

Georgia State University

ScholarWorks @ Georgia State University

ECON Publications

Department of Economics

5-2011

How should environmental policy respond to business cycles? Optimal policy under persistent productivity shocks

Garth Heutel

Georgia State University, gheutel@gsu.edu

Follow this and additional works at: https://scholarworks.gsu.edu/econ_facpub



Part of the [Economics Commons](#)

Recommended Citation

Heutel, Garth, "How should environmental policy respond to business cycles? Optimal policy under persistent productivity shocks" (2011). *ECON Publications*. 15.

https://scholarworks.gsu.edu/econ_facpub/15

This Article is brought to you for free and open access by the Department of Economics at ScholarWorks @ Georgia State University. It has been accepted for inclusion in ECON Publications by an authorized administrator of ScholarWorks @ Georgia State University. For more information, please contact scholarworks@gsu.edu.

How Should Environmental Policy Respond to Business Cycles? Optimal Policy under Persistent Productivity Shocks

Garth Heutel
University of North Carolina at Greensboro
PO Box 26165
Greensboro, NC 27402-6170
and NBER
336 334 4872
Fax: 336 334 5580
gaheutel@uncg.edu

May 2011

Abstract

How should environmental policy respond to economic fluctuations caused by persistent productivity shocks? This paper answers that question using a dynamic stochastic general equilibrium real business cycle model that includes a pollution externality. I first estimate the relationship between the cyclical components of carbon dioxide emissions and US GDP and find it to be inelastic. Using this result to calibrate the model, I find that optimal policy allows carbon emissions to be procyclical: increasing during expansions and decreasing during recessions. However, optimal policy dampens the procyclicality of emissions compared to the unregulated case. A price effect from costlier abatement during booms outweighs an income effect of greater demand for clean air. I also model a decentralized economy, where government chooses an emissions tax or quantity restriction and firms and consumers respond. The optimal emissions tax rate and the optimal emissions quota are both procyclical: during recessions, the tax rate and the emissions quota both decrease.

JEL codes: Q58, E32, Q54

Keywords: Climate change; Environmental policy

I would like to thank Don Fullerton, Bill Hogan, Stephen Holland, John Horowitz, Bidisha Lahiri, Billy Pizer, Rob Stavins, Dan Tortorice, Richard Zeckhauser, two anonymous referees, and various seminar participants for helpful comments, and the Kernan Brothers fellowship at the Harvard University Center for the Environment for funding.

Business cycles can substantially affect an economy, and many policies recognize economic fluctuations and are designed to adapt to them. For example, unemployment insurance is designed to have a stabilizing countercyclical effect since payments increase with unemployment. Environmental policy, though, typically has not been designed to respond to business cycles, likely because the scale of most environmental policies is small relative to the economy. However, addressing global climate change will require policies that dwarf conventional environmental policies in scale and scope. Can climate policy designers continue to ignore business cycles, or does climate policy require a more explicit integration with macroeconomic fluctuations?

This paper investigates how environmental policy optimally responds to business cycles. I develop a dynamic stochastic general equilibrium model with persistent productivity shocks and with pollution as a stock variable that negatively affects the economy. The model is calibrated to the US economy and to damages from carbon dioxide, the leading greenhouse gas. I numerically solve for the dynamic optimal response of policy to shocks. A welfare analysis compares the dynamic policy that optimally adapts to productivity shocks to the best static policy that holds emissions constant at its long-run optimal value but ignores shocks. I also model a decentralized economy, where firms and consumers optimize in response to a government policy of pollution taxes or quantity constraints, potentially under information asymmetry.

This analysis highlights the importance of allowing climate policy to adjust to business cycles in three ways. The first is from political economy reasons. Firms worry that a rigid cap will result in permit prices that are too high. Suggested methods of cost containment, for instance safety valve prices, are rather *ad hoc*. The suggested policy instrument in this paper represents the first-best response to cost fluctuations, which may be more politically appealing to both polluters and environmentalists. Policies have been politically lambasted for *not* adapting to economic fluctuations. The clearing price of emissions permits after the first auction of the Regional Greenhouse Gas Initiative (RGGI) in September 2008 was \$3.07 per ton of carbon dioxide, well below what was expected. A likely explanation is that the cap was set too high and inflexible given the unexpected downturn in the economy (Metcalf 2009). The first-best policy

here could alleviate political concerns on both sides of climate policy and eliminate the inclusion of *ad hoc* measures to adapt to business cycles.¹

Second, the welfare analysis in this paper shows that the net benefits of adapting to cycles are of a comparable magnitude to other environmental policies, although they may be smaller than those of getting the correct long-run policy. Third, these net benefits will not be uniformly distributed across all agents in the economy. Nearly all of the net benefits of the optimal policy come from smoothing abatement costs rather than pollution damages; this is because pollution is such a long-lived stock that cyclical fluctuations in emission do not substantially affect damages. Thus the beneficiaries of the cost saving from the optimal policy are those who bear the costs of the policy. Assuming forward shifting of the policy costs onto consumers, the burden of climate policy is strongly regressive (Metcalf et. al. 2010). Backward shifting of the policy costs onto firms may also make the distribution of these cost savings important. Electricity generation contributes 41% of all U.S. carbon dioxide emissions although electric utilities comprise only about 2.5% of the U.S. economy. If utilities bear roughly half the burden of the policy, then the cost savings from dynamically optimizing may represent a large fraction of their overall expenses.

Most of what we know already about how policy should respond to productivity shocks comes from studies that focus on optimal policy under information asymmetry over abatement costs; this includes Weitzman (1974) and its many extensions. These papers use a reduced form specification (typically quadratic) of costs and benefits of pollution. A total factor productivity shock can indirectly affect these reduced form functions. Instead of imposing a reduced form quadratic cost function and another reduced form quadratic benefit function, I begin with a standard dynamic stochastic general equilibrium (DSGE) model of real business cycles. I add to the standard model an externality that comes from an accumulated stock of pollution. The stochastic element of the model does not arise from shocks to abatement costs; rather, it comes from total factor productivity shocks. This creates a price and an income effect that can counter each other, such that the response of optimal policy is not apparent even in the absence of asymmetric uncertainty over the value of the shock.² This choice of modeling is also

¹ For instance, during the debate over a cap-and-trade bill in the US House of Representatives in 2009, three amendments were unsuccessfully offered to the bill that would have suspended the program if gasoline prices exceed \$5 per gallon, if electricity prices climb by 10%, or if the unemployment rate hits 15%.

² Kelly (2005) also considers the effects of productivity shocks on environmental policy, in a static setting.

advantageous in that I can draw upon a large prior RBC literature to calibrate the model. The reduced form shocks to abatement costs and benefits present in other models are likely due to productivity shocks, and here they are modeled directly as such. Furthermore, most of these earlier studies are static, and those that are dynamic do not model productivity shocks.³ Earlier papers are focused on information asymmetry. While this is not the focus of this paper, the model here is extended to include asymmetric uncertainty, when firms but not policymakers observe technology shocks.

The paper begins with an empirical examination of how carbon dioxide emissions in the United States respond to cyclical fluctuations in GDP. I estimate the elasticity of emissions with respect to GDP. This question to my knowledge has not been answered econometrically. The expected sign of this elasticity is positive, but its magnitude has not been measured. Using monthly data on GDP and carbon dioxide (CO₂) emissions, I find that emissions are significantly procyclical with an elasticity between 0.5 and 0.9. Thus, emissions are inelastic with respect to output. This result is robust to a number of empirical specifications. The purpose of the empirical section is twofold. First, a measurement of this elasticity is interesting in and of itself (during a recession of a specified magnitude, by how much do we expect emissions to drop?) Second, the estimated elasticity is used as a parameter in the calibrated model.

The model incorporates two offsetting effects from productivity shocks: a price effect and an income effect. A positive shock to productivity increases wealth, and so the income effect leads to higher demand for a clean environment and thus lower emissions. However, since capital is more productive after a positive productivity shock, the opportunity cost of spending on abatement instead of investing in capital is higher and thus abatement is relatively costlier. The price effect thus leads to lower demand for abatement and higher emissions.

After calibrating and solving the dynamic model, I find that the optimal policy response to an economic expansion is to increase emissions, because the price effect dominates the income effect. In fact, a policy that pegs emissions to GDP is a good approximation of optimal policy. This qualitative result is robust to sensitivity analysis over several parameters. Simulating the decentralized economy shows that both the optimal emissions tax rate and the optimal emissions quota are procyclical; they increase during an expansion. Optimal policy thus

³ See Fell et. al. (2008), Newell and Pizer (2003), Hoel and Karp (2002), and Pizer (2002) for examples of papers with a dynamic model of environmental policy but without modeling productivity shocks.

dampens the procyclicality of emissions: emissions rise during booms and fall during busts, but not by as much as they would without optimal policy. An emissions quota policy is strengthened during recessions (the quota is reduced), while an emissions tax policy is weakened (the tax rate is reduced). This is perhaps a political advantage of taxes over quotas, if it is difficult to strengthen environmental policy during recessions.

Several previous studies have considered the relationship between business cycles and environmental policy.⁴ Kelly (2005) compares prices and quantities in a static general equilibrium setting with total factor productivity shocks. Strand (1995) develops a model of optimal environmental policy that includes worker moral hazard, with business cycles operating through stochastic changes in output prices. Bouman et. al. (2000) develop a model to find the optimal time to invest in abatement under business cycles driven by preference shocks, and they find that the best time is during downturns. Fischer and Springborn (2009) is the only study besides this one that incorporates environmental policy into a real business cycle (RBC) model. Their focus is not on optimal *dynamic* policy, but rather on how various *static* policies perform in the presence of productivity shocks. They find that a cap on emissions dominates either an emissions tax or an intensity standard (restriction on emissions per unit output) because the cap reduces the volatility of all economic variables, not just emissions.

In the next section, I present some dynamic descriptive statistics and regression results regarding the relationship between business cycles and carbon dioxide emissions in the US. Section II presents a dynamic model of a social planner's problem, which is calibrated in Section III. The model is solved and simulation results are presented in Section IV. Section V presents a decentralized version of the dynamic model, where competitive firms and utility maximizing consumers react to government policy of an emissions tax. Section VI concludes.

I. Empirical Analysis

I investigate the relationship between business cycles and emissions using emissions data at the monthly or quarterly level. While the rest of this paper is a normative analysis of optimal policy, this section is purely a descriptive analysis of how carbon emissions respond to business cycles in an economy without optimal policy. In fact, no carbon policy was in effect during this

⁴ Chay and Greenstone (2003) estimate the effect of air pollution on infant mortality in a quasi-experimental design where the 1981-1982 recession affects emissions differentially across counties.

period of analysis. The estimated elasticity is used in the normative analysis; to know how emissions ought to vary with cycles one must know how they do (sub-optimally) vary without any policy. The lack of any carbon policy alleviates one endogeneity concern, since carbon policy will not itself be adjusting with the business cycle. What is measured is the "natural" response of carbon emissions to aggregate changes in output. Though the choice of firms and households in how much to pollute is endogenous (i.e. they respond to the business cycle), this choice is precisely what I want to measure.

Blasing et. al. (2004) provide estimates of monthly emissions of carbon dioxide from fossil fuel combustion in the United States from 1981-2003.⁵ Figure 1 plots seasonally adjusted monthly levels of nationwide emissions and monthly real GDP. Both series are normalized so that the starting value is 1 (in January 1981).⁶ Over the period of analysis, both GDP and carbon emissions grew, but at different rates, so that the carbon intensity of the economy declined. GDP increased by a factor of 2, while carbon emissions increased only 25%. Business cycle effects can also clearly be seen in this figure; recessions in the early 1980s, early 1990s, and early 2000s are reflected in the GDP curve. Concurrently with these recessions, carbon emissions appear to drop off. Similarly, when GDP is rising at a fast rate, carbon emissions appear to do the same. This happens during the expansionary period of the mid- to late-1990s.

While basic patterns can be seen from eyeballing Figure 1, a more thorough method of identifying cycles in output or in emissions is available. The Hodrick-Prescott (HP) filter is a commonly used method to detrend time series, separately identifying the trend component from the cyclical component.⁷ I aggregate the monthly emissions and GDP data to the quarterly level, take the natural log of the values, and use the HP filter with 1600 as the smoothness parameter λ .

The cyclical component of the detrended data is presented in Figure 2. The patterns in the GDP curve correspond to recessions: troughs in the early 1980s, the early 1990s, and the early 2000s. For the emissions curve, some of these troughs can be clearly identified

⁵ Available at: http://cdiac.esd.ornl.gov/trends/emis_mon/emis_mon_co2.html. These estimates are based on reported fuel consumption from the Energy Information Administration in the Monthly Energy Review for several types of fossil fuels. Each type of fuel is converted into its associated level of carbon dioxide emissions, based on fuel-specific carbon dioxide emissions factors from the EPA.

⁶ Emissions data unadjusted for seasonality exhibit a large peak during the winter months and a smaller peak during summer months. The data are seasonally adjusted using the Census Bureau's X-12-ARIMA program, available at <http://www.census.gov/srd/www/x12a/>.

⁷ See Hodrick and Prescott (1997). Among the many papers that analyze time series using the HP filter are Kydland and Prescott (1982), and Chang and Kim (2007). See Baxter and King (1999) for comparisons of alternative filters.

concurrently, especially in the early 1980s and early 1990s recessions. The correlation between the two series seems to dissipate starting in the late 1990s, though a drop in GDP in 2000 is accompanied by one in emissions as well.

The standard deviation of cyclical GDP is 1.31%, while the standard deviation of cyclical carbon emissions is 2.04%, so emissions are more variable than GDP. The correlation coefficient between the two time series is 0.5627, with a p-value less than 0.0001. The two time series are strongly correlated.⁸ Unsurprisingly, when output increases, so do carbon emissions.

These data can also be analyzed using time-series analysis. Unconditional correlations show that periods of higher GDP (in deviations from trend) tend to occur with periods of higher CO₂ emissions. What is the magnitude of this relationship? Are emissions elastic or inelastic with respect to GDP? Table 1 presents regression results to identify the magnitude of the relationship between CO₂ emissions and GDP. In column 1, I present regression results from a seasonal ARIMA(1,1,1)×(0,1,1)₁₂ regression of the log of emissions on the log of GDP.⁹ The regression results show that CO₂ emissions are inelastic with respect to GDP, with a coefficient of 0.758. An alternative method of dealing with the seasonal component of the emissions data is to perform the regression on the emissions data that have already been seasonally adjusted. This is done in Column 2, and the coefficient on GDP is almost identical to that in Column 1.¹⁰

Both of these regressions use first differencing to eliminate the trend from the two time series. Alternatively, a number of other filters are available to identify the trend. The most commonly used filter in the business cycle literature is the HP filter. Column 3 presents regression results where the dependent variable is the deviation from trend in the log of CO₂ emissions, determined by the HP filter, and the independent variable is the deviation from trend in the log of GDP. The smoothing parameter λ is set at 129,600, the standard value for monthly data. A least squares regression is performed, allowing for a Newey-West specification of the error term. The coefficient is consistent with the results from the ARIMA regressions; CO₂ emissions are inelastic with respect to GDP.

