# A Global Portfolio Strategy for Climate Change Technology Development

*Geoffrey J. Blanford*, John P. Weyant*

*Department of Management Science and Engineering, Stanford University*

**Abstract**

This paper presents a dynamic strategy model for climate change technology research and development (R&D). It is distinguished from previous models in two respects: first, we focus on long-term basic R&D conducted from a social perspective, rather than on short-term applied R&D motivated by individual firm profit; and second, we analyze of the optimal allocation of R&D resources across the technology space, rather than considering only the optimal level of investment. The decision problem is described analytically and optimality conditions are derived. A case study is described for the electricity sectors in the United States and China examining the allocation of R&D investment between programs in renewable technologies, carbon capture and sequestration, and fossil fuel combustion efficiency.

** Corresponding Author: G.J. Blanford, Doctoral Candidate, Department of Management Science and Engineering, Stanford University, 491 Terman Engineering Center, Stanford, California 94305-4026 USA, blanford@stanford.edu*

# 1 Introduction

The field of climate policy has examined many aspects of an optimal response to global climate change as a result of anthropogenic greenhouse gas emissions. Among the difficulties posed by this problem are long timeframes, pervasive uncertainty, extremely high potential costs of mitigation and damages, and a global scale requiring unprecedented institutional coordination. The majority of climate policy studies have focused on the optimal choice of market-based instruments to induce appropriate internalization of the costs of emissions in private energy supply and demand decisions. Many economic and physical processes influence these calculations, but one of the most central is the future state of technology. The process by which carbon-free or carbon-reducing energy technologies are made available is a powerful lever for lessening the economic impact of a given level of abatement, or conversely for reducing the environmental risks of a given level of abatement expenditure. However, the nature of this process is not well understood, and in particular the role of a strategy to control it has not been directly analyzed. A significant body of work has examined the extent to which technological change is accelerated by private actors in response to the introduction of an internalization policy, but this work does not address the current problem faced by a public research and development (R&D) manager[1] under uncertain policy conditions.

This paper proposes a novel formulation of a decision problem in R&D strategy. The problem is motivated by and applied to the context of global climate change, but is characterized in general by an aggregate R&D decision-maker with a social welfare objective, and technology diffusion markets subject to externalities in which private costs are minimized. Of particular interest in this study is the *allocation* of the R&D investment portfolio and its effect on the expected value of research outcomes. We also examine the relationship between technology strategy, defined as the choice of the allocation vector, and market-based internalization policies such as emissions permits or taxes. The dynamic nature of this problem is emphasized in our methodology in order to properly account for several structural features, including the time lag between R&D investment and technology diffusion, the uncertainty in both technological and environmental factors, and the extended time horizon for climate change induced damages.

The results presented in this paper encompass the analytics underlying the model and a discussion of the attributes of the problem we have included. We also describe how the model may be implemented to examine specific elements of a case study for the electric generation sectors in the United States and China. However, a full case study evaluation requires the use of an integrated assessment model, which is not presented here. The main purpose of this paper is to outline a methodology and identify the essential elements of the problem.

# 2 Background

Achieving a stabilization of atmospheric greenhouse gas concentrations without significant technological advance would be very expensive (Hoffert et al. 2002); thus assumptions about technology are critical to integrated assessment modeling of climate policy. Early studies assumed that technology improved exogenously over time, often along one of several tracks, or scenarios. Recently models have been developed to include in the firm's decision space investment in technology development, allowing for endogenous short-term, late-stage R&D to reduce the instantaneous cost of abatement. Important examples include Goulder and Schneider (1999), Goulder and Mathai (2000), Nordhaus (2002), and Popp (2002). In the same vein, other models link technological improvement to deployment as a representation of the learning effect, such as those described in Seebregts et al. (1999), Gritsevskii and Nakicenovic (2000), Grubler and Gritsevskii (2001), and van der Zwaan et al. (2002). Studies that incorporate endogenous technical change (ETC) have demonstrated that enhanced technology development (through either channel) can provide significant cost savings with respect to a fixed policy goal, usually stabilization of atmospheric concentrations.

However, the focus on ETC describes the stimulus of a contemporaneous price signal on learning or R&D by private actors. Long-term, radical developments are not a part of this

[1] We will use the term public R&D to refer not necessarily to R&D conducted by governments, but more generally to any R&D effort that generates knowledge which is treated as a public good. Therefore our manager is a hypothetical social optimizer. These distinctions will be discussed further below.

framework. Since these developments may require many years of basic research, they cannot be motivated by current emissions penalties, as in the ETC models. Their value must be measured in terms of expected prospective benefits in future uncertain conditions. Because of the limited appropriability of such benefits by private firms, and because of innovation market failures in general (see Arrow 1962), research of this type is most often conducted by governments, or with government funding. In this study we consider a broader portfolio that includes any basic research conducted with the intention of public dissemination. Thus any knowledge generated by the research portfolio is a public good and may be adopted by all private industry actors. This formulation complements the motivation behind the research, since it is concerned primarily with developing a technology that itself provides a public good, such as environmental mitigation. A similar idea is captured by the notion of a 'global energy innovation system' in Sagar and Holdren (2002).

Some studies have addressed the issue of basic R&D with public, system-wide benefits. Miketa and Schrattenholzer (2004) use two-factor learning curve to model the influence, and optimal level, of aggregate R&D in the diffusion of photovoltaics and wind. Schock et al. (1999) conduct a modeling experiment in which two technological scenarios are evaluated over a long time horizon and compared to a no-advance baseline. This paper calculated stabilization costs for several targets and found that expected savings with advanced technology measured in the trillions of dollars, but were decreasing in the stabilization target. The study used the expected savings to infer guidelines for current total R&D investment (which were significantly greater than observed levels), but made no recommendations for the allocation of investment to different technologies. This point is most often treated by simply advocating a diversified portfolio, a natural prescription for three main reasons. First, R&D investment is commonly assumed to exhibit decreasing returns to scale, especially in the near term. When marginal benefit is decreasing with investment, a single research program will not dominate in optimality. This result is demonstrated in both Dasgupta and Maskin (1987) and Loch and Kavadias (2003) in the context of private R&D portfolios. Second, diversification can reduce risk under uncertainty. This fundamental result in the finance literature was shown for the case of climate change R&D by Baker et al. (2003), in which an analytic comparative statics approach is used to examine the contributions of different types of technological advance with increases in risk. Third, diversification is important when there exist many different applications for the target technologies. Specific to the case of climate change, abatement can be achieved in many sectors, and many technologies may be used in a given sector, due in part to the heterogeneity of energy supply and demand constraints. While the case for diversification may be clear, the details of an optimally diversified climate change technology portfolio have not been substantiated, which is one of the primary goals of this work.

