

**PRICES IN EMISSIONS PERMIT
MARKETS: THE ROLE OF
INVESTOR FORESIGHT AND
CAPITAL DURABILITY**

Bryan K. Mignone

Brookings Energy Security Initiative

Working Paper 08-01

November 2008

Prices in Emissions Permit Markets: The Role of Investor Foresight and Capital Durability

Bryan K. Mignone^{1,2*}

¹Foreign Policy Program, The Brookings Institution, Washington, D.C., USA

²Centre for Applied Macroeconomic Analysis, The Australian National University, Canberra, Australia

This working paper is identical to Centre for Applied Macroeconomic Analysis Working Paper 31/2008 available at <http://cama.anu.edu.au/publications.asp>. The author acknowledges financial support from the John D. and Catherine T. MacArthur Foundation and thanks Richard Richels of the Electric Power Research Institute for making available the MERGE model code. In addition, the author acknowledges helpful discussions with Billy Pizer, Adele Morris and Paul Higgins.

* Corresponding author's E-mail: bmignone@brookings.edu

Prices in Emissions Permit Markets: The Role of Investor Foresight and Capital Durability

Abstract

Of the many regulatory responses to climate change, cap-and-trade is the only one currently endorsed by large segments of the scientific, economic and political establishments. Under this type of system, regulators set the overall path of carbon dioxide (CO₂) reductions, allocate or auction the appropriate number of emissions allowances to regulated entities and – through trading – allow the market to converge upon the least expensive set of abatement opportunities. As a result, the trading price of allowances is not set by the regulator as it would be under a tax system, but instead evolves over time to reflect the underlying supply and demand for allowances. In this paper, I develop a simple theory that relates the initial clearing price of CO₂ allowances to the marginal cost premium of carbon-free technology, the maximum rate of energy capital replacement and the market interest rate. This theory suggests that the initial clearing price may be lower than the canonical range of CO₂ prices found in static technology assessments. Consequently, these results have broad implications for the design of a comprehensive regulatory solution to the climate problem, providing, for example, some intuition about the proper value of a possible CO₂ price trigger in a future cap-and-trade system.

Keywords: Environmental regulation; instrument choice; climate policy; cap-and-trade; carbon tax; energy modeling; integrated assessment.

1. Introduction

Cap-and-trade is one of the few regulatory responses to climate change capable of simultaneously securing support from large segments of the scientific, economic and political establishments. The basic approach is attractive to the science policy community because binding quantitative limits (the "caps") are consistent with the notion of a scientifically-informed emissions reduction path (e.g. *Mignone et al., 2008; O'Neill et al., 2006; Schneider and Mastrandrea, 2005; Pacala and Socolow, 2004; Caldeira et al., 2003; O'Neill and Oppenheimer, 2002*). At the same time, the approach is generally endorsed by economists because the implied flexibility in compliance (the "trade" aspect, among others) fosters an economically efficient outcome by lowering compliance costs relative to more rigid "command and control" alternatives (e.g. *van Vuuren et al., 2006; Newell et al., 2005; Carlson et al., 2000*). Support from the political establishment, measured imperfectly in the US by the number of economy-wide cap-and-trade proposals introduced in the 110th Congress, is primarily a function of this underlying structural flexibility and the fact that it provides a natural platform for compromise among affected stakeholders.¹

In the wider taxonomy of regulatory solutions, a quantity-based system like cap-and-trade is the obvious alternative to a price-based mechanism like a tax. Both can be designed to yield economically efficient outcomes, but the two approaches differ in their respective responses to future technological uncertainty. Cap-and-trade provides certainty in the quantity of emissions mitigation at the expense of certainty in the resulting price of allowances, while a tax provides certainty about prices at the expense of certainty in the quantity of mitigation. This tradeoff implies that the optimal

¹ A comparison of various US legislative proposals is maintained by the World Resources Institute.

instrument choice ultimately depends on the underlying policy context and the perceived benefits associated with mitigating different forms of uncertainty (*Weitzman, 1974*).

