

March 2007 ■ RFF DP 07-11

Endogenizing Technological Change

*Matching Empirical Evidence to
Modeling Needs*

William A. Pizer and David Popp

1616 P St. NW
Washington, DC 20036
202-328-5000 www.rff.org

Endogenizing Technological Change: Matching Empirical Evidence to Modeling Needs

William A. Pizer and David Popp

Abstract

Technologies to reduce significantly fossil-fuel emissions currently are unavailable or only available at high cost. In light of this, the amount of research on the pace, direction, and benefits of environmentally friendly technological change has grown dramatically in recent years. This research includes empirical estimates of these effects and modeling exercises designed to simulate endogenous technological change in response to climate policy. Unfortunately, few attempts have been made to connect these two streams of research. This paper attempts to bridge that gap, reviewing both the empirical and modeling literature on technological change.

Our goal is to provide an agenda for how both empirical and modeling research in these areas can move forward in a complementary fashion. In doing so, we discuss how models used for policy evaluation can better capture empirical phenomena and how empirical research can better address the needs of models used for policy evaluation.

Key Words: endogenous technological change, climate change, CGE modeling

JEL Classification Numbers: Q55, Q28, O3, D58

© 2007 Resources for the Future. All rights reserved. No portion of this paper may be reproduced without permission of the authors.

Discussion papers are research materials circulated by their authors for purposes of information and discussion. They have not necessarily undergone formal peer review.

Contents

Introduction	1
The R&D Process	3
Key Features	4
What We Know about the Response of R&D to Policy	6
Modeling Approaches in Practice.....	7
Summarizing Behavior	9
Looking Forward	11
Learning by Doing	12
Key Features	13
Modeling Approaches in Practice.....	14
Summary and Looking Forward.....	15
Government R&D	16
What about Technology Diffusion?	18
The Timing of Diffusion.....	19
Diffusion across Regions	21
Conclusion	22
References	24
Figure	30

Endogenizing Technological Change: Matching Empirical Evidence to Modeling Needs

William A. Pizer and David Popp*

Introduction

The past decade has witnessed an explosion of research into the process of technological innovation. To a large degree, this research has proceeded in two distinct directions: exploring mechanisms to endogenize the process of productivity change in dynamic models used to evaluate public policy, particularly in the arena of climate change, and exploring various empirical, microeconomic questions surrounding the process of invention, innovation, and diffusion of individual technologies. There have been a number of studies reviewing these two trends separately (Gillingham et al. 2007; Löschel 2002; Popp 2005; Jaffe, Newell, and Stavins 2003). However, there has been little attempt to tie these two literatures together and, in particular, to suggest an agenda for future work. How can models used for policy evaluation better capture empirical phenomena, and how can empirical research better address the needs of models used for policy evaluation? The purpose of this paper is to examine these two questions.

In the area of climate change, it is hard to overstate the importance of estimating the expected path of technological change and the potential for public policies to influence that path. Kneese and Schultz (1978) argued that the potential to spur technological innovation may be the single most important criteria for environmental policy “over the long haul.” Their applications were typical environmental problems where attention was focused on achieving rapid improvements in environmental quality immediately. Solutions were at hand and interest in technological change focused on incremental improvements to existing technologies—improvements like higher efficiency, cleaner fuels and cheaper end-of-pipe treatment—that might arise but were not essential to achieving societal goals. The long haul attracted the attention of Kneese and Schultz but was not the focus of stakeholders and policymakers. In sharp contrast, climate change focuses squarely on the long haul.

* Authors are Senior Fellow, Resources for the Future, and Associate Professor, Syracuse University, respectively. The authors gratefully thank participants at the 2006 Workshop on Technological Change and the Environment at Dartmouth College, and Karen Fischer-Vanden in particular, for providing the genesis for this paper.

Addressing climate change means, for many people, stabilizing atmospheric concentrations of carbon dioxide and other greenhouse gases at a “safe” level.¹ There is no rapid improvement to be sought right now, and stabilization at almost any plausible level will require entirely different technologies for transportation and, likely, power generation.² This requires us to make predictions about new technologies—like Integrated Gasification Combined Cycle (IGCC) with capture and sequestration and hydrogen- and electric-fueled vehicles—rather than the kinds of incremental changes typically imagined, for example, with SO₂ scrubbers.³ Indeed, recent studies of various stabilization scenarios for the United States focused much of their attention on technological development and assumptions (Clarke et al. 2006)

This inherent link between technological development and successfully addressing climate change means that modeling technological change is no longer an academic exercise. Stern (2006) recently estimated that greenhouse gases could be stabilized at 550 ppmv CO₂e, or about 450 ppmv CO₂, for an estimated 1 percent of global world product.⁴ He included a notional confidence interval of as much as a 4 percent loss but also a 2 percent gain. What strikes some economists as an optimistic central estimate, and a very optimistic upper-end of the range, derives in part from assumptions about technological change and the potential for public policy to affect it (Byatt et al. 2006). Yet, the absence of empirical work coupled to modeling applications makes it difficult convincingly to say much about those results. Meanwhile, the report continues to influence public opinion (Stern 2007).