⁸ As it appears in the graph, the correlation between the two series is higher in the first half of the period (0.714) than in the second half (0.164), though still significantly positive in both.

⁹ The optimal lag lengths and differencing were determined by minimizing the Akaike Information Criterion (AIC) or Bayesian Information Criterion (BIC) (same result). Both series are nonstationary; an augmented Dickey-Fuller test finds strong evidence for unit roots in each series (for log of GDP, the test statistic is -0.033 with a p-value of 0.9557; for log of CO₂ emissions the test statistic is -0.708 with a p-value of 0.8446). Both series are integrated of order one.

¹⁰ In this regression, minimizing the AIC or BIC dictates an ARIMA (1,1,2) regression.

The HP filter is not the only method for detrending time series. Three additional filters are applied to the data, and regression results on the detrended data are presented in Columns 4-6. The Baxter-King filter is a bandpass moving average filter that filters out the trend as well as higher frequency components. The minimum and maximum periodicities are set at 18 and 96 months, respectively (Baxter and King 1999). The Christiano-Fitzgerald filter is a random walk filter, and the same minimum and maximum periodicities are used as in the Baxter-King filter. Finally, the digital Butterworth filter is a rational square wave filter. In regressions presented in Columns 4-6, I detrend the GDP and CO₂ emissions series with each of the three filters and run the same regression as in Column 3.

The purpose of these additional regressions is not to make any claim about which filtering method is preferable. Rather, I seek to demonstrate that the key result found in the first three columns of Table 1 is robust to a wide range of filtering methods. In fact, this is what I find, as can be seen from the regression coefficients in Table 1. The coefficient on the deviation from trend in GDP is consistently positive and between 0.5 and 0.9. The result that emissions are procyclical but inelastic with respect to GDP is thus quite robust. The elasticity is significantly positive in all columns and significantly less than one in all columns except the third and sixth. In addition, analogous regressions on data that are aggregated to the quarterly level yield results almost identical to those presented in Table 1. The results are also robust to varying the lag lengths in the ARIMA regressions and varying the smoothing parameter(s) in each of the filters, and all regressions are robust to including world oil spot prices and seasonally adjusted temperatures as exogenous regressors.¹¹ Similar results are found from panel data regressions on state-level annual emissions (presented in the Appendix).

II. Model – Centralized Economy

The static model presented in the Appendix formalizes intuition regarding two conflicting effects that push for more and less abatement during periods of high productivity, but it fails to account for dynamic considerations. At least three such considerations are built into the following dynamic model, corresponding to three state variables. First, the productivity shock a_t is likely to be autocorrelated. This autocorrelation of factor productivity shocks is in fact what drives RBC models. Thus, a high value of the shock in one period also serves as a signal about

¹¹ Killian (2009) and Balke et. al. (2010) consider the endogeneity of oil prices to economic activity.

the likely shock in subsequent periods. Second, capital k in the economy is a stock. Choosing to save more during one period leads to more resources available for consumption or abatement in subsequent periods, after depreciation. Third, the damages from pollution may come not just from emissions in the current period but from the total stock of emissions. This stock is a function of current and past levels of emissions, subject to a depreciation rate unique to the pollutant. For example, CO₂ has a half-life in the atmosphere of several decades. Ground-level SO₂ has a half-life of only a few days, so for the purposes of a quarterly business cycle model, it can be considered purely as a flow. Furthermore, for a global pollutant like CO₂, the stock is determined not just by domestic emissions but also by emissions from the rest of the world.

In this section I consider a centralized economy; that is, I model the social planner's problem. In the standard RBC model, this choice of modeling is justified by the fact that the economy lacks externalities and thus satisfies the first fundamental theorem of welfare economics, so the solution to the planner's problem coincides with the allocations of the competitive equilibrium. Here, the first fundamental theorem is not generally expected to be satisfied because of the externalities imposed by pollution. In this section, I present a model of optimal policy, in the case where a central authority can select investment, abatement, and consumption for the representative consumer. Later, I model a decentralized economy, where government can attempt to fix the inefficiencies associated with the pollution externality through a policy, such as an emissions tax or tradable permits. The model allows for exogenous growth in technology, consumption, or output, though it is not explicitly modeled. As with other RBC models, this model can be derived from a model that includes a constant growth rate; see the Appendix for an example of such an extension.¹² The assumption that technology is invariant makes the model stationary and allows analysis of fluctuations about the steady state.

Consider a representative agent, whose choices match those of a social planner, with access to production technology yielding potential output $a_t f(k_{t-1})$, a function of a current productivity shock a_t and capital carried over from last period, k_{t-1} . Output is affected negatively from the stock of pollution in period t , x_t . The total output y_t is $y_t = (1-d(x_t)) a_t f(k_{t-1})$, where d is an increasing function that takes values between 0 and 1 and represents the loss

¹² See King et. al. (1988). Newell and Pizer (2003) incorporate both growth and cyclical effects in their reduced form function for emissions abatement costs.

of potential output from pollution.¹³ In each period, the agent chooses quantities of consumption c_t , abatement z_t , and investment i_t , subject to the resource constraint determined by current production: $c_t + z_t + i_t \leq y_t$. Note that just one input is used in production. Labor is not included in the model, since this paper is not concerned with employment fluctuations.¹⁴ Income is not exogenous; it is determined by the agent's choice of investment in the previous period along with the current productivity.

The agent's utility function is defined over consumption: $U(c_t)$. The pollution stock decays at a linear rate equal to η : $x_t = \eta x_{t-1} + e_t + e_t^{row}$, where e_t is current-period domestic emissions and e_t^{row} is current-period emissions from the rest of the world. The policy-maker cannot choose the level of emissions from the rest of the world but can choose the level of domestic emissions. Domestic emissions are a function of total production y_t and abatement. Let $e_t = (1-\mu_t) \cdot h(y_t)$, where $\mu_t \in [0,1]$ is the fraction of emissions abated in period t , and h is the function determining how emissions are related to output, holding constant abatement. Emissions could thus increase more rapidly than output if h is convex, or they could increase less rapidly if h is concave. The fraction of emissions abated μ_t is determined by total abatement spending z_t via the equation $z_t/y_t = g(\mu_t)$; g thus relates the fraction of emissions abated to the fraction of output spent on abatement (the price of abatement is normalized to one). The stock of capital evolves with a decay rate δ : $k_t = (1-\delta)k_{t-1} + i_t$. Finally, the technology shock evolves according to a Markov process, so that the probability distribution of a_{t+1} is a function of a_t .

At the beginning of a period, k_{t-1} and x_{t-1} are already determined, and the realization of the productivity shock a_t occurs. Given those three state variables, the representative agent chooses abatement, consumption, and investment to maximize total expected discounted utility, subject to the constraints described above. This can be written as a dynamic programming problem in the following way.

$$V(k_{t-1}, x_{t-1}, a_t) = \max_{z_t, i_t, c_t} [U(c_t) + \beta E_t V(k_t, x_t, a_{t+1})], \text{ such that}$$

¹³ An alternative modeling choice is to include pollution as directly affecting the utility function, rather than production. This alternative may be more appropriate for conventional pollutants that directly affect health. Carbon dioxide and other greenhouse gases, by contrast, are expected to affect the production possibilities of the world economy (Nordhaus 2008). An earlier version of this paper, which employed this alternative modeling choice but calibrated the model to carbon dioxide emissions, yielded qualitatively similar results (available upon request).

¹⁴ Fischer and Springborn (2009) consider labor fluctuations in a DSGE model with productivity shocks and emissions policy.

$$\begin{aligned}
c_t + i_t + z_t &\leq y_t \\
k_t &= (1 - \delta)k_{t-1} + i_t \\
x_t &= \eta x_{t-1} + e_t + e^{row}_t \\
e_t &= (1 - \mu_t)h(y_t) \\
z_t &= g(\mu_t)y_t \\
y_t &= (1 - d(x_t))a_t f(k_{t-1})
\end{aligned}$$

The operator E_t represents the expectation of future values of a_{t+1} at period t , and the discount factor is β . The problem can be rewritten, taking into account the constraints, as a choice over only the non-stochastic variables k_t and x_t . Assuming the appropriate transversality conditions,¹⁵ first order conditions can be found for the choice of both variables. They are:

$$\begin{aligned}
0 &= -U'(c_t) \\
&+ \beta E_t U'(c_{t+1}) \cdot \{(1 - d(x_{t+1}))a_{t+1}f'(k_t)[1 - g(\mu_{t+1}) - y_{t+1}g'(\mu_{t+1})\left(\frac{1 - \mu_{t+1}}{h(y_{t+1})}\right)h'(y_{t+1})] + (1 - \delta)\} \\
0 &= U'(c_t) \left\{ -(1 - g(\mu_t))d'(x_t)a_t f(k_{t-1}) + y_t g'(\mu_t) \frac{h(y_t) + e_t h'(y_t) d'(x_t) a_t f(k_{t-1})}{h(y_t)^2} \right\} \\
&- \beta E_t U'(c_{t+1}) g'(\mu_{t+1}) \cdot \frac{\eta}{h(y_{t+1})}
\end{aligned}$$

The first equation is the first order condition for the choice of k_t , which is equivalent to the choice of investment this period i_t . The first term, $-U'$, is the marginal cost of an additional unit of investment, which is the foregone marginal benefit of an additional unit of consumption. The second (long) term is the marginal benefit of an additional unit of investment. It is not realized until next period, so it is discounted and taken as an expectation. It equals the marginal benefit of an additional unit of consumption (U') times the marginal effect that investment has on next period's consumption. This marginal effect is composed of the terms in the curly brackets in the second line of the equation: the gross return on capital, net of depreciation, plus the marginal environmental damage caused by the increase in capital.

The second equation above is the first order condition for the choice of the pollution stock x_t . The first part, on the top line, is the marginal effect on current consumption from an additional unit of x_t . It is composed of a negative term (the first half of the expression in the curly brackets), because more pollution reduces output and thus available consumption, and a

¹⁵ $\lim_{t \rightarrow \infty} \beta^t F_x(x_t, x_{t+1})x_t = 0$, where F is the agent's utility as a function of state variables x and F_x is the derivative with respect to any of the current period state variables x_t .

positive term (the second half of the expression in the curly brackets), because more pollution means less was spent on abatement and thus more available for consumption. The second part of the first order condition, on the second line, is the marginal effect on next period's consumption. It is negative, since the additional unit of x_t this period increases the amount of abatement necessary next period to achieve an equivalent level of emissions next period, thus reducing the budget for consumption next period.

III. Calibration

The sources for the calibration of the model fall into two main groups: macroeconomic parameters are taken from the RBC literature, and parameters related to emissions are taken from several studies that estimate the costs and benefits of emissions reductions. One additional parameter is estimated from the empirical results presented above. The model is calibrated to the US economy, and the pollutant considered is carbon dioxide, the leading greenhouse gas.

The model is calibrated, not estimated. Estimation of DSGE models typically relies on the first fundamental theorem of welfare economics to ensure that the optimal allocation, from the model, coincides with the competitive outcome, found in the data. Here, as mentioned earlier, this theorem cannot be invoked, since the pollution externality violates an assumption necessary for the theorem. Therefore, I calibrate the model using previously estimated parameter values where available, and I estimate one parameter that is not available.

The macroeconomic parameters include a functional form and parameterization for the production function, the process governing the productivity shocks, the capital decay rate, and the discount factor. All of these parameters are commonly used in RBC papers, and it is to those papers that I turn for their calibration. Specifically, the values I use here are those used in recent RBC literature, including Chang and Kim (2007), and pioneering papers, including Kydland and Prescott (1982).¹⁶

The production function is $f(k) = k^\alpha$, where $0 < \alpha < 1$ to accommodate positive but diminishing marginal returns. In most RBC models, two inputs to production are present: capital and labor. Then, the production function is Cobb-Douglas: $f(k,l) = k^\alpha l^{1-\alpha}$. The capital share of income is taken as the value α . Here, production has only one input, since I am unconcerned

¹⁶ These papers do not model carbon emissions, but that does not invalidate using their macroeconomic parameters here since firms and households in the economy were not responding to any carbon policy. See Bartz and Kelly (2008) for a calibration of a deterministic model similar to the stochastic model here.

with the labor market. In other words, the model here assumes fixed labor inputs and models only the response of capital inputs to production.¹⁷ The value used for α is 0.36, capital's share in national income.

Each period of time is one quarter. Capital depreciation δ is 2.5%, and the discount factor β is 0.98627, consistent with a quarterly rate of return to capital of 1 percent. This is equivalent to an annual discount rate of about 5% and annual capital depreciation of about 9.6%. Finally, the productivity shock evolves according to a Markov process with the following equation: $\ln a_t = \rho \ln a_{t-1} + \varepsilon_t$, where the persistence parameter ρ is 0.95, and ε_t is an IID shock distributed normally, with mean zero and standard deviation $\sigma_\varepsilon = 0.007$. The agent's utility function is isoelastic: $U(c) = \frac{c^{1-\varphi_c}}{1-\varphi_c}$. The coefficient of relative risk aversion φ_c is set to equal 2 (Stern 2008, Weitzman 2007).¹⁸

Next I use chemistry and economics studies to calibrate the functions regarding emissions. The decay equation for the stock of pollution in the atmosphere, $x_t = \eta x_{t-1} + e_t + e^{row}_t$, contains the parameter η , which can be calibrated from the half life of atmospheric carbon dioxide. This value is not precisely known, and various papers use different estimates.¹⁹ For my base case, I use 83 years, the value used in Reilly (1992). In sensitivity analysis, I see how this parameter affects the results. A half life of 83 years implies a quarterly parameter η of 0.9979. I assume that rest-of-world emissions e^{row}_t are maintained at a constant level e^{row} ; international emissions do not respond to domestic business cycles. The US is responsible for about one-fourth of global anthropogenic carbon emissions, so e^{row} is set at three times the steady-state value of e .

The damages from atmospheric carbon dioxide have been estimated in other papers, including in Nordhaus (1991, 2008) and Stern (2008). I calibrate $d(x)$ from the DICE-2007 model in Nordhaus (2008) as described in the Appendix. Fitting the function to a quadratic $d(x) = d_2x^2 + d_1x + d_0$ yields $d_2 = 1.4647e-8$, $d_1 = -6.6722e-6$, and $d_0 = 1.3950e-3$, where the units of x are gigatons of atmospheric carbon (GtC). For the 2005 value of atmospheric carbon

¹⁷ See Fischer and Springborn (2009) for an RBC model of pollution policy that incorporates endogenous labor.

