearly create inaccuracy in the

results. Later work formulated the test of robustness somewhat differently and found more instability in the MAC. Essentially, policies abroad create a shift in the baseline for a country through changes in prices in energy markets. If this shift is taken into account in the construction of the MAC, as the earlier analysis did, then the MAC appears stable. If, however, this shift is not accounted for, then the MAC shifts and, especially at low levels of CO₂ prices, this can lead to very large errors (in percentage terms) in predicted CO₂ prices.

Is there a best practice in how to construct MACs? We argue that it depends on how the MAC is to be used. Whether, for example, the baseline change from policy abroad is explicitly (or implicitly) taken into account will determine which of the approaches to MAC construction is more accurate. One approach to constructing MACs is to simulate different levels of abatement in all countries at the same time. This approach introduces the baseline shift into the MAC of each country by gradually changing the MAC slope at higher CO₂ prices—when no one is abating, there are no energy market effects from abroad but these become bigger as the abatement level becomes stronger everywhere. This could be appropriate for some purposes—in negotiations for example, if one imagined other countries matching your offer on how much to abate, then a MAC constructed in this manner would give you an accurate measure of the CO₂ price you could expect in your country. In general, however, the baseline shift will introduce some inaccuracy, and earlier analysis that demonstrated stability did so under a special case that may not be appropriate to the many uses to which MACs have been applied.

Given how MACs have come to be used, the robustness of MACs in one country to changes in policy in another is a relatively limited test. We also investigated their stability over time and whether there was path dependence—the extent to which a MAC in later years depended on the abatement level in earlier years. We examined the relationship of the MAC to MWC, and we extended MACs to include non-CO₂ GHGs. In general, these investigations revealed far greater inaccuracies and instabilities in MACs than the single period analysis of impacts of abatement in other countries. These findings suggest caution in applying MACs other than for the simplest of illustrations.

With regard to stability over time, MACs, at least those derived from the parent model we used, changed greatly from period to period. If this were solely the result of changing technological opportunities over time, then a set of MACs could be generated to represent each time period. However, we found strong path dependence—MACs in later years were strongly affected by the level of abatement simulated in earlier years. Thus, MAC analyses that consider dynamics of abatement—banking and borrowing—must be considered suspect.




Fig. 6 More realistic MACs for 2020

Given a variety of previous analyses comparing welfare derived from CGE models to measures of welfare derived from a MAC analysis, we expected MACs to be a poor indicator of MWC. What was surprising was how little correspondence there was between the MAC and the MWC. MWC was far above the MAC for some regions and far below for others, and the relationship could change substantially over the range of CO₂ costs represented in the MAC. MWCs were particularly sensitive to policies in previous years. Upon reflection, this result is not too surprising since saving and investment as it is affected by policies in earlier years will obviously carry over to affect welfare in later years as a completely separate influence from any mitigation policy in the later year. Extending the MACs to include other GHGs also substantially changed

the shape of the MAC, lowering the slope of the MAC at low CO₂ prices.

The use of the MACs has been popularized by McKinsey & Company, which published their original MAC for the USA in 2007 [11] and produced several updates later, including MACs for other regions.⁶ Many analysts have used these and other MACs. They offer an easy to understand visualization of how costs depend on the level of abatement. Unfortunately, unless one takes great care in understanding the exact conditions under which MACs are constructed and constructs them for the specific use in mind, it is very easy to misuse them. By misuse we

⁶ For a detailed critique of the McKinsey approach, see Kesicki and Ekins [12].

mean exercising MACs under conditions where the results they provide would differ substantially from the result one would get from running the parent model. There are of course great uncertainties in estimating costs, and different parent models will yield very different results and so perhaps the errors introduced by using simplified MACs are swamped by differences among the more complex models anyway. However, we can trace these differences to specific structural considerations and feedbacks in the parent model, and once one is aware of these processes and can model them, it seems a mistake to simply ignore these effects to avoid running the parent model. With cautions in mind, we present in the final section of the paper MACs derived from our EPPA model for various regions for 2020 under a specific set of assumptions about how policy will evolve over that time. In developing the MACs with an eye toward the possible evolution of policy over the next few decades—or at least the types of policy paths that are being investigated as we write this report—they may provide a rough indication of potential abatement costs if used carefully.