Three important concerns underlie the frustration and confusion surrounding the use of these and similar estimates in the area of public policy. One is whether the models cited in support of the estimates—models that include assumptions about technological developments responding to public policy—were intended to provide quantitative results or qualitative insights. The former, normative modeling exercises are designed for cost–benefit analysis and are the

¹ Economists advocating cost–benefit analysis (Nordhaus 1994; Pizer 2002; Tol 2006) would suggest pricing emissions to reflect marginal damages, which might or might not stabilize concentrations.

² While often ignored, geo-engineering to address climate with, for example, the emission of reflective particles in the atmosphere, would allow us to continue to use conventional technologies and atmospheric concentrations of carbon dioxide to continue to rise. See IPCC (2007).

³ In some ways, these are not entirely “new”—IGCC plants exist, CO₂ has been captured and sequestered for enhanced oil recovery for some time, and electric cars exist. However, these technologies (or proposed alternatives) remain unproven at the necessary scale and uncompetitive with cheap, conventional alternatives.

⁴ Note that concentrations already have risen from 315 ppmv CO₂ to 380 from 1956 to 2006—suggesting 450 will be breeched without significant action in the next 50 years.

appropriate basis for Stern-like assessments. The latter, positive exercises are designed to draw qualitative insights about behavior but may or may not be calibrated to match important quantitative benchmarks. While few models would define themselves explicitly as one or the other, it is nonetheless important to understand what kind of confidence the modelers themselves have in their results.

Second, there are many outstanding empirical questions in the literature on technological change, and modelers are not always transparent about how they have chosen to answer these questions. Even in cases where the modelers, in fact, may have a great deal of confidence in their results, users may wish to be aware of what beliefs regarding these unanswered questions are being fed into the model's assumptions.

Third, approaching this from the standpoint of theory, how much have the present suite of models sought to capture empirical features and recent research surrounding the innovation process? Diffusion of new technologies, for example, has been much studied in the empirical literature, but because many models do not identify individual, emerging technologies, their treatment of diffusion is quite crude.

To help shed some light on these issues, this paper walks through key features in both the literature on dynamic policy simulation and empirical studies of technological innovation and diffusion: the research and development process, learning by doing, the role of public versus private research, and technology diffusion. In each case, we consider what economists know about the issue, what modelers have done, what empirical questions the modelers' approaches raise, and what empirical features the modelers have ignored. We end by reviewing the agenda for empirical work and model development that we have accumulated in each section.

The R&D Process

New ideas are developed through research and development (R&D) projects. The process of R&D is costly, requiring the investment of capital by firms and government and the skills of highly trained research scientists. The economic influences on innovation can be divided into two broad categories. Economists refer to innovation responding to the needs of the marketplace as a demand-pull influence. In contrast, new research opportunities provided by scientific

advances are referred to as a technology-push influence.⁵ For example, advances in basic semiconductor materials were necessary to make solar energy feasible. Both of these affect the expected marginal rate of return (MRR) of innovation. Demand-pull influences raise the MRR by making the potential innovation more valuable. For example, innovation on pollution-control devices is only beneficial to a firm in the presence of environmental regulation, so that there will be a market for the device. Similarly, basic scientific advances increase potential returns by improving the likelihood that a research project will succeed. Finally, the innovative policy climate also is important. To earn profits from innovation, inventors must be able to appropriate the returns from their efforts. This can vary depending on factors such as the level of intellectual property rights and the type of innovation.

Key Features

When modeling innovation, the most important feature is that knowledge markets are imperfect. In most cases, new technologies must be made available to the public for the inventor to reap the rewards of invention. However, by making new inventions public, some, if not all, of the knowledge embodied in the invention becomes public knowledge. This public knowledge may lead to additional innovations or even to copies of the current innovation.⁶ These knowledge spillovers provide benefit to the public as a whole but not to the innovator.

Economists studying the returns to research consistently find that knowledge spillovers result in a wedge between private and social rates of return to R&D. As a result, private firms do not have incentives to provide the socially optimal level of research activity. Examples of such studies include Mansfield (1977, 1996), Pakes (1985), Jaffe (1986), Hall (1996), and Jones and Williams (1998). Typical results include marginal social rates of return between 30 percent and 50 percent. In comparison, estimates of private marginal rates of return on investments in physical capital range from 7 percent to 15 percent (Bazon and Smetters 1999; Jones and Williams 1998; Hall 1996). Popp (2001, 2003) finds similar results specifically looking at environmental R&D. Since firms make investment decisions based on their private returns, the wedge between private and social rates of return suggests socially beneficial research

⁵ For examples, see Scherer (1965), Schmookler (1966), Rosenberg (1974, 1982), and Mowrey and Rosenberg (1979).