The analysis of R&D investment is most accurately viewed as the valuation of a real option, as in Dixit and Pindyck (1994). Traditional option analysis has been applied to problems in the firm at the level of individual R&D projects, or groups of projects in which only the most successful can be adopted. For example, Zhu and Weyant (2003) considers investment decisions in new information technologies, and Hsu and Schwartz (2003) uses an options-based model to evaluate the impact of R&D subsidies in the pharmaceutical industry. These models consider projects with fixed costs, fixed probabilities of success, and in some cases multiple stages. The problem is to determine whether to invest in the next stage, or equivalently whether the option value generated by the investment under the current information set is greater than its cost. This setting is inadequate to represent the decision problem posed here for three key reasons. First, we assume, as is the case in the climate change context, that the R&D decision-maker conducts ongoing investment in multiple 'projects' (which may represent broad program areas), each of which can provide net social benefits even when all are successful. Moreover, these benefits can be realized as the result of intermediate developments while the research investment continues, and the relative timing of success in different program areas contributes to their ultimate value. Second, we assume that the costs and success rates are not fixed, but rather are governed by an innovation production function, so that the technological outcome distribution is a function of the level and allocation of investment. Third, because of the externality in the end-use market, future internalization policy is an important determinant of technology diffusion, and therefore the value of technological outcomes. Thus the allocation decision depends on expectations about this policy. While the traditional model allows for exogenous market uncertainty, the uncertainty is this case is actually endogenous, because the state of technology will influence the choice of an internalization policy in the future. In fact, expectations about the future state of technology also influence the current policy-making process. The model presented here extends the existing R&D options approach by incorporating these three factors.

An indirect effect of R&D investment, especially when modeled sequentially as in the proposed option framework, is the ability to learn about which research programs are more effective. Since returns from a given program are uncertain, the problem resembles a multi-armed bandit setting,[2] in which the manager must trade off exploration for exploitation. That is, as certain programs accumulate more successes than others, diversification behavior may give way to concentration on the most promising options. While we do not include in our model the ability to update beliefs about the efficacy of research investment,[3] the dynamic options framework will allow for the indirect benefits of sequential learning to influence investment allocation. Weitzman (1979) provides the following insight based on a reservation price model of sequential search:

> *Low-probability high-payoff situations should be prime candidates for early investigation even though they may have a smaller chance of ending up as the source ultimately yielding the maximum reward when search ends.*

This result is driven by the insensitivity of the reservation price to the lower tail of the reward distribution, since only the maximum sampled reward is captured. Although Weitzman's model describes the optimal sequencing of individual research programs, rather than parallel investigation modeled here, the intuition behind preferring risky options in early periods may also apply in the current context.

# 3 The Model

## 3.1 Formulation

We consider a model of sequential innovation with $T$ periods and $n$ research programs available in each period. Each research program represents a technology, or class of technologies, with the potential to mitigate an environmental externality in the end-use market. The decision variable is the level of aggregate investment in public research allocated to each program in each period, subject to a budget constraint which may or may not be binding. In a given period, each program has set of possible technological outcomes, which depend on the state of the program's technology at the beginning of the period. The probability distribution across the outcome space is conditioned on the level of investment in the program. This relationship will be referred to as the *innovation production function*. A fundamental assumption is that the relationship between research investment and the probability of success is increasing, but exhibits decreasing returns to scale. However, this does not necessarily imply that the relationship between research investment and the *value* of success has decreasing returns to scale, since the benefits of an advance may be highly nonlinear. Value is measured by the technologies' performance in the market, and by the realization of an uncertain damage function associated with an environmental externality.

The economic system into which the developing technologies are diffused evolves in parallel with the research programs. Given the state of the economic system at the beginning of a period, defined by the market shares of the technologies associated with the research programs, the state of the system at the beginning of the next period is modeled as the result of private cost minimizing decisions by the relevant industries. These decisions describe diffusion as a deterministic function of the state of technology in the current period (but *not* of the outcomes of the concurrent research programs), a set of exogenous parameters such as demand conditions and environmental damages, and the existence and extent of an internalization policy.[4] The policy

[2] The term 'multi-armed bandit' refers to a class of problems in which the decision-maker repeatedly chooses among a set of multiple options with uncertain returns and bases future decisions on observed results.

[3] Including updating would preclude the use of the dynamic programming structure, since previous levels of investment would affect current decisions. This issue is discussed below in the context of the internalization policy in Section 3.3.

[4] The diffusion of new technology is typically viewed as a highly uncertain process subject to attributes of the economic system not modeled here, including firm size, industry structure, and institutional

may be endogenous to the R&D decision problem (i.e., influenced by the state of technology), as discussed in the Section 3.3. The evolution of the economic system can also contribute to the state of technology in the next period. As discussed in Section 2, there is strong empirical evidence that increased deployment of a technology leads to improvement via learning effects. These effects, combined with the existence of research efforts outside the scope of the public decision-maker (which may or may not be proportional to deployment), admit the possibility of advance independent of the investment allocation. In particular, the probability of independent advance in a given technology is at least partially increasing in its market share. Thus the probability distribution across the outcome space is conditioned on both the innovation production function and the diffusion pattern in the economic system.

