Carrots and sticks for new technology: Abating greenhouse gas emissions in a heterogeneous and uncertain world

David A. Robalino * and Robert J. Lempert **

RAND, 1700 Main Street, Santa Monica, CA 90407-2138, USA

Received 17 November 1998; revised 20 August 1999

Many governments use technology incentives as an important component of their greenhouse gas abatement strategies. These "carrots" are intended to encourage the initial diffusion of new, greenhouse-gas-emissions-reducing technologies, in contrast to carbon taxes and emissions trading which provide a "stick" designed to reduce emissions by increasing the price of high-emitting technologies for all users. Technology incentives appear attractive, but their record in practice is mixed and economic theory suggests that in the absence of market failures, they are inefficient compared to taxes and trading. This study uses an agent-based model of technology diffusion and exploratory modeling, a new technique for decision-making under conditions of extreme uncertainty, to examine the conditions under which technology incentives should be a key building block of robust climate change policies. We find that a combined strategy of carbon taxes and technology incentives, as opposed to carbon taxes alone, is the best approach to greenhouse gas emissions reductions if the social benefits of early adoption sufficiently exceed the private benefits. Such social benefits can occur when economic actors have a wide variety of cost/performance preferences for new technologies and either new technologies have increasing returns to scale or potential adopters can reduce their uncertainty about the performance of new technologies by querying the experience of other adopters. We find that if decision-makers hold even modest expectations that such social benefits are significant or that the impacts of climate change will turn out to be serious then technology incentive programs may be a promising hedge against the threat of climate change.

Keywords: climate change, technology policy, uncertainty, agent-based modeling, exploratory modeling, social interactions

1. Introduction

Most policy-makers would very much like to use carrots as well as sticks to address the threat of climate change. Among the most popular carrots are technology incentives, such as tax credits and subsidies, that encourage economic actors to adopt new, low-greenhouse-gas emitting technologies. For instance, the Clinton Administration has proposed a three-stage plan to meet US commitments under the Kyoto framework, of which the first stage (1999–2003) focuses on voluntary actions and tax credits for energy efficient technologies. On the one hand, such incentive policies have a compelling logic. New technologies will likely be critical to any significant reduction of society's greenhouse gas emissions during the course of the 21st century. Technology incentive policies offer the potential for politically feasible actions that might make significant emissions reductions possible in the future.

There are, however, good reasons to take a jaundiced view of such incentives. Economic theory suggests that in the absence of market failures the most efficient policies for inducing innovation, as well as reducing emissions, are sticks such as carbon taxes or tradable emissions permits that impose on emitters the full social costs of greenhouse gases. By comparison, technology incentives may distort the market. While there has been much less practical experience with price-based sticks than with technology carrots, the latter have had a mixed record of achieving any practical success [1] independent of their relative efficiency. Finally, such incentive policies may be irrelevant given the scale of the climate problem. Certainly, the technology programs conducted to date by many governments have not stopped the inexorable rise in worldwide emissions of greenhouse gases.

This study addresses the question of whether and under what conditions technology incentives can provide a key building block for effective and feasible climate change policies. To date, the integrated assessment modeling community has been largely silent on this important question because of two methodological problems. First, technology incentives aim to influence the development and diffusion of new emissions-reducing technologies in the presence of market failures that would otherwise deflect this diffusion below the socially optimal path. Paramount among them are coordination failures [2,3] that occur in the presence of increasing returns to scale, imperfect information, and heterogeneous preferences regarding new technologies. It is difficult to represent such factors in a mathematical model that can be solved within the optimization techniques commonly used for climate change assessment.

Second, extreme uncertainty surrounds many of the factors associated with technology diffusion. These range from our understanding of micro processes to our judgments

* Now at: World Bank, 1818 H. St. NW, Washington, DC 20433, USA. E-mail: drobalino@worldbank.org.

** E-mail: lempert@rand.org.

about the fundamental macro question – what is the potential for new technologies to dramatically reduce the costs of future reductions of greenhouse gas emissions? Any judgments about the efficacy of technology incentives will depend, in part, on expectations about these uncertainties. Unfortunately, different stakeholder groups hold very different expectations, probably unresolvable in the near-term, about the potential of new technologies. Thus, traditional methods of prediction-based policy analysis will have difficulty adjudicating among alternative policy options.

We address these problems using two analytic innovations. First, we employ a new method of decision-making under extreme uncertainty, exploratory modeling [4,5] that allows us to compare alternative policies without requiring predictions of the future cost and performance of new technologies. The basic idea is to use simulation models to create a large ensemble of plausible future scenarios, where each member of the ensemble represents one guess about how the world works and one choice among many alternative strategies we might adopt to influence the world. We then use search and visualization techniques to extract from this ensemble of scenarios information that is useful in distinguishing among policy choices. These methods are consistent with the traditional, probability-based approaches to uncertainty analysis because when such distributions are available, one can lay them across the scenarios and thus calculate expected values for various strategies, value of information, and the like. However, in situations characterized by extreme uncertainty, the exploratory modeling method has an advantage. It allows us to express information at many levels of uncertainty within the same framework, and to draw upon tools from dynamic areas of today's technology – computer search and visualization – to help us extract knowledge from this collection of information and make policy arguments [6,7].

Second, we employ an agent-based model of technology diffusion, which provides a convenient platform for treating key factors, such as increasing returns to scale, the flow of imperfect information about technology performance among agents, and the heterogeneity in agents' preferences governing the adoption of new technologies. In particular, we consider heterogeneous economic agents that adopt technologies on the basis of expectations about future costs (influenced by expectations about the potential for increasing returns to scale) and performance (influenced by learning about the performance experience of other agents). Our model is similar in many respects to others recently proposed [8–10], though as we will discuss below it includes several commonly neglected factors that are important to our conclusions about policy choices. Agent-based models have been popular because they are a useful means to represent important features of economies, such as bounded rationality, and heterogeneity among economic actors, that are poorly treated in standard models. However, their use in practical policy analysis has been limited. In capturing such features, agent-based models also capture some of the indeterminism of the real world, so that a single model with a single set of input parameters might admit of a wide array of potential outputs. Thus, these models are hard to employ within traditional, prediction-based policy analysis, and are often relegated to building intuition about a problem rather than providing rigorous contributions to the policy debate. By framing and beginning to address the question – what is the set of all plausible agent behaviors consistent with choosing one policy over another? – exploratory modeling provides a means for making agent-based models useful for policy analysis.

In this study, we compare two potential near-term (that is, over the next one or two decades) carrot and stick policies for abating global emissions of CO_2: carbon taxes and technology subsidies.[1] Carbon taxes, the stick, can reduce emissions by inducing economic actors to use less energy and to choose lower-emitting energy technologies. Technology subsidies, the carrot, can reduce emissions by inducing early adopters to purchase new technologies. While theoretically the most efficient policy, carbon taxes can slow the economy and can be politically unattractive. Subsidies also need to be financed either through taxes or government borrowing and, in the absence of market failures, should be a more costly means of reducing emissions than the tax, since, by definition, if the optimal tax does not cause the agents to switch to an alternative technology, this technology must not be efficient. However, in a dynamic setting, the benefit of a technology subsidy may be higher than its cost, if it accelerates the diffusion of an environmental friendly technology. But policy makers do not usually know ex ante whether the diffusion of low-emitting technologies can be accelerated or the benefits of doing so. This is the problem that we analyze.

Although our model is highly stylized, it provides useful guidance on how carbon-taxes and technology subsidies should be combined in the presence of uncertainty about the severity of future climate change and the potential for near-term actions to affect the future costs of greenhouse gas emissions. We find that adaptive-decision strategies for greenhouse gas abatement that combine the sticks of carbon taxes with the carrots of technology incentives perform better than strategies employing carbon taxes alone, if the agents have heterogeneous preferences and there are either significant opportunities for increasing returns or significant opportunities for learning about technology performance among the populations of actors. Our results also suggest that this combined strategy is more robust than the tax-only policy, in that policy makers need to ascribe about

[1] Much of the political debate currently focuses on emissions trading rather than carbon taxes as a means of placing a cost on carbon emissions. We focus on taxes here because they are analytically somewhat simpler. There are important differences, both theoretical and practical, between taxes and emissions trading, but the results of this study should apply similarly to both. Governments propose and pursue many different types of technology incentive programs, including tax credits, subsidies, government procurements, information programs, technology transfer programs, mandates, etc. We focus on subsidies here for analytic simplicity. It is less clear how the results of this study apply to other types of incentive programs.

1 : 2 odds that the externalities are important in order to prefer this combined strategy. Finally, our analysis confirms the importance of carbon taxes. We replicate the standard result that carbon taxes are the most efficient policy when opportunities for increasing returns and learning are not significant and find no cases where technology subsidies alone are a viable response to climate change.

2. Landscape of plausible futures

This study compares the performance of alternative greenhouse gas-emissions reductions strategies against a wide range of plausible futures which represent our uncertainty about the damages due to climate change, the extent to which new technologies may reduce the costs of future emissions reductions, and the role that heterogeneity and social interactions have in economic agents' technology choices. In this section, we describe how this landscape of plausible futures has been generated. To do so, first we describe the agent-based model of technology diffusion that is used to formalize the uncertainties. We next describe how we constrain the set of plausible input parameters by requiring the model to reproduce the current market mix of energy technologies, current greenhouse gas emissions, and the historic range of rates of technology diffusion. Finally we describe the landscape of plausible futures.

2.1. Model

As shown in figure 1, our model considers a population of agents, each a producer of a composite good that is aggregated as total GDP, using energy as one key input. Each time period the agents make two key choices that affect the evolution of the system as a whole: first they choose among several energy-generation technologies and second, given their chosen technology, they decide how much energy to consume. (We can think of this as choosing a production function and choosing where to operate on that production function.) The technologies differ in cost, carbon emissions per unit of energy, and the performance each agent obtains from them. The agents choose among these technologies in order to maximize a utility function which depends on the cost (including the cost of capital and those costs resulting from carbon taxes) and performance of each technology. In general, the agents do not have perfect information about future costs and the performance of new technologies, so they make this choice based on their expectations and their level of risk aversion. In addition, the agents are heterogeneous, that is, they differ from one another in terms of the cost/performance tradeoffs expressed in their utility functions, their level of risk aversion, their initial expectations about technology performance, and the actual performance they will get from various technologies.

Aggregated across the agent population, the agents' technology and energy consumption choices have a variety of important consequences. First, each agent pays a certain price for energy and generates a certain amount of emissions. The average energy cost and the total carbon emissions can affect the rate of economic growth in an otherwise exogenously expanding global economy. Second, the agents' technology choices affect future costs of new technologies as well as the agents' expectations about these costs and the technologies' performance. When a new technology is introduced there is often great uncertainty about how well it will serve the needs of a variety of potential users. Agents can estimate the performance of a new technology by querying other agents who have used it. Thus, the diffusion rate can depend reflexively on itself, since each user generates new information that can influence the adoption decisions of other potential users. In addition, the cumulative number of agents that have used a technology can affect the cost, and expectations about the cost, of that technology through increasing returns to scale.