⁶ Intellectual property rights, such as patents, are designed to protect inventors from such copies. However, their effectiveness varies depending on the ease in which inventors may “invent around” the patent by making minor modifications to an invention. See, for example, Levin et al. (1987).

opportunities are being ignored by firms because they are unable to fully capture the rewards of such innovations.

These spillovers have two important implications for models of environmental policy. First, it means that even if policies to correct the environmental externalities are in place, the level of environmental R&D still will be suboptimal. Because they ignore the positive spillovers created by R&D, firms will underinvest in research activity. This suggests two possible avenues through which policy can encourage the development of environmentally friendly technologies: correcting the environmental externality or correcting knowledge-market failures. Ideally, policies addressing knowledge-market failures address these issues for the economy as a whole, rather than just for environmental innovation. Examples of policies designed to encourage additional R&D include patent protection, broad government R&D subsidies, and generic R&D tax credits. Indeed, in studies on the potential effects of energy R&D for climate-change mitigation, Schneider and Goulder (1997) and Otto et al. (2006) show that policies to address knowledge spillovers are more effective if they address all knowledge spillovers, rather than focusing on R&D pertaining to alternative energy. Moreover, papers such as Popp (2006a) and Newell and Fischer (2005) find that using R&D subsidies alone, without environmental policies also in place, is not an effective climate-policy measure. While R&D subsidies encourage the creation of new technologies, they do not create incentives for the adoption of new technologies.

A second important consequence of knowledge spillovers is that they create potentially high opportunity costs for increasing environmental R&D. Evidence of high social returns to R&D is found not just in the environmental sector but throughout the entire economy. As such, any environmental R&D that comes at the expense of other R&D will have high opportunity costs. For example, consider an economy in which aggregate output, Y_t , is devoted to consumption, C_t , investment in physical capital, I_t , or two types of research: environmental, $R_{E,t}$, and other, $R_{O,t}$, so that:

$$(1) Y_t = C_t + I_t + R_{E,t} + R_{O,t}$$

When environmental R&D increases by one dollar, one less dollar is available for the other activities. The opportunity cost of the first two simply is valued at one dollar. However, if the social rate of return is four times higher for R&D than for physical capital investments, giving up a dollar of other R&D activity has the same future effect as giving up four dollars of physical investment.

This is important because empirical work suggests that at least some environmental R&D will displace other forms of R&D. Research activities are carried out by highly trained scientists

and engineers. Since years of training are needed to enter the field, the supply of scientists and engineers available at any one time is inelastic—it cannot be increased quickly when new research needs arise. For example, Goolsbee (1998) finds that one of the chief beneficiaries of R&D tax subsidies are scientists and engineers, who receive larger wages when subsidies are increased. Similarly, Popp (2004a) presents evidence that approximately one-half of the energy R&D spending that took place in the 1970s and 1980s came at the expense of other R&D.

What We Know about the Response of R&D to Policy

There is a growing body of empirical literature that links environmental policy to innovation. Most of these studies have focused on estimating the direction or magnitude of the relationship between policy and innovation. Using international patent data, Lanjouw and Mody (1996) find a correlation across nations between patenting in environmental fields and pollution-abatement expenditures. Jaffe and Palmer (1997) look at the same relationship across U.S. industries. Comparing pollution-abatement data and innovation data across these industries, they estimate a positive elasticity of pollution-control R&D with respect to pollution-control expenditures of 0.15. However, when estimating the same relationship using patents as the measure of innovation, they do not find evidence of a relationship between patenting and pollution-abatement expenditures. Finally, Brunneimer and Cohen (2003) find that patents increase by just 0.04 percent when pollution-abatement costs and expenditures increase by \$1 million. While both of these papers find small effects, they also use aggregate measures of technological change. Each looks at innovation within specific industries. Jaffe and Palmer include all patents and all R&D associated with an industry, whether or not they are environmental technologies. Thus, their results show a small increase in overall inventive activity, but do not show if additional environmental R&D takes place at the expense of other R&D. Brunneimer and Cohen focus specifically on environmental technologies but group several such technologies together for each industry. Thus, the effects of innovation on specific technologies may be masked by stagnant trends in other technologies.

