State-Contingent Pricing as a Response to Uncertainty in Climate Policy*

Ross McKitrick
Department of Economics
University of Guelph
rmckitri@uoguelph.ca

July 2011

Abstract

Uncertainties over the future path of global warming and the underlying severity of the problem make derivation of an intertemporally-optimal emissions price on carbon dioxide both theoretically and politically very difficult. A number of methods for dealing with uncertainty have dominated the economics literature to date. These involve trying to derive an emissions price or insurance premium to which agents are expected to make a long term commitment. This chapter explores an alternative approach based the concept of state-contingent pricing, in which agents commit to a pricing *rule* rather than a *path*. The rule connects current values of the emissions price to observed temperatures at each point in time. In essence, if the climate warms, the tax goes up, and vice versa. A derivation is provided showing how such a rule yields an approximation to the unknown optimal dynamic externality tax, yet can be computed using currently-observable data. A recently-proposed extension coupling the state-contingent tax with a tradable futures market in emission allowances would yield not only a feasible mechanism for guiding long term investment, but an objective prediction market for climate change. The advantage of the state-contingent approach for facilitating coalition-formation is also discussed, as are directions for research.

* This is a chapter prepared for *Handbook on Energy and Climate Change*, Roger Fouquet, ed., Cheltenham: Edward Elgar.

State-Contingent Pricing as a Response to Uncertainty in Climate Policy

1. Introduction

Suppose we have a time machine that allows us to visit the year 2040 just long enough to collect some climate data. Figure 1 shows the post-1979 globally-averaged lower tropospheric air temperature anomaly averaged over the two satellite series developed by, respectively, Spencer and Christy (1990) and Mears and Wentz (2005). This is only one of many data series people use to try and represent the global climate as a univariate time series, but it will do for the current illustration. Figure 1 shows the observed data from 1979 up to the end of 2010 (shown by the vertical line), and then runs the series forward using assumed trends and random numbers to conjecture two quite different futures. In the gray dots the next three decades exhibit continued variability but no upward trend, and even a slight downward trend. The black dots show variability and a strong upward trend. Now suppose that, given an identical future greenhouse gas emissions trajectory the data we collect in 2040 will look like one of those two paths. If we could find out which one would be observed, would it affect today's policy choices?

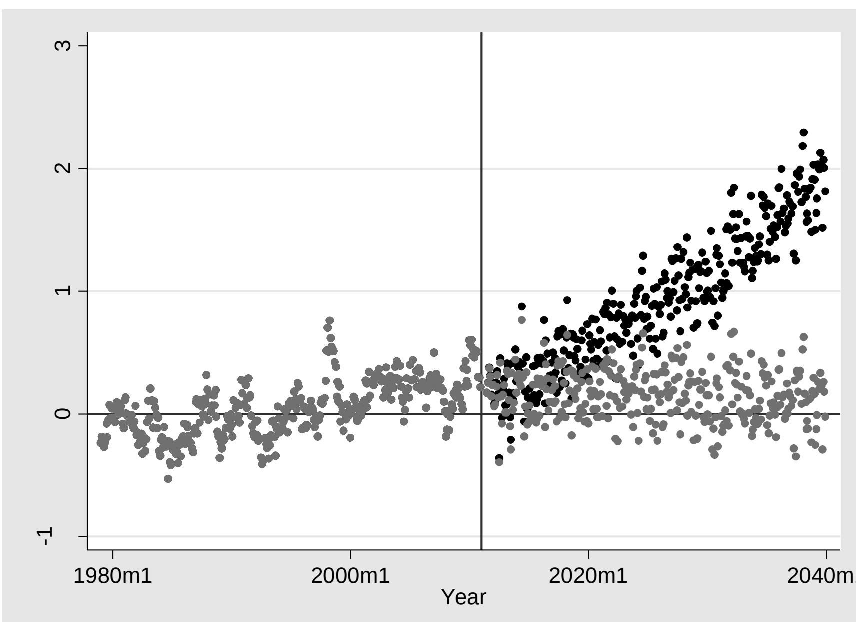

Figure 1. Two conjectured atmospheric warming paths 1979-2040. Left of the vertical line are observations from weather satellites of the global average lower troposphere temperature anomaly. Right are conjectured using trends and random numbers.

Obviously the answer is yes. The fact that we do not know what the graph will look like has led to longstanding and well-known political difficulties in devising policy strategies. In this chapter I will critically review the main current approaches to dealing with the uncertainty, and then propose an alternative that I believe is more likely to lead to the right policy outcome than any others currently being examined. Briefly stated, my argument is as follows.

- Forecast-based proposals (such as from Integrated Assessment Models, or IAMs) for making optimal climate policy decisions effectively assume we can agree what the data to the right of the line will probably look like, and we only need to resolve the time-sequence of emission pricing. While the optimal time-sequence is a significant puzzle to be solved, framing the issue in this way assumes away all the real uncertainty that makes the problem difficult in the first place. If we make a commitment to a long-term policy based on IAM analysis, when we get to, say, 2040, there is a high probability we will realize that we followed the wrong emission pricing path.
- Bayesian updating and other learning strategies involve placing bets on the unknown future then observing the effects of the policy decisions and revising our strategy when we have learned enough to figure out if the bet was right or wrong. The main lesson of these approaches is that in the climate case, this kind of learning will be too slow to be of any use in guiding policy now or in the foreseeable future. Consequently, when we get to 2040, we will likely not know if we were on the correct path or not.
- Each of the futures in Figure 1 implies a corresponding optimal emissions price path, which, for instance, might look something like those in Figure 2. If we knew the future temperatures with certainty, we would, in principle, be able to work out the optimal emission tax path.
- The state-contingent approach involves starting an emissions tax at the current best guess as to its optimal level, then specifying a rule that updates it each year based on the observed climate state. As of the present we do not know what the path will look like, but if we choose the rule correctly, *we can know today that as of 2040 we will have followed the closest possible approximation to the optimal price path*. Furthermore, the greatest economic gains will accrue to agents that make the most accurate forecasts about the climate state, and hence, the emissions price.
- Under a state-contingent pricing rule, the need for accurate forecasts of the future tax for the purpose of guiding investment decisions will create market incentives to make the maximum use of available information and the most objective climate forecasts to guide optimal behavioural responses to the policy. Consequently, while the state-contingent price path will only track the actual optimum within error bounds, there is no information currently available that could identify a better price path than the one generated by information markets induced by a state-contingent pricing rule.

The rest of this chapter explains these ideas in more detail.