¹⁸ Kelly (2005) finds, in a similar though static model, that the value of the coefficient of relative risk aversion affects the choice of quotas vs. taxes. Under his modeling assumptions, a quota dominates when the coefficient exceeds one-half.

¹⁹ Falk and Mendelsohn (1993) and Nordhaus (1991) use a decay rate implying a half life of 139 years. Reilly (1992) assumes a half life of 83 years. Moore and Braswell (1994) estimate the half life of atmospheric CO₂ under a range of different assumptions, and consequently find a range of answers, from 19 to 92 years.

mass of about 800 gigatons, this creates an output loss of 0.26%. A doubling of the carbon mass to 1600 GtC creates an output loss of 2.54%.²⁰

The abatement cost function $g(\mu_t)$ gives the cost, as a ratio of total output, of reducing the fraction μ_t of baseline emissions. This function is taken directly from Nordhaus (2008). The function form used is $g(\mu) = \theta_1 \mu^{\theta_2}$. The calibrated value of the exponent θ_2 is 2.8, indicating a convex cost function. The coefficient, θ_1 , is actually a function of time in Nordhaus (2008), where each time period is ten years. I use the initial value (calibrated to 2005) of 0.05607, though this value decreases over time slightly (it drops to 0.0392 in 50 years).²¹

Finally, the function mapping output to emissions controlling for abatement, h , is estimated from the data described in section I. I impose an iso-elastic form to the function, so that $e_t = (1-\mu_t)h(y_t)$ and $h(y_t) = y_t^{1-\gamma}$. The parameter γ is calibrated based on an estimation of the log of emissions on the log of output, where the regression coefficient is $1-\gamma$. The assumption behind this calibration is that μ_t is constant and unaffected by y_t , which is unlikely to fail given that no climate policy existed during the period examined. Regression results are presented in Table 1. The coefficient on the log of GDP ranges from 0.55 to 0.88. For the base case value of γ , I use the results from the initial ARIMA regression on the seasonally adjusted series from column 2 of Table 1, giving an exponent for h of $1-\gamma = 0.696$. I vary this parameter in sensitivity analysis.

Table 2 describes the parameters in the model, gives their calibrated base case values, and lists the source of each parameter calibration.

IV. Model Solution and Simulation

I solve the model by log-linearizing about the steady state and analytically solving the system of linear rational expectations equations. This solution method is fast (taking less than one second on a typical PC) and it removes the need for discretizing or approximating. However, it comes as a cost: the non-linear model is approximated as a system of linear equations. This will not be a large problem here since I focus on small fluctuations about the

²⁰ Because d is defined over gigatons of carbon and the model yields pollution stock in arbitrary units, these coefficients are scaled to keep the proportional output loss consistent.

²¹ See the Appendix for an extension to the model allowing for growth in abatement technology.

steady state.²² The system of linear equations is solved using the Anderson-Moore algorithm (AMA).²³ The system of linear equations and the method for implementing the AMA are presented in the Appendix. The Matlab code is available on the author's website.²⁴

The solution provides a set of matrices that describes how the choice variables (c_t, k_t, x_t) respond to different values of the state variables (x_{t-1}, k_{t-1}, a_t) . This solution is difficult to interpret, but the results can be summarized graphically in two ways. First, one can examine impulse response functions: given a non-zero value of ε_t in period $t = 0$ and no non-zero values in any other periods, what is the response path of the choice variables? Second, one can simulate a series of shocks ε_t and analyze the response of variables to those shocks. This simulates the actual business cycles in the economy and how policy can optimally respond.

Base Case Simulation

Impulse response functions for the base case parameters are plotted in Figures 3 and 4. Both figures come from the same simulation of 100 periods, but they plot different variables (except for the productivity shock a , which is plotted in both figures). In the simulation, an innovation to the productivity shock ε_t occurs in period 1. The size of this shock is 1%, about one-and-a-half times σ_ε , the standard deviation of the innovation in the calibrated model. Figure 3 plots the proportional deviation from the steady state value of the productivity shock a along with economic variables traditionally seen in RBC models: output y , capital stock k and consumption c . Figure 4 repeats the plot the productivity shock a , while adding the response of three variables related to pollution: abatement z , single-period emissions e and pollution stock x . The path of the productivity shock in both of these figures is exogenous; it decays at the rate $\rho = 0.95$. The value of output y is also decreasing along with productivity. Output is not perfectly coincidental with productivity, since the choice between allocating resources between savings and consumption is altered by the productivity shock. If the capital stock and pollution

²² See the conclusion for a discussion of the interpretation of these results given that the economy is likely not on a steady state growth path when it comes to greenhouse gas emissions.

²³ Anderson and Moore (1985). I use a Matlab application provided by the Boston Fed (<http://www.bos.frb.org/economic/econbios/fuhrer/matlab.htm>). Anderson (2008) compares the practicality of several different methods for solving linear rational expectations models and finds that the AMA provides the highest accuracy and significant gains in computational efficiency.

²⁴ I also solve the model without linearizing by using value function iteration (VFI). The solutions under this method are nearly identical to those from the AMA, suggesting that the linearization does not introduce appreciable error, given that I am looking only at small deviations. Later, when I consider a growth model that allows for an emissions elasticity that changes with income, I use VFI. VFI code is also available online.

stock were kept fixed at the steady state value, then the a and y curves would be coincidental. Figure 3 shows, though, that capital is responding to the increased productivity. More is being invested, and thus capital is higher than it is in the steady state. Because capital is a stock good, its peak does not coincide with the productivity peak in period 1. Capital peaks around period 40 (year 6). Similarly, consumption is higher thanks to the increased productivity, but its peak occurs around period 15. The lag between the peak of productivity and the peak of consumption is not because consumption is a stock; it is not. Rather, it is a function of the resource constraint in each period, which is dependent on capital, a stock.

Figure 4 shows how optimal emissions policy responds to the productivity shock. The path of productivity is included again in Figure 4 for comparison. Note that the y-axis scale is increased. The increased productivity induces an increase in abatement expenditure z . Although more is being spent on abatement, the increase in output means that emissions will be higher for a fixed level of abatement. It is thus unclear in which direction emissions will go, but Figure 4 shows that emissions e are higher after the productivity shock. This demonstrates the fact that, in the base case calibration of the model, the price effect dominates the income effect. Higher productivity yields more income, and the absence of pollution is a normal good, so less emissions will be demanded. On the other hand, abatement is costlier with high productivity, so the price effect causes more emissions to be demanded.²⁵ Finally, note that the curve for pollution x exhibits two striking features: it is of a much smaller magnitude than the other curves, and it is more persistent. Even after 100 periods it has not quite yet reached its peak. Both of these features are due to the fact that the decay rate for CO₂ is so low ($\eta = 0.9979$), so a change in any one year's emissions has little effect on the total stock, and deviations from the steady state of emissions take a long time to decay.

What do these impulse response functions say about how optimal policy responds to productivity shocks, for the base case given by these parameters? Three qualitative lessons can be learned. First, optimal emissions are procyclical; that is, periods of higher productivity are allowed higher emissions. Second, to a first order, a policy that pegs allowable emissions to GDP is optimal. Comparing the curve for y in Figure 3 to the curve for e in Figure 4, it

²⁵ The income effect being dominated is unsurprising given that it comes from a temporary though persistent shock to income, which is small compared to permanent income. A two standard deviation positive shock to productivity increases total discounted output by less than one-hundredth of a percent, compared to total discounted output with no positive shock.

appears that these two values are almost identical in shape.²⁶ Third, the magnitudes of the adjustments to the pollution stock are quite small, because of the low decay rate of emissions.

Next, I simulate a draw of productivity shock innovations ε_t from its calibrated distribution and see how the economy responds to that series of shocks. This is a simulation of an actual RBC economy, as opposed to the impulse response functions, which are abstractions designed to see the qualitative nature of responses. These simulations cannot explicitly show how individual policy variables respond to particular productivity shocks, but they can provide summary statistics of the nature of the business cycles under the optimal solution, such as the standard deviations of output and emissions, or the covariance of output and emissions.

Each simulation is of course dependent on the particular draw of shocks. Figure 5 presents the results from one such draw. The curve marked with squares is output. For this draw of shocks, after an initial 20 periods of fluctuations near zero, the economy experiences a long expansion followed by a long recession. The curve marked with circles is emissions; it closely follows output. This conforms with the impulse response functions, where the curves for k and e were close to each other. The curve marked with diamonds in Figure 5 is the capital stock. The lag between the peaks of capital and the peaks of output is consistent with the lagged peak found for capital in the impulse response function; it is because capital is stock that accumulates over time. Finally, the curve marked with triangles is the pollution stock. It appears almost flat, though it actually fluctuates somewhat. As in the impulse response functions, its small magnitude is due to its very low decay rate.

Finally, the base case simulation results can be used to conduct welfare analysis. In particular, how valuable is it to society to enact the dynamically optimal policy rather than proceed with a static policy that fixes the level of emissions at the steady state value? To answer this, I simulate the dynamically optimal economy as well as a sub-optimal economy where emissions are kept constant at the steady-state level.²⁷ The improvement that the dynamically optimal policy provides over the static policy is the compensating variation between the two policies, defined in terms of a consumption metric. This compensating variation is defined as the

²⁶ Setting an aggregate cap on emissions as a function of aggregate output differs from allocating permits across firms based on firm-level output. The latter, output-based allocation schemes, can create distortions (Bushnell and Chen 2009), mitigate pre-existing distortions (Fischer and Fox 2007), or affect distributional outcomes (Fullerton and Heutel 2010).

²⁷ Because I do not have to linearize the objective function in these simulations, I avoid the "pitfalls" of welfare analysis identified in Kim and Kim (2007).

level of consumption c_v such that the representative agent is indifferent between the dynamically optimal policy and the static policy with an additional c_v units of consumption given to him in each period.²⁸ The compensating variation is a function of the particular draw of shocks, so I run 250 simulations and use the median value found. When calibrated to total level of domestic consumption expenditures, the compensating variation is about \$950 million per year (in 2005 US\$). This is a comparable magnitude to the estimated net benefits of other environmental policies. For example, Carlson et. al. (2000) estimate that the cost savings from the U.S.'s sulfur dioxide cap-and-trade program compared to a command-and-control program is \$700-\$800 million annually. Oates et. al. (1989) find net benefits from an improvement in emissions standards for the pollutant total suspended particulates (TSP) in Baltimore, MD, on the order of \$20-\$50 million. For comparison, Nordhaus's (2008) calculation of the net benefits from climate policy that is optimal in the long-run but does not accommodate business cycles implies annual domestic net benefits of approximately \$7 billion.²⁹ As mentioned in the introduction, although small as a fraction of total GDP, these net benefits are likely to be non-uniformly distributed across the economy, such that for some households (especially the poorest) and some firms (especially electric utilities) the cost savings from the optimal policy are quite substantial.

Sensitivity Analysis

The results presented thus far are taken from simulations using the base case parameter values, listed in Table 3. Sensitivity analysis can be done by varying these parameters and seeing the effect on policy. I first perform sensitivity analysis on γ , the parameter that defines the elasticity of emissions with respect to output when abatement is held constant. This parameter was estimated, and its base case value was 0.696 (Table 1). The estimate from Table 1, though, varies across columns. Although all columns of that table find that emissions are inelastic, in Figures 6 and 7, I vary $1 - \gamma$, the elasticity, from 0.25 to 1.2. When $1 - \gamma > 1$, emissions are elastic. Figure 6 plots the deviation from steady state in emissions resulting from a

²⁸ The same method is found in Lucas (1987). Tallarini (2000) finds higher proportional welfare differences with increased risk aversion.

²⁹ This back-of-the-envelope calculation comes from Nordhaus's (2008) reported global net-present-value benefit of \$3 trillion. Allocating the benefits uniformly across all years, at an annual discount rate of 5% this implies annual global net benefits of \$150 billion. Allocating the U.S.'s share of those benefits proportional to population (4.6% of world population) yields the reported figure.

one-percent positive productivity shock under the different values of γ . Higher values of $1 - \gamma$ imply a higher deviation from steady state of emissions in response to the shock. Even if the change in output was no different between the five simulations shown in Figure 6, these results are expected, since the parameter that is varied gives the response of emissions to changes in output. Optimal policy dampens the cyclicity of emissions at a higher rate the higher is the elasticity of emissions with respect to output.

In Figure 7, I plot the effect of parameter value deviations on μ , the fraction of output that is spent on abatement. Here the impact of γ qualitatively can be seen. When the elasticity $1 - \gamma$ is 0.5 or 0.25, μ decreases after a positive productivity shock. This differs from when $1 - \gamma$ is at its base case value of 0.696 or when it is 1 or greater; in those cases μ increases after an increase in productivity. If $1 - \gamma \leq 0.5$ and the economy experiences a positive productivity shock, the optimal response is to reduce the fraction of emissions abated. Note, however, that emissions rise after the shock under all of the alternative values of γ . When emissions are very inelastic to output ($1 - \gamma$ is low), then an increase in output caused by a positive productivity shock results in a relatively low emissions increase, holding abatement constant. The optimal response of abatement spending, then, will be relatively low. When emissions are inelastic enough ($1 - \gamma \leq 0.5$), then the optimal response is to actually reduce abatement spending as a fraction of total output.

Each of the simulations represented in Figures 6 and 7 assumes that the elasticity parameter is constant throughout a simulation. Some evidence suggests that this elasticity can vary with total output. Annual CO₂ emissions data from the U.S. can be used to show that responses of emissions to business cycle shocks were much greater in the distant past (1870-1938) than more recently (1939-2007).³⁰ Accommodating this feature of the economy is not possible in the above solution method, which assumes that deviations about any growth path are identical to deviations about the steady state. Instead, to investigate the effect of a changing elasticity, I solve the model using value function iteration (VFI) and simulate fluctuations about a growth path. I allow the elasticity of emissions with respect to output to decrease with output, and then examine the optimal emissions path of fluctuations about the growth path. The results

³⁰ These data are available at http://cdiac.ornl.gov/trends/emis/tre_coun.html. I thank an anonymous referee for providing the calculations. Similarly, one may worry that the emissions elasticity varies by nation. Gregg et. al. (2008) finds this comparing CO₂ emissions from China and the US. This paper, however, focuses on domestic emissions and policy.

along the growth path (available upon request) reinforce the intuition behind the results presented in Figure 6. In earlier, poorer periods, when the emissions elasticity is larger, the optimal response of emissions to productivity shocks is larger, compared to later, richer periods, when the emissions elasticity is smaller.