Because the R&D decision-maker is concerned with maximizing social benefit, the problem's objective function includes measures of the utility of consumption, economic abatement costs, and environmental damages associated with the state of the economic system and the exogenous parameters. Consumption utility is defined as the aggregate utility in the system associated with private consumption *before* any environmental mitigation is enforced. This baseline is included to capture the value of technological development independent of its impact on abatement costs. These costs are measured as the reduction in aggregate utility of consumption as the result of an optimal internalization policy (discussed below), but are calculated net of any emissions penalties imposed by the policy.[5] Finally, unless the policy-maker finds it optimal to abate fully, there will remain a certain level of emissions associated with the policy-controlled economic system, whose environmental costs depend on the realization of an uncertain damages function. From the perspective of the R&D decision-maker, value is calculated as expected utility net of the sum of economic and environmental costs, taking into account the policy-maker's choice. [6] The decision-maker's problem is therefore to allocate investment across the programs in each period, and at each node in the outcome space, so as to maximize the expected value of the system's evolution over the time horizon. The solution is found with a dynamic programming recursion, and thus takes the form of a dynamic, adaptive strategy that specifies optimal investment at the present moment, and in every possible future scenario.

## 3.2 Specification

The notation presented in this section is summarized in a table in the Appendix. The state of research program $i$ at the beginning of period $t$ is represented by $x_i^t$ (for example, $x_i^t$ might represent the cumulative number of successes in program $i$, or the value of a physical parameter such as efficiency or per unit capital cost). The state of technology at time $t$ is given by the vector $x^t = \{x_1^t, \ldots, x_n^t\}$. Because the state of technology determines the diffusion pattern, and because successive diffusion patterns depend on each other, the technology state *path*, in addition to the instantaneous state, is also a relevant variable. Let $\Theta_T$ represent the set of all possible paths for the state of technology over the time horizon $T$, with $\theta_T \in \Theta_T$ representing an instance of the vector $\{x^1, \ldots, x^T\}$. Since it will be necessary to evaluate intermediate positions on the technology path, define a partial path $\theta_t = \{x^1, \ldots, x^t\}$ for any $t < T$. The notation $\Theta_t$ is used to represent all possible partial technology paths up to time period $t$, and $\Theta_t(\theta_{t_0})$, with $t_0 < t$, represents the subset of $\Theta_t$ that includes the partial path $\theta_{t_0}$. In other words, $\Theta_t(\theta_{t_0})$ is the set of states in $\Theta_t$ that can be reached from the state $\theta_{t_0}$.

Diffusion is modeled as a deterministic function of the state of last period's economic system (which determines the initial state), the current state of technology, the current state of an exogenous demand parameter, $w^t$, and the current internalization policy, $p^t$. The parameter $w^t$ encapsulates socioeconomic trends such as population growth, economic development, income convergence, and other non-technological factors. Its future evolution is uncertain, but

barriers. The representation here was chosen to focus the analysis on technological, environmental, and policy uncertainties.

[5]In some cases, such as grandfathered permits, these penalties are in fact costless. In other cases, such as auctioned permits or taxes, the penalties are simply a transfer and do not affect the net social cost of the policy. Moreover, some economists have suggested that emissions taxes can actually reduce social costs by offsetting revenues from other, more distortionary taxes. This so-called 'double dividend' effect is not considered here.

[6]The value associated with a *particular* R&D investment is calculated as the difference between expected value with and without the investment.

in each period its current state is revealed, and beliefs about the probabilities of future states are updated. The policy variable $p^t$ represents the price or penalty levied on each unit of the externality-associated emission. We will discuss below whether this variable is exogenous or endogenous, and how it is determined. The state of the economic system in period $t$ is represented by $y^t$, and its evolution is defined by the diffusion mapping $m$:

$$y^t = m(x^t, w^t, p^t; y^{t-1}) \tag{1}$$

The allocation of investment to program $i$ in period $t$ is given by $\alpha_i^t$, and $\alpha^t = \{\alpha_1^t, \ldots, \alpha_n^t\}$ represents the overall investment allocation in period $t$. Total R&D investment in period $t$ is denoted $A^t = \sum \alpha_i^t$; the budget for period $t$ is denoted $B^t$. The vector $\alpha^t$ partially determines the probability distribution for $\theta_{t+1}$ over $\Theta_{t+1}(\theta_t)$ according to the innovation production function. However, the actual density function over the technology outcome space also depends on the possibility of exogenous advance, which may depend in turn on the diffusion pattern of previous periods. Let the function $\ell(y^t)$ represent the contribution of present deployment of a technology to its future progress. Accordingly we have the technology outcome density function $f$, which includes both sources of potential advance:

$$f(\theta; \alpha^t, \ell(y^t), \theta_t) \equiv \Pr(\theta_{t+1} = \theta), \forall \theta \in \Theta_{t+1}(\theta_t) \tag{2}$$

The baseline utility of consumption and economic abatement costs incurred in each period are deterministic functions of the diffusion, denoted by $U^t(y^t)$ and $C^t(y^t)$, respectively. However, the environmental costs are subject to uncertainty. In the case of climate change, the damage function is linked to the stock of greenhouse gases, not to their emission, so that environmental costs generated by emissions are neither fully incurred, nor their extent fully known, in the current period. Let $e^t(y^t)$ represent the emissions in period $t$, and let $D(e^1, \ldots, e^T)$ denote the sum of discounted damages associated with the emissions path, which is not known with certainty until the final period. Let $z^t$ represent the exogenous state of knowledge about the externality's damage function, that is, the distribution at time $t$ over possible underlying states of $D$. With time this distribution is assumed to decrease in spread, though it may vary in mean, as more information is obtained.

As with the state of technology, the path of each state variable is influential via its intertemporal effects on diffusion. Therefore we define path variables $\omega^t = \{w^1, \ldots, w^t\}$ and $\psi^t = \{z^1, \ldots, z^t\}$. As before, $\Omega_t$ and $\Psi_t$ are respectively the set of all possible paths to time $t$, and $\Omega_t(\omega_s)$ and $\Psi_t(\psi_s)$ are the respective reachable sets of paths to time $t$ from a particular path to time $s$.