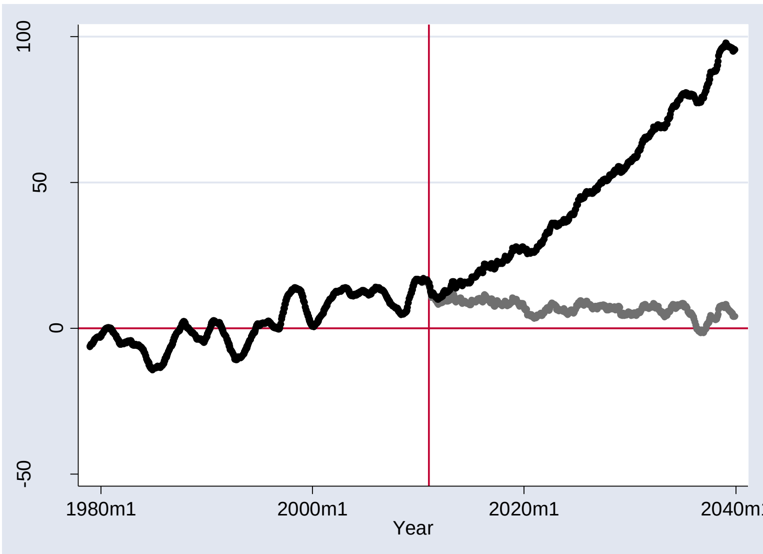

Figure 2. Possible optimal emission price paths corresponding to future warming scenarios in Figure 1.

2. Uncertainty and inertia problems in carbon dioxide emission pricing

Sources of uncertainty

In one respect, analysis of carbon dioxide (CO_2) emission pricing is simple compared to other air emissions such as sulphur dioxide (SO_2) or particulate matter (PM). Since there is no CO_2 scrubber technology, knowing the amount of fuel consumed yields a close estimate of the total CO_2 released, whereas fuel consumption can be quite uncorrelated with other emissions depending on the pollution controls and combustion technology in place. For that reason, CO_2 emissions are easily represented in empirical and computational economic models, as long as the consumption levels of coal, oil and natural gas are resolved. However, the time element that connects CO_2 to its external costs is considerably more complex than for other emissions. SO_2 and PM do not stay aloft very long after release (days or weeks), and investment in a scrubber today will yield potentially large emission reductions within a year, so from a planning point of view the time path of control policies for these pollutants can be considered as a sequence of short-run decisions.

In the CO_2 case, however, time complicates the planning problem in several ways.

a) The atmospheric residency of CO_2 is measured in decades, so emissions today could potentially have effects many years into the future, and each year's emissions have marginal effects that accumulate with those of other years;
b) The response of the climate system to changes in the atmospheric stock of CO_2 may be slow, especially if the ocean acts as a flywheel, delaying effects for decades or centuries;

c) Since there are no scrubbers, emission reductions must take the form either of changes in combustion efficiency, fuel-switching or reductions in the scale of output, all of which take time to plan and implement;
d) Economically-viable technology for generating electricity and converting fossil fuels into energy is subject to innovation over time, and while it is reasonable to assume that some innovations and efficiency gains will be realized, the effects and timing of changes can only be conjectured;
e) Since the stock of CO_2 mixes globally, actions of individual emitters are negligible. The only policies that would affect the atmospheric stock must be coordinated among all major emitting nations, and such processes are slow and subject to uncertain success.

Parson and Karwat (2011) examine these issues under the headings of *inertia* and *uncertainty*. If we faced only inertia, we could sit down today and devise an optimal intertemporal policy plan today that would yield the right sequence of interventions at the right time through the future. This is something like what IAMs do: assuming that we know the important parameters of the system, we can solve for the optimal intertemporal emission pricing path. On the other hand, if we faced only uncertainty, we could make short-term decisions on the expectation that new decisions would be made at each point in the future as circumstances change. It might be argued that this is more like what climate policy has been in practice for the past 20 years: a series of short-term decisions that resolve momentary political pressures, but which do not seem rooted in an overall intertemporal plan. Faced with both uncertainty and inertia, Parson and Karwat conclude that sequential decision-making is necessary, though they do not spell out how such a process would work in practice. The state-contingent approach, it will be shown, attempts to create a formal structure for sequential decision-making in light of both uncertainty and inertia.

With regard to climate change there are two very large sources of anxiety that have fueled decades of intense controversy. On one side are those who believe the threat from CO_2 and other greenhouse gas emissions are substantial, and who fear that inadequate policy actions are being taken, so that future generations will experience serious welfare losses due to global warming. On the other side are those who believe the threat from CO_2 and other greenhouse gas emissions are small, and who fear that implementation of policies sufficiently stringent to achieve large emission reductions will impose costs on current and future generations far larger than any benefits they yield. For the first group, the fear is that by the time enough information is obtained to resolve uncertainty about the environmental effects of CO_2 it will be too late to avert intolerable environmental damages. For the second group, the fear is that if we act now to try and prevent such damages we will have incurred intolerable economic costs by the time they are shown to have been unnecessary.

So-called “no regrets” policies are sometimes invoked to try and make this wrenching dilemma disappear, but they are irrelevant to the discussion. There is a strain of argument that says, in light of the threat of catastrophic (or even somewhat harmful) global warming, we must act, and the actions we propose would actually make us better off by saving energy and reducing air pollution anyway, so on balance it is better to implement them. This argument fails once the details are examined. The scale of emission reductions necessary to substantially affect the future stock of global atmospheric CO_2 is quite large, namely worldwide reductions of some 50% or more, and marginal local changes in energy efficiency would not begin to be sufficient. Improvements in energy efficiency that actually make consumers and firms better off are automatically adopted by rational economic decision-makers anyway, yet CO_2 emissions continue to rise globally as population and income rise. And air pollution is already subject to regulation throughout the developed (and much of the underdeveloped) world. If we assume that households, firms and governments have already made reasonably efficient decisions as to energy efficiency and pollution reduction, further large-scale reductions in CO_2 emissions must be, on net, costly. In other words, policies that might have a trivial cost will only have trivial climatic effects. The policies that actually have an effect on the climate must entail a large economic cost. The dilemma is real.

Integrated assessment models and pseudo-optimal solutions

The IAM approach of Nordhaus (2007) and coauthors yields a solution that can be described as "pseudo-optimal." It assumes the modeler knows the key parameters that govern the economy and the climate, and solution of the model yields a smooth policy "ramp" in the form of an escalating tax on CO_2 emissions over time. This solution can only be considered optimal if we assume the model parameters are correct. But strong assumptions about key functional forms and parameter values are not put to the test by implementation of the policy. If decision-makers were to commit to a policy path based on the IAM analysis, it would amount to acknowledging the inertia but not the uncertainty in the policy problem. The lack of recognition the extent of uncertainty in the IAM approach is one of the bases of the criticism of Weitzman (2009).