The political bias toward quantity certainty (and thus away from price certainty) raises an additional political concern that prices (and hence costs) may be volatile or higher than anticipated once legislation is actually enacted. For this reason, the success of future cap-and-trade proposals in the US will likely hinge on whether or not policymakers feel confident in the projected trajectory of carbon prices associated with a given set of targets. This confidence, in turn, depends on the extent to which available estimates of the carbon price path are found to be credible, comprehensible and robust. In fact, the inherent uncertainty in prices has already led some to consider more transparent "cost containment" provisions, some of which would explicitly cap the price of allowances in the future carbon market (*Jacoby and Ellerman, 2004; McKibbin and Wilcoxon, 2004; Pizer, 2002; Roberts and Spence, 1976*). Since the value of any such price trigger would need to be set in advance by the regulator, confidence about the *expected* carbon price is, paradoxically, even more essential for the design of secondary price instruments designed to mitigate uncertainty in prices.²

Despite the obvious need for clear information about prices, the available literature remains difficult to unravel. Technology analyses often provide static estimates of the carbon price required to make a given technology attractive relative to its most obvious competitor or set of competitors. As an example, estimates of the carbon price required to generate investment in carbon capture and storage (CCS)

² To date, the only US legislative proposal to incorporate this mechanism is the so-called Low Carbon Economy Act of 2007 (S. 1766), introduced by US Senators Jeff Bingaman and Arlen Specter. This proposal includes a price ceiling that would begin at \$12 per ton CO₂ in the first compliance year and rise at 5% per year thereafter. Full text of the legislation is available at: <http://www.thomas.gov>.

technology typically hover around \$30-50 per ton CO₂ (e.g. *Deutch et al., 2007*), with estimates of first-plant costs often even higher. At the same time, estimates from computable economic models tend to show a price path in which the initial price is significantly lower than these values but in which the trajectory increases rapidly over time (e.g. *Clarke et al., 2007*). These differences alone are sufficient to generate confusion about the initial range of prices that may be required to achieve specific emissions reduction goals, but such confusion is further exacerbated by the observed variability among estimates from *within* the economic modeling community. While these differences can often be explained qualitatively in terms of differences in underlying technology assumptions (*Peace and Weyant, 2008*), relatively little effort has been devoted to understanding this variability at a more quantitative level.

In order to provide a more mechanistic explanation for the expected price path and for the observed variability among model estimates, I derive, in this paper, a simple theory relating the initial price of allowances under a cap-and-trade system to three fundamental characteristics of the energy-economic system: the assumed (static) marginal cost premium of carbon-free technology, the maximum rate of energy capital replacement and the market interest rate. I then show that this theory captures the essential functional dependence observed in one well-documented computable general equilibrium (CGE) model.

The theory developed here suggests that the expected initial price required to meet a declining emissions path is less than the statically-estimated marginal cost premium of carbon-free technology by an amount that increases with both the market interest rate and the amount of assumed inertia in the energy capital stock. These results not only provide intuition about the observed variability between estimates from economic assessment models but also potentially resolve much of the apparent disagreement between static engineering-based estimates and more dynamic economic

model-based estimates. In addition, this theory offers specific, quantitative guidance to policymakers concerned with the design of secondary price instruments in a cap-and-trade system.

2. Expected Prices in Emissions Permit Markets

In a price-based system, the emissions price is set by the regulator based upon some presumed understanding of the "social cost" of emissions. In other words, in this approach, the price is effectively a Pigouvian tax, intended to internalize the environmental or societal damage associated with a continued increase in emissions (*Baumol, 1972*). In a quantity-based system, on the other hand, knowledge about damages (and about the implied benefits of mitigation) may inform the choice of emissions targets (e.g. *O'Neill and Oppenheimer, 2002*), but the resulting market price of allowances in such a system reflects the marginal cost of compliance rather than the marginal benefit or "social cost" of emitting. More specifically, the price of allowances in a cap-and-trade system is directly related to the cost of achieving particular reduction goals, given explicit assumptions about available technology, but only indirectly (to the extent that damages inform targets and targets determine costs) to any knowledge about larger societal impacts.