This model is in accord with other models of induced technology change introduced in the climate change literature in recent years (for instance, see recent reviews by Azar and Dowlatabadi [11] and Grübler, Nakicenovic and Victor [12]). Mattsson [8] proposes a model where the investment costs for new emissions-reducing technologies are related to accumulated experience through learning curves. He calculates the optimal emissions-reduction path, which depends on the *exogenous* parameters characterizing the learning rate. Grübler and Gritsevskii [9] consider the potential effects of cost reductions in new technologies from learning by doing (through investment and R&D) and of uncertainty about these reductions. They show that technologies that are economically unattractive at present, but perhaps attractive in the future, can diffuse into the market if policy-makers invest up front in R&D and promote niche market applications. More recently, Gritsevskii and Nakicenovic [10] have introduced a model to examine, similarly to ours, the effects of learning-by-doing, uncertainty, and technology spillovers (due to technology clusters rather than learning among heterogeneous agents) using stochastic optimization methods. They argue that these effects are crucial in shaping technology paths over the next few decades, find a wide, bimodal range of "base-case" emissions scenarios, and suggest that research and development policies can have the largest impacts if they focus on related clusters of technologies. Similarly to these studies, our model considers the role that uncertainty and learning-by-doing may have in shaping the diffusion of new technologies. However, we also consider the learning process through which economic agents learn about particular performance characteristics of new technologies. Our study only considers the effects of policies on the diffusion of fuel-switching technologies. We do not consider the effects of policies on innovation (e.g., as in work by Goulder, Mathai, and Schneider [13,14]) nor the rate of improvements in energy-efficiency.

As described in detail in section 3.2, we find that the dynamics of the system shown in figure 1, in the absence of carbon taxes and technology subsidies, depends

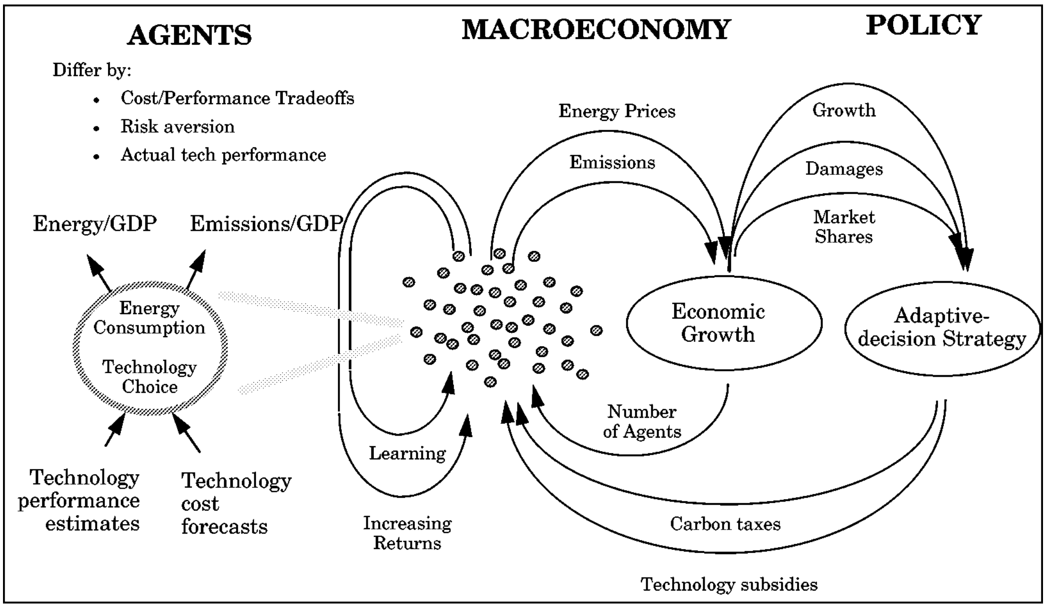

Figure 1. Agent-based model of technology diffusion used in this study. Economic agents choose among alternative technologies on the basis of forecasts of cost and performance. The forecasts are influenced by learning among the agents and potential price decreases due to increasing returns to scale. The agents have heterogeneous initial expectations about technology performance and heterogeneous preferences for technology cost/performance tradeoffs. The agents' choices influence the level of energy prices and of greenhouse gas emissions, which both influence the rate of economic growth. Policy decisions about the level of carbon taxes and technology subsidies, which depend on observations of economic growth, damages and technology diffusion, also influence the agents' technology choices.

primarily on several key factors: the extent to which economic growth is affected by changes in energy costs and by increases in CO_2 emissions, the extent to which economic actors evaluate new technologies using different cost/performance preferences and different degrees of risk aversion, the rate at which reliable information about the performance of new technologies flows through a population of economic actors, the extent to which cumulative production volumes reduces the cost of new technologies, and the economic actors' initial expectations about the cost and performance of new technologies.

2.1.1. Economic growth

Given the focus of our study, we use a very simple representation of economic growth. We assume that the world economy is in a *steady state* where output per capita grows at some exogenous rate that can be modified by changes in the price of energy and any damages due to climate change.

We write the Gross Domestic Product (GDP) in each of two world regions, the OECD and the Rest of the World (ROW), with the difference equation:

$$\mathrm{GDP}_g(t) = \left[\mathrm{GDP}_g(t-1)\left(\gamma_g - \phi_{xg}\dot{C}_g(t) - \phi_{sg}\dot{S}_g(t)\right)\right] \times \left\{1 - \kappa_0\left[\frac{\mathit{Conc}(t)}{\mathit{Conc}(1765)}\right]^{\kappa_1}\right\}, \tag{1}$$

where γ_g represents the exogenous growth rate in the regions $g =$ OECD and ROW; $\dot{C}_g(t)$ is the growth rate of the average cost of energy per unit of output in region g including the costs of energy-producing technologies, any carbon tax imposed in order to reduce CO_2 emissions, and any subsidies on new, low-emitting technologies; $S_g(t)$ is the per unit of output cost of the subsidy; and ϕ_{xg} and ϕ_{sg} are the corresponding elasticities. The last term in equation (1) represents the share of GDP that is lost as a result of climate change. For simplicity, we express these damages due to climate change as a geometric function of *Conc*(t), the atmospheric concentration of CO_2 (see Cline [15] and Nordhaus and Yang [16]), the parameters κ_0 and κ_1, and the pre-industrial concentration *Conc*(1765) $=$ 280 ppm.[2]

The concentration is given by the simple difference equation [16]

$$\mathit{Conc}(t) = \kappa_2 E(t) + (1-\kappa_3)\mathit{Conc}(t-1), \tag{2}$$

where $E(t)$ are the global carbon emissions, $\kappa_2 = 0.90$ is the marginal atmospheric retention ratio of CO_2 emissions, and $\kappa_3 = 0.005$ is the rate of transfer of CO_2 from atmosphere to other reservoirs. In our model, worldwide emissions of carbon dioxide are given by

$$E(t) = \sum_g \sum_{i=1}^{N_g(t)} s(t) n_{g,j(i)}(t) m_{j(i)}, \tag{3}$$

where $N_g(t) = \mathrm{GDP}_g(t)/s(t)$ is the number of agents in each region, each producing $s(t)$ units of output. The CO_2 emissions intensity $m_{j(i)}$ (carbon emitted per unit energy consumed) is determined by agent i's choice of technology j. As described in the appendix, the energy intensity $n_{g,j(i)}(t)$, the energy agent i requires to produce one unit

[2] In earlier work [55], we used a damage function based on the flow, rather than the stock, of carbon emissions which produced results relatively less favorable to the combined carbon tax and subsidy strategy than we find here.

of GDP with technology j, represents the agent's choice of energy consumption and is determined by the cost of energy (inclusive of all taxes and subsidies), the elasticity of substitution, and improvements in energy efficiency.

2.1.2. Agents' technology choices

As shown in figure 1, we pay particular attention in this study to the process by which economic agents choose among alternative technologies. The economic literature generally uses one of two ways to describe such choices. The standard approach assumes that agents make choices in order to maximize a given utility or pay-off function. This implies that the agent chooses from a *known* set of alternatives by computing ex ante the pay-off of each alternative. Under conditions of uncertainty, the approach assumes agents know the probability distribution of the pay-off and choose the alternative that provides on average maximum risk-adjusted pay-off [17,18]. The second approach replaces maximizing behavior by trial and error processes and adaptive behaviors, based on the assertion that real economic agents have only a limited capacity to gather and process all the information that would actually be necessary to solve a real-world optimization problem, and rather use rules of thumb [19]. While we believe the later approach to be a more general description, we use the first in this study for two reasons. First, we consider technologies with long lifetimes and high switching costs, so economic agents have a strong incentive to approximate as closely as possible the optimal choice given their objectives, preferences and constraints. Second we work with a small set of alternatives (three technologies), and thus we expect that the agents' information-gathering problems should be small.

Each of the technologies in our study are represented by three factors. The cost, which can drop over time due to increasing returns, and the emission intensities (the quantity of CO_2 emitted when generating one unit of energy) are intrinsic to the technology. The performance of the technology also depends on the agent using it. Technologies differ according to characteristics such as size of the equipment, noise generated when operating, resistance to changes in temperature, and reliability under different types of applications and intensities of use. The importance of such characteristics depends on the particular characteristics of the user, such as firm size, technical skills, and environmental conditions [20]. For instance, many renewable energy sources produce intermittent power but can operate independently from an electric grid. Thus, in our model, the performance of each technology is represented by a distribution of performance values across the population of agents.

With these considerations, we assume that our agents choose among three technologies in order to maximize an intertemporal expected utility. The agents have imperfect information, so they estimate utility on the basis of their expectations about technology performance and costs. We define $\langle U_{i,g,j}(\tau, T_i^{\text{life}})|t\rangle$ as the agent i in region g's estimate at time t of the risk-adjusted pay-off it would gain by using energy conversion technology j from some time $\tau > t$ through the end of the technology's lifetime, T_i^{life}. We write this risk-adjusted pay off using the Cobb–Douglas functional form

$$\begin{aligned}\left\langle U_{i,g,j}\left(\tau, T_i^{\text{life}}\right)\middle|t\right\rangle = {}&\left\langle \mathit{Performance}_{i,g,j}^{\alpha_i}|t\right\rangle \\ &\times \left\langle \mathit{Cost}_{i,g,j}\left(\tau, T_i^{\text{life}}\right)^{\alpha_i - 1}\middle|t\right\rangle \\ &- \lambda_i(\mathit{Var}_{\text{Performance}} + \mathit{Var}_{\text{Cost}}). \qquad (4)\end{aligned}$$

The first term, $\langle \mathit{Performance}_{i,g,j}|t\rangle$, is the agent i's expectation at time t of the performance it will get from the technology j, which it forms on the basis of its own past experience with the technology and from the experience of other agents that have used it [21]. As described in detail in the appendix, this term depends on the rate at which the agents sample the experience of others in order to learn about the performance of new technologies. This rate is given by the parameter ϑ, defined in the appendix, which represents the fraction of the agent population sampled by each agent each time period. The second term, $\langle \mathit{Cost}_{i,g,j}(\tau, T_i^{\text{life}})|t\rangle$, is the expected cost of using the technology over its lifetime, which depends on projections of future use and estimates of the potential for cost reductions from learning-by-doing. The projections and estimates are derived from observations of past trends in usage and cost of the technology. As described in the appendix, this term depends on the potential for cost reductions due to increasing returns to scale, represented here by the learning curve parameters β_j. The third term represents the agent's risk aversion taken as a function of the variance of the estimates of technology performance and future costs.

Equation (4) addresses the role heterogeneity may play in affecting the diffusion of new technologies across the system. First, the agents have different preferences regarding new technologies, as reflected in different values for the exponents α_i, representing the cost/performance tradeoffs, and in the risk aversion coefficient λ_i. We assume that both the α_i and the λ_i are normally distributed across the agent population, with means α and λ, respectively, and with variance υ. Second, the agents have different expectations regarding the performance and cost of each technology. Each agent's expectations of performance are private; that is, they apply only to that agent, since in fact each agent will gain a unique performance from each technology. While the cost forecasts are public, that is shared in common by all the agents, each agent in general will have different planning horizons, determined by the remaining lifetime of the technology they are already using. As we illustrate in section 3, both types of heterogeneity significantly affect observed market shares and the dynamics of the diffusion process.[3]

[3] Bassanini and Dosi [56] have formally examined the effects of heterogeneity on technology diffusion in models simpler than that used here. In a system with two technologies, constant prices, and non-switching behavior (agents choose a technology once and cannot subsequently change their minds) these authors find that market domination can result from "chance" and that heterogeneity diminishes the probability of

2.2. Choosing model parameters

In order to compare the performance of alternative climate change policies using this agent-based model we must choose the model's parameters. This task is complicated by the large number of parameters for the model, over 40 as shown in table 1, as well as the great uncertainty surrounding many of them, such as those describing the damages due to climate change, the potential for cost reductions in new energy technologies, and the heterogeneity of the agent population.