Bayesian learning models

Kelly and Kolstand (1999) and Leach (2007) introduced learning into the IAM framework by supposing that we observe the response of the climate to policy innovations, and then we use such information in a Bayesian updating routine. The goal is to accumulate enough information that the policymaker can decide, at 5 percent statistical significance, whether or not to reject the hypothesis that the correct policy is being implemented. Uncertainty and inertia interact in an interesting way: uncertainty about even one or two key inertia (lag) parameters is sufficient to delay for hundreds of years the identification of an expected-optimal policy rule. With only two model parameters subject to uncertainty, Leach (2007) showed the learning time ranges from several hundred to several thousand years, depending on the base case emissions growth rate. An expanded version of the model, incorporating simple production and an intertemporal capital investment structure, not only yields a time-to-learn measured in centuries, even when most model parameters are assumed known, but depending on which of several climate data sets are used to form the priors, the policy path may never converge on the correct target.

It is an illusion to suppose that the IAM, or pseudo-optimal approach, is better, because we apparently follow an optimal path from the outset. The difference between them is that in the Bayesian approach we eventually learn if we are on the wrong path and in the IAM approach we never do.

Insurance and fat tails

Weitzman (2009) looked at the global warming problem as one of trying to price an insurance contract when there is a nontrivial probability of extreme damages. Geweke (2001) had shown that a basic insurance problem can become degenerate if a few features of the set-up are chosen in a particular way. If the risk is distributed normally and utility is of the constant relative risk aversion form, and the change in consumption over the insured interval is expressed as e^C where C is future consumption relative to current consumption, then the expected cost of insuring future consumption under some general conditions can be shown to take the form of a moment generating function for the *t* distribution, which does not exist, or is infinitely large, making it impossible to place a finite value on a full insurance contract. Weitzman's adaptation of this model to the climate case depends on some specific assumptions, some of which are conventional and some of which are not. One unusual assumption is that there is a possibility of infinite (+ or -) climate sensitivity, or in other words, that while the possibility of an extreme change in the climate (twenty degrees or more) may be small, it cannot be ruled out, no matter how large. To perform conventional cost-benefit analysis it is necessary either to truncate the range of climate sensitivities or assume that the distribution has "thin tails." But, as Weitzman points out, this implies that the optimal insurance policy depends on assumptions about the distribution of possible climatic changes in regions where there are too few observations to know for sure. Hence cost-benefit analysis using IAMs assumes away extreme risks, and cannot therefore provide an economic case for ignoring them. Nordhaus (2009), Pindyck (2011) and others have critiqued the Weitzman model, especially for its assumption of infinite marginal utility as consumption gets very low.

The state-contingent approach

McKitrick (2010) proposed an alternative approach to the pricing of complex intertemporal externalities which focuses on developing an adaptive pricing rule, rather than a long term emissions path. In the standard economic model of pollution pricing, current damages are a direct function of current emissions:

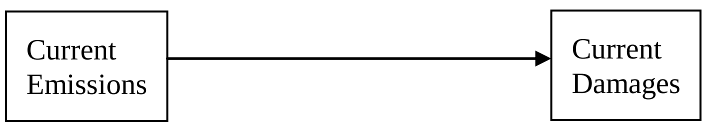

In simple problems of this sort, the solution is to impose an emissions price

$$\tau = D'(e_t) \qquad (1)$$

where *D* is the damage function and e_t is total current emissions, which are assumed to be observable. In the presence of inertia, emissions may have lagged effects, and the length of the delay may itself be unknown. If so, then we are currently experiencing the effects not only of present-day emissions, but also of emissions that have occurred at some point in the past. A lag process may arise from stock effects in which the decay rate is less than 100% each period. But here we are interested in the case in which, instead of directly causing damages, emissions affect some aspect of the environment (such as the average air temperature, possibly with lags due to geophysical processes), and those changes cause damages. In that case the above diagram would be redrawn as follows.

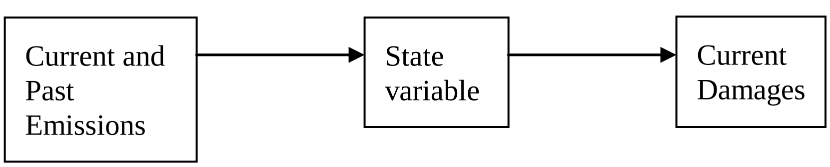

Emissions affect an observable state variable $s(t)$ and current-period damages are a function of the current state variable $D(s(t))$. The state variable, in turn, is determined as

$$s(t) = s(e_t, e_{t-1}, ..., e_{t-k}) \quad (2)$$

where the lag length *k* is unknown. Equation (2) is sufficiently general to handle stock-flow problems. If the stock of emissions is denoted M_t, then the state may be described as the function $s(M_t)$, where $M_t = \delta M_{t-1} + e_t$ and δ is the stock decay rate. But since the emissions flow is the policy target it will be necessary to include e_t in the specification of (2). A form like $s(M_t)$ would not lead to a policy-relevant conclusion since the policy maker cannot control M_t (or past values of e_t).

Note also that $s(t)$ is not, in most cases, the same as the stock of the pollution (i.e. the atmospheric concentration). It is therefore necessary to identify the correct state variable. Hsu (2011) discusses the state-contingent pricing option for carbon dioxide emissions and proposes a list of possible candidates for the state variable. The list will be critically evaluated below, but for the moment our point is simply that for a complex issue like global warming it can be difficult to agree on what to measure, and not all proposed values of the state variable make sense for the purpose of configuring an optimal policy mechanism.

Setting that aside for now, given a definition of *s* the Discounted Present Value (DPV) of damages is

$$V(t) = \sum_{j=0}^{\infty} \beta^j D(s(t+j)) \quad (3)$$

where β^j is the discount factor *j* periods ahead. The optimal emissions price is:

$$\tau(t) = \frac{\partial V(t)}{\partial e_t} \quad (4).$$

The influence of current and past emissions on the state variable is complex and uncertain. While this adds to the difficulty of determining how current emissions ought to be priced, it also implies that the state variable contains information about the effect of emissions over time, and this information can be used to reduce uncertainty. The next section explains how to use observations on *s* to approximate $\tau(t)$.

3. Derivation of the State-Contingent Pricing Rule

Required Assumptions

Following McKitrick (2010) the derivation requires a number of assumptions.

(A1) ASSUMPTION 1: The state function *s* is homogeneous of degree *c*.

This implies

$$s(\lambda e_t,...,\lambda e_{t-k}) = \lambda^c s(e_t,...,e_{t-k}) \quad (5).$$

We do not necessarily assume linear homogeneity.

(A2) ASSUMPTION 2: Over the interval $(t-k,...,t+k)$, ∇s is locally autonomous, that is, $\frac{\partial s(t+i)}{\partial e_t} = \frac{\partial s(t)}{\partial e_{t-i}}$ for all $i = 0,...,k$.

This imposes slightly more structure on *s*, as it implies the marginal effects over a lag of length *i* are independent of *t*. For example, suppose *s* depends on emissions out to three lags:

$$s(t) = s(e_t, e_{t-1}, e_{t-2}, e_{t-3}).$$

At time *t*, the partial derivative with respect to the second argument is $\frac{\partial s(t)}{\partial e_{t-1}}$. At time *t*+1, the partial derivative with respect to the second argument is $\frac{\partial s(t+1)}{\partial e_t}$. (A2) requires that these partial derivatives be equal. This will be true, for instance, if *s* is a function of a weighted sum or moving average of the *e*'s. It will not be true if *s* is nonlinear in the individual *e*'s, in which case it will be true only approximately, where the approximation will depend on how much "curl" *s* has over time.