### 3.3 Interaction with Policy

Since this study focuses on the perspective of the public R&D decision-maker, $\alpha$ represents the primary decision variable. However, an interesting property of this problem is that the optimal selection of $\alpha$ depends on the selection of the policy variable $p$. Assuming there is a second policy-maker player that is optimally choosing $p$, this player will consider expectations about future abatement opportunities when determining the optimal current level of abatement. These expectations should be predicated on the concurrent investment decisions of the technology strategist. However, the technology strategist must consider expectations about the future state of the economic system when determining the optimal allocation of current research investment, and because of the dynamic nature of technology diffusion, this state depends on the concurrent actions of the policy-maker.[7] Therefore in theory, the policy variable and the technology strategy must be solved for simultaneously. Since both players have the same objective function, the maximization of social welfare, the equilibrium solution will coincide with the social optimum.

[7]Off-equilibrium deviation by one player from her optimal strategy alters the optimal strategy of the other player. For example, if a (real or perceived) budget constraint prevents optimal allocation of the research budget in early periods, a different policy would be optimal, which in turn would affect the optimal research strategy. Similarly, if there were a (real or perceived) upper limit on the economic costs introduced by a policy, causing the expectation of a sub-optimal level of internalization at some points in the outcome space, a new equilibrium would emerge. These issues are left for future work.

Although there exists a simultaneous equilibrium $(\alpha^*, p^*)$, which represents an optimal strategy set in the current state and in each future state of the world, it is difficult to compute. First consider the policy variable alone. Some integrated assessment models are designed to calculate the optimal policy trajectory over time (for example, the DICE model in Nordhaus 1994), but these calculations assume fixed, exogenous paths for technology and all other parameters. In the dynamic setting proposed here, we impose a diverging state space and seek the policy that maximizes expected future value at each node. Such a setting is only useful, however, when the optimization at a given node depends only the current choice and future choices (which have been determined via the recursion solution technique). In the case of an internalization policy, the optimal choice depends also on the policy in previous periods, so that the separability required for a dynamic solution is not achieved. The optimal choice in the final period would be a function of the optimal choices in all preceding periods, so that the entire state space would have to be solved simulataneously. Now consider the technology strategy variable alone. A dynamic setting is appropriate in this case, because the technology state space can be discretized. Considering each partial path of technology as a different, the optimal investment allocation in a given state can be determined independent of previous investment allocations, since only the technology path matters. This 'trick' cannot be used in the case of the policy variable unless either the state of technology diffusion or the policy variable itself is discretized, but neither lends itself to such a representation. So, although a dynamic approach to the determination of an optimal policy is intuitively attractive, it cannot be suitably modeled in this way. Combining both into a simultaneous calculation only exacerbates the difficulty, because the optimization of both variables would depend on previous choices.

For the purposes of the current study, we wish to preserve the dynamic programming structure and emphasize the technology strategy problem. However, the policy variable is instrumental in determining the value of technology development. Therefore we include it in the model, but with the assumption that it is chosen *myopically* in each period as a function only of the current state variables. In particular, we must assume that the policy optimization does not take the implications of the current research portfolio about the future state of technology into account. Two cases could be considered. In the first case, the policy variable is determined by the parameters $\psi_t$, the state path of knowledge about the damage function, and $\omega_t$, the development path. Thus it is independent of the state of technology, and therefore the current state of the economic system. This is a particularly naive assumption, since the optimal policy should depend on both the marginal abatement cost and the marginal damage curves, but it is made implicitly by any technology study with an exogenous policy variable, and will serve as a useful reference point. In the second case, the optimal policy is calculated as a function of $\theta_t$, the state path for technology, as well as the two other state variables. It is again independent of expectations about the future based on $\alpha^t$, but because the state of economic system is fully determined by the three state variables, the policy implicitly takes this into account as well. A variation of this type of policy regime that is often used in integrated assessment models is the enforcement of a certain level of emissions in each period that corresponds to a stabilization target, such as the emission paths described in Wigley et al. (1996). Such an approach simply minimizes the cost of achieving the current period's allotted emission reductions and does not attempt an intertemporal optimization. In this case the parameter $z^t$ would effectively represent the desired target.

## 3.4 Dynamic Programming Recursion

The problem is solved using a dynamic programming recursion over the finite time horizon. In the final period, each possible state is evaluated. A state is represented by $s_T = (\theta_T, \omega_T, \psi_T)$, the three complete path variables, and the state space in the final period is denoted $\mathcal{S}_T = \Theta_T \times \Omega_T \times \Psi_T$. Thus for a given element $s_T \in \mathcal{S}_T$, the evolution of the economic system $\{y^t\}$, including emissions and the internalization policy, may be determined in each period of the path, and an absolute value $V^T$ assigned:[8]

$$V^T(s_T) = \sum_{t=1}^{T} (1-\delta)^t \left(U^t(y^t) - C^t(y^t)\right) - D(e^1, \ldots, e^T) \tag{3}$$

[8] $\delta$ indicates the per period discount rate.

From the absolute value assigned to the entire path, the dynamic programming recursion can be established. Let $V^t(s_t)$ equal the value of any partial path through the state space to time $t$. This value is determined by optimizing the choice of the R&D investment allocation vector. At any node in the state space, the allocation vector is chosen to maximize the expected value of the next period's state, $V^{t+1}$, less the cost of the investment. Since $V$ is based on a discounted sum of utility and costs across all periods, the research investment must also be discounted back to the present. Therefore all values of $V$ are in terms of present value throughout the dynamic program.

$$V^t(s_t) = \max_{\alpha^t} \left( \sum_{\theta \in \Theta_{t+1}(\theta_t)} f\left(\theta; \alpha^t, \ell(y^t), \theta_t\right) \mathrm{E}_{\omega_{t+1}|\omega_t} \mathrm{E}_{\psi_{t+1}|\psi_t} V^{t+1}(s_{t+1}) - (1-\delta)^t A^t \right) \quad (4)$$

The solution will take the form of a policy $\alpha^*$, consisting of an optimal allocation vector $\alpha^{t*}(s_t)$ for every $s_t \in \mathcal{S}_t$ and for $t = 1, \ldots, T-1$. Because of the one-period lag between R&D investment and outcomes, there is no investment in the final period. In the first period, there is a single optimal investment allocation, since $|\mathcal{S}_1| = 1$, the current state.