Traditionally, one would address this model uncertainty by choosing probability distributions for the model parameters and then propagating these distributions through the model to calculate probability distributions for model outcomes as a function of the choice of policy. This approach is questionable here for two reasons. First, good information does not exist to provide accurate estimates of the probability distributions for key parameters in the model. Indeed, much of the political debate on climate change can be characterized as very different expectations about such unknown parameters among different stakeholders [6,7].

Second, there is no reason to believe that this agent-based model, or *any* model, would provide an accurate description of the likelihood of different climate change futures. Mathematically, it is well known that the time evolution of many processes is unpredictable, and common sense suggests that long-term predictions are usually wrong. Nonetheless, we can reasonably expect a model to propagate certain constraints into the future and to track the implications of a wide variety of assumptions. For instance, we cannot know whether or not new technology will dramatically lower the cost of future greenhouse gas emissions, but we know that if such technology becomes available, it will diffuse according to certain economic rules.

Exploratory modeling [4,5] is a method for decision-making under conditions of extreme uncertainty that allows us to distinguish among alternative policy strategies using models which do not predict the future. Figure 2 shows a schematic of the exploratory modeling method as we use it here. In the first step, we create what we call a *landscape of plausible futures*, that is, the set of all plausible paths into the future consistent with available information. To do so we use the agent-based model which describes alternative technology diffusion paths through an economy characterized by uncertainty, heterogeneity, imperfect information, and the potential that increased use might lower the cost of new technologies. Not all outputs from this model are plausible, however, because both the plausible range of values for individual input parameters and the plausible range of behavior of the model outputs are constrained by available information. For instance, we know that the individual parameters representing the elasticities in equation (1) lie within a certain range. We also know that historically energy technologies have shown a wide, but finite range of diffusion speeds. If we thus constrain the diffusion speeds output by the model to lie within this range, we have a constraint on allowable combinations of input parameters.

observing domination. Our simulations results, presented in section 4, suggest that these results hold in our model as well.

The set of input parameters for the model that satisfy all the constraints defines what we call an *uncertainty space*. The model outputs resulting from these parameters defines the *landscape of plausible futures*. To find these plausible futures, we conduct a search over the model, looking for allowable sets of input parameters. This search is the subject of this section. In the next stage of the analysis, we define a set of alternative strategies and calculate their performance across this landscape of plausible futures to construct a large ensemble of alternative scenarios. We can then conduct searches across these scenarios to compare the performance of the alternative strategies.

In order to create the landscape of plausible futures, we must define ranges for the model inputs and constraints on the model outputs. Given the simplicity of the model and the types of data readily available, we chose three constraints for the model outputs. The model should reproduce current (1995) market shares for energy technologies; reproduce current levels of carbon emissions, energy intensities, and emissions intensities; and generate diffusion rates no faster than 20 years (from 1% to 50% penetration), which is faster than the rates historically observed for energy technologies [22]. The first two constraints guarantee that our model is consistent with current data. The third forces the model to be consistent with one (of the few) binding constraints history places on future technologies. (See [12] for a discussion of reoccurring patterns in technology diffusion.)

In order to define the ranges for the input parameters, it is convenient to divide them into two classes, macroeconomic parameters and microeconomic parameters. The former describe the growth of the economy and its response to changes in the cost of energy and damages due to climate change. These parameters include the exogenous rate of economic growth, the elasticity of economic growth with respect to the cost of energy, the elasticity of economic growth with respect to the cost of the subsidy, the parameters characterizing damages and the concentration of carbon in the atmosphere, the parameters defining the energy demand functions, and those used to simulate exogenous improvements in energy-efficiency. The plausible ranges we assign to these parameters and the sources for this information are shown in table 1.

The microeconomic parameters describe the cost and performance of energy technologies and the behavior of the population of consumers of energy technologies. In this preliminary study we confine ourselves to parameters describing only three generic types of technologies: high emission intensity systems, such as coal fired power plants, which at present provide the bulk of the world's energy; medium emissions intensity systems, such as natural gas powered combustion, which provide a significant minority of the world's energy; and low emissions intensity systems,

Table 1
Model parameters.

Parameter	Description	Range	Reference
	Macroeconomic parameters		
γ_{OECD}, γ_{ROW}	Exogenous economic growth rate	2%, 4%	World Bank [57]
$\phi_{x,\mathrm{OECD}}$, $\phi_{x,\mathrm{ROW}}$	Elasticity of economic growth with respect to the cost of energy	[0.00,0.1]	Dean and Hoeller [58]
$\phi_{s,\mathrm{OECD}}$, $\phi_{s,\mathrm{ROW}}$	Elasticity of economic growth with respect to the cost of the subsidy	[0.00,1]	
κ_0 κ_1		0.3 [0–3]	
Conc(1765) κ_2 κ_3	Parameters governing the accumulation of CO_2 and damages	750 0.9 0.005	Houghton et al. [60] Cline [15]
	Parameters governing the demand functions for energy:		
$a_{g,j}$	Scale parameter	n.a	Endogeneous
ε	Elasticity of substitution	0.5	Burniaux et al. [59]
$\dot{a}_{\mathrm{OECD}}$, $\dot{a}_{\mathrm{ROW}}$	Exogenous technological progress (energy efficiency gains)	0.02, 0.016	Defined to reproduce Manne and Richels [61] projections with no diffusion of new technologies
	Microeconomic parameters		
Technologies			
q_j	Mean of the performance distribution for technology j	100	Fixed
ν_j	Variance of the performance distribution	$j=1$: [1,5] $j=2$: [1,10] $j=3$: [10,50]	Defined to reproduce different levels of uncertainty
$C_{g,j}(0)$	First unit cost	$j=1$: 9 $/Gj $j=2$: 8 $/Gj $j=3$: 19 $/Gj	Manne and Richels [61]
m_j	Emission intensity	$j=1$: 0.038 Tc/Gj $j=2$: 0.023 Tc/Gj $j=3$: 0.001 Tc/Gj	Tester et al. [62]
T^{life}	Lifetime	[5,40] years	
β_j	Level of increasing returns to scale	$j=1$: $[-0.2,0.05]$ $j=2$: $[-0.3,0.05]$ $j=3$: [0,0.6]	
γ_j	Switching costs	[0,1]	
Agents			
α	Means of the distributions of preferences with respect to performance and costs	[0,1]	
λ	Mean of the distribution of risk aversion coefficients	[0,1]	
υ	Variance of the distribution of preferences (heterogeneity)	[0,2]	
ρ_g	Learning error (performance)	[1, 10]	
ρ_c	Learning error (price)	[1]	
ϑ	Sample size (social interactions)	[0%,20%]	

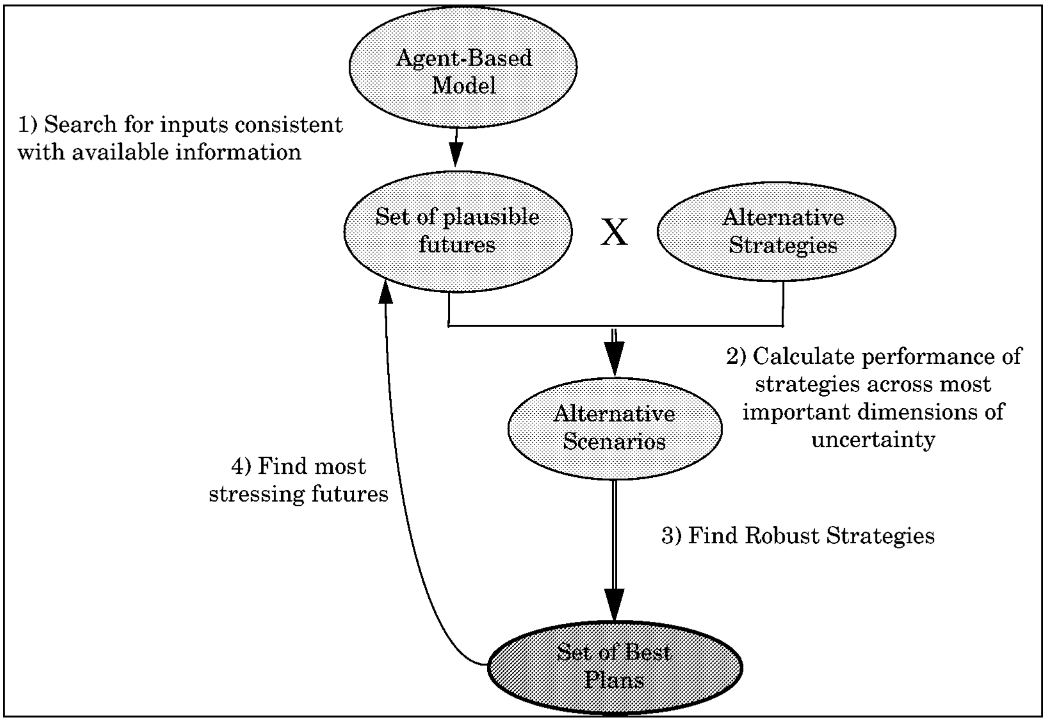

Figure 2. Schematic of the exploratory modeling method. We first conduct a search across the space of input parameters for the agent-based technology diffusion model, looking for those combinations of model inputs which are themselves and which produce model outputs consistent with available data (step 1). We next examine the performance of our alternative adaptive-decision strategies across the most important dimensions of this set of plausible futures (step 2) and search for those strategies that best meet the decision-makers' goals (step 3). Finally, we search across the entire set of plausible futures, looking for counter-examples to our conclusions (step 4).

Table 2
Range of parameter values that give model results.