(A3) ASSUMPTION 3: The function $s(t)$ is locally linear in e_t, and current period emissions must have a non-zero effect on $s(t)$.

This assumption states that in the direction e_t, $\frac{\partial s}{\partial e_t} = \nu$ in the neighbourhood of e_t, where ν is a positive constant and may be arbitrarily small.

(A4) ASSUMPTION 4: At each time t, the damage function $D(t)$ can be approximated by a step-wise quadratic (in emissions) function, such that at time t, $\frac{\partial D(t+j)}{\partial e_{t+j}} \approx \delta e_t$ for $j = 0,\ldots,k$.

This is different from, and slightly more restrictive than, assuming D is quadratic. It states that the slope can be extrapolated forward at time t over k-1 subsequent periods, and that the extrapolation will be re-set each period. (A4) can be pictured as an approximation for D by a porcupine-like pattern of tangent lines, as shown in Figure 3. The approximation will get worse the larger is k. However, in cases where k is large, effects mix slowly across long time spans and we would not expect total damages to be strongly nonlinear (convex) in current emissions, so the extrapolation is automatically used across a shorter interval in circumstances where it is less accurate over long intervals. In other words, the larger is k (implying greater potential inaccuracy in the slope extrapolation), the less is the likely curvature in D (implying less inaccuracy in the slope extrapolation). If D has a strong curvature then current emissions must have a rather strong immediate effect, so we expect the tangent lines in Figure 3 to be shorter, mitigating the extrapolation error. (A4) also implies that marginal damages are zero when emissions are zero, which is a common and reasonable assumption.

(A4) combined with (A3) implies

$$\frac{\partial D(t+j)}{\partial s_{t+j}} = \theta e_t \text{ for } j = 0,\ldots,k \qquad (6)$$

where $\theta = \delta / \nu$, a positive constant. Neither c nor θ are typically encountered in ordinary environmental policy models. θ measures the change in damages due to a change in the state variable, per unit of emissions. c is the degree of homogeneity of the state variable, and thus equals the sum of the partials of $s(t)$ with respect to k emission lags, with each term multiplied by $e_{t-k} / s(t)$. If we assume emissions are constant for k periods, the product $c\theta$ can be shown to equal the sum of the partials of the damage function with respect to k lags of e, all divided by $s(t)$. In other words, $c\theta$ can be thought of as, approximately, the marginal damage rate, or marginal damages proportional to the value of the state variable.

The next assumption eliminates the role of discounting when summing marginal damages over the lag interval k. While it introduces a form of imprecision in the tax instrument, it is conservative in the sense that it will tend to overstate rather than understate the value of the sequence of damages.

(A5) ASSUMPTION 5: In the evaluation of τ using equation (3), β^j is set equal to unity for $j = 1,\ldots,k$.

The discounted present value of damages is defined by Equation (3). Since emissions only have an effect for up to k lags we can write this as

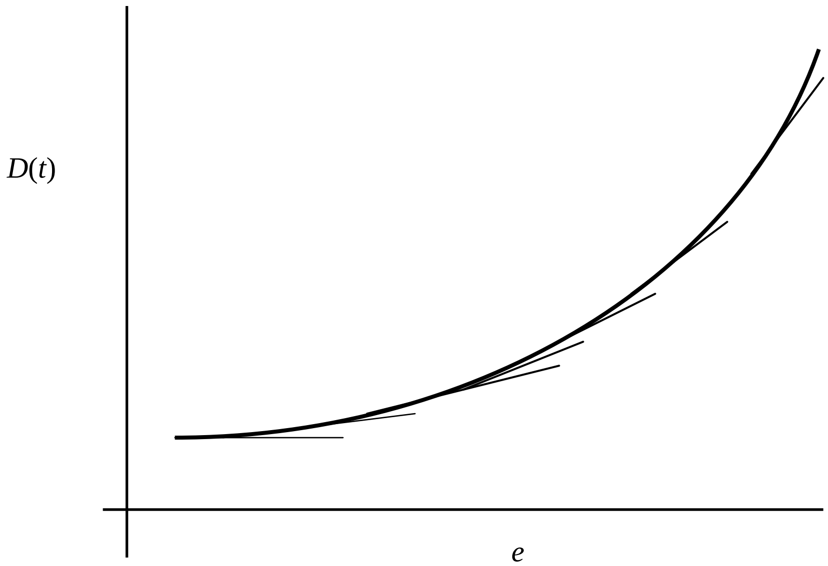

FIGURE 3: Approximation of damage function embodied in Assumption A4

$$
\begin{aligned}
\frac{\partial V(t)}{\partial e(t)} &= \sum_{j=0}^{k} \beta^{j} \frac{\partial D(s(t+j))}{\partial s(t+j)} \frac{\partial s(t+j)}{\partial e_t} \\
&= \theta e_t \sum_{j=0}^{k} \beta^{j} \frac{\partial s(t+j)}{\partial e_t} && \text{by (6)} \\
&= \theta e_t \sum_{j=0}^{k} \frac{\partial s(t)}{\partial e_{t-j}} && \text{by (A2) and (A5)} \qquad (7).
\end{aligned}
$$

Now denote, respectively, moving sums, moving averages and weights of e_t as $E_t \equiv \sum_{i=0}^{k} e_{t-i}$, $\bar{e}_t \equiv E_t/(k+1)$ and $\varepsilon_i(t) \equiv e_{t-i}/E_t$. Note that $\sum \varepsilon_i = 1$.

(A1) implies (by Euler's theorem)

$$c \cdot s(t) = \sum_{j=0}^{k} e_{t-j} \frac{\partial s(t)}{\partial e_{t-j}}$$

Hence

$$c \cdot s(t) = E_t \sum_{j=0}^{k} \varepsilon_{t-j} \frac{\partial s(t)}{\partial e_{t-j}} \qquad (8).$$

Equation (8) states that $cs(t)/E_t$ equals a weighted average of the partial derivatives of *s*. In situations where the emissions do not vary too much in percentage terms, over the *k*-period interval an unweighted mean will be a reasonable approximation to the weighted mean, hence:

(A6) ASSUMPTION 6: $c \cdot s(t)/E_t \approx \frac{1}{k+1} \sum_{j=0}^{k} \frac{\partial s(t)}{\partial e_{t-j}}$.

Combining (A6) and Equation (7) yields

$$\frac{\partial V(t)}{\partial e_t} \approx \theta e_t (k+1) c \cdot s(t)/E_t$$

Hence

$$\tilde{\tau}_t = \gamma \frac{e_t}{\bar{e}_t} s(t) \qquad (9)$$

where $\gamma = c\theta$ and ~ denotes the approximation to the optimum.