In the above formulation, a budget constraint is not applied. Using the assumption that the innovation production function is concave for all research programs, a finite solution to the maximization will exist without a budget. It may be interesting to investigate the trajectory of optimal total investment, along with its allocation, through time and at the various decision nodes. Once the unconstrained solution is obtained, a budget constraint can be applied (i.e. a positive shadow price imposed on investment levels at optimality) and the solution modified by reducing investment in every program so as to equalize marginal benefit at the implicit shadow price.

### 3.5 First Order Conditions

The optimization problem at an arbitrary stage in the dynamic program may be characterized as follows, with time subscripts removed:

Maximize

$$V_0 = \sum_{\theta \in \Theta(\theta_0)} f(\theta; \alpha, \ell(y^0), \theta_0) \mathrm{E}_{\omega,\psi} V(\theta, \omega, \psi) - (1-\delta) \sum_i \alpha_i \quad (5)$$

Therefore in the absence of a budget constraint, the first order condition for $\alpha_i$ is:

$$\frac{\partial V_0}{\partial \alpha_i} = \sum_{\theta \in \Theta(\theta_0)} \frac{\partial f}{\partial \alpha_i}(\theta; \alpha, \ell(y^0), \theta_0) \mathrm{E}_{\omega,\psi} V(\theta, \omega, \psi)) = 1 - \delta \quad (6)$$

In this case, investment in each program is increased until the marginal benefit equals the discounted opportunity cost. Although the criterion for the marginal condition is independent of investment in other programs, the benefit function depends on these variables, so the $n$-equation system must be solved simultaneously. The balance between research programs will depend on the response in $f$ to changes in investment, the influence of the various programs on the value function, and interdependence among the programs. When a budget constraint is applied, the first order conditions are adjusted slightly to include the budget's shadow price, $\lambda$, to which the marginal benefit in each program must be equalized.

$$\frac{\partial V_0}{\partial \alpha_i} = \frac{\partial V_0}{\partial \alpha_j} = \lambda + 1 - \delta, \forall i, j, \quad (7)$$

$$\sum_i \alpha_i \leq B \quad (8)$$

In optimality, the usual complementary slackness condition holds. If the budget constraint is binding, $\lambda > 0$; when the constraint is slack, the optimal level of investment is less than $B$, and $\lambda = 0$.

## 3.6 Example

To better understand the optimality conditions, a simple example is analyzed. Let $n$ and $T$ both equal 2, so that there is only one decision, the allocation of the budget between two research programs at the beginning of the time horizon. Suppose each program may succeed or fail (denoted by $x_i = 1$ or 0. respectively), and the probability of success in program $i$ is determined by a function $f_i(\alpha_i)$. This function may include learning effects, since any such effects would be fully determined by the economic system's evolution in the first period, which we assume depends only on initial conditions. Similarly, suppose the exogenous parameters (considered in aggregate for simplicity) may assume either a high level with probability $\mu$ or a low level with probability $1 - \mu$. Thus there are eight possible scenarios for the evolution of the economic system in the second period, or eight instances of $s_2$. The absolute value function $V^2$ is defined as follows:

| $\theta_2$ | $\omega, \psi$ | |
|---|---|---|
| | High | Low |
| $\{(0,0),(0,0)\}$ | $\mathcal{U}^h$ | $\mathcal{U}^l$ |
| $\{(0,0),(1,0)\}$ | $\mathcal{U}^h + \mathcal{V}_1^h$ | $\mathcal{U}^l + \mathcal{V}_1^l$ |
| $\{(0,0),(0,1)\}$ | $\mathcal{U}^h + \mathcal{V}_2^h$ | $\mathcal{U}^l + \mathcal{V}_2^l$ |
| $\{(0,0),(1,1)\}$ | $\mathcal{U}^h + k^h(\mathcal{V}_1^h + \mathcal{V}_2^h)$ | $\mathcal{U}^l + k^l(\mathcal{V}_1^l + \mathcal{V}_2^l)$ |

Here $\mathcal{U}^h$ and $\mathcal{U}^l$ represent the net welfare over the entire time horizon in the baseline case with no technological advance when the level of the exogenous parameter is high and low, respectively. $\mathcal{V}_i^h$ and $\mathcal{V}_i^l$ represent the *increase* in net welfare in the case of an advance in *only* technology $i$ in the two exogenous states.[9] Finally, the scalars $k^l$ and $k^h$ represent the nature of the additivity of welfare gain when *both* programs succeed. When the $k > 1$, the technologies associated with the programs are complements; when $k < 1$, subsitutes.[10]

Subtracting the term $\mathrm{E}_{\omega.\psi}\mathcal{U} = \mu\mathcal{U}^h + (1-\mu)\mathcal{U}^l$ from each to normalize the value of the no advance case to zero, we evaluate Equation 5 to obtain an objective function as follows:

$$\begin{aligned} V \quad = \quad & f_1(\alpha_1) * (1 - f_2(\alpha_2)) * (\mu\mathcal{V}_1^h + (1-\mu)\mathcal{V}_1^l) + \\ & (1 - f_1(\alpha_1)) * f_2(\alpha_2) * (\mu\mathcal{V}_2^h + (1-\mu)\mathcal{V}_2^l) + \\ & f_1(\alpha_1) * f_2(\alpha_2) * (\mu k^h(\mathcal{V}_1^h + \mathcal{V}_2^h) + (1-\mu)k^l(\mathcal{V}_1^l + \mathcal{V}_2^l)) \end{aligned} \quad (9)$$

To solve the problem in this example, $V$ is maximized subject to $\alpha_1 + \alpha_2 = B$. Note that no discount factor is required since all the value nodes occur in the same period. The objective may be reduced to an equivalent form:

$$V = f_1(\alpha_1)[\mathrm{E}_{\omega,\psi}\mathcal{V}_1 + f_2(\alpha_2)K] + f_2(\alpha_2)\mathrm{E}_{\omega,\psi}\mathcal{V}_2 \quad (10)$$

Where

$$K = \mu k^h(\mathcal{V}_1^h + \mathcal{V}_2^h) + (1-\mu)k^l(\mathcal{V}_1^l + \mathcal{V}_2^l) - \mathrm{E}_{\omega,\psi}(\mathcal{V}_1 + \mathcal{V}_2) \quad (11)$$

$K$ quantifies the impact of the technologies' success on each other. If the technologies are complements, $K$ is positive, and if they are substitutes, $K$ is negative.[11] When the technologies are independent, $K$ vanishes, and the objective function is separable in the decision variables.