Acknowledgments We thank Denny Ellerman for his valuable comments. We are also thankful to Jaemin Song for research assistance on an earlier version of this paper. The EPPA model used in the analysis is supported by a consortium of government, industry, and foundation sponsors of the MIT Joint Program on the Science and Policy of Global Change (<http://globalchange.mit.edu>).

References

- Paltsev, S., Reilly, J., Jacoby, H., Eckaus, R., McFarland, J., Sarofim, M., Asadoorian, M., Babiker, M. (2005). The MIT Emissions Prediction and Policy Analysis (EPPA) model: Version 4. MIT Joint Program on the Science and Policy of Global Change Report 125, August. http://web.mit.edu/globalchange/www/MITJPSPGC_Rpt125.pdf. Accessed 12 Oct 2011
- Ellerman, A. D., & Decaux, A. (1998) Analysis of post-Kyoto CO₂ emissions trading using marginal abatement curves. MIT Joint Program on the Science and Policy of Global Change Report 40, October. http://web.mit.edu/globalchange/www/MITJPSPGC_Rpt40.pdf. Accessed 12 Oct 2011.
- Klepper, G., & Peterson, S. (2006). Marginal abatement cost curves in general equilibrium: The influence of world energy prices. *Resource and Energy Economics*, 28(1), 1–23.
- Gurgel, A., Paltsev, S., Reilly, J., Metcalf, G. (2007) U.S. greenhouse gas cap-and-trade proposals: Application of a forward-looking computable general equilibrium model. MIT Joint Program on the Science and Policy of Global Change Report 150, June. http://globalchange.mit.edu/files/document/MITJPSPGC_Rpt150.pdf. Accessed 12 Oct 2011.
- Morris, J., Paltsev, S., Reilly, J. (2008). Marginal abatement costs and marginal welfare costs for greenhouse gas emissions reductions: Results from the EPPA model. MIT Joint Program on the Science and Policy of Global Change Report 164, November. http://globalchange.mit.edu/files/document/MITJPSPGC_Rpt164.pdf. Accessed 12 Oct 2011.
- Goulder, L. (1995). Environmental taxation and the ‘double dividend’: A reader’s guide. *International Tax and Public Finance*, 2, 157–183.
- Metcalf, G., Babiker, M., Reilly, J. (2004). A note on weak double dividends. *Topics in Economic Analysis & Policy*, 4(1), Article 2.
- Paltsev, S., Reilly, J., Jacoby, H., & Tay, K. (2007). How (and why) do climate policy costs differ among countries? In M. Schlesinger, H. Kheshgi, J. Smith, F. C. de la Chesnaye, J. Reilly, T. Wilson, & C. Kolstad (Eds.), *Human-induced climate change: An interdisciplinary assessment*. Cambridge: Cambridge University Press.
- Webster, M., Paltsev, S., & Reilly, J. (2010). The hedge value of international emissions trading under uncertainty. *Energy Policy*, 38(4), 1787–1796.
- Paltsev, S., Reilly, J., Jacoby, H., Gurgel, A., Metcalf, G., Sokolov, A., & Holak, J. (2008). Assessment of U.S. GHG cap-and-trade proposals. *Climate Policy*, 8(4), 395–420.
- Creyts, J., Derkach, A., Nyquist, S., Ostrowski, K., Stephenson, J. (2007). Reducing U.S. greenhouse gas emissions: How much at what cost? *U.S. Greenhouse Gas Abatement Mapping Initiative, Executive Report, December*. New York: McKinsey & Company.
- Kesicki, F., & Ekinsm P. (2011). Marginal abatement cost curves: A call for caution. *Climate Policy*, in press.