Although most analysts agree that the price of CO₂ allowances in a future cap-and-trade system will necessarily reflect the underlying cost of compliance, the same analysts often disagree about the actual range of CO₂ prices that such a system would generate. Static technology assessments that compare the marginal cost of emerging low-carbon technology to the cost of existing technological competitors generally show that the carbon price needed to induce technological substitution would be relatively high, implying that the initial trading price of allowances would also be high, if the emissions caps were stringent enough to require such substitution in the early years of a new program.

On the other hand, assessments with more complex energy-economic models often yield significantly lower estimates of the requisite initial carbon price, even when technology assumptions are comparable. The reason is that economic models tend to view the problem through the lens of exhaustible resource theory. In this case, policies that combine a strategic schedule of emissions limits with sufficient "when-flexibility" (the opportunity for regulated entities to shift abatement across time) are functionally identical to policies in which the cumulative amount of allowable emissions is fixed, effectively making CO₂ an "exhaustible" resource. Well-known economic theory suggests that the price of such a resource would increase at the market interest rate (*Hotelling, 1931*), meaning that the price path of CO₂ should be uniquely defined by two parameters: its initial value and the relevant rate of interest. While all economic analyses generally agree on the range of plausible interest rates (almost all use annual rates between 5 and 8 percent), there is far less agreement about what initial range of prices would be supported by a real cap-and-trade policy. In what follows, I develop a simple theory relating this initial value to a relatively small number of fundamental assumptions about the energy-economic system.

Toward this end, suppose that investors must choose between two different supply technologies ("T₁" and "T₂") to meet existing energy demand. In addition, suppose that the former is inexpensive (with marginal cost C₁) but rather carbon-intensive, while the latter is relatively more expensive (with marginal cost C₂, and C₂ > C₁) but carbon-free. Finally, suppose that the adoption of a climate policy gives rise to a unique carbon price (P) in the economy that rises at the interest rate (r), so that $P(t) = e^{rt}$, with the initial condition that $P(t=0) = P_0$. For simplicity, we may multiply the carbon price (whose natural units are dollars per ton CO₂) by the carbon coefficient (α_1) of T₁ (in units of tons CO₂ per unit energy) so that P is reported in the same units as C₁ and C₂, namely dollars per unit energy. In this case, P is the carbon charge associated with the use of T₁.

If the economy can readily switch between these two technologies, it will generally choose to deploy T_1 when the carbon price is low and T_2 when the carbon price is high. Since the carbon price rises over time, investors will presumably wish to switch from T_1 to T_2 at some particular time (τ) in the future. In fact, if the energy capital stock were sufficiently malleable, then investors would instantaneously switch from T_1 to T_2 at the exact moment when the cost of deploying the first technology, taking into account the implied carbon charge, exceeded the cost of deploying the second technology. Mathematically, this condition can be expressed as:

$$C_1 + P_0 e^{r\tau} > C_2 \quad (1)$$

Given explicit assumptions about C_1 , C_2 , P_0 and r , one could easily solve for τ .

Of course, in the real world, the capital stock cannot be replaced instantaneously, meaning that an initial decision to substitute T_2 for T_1 would not immediately transform the energy system, but would merely begin the slow replacement process. In this case, the rate of turnover would depend on the assumed time constant of the capital stock (ϵ), typically several decades, but perhaps 50 years or more for large-scale infrastructure associated with electric power generation.

In the context of a forward-looking economy in which investors make decisions based on expectations about the future, the existence of inertia in the capital stock has profound implications. Most importantly for our purposes, it means that the process of technological substitution may begin before the condition expressed in (1) is actually satisfied. In effect, forward-looking investors will be willing to pay higher energy costs in the present (by switching from T_1 to T_2 "prematurely") in order to avoid being "stuck" with T_1 in later years, when the exponentially increasing carbon charge is very high. In fact, an optimal transition strategy would exactly balance the cost of premature switching with the benefit of avoided carbon charges in the future, once the appropriate

weights have been applied to account for standard time preferencing. To distinguish this strategy from the one pursued in the absence of inertia, we may label the time at which substitution begins τ^* , where $\tau^* < \tau$.