Equation (9) is the state-contingent pricing rule. It is an easily-calculated approximation to the marginal damages of the complex intertemporal externality shown in equation (3). Comparing it to equation (4), the trick to its usefulness has to do with how knowledge of the future is represented. In a system with inertia, emissions today will have an effect on the state variable over an uncertain span into the future. But this means that the current value of the state variable must also reflect the influence of past emissions. If the set of lag relationships extending over the past is structurally similar to the set of lag relationships that will extend into the future, then the current observation of *s* contains information about how past emissions affected today's climate, and hence how today's emissions will affect the future climate. Equation (9) uses that information to guide the emission price path.

Note that although (9) only uses information dated at time *t*, the underlying form of the tax is determined by equation (4). That means that, in principle, equation (9) charges firms today for the future value of their emission damages as well as the current value.

There remains one assumption to invoke: $\bar{e}_t$ is not known exactly since it depends on the unknown lag length *k*, so a lag length must be selected. However, trailing averages are smoothing devices, so unless the emissions series is extremely volatile, $\bar{e}_t$ will be relatively stable across a range of choices of lag length.

Figure 4, taken from McKitrick (2010), shows the implied value of the state-contingent emissions tax over the 1979-2010 interval, using the mean temperature of the tropical troposphere (see next section) and calibration of the free parameter to yield a value of $15 per tonne in 2002. The thick line is a 3-year moving average.

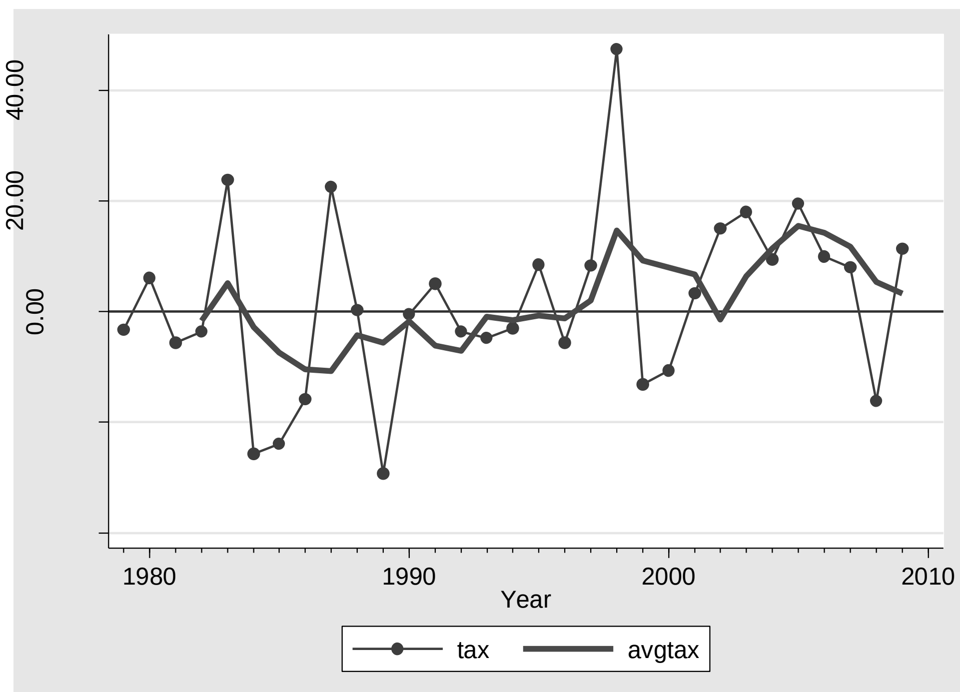

FIGURE 4: Value of state-contingent tax on greenhouse gas emissions since 1979. 'avgtax' denotes 3-year moving average.

McKitrick (2010) presented synthetic examples calibrated to the stylized facts of global warming, in which a series of complex and varied simulations into the future are shown, and in each case the state-contingent mechanism $\tilde{\tau}_t$ closely followed the unobservable optimum, even when known errors in approximation were introduced, with correlations typically in the >95% range.

Choice of State Variable

Assumption (1) requires that the state variable be chosen carefully, so that while it represents the factor that influences damages, as much as possible it represents the influence of emissions rather than exogenous factors. To give an example, while forest fires are costly, and potentially influenced by global warming, it would not make sense to use the number and extent of forest fires as the state variable for global warming since they are mainly influenced by causes unrelated to greenhouse gas emissions. Denote these other causes as $x(t)$. If the state variable for forest fires is $s(t) = \alpha x(t) + (1-\alpha)e_t$ and α is close to 1, then Assumption (1) can easily be seen fail.

On the other hand, even though Assumption (1) would certainly hold if we simply used total emissions (or the atmospheric stock) as the state variable, this would not make sense unless it were known that the sensitivity term $ds(t)/de(t) = 1$ for all t, which is not the case, or at least it is not known to be the case. Imposing it by assumption means we are assuming away all the uncertainties over the magnitude of $ds(t)/de(t)$. In the extreme, it would imply that e_t should still be priced the same regardless of whether $ds(t)/de(t) = 0$ or $ds(t)/de(t) \to \infty$, which is clearly not credible. CO_2 has not historically been regulated as an air pollutant because it is naturally occurring and harmless for humans except at extremely high concentrations. It is only a candidate for policy intervention if it turns out to have a strong effect on the climate by changing atmospheric temperatures. So the state variable should be chosen to capture that contingency.

Hsu (2011) correctly points out that the state variable must be non-manipulable and reliably, regularly and uncontroversially measurable. He then proposes a basket of measures arising from the tort principle of focusing on the harms done by climate change, which would include (a) the global mean temperature, (b) counts of days of "unusually" high or low temperatures, (c) counts of extreme rainfall or drought events, (d) rises in sea level, (e) ocean acidity and (f) numbers of hurricanes above a certain intensity level. The weighting scheme is not specified.

McKitrick (2010), by contrast, proposes just one measure, the mean temperature of the tropical troposphere. This is based on the argument that, as well as being non-manipulable and reliably measurable, the state variable must represent an index of the unique effect of emissions with as rapid a response time as possible. The climate modeling work surveyed in the 2007 IPCC Report (IPCC 2007) and Karl et al. (2006) clearly points to the tropical troposphere as a place in the climate system where an unusually strong and rapid signal of the effects of greenhouse gases should be observable. Thorne et al. (2011) and Fu et al. (2011) provide updated analyses emphasizing the importance of the response to greenhouse forcing in the tropical troposphere as a metric of climate sensitivity, since this is where climate model behaviour is strongly constrained. The Arctic surface temperature is also believed to be strongly influenced by greenhouse gases, but is also strongly influenced by natural phenomena and, being an oceanic region, is subject to poor spatial sampling.