	N	Mean	Std	Min	Max
Selected endogenous variable					
GDP (pv to 2045) 10^{12} US$	1611	89.38879	0.6321316	85.95	95.14
Variance GDP	1611	0.0691186	0.1807702	0	3.58
Emissions (2045) 10^9 tons carbon	1611	17.86817	7.139093	3.43	91.52
LE Mrkt Share (2045)	1611	0.2116574	0.175137	0	1
Variance LE Mrkt Shr	1611	0.0011767	0.0008317	0	0.0044
HE Mrkt Shr (1995)	1611	0.7018187	0.2370247	0	1
ME Mrkt Shr (1995)	1611	0.2669522	0.2331999	0	1
LE Mrkt Shr (1995)	1611	0.0312477	0.0233843	0	0.09
Parameter					
υ	1611	0.4617517	0.2546615	0	1
α	1611	0.3641409	0.2465205	0	1
λ	1611	0.1782104	0.1416057	0	0.948
ν_1	1611	1.642264	0.8030245	1	4.845
ν_2	1611	4.864824	2.265987	1.05	10
ν_3	1611	23.44834	8.575531	10	49.4
β_1	1611	0.0299963	0.0203724	0	0.05
β_2	1611	0.0273085	0.0192474	0	0.05
β_3	1611	0.1872024	0.1201763	0	0.582
ϑ	1611	0.0189926	0.0149588	0	0.126
ρ_q	1611	5.129235	2.146752	1.008	10
κ_1	1611	0.5011173	0.6329159	0	3
$\phi_{x,\mathrm{OECD}}$, $\phi_{x,\mathrm{ROW}}$	1611	0.0164333	0.0150177	0	0.145
γ_j	1611	0.4586189	0.2518373	0	1
T^{life}	1611	20.81891	8.306418	5.2	40

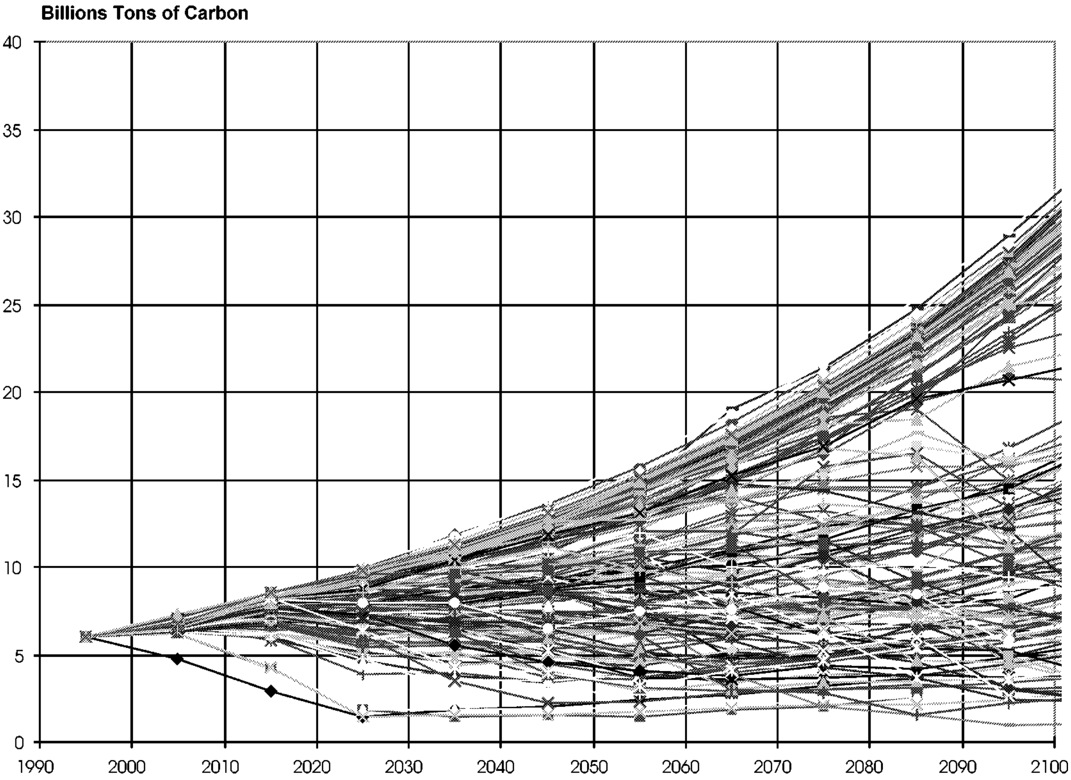

Figure 3. Set of plausible greenhouse gas emissions trajectories over the 21st century. The trajectories shown are a random sample of those calculated using the full range of plausible inputs to our agent-based model of technology diffusion.

such as renewable, biomass, and/or new nuclear power facilities, which at present are not in widespread use, but may be significant energy sources in the future. This is clearly a very coarse grouping, similar to that used in [12], but it is sufficient to draw policy conclusions about the appropriate mix of carbon taxes and technology subsidies. The ranges and sources for these parameters are also shown in table 1.

We can now search across the space of allowable model inputs looking for those combinations that give model outputs consistent with our constraints. In this study we found it convenient to use a genetic algorithm search, similar to that used by Miller [23]. Our search generated 1,611 plausible sets of input parameters that vary widely in key parameters such as the level of increasing returns to scale, agents' heterogeneity, uncertainty regarding new technologies, and future damages due to climate change. This process did, however, rule out as implausible sets of points with very small levels of uncertainty and heterogeneity regarding expectations about the performance of new technologies (such points do not satisfy the constraint on initial market shares), sets of points where the agents' utility functions are largely insensitive to costs, and sets of points with both very high learning rates and levels of increasing returns to scale. The range of variation of the restricted parameters is summarized in table 2.

Using this plausible set of input parameters, we project society's future carbon dioxide emissions, which we plot as a landscape of plausible futures in figure 3. Hence, figure 3 provides our model's best estimate forecast of global emissions of carbon dioxide over the course of the next century. The model suggests that the plausible levels of emissions could vary by over an order of magnitude, from a 1 GT/year to over 30 GT/year. This distribution of estimates is similar to that of Gritsevskii and Nakicenovic [10] and spans the lower part of the range of 400 published scenarios collected by the IPCC Special Report on Emissions Scenarios [24,25]. We generate a narrower range because ours are driven largely by different assumptions about technology, while the scenarios in the literature span a range of assumptions about both demand and technology. While our model is very simple and ignores a wide variety of important factors, there is no reason to believe that adding complexity should reduce this range of plausible futures. Imagine that we had a highly complex, "perfect" model that could reproduce all past trends to any desired level of detail. Unless history is completely deterministic, such a model would need to be able to generate many possible pasts, of which the actual past was only one example [26]. Even when run through the crucible of the past, such a "perfect" model should spread out into a wide cone of plausible futures as it moves forward in time. Figure 3 represents what is probably the most concrete fact we know about the future a century hence – it is in fact highly unpredictable no matter how perfectly it is modeled.

Table 3
Parameters defining adaptive decision strategies.

Parameter	Carbon tax policy	Subsidy policy	Combined policy
Tax_0	100 \$/ton	0	100 \$/ton
d_{tax}	5%	0	5%
Min_{tax}	0% economic growth	0	0% economic growth
Sub_0	0	40% mrkt share	40% mrkt share
Max_{sub}	0	50% mrkt share	50% mrkt share
Min_{sub}	0	20% mrkt share	20% mrkt share
L_{sub}	0	15 years	15 years

3. Robust strategies

We cannot predict the future, but it may often be possible to determine the best actions available to shape the future to our liking. As shown in figure 2, the next step of an exploratory modeling analysis is to define a set of strategies, compare these strategies against the landscape of plausible futures, and create a ensemble of scenarios. We can then examine these scenarios to find the strategies that best meet our goals.

In this study we aim to evaluate how policy-makers ought to combine two different types of policy levers – carbon taxes and technology subsidies – given the great uncertainty of the climate change problem, the possibility that if climate change is a serious problem, society will have to make much larger emissions reductions than are currently seen as possible, and the potential that early policy actions may improve society's ability to make such large reductions. We will look for strategies that are robust, that is, ones that are effective over a very wide range of expectations about the future, because such strategies may be the only ones that will be politically viable given the diverse array of stakeholders who will be impacted by any climate change policies. We expect that such robust strategies will be by necessity adaptive [6,7,27], that is, they will adjust themselves over time in response to new information generated by observations of the climate and economic systems.

3.1. Defining adaptive strategies

Carbon taxes have two main effects. First, they increase the cost of energy and thus reduce its use. Second, taxes provide incentives for economic actors to develop and choose energy technologies with relatively lower emissions by reducing the cost differential between low- and high-emitting technologies. Overall, carbon taxes have competing effects on economic growth: slowing growth by increasing the marginal cost of production and speeding growth by reducing damage-causing carbon emissions.

Technology subsidies have the potential to speed up technology diffusion by increasing the number of early users, which can reduce costs via learning-by-doing and speed the spread of information about the performance of these technologies. However, subsidies must be financed out of taxes or savings, which slows economic growth and may have little impact on emissions if the subsidized technologies do not turn out to have desirable cost/performance characteristics.

As shown in figure 1, taxes and subsidies affect our model of technology diffusion in two ways. First, they influence the adoption decisions of individual agents by changing their technology cost forecasts. As described in the appendix, subsidies decrease the expected cost of using non-emitting technologies and taxes increase the cost of using high and medium emitting technologies in proportion to their carbon intensities. In addition, taxes and subsidies directly affect the rate of economic growth by increasing the cost of energy to the economy and through the effect of the government spending required, respectively.

Taxes and subsidies in this study are adaptive, that is, they can change over time in response to observations of the rate of economic growth, the damages due to climate change, and the market share of low-emitting technologies. In our model, the tax begins with an initial level Tax_0 per ton of emitted carbon. The tax grows at an annual rate d_{tax} but is adjusted if one of two conditions hold: (a) the cost of the tax is greater than the marginal cost of emissions of carbon dioxide (expressed as a percentage of GDP), or (b) the global economic growth rate is below some minimum rate, $\partial(\text{GDP}_{\text{OECD}} + \text{GDP}_{\text{ROW}})/\partial t < Min_{\text{tax}}$. In the latter case, the tax drops to Tax_0 and begins to grow again. This description of a steadily growing tax is consistent with the optimum taxes described in the literature, and the stopping condition represents a way in which political conditions may force a tax to terminate.

The implementation of the subsidy is in general more complicated. The potential benefits of the subsidy depend on the existence of spillover effects in the diffusion of subsidized technologies, that are usually unknown ex ante. The costs depend on how the subsidy is financed. In this paper we take a very simple approach. The subsidy begins at a level Sub_0 percent of the cost of the subsidized technology. This subsidy stays at a constant level over time until either: (i) the market share for low emitting technologies goes above a level Max_{sub}, or (ii) the market share fails to reach a minimum level, Min_{sub}, after L_{sub} years. If either of these conditions is met, the subsidy is permanently terminated. Basically the subsidy is terminated once policy makers observe that the technology succeeds or that it fails to diffuse over a long period of time.

For simplicity, we consider three alternative strategies in this study. The Tax-Only strategy, the Subsidy-Only strategy and a combination of the Tax- and Subsidy-Only strategies (see table 3). Each strategy is characterized by the seven parameters: Tax_0, d_{tax}, Min_{tax}, Sub_0, Max_{sub}, Min_{sub}, and L_{sub}. The Tax- and Subsidy-Only strategies were chosen by searching for the best tax and the best subsidy strategy at the point in uncertainty space characterized by the average value for each of the model input parameters shown in table 1. Preliminary explorations suggest that results presented in the following sections are reasonably insensitive to this simplification.

3.2. Generating scenarios

In order to choose among the alternative strategies, we must compare their performance against a variety of plausible futures, as shown in figure 2. We define a scenario as one combination of a strategy and a plausible future represented by a set of input parameters to the model. It is impractical to examine all the possible scenarios because the computations would take too long and, more importantly, it would be difficult to interpret so many dimensions (on the order of 15) of possible actions and plausible uncertainties using the visualization tools at our disposal. Thus, we will first examine a strategically chosen subset of scenarios that we will use for generating hypotheses. As discussed in section 3.4, we will then test these hypotheses by searching back across the full range of scenarios.

We reduce the dimensionality of the uncertainty space by using econometric techniques to find those dimensions most important to the questions of interest. In particular, we want to find those input parameters that have the biggest impacts, individually and in combination, on the future market share of low-emitting technologies. This market share is a key model outcome, because it determines whether or not low cost greenhouse gas abatement is available in the model.

We begin by identifying the market share of low emitting technologies in the year 2045 as a field variable [2,28] and calculate its mean and variance for each of our set of 1,611 uncertainty space points. Then we estimate the probability that the mean will be above 20% of the market in year 2045, as a function of the model parameters. Since the appropriate functional form is unknown we use the second order linear expansion

$$\Pr(\mu_{2045} > 0.20) = \Phi\left(\psi_0 + \sum_{i=1}^{k} \psi_i \theta_k + \sum_{i=1}^{k} \sum_{j=1}^{k} \psi_{k+ij} \theta_i \theta_j\right), \quad (5)$$

where μ_{2045} is the mean of the distribution, θ is a vector of model parameters, ψ is a vector of coefficients to be estimated, and Φ is the logistic probability distribution. Because our simulations showed that the variance of the distribution also changes as a function of the parameters, each observation included in the estimation data set was weighted by its variance.