A solution for the optimal transition path (i.e. the value of τ^*) can be found by considering the above investment choice in the framework of a cost minimization problem. In this simplified two-technology world in which investors are forward-looking and the capital stock slow to overturn, we may assume that investors wish to minimize the net present value of the total (cumulative) cost of energy production. This total cost can be expressed mathematically as follows:

$$T = \int_0^{\tau^*} E(t) \cdot (C_1 + P_0 e^{rt}) \cdot e^{-rt} dt + \int_{\tau^*}^{\infty} E(t) \cdot (C_1 + P_0 e^{rt}) \cdot e^{-\varepsilon(t-\tau^*)} e^{-rt} dt + \int_{\tau^*}^{\infty} E(t) \cdot C_2 \cdot (1 - e^{-\varepsilon(t-\tau^*)}) \cdot e^{-rt} dt \quad (2)$$

In the window between the start of the policy and the time at which technology substitution begins ($0 < t < \tau^*$), the marginal cost of energy in any period, taking into account the instantaneous carbon charge, is $C_1 + P_0 e^{rt}$. When this adjusted marginal cost is multiplied by total energy production and then integrated over the relevant window, the total cost, expressed in net present value terms, is given by the first integral on the right-hand side of (2).

After substitution begins ($t > \tau^*$), the marginal cost of energy in any period is effectively the average of $C_1 + P_0 e^{rt}$ and C_2 , weighted by the share that each technology contributes to total energy production. We may assume that the T_1 share declines at the rate ε , while T_2 expands to make up the difference. Thus the total integrated cost, in net present value terms, for the period when ($t > \tau^*$) is given by the sum of the second and third integrals on the right-hand side of (2).

By expanding the integrals in (2) and then solving for the value of τ^* that minimizes total cost (a complete derivation is provided in the appendix), one finds that:

$$\frac{P_0}{(C_2 - C_1)} = \frac{\varepsilon}{r + \varepsilon} \cdot e^{-r\tau^*} \quad (3)$$

The term on the left-hand side of (3) is the ratio of the initial carbon charge to the marginal cost premium of carbon-free technology. In the expression above, P_0 , C_1 and C_2 are all reported in units of dollars per unit energy. Since the carbon price is usually reported in dollars per ton CO_2 , a more useful expression can be found by dividing both the numerator and the denominator on the left-hand side of (3) by the carbon coefficient α_1 . In that case, P_0 is once again the carbon price reported in the familiar units of dollars per ton CO_2 , and the denominator is the effective technology "crossover price" in the same units.³ In the remainder of this paper, we report P_0 , C_1 and C_2 in these "per ton" units to facilitate interpretation in the policy context.

With these modifications, the expression in (3) shows that the ratio of the initial carbon price to the statically-estimated technology crossover price is always less than 1. In other words, this result suggests that, in a forward-looking competitive economy in which capital replacement is gradual, static price estimates tend to overestimate the initial carbon price required to motivate technology substitution. The magnitude of this overestimation increases with the interest rate (r) and decreases with the capital turnover rate (ε).

³ For example, suppose that the marginal cost difference between the two technologies ($C_2 - C_1$) is \$3 per GJ. If the carbon coefficient (α_1) of T_1 is 0.06 tons CO_2 per GJ, while the carbon coefficient (α_2) of T_2 is 0, then the carbon price at which T_2 would become competitive with T_1 in a static world would be $(\$3/\text{GJ})/(0.06 \text{ tons}/\text{GJ}) = \50 per ton CO_2 . This is the static technology crossover price.