The Hsu basket approach includes measures that can be criticized on the grounds mentioned above. The tort concept is not the appropriate criterion for picking a state variable, instead the connection to greenhouse gases is. The tort concept matters when it comes to estimating parameters of the damage function, but that is a different step.

Items (b), (c) and (f) are sub-grid scale weather phenomena with poorly-understood and controversial connections to greenhouse gases. The history of global weather shows that extended intervals of elevated drought conditions or excessive moisture can and do arise purely due to natural variability, so these measures do not satisfy Assumption (1). In other words, such index terms could go up (or down) for extended periods even if greenhouse gases turn out to have no effect on the climate. Item (e) is not an effect of global warming, it is potentially a measure of the atmospheric CO_2 level since acidification (or reduced alkalinity) would occur if large enough quantities of CO_2 were to dissolve in the world's oceans. While use of this measure would make sense if the state variable were ocean alkalinity (and the damage function were defined accordingly), if the damages are connected to atmospheric warming the use of ocean alkalinity would, at best, be equivalent to making emissions the state variable, which is a flawed concept for the reason noted above. Item (d) has a somewhat controversial connection to global warming, as Hsu notes, with sharply varying projections as to the rate of rise over the coming century. The main deficiency as a state variable is that sea levels have apparently followed a steady upward trend for centuries, and what is at issue is not whether this will continue but whether global warming will cause it to accelerate. So the more relevant concept would be the acceleration of sea level rise. But the oceanic system is so slow to respond that any acceleration may not be apparent for a long time, due to the inertia involved. Finally, item (a) suffers from controversy and data quality problems (see, for instance, de Laat and Maurellis 2006, McKitrick and Nierenberg 2010, Fall et al. 2011), especially over the ocean (Thompson et al. 2008, Christy et al. 2001), as well as being subject to many influences other than greenhouse gases, which is why the satellite-based mean temperature of the tropical troposphere would be a more accurate state variable in this context (see discussion in McKitrick 2010).

These are all very large issues and the above summary paragraph is not sufficient to dispose of them all. Indeed one of the benefits of the state-contingent approach is that it forces a discussion on the issue of how climate change ought to be measured. The fact that it is hard to come up with a simple answer provides some needed context to the policy discussion, namely that there is work to be done simply to clarify what we are talking about before supposing we are in a position to measure the costs and benefits of policy. If we do not agree on how to measure global warming, how would we know if a policy, once enacted, was making a difference?

4. A permits-based prediction market to generate future prices

Use of a state-contingent, myopic pricing rule does not mean that we only put a price on emissions after the damage is done, since the approximation is to the discounted present value of current and future emissions. Also, businesses are forward-looking and investment plans are based not only on today's prices, but on expectations of future prices. While those forecasts may be wrong, there is no incentive for firms to make systematically biased forecasts, since in the future the actual prices will be revealed and mistakes will be costly. However, the basic proposal in McKitrick (2010) does not reveal the expectations in the market, except as they are implicit in firms' investment decisions.

In order to reveal the future price expectations, Hsu (2011) proposed extending the tax rule of McKitrick (2010) to include a tradable exemption permits market for future periods. The permits market would allow firms in year *t* to purchase a per-tonne exemption from paying the tax in years *t*+1, *t*+2, etc., out to the limit of the forecast horizon. Hence it would establish, in effect, a prediction market for the state-contingent tax, and by implication, the state variable. As with any prediction market, the maximum financial rewards would accrue to those who make the best forecasts, so the published list of price futures would then imply an optimal and objective set of climate forecasts. If agents believe the permit price for year *t*+10 (say) is below the likely value of the tax rate, or in other words if agents believe people are underestimating the amount of likely warming between now and then, they will have an incentive to buy permits, and will thereby drive up the price. And likewise, agents who believe warming is being overestimated will have an incentive to sell permits, or short the market.

Hsu (2011, pp. 33-35) captures well the incentives this would create for basing forward investment behaviour on the most accurate possible scientific information about climate:

> The price bid by emitters for say, permits to emit in 2020, would speak volumes about private expectations of the consequences of climate change, free from, as climate skeptics claim, conspiracies by climate scientists to shore up their research grant fiefdoms, or desires by radical environmentalists who really wish to use climate change as an excuse for imposing environmental restrictions. ….Without an obvious ideological horse in the race, emitters like [American Electric Power, AEP] will brutally and honestly evaluate the credibility of climate science, and spend its climate investigation money carefully. It is the participation of large emitters in a cap-and-trade program for emissions futures that is likely to make or break the credibility of climate science. In essence, this proposal uses markets to turn the evaluation of climate science over to those emitters that will potentially rely on those permits for their emitting operations. And such a liability could be very significant: in 2005 AEP emitted approximately 161 million tons of CO_2; if one assumed a very modest carbon tax that was set to five dollars per ton at current climate outcomes, AEP's annual carbon tax liability would be about $805 million. If climate outcomes increased by say, twenty-five percent, its annual carbon tax liability would top one billion dollars. All 101 electricity generators in the EPA's Egrid database would have a combined current carbon tax liability (assuming a rate of five dollars per ton of CO_2) of $8.75 billion. Environmental advocates may chafe at the notion that the greatest greenhouse gas emitters will have such a large say in evaluating the quality of climate science, but $8.75 billion is a lot of impetus for honestly evaluating climate science.

Some of the concern expressed by authors like Stern (2006) and Weitzman (2009) is that the climate may be subject to nonlinear effects of greenhouse gases (Hsu 2010 also raises this concern). Future greenhouse warming may be subject to sudden, rapid acceleration into catastrophic levels of environmental damage. But when we consider how to integrate this possibility into current decision-making, we confront the same dilemma as before, namely between the fears of highly costly potential environmental damages and highly costly potential policy mistakes. The prospect of sudden catastrophic change only amplifies the magnitude of each of these two apparently awful options. But now imagine a state-contingent tax is implemented along with a thirty-year sequence of prediction markets. Suppose initially the prediction markets show a slow, smooth ramp in CO_2 emission prices. This would indicate the market discounts the possibility of a serious bifurcation or (so-called) “tipping point” into catastrophic change. But to the extent that anyone can construct a *credible* scientific argument that a bifurcation is approaching, they would know that those emission futures are underpriced and they be able to invest profitably in them, knowing that once their scientific arguments were digested by the prediction market, there would be a predictable ramp in prices. The prediction market would be the most reliable instrument available for generating a rational signal of a coming bifurcation. Again, the reason is that firms would have a strong financial incentive to get the climate science right, to whatever extent possible, and not to make a systematic forecast error or adopt a disingenuous view of the underlying problem. If an objective information-processing system like a prediction market existed, research warning of a climate nonlinearity would be “brutally and honestly” evaluated, and could potentially be the basis for a spike in emission price futures, thereby providing an instantaneous signal of the gravity of the threat. For this reason, while the state-contingent price mechanism is not guaranteed to resolve the uncertainty at the core of Weitzman’s analysis, it provides incentives for the maximum possible resolution; in other words, no other policy path could provide a more objective basis for forming expectations about possible future catastrophes.