As hypothesized, parameters related to the potential for costs reductions of the technologies, as well as those defining the distribution of preferences and the number of social interactions, appear to be important (i.e., their coefficients are statistically significant) in determining the future market share of low-emitting technologies. In table 4 we present the estimated coefficients of the model parameters. We focus our attention on β_3, the level of increasing returns to scale for the low-emitting technology; ϑ, the learning rate among agents; the risk aversion λ, and the heterogeneity of the agents' preferences υ. We observe that there are important interactions among these parameters. Some parameters enhance each other. For instance, increasing returns to scale and the sample size work together to increase the market share. The level of heterogeneity of the population also enhances the positive effect of these two parameters. This is so because heterogeneity increases the number of potential early adopters of the new technology. However, more heterogeneity is not always good since it can also slow the later stages of diffusion by increasing the number of late adopters, that is, agents with a strong preferences for traditional technologies or high risk aversion. Also, the negative effect of high uncertainty (variance) with respect to the new technology increases with the level of heterogeneity.

It is important to note that some of the effects are nonlinear. For instance the negative square term for the sample size parameter suggests that the positive effect diminishes after a given value where the total effect reaches a maximum. This result is consistent with Lane [29] who finds that a higher sample size in the case of Bayesian optimizers does not necessarily improve social outcomes. In addition, some parameters counteract each other. For example, the positive effect of the sample size is reduced in the presence of increasing returns in the production of medium emitting technologies. Our results also suggest that the switching cost γ is a critical determinant of the variance of the future market share. Low switching costs increase the turnover between technologies and make the model dynamics extremely sensitive to early stochastic variations (see Grubb et al. [30], for a discussion on the role of inertia in climate change policy).

3.3. When technology subsidies are useful

We now examine the performance of the adaptive decision strategies defined above in the sub-region of the landscape of plausible futures spanned by the vector $S = \{\beta_3, \upsilon, \vartheta, \lambda, \kappa_1\}$. The first four dimensions are the most important determinants of the dynamics of technology diffusion in our model and the fifth captures the damages due to climate change. We applied our adaptive decision strategies in a grid of points defined by combinations of the parameters in S, while the other parameters were kept constant at their mean values. This process generates an ensemble of scenarios (see figure 2), that we search to identify the factors that determine the relative performance of one strategy over another. This relative performance is measured by the *regret* of each decision strategy. The regret is defined as the expected difference between the performance of a strat-

Table 4
Econometric estimates of the effect of model parameters in the market share of new emitting technologies (selected interactions).

	Coef.	Std. err.	z	$P > \|z\|$	(95% conf. interval)	
υ	12.49103	3.873437	3.225	0.001	4.899229	20.08282
α	7.32531	1.40877	5.2	0	4.564172	10.08645
λ	−1.629188	6.250245	−0.261	0.794	−13.87944	10.62107
ν_1	−0.1703135	0.3372748	−0.505	0.614	−0.8313599	0.490733
ν_2	0.1379413	0.1301188	1.06	0.289	−0.1170869	0.3929694
ν_3	−0.069668	0.0334578	−2.082	0.037	−0.1352442	−0.0040919
β_1	6.41204	36.27177	0.177	0.86	−64.67933	77.50341
β_2	36.95832	35.85846	1.031	0.303	−33.32297	107.2396
β_3	−1.859102	5.218475	−0.356	0.722	−12.08713	8.368921
ϑ	71.76872	39.879	1.8	0.072	−6.392691	149.9301
ϕ_g	0.8488508	42.04792	0.02	0.984	−81.56355	83.26125
γ	−2.061293	1.936609	−1.064	0.287	−5.856978	1.734392
T^{life}	0.1041627	0.06325	1.647	0.1	−0.019805	0.2281305
υ^2	−0.5951488	1.641682	−0.363	0.717	−3.812787	2.622489
α^2	−4.799885	1.266893	−3.789	0	−7.28295	−2.31682
λ^2	3.463986	3.506615	0.988	0.323	−3.408854	10.33683
β_1^2	−531.9518	468.2693	−1.136	0.256	−1449.743	385.8392
β_1^2	131.5067	421.3547	0.312	0.755	−694.3334	957.3468
β_3^2	8.934308	5.294473	1.687	0.092	−1.442669	19.31128
ϑ^2	−1037.171	227.8347	−4.552	0	−1483.719	−590.6235
ρ^2	0.0380914	0.0142229	2.678	0.007	0.0102151	0.0659678
γ^2	2.994101	1.254874	2.386	0.017	0.5345937	5.453609
$T_3^{\text{life}^2}$	−0.0000981	0.0010801	−0.091	0.928	−0.002215	0.0020189
$\upsilon\alpha$	−10.26669	2.082545	−4.93	0	−14.3484	−6.184974
$\alpha\nu_3$	−0.3600053	0.066708	−5.397	0	−0.4907506	−0.22926
$\upsilon\beta_3$	11.28381	4.648449	2.427	0.015	2.173014	20.3946
$\upsilon\vartheta$	220.9146	41.36304	5.341	0	139.8445	301.9847
$\upsilon\rho$	−1.146068	0.2303033	−4.976	0	−1.597454	−0.694682
$\lambda\rho$	−0.8403946	0.4178413	−2.011	0.044	−1.659348	−0.0214408
$\beta_1\rho$	10.21623	2.534978	4.03	0	5.247767	15.1847
$\beta_2\vartheta$	−949.2111	426.5594	−2.225	0.026	−1785.252	−113.17
$\beta_3\vartheta$	232.8181	71.52681	3.255	0.001	92.62811	373.0081

egy at a particular uncertainty space point – measured as the present value of the GDP in a particular scenario from 1995 to 2045 – and the performance of the strategy with the best performance at that uncertainty space point (see [27] for an example of the use of regrets to find strategies robust across multiple scenarios).

The results suggest not only that increasing returns to scale and learning feedbacks cause the Combined strategy to dominate the Tax-Only strategy (not a surprising result), but also that this dominance is conditioned by the degree of risk aversion of the population of agents and the level of heterogeneity. The role of heterogeneity is an interesting phenomenon. Indeed, if the mean of the distribution of preference does not change, one should not expect changes in aggregate behavior (i.e., the mean representative agent does not change). However, in the presence of interactions between agents, changes in the variance of the distribution do affect aggregate behavior and therefore policy choices.

We illustrate these ideas in figure 4, where we compare the regret of the Tax-Only strategy policy (lighter surface) with the regret of the Combined strategy (darker surface), as a function of the heterogeneity and risk aversion of the agents. For this figure we have assumed moderate increasing returns to scale ($\beta_3 = 0.2$), a moderate level of social interactions ($\vartheta = 5\%$), and moderate damages due to climate change ($\kappa_1, \kappa_{\text{ROW}} = 1.3$). The figure shows that the Tax-Only strategy is preferable in a world where the agents are homogenous ($\upsilon = 0$). As the heterogeneity of agents' preferences increases, the Combined tax and subsidy strategy quickly becomes more attractive. However, the effect is non-linear; the difference between the Tax-Only and the Combined strategies becomes independent of heterogeneity once the heterogeneity is larger than about $\upsilon = 0.4$, when the preference of an extremist agent is about twice the preference of an average agent. Heterogeneity favors the Combined strategy because it creates a number of potential early adopters that are well disposed to use the new low-emitting technology. The subsidy encourages many of these agents to adopt, thus generating learning and cost reductions above and beyond the social benefit gained by any individual adopting agent.

The agents' risk aversion also affects the choice between the Tax-Only and Combined strategies. Both high and low levels ($\lambda < 0.1$ and $\lambda > 0.4$) increase the desirability of the latter strategy, while intermediate levels of risk aversion favor the tax. When agents are risk neutral, the low-emitting technologies are likely to diffuse independent of any adverse initial experience with the technology. However, the

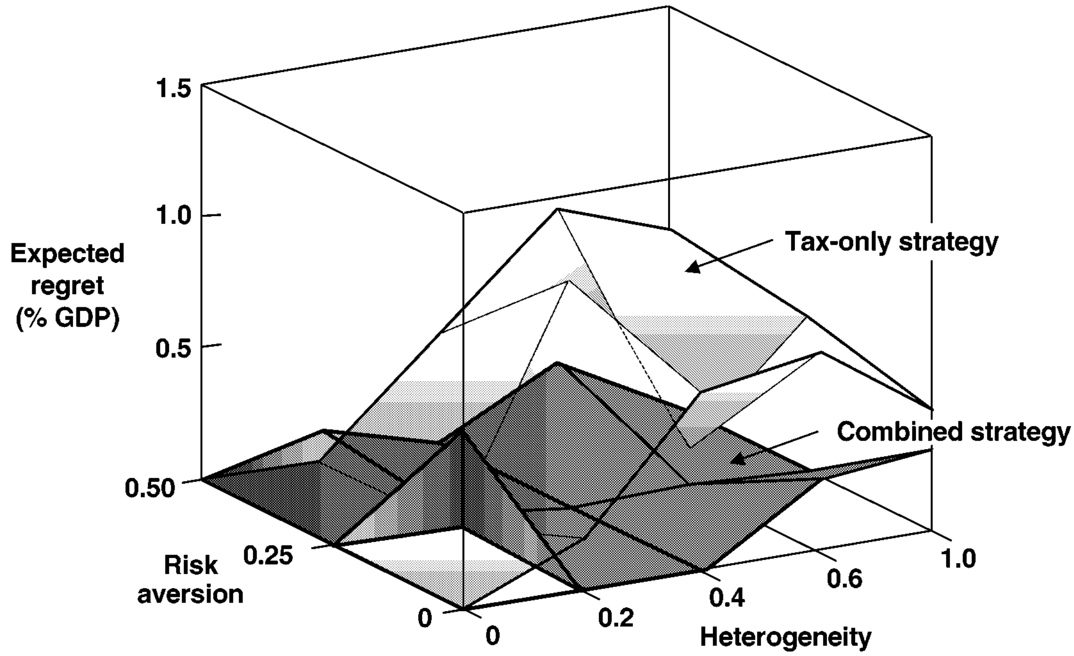

Figure 4. Expected regret of the Tax-Only (lighter surface) and Combined tax and technology subsidy (darker surface) adaptive-decision-strategies. The surfaces are displayed as a function of the risk aversion and heterogeneity of the agents' population. All other input parameters are held constant at their mean values.

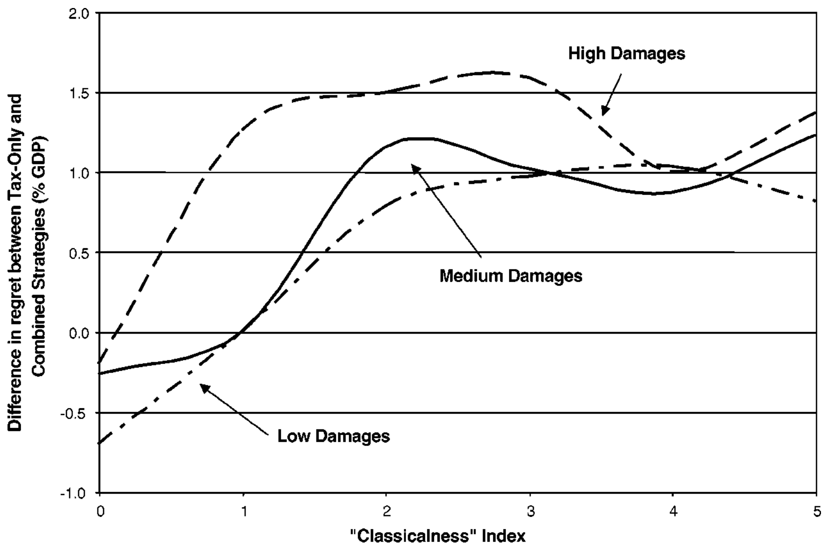

Figure 5. Difference in expected regret between the Tax-Only and Combined strategies as a function of the damages due to climate change, given by the damage function coefficient, and the "classicalness" index, a composite index of the level of increasing returns to scale, the speed of learning, the risk aversion and the heterogeneity of the agents' preferences.

subsidy speeds up this diffusion at a very low cost (given increasing returns to scale). When the agents are highly risk adverse, adverse early experience with a technology from a small number of agents may greatly delay diffusion of the new technology. In such cases, the subsidy is more expensive, but may provide the important benefit of launching the diffusion process as the *uncertainty* regarding the characteristics of new technologies diminishes.