The next parameter that is varied, in Figures 8 and 9, is the coefficient of relative risk aversion φ_c . The base case version is calibrated to 2. In Figures 8 and 9 I allow this value to be 0.5 (low risk-aversion), 1 (log utility), or 3 (high risk aversion), in addition to the base case value. Kelly (2005) finds, in a model similar to the one here but static, that the relative risk aversion of the representative agent has a substantial impact on the choice between price and quantity instruments. This suggests that the same coefficient may impact the tradeoff between the income and price effects in this model. Figure 8 shows that varying this parameter has little impact on the path of emissions. Figure 9 plots the response of consumption. With higher risk aversion, the agent smooths consumption more in response to the positive shock compared to when risk aversion is low.

One parameter of interest is the decay rate of the pollution stock, η . In the base case this value is 0.9979, corresponding to a half-life of 83 years, but this value is uncertain. How sensitive is the optimal policy to this parameter? Figure 10 suggests that the answer is not very sensitive. It plots the curve for e , but it presents the responses under six different values of η , going down all the way to zero, representing a pure flow pollutant like SO₂ or NO_x. Surprisingly, the results are not very sensitive to this parameter. Even when pollution is purely a flow, the optimal emissions path is almost identical to the one in the base case, in which CO₂ has a long half-life.³¹ Figure 11 shows the optimal impulse response functions for the pollution stock. Here, the differences are clear, and they are due to the different decay rates of the pollution stock. Intuitively, one might think that the income effect should have more bite for a flow pollutant ($\eta = 0$) than for a stock pollutant. For a stock pollutant, society does not care much about fluctuations in single-period emissions since they have little effect on the value of the stock. Thus the income effect of cyclical fluctuations in emissions is higher for a flow pollutant, and optimal emissions ought to be lower when society is richer. However, there is a counteracting component of the income effect: for a stock pollutant, the effect of a fluctuation in

³¹ Throughout this analysis, the impulse response functions plotted are in proportional change from the steady-state value, and the steady-state values of the variables differ under each parameter set.

single-period emissions remains longer than for a flow pollutant. Though society cares less in each period about the damages from an increase in emissions for a stock pollutant, it cares about those damages for more periods. As η varies, those two counteracting components of the income effect approximately cancel each other out.

Finally, I investigate the assumption about emissions from the rest of the world, e^{row} . In the base case, this level of emissions is constant and is equal to three times the steady-state level of domestic emissions. I consider two alternate specifications of e^{row} , one in which it is just equal to domestic steady-state emissions and another in which it equals six times domestic steady-state emissions. If the level of pollution from the rest of the world is really high, policy should be very responsive to productivity shocks, since the marginal benefits of abatement will be higher. On the other hand, the impact of domestic policy is less important when e^{row} is very high, and thus there is less need for responsiveness to shocks. On net, the impact of e^{row} on the impulse response functions is not noticeable. Optimal policy is slightly more responsive to shocks when e^{row} is smaller; the response of abatement spending one period after a 1% productivity shock is a 1.42% increase when e^{row} equals domestic emissions and is a 1.38% increase when e^{row} is six times domestic emissions.

Under all parameter values studied, a positive productivity shock causes optimal emissions to increase relative to the steady state. In other words, the price effect dominates the income effect.³² Can the income effect dominate the price effect? The price effect is driven in part by the convexity of the abatement cost function g – this makes it very costly to increase abatement. In a simulation where all parameters are kept at their base case values except $\theta_2 = 0.1$, so that g is concave, the income effect dominates the price effect. After a positive productivity shock, emissions initially rise. But after 40 periods, optimal emissions fall below their steady-state value, suggesting that the price effect dominates in the early periods after a shock, but eventually the income effect catches up. If, in addition to changing θ_2 to 0.1, we also change $1 - \gamma$ to 0.2, the income effect dominates starting just 12 periods after the shock.

All of the simulations come from a shock to total factor productivity, which creates both an income and a price effect. This is a departure from much of the literature that models a reduced form abatement cost function with potential shocks to the slope of that function. In that

³² Figure 7 shows that when the emissions elasticity $(1 - \gamma)$ is low enough, the fraction of output spent on emissions (μ) decreases after a positive productivity shock. But, as Figure 6 shows, the absolute level of emissions still increases.

case, when the value of the shock is known by both firms and regulators, the policy implications are clear since only a price effect arises: when abatement costs more, optimal policy is to abate less. This model can include shocks directly to abatement costs, and it can make those shocks autocorrelated, by including a shock term in the linearized equation for abatement costs (z). When such a shock is simulated in this model, the expected result is reached. During periods of high abatement cost, optimal policy is to abate less. However, the focus of this paper is to extend beyond a model with direct shocks to abatement, and to consider instead aggregate productivity shocks, where the direction of the response of optimal policy is not as clear.³³

V. Decentralized Economy

The model above represents a social planner's problem: the representative agent chooses investment, consumption, and emissions simultaneously to maximize total discounted utility. Thus, all externalities are internalized. If the goal is to find optimal environmental policy in response to correlated productivity shocks, then this model is helpful to find the levels of these variables that would maximize utility in the presence of these shocks. However, the model does not explicitly specify the particular policy that a government can take to achieve these first best results in the presence of the environmental externality. For this problem, I turn to an extension of the model where decision-making is decentralized among competitive firms, utility-maximizing individuals, and a benevolent government. The government can choose an emissions tax levied on firms; the appendix demonstrates how this policy is equivalent to a quantity restriction.³⁴

First consider the behavior of the representative firm.³⁵ It seeks to maximize profits by choosing the appropriate level of capital and abatement. Its profit function is

$$\pi_t = y_t - \tau_t e_t - r_t k_{t-1} - z_t,$$

where τ_t is the emissions tax rate and r_t is the endogenously determined cost of capital. The price of output and the price of abatement are both normalized to one. The firms maximize profit subject to the emissions function $e_t = (1-\mu_t) \cdot h(y_t)$ and the abatement cost function z_t

³³ A potentially interesting extension could consider if the abatement technology shocks were correlated with the TFP shocks, in which case the two effects could either support each other or cancel each other out, depending on the sign of the correlation.

³⁴ These are by no means the only policies that governments could actually enact. Governments may subsidize abatement, tax consumption, output, or capital, etc. The emissions tax and quota, though, directly address the relevant externality.

³⁵ I do not model entry or exit, the equilibrium number of firms, or firm size.

$=y_t g(\mu_t)$. An externality arises because the firm does not take into account its emissions' impact on the pollution stock and on productivity; i.e. it takes x_t as given.³⁶

The firm's profit maximizing behavior sets the marginal value product of capital equal to the rental rate and the marginal value product of abatement equal to its price, normalized to one. Thus,

$$r_t = y_t f'(k_{t-1})/f(k_{t-1})[1 - \tau_t(1-\mu_t)h'(y_t) - g(\mu_t)],$$

$$\tau_t h(y_t) = y_t g'(\mu_t).$$

Next, consider the behavior of the representative consumer. The consumer is the owner of capital and rents it out at the market rate r_t . The emissions tax is determined exogenously and hence taken as given. Consumption and investment are chosen subject to a budget constraint

$$c_t + i_t \leq r_t k_{t-1} + \tau_t e_t + \pi_t.$$

The consumer's income in each period comes from three sources: the rental income from the capital owned by the consumer and rented by the firm at rate r_t , the emissions tax revenues collected by the government and returned to the consumer, and the firm profits, since the firm is owned by the consumer. The consumer does not directly spend anything on abatement; that decision is made by the firm and taken as given by the consumer. The consumer's maximization problem is

$$\max_{\{c_t\}, \{k_t\}} \sum \beta^t U(c_t), \text{ such that}$$

$$c_t \leq r_t k_{t-1} + \tau_t e_t + \pi_t - k_t + (1 - \delta)k_{t-1},$$

where the firm's decisions determine r_t , e_t , and π_t , and thus those are all exogenous to the consumer.

The consumer's problem can be expressed as choosing a path $\{k_t\}$ to maximize total discounted utility. The first order condition for the choice of c_t is

$$-U'(c_t) + \beta E_t U'(c_{t+1}) \cdot [r_{t+1} + (1 - \delta)] = 0$$

Emissions, profits, and the rate of return are all considered exogenous by the consumer. Thus the first order condition above simply sets the marginal cost of an additional unit of investment (foregone consumption this period) equal to its marginal benefit (the expected value of additional consumption next period).

³⁶ In this model with one consumer owning the one firm, the pollution externality may be internalized if the consumer controls the behavior of the firm and departs from profit-maximizing behavior because of pollution; see Gordon (2003).

A benevolent government chooses the optimal tax rate given the behavior of the firm and the consumer. In other words, the government solves a Ramsey problem by choosing $\{\tau_t\}$ to maximize total expected discounted utility, given the constraints from the firm's profit maximizing behavior and the constraints from the consumer's profit maximizing behavior.³⁷ The government considers how its choice of τ_t affects consumer and firm choices and thus social welfare.

The government's problem is

$$\begin{aligned} & \max_{\{c_t\}\{k_t\}\{y_t\}\{\tau_t\}} \sum \beta^t U(c_t) \\ & \text{Such that, for all } t, \\ & -U'(c_t) + \beta E_t U'(c_{t+1}) \cdot [r_{t+1} + (1 - \delta)] = 0 \\ & r_t = y_t f'(k_{t-1}) / f(k_{t-1}) [1 - \tau_t (1 - \mu_t) h'(y_t) - g(\mu_t)] \\ & \tau_t h(y_t) = y_t g'(\mu_t) \\ & c_t = r_t k_{t-1} + \tau_t e_t + \pi_t - k_t + (1 - \delta) k_{t-1} \\ & \pi_t = y_t - \tau_t e_t - r_t k_{t-1} - z_t \\ & e_t = (1 - \mu_t) \cdot h(y_t) \\ & x_t = \eta x_{t-1} + e_t + e^{row}_t \\ & z_t = y_t g(\mu_t) \\ & y_t = (1 - d(x_t)) a f(k_{t-1}). \end{aligned}$$

The series of constraints can be reduced in number by exploiting the fact that many of the constraints and first order conditions provide an implicit mapping of variables. The firm's first order condition for abatement expenditure defines $\mu_t = \mu(\tau_t, y_t)$. Then, because both e_t and z_t depend only on μ_t and y_t , I can use the functions $e_t = e(\tau_t, y_t)$ and $z_t = z(\tau_t, y_t)$. Finally, the firm's first order condition for capital demand defines $r_t = r(\tau_t, y_t, k_{t-1})$. Making these substitutions, along with the resource constraint for c_t , yields a government problem with four constraints per period. The Lagrangian for this problem is:

³⁷ The government is able to dynamically adapt the policy in light of new information on productivity shocks and state variables. This is thus a "feedback" policy, as defined in Hoel and Karp (2002). The alternative is an "open-loop" policy, where the government must choose the entire policy trajectory at the initial period. This other extreme, as well as other policy options along the spectrum between these two extremes, is considered later, when I allow for information asymmetry between regulators and firms/consumers. I assume that the government can commit to future tax rates; see Kydland and Prescott (1977). The government's optimal policy can be thought of as a Taylor rule for emissions policy.

$$\begin{aligned}
& \sum \beta^t U(y_t - k_t + (1 - \delta)k_{t-1} - z(\tau_t, y_t)) \\
& + \sum \lambda_t [-U'(y_t - k_t + (1 - \delta)k_{t-1} - z(\tau_t, y_t)) \\
& \quad + \beta U'(y_{t+1} - k_{t+1} + (1 - \delta)k_t - z(\tau_{t+1}, y_{t+1})) \cdot (r(\tau_{t+1}, y_{t+1}, k_t) + 1 - \delta)] \\
& + \sum \zeta_t [x_t - \eta x_{t-1} + e_t^{row} + e(\tau_t, y_t)] \\
& + \sum \omega_t [y_t - (1 - d(x_t))a_t f(k_{t-1})]
\end{aligned}$$

The first constraint is from the consumer's first order condition, the second is the pollution evolution equation, and the third is the production function. The Lagrangian multipliers for each constraint in period t are λ_t , ζ_t , and ω_t , respectively. The first order equations for this problem are presented in the Appendix. The solution can be found computationally by log-linearizing the system and using the AMA.

Figure 12 presents simulation results from a decentralized economy. Parameters are kept at the base case values, and the economy faces the same set of productivity shocks as in Figure 5. Figure 12 shows the cyclical components of output y , emissions e , and the emissions tax τ for a 100 period simulation. As the Appendix shows, any outcome under the tax policy can also be achieved by a regulator who has access to a quantity policy, for instance a cap-and-trade scheme. This equivalence holds when the regulator has symmetric information about all state variables; later I consider information asymmetry. Thus the simulation in Figure 12 can also be interpreted as modeling a quantity policy where the quota is optimally adjusted; emissions e in Figure 12 would just be set equal to the chosen level of the quota. As in the centralized economy, the optimal level of emissions closely follows output. During expansions (such as the long one that occurs between about periods 20 and 60), the optimal response of emissions rises, because abatement is relatively more expensive when productivity is higher.

However, this is not achieved by reducing the emissions tax during expansions, as might be intuited. The emissions tax is also procyclical. How are higher emissions achieved in expansionary periods from higher emissions taxes? Because productivity shocks affect abatement costs, the marginal cost of abatement is higher during expansions. Also, the stock of capital responds positively to productivity shocks, and a higher level of capital increases baseline emissions. Without an increase in the emissions tax during expansions, emissions would increase more than optimally during these periods, and a countercyclical tax policy would exacerbate the problem. The simulation under the centralized economy shows that optimal emissions policy is procyclical, with emissions rising during expansions. However, the tax

policy simulation shows that in an important way the policy is countercyclical, since the emissions tax rate increasing during expansions means that emissions are lower than they otherwise would be during expansions. The optimal path of emissions is procyclical, but not as procyclical as would hold without policy responding optimally to shocks.

The following three observations can be garnered from this simple extension of the model to a decentralized economy. First, although emissions follow output, rising during expansions and falling during recessions, this is achieved with a tax rate that rises during expansions and falls during recessions. Because abatement is costlier during expansions, the emissions tax must rise to keep emissions from overshooting their optimal trajectory. Second, the volatilities of the emissions tax and emissions itself are about equal. One may think that an emissions tax would have to vary less with productivity fluctuations than would emissions, though this turns out to not be the case.³⁸ Third and finally, a potentially interesting political economy observation can be made. During recessions it may be politically difficult or infeasible to get an emissions policy strengthened; consumers and producers will likely lobby against anything that may increase costs. If the policy is an emissions quota, then optimal policy is strengthened during recessions (the quota is reduced). However, if the policy is an emissions tax, it is weakened during recessions (the tax is reduced). Weakening policy during recessions may be more politically feasible than strengthening policy, perhaps suggesting an advantage of taxes over quotas in their ability to respond to productivity fluctuations.³⁹ As mentioned in the introduction, the permit price in the initial RGGI auction was below expectations. The analysis here suggests that the cap should have been lowered because of the recession. If the policy had instead been a tax, it may have been politically easier to lower the tax rate (optimally) during the recession.