[9] Technological advance is assumed to have a strictly positive welfare impact.

[10] While $k^l$ and $k^h$ are assumed to be either both greater than 1 or both less than 1, they need not, and likely will not in most cases, be identical.

[11] Enforcing a lower bound on $k^l$ and $k^h$ such that the value when both technologies succeed is at least as great as the value in either individual success scenario limits the substitution effect so that $K > -E\mathcal{V}_1$ and $K > -E\mathcal{V}_2$.

From the form in Equation 10, the first order condition in Equation 7 (accompanied by the budget constraint) is readily derived:

$$f_1'(\alpha_1)[\mathrm{E}_{\omega,\psi}\mathcal{V}_1 + f_2(\alpha_2)K] = f_2'(\alpha_2)[\mathrm{E}_{\omega,\psi}\mathcal{V}_2 + f_1(\alpha_1)K] \tag{12}$$

This differential equation describes the optimal choice of $\alpha_1$ and $\alpha_2$. For further insight into its meaning, let us suppose that the innovation production function $f$ takes the following form:

$$f_i(\alpha_i) = \rho_i(1 - e^{\frac{-\alpha_i}{\beta}}), \text{ for some } \rho_i < 1 \tag{13}$$

This form satisfies the decreasing returns to scale assumption, and is a suitable probability mapping, since $f_i(0) = 0$, and the limit as $\alpha_i$ approaches infinity is $\rho_i < 1$. This parameter represents the best possible success rate achievable in the program, and $\beta$ is a scaling parameter common to both programs. Note that this functional form assumes no exogenous probability of advance from learning effects. Substituting Equation 13 into Equation 12, the first order condition is reduced to:

$$\alpha_2 - \alpha_1 = \beta \ln\left(\frac{\rho_2\mathrm{E}\mathcal{V}_2 + \rho_1\rho_2 K}{\rho_1\mathrm{E}\mathcal{V}_1 + \rho_1\rho_2 K}\right) \tag{14}$$

Thus when the budget constraint is binding, the optimal solution is defined by a simple linear system. Moreover, the research programs may be rank-ordered according to the value $\rho_i\mathrm{E}\mathcal{V}_i$, and the first-ranked program receives a stricly larger share of the budget. In the case of a tie, the difference evaluates to zero, and an equal allocation is optimal. If the budget is less than the constant on the RHS of Equation 14, the first-ranked program receives the entire allocation. When the budget is greater than this amount, all investment up to the level of the constant is allocated to the first-ranked program, and the remainder is split evenly. When the budget constraint is not binding, the optimal level of investment is determined by setting the marginal benefit expression (either side of Equation 14) equal to $1 - \delta$.

The threshold constant, or difference between the investment in the two programs, also has an intuitive interpretation. When the expected values of the technology developments on their own in the best case (i.e. with probability $\rho_i$) are equal, the optimal values of $\alpha_1$ and $\alpha_2$ are also equal, regardless of $K$, the level of interaction. When one expected best case value is larger, that program receives more funding, but the difference in funding depends not only on the difference in expected value, but also on $K$. All else equal, the difference is decreasing in $K$ when the technologies are complements and increasing in $-K$ when they are substitutes. These comparative statics suggest that the stronger is a complementarity effect, the more evenly investment should be allocated, but the stronger is a substitution effect, the technology that appears more favorable on its own receives an increasingly larger share of the investment.

When more than two programs are considered, the expression for marginal benefit (which is to be equalized across programs) will include a $K$ term for each combination of successful programs. However, the insights from the two-program case, in terms of a rank-ordering and successive inclusion in the portfolio up to the budget constraint, will continue to apply. This example therefore primarily illustrates the effects on portfolio diversification of the decreasing returns to scale assumption. Also, the influence of heterogeneous applications is also demonstrated, to the extent that programs aimed at different applications will have values of $K$ close to zero, avoiding the concentration caused by strong subsitution effects. However, the effects of risk and uncertainty are not apparent here, because these require multiple periods to manifest. When more than two periods are considered, the model in the example is applied in the penultimate period, and its optimal solutions become the value nodes for the preceding period. The results suggested by the optimality confition in Equation 14 will be complicated in this case by the ability to 'explore' in early periods. Finally, when research programs may result in more than two outcomes, the interaction space is expanded as when more programs are added. This extension may be handled similarly by including additional $K$ terms, but in none of these cases is the model fundamentally altered.

# 4 Implementation

## 4.1 Technologies

As a first step in implementing the methodology described here, we focus on the electric power generation sector. Carbon intensity reductions can be achieved in this sector by three broad classes of technological developments. First, renewable generation technologies can be improved. This category refers to generation technologies with zero net carbon emissions, and for the purposes of this study includes wind, photovoltaic cells, solar thermal, biomass, and fuel cells using a non-fossil fuel source, such as biologically produced hydrogen.[12] Developments in this group of technologies might take many different forms, but the economic impact in most cases will be a reduction in capital cost per unit output. Such a shift could lead to expanded deployment of renewable technologies even without a policy incentive, but their competitive advantage would be much greater in the presence of a carbon constraint.