The ratio expressed in (3) also increases with the stringency of mitigation, as one would expect. The condition $\tau^* = 0$ reflects the most stringent mitigation because it implies that technology substitution begins immediately. The condition $\tau^* > 0$ reflects less stringent mitigation because it implies that the start of substitution is postponed by some number of years. In the immediate mitigation limit in which $\tau^* = 0$, the expression in (3) reduces to:

$$\frac{P_0}{(C_2 - C_1)} = \frac{\varepsilon}{r + \varepsilon} \quad (4)$$

This expression suggests that two additional limits are worth exploring. First, when the interest rate is small so that $r \ll \varepsilon$, the ratio of the initial carbon price to the crossover price approaches 1. This is the static limit, in which the carbon price is assumed to remain constant over time. In this case, there is no tendency to anticipate higher prices, meaning that the initial price must equal the crossover price associated with the $T_1 \rightarrow T_2$ transition in order for substitution to be desirable. Secondly, in the limit where the rate of capital turnover is assumed to be high, so that $\varepsilon \gg r$, the ratio also approaches 1. This limit is mathematically equivalent to the one above, but the conceptual explanation here is different. This limit results from the fact that, if the capital stock were perfectly malleable, so that it could be effectively replaced instantaneously, then there would again be no gains to anticipation.

When the ratio in (4) is significantly less than one, then the amount by which the initial carbon price is suppressed relative to the technology crossover price depends on the exact values of r and ε (see **Figure 1**). In all cases, a higher interest rate implies a lower initial carbon price, because the steep slope of the price path yields greater gains to anticipation in the form of reduced future carbon charges. On the other hand, a higher rate of capital turnover implies a higher initial carbon price for the reasons mentioned above, namely that the ability to quickly turnover the capital stock lowers the gains to anticipation.

3. Comparison between Theoretical and Computational Results

Any simple theory like the one above will necessarily understate the complexity of real market behavior. One way to evaluate the usefulness of such a theory is to explicitly compare theoretical predictions to numerical output from a computable energy-economic model. While any such model will also understate the complexity of real markets, good agreement between theoretical and computational results would demonstrate that the theory is sufficiently sophisticated to capture essential model behavior and to provide guidance to those interpreting such models and developing policy recommendations based upon them.

With this motivation in mind, we compare results from the theory described in Section 2 to output from a well-known optimal growth model, the so-called MERGE model (*Manne et al., 1995*). To facilitate comparison, we run the model in a simplified configuration in which we focus on a single region (the United States) and reduce the number of available energy technologies to four: (a) coal in the power sector; (b) a generic carbon-free technology in the power sector (which could be viewed as some combination of nuclear, CCS and/or other renewables); (c) oil in the fuels sector; and (d) a generic technology in the fuels sector. For simplicity, the values for the marginal costs of both carbon-free technologies are chosen arbitrarily so that the static crossover price between carbon-intensive and carbon-free energy is \$50 per ton CO₂ in both sectors.

Under business-as-usual (i.e. when no policy constraints are applied), economic growth is assumed to proceed at the rate of 3% per year, while assumed improvements in energy intensity hold the growth in underlying energy demand to about 1% per year. Without explicit constraints on carbon, coal and oil continue to dominate the power and fuels sectors, respectively. When policy constraints are applied, the new optimal growth path is deflected from business-as-usual. As above, a sufficiently large price on carbon,

whether applied directly as a tax or indirectly through quantity limits, generates a preference for carbon-free energy over carbon-intensive energy in both sectors, leading to gradual penetration of these technologies into the energy system.

In what follows, we evaluate output from 81 individual simulations with the modified MERGE model, representing combinations of nine different choices of the interest rate and nine different choices of the capital decline rate (with values of r and ϵ both ranging from 2% to 10%). In each of these 150-year simulations, we find the value of the initial carbon price (P_0) that is just large enough to induce substitution from carbon-intensive to carbon-free energy in the first model period, given particular values for the interest rate (r) and the capital decline rate (ϵ), and assuming that the carbon price rises at the interest rate for the duration of the simulation.