These observations must be made with an important caveat. While information asymmetry between regulators and consumers or firms is the key component of the "prices

³⁸ Intuition might prescribe that, since the pollution stock is not varying by much, the marginal damages from emissions are not varying by much, and thus the emissions tax should not vary by much, compared to the quantity restriction. However, the government's first order condition for the choice of the emissions tax equates marginal costs with marginal benefits, where marginal costs are from foregone consumption. Since consumption is changing with the business cycle, the marginal cost of the emissions tax is also changing. It is this variance in consumption, not in the pollution stock, that leads to the variance in the emissions tax.

³⁹ Economically, comparing the tax rate with the quota quantity is an apples-to-oranges comparison, but politically these are the two values that are likely to be compared to each other. The argument here is that, *once in place*, quotas are hard to reduce in a recession. But others have argued that quotas are more politically feasible than taxes (Stavins 2008).

versus quantities" literature pioneered by Weitzman (1974), the specification here has no such information asymmetry. Firms, consumers, and the government all face uncertainty about future values of productivity, but each agent's information set is identical: they observe lagged values of the stock variables k_{t-1} and x_{t-1} and the current productivity shock a_t . To the extent that the simulation results here suggest that prices may be preferable to quotas due to the fact that the optimal tax drops during recessions and rises during expansions, this omits considerations of information asymmetries. Previous studies find that those considerations tend to support taxes over quotas as the optimal instrument for CO₂ (e.g. Newell and Pizer 2003); this study suggests an alternate reason for the same conclusion.

The model can be extended to consider the effects of information asymmetry on optimal policy in a decentralized economy. If government has to set its policy each period before it observes the state, then the government's first-order conditions will not hold. Instead, the government may choose the policy based on whatever information set it has available. For example, the tax chosen in period t may be a function of the value of the productivity shock observed in period $t-1$, or of the value of output in period $t-1$, or a combination of both. In this case of limited information, the government's first order conditions in the model are replaced by an equation relating the current emissions tax to the specified variables. The resulting system, since it is still linear in proportional deviations from steady state values, can be solved with AMA, and the dynamics can be evaluated.

Given this information constraint placed on the government, I can solve for the second-best policy option. But with the information asymmetry, the tax policy and a quota policy are no longer equivalent. I can solve separately for the second-best policy using either instrument. For example, in the case of a tax policy, suppose that the government is constrained to set the tax in period t , in proportional deviation from its steady-state value, as a linear function of the deviation of output in period $t-1$.⁴⁰ The policy is $\tau_t = d y_{t-1}$, where d is a policy parameter representing how the tax should respond to lagged output. The optimal policy is found by choosing the value of d that minimizes the difference between the impulse response function of consumption under the unconstrained and constrained scenarios. For the tax case, this value is 1.415; the optimal tax in period t should be about 142% of the deviation from steady-state output in period $t-1$. For a quantity policy with the same information constraint, the

⁴⁰ That is, the regulators observes neither a_t , k_{t-1} , nor x_{t-1} directly.

government's policy is $q_t = d \cdot y_{t-1}$, and the optimal value for d is 0.658. As in the optimal solution with no information asymmetry, both the optimal tax and the optimal policy are procyclical. The optimal quantity policy varies significantly less than does the optimal tax policy. Though these second-best policy rules are obviously inferior to the first-best rules found in the centralized model, they may be the most relevant policy prescriptions of the model, since indexing emissions targets to GDP has been suggested as a viable policy option (Newell and Pizer 2008, Pizer 2005).⁴¹

VI. Conclusion

Recessions can motivate policymakers to enact significant policy changes, or they can prevent them from doing so; the failure of a prominent climate bill in the US Senate in June 2008 may be attributed in part to fears of increased energy costs in the face of an upcoming recession. To answer the question posed in this paper's title, a dynamic stochastic general equilibrium real business cycle model that features damages from the stock pollutant carbon dioxide is calibrated and solved. Climate change abatement is costly, and since damages are roughly smooth over the business cycle, it makes sense to smooth costs out over the business cycle. A first-best solution to the social planner's problem finds that the optimal level of emissions increases with productivity. Thus, a quantity policy is relaxed during economic expansions and tightened during recessions. This result is attributed to an income effect, in which economic expansions create a higher demand for clean air, being outweighed by a price effect, in which achieving a particular level of emissions is costlier during an economic expansion because of increased productivity. This result is robust to various parameterizations of the model. Finally, I model a decentralized economy, where firms maximize profits, consumers maximize utility, and the government chooses a tax to maximize social welfare. This model shows how asymmetric information affects optimal policy and the choice between prices and quantities.

The response of policy to business cycle fluctuations may be second-order relative to questions about secular trends in emissions, abatement, and climate change. However, it is still important. The compensating variation (\$950 million per year) is modest but not negligible, especially considering that the cost savings from the optimal policy are not uniformly distributed

⁴¹ In principle one could calculate the welfare difference between the second-best linear tax policy and the second-best linear quota policy. However, in practice the welfare ordering between the policies is dependent on the draw of shocks.

across the economy – poorer households and carbon-intensive firms are more likely to benefit more than proportionally. The compensating variation calculation does not compare the dynamically optimal policy to no policy, but it compares the dynamically optimal policy to the best static policy. Given a cap-and-trade scheme, the cost of adjusting the cap with the business cycle is likely low relative to its benefit. A more complete analysis would incorporate both secular and cyclical features. Since virtually all prior research incorporates only secular trends, first considering only cyclical effects is a good place to start.

The model solves for a steady state and focuses on deviations about that equilibrium. Is the economy at a steady state regarding carbon policy? Though the basic model here does not consider growth, the Appendix presents an extension to the model that includes exogenous growth in production technology and abatement technology. Along the steady state growth path of this model, pollution is constant, since these two growth rates just offset each other. In an alternate model with a different specification of damages, the pollution stock can be increasing or decreasing, depending on parameter values. Realistically, though, a growth path with an increasing stock of atmospheric greenhouse gases is probably unsustainable, given the predictions from climate scientists about the scale of likely damages from even the current level of greenhouse gas concentrations. Thus, a steady state growth path must involve a reduction in the pollution stock over time. But, carbon emissions and stocks are currently increasing, not decreasing, both domestically and globally.

If the economy is not currently on a steady state growth path, how does one interpret the model and results? One interpretation is that this model only gives policy recommendations regarding the economy once it has reached a steady state growth path. A second interpretation is that the model can also give policy recommendations along the transition path to the steady state (i.e., right now). This second interpretation requires the assumption that cycles about the transition behave the same as cycles about the steady state growth path, once each respective trend is filtered out. For example, this assumption requires that the elasticity of detrended emissions with respect to detrended GDP is the same along the transition path as it has been in the past. But any model about carbon policy must make assumptions about future technologies, either cyclical models like this one or secular models like Nordhaus (2008). Predictions of future technologies are inherently laden with uncertainty; one can think of this model's predictions as a central estimate and use sensitivity analysis to examine alternate assumptions. Dropping the

assumption of a constant elasticity and solving the model through value function iteration shows that when the elasticity is higher, the sensitivity of optimal emissions to the productivity shock is greater. This suggests that the qualitative results are robust, though the quantitative results may be subject to inaccuracy near term.

The model makes several simplifying assumptions that can be relaxed to answer other important questions. Because it contains only one representative agent, the model says nothing about distributional issues. Because it contains only one representative firm, it cannot address the cost advantages of taxes or tradable permits over command and control policies in the presence of heterogeneous abatement costs. By including banking and borrowing of emissions permits the model could analyze how those policy options affect the government's ability to optimally respond to cycles. Labor is not included as an input to production, so the effect of business cycles on employment is not considered here, though that area has been studied extensively in the RBC literature, and in the context of climate policy in Fischer and Springborn (2009). Behavioral anomalies are not modeled here; the representative agent is rational. This may be relevant to how policy could respond to business cycles. For example, loss aversion may affect optimal policy during recessions.⁴² Technology can grow in the model, but not endogenously. Lastly, the model presented is a quite basic DSGE business cycle model where the only stochastic element is an autocorrelated productivity shock. Other elements sometimes included in DSGE models, including intermediate goods producers, Keynesian price dynamics, and monetary effects, are omitted.⁴³

This paper presents the first study of how productivity shocks affect optimal environmental policy in an RBC context. Adapting to business cycles is important, perhaps more so for a policy as significant as climate change policy. As regional and federal regulation to mitigate climate change moves ahead, the implications of business cycles on policy choices are likely to be important components to achieving efficiency.

⁴² To address a similar concern, Fujii and Karp (2008) develop a method for computing optimal long-run climate policy paths under a non-constant pure rate of time preference.

⁴³ Gali (1999, 2004) argues that technology shocks play only a limited role, if any, in explaining business cycles, though this is disputed by Christiano et. al. (2003).

References


- Anderson, Gary and George Moore. "A Linear Algebraic Procedure for Solving Linear Perfect Foresight Models." *Economic Letters*, Vol. 17, No. 3 (1985), p. 247-252.
- Anderson, Gary. "Solving Linear Rational Expectations Models: A Horse Race." *Computational Economics*, Vol. 31, No. 2 (March 2008), p. 95-113.
- Balke, Nathan, Stephen Brown, and Mine Yucel. "Oil Price Shocks and U.S. Economic Activity: An International Perspective." Resources for the Future Discussion Paper 10-37 (July 2010).
- Bartz, Sherry, and David Kelly. "Economic Growth and the Environment: Theory and Facts." *Resource and Energy Economic*, Vol. 30, No. 2 (May 2008), p. 115-149.
- Baxter, Marianne and Robert King. "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series." *Review of Economics and Statistics*, Vol. 81, No. 4 (November 1999), p. 575-593.
- Blasing, T.J., Christine Broniak, and Gregg Marland. "Estimates of Monthly Carbon Dioxide Emissions and Associated $\delta^{13}\text{C}$ Values from fossil-fuel consumption in the U.S.A." In *Trends: A Compendium of Data on Global Change*, Carbon Dioxide Information Analysis Center, Oak Ridge National Laboratory, U.S. Department of Energy, Oak Ridge, TN, U.S.A. (2004).
- Bouman, Mathijs, P.A. Gautier and M.W. Hofkes. "Do Firms Time Their Pollution Abatement Investments Optimally?" *De Economist*, Vol. 148, No. 1 (March 2000), p. 71-86.
- Bushnell, James and Yihsu Chen. "Regulation, Allocation, and Leakage in Cap-and-Trade Markets for CO₂." NBER Working Paper #w15495 (November 2009).
- Carlson, Curtis, Dallas Burtraw, Maureen Cropper and Karen Palmer. "Sulfur Dioxide Control by Electric Utilities: What are the Gains from Trade?" *The Journal of Political Economy*, Vol. 108, No. 6 (December 2000), p. 1292-1326.
- Chang, Yongsung, and Sun-Bin Kim. "Heterogeneity and Aggregation: Implications for Labor-Market Fluctuations." *American Economic Review*, Vol. 97, No. 5 (December 2007), p. 1939-1956.
- Chay, Kenneth and Michael Greenstone. "The Impact of Air Pollution on Infant Mortality: Evidence from Geographic Variation in Pollution Shocks Induced by a Recession." *Quarterly Journal of Economics*, Vol. 118, No. 3 (August 2003), p. 1121-1167.
- Christiano, Lawrence, Martin Eichenbaum and Robert Vigfusson. "What Happens after a Technology Shock?" NBER Working Paper #9819 (July 2003).

- Cline, William. "The Economics of Global Warming." Washington, DC: Institute for International Economics (1992).
- Falk, Ita, and Robert Mendelsohn. "The Economics of Controlling Stock Pollutants: An Efficient Strategy for Greenhouse Gases." *Journal of Environmental Economics and Management*, Vol. 25, No. 1 (July 1993), p. 76-88.
- Fell, Harrison, Ian MacKenzie and William Pizer. "Prices versus Quantities versus Bankable Quantities." Resources for the Future Discussion Paper 08-32 (July 2008).
- Fischer, Carolyn and Alan Fox. "Output-Based Allocation of Emissions Permits for Mitigating Tax and Trade Interactions." *Land Economics*, Vol. 83, No. 4 (November 2007), p. 575-599.
- Fischer, Carolyn and Michael Springborn. "Emissions Targets and the Real Business Cycle: Intensity Targets vs. Caps or Taxes." Resources for the Future Discussion Paper 09-47 (November 2009).
- Fujii, Tomoko and Larry Karp. "Numerical Analysis of Non-Constant Pure Rate of Time Preference: A Model of Climate Policy." *Journal of Environmental Economics and Management*, Vol. 56, No. 1 (July 2008), p. 83-101.
- Fullerton, Don and Garth Heutel. "The General Equilibrium Incidence of Environmental Mandates." *American Economic Journal: Economic Policy*, Vol. 2, No. 3 (August 2010), p. 64-89.
- Gali, Jordi. "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review*, Vol. 89, No. 1 (March 1999), p. 249-271.
- Gali, Jordi. "On the Role of Technology Shocks as a Source of Business Cycles: Some New Evidence." *Journal of the European Economic Association*, Vol. 2, No. 2-3 (April-May 2004), p. 372-380.
- Gordon, Roger. "Do Publicly Traded Corporations Act in the Public Interest?" *Advances in Economic Analysis and Policy*, Vol. 3, No. 1 (2003).
- Gregg, Jay, Robert Andres, and Gregg Marland. "China: Emissions Pattern of the World Leader in CO₂ Emissions from Fossil Fuel Consumption and Cement Production." *Geophysical Research Letters*, Vol. 35, No. 8 (April 2008), L08806.
- Hodrick, Robert, and Edward Prescott. "Postwar US Business Cycles: An Empirical Investigation." *Journal of Money, Credit, and Banking*, Vol. 29, No. 1 (February 1997), p. 1-16.