A second broad technological class is carbon capture and sequestration. This category is considered separately from other renewable technologies because of its interaction with conventional fossil generation. Developments in this category can decrease the cost of carbon separation and capture associated with both existing and future fossil generation technologies, as well as establish the viability of geologic storage reservoirs. However, these technologies will never be profitably implemented without a carbon constraint, so their economic impact is highly sensitive to policy uncertainty. On the other hand, in the case of a policy incentive sufficiently large to induce adoption of sequestration services, the availability of this technology could significantly erode the marketability of competing renewable technologies.

The final technological area under consideration is efficiency improvements in fossil fuel generation. Such improvements could take the form of more efficient combustion processes in conventional configurations, or the introduction of new configurations, such as a fuel cell with hydrogen produced from a fossil source. In either case, technological developments increase the amount of output available from a fuel input, possibly at the expense of other inputs. As long as this combination results in a reduced levelized cost, the new developments will be implemented widely regardless of policy conditions. Because the developments would also encompass reduced emissions per unit output, their diffusion and value will be increasing in the carbon constraint; however, there exists a finite limit to the amount of abatement they can achieve on their own.

Because the value of a technology portfolio is linked to the diffusion of developments, the pre-existing characteristics of the market landscape must be considered. Many of these are region-specific, so it is important to differentiate geographically as well as technologically. To fully assess the R&D strategy problem, all sectors, technologies and regions should be considered simultaneously to identify the greatest potential for effective cost reduction. In the ase study that accompanies this paper, we demonstrate the methodology by examining a small but important subset of the entire space. The three technology classes of electric power generation are considered in the context of two distinct global regions: the United States and China. These two regions are both large greenhouse gas emitters, but they have several important differences. First, while the US is single largest emitter today, China's expected growth rate is considerably higher, and it will likely overtake the emissions lead in one or two decades. Second, China's existing electric generation infrastructure is predominantly coal-based, and it has an enormous resource base of this fuel, which ranks first in carbon intensity. The United States is also widely invested in coal generation, but its plants are more modern and subject to stricter environmental regulations. Of particular importance is the incompatibility of most of China's coal plants with potential carbon capture systems. Third, the cost of capital is higher in China than in the United States, which affects the financing and diffusion of new technologies. These factors suggest that the most effective suite of mitigation technologies may be different for the two regions, and this will influence the allocation of R&D investment.

[12]Other zero carbon generation technologies have been omitted. The extent of hydroelectric and geothermal generation is assumed to be constrained by physical resource availability and therefore relatively unresponsive to technological improvements. Similarly, the use of nuclear generation is more sensitive to political constraints than to the state of technology.

### 4.2 Integrated Assessment

The model described in Section 3 outlines the decision-making process assuming that the absolute values of each path are known. This process is an innovative adaptation of the usual mode of integrated assessment, but its implementation as a policy tool depends nonetheless on an ancillary model to provide the necessary value inputs. Many such models exist and have been widely used in the analysis of climate policy options. However, the application of a typical integrated assessment model (IAM) in this context is problematic. First, the dimensionality of the dynamic programming approach requires a very large number of model runs, depending on the granularity of the various probability distributions. Most IAMs are top-down in nature, but because they are designed for moderately fast computation of only a single scenario, they are able to include a considerable amount of detail, which could make repeated iterations cumbersome. A second consideration is the representation of both the potential states of technology and their diffusion. Most models have sufficient technological detail to allow for the encoding of progressively advanced states at different times, but there are different methods for determining the technology mix. A key feature of diffusion is its dependence on the inertia of the current system. Because diffusion is modeled here as the outcome of cost-minimizing decisions in each time period based on the existing infrastructure, available new technology, and current policy conditions, capital vintage accounting will be crucial to the diffusion pattern. An example of a modeling structure that coincides well with the diffusion process proposed here is the EPPA model, described in McFarland et al. (2004). A model with similar treatment of technological diffusion, but pared down to facilitate rapid computation, is under development.

## 5 Conclusion

This paper highlights the need to study the basic public R&D investment allocation problem. Despite the sensitivity of climate policy approaches to the future state of technology, there has been little attention focused on the role of basic public research in determining that state. The problem is defined here as distinct from other R&D problems by virtue of ongoing, interrelated projects, a flexible innovation production function, and an interdependence with an environmental externality. We present a model for systematically evaluating the effects of a technology strategy in this context, and demonstrate the dynamic optimization of this strategy. The analytic illustration of the decision problem shows that the potential for deployment, timing of advance and uncertainty resolution, and the innovation production relationship all contribute to the determination of the optimal allocation. The paper identifies and characterizes the dependencies and provides intuitive and straightforward optimality conditions. In particular, the interaction of the technology strategy with the optimal internalization policy is explored in depth. Finally, a simple application is used to develop insights into the effects of decreasing returns to scale and technology correlation on a single-period diversification decision. We also describe the motivations for diversification based on risk management.

The decision model presented here demonstrates how existing modeling frameworks can be used to analyze the R&D problem. Future work will include the incorporation of IAM-based estimates of the different state paths so that the methodology can be applied to realistic policy analysis. Based on these inputs, recommendations can be offered for specific technologies based on their performance, interaction with other technologies, and relationship to the current state of the energy system.

## Acknowledgements

The funding for this research has been provided by the Global Climate and Energy Project at Stanford University. Please visit gcep.stanford.edu for more information about our research and sponsors. The authors would also like to thank Leon Clarke, Jim Sweeney, Chris Edwards, and Erin Baker for helpful discussions during the course of this work.