In each of these studies, pollution-abatement expenditures serve as a proxy for the stringency of environmental regulation. Other papers use energy prices and related regulations as the mechanism that induces innovation. Although the observed price changes might not be policy-related, the results also can be applied to situations where policy affects prices, such as a carbon tax. Newell et al. (1999) use an approach closely related to hedonic techniques to study the effect of both energy prices and energy-efficiency regulations on technological advances in energy efficiency for air conditioners and natural gas water heaters. They find that energy prices

have the largest inducement effect, although they do not estimate a price elasticity. Popp (2002) uses patent data to study the relationship between energy prices and innovative activity. Using a distributed lag model, he estimates a long-run elasticity of energy R&D with respect to energy prices of 0.35. The median lag occurs in 4.9 years, suggesting that more than one-half of the full effect on R&D of an energy-price increase will have been experienced after just five years. Thus, prices (or other regulations that increase the cost of using fossil fuels) can be expected to stimulate new research quickly.

Modeling Approaches in Practice

That innovation responds to incentives suggests that modeling accurately induced technological change is important, as the path of climate-friendly technology will vary in scenarios with and without policy. As a result, models used to evaluate the cost and effects of climate-change mitigation policies increasingly have explored ways to endogenize the effect of R&D on technology and mitigation costs. They generally have included a variable representing R&D or “knowledge” that influenced economic behavior through one of three routes: a direct effect on the level of emissions intensity, a reduction in the mitigation cost function, or productivity gains in sectoral production functions (with sector-specific R&D stocks in multi-sectoral models). At the same time, these models have pursued a range of approaches to matching different pieces of empirical evidence described above.

For example, Nordhaus (2002) assumes that the rate of growth in emissions intensity is a function of the annual R&D directed at carbon saving:

$$(2) \quad \dot{\sigma}_t / \sigma_t = \Psi_1 R_t^{\Psi_2} - \Psi_3$$

where σ_t is the industrial carbon energy/output ratio at year t (implying $\dot{\sigma}_t / \sigma_t$ is the rate of change of the carbon energy-output ratio), R_t is the R&D inputs into the carbon-energy sector in year t , and the Ψ_i are parameters (calibrated assuming optimized R&D in the past). Importantly, he (a) assumes that the cost of a dollar of R&D is four dollars of output—reflecting the opportunity cost of displaced R&D in other sectors—and (b) calibrates the model so that in the baseline without climate policy, the return to additional R&D also is four times higher than ordinary investment. Buonanno, Carraro, and Galeotti (2003) take a similar approach in that they focus on R&D affecting emissions intensity but instead use a stock of accumulated R&D (versus the flow above) to affect the level of emissions intensity (versus the rate of change above). Their R&D stock also influences overall productivity, building in significant spillovers among different types of R&D. Given these different approaches to modeling spillovers and opportunity

costs, it may not be that surprising that Nordhaus finds relatively small effects from endogenizing R&D, while Buonanno et al. find much larger effects (see Edenhofer et al. 2006).

Popp (2004a) also uses a knowledge stock that reduces emissions intensity. But his ENTICE model allows additional substitution—both out of energy as a whole and out of carbon-emitting energy and into an alternative, carbon-free technology represented by the knowledge stock. He uses this model to focus in on the effects of crowding out on the gains from induced innovation in climate-change models. The base case of ENTICE assumes that one-half of new energy R&D crowds out other R&D. In this case, induced technological change increases welfare by 9.4 percent. Assuming no crowding out increases the welfare gains from induced technological change to as much as 45.3 percent. Similarly, assuming full crowding out of R&D reduces welfare gains to as little as 1.9 percent.⁷

The preceding approach can be viewed as an evolution toward knowledge as an ordinary factor of production versus an exceptional device for lowering costs. Goulder and Mathai (2000) are a parallel step on that journey. They assume that accumulated R&D—knowledge— influences mitigation costs. Their cost function, therefore, depends on both the abatement level and the stock of knowledge: $C(A_t, H_t)$ where A_t is the abatement level, H_t is the stock of knowledge, and $C()$ is the abatement cost function. They make particular assumptions about functional form ($C()$ is proportional to H^1) and knowledge accumulation (including adjustment costs). Finally, they calibrate the model so that when knowledge accumulation is allowed, it lowers costs by 30 percent (based on work by Manne and Richels [1992]). Not surprisingly, they find that endogenizing technological change lowers costs by about 30 percent.