- Hoel, Michael and Larry Karp. "Taxes versus Quotas for a Stock Pollutant." *Resource and Energy Economics*, Vol. 24, No. 4 (November 2002), p. 367-384.
- Kelly, David. "Price and Quantity Regulation in General Equilibrium." *Journal of Economic Theory*, Vol. 125 (2005), p. 36-60.
- Killian, Lutz. "Not all Oil Price Shocks are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market." *American Economic Review*, Vol. 99, No. 3 (June 2009), p. 1053-1069.
- Kim, Jinill and Sunghyun Kim. "Two Pitfalls of Linearization Methods." *Journal of Money, Credit and Banking*, Vol. 39, No. 4 (June 2007), p. 995-1001.
- King, Robert, Charles Plosser and Sergio Rebelo. "Production, Growth and Business Cycles. I. The Basic Neoclassical Model." *Journal of Monetary Economics*, Vol. 21, No. 2 (March 1988), p. 195-232.
- Kydland, Finn, and Edward Prescott. "Rules Rather than Discretion: The Inconsistency of Optimal Plans." *Journal of Political Economy*, Vol. 85, No. 3 (June 1977), p. 473-492.
- Kydland, Finn, and Edward Prescott. "Time to Build and Aggregate Fluctuations." *Econometrica*, Vol. 50, No. 6 (November 1982), p. 1345-1369.
- Lucas, Robert. *Models of Business Cycles*. Yrjo Jahnsson Lectures Series. London: Blackwell (1987).
- Metcalf, Gilbert. "Market-based Policy Options to Control U.S. Greenhouse Gas Emissions." *Journal of Economic Perspectives*, Vol. 23, No. 2 (Spring 2009), p. 5-27.
- Metcalf, Gilbert, Aparna Mathur and Kevin Hassett. "Distributional Impacts in a Comprehensive Climate Policy Package." NBER Working Paper #16101 (June 2010).
- Moore, Berrian and B.H. Braswell. "The Lifetime of Excess Atmospheric Carbon Dioxide." *Global Biogeochemical Cycles*, Vol. 8, No. 1 (1994), p. 23-38.
- Newell, Richard and William Pizer. "Regulating Stock Externalities under Uncertainty." *Journal of Environmental Economics and Management*, Vol. 45, No. 2S (March 2003), p. 416-432.
- Newell, Richard and William Pizer. "Indexed Regulation." *Journal of Environmental Economics and Management*, Vol. 56, no. 3 (November 2008), p. 221-233.
- Nordhaus, William. "To Slow or not to Slow: The Economics of the Greenhouse Effect." *The Economic Journal*, Vol. 101, No. 407 (July 1991), p. 920-937.


- Nordhaus, William. "A Review of the Stern Review on the Economics of Climate Change." *Journal of Economic Literature*, Vol. 45, No. 3 (September 2007), p. 686-702.
- Nordhaus, William. *A Question of Balance: Weighing the Options on Global Warming Policies*, New Haven and London: Yale University Press (2008).
- Oates, Wallace, Paul Portney, and Albert McGartland. "The Net Benefits of Incentive-Based Regulation: A Case Study of Environmental Standard Setting." *American Economic Review*, Vol. 79, No. 5 (December 1989), p. 1233-1242.
- Pizer, William. "Combining Price and Quantity Controls to Mitigate Global Climate Change." *Journal of Public Economics*, Vol. 85, No. 3 (September 2002), p. 409-434.
- Pizer, William. "The Case for Intensity Targets." *Climate Policy*, Vol. 5, No. 4 (2005), p. 455-462.
- Reilly, John. "Climate-Change Damage and the Trace-Gas-Index Issue." In Reilly, John and Margot Anderson (eds.), *Economic Issues in Global Climate Change: Agriculture, Forestry, and Natural Resources*, Boulder and Oxford: Westview Press (1992).
- Stavins, Robert. "A Meaningful U.S. Cap-and-Trade System to Address Climate Change." *Harvard Environmental Law Review*, Vol. 32, No. 2 (2008), p. 293-371.
- Stern, Nicholas. "The Economics of Climate Change." *American Economic Review*, Vol. 98, No. 2 (May 2008), p. 1-37.
- Strand, Jon. "Business Fluctuations, Worker Moral Hazard, and Optimal Environmental Policy." In Dixon, Huw and Neil Rankin (eds.), *The New Macroeconomics: Imperfect Markets and Policy Effectiveness*, Cambridge: Cambridge University Press (1995), p. 214-33.
- Tallarini, Thomas. "Risk-Sensitive Real Business Cycles." *Journal of Monetary Economics*, Vol. 45, No. 3 (June 2000), p. 507-532.
- Weitzman, Martin. "Prices Versus Quantities." *Review of Economic Studies*, Vol. 41, No. 4 (October 1974), p. 477-491.
- Weitzman, Martin. "A Review of the Stern Review on the Economics of Climate Change." *Journal of Economic Literature*, Vol. 45, No. 3 (September 2007), p. 703-724.

Figure 1



Note: Emissions data are from Blasing et. al. (2004) and represent total carbon emissions from fossil fuel combustion at the monthly level. They are seasonally adjusted using the X-12-ARIMA program. GDP data are from the Bureau of Economic Analysis and are seasonally adjusted. Both series are normalized to the January 1981 levels.

Figure 2



Note: The values are the cyclical residuals from applying the HP filter to quarterly carbon emissions and GDP data.

Figure 3




Figure 4




Figure 5




Figure 6




Figure 7




Figure 8

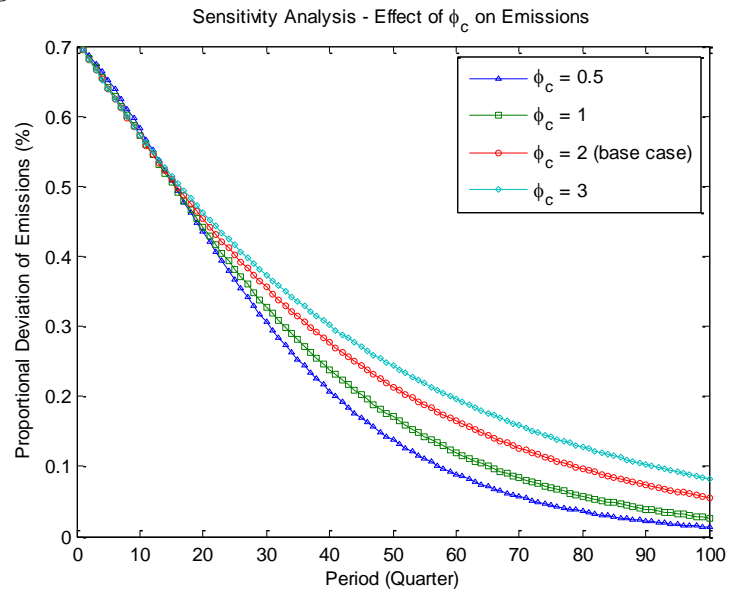


Figure 9




Figure 10




Figure 11




Figure 12




Table 1

CO₂ Emissions and GDP Regressions, Monthly						
	(1)	(2)	(3)	(4)	(5)	(6)
	Seasonal ARIMA	Seasonally Adjusted	HP detrend	BK detrend	CF detrend	BW detrend
Ln_gdp	0.758*** (0.175)	0.696*** (0.181)	0.859*** (0.126)	0.639*** (0.147)	0.545*** (0.00611)	0.723** (0.280)
Observations	263	275	276	180	276	276

Notes: Standard errors in parentheses. The dependent variable in each regression is the log of carbon dioxide emissions, adjusted or detrended as labeled. Column 1 presents results from a seasonal ARIMA(1,1,1)×(0,1,1)₁₂ regression of log of CO₂ (not seasonally adjusted) on log of GDP. Column 2 presents results from an ARIMA(1,1,2) regression of log of CO₂ (seasonally adjusted) on log of GDP. AR, MA, and cointegrating equation terms are not reported for either regression. Columns 3-6 present results from least squares regression of detrended emissions on detrended output allowing for Newey-West standard errors, where the series are detrended by either Hodrick-Prescott (Column 3), Baxter-King (Column 4), Christiano-Fitzgerald (Column 5), or Butterworth (Column 6) filters. Optimal lag lengths are all determined by minimizing Akaike Information Criterion. Constant term is omitted from regressions using detrended series. *** p<0.01, ** p<0.05, * p<0.1

Table 2

Base Case Parameter Values			
Parameter	Value	Description	Source
α	0.36	Curvature of production function: $f(k) = (k)^\alpha$	Chang and Kim (2007), Kydland and Prescott (1982)
β	0.98267	Quarterly discount rate	
δ	0.025	Quarterly capital depreciation	
ρ	0.95	Persistence of productivity shock	
σ_ε	0.007	Standard deviation of IID productivity innovation	
η	0.9979	Pollution decay	Reilly (1992)
θ_1	0.05607	Abatement cost	Nordhaus (2008)
θ_2	2.8	function: $g(\mu) = \theta_1 \mu^{\theta_2}$	
d_2	1.4647e-8	Pollution damages	
d_1	-6.6722e-6	function: $d(x) = d_2 x^2 +$	
d_0	1.3950e-3	$d_1 x + d_0$	
φ_c, φ_x	2	Coefficient of relative risk aversion	Stern (2008), Weitzman (2007)
$(1-\gamma)$	0.696	Elasticity of emissions with respect to output	Estimated from monthly emissions and GDP data, see Table 1

Appendix – Not for Publication

A1: State-level Regressions

The carbon emissions data analyzed in the paper are at the national level. Blasing et. al. (2004) also provide estimates of state-level carbon emissions from 1960-2001. The state-level data are annual, not monthly, limiting the extent to which cyclical behavior can be studied. However, when emissions are disaggregating by state, between-state heterogeneity in output and emissions can be used to analyze further the relationship between these two variables. Table A1 presents regression results from a model where state-level logs of carbon emissions are regressed on state-level logs of real GDP. State-level GDP is available from 1963 on from the Bureau of Economic Analysis. In column 1, a simple OLS regression of the log of emissions on the log of GDP shows a positive correlation between the two. Of course, both series are trending upwards, which could bias this coefficient. Therefore, in column 2, I include a set of state and year fixed effects. The coefficient on the log of GDP is smaller but still significantly positive. Finally, in column 3 I account for autoregressive error terms. Overall, the state-level (though annual) data provide more compelling evidence of the observation made from examining the national-level data: carbon emissions are higher in places and periods of higher output, and the relationship is inelastic.

A2: Static Model

The advantage of presenting this simpler model is the intuition it provides regarding the tradeoffs inherent in abatement choice. These tradeoffs are present in the dynamic model, though the use of a static environment allows me to find analytical solutions for optimal policy, and these solutions conveniently demonstrate the tradeoffs. To show how income effects and price effects differentially influence policy decisions, I include in this static model both a productivity shock and an income shock.

Consider a representative agent maximizing net output y , where $y = (1-d(e))f(ak)$. Gross output is $f(ak)$, e is emissions, and $d(e)$ is the damage function from emissions, expressed as the fraction of gross output lost. Let $d(e)$ be increasing and convex. In the production function $f(ak)$, k represents the inputs used in production (capital), and a is an exogenous shock to productivity. Let $f'(ak) > 0$ and $f''(ak) \leq 0$, so that the marginal

productivity of capital is positive and non-increasing. The agent is endowed with income b . The income can be distributed into two uses: production and abatement. The resource constraint is thus $k + z \leq b$, where z is the amount of abatement. This economy has two sources of exogenous variation: a is a productivity shock, and b is an income shock. The level of emissions is determined by both production and abatement: $e = g(z)f(ak)$, where $g(z)$ is an abatement function. Suppose that $g(z)$ is decreasing and convex, so that emissions decrease with more abatement but at a decreasing rate.

The representative agent's problem is to choose k and z to maximize utility $(1-d(e))f(ak)$, subject to the resource constraint $k + z \leq b$. Given that the resource constraint must bind, the problem can be written as the choice over just k , and the corresponding first order condition is

$$af'(ak)[1 - d(g(b-k)f(ak))] - d'(g(b-k)f(ak))[-g'(b-k)f(ak) + g(b-k)af'(ak)]f(ak) = 0.$$

The first order condition ensures that the marginal benefit of an additional unit of capital, $af'(ak)[1-d]$, equals the marginal cost, which is the reduced output emissions resulting from the marginal switch from abatement to capital, given in the second term of the above equation.

The first order condition can be used to consider comparative statics. Specifically, how is the choice between capital k and abatement z dependent on the productivity shock a and the income shock b ? For simplicity, assume that the production function f is linear. The relationship between the shocks and the optimal allocation are consistent with the counteracting income and price effects described above.

Applying the implicit function theorem to the first order condition, after setting the production function to be linear ($f(ak) = ak$), the effect of a change in the income shock b on capital can be written as

$$\frac{dk}{db} = \frac{-ak(d''g'ak(g - g'k) + d'(g' - kg'')) - d'akg'}{a[(g - g'k)^2(ad''k(g - g'k) + 2d') + kd'(kg'' - 2g')]}.$$

In this expression, d' is shorthand for $d'(e)$ evaluated at the equilibrium, and likewise for first and second derivatives of g . The denominator is positive. The numerator is also positive, thus indicating that an increase in b leads to an increase in k . The effect on abatement z can be solved by taking advantage of the resource constraint: $k + z = b$. This implies that $dz/db = 1 - dk/db$. Using the expression for dk/db in this expression for dz/db , it can be shown that

$$\frac{dz}{db} = \frac{a(g - g'k)(agd''k + 2d')}{a[(g - g'k)^2(ad''k(g - g'k) + 2d') + kd'(kg'' - 2g')]}.$$

Again, the numerator and the denominator are both positive, so an increase in b leads to an increase in both k and z .

This shock (b) is purely to the amount of resources available, and not to productivity, and thus it creates only an income effect. When income increases, the optimal response is to distribute additional resources to both production and abatement.

The effect on emissions is found from the emissions equation $e = f(ak)g(z)$, so that $de/db = g'(z)f(ak)dz/db + g(z)af'(ak)dk/db$. Substituting in the expressions for dz/db and dk/db and setting $f(ak) = ak$ yields:

$$\frac{de}{db} = \frac{a^2k^2d'(gg'' - 2g')}{a[(g - g'k)^2(ad''k(g - g'k) + 2d') + kd'(kg'' - 2g')]}.$$

The denominator is positive, but the numerator is of ambiguous sign. Thus, a positive income shock could increase emissions; this would happen if clean air is an inferior good.

Next, consider the effect of the productivity shock, a , on emissions. This shock provides both an income and a price effect. When a is high, the representative agent has a more productive technology and is thus richer. If clean air is a normal good, then an income effect pushes for lower emissions. However, a high value of a also means that capital is more productive and therefore the opportunity costs of investing in abatement instead of production are higher; in other words, abatement is more expensive. A price effect pushes for higher emissions. (The relevant price is the price of abatement, which increases with positive productivity shocks.) The net effect can be found from taking the derivative of the emissions equation: $de/da = af'(ak)g(z)dk/da + kf'(ak)g(z) + f(ak)g'(z)dz/da$. Substituting in the appropriate expressions and simplifying terms yields:

$$\frac{de}{da} = \frac{ak^2d'[g(kg'' - g') - (g')^2]}{a[(g - g'k)^2(ad''k(g - g'k) + 2d') + kd'(kg'' - 2g')]}.$$

The denominator is positive, but the numerator's sign is ambiguous (the first term in the square brackets is positive, and the next term including the minus sign is negative). However, this expression can be rewritten explicitly incorporating the derivative de/db :

$$\frac{de}{da} = k \frac{de}{db} + \frac{ak^2d'g'(kg'' - g)}{a[(g - g'k)^2(ad''k(g - g'k) + 2d') + kd'(kg'' - 2g')]}.$$

The first term, which includes and is of the same sign as de/db , is the income effect. The rest of the expression is the price effect. It is unambiguously positive. The net effect thus combines these two effects. Which effect dominates is determined by the utility function and abatement technology function.