## Appendix: Notation

**Variables**

| | | | |
|---|---|---|---|
| $T$ | | | number of time periods |
| $n$ | | | number of research programs |
| $\alpha^t$ | $=$ | $\alpha_i^t, \ldots, \alpha_n^t$ | allocation of research investment in period $t$ |
| $A^t$ | $=$ | $\sum_i \alpha_i^t$ | total research investment in period $t$ |
| $B^t$ | | | budget constraint for period $t$ |
| $x^t$ | $=$ | $x_i^t, \ldots, x_n^t$ | state of technology in period $t$ |
| $w^t$ | | | state of socioeconomic development in period $t$ |
| $z^t$ | | | state of knowledge about environmental damages at period $t$ |
| $\theta_t$ | $=$ | $x^1, \ldots, x^t$ | time path of state of technology to period $t$ |
| $\omega_t$ | $=$ | $w^1, \ldots, w^t$ | time path of state of development to period $t$ |
| $\psi_t$ | $=$ | $z^1, \ldots, z^t$ | time path of state of damages knowledge to period $t$ |
| $s_t$ | $=$ | $(\theta_t, \omega_t, \psi_t)$ | dynamic program state of all three path variables |
| $y^t$ | | | state of economic system in period $t$ |
| $p^t$ | | | internalization policy in period $t$ |
| $\delta$ | | | per period discount rate |

**Sets**

| | | | |
|---|---|---|---|
| $\Theta_t$ | | | set of all possible technology paths to period $t$ |
| $\Theta_t(\theta_{t_0})$, for $t_0 < t$ | | | set of all possible technology paths to period $t$ that include $\theta_{t_0}$ |
| $\Omega_t$ | | | set of all possible development paths to period $t$ |
| $\Psi_t$ | | | set of all possible damages knowledge paths to period $t$ |
| $\mathcal{S}_t$ | $=$ | $\Theta_t \times \Omega_t \times \Psi_t$ | set of all possible dynamic program states at time $t$ |

**Functions**

| | |
|---|---|
| $m(x^t, w^t, p^t; y^{t-1})$ | state of economic system in period $t$ (diffusion function) |
| $\ell(y^t)$ | learning effects from deployment in period $t$ |
| $f(\ .\ ; \alpha^t, \ell(y^t), \theta_t)$ | probability density for state of technology in period $t+1$ |
| $e^t(y^t)$ | emissions in period $t$ |
| $U^t(y^t)$ | utility of consumption without a policy in period $t$ |
| $C^t(y^t)$ | economic abatement costs in period $t$ |
| $D(e^1, \ldots, e^T)$ | discounted environmental damages of emissions time path |
| $V^t(s_t)$ | dynamic value of optimal technology strategy in state $s_t$ |

## References


Arrow, K. (1962). Economic welfare and the allocation of resources for invention. In R. Nelson (Ed.), *The Rate and Direction of Innovative Activity*. Princeton University Press.

Baker, E., L. Clarke, and J. Weyant (2003). R&D as a hedge against climate damages. *Submitted to Energy Journal*.

Dasgupta, P. and E. Maskin (1987). The simple economics of research portfolios. *The Economic Journal 97*, 581–595.

Dixit, A. and R. Pindyck (1994). *Investment Under Uncertainty*. Princeton University Press, Princeton, NJ.

Goulder, L. and K. Mathai (2000). Optimal $CO_2$ abatement in the presence of induced technological change. *Journal of Environmental Economics and Management 39*(1), 1–38.

Goulder, L. and S. Schneider (1999). Induced technological change and the attractiveness of $CO_2$ abatement policies. *Resource and Energy Economics 21*, 211–253.

Gritsevskii, A. and N. Nakicenovic (2000). Modeling uncertainty of induced technological change. *Energy Policy 28*.

Grubler, A. and A. Gritsevskii (2001). A model of endogenous technological change through uncertain returns to learning (R&D and investments). *Ukranian Economic Review*.

Hoffert, M., K. Caldeira, G. Benford, D. Criswell, C. Green, H. Herzog, A. Jain, H. Kheshgi, K. Lackner, J. Lewis, H. Lightfoot, W. Mannheimer, J. Mankins, M. Manuel, L. Perkins, M. Schlesinger, T. Volk, and T. Wigley (2002). Advanced technology paths to global climate stability: Energy for a greenhouse planet. *Science 298*, 981–987.

Hsu, J. and E. Schwartz (2003). A model of r&d valuation and the design of research incentives. National Bureau of Economic Research, Working Paper 10041.

Loch, C. and S. Kavadias (2003). Dynamic portfolio selection of npd programs using marginal returns. *Management Science 48*(10), 1227–1241.

McFarland, J., J. Reilly, and H. Herzog (2004). Representing energy technologies in top-down economic models using bottom-up information. *Energy Economics 26*(4), 685–707.

Miketa, A. and L. Schrattenholzer (2004). Experiments with a methodology to model the role of r&d expenditures in energy technology learning processes; first results. *Energy Policy 32*, 1679–1692.

Nordhaus, W. (1994). *Managing the Global Commons: The Economics of Climate Change*. Cambridge, MA: MIT Press.

Nordhaus, W. (2002). Modeling induced innovation in climate-change policy. In A. Grubler, N. Nakicenovic, and W. Nordhaus (Eds.), *Technological Change and the Environment*. Washington, D.C.: Resources for the Future.

Popp, D. (2002). ENTICE: Endognenous technological change in the DICE model of global warming. Syracuse University, Submitted to *Journal of Environmental Economics and Management*.

Sagar, A. and J. Holdren (2002). Assessing the global energy innovation system: Some key issues. *Energy Policy 30*, 465–469.

Schock, R., W. Fulkerson, M. Brown, R. S. Martin, D. Greene, and J. Edmonds (1999). How much is energy research and development worth as insurance? *Annual Review of Energy and the Environment 24*, 487–512.

Seebregts, A., T. Kram, G. Schaeffer, A. Stoffer, S. Kypreos, L. Barreto, S. Messner, and L. Schrattenholzer (1999). Endogenous technological change in energy systems models: Synthesis of experience with eris, markal, and message. Technical report, Netherlands Research Foundation ECN.

van der Zwaan, B., R. Gerlagh, G. Klaassen, and L. Schrattenholzer (2002). Endogenous technological change in climate change modeling. *Energy Economics 24*, 1–19.

Weitzman, M. (1979). Optimal search for the best alternative. *Econometrica 47*(3), 641–654.

Wigley, T., R. Richels, and J. Edmonds (1996). Economic and environmental choices in the stabilization of atmospheric $CO_2$ concentrations. *Nature 379*(6562), 240–243.

Zhu, K. and J. Weyant (2003). Strategic decisions of new technology adoption under assymetric information. *Decision Sciences 34*(4), 643–675.