For each choice of r and ϵ , we then compare the theoretically-derived ratio, $P_0/(C_2-C_1)$ using equation (4) to the value obtained using the model methodology described above. Panel (A) of **Figure 1** shows this ratio (in percent) as a function of r and ϵ , using the model, while panel (B) shows the same ratio, in the same parameter space, as given by equation (4). In these two panels, a value of 50% implies that the initial carbon price need only be half the statically-estimated crossover price to induce technological substitution. Assuming a crossover price of \$50 per ton CO₂ as above, this implies that the initial carbon price would be \$25 per ton CO₂. Alternatively, given a realistic interest rate of 6% and an upper limit on the annual capital turnover of 3%, the ratio is closer to 33%, meaning that, using the same technology assumptions, the initial carbon price would be only \$17 per ton CO₂, .

Quick inspection of panels (A) and (B) shows that the model results agree well with theoretical predictions. Indeed, the differences between them, shown explicitly in panel (C) (in percent) are less than a few percentage points for most combinations of r and ϵ , with only 4 of the 81 simulations yielding differences greater than 5 percent.

Such differences can be explained by recognizing that the time horizon in the numerical model is artificially truncated. In the real world, planning horizons are obviously finite, so these results suggest that the planning horizon itself may be an interesting sensitivity to explore in future work.

4. Conclusions

The theory developed in this paper makes only two assumptions about the energy system. First, it assumes that capital replacement is a gradual process, and secondly, it assumes that investors are fundamentally forward-looking. When both of these assumptions hold, then the initial carbon price required to induce technological substitution is suppressed relative to the statically-estimated technological crossover price. Our results demonstrate that the actual amount by which the carbon price is suppressed increases with the market interest rate and decreases with the capital replacement rate.

If either of these two assumptions is discarded, then the initial carbon price converges to the technology crossover price. As such, these results show how common outputs generated by two distinct energy analysis communities (engineering and economics) are related to one another. They also make explicit the frequent qualitative observation that model-projected differences in the initial price of CO₂ depend on differences in underlying technology cost assumptions. This theory shows that the relationship between the projected market price and the underlying marginal cost premium is linear, with a constant of proportionality always less than 1.

In the policy context, this theory shows how the choice of a CO₂ price trigger in a cap-and-trade system could be informed by rather basic assumptions about the energy system, mitigating the criticism that any such mechanism would necessarily be arbitrary. More specifically, if the purpose of such a mechanism is to hedge against the

risk that prices will be higher than anticipated, then the trigger should be set at a value that reflects current expectations about the future. The theory provided here offers a way to calculate this value based on available estimates of future technology costs.

The results of this paper may therefore prove to be useful in two different settings. Within the energy analysis community, the theory provides greater mechanistic insight into the range of observed price estimates and shows how static and dynamic estimates are related to one another, potentially resolving some confusion about prices within the research community. In the climate policy context, this theory shows how the initial price of CO₂ can be tied to "measurable" quantities in the energy system, thus providing concrete guidance to those concerned with the design of secondary price instruments in a cap-and-trade system.

Mathematical Appendix

Expanding the integrals in (2), we find:

$$\begin{aligned} \Gamma = & \int_0^{\tau^*} E(t) \cdot C_1 \cdot e^{-rt} dt + \int_0^{\tau^*} E(t) \cdot P_0 dt + \int_{\tau^*}^{\infty} E(t) \cdot C_1 \cdot e^{-(r+\varepsilon)t+\varepsilon\tau^*} dt + \\ & \int_{\tau^*}^{\infty} E(t) \cdot P_0 \cdot e^{-\varepsilon(t-\tau^*)} dt + \int_{\tau^*}^{\infty} E(t) \cdot C_2 e^{-rt} dt - \int_{\tau^*}^{\infty} E(t) \cdot C_2 e^{-(r+\varepsilon)t+\varepsilon\tau^*} dt \end{aligned} \quad (\text{A1})$$