The tradable futures market would not allow firms to evade paying the tax on future emissions, since they would have to buy the exemption permits. Instead it would allow firms to trade on changes in expectations about the future path of the state variable and the emissions price. This illustrates another distinction between the state-contingent approach and the IAM approach. In the latter, it is assumed that we possess correct, unbiased scientific information regarding climate, and all we require is a mechanism to implement the optimal price. No allowance is made for the possibility that key parameter estimates (such as climate sensitivity) might be biased due to distorted incentives for scientists who prepare such forecasts. In the IAM approach, implementation of the policy does not induce an improvement in the scientific basis of policy. But the approach in this chapter does not assume we have correct scientific information. It works even if the information we have today is incorrect, either through technical inadequacy or researcher bias. The policy mechanism rewards those agents that eliminate bias and inaccuracy in their forecasting work, by allowing them to trade in futures markets on the difference between the current market price and their expectations of how it will evolve.

Some implementation issues that need to be addressed for other types of policy would also need to be addressed in the present analysis. For instance, ideal implementation would be at the global level with each country imposing the tax based on damage valuations over the whole world, rather than at the national level. But incentives favour free-riding, making it difficult to ensure global participation. On the other hand, the tax approach has the advantage that revenues can be retained domestically to reduce the excess burden of the tax system, reducing the macroeconomic cost of implementation, and the state-contingent nature of the policy provides reassurance that the stringency will be increased only if the underlying problem is shown objectively to merit such tightening, and both these features may increase incentives for multilateral cooperation compared to alternative policies.

Another potential issue is that as new information becomes available, policymakers may decide the initial calibrated value of the tax was incorrect and must be revised. This too would be a problem for any policy mechanism. In the state-contingent case an argument can be made that incentives favour efficient acquisition of the information required to optimally calibrate the tax. In this regard the feedback between policy and science would be particularly fruitful, since emitters subject to the tax will have an incentive to pay for the best, most objective information they can get, and there is no assumption that the forecasts made prior to implementing the policy are correct, or will be validated in the future.

5. Further Extensions to the Basic Concept

Endogenous emissions response

Since no information about abatement costs are used in deriving the tax *t*, it may seem that it cannot be a complete policy prescription. The tax paths derived in integrated assessment models are solutions to a two-sided optimization problem, with intertemporal damages netted against intertemporal abatement costs. However, it is important to bear in mind that the formula above does not prescribe a policy *path*, it yields a *rule* that ties the tax rate to the environmental state. The actual path of taxes over time will be determined by the evolution of the state variable, and the ensuing level of abatement will be determined by emitters who respond to the current and expected future tax rates according to their current and future marginal abatement costs. If the capital stock is highly variable then firms will respond to current emission tax rates as they would to any variable input costs. If capital is fixed and time-to-build lags are long, firms will need to form forecasts of the future values of the tax rate, which in turn will depend on future values of the temperature variable, and the usual structure of optimal investment under uncertainty will ensue.

The simulations in McKitrick (2010) assume that the emissions tax $\tau(t)$ is a function of the future state $s(t)$, but not vice-versa. The coherence between the approximate tax, given by equation (9), and the actual optimum is demonstrated with this assumption in force. However, it is likely to be the case that $s(t)$ is a function of $\tau(t)$ as well. This will certainly be true if emission tax increases reduce future temperatures, but it will also be true as long as they reduce future emissions since $e(t)$ enters equation (9) directly, as well as through $s(t)$, and the emissions path in McKitrick's simulations are exogenous. The rationale is that emissions are added up globally whereas the tax is imposed by single governments, and no one country can do much to reduce global total emissions through unilateral action. However, if all countries (or a substantial majority) were to enact the tax, total emissions would be affected by the path of $\tau(t)$. In that situation it has not been shown that the state-contingent pricing rule would yield a stable approximation to the true optimum. This is something that needs to be addressed in subsequent research.

Coalition-formation

An interesting feature of the state-contingent tax is its potential ability to appeal to a broad coalition of interests.[1] People with conflicting expectations about the future evolution of the state variable will nevertheless each expect to observe his or her preferred policy path. Those who think emissions have no effect on climate will expect low emission taxes to prevail in the future, and those who think they have strong effects will expect the tax to increase rapidly. Since each agent expects to get his or her preferred outcome, it may be easier to get agreement for implementation. One of the challenges of climate policy is the need to get agreement at the global level. Different regions have different views on the urgency of the problem and how it compares to their domestic economic priorities, which makes it all but impossible to get agreement on emission targets, or to ensure compliance with earlier agreements. Asking policymakers around the world to agree on a state-contingent tax might be easier. The tax revenue would stay within each country, reducing the burden of inequality across different nations. And during the negotiations, there would be no reason for countries that took opposing views on the likely future path of temperatures to take opposing views on whether the tax is desirable, since each party will expect to get what they consider to be the "correct" outcome.

Suppose the stringency, and hence costliness, of a policy can be summarized as a parameter z, where a higher value of z corresponds to a more stringent policy. A potential voter (person i) has a private view of the optimal value of z given his or her beliefs about the marginal effect of emissions on $s(t)$, which we denote by s^i. Their preferred policy is thus $z_i(s^i)$. If $z < z_i(s^i)$ then the proposed policy is deemed too lax, and vice versa.

Typical median voter models only require $z \geq z_i(s^i)$ to ensure voter i's support, namely people are satisfied as long as z equals or exceeds their preferred policy. But suppose the voter's support for a policy z declines based on the distance ($z > z_i(s^i)$), i.e. z can be too strict even for someone who prefers a relatively high value. In this case, obtaining majority support can be difficult since it faces two-sided opposition. For example, a moderate emissions price might be opposed by those who prefer it to be much higher, as well as much lower. But proposals to adjust z up or down may alienate as many supporters as they would attract, making it impossible to get a majority.

Suppose a potential voter uses a quadratic loss function like $L_i = [z_i(s^i) - z]^2$ to determine his or her degree of opposition to the policy. Then the greater the variance of beliefs about s^i, the smaller the coalition of support for any policy. The state-contingent approach can potentially alleviate this problem, however. The policymaker no longer proposes a fixed value of z, but instead proposes a function of the observed state $z(s(t))$ over time. Each agent will then expect future values of s to be correlated with s^i, hence the sum of the expected loss terms L_i will be smaller than before, even if the variance of beliefs about s^i remains large. Intuitively, by proposing a policy target that is dependent on the actual future state, each agent "expects" to get his or her preferred outcome. The one who expects the emissions to have a large effect expects the policy to end up being stringent, while the one who expects emissions to have little effect expects the policy to end up being lax. Consequently both types of agents expect small losses from

[1] As an anecdotal illustration, Hsu (2011) and McKitrick (2010) hold very different views on the underlying threat of global warming, yet both advocate the same policy mechanism, albeit with different recommendations as to the appropriate state variable.

the policy, and have an equally strong incentive to support it, even though they have conflicting views about the form it will actually take.