Finally, multi-sector general equilibrium models allow one to move away from mitigation cost functions and to think concretely about R&D influencing costs and productivity in particular sectors, such as “carbon free energy” or even a particular energy-generating technology. Goulder and Schneider (1999) do this in a seven-sector model where spillovers exist within, but not among, each sector. That is, each sector has a top-level constant elasticity of substitution production function of the form

⁷ Popp (2004a) presents upper and lower bounds of the influence of opportunity costs. The upper-bound estimates assume that R&D levels change as the opportunity costs change. The lower-bound estimates change the opportunity costs but constrain R&D to remain at the base-case levels. The later can be interpreted as the gains from removing the assumption of partial crowding-out. They are not the gains that would result if the government intervened to remedy the problem of partial crowding-out.

$$(3) Q = \gamma(\bar{H})(\alpha_H H^\rho + \alpha_G G^\rho)^{1/\rho}$$

where Q is total output, H is appropriable sector-specific knowledge, G is an aggregate of other inputs, and \bar{H} is non-appropriable sector-specific knowledge (α 's and ρ are parameters). Both H and \bar{H} accumulate research R_t , with a parameter β governing how much spills over into the non-appropriable stock.

They have a specific R&D sector, which also creates the possibility of crowding out as marginal R&D costs increase with supply. Without spillovers, they find that endogenizing technological change creates the possibility of lowering the cost of achieving a given emissions target. With spillovers, this effect can be muted or reinforced depending on the pattern of spillovers across sectors. In particular, they find that asymmetries in spillovers matter more than other specification choices.

Summarizing Behavior

The R&D modeling discussion is organized around two overarching questions: the form of the model (whether it focuses on emissions intensity, abatement cost, or sectoral productivity) and the parameterization of R&D effects. That is, the assumptions about the return to R&D investments (e.g., spillovers) and the opportunity costs of the endogenous R&D (e.g., crowding out). Regardless of form, the parameterization determines how endogenizing technological change affects the results relative to an exogenous technology assumption.

The single most important feature in these models is how the model accounts for imperfections in knowledge markets that produce high social rates of return. Models that omit these imperfections fail to capture both above-market-rate returns to climate-related R&D and above-market-rate returns to other R&D that exist. In cases where they are calibrated to have such returns without these imperfections, allowing the models to then equilibrate R&D and market returns will either: 1) overestimate the level of innovative activity that will be induced by climate policy; or 2) implicitly assume that appropriate government policies to correct R&D market imperfections will accompany climate policy, so that social and private rates of return are equalized. Thus, while such results may not be incorrect, it is important to note that results from models that do not explicitly include R&D market imperfections would not be obtained without government support for R&D. Moreover, these models not only assume that government provides R&D support but that it provides this support in an optimal manner so the gap between private and public returns is closed.

To aid understanding of how spillovers are included in models, we can imagine two summary statistics that quantify the key assumptions. First, the magnitude of spillovers can be measured by a simulation showing how an incremental dollar of environmental R&D changes welfare, instituted by an exogenous, “manna-from-heaven” increase in the environmental R&D budget. Positive spillovers would suggest, for example, that an additional \$1 million of R&D should improve welfare by more than \$1 million.⁸ Second, the model should be simulated with a climate policy in place and the environmental R&D budget both frozen at the reference-case level and determined endogenously. Here, it would be useful to calculate the change in welfare divided by the change in environmental R&D expenditures. This would represent an estimate combining both positive spillover effects and the opportunity costs of R&D. A positive number would indicate that spillovers from environmental R&D exceed any opportunity costs; a negative number would indicate they do not. In simpler models, these measures would be synonymous with the parameter choice. In the multi-sector models, the latter question tends to be lurking in assumptions about the increasing cost of R&D and the pattern of R&D returns in various sectors. This has the appearance of being more structural but begs the question of the degree of spillovers across sectors, something Goulder and Schneider (1999) ignore and something about which the empirical literature is silent.

A second pertinent question is how the marginal effects noted above evolve as a simulation deviates from the baseline. Put another way, what is the longer term potential for innovative activity? The original literature (see, for example, Leontief 1966) on computable general equilibrium modeling emphasized its usefulness for analyzing relatively small changes from a baseline, yet what climate-change models frequently are asked is how an entirely different future might evolve. Is the potential for large rates of return unlimited or will large increases in R&D result in diminishing returns? For example, Popp (2006b) finds evidence that diminishing returns to R&D are most prominent when large spikes in R&D occur. In the case of technological change, it would be useful to know how different model forms and parameterizations translate into long-term potential for cost-saving or productivity-enhancing R&D.

⁸ While the empirical literature predicts returns of roughly a 4:1 ratio, note that this will be slightly higher if the welfare calculation in the model includes avoided damages from clean technology. That is because studies finding a 4:1 ratio of social to private returns for R&D focus on returns measurable in markets. Any environmental benefits that occur will be in addition to this 4:1 ratio.

This question of deviations from the baseline recalls our earlier remark on the importance of how well these models have been calibrated to match important quantitative benchmarks. Do they match historic rates of productivity improvement in the baseline simulations? Do historic responses to price changes match simulated behavior? While such calibrations are less important for qualitative conclusions from simulation modeling, they become increasingly important if simulations are used for normative policy evaluation.