A3: Growth Model

Consider an extension of the centralized economy model from section 2, where production technology and abatement technology both grow at a constant rate. A constant rate of growth in production technology implies that the production function includes an additional coefficient: gross output (not accounting for losses from pollution) is $a_t m_t k_{t-1}^\alpha$, where m_t captures an exogenous secular growth rate in productivity. Suppose that the growth rate is $\kappa > 1$, so that $m_{t+1} = \kappa m_t$. A constant growth rate in abatement technology is captured by an additional coefficient in the abatement cost function, q_t , representing the growth in abatement technology. If abatement technology is becoming cheaper, then q_t will be shrinking. Suppose that $q_{t+1} = q_t/\lambda$, where $\lambda > 1$ is the exogenous growth rate in abatement technology.

A balanced growth path is a solution to the model where each endogenous variable is growing (or shrinking) at a constant rate, though each variable need not be growing at the same rate. To find a balance growth path, the centralized model from section 2 is altered in two ways. First, let $e_t = q_t p(z_t) h(y_t)$, where $p(z_t)$ represents the effectiveness of abatement technology and is a decreasing function (more abatement = less emissions). The exogenous growth in abatement technology is still represented by q_t declining at a constant rate λ . Suppose that p takes the form of a power function with a negative exponent: $p(z_t) = z_t^{-\theta_3}$, where $\theta_3 > 0$. Second, let the effect of emissions on output be given by $y_t = m_t d(x_t) a f(k_{t-1}) = m_t x_t^d a_t k_{t-1}^\alpha$. Now, $d(x_t)$ is a decreasing function whose power $d < 0$.⁴⁴ The pollution stock equation, $x_t = \eta x_{t-1} + e_t$, can be used to show that emissions e_t and pollution x_t must be growing at the same rate. The budget constraint, $c_t = y_t - k_t + (1-\delta)k_{t-1} - z_t$ shows that c_t , y_t , k_t , and z_t are all growing at the same rate (i.e. $r_c = r_y = r_k = r_z$). The production function $y_t = m_t x_t^d a_t k_{t-1}^\alpha$ shows that $r_y = \kappa r_x^d r_k^\alpha$. The emissions equation $e_t = q_t z_t^{-\theta_3} y_t^{1-\gamma}$ shows that $r_e = 1/\lambda r_z^{-\theta_3} r_y^{1-\gamma}$. Finally, combining these with the first order conditions from this problem, it can be shown that on a balanced growth path the

⁴⁴ Without making these two changes, the balanced growth path is more restrictive, requiring growth rates of 1 (no growth) for all variables.

rate of growth in consumption, output, and capital is $r_c = \kappa^{1/(1-\alpha)}$. The rate of growth in emissions and the pollution stock is 1 (no growth).

In an alternate specification of the model where damages from pollution are in the utility function (see footnote 11), a more general solution is found where the pollution stock can grow or decline along the balanced growth path.

A4: Calibration of Damage Function d

Nordhaus (2008) includes a complex carbon cycle model, where the mass of carbon in the atmosphere, the upper oceans, and the lower oceans dynamically evolve according to anthropogenic and natural emissions, and these carbon concentrations enter into a calibrated function determining atmospheric and ocean temperatures. Here, I abstract from those considerations and model the pollution stock directly affecting output. This ignores some potentially important effects regarding the timing of abatement. However, as I show in the simulation results, deviations in emissions levels from business cycles have very little effect on the pollution stock because of the slow decay rate of CO₂. The damages from pollution do not change significantly with the business cycle and thus do not affect optimal abatement of CO₂; for this pollutant it is solely due to the cost side that optimal abatement varies. Adding an additional delay in the effect on output of current-period emissions, such as would be the case if the evolution of temperature were modeled, would only further diminish the already negligible cyclical effects of emissions damages.

The damage function d is calibrated in the following way. For a given level of the pollution stock, I use the equations in Nordhaus (2008) relating the atmospheric mass of carbon to evolving atmospheric and ocean temperatures to find the path of these temperatures over the extent of the period of simulation. Nordhaus (2008) also provides a function relating atmospheric temperature to cost in terms of fraction of output lost.⁴⁵ This represents global damages, but since my model considers optimal domestic policy I want US damages. I convert the global damage function using regionally disaggregated estimates for damages from temperature changes that are available in the model's documentation.⁴⁶ Even when expressed as

⁴⁵ The model's equations are provided in Nordhaus (2008), and the parameter values are available on his website: <http://nordhaus.econ.yale.edu/DICE2007.htm>.

⁴⁶ The documentation is available here: http://nordhaus.econ.yale.edu/Accom_Notes_100507.pdf. The relevant information is in the table on p. 24

a fraction of output, domestic damages from temperature changes are less than global damages, a fact that is consistent with the EPA's Technical Support Document on benefits of greenhouse gas abatement.⁴⁷ For a given level of the pollution stock (assumed to be constant for this calibration), I solve for the path of temperature and the associated domestic damages in each period of the simulation. I then discount those damages to come up with a measure of the current cost for a particular level of the pollution stock. This is done for 100 values of the pollution stock (ranging from 600 to 1200 GtC), and the results are shown in Figure A1. The y-axis plots the value of Ω , which equals one minus the fraction of domestic output lost from climate change damage (i.e. $1 - d(x)$). It ranges from about one at a pollution stock level of 600 GtC (equivalent to about 280 ppm CO₂, the pre-industrial level), to about 0.987 at 1200 GtC (equivalent to about 580 ppm CO₂). The 2005 level is about 800 GtC or 380 ppm CO₂. This curve is fit to a quadratic function via least squares: $d(x) = d_2x^2 + d_1x + d_0$. The fitted coefficients are reported in the text and in Table 2.

A5: Linearizing and Solving the Dynamic Model

The equations describing the model as presented in Section 2 and parameterized in Section 3 are the following:

$$\begin{aligned}
 \ln a_t &= \rho \ln a_{t-1} + \varepsilon_t \\
 c_t &= y_t - i_t - z_t \\
 i_t &= k_t - (1-\delta)k_{t-1} \\
 x_t &= \eta x_{t-1} + e_t + e^{row}_t \\
 e_t &= (1-\mu_t)y_t^{1-\gamma} \\
 z_t &= \theta_1 \mu_t^{\theta_2} y_t \\
 y_t &= (1-d_2x_t^2 - d_1x_t - d_0)a_t k_{t-1}^\alpha \\
 -\frac{1}{c_t^{\phi_c}} + \frac{\beta}{c_{t+1}^{\phi_x}} \{ &\alpha y_{t+1} k_t^{-1} [1 - \theta_1 \mu_t^{\theta_2} - \theta_1 \theta_2 \mu_t^{\theta_2-1} (1-\gamma) y_{t+1}^{\gamma-1} \mu_{t+1}^{\theta_2-1} e_{t+1}] + (1-\delta) \} = 0
 \end{aligned}$$

⁴⁷ The document can be found at www.regulations.gov searching for "Technical Support Document – Benefits." See Table 1 on p. 12.

$$\frac{1}{c_t^{\varphi_c}} \{ -(1 - \theta_1 \mu_t^{\theta_2}) (2d_2 x_t + d_1) a_t k_{t-1}^\alpha + \theta_1 \theta_2 \mu_t^{\theta_2-1} y_t^\gamma + \theta_1 \theta_2 (1 - \gamma) \mu_t^{\theta_2-1} y_t^{\gamma-1} e_t (2d_2 x_t + d_1) a_t k_{t-1}^\alpha \} - \frac{1}{c_{t+1}^{\varphi_c}} \beta \theta_1 \theta_2 \eta \mu_{t+1}^{\theta_2-1} y_{t+1}^{\gamma-1} = 0$$

The first equation describes the Markov process governing the productivity shock; the second through seventh equations describe the resource constraints and technologies. The last two equations are the parameterized and simplified versions of the first order conditions for the choice of capital k_t and pollution x_t .

The model is solved by linearizing about the steady state, so first a steady state solution must be found. This is done by dropping the time subscripts from each of the five equations governing the dynamic model and solving. Additionally, the steady state value for the productivity shock is set to 1, and the steady state value of rest-of-world emissions e^{row} is set to three times the value of domestic emissions, so that the steady state version of the fourth equation is $x = \eta x + e + 3e$. An analytical solution for the remaining steady state values cannot be found, but the system of equations can be reduced to just three equations in three unknowns and solved computationally.

The nine equations can then be log-linearized about the steady state values of the variables. Using a tilde to denote a proportional deviation from the steady state value of a variable, and a bar to denote the steady state value itself, the six linearized equations are

$$\begin{aligned} \tilde{a}_t - \rho \tilde{a}_{t-1} &= \varepsilon_t \\ \bar{c} \tilde{c}_t + \bar{i} \tilde{i}_t + \bar{z} \tilde{z}_t - \bar{y} \tilde{y}_t &= 0 \\ \bar{k} \tilde{k}_t - \bar{k} (1 - \delta) \tilde{k}_{t-1} - \bar{u}_t &= 0 \\ \tilde{x}_t - \eta \tilde{x}_{t-1} - (1 - \eta) \tilde{e}_t &= 0 \\ \bar{e} \tilde{e}_t - \bar{y}^{1-\gamma} (1 - \gamma) \tilde{y}_t + \bar{\mu} \bar{y}^{1-\gamma} (\tilde{\mu}_t + (1 - \gamma) \tilde{y}_t) &= 0 \\ \tilde{z}_t - \theta_2 \tilde{\mu}_t - \tilde{y}_t &= 0 \\ \bar{y} \tilde{y}_t - (1 - d_2 \bar{x}^2 - d_1 \bar{x} - d_0) \bar{k}^\alpha (\tilde{a}_t + \alpha \tilde{k}_{t-1}) + \bar{k}^\alpha (2d_2 \bar{x}^2 + d_1 \bar{x}) \tilde{x}_t &= 0 \\ \varphi_c \tilde{c}_t + \beta (-\varphi_c) \tilde{c}_{t+1} \{ \alpha \bar{y} \bar{k}^{-1} [1 - \theta_1 \bar{\mu}^{\theta_2} - \theta_1 \theta_2 (1 - \gamma) \bar{y}^{\gamma-1} \bar{\mu}^{\theta_2-1} \bar{e}] + 1 - \delta \} \\ &+ \beta \{ \alpha \bar{y} \bar{k}^{-1} [1 - \theta_1 \bar{\mu}^{\theta_2} - \theta_1 \theta_2 (1 - \gamma) \bar{y}^{\gamma-1} \bar{\mu}^{\theta_2-1} \bar{e}] (\tilde{y}_{t+1} - \tilde{k}_t) \} \\ &+ \beta \{ \alpha \bar{y} \bar{k}^{-1} [-\theta_1 \bar{\mu}^{\theta_2} \theta_2 \tilde{\mu}_{t+1} \\ &- \theta_1 \theta_2 (1 - \gamma) \bar{y}^{\gamma-1} \bar{\mu}^{\theta_2-1} \bar{e} ((\gamma - 1) \tilde{y}_{t+1} + (\theta_2 - 1) \tilde{\mu}_{t+1} + \tilde{e}_{t+1})] \} = 0 \end{aligned}$$

$$\begin{aligned}
& -\varphi_c \tilde{c}_t \{ -(1 - \theta_1 \bar{\mu}^{\theta_2}) (2d_2 \bar{x} + d_1) \bar{k}^\alpha + \theta_1 \theta_2 \bar{\mu}^{\theta_2 - 1} \bar{y}^\gamma \\
& \quad + \theta_1 \theta_2 (1 - \gamma) \bar{y}^{\gamma - 1} \bar{\mu}^{\theta_2 - 1} \bar{e} (2d_2 \bar{x} + d_1) \bar{k}^\alpha \} \\
& \quad + \{ \theta_1 \bar{\mu}^{\theta_2} (2d_2 \bar{x} + d_1) \bar{k}^\alpha \theta_2 \tilde{\mu}_t + (-(1 - \theta_1 \bar{\mu}^{\theta_2}) 2d_2 \bar{x} \bar{k}^\alpha) \tilde{x}_t \\
& \quad + (1(1 - \theta_1 \bar{\mu}^{\theta_2}) (2d_2 \bar{x} + d_1) \bar{k}^\alpha (\tilde{a}_t + \alpha \tilde{k}_{t-1}) \\
& \quad + \theta_1 \theta_2 \bar{\mu}^{\theta_2 - 1} \bar{y}^\gamma ((\theta_2 - 1) \tilde{\mu}_t + \gamma \tilde{y}_t) \\
& \quad + \theta_1 \theta_2 \bar{\mu}^{\theta_2 - 1} \bar{y}^{\gamma - 1} \bar{e} (2d_2 \bar{x} + d_1) \bar{k}^\alpha ((\theta_2 - 1) \tilde{\mu}_t + (\gamma - 1) \tilde{y}_t + \tilde{e}_t + \tilde{a}_t + \alpha \tilde{k}_{t-1}) \\
& \quad + \theta_1 \theta_2 (1 - \gamma) \bar{\mu}^{\theta_2 - 1} \bar{y}^{\gamma - 1} \bar{e} 2d_2 \bar{x} \bar{k}^\alpha \tilde{x}_t \} - \beta \theta_1 \theta_2 \eta \bar{\mu}^{\theta_2} \bar{y}^{\gamma - 1} (-\varphi_c \tilde{c}_{t+1} \\
& \quad + (\theta_2 - 1) \tilde{\mu}_{t+1} + (\gamma - 1) \tilde{y}_{t+1} = 0
\end{aligned}$$

This set of linear equations can be solved by the AMA, which is suited for solving linear rational expectations models. The model is a rational expectations model in the sense that all of the variables dated $t+1$ are actually expectations of those values, though the expectations operator is dropped in the equations written above. The structure of these models is

$$\sum_{i=-\tau}^0 G_i x_{t+i} + \sum_{i=1}^{\theta} F_i E_t x_{t+i} = \varepsilon_t .$$

The constants $\tau > 0$ and $\theta > 0$ are the number of lags and leads, respectively. In this model, both of these values are equal to 1. The matrix G_i is the coefficient matrix on the lagged and contemporaneous values of the variables x , and F_i is the coefficients on the future values. The shock is represented as ε_t and has expectation zero. Taking expectations and simplifying the above equation yields

$$\sum_{i=-\tau}^{\theta} H_i E_t x_{t+i} = 0,$$

where H_i are referred to as the structural coefficient matrices. These matrices are input into the matlab code for AMA, and a set of solution matrices is output. The solution can be used to find the reduced form of the structural model:

$$x_t = \sum_{i=-\tau}^{-1} B_i x_{t+i} + B_0 \varepsilon_t .$$

The matrices B_i and B_0 are used to find impulse response functions to a technology shock (ε_t), or to simulate cycles. The matlab code is available on the author's website.