We have similarly compared the performance of the Tax-Only and Combined strategies in other subsets of S. In each case we find, in accord with most economic analyses, that the Tax-Only policy is best in what we can think of as the classical limit, where heterogeneity, increasing returns, and endogenous learning are all small. When these effects are non-negligible, however, we find that the Combined strategy is best. We summarize these results in figure 5, which compares the performance of the Tax-Only and Combined strategies as a function of the damages due to climate change, coefficient of the damage function, and the "classicalness" of the economy, a composite index of

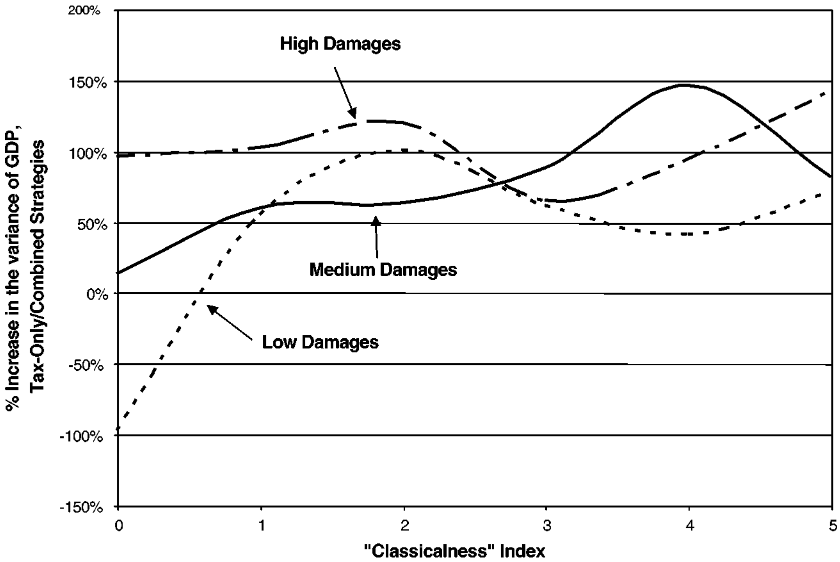

Figure 6. Percent increase in the variance of the expected GDP under the Tax-Only strategy relative to the Combined strategy as a function of the damages due to climate change, given by the damage function coefficient, and the "classicalness" index.

the extent to which heterogeneity, risk aversion, increasing returns, and endogenous learning are small.[4]

Our results are consistent with those of Goulder and Schneider [14,31], who examine greenhouse gas abatement policies using a general equilibrium model for the United States that takes into account incentives to invest in research and development, knowledge spillovers, and the functioning of research and development markets. They find that the tax should be accompanied by a subsidy *only* when there are spillovers benefits from research and development. Here we show that this result can be more general. Many types of spillovers, such as those resulting from increasing returns to scale, network externalities, or *non-R&D knowledge spillovers* from users to non-users of new technologies may justify a subsidy. However, the presence of spillovers is not a sufficient condition. The level of spillovers is an important consideration. Furthermore the level of spillovers that justifies the subsidy depends on the degree of heterogeneity of agents preferences and their attitude towards risk. Our results suggest, all other things being equal, that the critical level to justify the subsidy decreases, but non-monotonically, as heterogeneity and risk aversion increase.

We also find that the Combined strategy reduces the variability of GDP, in addition to increasing its expected value, relative to that of the Tax-Only strategy in the non-classical regions of the uncertainty space, as shown in figure 6. This stochastic behavior of our results for any given set of parameter input values is due to the random distribution of characteristics across the initial agent population and each agent's random choice of other agents to query for information. The Tax-Only strategy fails to generate a pool of early adopters that increase the quantity and quality of the information that would allow all the agents to make better technology adoption choices. In cases with significant potential for learning and increasing returns, the dynamics of the economy thus become more sensitive to chance, in particular the initial distribution of agents' preferences and expectations about new technologies, for the Tax-Only than the Combined strategy.

3.4. Robustness of the adaptive decision strategies

A decision-maker's choice between the Tax-Only and Combined strategies should depend on their assessment of the likelihood of a significant potential for learning and increasing returns, as well as the likelihood of observing heterogeneous preferences and risk aversion. Yet, these likelihoods are themselves important uncertainties. Hence, in our analysis we proceed backwards, by asking the question: what likelihoods would cause a decision-maker to prefer one policy to another? An adaptive decision strategy that has low expected regrets over most regions of the *probability space* is said to be a *robust* strategy.

Figure 7 compares the expected value of the Combined and Tax-Only strategies as a function of the probability of observing a non-classical world, CI $= 0$, and the probability of observing low damages, given by $\kappa_1 = 1$ in equation (1), equivalent to a 0.3% reduction in GDP resulting from a doubling in atmospheric concentrations relative to pre-industrial levels. We created this figure by weighting the points in figure 5 with classical index 0 by $w_1 \in [0, 1]$, and those points with classical index 1 through 5 by $(1 - w_1)/5$. Similarly we weighted the points with damage index 1 by $w_2 \in [0, 1]$ and those points with damage indices 2 and 3 by $(1 - w_2)/2$ (see [7]). The region in the lower-left-hand corner represents those expectations about the state of the world such that the expected value of

[4] This classicalness index CI $= \{0, i, \ldots, 4, 5\}$ characterizes the four parameters $(\beta_3, \vartheta, \upsilon, \lambda)_i$, $i \in$ CI, where each parameter z of the four is given by: $z_i = z_{\min} + i(\frac{z_{\max} - z_{\min}}{5})$, and $[z_{\min}, z_{\max}]$ determines the range of variation of the parameter.

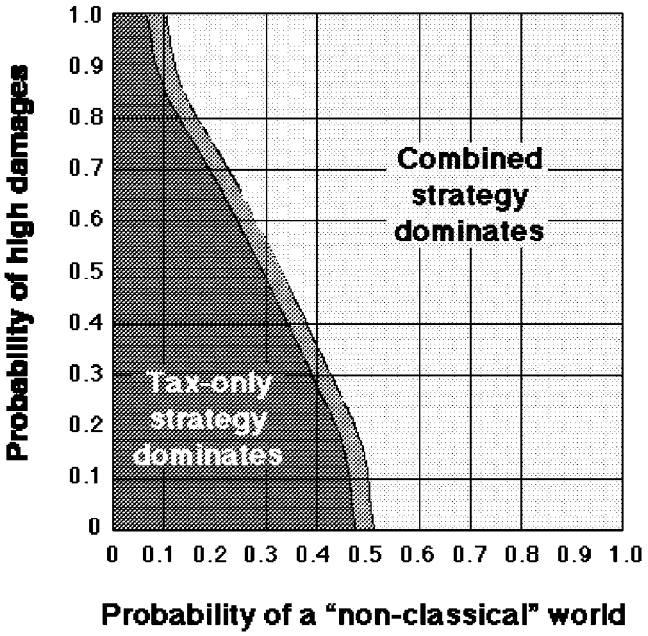

Figure 7. Regions in probability space where the expected GDP resulting from the Tax-Only strategy is greater than that of the Combined strategy as a function of the probability of a classical world and low damages due to climate change, as defined in the text.

the Tax-Only strategy is favored and the region in the upper-right represents those expectations such that the Combined strategy is favored.

It is difficult to make definitive policy statements based on the results in figure 7 because we cannot directly relate all the components of our classical index to the real world. In particular, our model of learning due to random sampling among agents is almost certainly a poor representation of the actual networks that exist among economic actors. Of the components of our classicalness index, only the increasing returns parameter and probably the degree of risk aversion can be supported by real data. With respect to the former we find that increasing returns to scale in the order of 0.1, which is one sixth of that estimated for natural gas turbines [32], justify the use of subsidies, *if* the population is heterogeneous and risk averse. In figure 4, the highest level of risk aversion, 0.5, corresponds to a curvature of the utility function that is consistent with the one estimated by models of health insurance demand [33–35].

The interpretation of the degree of heterogeneity and the sample size is more complicated. Nonetheless we try to provide some insights. In figure 4, a level of heterogeneity of 0.4, implies for example that if the distribution of the risk aversion parameter across the real population has a mean of 0.5 then the standard deviation would be approximately 0.12. So, our assumptions regarding risk aversion and heterogeneity, seem consistent with empirical evidence. Finally, in our simulation a sample size of approximately 1% of the population (in figure 5 the sample size varies between 0 and 12%) implies that within 5–7 years, an agent's estimate of the variance of its expectations will converge to zero. This can be interpreted as the agent been extremely confident about its beliefs about a given product or technology. While we don't have evidence about how long it takes on average for real agents to feel confident about the quality of a given product, we believe that five to seven years could be a high range for this time period.

These heuristic considerations, suggest that the conditions necessary for the Combined strategy to be preferred to the Tax-Only strategy may exist in the real world, and that the Combined strategy may be the most robust strategy.

3.5. Back search

These results are based on a search over the restricted region of the uncertainty space given by the set S. Yet, it may be the case that our conclusions cannot be extended to other regions of the uncertainty space. To test this possibility we implement a search over the entire range of plausible futures defined by the parameters in table 1, looking for points for which the Tax-Only strategy is superior to the Combined. We use a genetic algorithm search routine similar to that used to find the landscape of plausible futures in figure 3. While such a search cannot guarantee the absence of counter examples to our conclusions, it provides strong corroborating evidence that we are interpreting our model results fairly. We examined roughly 5,000 uncertainty points. The only cases found where the Tax-Only strategy performs better than the Combined, other than those shown in figure 5, are: (a) cases with no learning, (b) cases with small levels of learning (e.g., the point $\beta_3 = 0.05$, $\vartheta = 0.01$) *and* homogenous preferences ($\upsilon = 0$), or (c) cases where the economic-growth-elasticity of the cost of energy (ϕ_{xg} in equation (1)) is equal to zero. At other points, usually with low switching costs, we observe arbitrarily high variability in the outcome variable, so that the differences in performance have not been able to be detected with the current number of Monte Carlo simulations.

4. Conclusions

Many governments are exploring policies designed to encourage potential early adopters to deploy new carbon emissions-reducing technologies. Such policies range from tax credits and subsidies on certain classes of technologies, to early action credits, to the Clean Development Mechanisms under the Kyoto Protocol. This study suggests that under particular, likely to be met, conditions such actions can be an important part of a successful climate change strategy. We find that a Combined strategy of carbon taxes and technology incentives, as opposed to carbon taxes alone, is the best approach to greenhouse gas emissions reductions if society has even modest expectations that the diffusion of new, emissions-reducing technology will be important in reducing the future costs of emissions abatement and that there are broad social benefits to the early adoption of such technologies by a small number of early users. In our model, these broad social benefits arise from two main sources: (a) cost reductions due to increasing returns to scale, and (b) improvements in the quantity and quality of the flows of information about the performance of new technologies that lead economic agents to better coordinate their choices. If such factors are important, early adopters generate cost reductions and public information with social value far beyond that which can be captured by the early adopters themselves. If the agent population also has heterogeneous preferences – that is, some agents are more prone to adopt a new technology than others are – and if some of these agents are risk averse, then technology incentives may provide an important complement to carbon taxes. While taxes raise costs for all agents, those inclined towards the new technologies as well as those highly disinclined towards them, subsidies focus on only those most inclined towards early adoption. This difference is of course the basis for the political popularity of incentives over taxes. This study suggests that under certain conditions, incentives can be not only popular, but cost-effective as well.