Looking Forward

A key observation from the preceding discussion is that often it is unclear what simulation models are calibrated to deliver, how far simulations are deviating from baseline behavior, and how the technology parameterization translates into intermediate conclusions about the returns to R&D and opportunity costs. A very useful near-term research agenda for the modeling community would be to clearly articulate the following summary features of models featuring endogenous technological change going forward:

1. Baseline productivity growth (e.g., per capita income).
2. Near-, medium-, and long-term elasticity of energy demand, in aggregate, and for different sectors, if present.
3. For policy simulations, report the range of variation in key relative prices (among various energy sources and aggregate consumption) and compare to historic price variation.
4. Elasticities of energy R&D with respect to changes in energy prices. Are these consistent with existing empirical evidence?
5. Clearly state assumptions used about R&D markets. Are spillovers modeled (either explicitly or implicitly)? What is the source of new energy R&D spending?
6. Estimate the marginal returns to energy R&D. Do this both through an exogenous, incremental increase in energy- or carbon-related-R&D, as well as by simulating climate policy with and without energy R&D frozen at reference levels (see comment in text).
7. Estimate the potential for long-term returns to energy R&D. As before, do this both through an exogenous, doubling (or other nonincremental change) in energy- or carbon-related-R&D, as well as by simulated climate policy with and without

energy R&D frozen at reference levels (but here considering a more aggressive, nonincremental climate policy).

Multi-sector general equilibrium models present additional data challenges. In particular, as those models contemplate specific knowledge stocks in different sectors—and often in different geographic regions—they require information quantifying spillovers not just in a given domestic industry, where most of the empirical literature has focused, but across industries and countries or regions. This is particularly important for energy R&D, since energy is an input to all production processes. Thus, fundamental research on new energy technologies may serve as a general-purpose technology that provides returns throughout the economy.⁹ This would be a useful area for additional empirical work.

Another area for future research is the role of technological change in baseline simulations. While climate policy increases incentives for emissions-reducing R&D, some emissions R&D can occur even without policy. This is particularly true for R&D that improves energy efficiency, which is encouraged by higher energy prices that may or may not be the result of environmental policy (see, for example, Popp 2002). Moreover, while large-scale climate policies only have recently taken hold (such as trading of carbon emissions permits in the European Union), several smaller policies have helped encourage the development of climate-friendly technologies. Policies such as government R&D subsidies, investment tax credits for renewable energy, and renewable portfolio standards all play a role in encouraging the development of these technologies. Any model that calibrates energy markets to current data trends will include technological change induced by current market forces and policies in the baseline. Given this, baseline runs are more appropriately interpreted as runs assuming extensions of existing policies rather than no policy. Moreover, rather than considering all technological change to occur as a result of climate policy, the appropriate question is what is the marginal increase in technology that occurs because of newly introduced climate policy.

Learning by Doing

The popular alternative (or compliment) to modeling technological change as a function of R&D is to model technological progress—particularly the lowering of costs of an existing technology—as a response to increased production, or learning by doing (LBD). As noted, the

⁹ Helpman (1998) provides an introduction to general-purpose technologies.

consequence of endogenous versus exogenous technological change through explicit, additional R&D expenditures depends on the relative importance of spillovers versus crowding out. The consequence of endogenous versus exogenous technological change through LBD is unambiguously positive: there is no opportunity cost other than the current cost of production. This suggests that assumptions of one mechanism or another will have important implications for model results. Most importantly, the potential of crowding out of R&D implies a higher opportunity cost to technological change that progresses via R&D, rather than by LBD. In contrast, investing in new energy technology (as in the LBD framework) does not include this added social cost.

Key Features

Because the social costs of technological change are lower in the LBD framework, models relying on LBD tend to provide larger estimates of the potential gains from endogenous technological change. Thus, an important empirical question is the relevant importance of each type of learning. Unfortunately, few empirical studies address this question. Such research is limited both by data availability (as measures of costs, technology usage, and matching R&D data would be needed) and concerns over endogeneity of key variables. A typical learning curve model regresses costs as a function of cumulative capacity:

$$(4) \text{ Cost} = A \cdot \text{CumCap}_i^{-\alpha}$$

where *Cost* is the cost of investment of a new installation, and *CumCap* is the cumulative installed capacity of (or sometimes cumulative output produced by) the technology prior to time *t*.