A6: First Order Conditions in Decentralized Model

The government's problem under a tax policy has four first-order conditions, for the government's (constrained) choice of τ_t , x_t , y_t , and k_t .

$$\begin{aligned}
& -U'(c_t)z_\tau(\tau_t, y_t) + \lambda_t U''(c_t)z_\tau(\tau_t, y_t) \\
& \quad + \lambda_{t-1} \left(U''(c_t)(-z_\tau(\tau_t, y_t))(r(\tau_t, y_t, k_{t-1}) + 1 - \delta) + U'(c_t)r_\tau(\tau_t, y_t, k_{t-1}) \right) \\
& \quad + \zeta_t(-e_\tau(\tau_t, y_t)) = 0 \\
& \quad \quad \quad \zeta_t + \zeta_{t+1}(-\eta)\beta + \omega_t a_t f(k_{t-1})d'(x_t) = 0 \\
& U'(c_t) \left(1 - z_y(\tau_t, y_t) \right) + \lambda_t (-U''(c_t) \left(1 - z_y(\tau_t, y_t) \right) \\
& \quad + \lambda_{t-1} \left(U''(c_t) \left(1 - z_y(\tau_t, y_t) \right) (r(\tau_t, y_t, k_{t-1}) + 1 - \delta) \right. \\
& \quad \left. + U'(c_t)r_y(\tau_t, y_t, k_{t-1}) \right) + \zeta_t \left(-e_y(\tau_t, y_t) \right) + \omega_t = 0 \\
& -U'(c_t) + \beta U'(c_{t+1})(1 - \delta) + \lambda_{t+1}(-U''(c_{t+1})(1 - \delta)\beta) \\
& \quad + \lambda_t (U''(c_t) + \beta U''(c_{t+1})(1 - \delta))(r(\tau_{t+1}, y_{t+1}, k_t) + 1 - \delta) \\
& \quad + \beta U'(c_{t+1})r_k(\tau_{t+1}, y_{t+1}, k_t) + \lambda_{t-1}(-U''(c_t)(r(\tau_t, y_t, k_{t-1}) + 1 - \delta)) \\
& \quad - \omega_{t+1}\beta(1 - d(x_{t+1}))a_{t+1}f'(k_t) = 0
\end{aligned}$$

Here each of the Lagrange multipliers are divided through by β^t to eliminate redundant β coefficients throughout the equations. Also, the first order conditions contain the partial derivatives of the e , r , and z functions with respect to y , τ , and k . Because the government moves first and therefore accounts for the impact of its actions on the actions of the firm and the consumer, it considers how its actions affect e , r , and z , and hence these derivatives appear. Each of the derivatives is evaluated using one of the firm's first order conditions or the equations governing emissions and abatement.

Combining the equation for abatement expenditures, $z_t = y_t g(\mu_t)$, with the firm's first order condition for choice of abatement, $\tau_t h(y_t) = y_t g'(\mu_t)$, and using the chosen function forms yields:

$$z(\tau_t, y_t) = \theta_1 \left(\frac{1}{\theta_1 \theta_2} \right)^{\frac{\theta_2}{\theta_2-1}} \tau_t^{\frac{\theta_2}{\theta_2-1}} y_t^{\frac{\theta_2(1-\gamma)-1}{\theta_2-1}},$$

and thus

$$z_\tau(\tau_t, y_t) = \frac{\theta_2}{\theta_2 - 1} \cdot \frac{z_t}{\tau_t}$$

$$z_y(\tau_t, y_t) = \frac{\theta_2(1-\gamma) - 1}{\theta_2 - 1} \cdot \frac{z_t}{y_t}$$

The equation for emissions, $e_t = (1-\mu_t)y_t^{(1-\gamma)}$, along with the firm's abatement choice equation and function forms yields:

$$e(\tau_t, y_t) = y_t^{1-\gamma} - \left(\frac{1}{\theta_1\theta_2}\right)^{\theta_2-1} \tau_t^{\frac{1}{\theta_2-1}} y_t^{\frac{-\gamma}{\theta_2-1}+1-\gamma}$$

and thus

$$e_\tau(\tau_t, y_t) = \frac{-1}{\theta_2 - 1} \frac{\mu_t y_t^{1-\gamma}}{\tau_t}$$

$$e_y(\tau_t, y_t) = (1-\gamma)y_t^{-\gamma} - \left(\frac{-\gamma}{\theta_2 - 1} + 1 - \gamma\right) \mu_t y_t^{-\gamma}$$

Finally, the firm's first order condition for choice of capital input defines the capital rental rate:

$$r(\tau_t, y_t, k_{t-1}) = \alpha y_t k_{t-1}^{-1} [1 - \tau_t (1 - \mu_t) (1 - \gamma) y_t^{-\gamma} - \theta_1 \mu_t^{\theta_2}]$$

and thus

$$r_\tau(\tau_t, y_t, k_{t-1}) = -\alpha(1-\gamma)y_t^{1-\gamma}k_{t-1}^{-1} + \alpha(1-\gamma)\left(1 + \frac{1}{\theta_2 - 1}\right)\mu_t y_t^{1-\gamma}k_{t-1}^{-1} - \frac{\alpha\theta_1\theta_2}{\theta_2 - 1} y_t k_{t-1}^{-1} \mu_t^{\theta_2} \tau_t^{-1}$$

$$r_y(\tau_t, y_t, k_{t-1}) = \alpha k_{t-1}^{-1} - \alpha(1-\gamma)^2 \tau_t y_t^{-\gamma} k_{t-1}^{-1} + \alpha(1-\gamma)\left(1 - \gamma - \frac{\gamma}{\theta_2 - 1}\right) \tau_t \mu_t y_t^{-\gamma} k_{t-1}^{-1}$$

$$- \alpha\theta_1 \left(1 - \frac{\theta_2\gamma}{\theta_2 - 1}\right) k_{t-1}^{-1} \mu_t^{\theta_2}$$

$$r_k(\tau_t, y_t, k_{t-1}) = -\alpha y_t k_{t-1}^{-2} + \alpha(1-\gamma)\tau_t y_t^{1-\gamma} k_{t-1}^{-2} - \alpha(1-\gamma)\tau_t \mu_t y_t^{1-\gamma} k_{t-1}^{-2} + \alpha\theta_1 y_t k_{t-1}^{-2} \mu_t^{\theta_2}$$

These expressions are substituted into the first order conditions presented above. The system is log-linearized and solved as in the case of the centralized economy.

A7: Quantity Policy

Suppose that instead of a tax levied on emissions, the government can mandate the level of emissions, q_t , that the firm must produce. With one representative firm, the quantity constraint is equivalent to a cap-and-trade scheme.⁴⁸ The firm's profit function is $\pi_t = y_t - r_t k_{t-1} - z_t$, which

⁴⁸ The policy does not allow for banking or borrowing of permits, or other suggested policies like a "safety valve" option or an allowance reserve. These options may partially alleviate the need for the government to adjust its

omits the tax payments. Firms do not pay for the emissions permits. The firm's constrained optimization problem is to maximize $y_t - r_t k_{t-1} - z_t$ such that $(1-\mu_t) \cdot h(y_t) = q_t$ and $z_t = y_t g(\mu_t)$. Solving this problem using the Lagrangian method yields the following equations:

$$r_t = y_t f'(k_{t-1})/f(k_{t-1}) \{1 - g(\mu_t) - (1-\mu_t)h'(y_t) \cdot y_t g'(\mu_t)/h(y_t)\},$$

$$(1-\mu_t) \cdot h(y_t) = q_t.$$

The second equation is just the quantity constraint. The first equation is analogous to the first equation under the tax policy, and it demonstrates that the shadow price on a unit of emissions created by the quantity constraint is $y_t g'(\mu_t)/h(y_t)$. The consumer's problem is identical to the consumer's problem under the tax policy.

The government's problem in the quantity policy is

$$\max_{\{c_t\}, \{k_t\}, \{x_t\}, \{q_t\}} \sum \beta^t U(c_t)$$

Such that, for all t ,

$$-U'(c_t) + \beta E_t U'(c_{t+1}) \cdot [r_{t+1} + (1-\delta)] = 0$$

$$r_t = y_t f'(k_{t-1})/f(k_{t-1}) \{1 - g(\mu_t) - (1-\mu_t)h'(y_t) \cdot y_t g'(\mu_t)/h(y_t)\}$$

$$q_t = e_t$$

$$c_t = r_t k_{t-1} + \pi_t - k_t + (1-\delta)k_{t-1}$$

$$\pi_t = y_t - r_t k_{t-1} - z_t$$

$$e_t = (1-\mu_t) \cdot h(y_t)$$

$$x_t = \eta x_{t-1} + e_t + e^{row}_t$$

$$z_t = y_t g(\mu_t)$$

$$y_t = (1-d(x_t)) a_t f(k_{t-1}).$$

The associated Lagrangian is

$$\begin{aligned} & \sum \beta^t U(y_t - k_t + (1-\delta)k_{t-1} - z(q_t, y_t)) \\ & + \sum \lambda_t [-U'(y_t - k_t + (1-\delta)k_{t-1} - z(q_t, y_t)) \\ & + \beta U'(y_{t+1} - k_{t+1} + (1-\delta)k_t - z(q_{t+1}, y_{t+1})) \cdot (r(q_{t+1}, y_{t+1}, k_t) + 1-\delta)] \\ & + \sum \zeta_t [x_t - \eta x_{t-1} + e_t^{row} + e(q_t, y_t)] \\ & + \sum \omega_t [y_t - (1-d(x_t)) a_t f(k_{t-1})] \end{aligned}$$

policy (quota level) to cycles. However, it will not eliminate this need, since the firm's objective will still differ from the social optimum.

Here the constraints can be manipulated to express z_t , r_t , and e_t in terms of q_t , y_t , and k_{t-1} . The Lagrangian multipliers are defined as in the tax case.

The first order conditions are

$$\begin{aligned}
& -U'(c_t)z_q(q_t, y_t) + \lambda_t U''(c_t)z_q(q_t, y_t) \\
& \quad + \lambda_{t-1} \left(U''(c_t) \left(-z_q(q_t, y_t) \right) (r(q_t, y_t, k_{t-1}) + 1 - \delta) + U'(c_t)r_q(q_t, y_t, k_{t-1}) \right) \\
& \quad + \zeta_t \left(-e_q(q_t, y_t) \right) = 0 \\
& \quad \quad \quad \zeta_t + \zeta_{t+1}(-\eta)\beta + \omega_t a_t f(k_{t-1})d'(x_t) = 0 \\
& U'(c_t) \left(1 - z_y(q_t, y_t) \right) + \lambda_t (-U''(c_t) \left(1 - z_y(q_t, y_t) \right) \\
& \quad + \lambda_{t-1} \left(U''(c_t) \left(1 - z_y(q_t, y_t) \right) (r(q_t, y_t, k_{t-1}) + 1 - \delta) \right. \\
& \quad \left. + U'(c_t)r_y(q_t, y_t, k_{t-1}) \right) + \zeta_t \left(-e_y(q_t, y_t) \right) + \omega_t = 0 \\
& -U'(c_t) + \beta U'(c_{t+1})(1 - \delta) + \lambda_{t+1}(-U''(c_{t+1})(1 - \delta)\beta) \\
& \quad + \lambda_t (U''(c_t) + \beta U''(c_{t+1})(1 - \delta)(r(q_{t+1}, y_{t+1}, k_t) + 1 - \delta) \\
& \quad + \beta U'(c_{t+1})r_k(q_{t+1}, y_{t+1}, k_t)) + \lambda_{t-1}(-U''(c_t)(r(q_t, y_t, k_{t-1}) + 1 - \delta)) \\
& \quad - \omega_{t+1}\beta(1 - d(x_{t+1}))a_{t+1}f'(k_t) = 0
\end{aligned}$$

These differ from those under the tax policy only in the fact that r , e , and z are functions of q rather than τ .

The equations defining r_t and z_t in terms of q_t , y_t , and k_{t-1} are

$$r(q_t, y_t, k_{t-1}) = \alpha y_t k_{t-1}^{-1} (1 - \theta_1 \mu_t^{\theta_2} - \theta_1 \theta_2 (1 - \gamma) \mu_t^{\theta_2 - 1} (1 - \mu_t))$$

$$z(q_t, y_t, k_{t-1}) = \theta_1 (1 - q_t / y_t^{1-\gamma})^{\theta_2} y_t$$


$$e(q_t, y_t) = q_t$$

These equations can be used to calculate the derivatives needed in the first-order conditions. It can be shown that the four first-order conditions of the quantity policy are equivalent to the first-order conditions of the tax policy; that is, any solution to the tax policy's equations is also a solution to the quantity policy's equations. In particular, the second equation in the quantity policy is identical to the second equation in the tax policy. The first equation in the quantity policy is equal to the first equation in the tax policy times a factor (it can be shown that $\frac{z_q}{z_\tau} =$

$\frac{r_q}{r_\tau} = \frac{e_q}{e_\tau}$, making the two equations equivalent). The same is true for the fourth equation in each

policy. It can be demonstrated numerically that the third equations under each policy are equivalent to each other.

Figure A1



Notes: Y-axis values are simulated losses in output from different atmospheric concentrations of carbon based on Nordhaus (2008).

Table A1

	State-Level Annual CO₂ Emissions and GDP		
	(1)	(2)	(3)
	lnemis	lnemis	lnemis
lngdp	0.874*** (0.057)	0.666*** (0.128)	0.113*** (0.033)
Constant	-7.038*** (0.67)	-4.182*** (1.35)	1.863*** (0.014)
Observations	1989	1989	1938
R-squared	0.77	0.98	0.65
State fixed effects?	No	Yes	Yes
Year fixed effects?	No	Yes	Yes
AR(1) error term?	No	No	Yes

Note: Standard errors are in parentheses. Standard errors are clustered at the state level in columns 1 and 2.

*** p<0.01, ** p<0.05, * p<0.1

Data are logs of annual state-level carbon emissions and GDP.