We come to these conclusions using two analytic innovations, an agent-based model of technology diffusion and an exploratory modeling approach to decision-making under conditions of extreme uncertainty. Agent-based models provide a convenient framework for modeling heterogeneity and imperfect information. Using such a model, we can examine the effects of endogenous technology diffusion, which are often missed in analytic studies of climate change policy and which are crucial to distinguish between the Tax-Only and Combined strategies considered here. While agent-based models are finding wide use in many studies of social systems, they have to date had limited impact on policy research because they do not fit easily into the standard tools of policy analysis. Thus, agent-based models are most often used to illuminate the evolution of several possible paths into the future as a means of helping decision-makers build insight about a problem. While useful, this approach is anecdotal but not systematic. One never knows what new and perhaps contradictory insights reside in the cases not considered. Exploratory modeling provides a fundamentally different alternative to this “flight-simulator” approach. We let the computer search through a huge number of plausible scenarios generated by the agent-based model, looking for those that distinguish one policy choice from another. We are thus able to place each scenario in context and move from insights to arguments. In this study, we map out those conditions that favor a Combined emissions-reduction strategy over a Tax-Only strategy and thus identify the factors that would justify the former policy. We believe that this approach is a powerful and widely applicable way to exploit the information contained in agent-based models.

Nonetheless, the models and data used in this study are very crude and thus can offer only general recommendations to decision-makers. Significant steps remain before our results can be translated into more specific policy recommendations. For instance, our model of learning due to random sampling among agents is almost certainly a poor representation of the actual networks that exist among economic actors. Thus, while we can relate our model of increasing returns to data on actual technologies and say that learning factors similar to those for natural gas turbines justify a Combined strategy, we cannot relate our model parameters on learning to any real world data. Similarly, our model of technology is sufficiently aggregated so that it is difficult to relate our subsidy to specific recommendations for spending levels or other questions about implementation. Thus, while we argue that technology incentives are likely to be an important part of any climate change strategy, we have not answered the question as to whether the subsidies currently in place and proposed by governments are the correct type of technology incentive or whether the monetary amount is sufficient, too much, or too little. We believe, however, that the methods laid out in this paper provide a powerful framework for addressing such questions.

Appendix: agent decision-making

As discussed in section 2, the model of technology diffusion used in this study pays particular attention to the process by which agents make decisions about the type of energy producing technology they employ and the amount of energy they use. This appendix describes in detail how we model these choices.

A1. Energy consumption

Global emissions of carbon dioxide in our model are determined by the energy intensities, $n_{g,j}(t)$, for each agent, as seen in equation (3). Each agent chooses to consume the amount of energy that will minimize its cost for producing one unit of output, so that its energy intensity depends on both its choice of energy technology and the (exogenous) state of conservation technology used in region g. Assuming that agents have a CES production function (as in [36]), we write

$$n_{g,j}(t) = \frac{a_{g,j}(t)}{\left\{[1 - \overline{S}_j(T_i^{\text{adopt}})]C_j(T_i^{\text{adopt}}) + \mathit{Tax}(t)m_j\right\}^{\varepsilon}}, \tag{A1}$$

where ε is the elasticity of substitution, and $a_g(t)$ is an energy-efficiency coefficient proportional to the Autonomous Energy Efficient Improvement Index (AEEI) used in other climate change studies. As defined after equation (A8) the terms in brackets represent the cost in constant 1997 dollars (inclusive of all taxes and subsidies) of producing one unit of energy[5] with technology j adopted in the year T_i^{adopt}.

[5] This formulation implies that the prices of output and other inputs remain constant.

A2. Technology choice

As shown in equation (4), the agents choose among energy technologies in order to maximize their intertemporal expected utility. The long lifetimes of each technology create two distinct decision problems for the agents, depending on whether or not an agent is currently using a technology. An agent may not be using a technology because it is a new agent, or because the lifetime of its previous technology has expired. In either case, agent i will choose technology j such that

$$\langle U_{i,g,j}(t, T_i^{\text{life}})|t\rangle > \langle U_{i,g,j'}(t, T_i^{\text{life}})|t\rangle, \quad \forall\, j \neq j'. \tag{A2}$$

More commonly, an agent currently using some technology must decide whether to switch to another. To solve this problem, we define $\langle U_{i,g,j\to j'}(t,\tau,T_i^{\text{life}})|t\rangle$ as the expected utility agent i, at time t, estimates it will derive from using technology j until time τ, and then using technology j' until time T_i^{life}. This expectation is given by

$$\begin{aligned}\langle U_{i,g,j\to j'}(t,\tau,T_i^{\text{life}})|t\rangle &= \langle U_{i,g,j}(t,\tau)|t\rangle \\ &+ (1+r)^{-(\tau-t)}\langle U_{i,g,j'}(\tau, T_i^{\text{life}})|t\rangle + \Gamma_{i,g,j}(\tau),\end{aligned} \tag{A3}$$

where $\Gamma_{i,g,j}(\tau)$ is the cost of abandoning technology j at time τ before its useful life is expired, and r is the discount factor. Thus, the present value at time t of the expected utility of the technology switch $j \to j'$ at time τ is given by the expected utility of using technology j from time t to time τ, plus the discounted expected utility of using technology j' from τ up to T_i^{life}, plus the switching penalty. We take T_i^{life} as the planning horizon, since at that time technology j will need to be replaced.

Thus, agent i using technology j will shift to technology j' if only and only if

$$\begin{aligned}&\langle U_{i,g,j'}(t, T_i^{\text{life}})|t\rangle > \langle U_{i,g,j}(t, T_i^{\text{life}})|t\rangle \quad \text{and} \\ &\langle U_{i,g,j'}(t, T_i^{\text{life}})|t\rangle > \langle U_{i,g,j\to j''}(t,\tau, T_i^{\text{life}})|t\rangle, \\ &\quad \forall\, j'' \in J,\ t < \tau < T_i^{\text{life}}.\end{aligned} \tag{A4}$$

The two conditions in equation (A4) are the profitability and arbitrage conditions, respectively. The former implies that an agent will only adopt a new technology j' if its expected utility is higher than that of currently employed technology j. The latter implies that an agent will adopt a new technology at time t if and only if there is no alternative technology or switching time that will generate a higher expected utility. The first term dominates when costs are rising; the second when they are falling [3,37,38].

In our simulations at each time period we solve numerically this intertemporal problem for every agent on the basis of its expectations about performance and costs. In sections A2.1 and A2.2 we will describe how, respectively, the agents construct these expectations.

A2.1. Learning about technologies

Different mechanism have been implemented to simulate learning within socioeconomic systems: genetic algorithms or classifier systems [19,39–42]; least squares learning [43–45], or Bayesian updating [46–50]. The appropriate choice is of course dependent on what is being learned. In our study agents learn about the performance of alternative technologies. So, we use an information contagion model, similar to [29] and [46], to describe the flow of information about the performance of new technologies through a population of potential users. This model assumes that agents form their expectations about technology performance from Bayesian updates based on their own past experience (if any) using the technology, and from the experience of other agents that have used it. As mentioned previously, our model differs from Arthur and Lane's in that we consider a heterogeneous population of agents, use more than two technologies, and do not assume that the relative costs are constant over time.

We begin by defining $q_{i,g,j}$ as the time-independent[6] performance agent i will achieve if it uses technology j. To capture agents' heterogeneity we assume that performances achieved for any technology j are normally distributed among the agents, so that $q_{i,g,j} \sim N(\bar{q}_{g,j}, \nu_{g,j})$, where $\bar{q}_{g,j}$ and $\nu_{g,j}$ are the mean and variance of the distribution. At any time t, the agents have imperfect information about the performance of the technologies where $\mu^q_{i,g,j}(t)$ and $\nu^q_{i,g,j}(t)$ are the mean and variance of agent i's estimate of $q_{i,g,j}$. Each agent improves its estimates by observing its own and other agents' previous experience using the technology. Thus, we assume that agent i's observation of the performance of technology j used by agent i' is given by

$$Z_{i,i',j}(t) = q_{i',g',j} + \omega_q(t), \tag{A5}$$

where $\omega_q(t) \sim \{q_{i,g,j} - q_{i',g',j}, \rho_q[1 + (q_{i,g,j} - q_{i',g',j})^2]\}$ is a normally distributed random variable representing the measurement error. We assume for simplicity that ρ_q is the same for all the agents and all the technologies. When $i = i'$, equation (A5) reduces to the form used by Arthur and Lane in their study of information flow across a homogenous population of agents. We have further assumed here that in a heterogeneous population of agents with different characteristics, when estimating how another agent's experience with a technology would apply to itself, an agent can observe these characteristics and partially compensate for them; however the error in the observation increases proportionally to the square of the difference between the agents. Based on its observations and its prior expectations, each agent can estimate the performance it expects from the technology using a discrete-time Kalman filter [51]

$$\begin{aligned}\mu^q_{i,g,j}(t+1) &= \mu^q_{i,g,j}(t) \\ &+ \left\{\frac{\nu^q_{i,g,j}(t)}{R + \nu^q_{i,g,j}(t)}\right\}\left[\mu^q_{i,g,j}(t) - Z_{i,i',j}(t)\right], \\ \nu^q_{i,g,j}(t+1) &= \left\{\frac{R\nu^q_{i,g,j}(t)}{R + \nu^q_{i,g,j}(t)}\right\},\end{aligned} \tag{A6}$$

[6] We are thus neglecting the ability of a technology to improve over time. We capture similar effects, however, through the effects of learning-by-doing on the price.

where $R = \rho_q[1 + (q_{i,g,j} - q_{i',g',j})^2]$. Note that the knowledge agents have about technologies they have not used comes exclusively from observations of other agents' experience.

This formalization addresses the problem of how agents may use the information generated by their own experience and the experience of other agents to learn about a given technology. However, it does not say anything about how much information a given agent demands or, even more importantly, how he/she finds the providers of such information. In this first version of the model we make two simplifying assumptions. First, we work within a global interaction framework [52]. Thus every agent has the same probability to be interviewed by another agent: interactions are chance driven. This assumption rules out the possibility that agents create networks that increase the efficiency of the learning process. These networks are known to be important determinants of the diffusion process [53], and have important policy implications since they reduce coordination failures.

Our second assumption is that the number of social interactions (the demand for information) is exogenously determined and is the same for all the agents. In practice the number of interactions that an agent decides to have would depend on the value of the additional unit of information (e.g., by how much it reduces its uncertainty), and its cost. Indeed, the fact that an agents' value of information does not take into account the social value is an important source of market failure. To compensate for this shortfall in this study we consider a wide range of interactions thresholds.

As we discuss in section 4, the type of learning process that we have described may generate externalities which have important policy implications. First, because late adopters learn from the experience of early adopters, there is an externality deriving from the adoption of technology by early users. Second, expectations may be inaccurate. When the sample size is small, the potential for initial, abnormally bad experiences with a new technology can significantly delay its subsequent diffusion and therefore increase abatement costs.