Recently, papers by Klaassen et al. (2003), Söderholm and Sundqvist (2003), and Söderholm and Klaassen (2007) have extended this basic model to estimate “two-factor” learning curves for environmental technologies. These two-factor curves model cost reductions as a function of both cumulative capacity (LBD) and R&D (learning by searching, or LBS). In these models, R&D typically is aggregated into a stock of R&D capital to be comparable with the notion of cumulative capacity. Thus, endogeneity is a concern, as we would expect both investments in capacity to be a function of past R&D expenditures and R&D expenditures to be influenced by capacity, which helps determine demand for R&D. Söderholm and Sundqvist (2003) address this endogeneity in their paper and find learning-by-doing rates around 5 percent

and learning-by-searching rates around 15 percent, suggesting that R&D, rather than LBD, contributes more to cost reductions.¹⁰

This work is preliminary and needs to be replicated with other data sets before it can be considered a definitive answer to the relative role of LBD versus LBS. However, it can be used to provide some guidelines for modeling technological change. First, it suggests that both mechanisms (LBD and LBS) are relevant. As of yet, however, no study has addressed potential interactions between LBD and LBS. For example, experience may reveal which attributes of a technology need attention and thus inform the direction of future R&D. Second, it suggests that the relative importance of each must be considered when calibrating models including both types of learning. The main lesson here is to avoid double counting. A learning-by-doing rate of just five percent is lower than typically reported in the LBD literature, where learning rates of 15 percent to 20 percent are common. A simple, one-factor LBD curve shows the relationship between capacity and costs but does not address causation. A two-factor curve begins to address this causation. As such, while the combined effect of LBD and LBS in a two-factor model may be comparable to learning rates from a one-factor model, the individual components should be smaller.

Modeling Approaches in Practice

Despite concerns about causality, the empirical tractability of learning curves has led to the wide use of learning-induced technological change throughout the literature, particularly in disaggregated, “bottom-up” models (including the NEMS model used by the Energy Information Administration; EIA 2003). Disaggregated models are well-suited for incorporating learning because of their rich technology specificity, which more easily lends itself to a learning curve for each technology. Some more aggregated models also use learning, but it is not as common. One reason is that learning tends to be thought of as a technology-specific phenomenon and is therefore harder to apply in the typical aggregation of a “top-down” model.

The most common way to capture learning-induced technological change in climate policy models is based on an exponential relation between unit cost and cumulative output, as in

¹⁰ To avoid problems with endogeneity, Nemet (2006) uses simulation techniques to decompose cost reductions for PV cells into seven categories. Plant size (e.g., returns to scale), efficiency improvements, and lower silicon costs explain the majority of cost reductions. Notably, most of the major improvements in efficiency come from universities, where traditional LBD through production experience would not be a factor. Learning from experience (e.g., through increased yield of PV cells) plays a much smaller role.

(4) above. This implies that a doubling of experience will reduce specific costs by a factor of $2^{-\alpha}$, also known as the progress rate. This formulation requires only the output and cost history to parameterize the learning function. Two-factor learning models, where both cumulative capacity and cumulative R&D enter the cost function, can be represented by appending $(R\&D)^{-\beta}$ to (4).

A recent summary (Section 11.5, IPCC WGIII 2007) of four bottom-up energy models reported single-factor learning models with rates between 4 percent and 19 percent—and in the case of solar photovoltaics in the MESSAGE model, 28 percent. Among those models that can include two-factor learning models, the second factor tends to be additive; that is, the learning-by-doing rate remains the same, while the effect of R&D is added. Doubling R&D is modeled to have between a 1 percent and 11 percent effect on costs.

Goulder and Mathai (2000) provide one example of an aggregate approach to modeling LBD. Rather than modeling cumulative knowledge to be a function of R&D investment, they assume it is a function of abatement. With this specification, current abatement acts as a learning investment in knowledge, analogous to R&D investment. The result is similar to many other learning-by-doing studies in that both the optimal carbon tax is lower at all points in time and that for any given carbon tax, there may be considerably more total abatement. And, as is the case with ordinary learning models, the effect of LBD on the slope of the optimal path of abatement is ambiguous, as learning both motivates earlier abatement to further lower future costs and later abatement to then take advantage of those lower costs.

Summary and Looking Forward

Much like the R&D discussion, an important agenda for the modeling community is clearer reporting of benchmark model results and responsiveness and comparison to historical, empirical data. Among the seven features noted earlier, the first three remain relevant: baseline productivity growth; elasticity of energy demand over the near-, medium-, and long-term; and, for policy simulations, a comparison of the range of simulated, relative prices compared to historic variation. In models that combine R&D and LBD, it is particularly important to conduct this type of benchmarking to ensure that the model is not double counting the potential gains from technological change and continues to replicate key benchmark features.

A key challenge for moving our understanding of LBD forward is disentangling the causal forces that lead to lower costs. While the statistical correlation between experience and lower costs is strong, understanding the causes of cost reductions is necessary to make informed policy decisions. Unfortunately, separately identifying the various learning mechanisms is

difficult. Learning through experience is confounded with more conventional learning through R&D, where R&D may be poorly measured and typically related to output and capacity. An important area for further work is to try and disentangle these effects.