Making the simplifying assumption that the overall demand for energy remains constant over time (i.e. that $E(t) = E$), we find upon integrating that:

$$\begin{aligned} \Gamma = & \left[\frac{-E \cdot C_1}{r} e^{-rt} \right]_0^{\tau^*} + [E \cdot P_0 \cdot t]_0^{\tau^*} + \left[\frac{-E \cdot C_1}{r + \varepsilon} e^{-(r+\varepsilon)t+\varepsilon\tau^*} \right]_{\tau^*}^{\infty} \\ & \left[\frac{-E \cdot P_0}{\varepsilon} e^{-\varepsilon(t-\tau^*)} \right]_{\tau^*}^{\infty} + \left[\frac{-E \cdot C_2}{r} e^{-rt} \right]_{\tau^*}^{\infty} - \left[\frac{-E \cdot C_2}{r + \varepsilon} e^{-(r+\varepsilon)t+\varepsilon\tau^*} \right]_{\tau^*}^{\infty} \end{aligned} \quad (\text{A2})$$

Plugging in boundary values, we get:

$$\begin{aligned}
T = & -\frac{E \cdot C_1}{r} e^{-r\tau^*} + \frac{E \cdot C_1}{r} + E \cdot P_0 \cdot \tau^* + \\
& \frac{E \cdot C_1}{r + \varepsilon} e^{-r\tau^*} + \frac{E \cdot P_0}{\varepsilon} + \frac{E \cdot C_2}{r} e^{-r\tau^*} - \frac{E \cdot C_2}{r + \varepsilon} e^{-r\tau^*}
\end{aligned} \tag{A3}$$

Finally, upon collecting terms, we find:

$$T = \frac{\varepsilon \cdot E \cdot (C_2 - C_1)}{r \cdot (r + \varepsilon)} e^{-r\tau^*} + \frac{E \cdot C_1}{r} + E \cdot P_0 \cdot \tau^* + \frac{E \cdot P_0}{\varepsilon} \tag{A4}$$

The optimal choice of τ^* can be found by differentiating (A4) with respect to τ^* and setting the result equal to 0. This yields the following expression:

$$\frac{\partial T}{\partial \tau^*} = \frac{-\varepsilon}{r + \varepsilon} \cdot E \cdot (C_2 - C_1) \cdot e^{-r\tau^*} + E \cdot P_0 = 0 \tag{A5}$$

Dividing through by E and re-arranging terms, we get:

$$\frac{P_0}{(C_2 - C_1)} = \frac{\varepsilon}{r + \varepsilon} \cdot e^{-r\tau^*} \tag{A6}$$

Equation (A6) is restated as equation (3) in the main text.

Figure 1


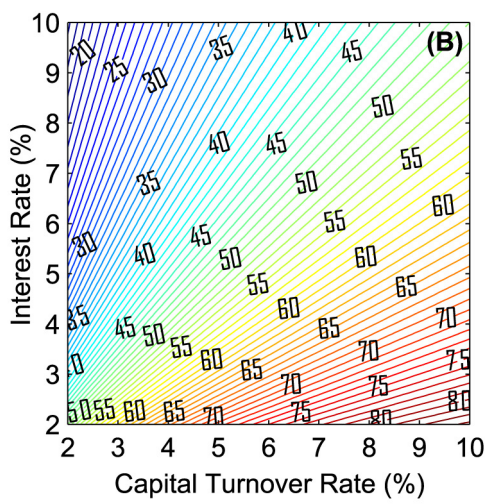
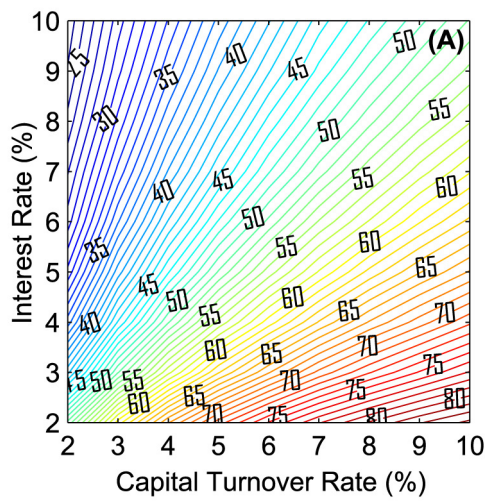