A2.2. Costs dynamics and agents' forecasts

An important and well-documented characteristic of the diffusion of new technologies is the decline in costs as the number of adopters increases [20,53,54]. Hence, we assume that the actual production cost of each energy technology drops as the number of users increases according to the standard logistic expression

$$C_{g,j}(t) = C_{g,j}(0)N_j(t)^{-\beta_j}, \tag{A7}$$

where $C_{g,j}(0)$ is the known first-unit cost, $N_j(t)$ is the cumulative number of units of the technology that have been adopted by year t, and β_j is the unknown learning factor. A β_j greater than zero implies increasing returns to scale resulting from learning by doing.

Agents' technology choices are in part based on their expectations about future costs. These expectations are based on observations of past costs. More precisely, agents generate expectations about the learning coefficient and the future number of users of a given technology through Bayesian updating. For simplification we have assumed that there is some common forecast, freely available to all the agents. This forecast is generated by a discrete Kalman filter [51], applied to the log linear equivalent of equation (A7).

On the basis of this common forecast agents compute $\langle C_{g,j}(T_i^{\text{adopt}}, T_i^{\text{life}}|t)\rangle$, the expected present value in year t of the cost of producing one unit of output over the period T_i^{adopt} to T_i^{life} with technology j. Noting that the cost of the energy needed to produce one unit of output in some future year τ is given by the product of $n_{g,j}(\tau)$ and the annual energy cost, we can use equation (A1) to write the expectation of this cost as

$$\begin{aligned}
&\left\langle C_{g,j}\left[C_{g,j}\left(T_i^{\text{adopt}}\right), T_i^{\text{adopt}}, T_i^{\text{life}}\middle|t\right]\right\rangle \\
&\quad = \left\langle \sum_{\tau=T_i^{\text{adopt}}}^{T_i^{\text{life}}} (1+r)^{-(\tau-t)} \right. \\
&\quad \left. \times \frac{a_{g,j}(\tau)}{\{[1 - S_j(T_i^{\text{adopt}})]C_{g,j}(T_i^{\text{adopt}}) + \mathit{Tax}(\tau)m_j\}^{\varepsilon-1}} \right\rangle, \quad \text{(A8)}
\end{aligned}$$

where $C_{g,j}(T_i^{\text{adopt}})$ is the annual cost per unit of energy for technology j adopted in the year $T_i^{\text{adopt}}, S_j(T_i^{\text{adopt}})$ is the price subsidy available for technology at time of adoption, and $\mathit{Tax}(\tau)$ is a tax per unit of CO_2 at time τ. In this study we assume each agent has perfect information of all the terms in equation (A8) but $C_{g,j}(T_i^{\text{adopt}})$.

Notice that in this model increasing returns to scale make the diffusion process sensitive to the adoption decisions of the early users. Also, pessimistic initial price expectations due randomness can delay or abort adoption. The same is true of optimistic expectations; agents will prefer to wait until the price drops further before using the new technology.

Acknowledgements

The authors would like to thank Steve Bankes, Giovanni Dosi, Agnes Festres, Robert Klitgaard, Julia Lowell, David Ortiz, Steven Popper, Susan Resetar, and two anonymous reviewers for their helpful comments and suggestions. This research was supported in part by the US National Science Foundation under grant DMS-96343000 and by the US Department of Energy under grant DE-FG03–96ER62282.

References

[1] L. Cohen and R. Noll, *The Technology Pork Barrel* (The Brookings Institution, Washington, DC, 1991).

[2] S.N. Durlauf, in: *The Economy as an Evolving Complex System II*, Vol. 27, eds. B.W. Arthur, S.N. Durlauf and D.A. Lane (Santa Fe Proc., NM, 1997).

[3] P. Stoneman and P. Diederen, The Economic Journal 104 (1994) 918–930.
[4] S.C. Bankes, Exploratory modeling for policy analysis, Operations Research 41(3) (1993) 435–449 (also published as RAND RP-211).
[5] S.C. Bankes and J. Gillogly, in: Exploratory modeling search through spaces of computational experiments, *Proc. 3rd Annual Conf. Evolutionary Programming*, eds. A.V. Sebald and L.J. Fogel (World Scientific, Singapore, 1994) pp. 353–360.
[6] R.J. Lempert and M.E. Schlesinger, Robust strategies for abating climate change, Climatic Change, forthcoming.
[7] R.J. Lempert, M.E. Schlesinger and S.C. Bankes, When we do not know the costs of the benefits: adaptive strategies for abating climate change, Climatic Change 33 (1996) 235–274.
[8] N. Mattsson, Internalizing technological development in energy systems models, Thesis paper, Energy Systems Technology Division, Chalmers University of Technology, Goteborg, Sweden (1997).
[9] A. Grübler and A. Gritsevskii, Environmentally Compatible Energy Strategies Project, IIASA, Laxenburg, Austria (1998).
[10] A. Gritsevskii and N. Nakicenovic, Modeling uncertainty of induced technological change; paper presented at the Pew Center Workshop on Economics and Integrated Assessment of Climate Change, Washington, DC (21–22 July 1999).
[11] C. Azar and H. Dowlatabadi, A review of technical change in assessment of climate policy, *Annual Review of Energy and Environment* (1999).
[12] A. Grubler, N. Nakicenovic and D.G. Victor, Dynamics of energy technologies and global change, Energy Policy 27 (1999) 247–280.
[13] L.H. Goulder and K. Mathai, Optimal CO_2 abatement in the presence of induced technological change, Working paper, Stanford University (1997).
[14] L.H. Goulder and S.H. Schneider, Induced technological change, crowding Out, and the attractiveness of CO_2 emissions abatement, *Resource and Environmental Economics* (1998), submitted.
[15] W. Cline, *The Economics of Global Warming* (Institute for International Economics, Washington, DC, 1992).
[16] W.D. Nordhaus and Z. Yang, The American Economic Review 86(4) (1996).
[17] M. Aoki, American Economic Review (1986).
[18] M. Aoki, Journal of Economic Literature (1990).
[19] L. Marengo, Journal of Evolutionary Economics 2 (1992).
[20] S. Davies, *The Diffusion of Process Innovations* (Cambridge University Press, Cambridge, 1979).
[21] G. Ellison and D. Fudenberg, Quarterly Journal of Economics 440 (1995) 93–125.
[22] A. Grubler, *The Rise and Fall of Infrastructures* (Physica Verlag, Heidelberg, 1990).
[23] J.H. Miller, Active Nonlinear Tests (ANTs) of complex simulations models, Management Science 44(6) (1998) 820-30.
[24] T. Morita and H.-C. Lee, Appendix to emissions scenarios database and review of scenarios, *Mitigration and Adaptation Strategies for Global Change*, Vol. 3 (1998) pp. 121–131.
[25] N. Nakicenovic, V. Nadejda and T. Morita, Emissions scenarios database and review of scenarios, *Mitigration and Adaptation Strategies for Global Change*, Vol. 3, (1998) pp. 95–120.
[26] L. Taylor, *Income Distribution, Inflation and Growth* (MIT Press, 1990) ch. I.
[27] R.J. Lempert, M.E. Schlesinger, S.C. Bankes and N.G. Andronova, The impacts of climate variability on near-term policy choices and the value of information, Climatic Change, forthcoming.
[28] M. Aoki, The Japanese Economic Review 46 (1995) 148–165.
[29] D.A. Lane, in: *The Economy as an Evolving Complex System II*, Vol. 27, eds. B.W. Arthur, S.N. Durlauf and D.A. Lane (Santa Fe Proc., 1997).
[30] M. Grubb, T. Chapuis and M. Ha-Duong, Energy Policy 23 (1995) 417–431.
[31] S.H. Schneider and L.H. Goulder, Nature 389 (1997).
[32] L. Christiansson, Diffusion and learning curves of renewable energy technologies, Working paper, IIASA, Luxembourg, Austria (1995).
[33] J.W. Pratt, Econometrica 32 (1991) 122–136.
[34] M.S. Marquis and M.R. Holmer, Choice Under Uncertainty and the Demand for Health Insurance, RAND N-2516-HHS, Santa Monica (1986).
[35] C.J. Cichetti and J.A. Dubin, Journal of Political Economy 102 (1994) 169–86.
[36] P. Hoeller and J. Coppel, Energy taxation and prices distortions in Fossil fuel markets: some implications for climate change policy, Working paper, Economics Department, no. 110, OECD, Paris (1992).
[37] A.B. Jaffe and R.N. Stavin, Energy Journal 15 (1994) 43–65.
[38] A.B. Jaffe and R.N. Stavins, Journal of Environmental Economics and Management 23 (1995).
[39] J.H. Holland et al., *Induction* (MIT Press, Cambridge, Mass, 1986).
[40] J.H. Holland, *Hidden Order How Adaptation Builds Complexity* (Addison-Wesley Publishing Co., 1995).
[41] L. Marengo, Revue Internationale de Systemique 7(5) (1993).
[42] L. Dibiaggio and P. Musso, Learning and organizational design; paper presented at a conference on *Economics with Heterogeneous Interacting Agents*, Ancona, Italy (May 1998).
[43] A. Marcet and T. Sargent, American Economic Review, Papers and Proceedings 78 (1988) 168–172.
[44] A. Marcet and T. Sargent, Journal of Economic Theory 48 (1989) 337–368.
[45] A. Marcet and T. Sargent, Journal of Political Economy 97 (1989) 1306–1322.
[46] B.W. Arthur and D.A. Lane, Information constriction and information contagion, SFI economics research program, Working paper, Santa Fe Institute (1991).
[47] L.E. Blume and D. Easley, Journal of Economic Theory 34 (1984) 116–129.
[48] M. Bray and D. Kreps, in: *Arrow and the Ascent of Modern Economic Theory*, ed. G. Feiwel (New York University Press, New York, 1987).
[49] M. Feldman, International Economic Review 28 (1987) 635–650.
[50] R. Townsend, International Economic Review 19 (1978) 481–494.
[51] F. Lewis, *Optimal Estimation* (Wiley, New York, 1986).
[52] A.P. Kirman, in: *The Economy as an Evolving Complex System II*, Vol. 27, eds. B.W. Arthur, S.N. Durlauf and D.A. Lane (Santa Fe Proc., 1997).
[53] P. Stoneman, *The Economic Analysis of Technology Policy* (Clarendon Press, Oxford, 1987).
[54] E. Mansfield, Econometrica (1961).
[55] D.A. Robalino and R.J. Lempert, Applying a multiagent-model of technology diffusion to the analysis of climate change policy; paper presented at a conference on *Economics with Heterogeneous Interacting Agents*, Ancona, Italy (May 1998).
[56] A.P. Bassanini and G. Dosi, Competing technologies, international diffusion and the rate of convergence to a stable market structure, Interim report from the International Institute for Applied Systems Analysis (1998).
[57] World Bank, World Development Report, Washington, DC, 1996.
[58] A. Dean and P. Hoeller, Costs of reducing CO_2 emissions. Evidence from six global model, Working papers, Department of Economics, no. 122, OECD, Paris (1992).
[59] J.M. Burniaux, J. Martin, G. Nicoletti and J. Oliveira, *GREEN, A Multi-Sector, Multi-Region General Equilibrium Model for Quantifying the Costs of Curbing CO_2 Emissions*, A Technical Manual, OCDE, Paris (1992).
[60] J.T. Houghton, G.J. Jenkins and J.J. Ephraums, Report of the intergovernmental panel on climate change, Cambridge University Press (1995).
[61] A.S. Manne and R.G. Richels, *Buying Greenhouse Insurance: The Economic Costs of Carbon Dioxide Emissions Limits* (MIT Press, Cambridge, 1992).
[62] J.W. Tester, D.O. Wood and N.A. Ferrari, eds, *Energy and the Environment in the 21st Century* (MIT Press, Cambridge, 1991).