6. Conclusions

Uncertainties over the future path of global warming and the underlying severity of the problem make derivation of an intertemporally-optimal emissions price on carbon dioxide both theoretically and politically very difficult. IAM-based approaches to the problem assume knowledge of a lot of key parameters, while learning models suggest it will take too long to resolve those uncertainties to be of much use in the current debate. If all conceptually-possible climate risks are considered, it may be impossible to place a finite value on full insurance against those risks without arbitrarily truncating the range of extreme outcomes being considered.

A fundamental problem with the existing analyses of carbon dioxide pricing is that agents and policymakers cannot commit to a long term emissions price. The issue is polarized such that fear of two very large potential mistakes seems to have paralyzed the decision-making process, namely fear of climate catastrophe due to failure to act, and economic catastrophe from inept action. The nature of the climate issue makes these fears justified.

This chapter explores an alternative approach based the concept of state-contingent pricing, in which agents commit to a pricing *rule* rather than a *path*. The rule connects current values of the emissions price to observed temperatures at each point in time. In essence, if the climate warms, the tax goes up, and vice versa. A derivation is provided showing how such a rule yields an approximation to the unknown optimal dynamic externality tax, yet can be computed using currently-observable data. A recently-proposed extension coupling the state-contingent tax with a tradable futures market in emission allowances would then yield not only a feasible mechanism for guiding long term investment, but an objective prediction market for climate change.

There are many potential advantages of the state-contingent approach. For one thing, people with divergent views on the nature of the climate issue can still potentially commit to the same instrument, since each one would expect to get his preferred outcome. The rule is structured such that, however the future unfolds, in retrospect we will know that we followed a reasonably good approximation to the optimum, and the incentives along the way favour the use of unbiased forecasts of the pricing path to guide investment decisions. Consequently there are informational, as well as theoretical and practical, advantages to the state-contingent approach, which make it worth exploring in more depth as a potentially viable tool for implementing sound climate policy.

7. References

Christy, John R., David E. Parker, Simon J. Brown, et al. 2001 Differential trends in tropical sea surface and atmospheric temperatures since 1979. *Geophysical Research Letters* 28, no. 1: PAGES 183–186.

de Laat, A. T. J. , and A. N. Maurellis. 2006. Evidence for influence of anthropogenic surface processes on lower tropospheric and surface temperature trends. *International Journal of Climatology* 26 (June 1): 897-913.

Fall, S., A. Watts, J. Nielsen-Gammon, E. Jones, D. Niyogi, J. Christy, and R.A. Pielke Sr., 2011: Analysis of the impacts of station exposure on the U.S. Historical Climatology Network temperatures and temperature trends. *J. Geophys. Res.*, in press.

Fu, Qiang, Syukuro Manabe and Celeste M. Johanson (2011) "On the warming in the tropical upper troposphere: Models versus observations" *Geophysical Research Letters* VOL. 38, L15704, doi:10.1029/2011GL048101, 2011.

Geweke, John 2001 "A note on some limitations of CRRA utility," *Economic Letters, 71, 2001, 341–345.*

Hsu, Shi-Ling. 2011 A Prediction Market for Climate Outcomes (February 26, 2011). Available at SSRN: http://ssrn.com/abstract=1770882

IPCC. 2007. *Climate Change 2007: The Physical Science Basis.* Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change. Solomon, S., D. Qin, M. Manning, Z. Chen, M. Marquis, K.B. Averyt, M. Tignor and H.L. Miller (eds.). Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.

Karl, T. R., Susan J. Hassol, Christopher D. Miller, and William L. Murray. 2006. *Temperature Trends in the Lower Atmosphere: Steps for Understanding and Reconciling Differences.* Synthesis and Assessment Product. Climate Change Science Program and the Subcommittee on Global Change Research. http://www.climatescience.gov/Library/sap/sap1-1/finalreport/sap1-1-final-all.pdf. Accessed August 3 2010.

Kelly, D. L., and C. D. Kolstad. 1999. Bayesian learning, growth, and pollution. *Journal of Economic Dynamics and Control* 23, no. 4: 491-518.

Leach, A. J. 2007. The climate change learning curve. *Journal of Economic Dynamics and Control* 31, no. 5: 1728-1752.

McKitrick, Ross R. (2010) A Simple State-Contingent Pricing Rule for Complex Intertemporal Externalities. Energy Economics doi:10.1016/j.eneco.2010.06.013

McKitrick, Ross R. and Nicolas Nierenberg 2010 Socioeconomic Patterns in Climate Data. *Journal of Economic and Social Measurement*, 35(3,4) pp. 149-175. DOI 10.3233/JEM-2010-0336.

Mears, C. A., and F. J. Wentz. 2005. The effect of diurnal correction on satellite-derived lower tropospheric temperature. *Science* 309: 1548-1551.

Nordhaus, William D. (2009) "An Analysis of the Dismal Theorem" Cowles Foundation Discussion Paper No. 1686.

Nordhaus, William D. 2007. To Tax or Not to Tax: Alternative Approaches to Slowing Global Warming. *Rev Environ Econ Policy* 1, no. 1 (January 1): 26-44. doi:10.1093/reep/rem008.

Parson, Edward A. and Darshan Karwat (2011) Sequential Climate Change Policy *WIREs Climate Change* DOI: 10.1002/wcc.128.

Pindyck, Robert (2011) "Fat Tails, Thin Tails and Climate Change Policy" *Review of Environmental Economics and Policy* 2011, forthcoming.

Spencer, R. W., and J. R. Christy. 1990. Precise monitoring of global temperature trends from satellites. *Science* 247, no. 4950: 1558-1562.

Stern, Nicholas. 2008. *Stern Review on the Economics of Climate Change* London, UK: Her Majesty's Treasury.

Thompson, David W. J., John J. Kennedy, John M. Wallace, and Phil D. Jones. 2008. A large discontinuity in the mid-twentieth century in observed global-mean surface temperature. *Nature* 453, no. 7195 (May 29): 646-649. doi:10.1038/nature06982.

Thorne, Peter W. , Philip Brohan, Holly A. Titchner, Mark P. McCarthy, Steve C. Sherwood, Thomas C. Peterson, Leopold Haimberger, David E. Parker, Simon F. B. Tett, Benjamin D. Santer, David R. Fereday, and John J. Kennedy (2011) "A quantification of uncertainties in historical tropical tropospheric temperature trends from radiosondes." *Journal of Geophysical Research-Atmospheres* VOL. 116, D12116, doi:10.1029/2010JD015487, 2011.

Weitzman, Martin L. (2009) On Modeling and Interpreting the Economics of Catastrophic Climate Change *Review of Economics and Statistics* 91(1) 1—19.