Figure 1 Caption

Panels (A) shows the ratio of the initial carbon price to the technological crossover price (in percent) as a function of the interest rate and the assumed rate of capital decline, as predicted by the model described in Section 3. Panel (B) shows the same ratio, in the same parameter space, as predicted from the theory described in Section 2. Panel (C) shows the simple (absolute) difference between panels (A) and (B), also in percent.

References

- Baumol, W J. On taxation and the control of externalities. *American Economic Review* 1972; 62; 307-322.
- Caldeira K, Jain A K, Hoffert M I. Climate sensitivity and the need for energy without CO₂ emission. *Science* 2003; 299; 2052-2054.
- Carlson C, Burtraw D, Cropper M, Palmer K L. Sulfur dioxide control by electric utilities: what are the gains from trade? *Journal of Political Economy* 2000; 108; 1292-1326.
- Clarke L E et al. Scenarios of Greenhouse Gas Emissions and Atmospheric Concentrations. United States Climate Change Science Program Synthesis and Assessment Product 2.1A, 2007. Available at: <http://www.climatechange.gov>.
- Deutch J et al. The Future of Coal. Massachusetts Institute of Technology, 2007. Available at: <http://web.mit.edu/coal>.
- Hotelling H. The economics of exhaustible resources. *Journal of Political Economy* 1931; 39; 137-175.
- Jacoby H D, Ellerman A D. The safety valve and climate policy. *Energy Policy* 2004; 32; 481-491.

- Manne A, Mendelsohn R, Richels R. A model for evaluating regional and global effects of GHG reduction policies. *Energy Policy* 1995; 23; 17-34.
- McKibbin W J, Wilcoxon, P J. Estimates of the costs of Kyoto: Marrakesh versus the McKibbin-Wilcoxon blueprint. *Energy Policy* 2004; 32; 467-479.
- Mignone B K, Socolow R H, Sarmiento J L, Oppenheimer M. Atmospheric stabilization and the timing of carbon mitigation. *Climatic Change* 2008; 88; 251-265.
- Newell R, Pizer W, Zhang J F. Managing permit markets to stabilize prices. *Environmental and Resource Economics* 2005; 31; 133-157.
- O'Neill B C, Oppenheimer M. Dangerous climate impacts and the Kyoto Protocol. *Science* 2002; 296; 1971-1972.
- O'Neill B C, Oppenheimer M, Peterson A. Interim targets and the climate treaty regime. *Climate Policy* 2006; 5; 639-645.
- Pacala S, Socolow R. Stabilization wedges: solving the climate problem for the next 50 years with current technologies. *Science* 2004; 305; 968-972.
- Peace J, Weyant J. Insights not numbers: the appropriate use of economic models, Pew Center on Global Climate Change White Paper, 2008. Available at: <http://www.pewclimate.org/white-paper/economic-models-are-insights-not-numbers>.
- Pizer W. Combining price and quantity certainty controls to mitigate global climate change. *Journal of Public Economics* 2002; 85; 409-434.
- Roberts, M J, Spence M. Effluent charges and licenses under uncertainty. *Journal of Public Economics* 1976; 5; 193-208.

Schneider S H, Mastrandrea M D. Probabilistic assessment of "dangerous" climate change and emissions pathways. *Proceedings of the National Academy of Sciences of the United States of America* 2005; 102; 15728-15735.

van Vuuren D P, Weyant J, de la Chesnaye F. Multi-gas scenarios to stabilize radiative forcing. *Energy Economics* 2006; 28; 102-120.

Weitzman, M L. Prices versus quantities. *Review of Economic Studies* 1974; 41; 477-491.