Government R&D

We have focused primarily on the incentives faced, and activities conducted, by private firms. However, governments also play an important role in energy R&D. For example, in 2003, U.S. industry spent \$1.8 billion on energy R&D (National Science Foundation, 2006), while the U.S. Department of Energy spent \$2.5 billion (U.S. Department of Energy 2004).¹¹ This government investment plays several roles, each of which has important implications for climate-policy models.

First, note that government R&D can help to compensate for underinvestment by private firms. Unlike firms, the government is in a position to consider social returns when making investment decisions. In addition, government R&D tends to have different objectives than private R&D. Government support is particularly important for basic R&D, as long-term payoffs, greater uncertainty, and the lack of a finished product make it difficult for private firms to appropriate the returns of basic R&D.

This combination of public and private efforts complicates efforts to model R&D. Focusing on private efforts suggests that a wedge should exist between social and private returns to R&D. Government R&D aims, at least in part, to close this gap. However, there is little empirical evidence specifically on the returns to government R&D. This is due, in part, to the nature of government projects, which are often more basic and long term in nature, making estimating returns difficult. Given this, estimating the gap between private and social rates of return that exist after accounting for both private and public energy R&D spending is an important question for future research.¹²

If government R&D successfully narrows the gap between private and social rates of return, this effect should be acknowledged in baseline models. However, if that is the case, it is

¹¹ Note that this is not the total Department of Energy budget, much of which goes to defense projects. Rather, it is Department of Energy budget items focused on energy supply or energy efficiency.

¹² Continued calls for additional government support for energy R&D suggest that government funding has not closed this gap completely.

also important for models to acknowledge that such policy interventions will need to continue for the simulated results to hold. That is, if a baseline simulation implicitly includes existing government R&D efforts as part of its calibration, policy simulation must assume that this government R&D continues. Furthermore, any additional government R&D that occurs in a policy simulation properly should be credited as a marginal increase over baseline government R&D efforts.

In addition to correcting for underinvestment by private firms, many government R&D projects aim to improve commercialization of new technologies (or “transfer” from basic to applied research). Such projects typically combine basic and applied research and often are government/industry partnerships (National Science Board 2006). For example, the United States passed several policies in the 1980s specifically designed to improve transfer from the basic research done at government and university laboratories to the applied research done by industry to create marketable products.¹³ As such, this technology transfer can be seen as a step between the processes of invention and innovation. Empirical work suggests that these policies have had the desired effect. For example, Dai, Popp, and Bretschneider (2005) find that the Bayh-Dole Act, which facilitated university patenting, led to more university patents but not more R&D, suggesting that universities responded to the Bayh-Dole Act by patenting results that previously would have gone unpatented.¹⁴ Related to energy, Popp (2006b) finds that energy patents spawned by government R&D are cited more frequently than other energy patents. More importantly, after passage of these technology-transfer acts, the “children” of these patents (the privately held patents that cite government patents) are the most frequently cited patents, suggesting that transferring research results from the government to private industry produces valuable research results.

The role of government R&D for commercialization is not explicitly modeled in climate-policy simulations, as most models do not include the necessary level of detail. However, they provide important lessons for interpretation. In particular, these results suggest that model results on the potential of technological change should be considered a best-case scenario. For such results to be realized in practice, policymakers need to be aware of the lessons from the

¹³ Examples include the Stevenson-Wylder Technology Innovation Act of 1980, the Bayh-Dole Act of 1980, and the Federal Technology Transfer Act of 1986.

¹⁴ The hope of the Bayh-Dole Act was that these patents would increase licensing opportunities, so that private firms would help to market university inventions. See Sampat (2006) for further discussion.

microeconomic literature on technology transfer and ensure policies are in place to make optimal use of the results of both public and private R&D.

In the end, simulation models used for policy analysis rarely include government R&D, even though government R&D tends to be a frequent policy lever. But, as this section highlights, we have little empirical evidence that could be used to model the interaction of public R&D, private R&D, and eventual technological change. And, even if such historical evidence existed, it might be hard to extrapolate given the wide variety of government R&D activities, management structures, and likely effectiveness over time. Perhaps the best guidance for the moment is to recognize that current trends in technological change exist alongside substantial government expenditures that likely need to continue in order to maintain those baseline trends. Moving forward, it would be valuable if future empirical work identified the effect of significant changes in government R&D, distinct from private R&D and other relative price changes.

What about Technology Diffusion?

The process of technological change ends with diffusion. Technological advances are of little use unless society makes use of the innovation. Diffusion refers to the process by which a new technology gradually penetrates the relevant market. Often times, a technology that appears to surpass competing technologies in performance and cost will not immediately be chosen over existing technologies. A key question is whether this slow diffusion is a result of rational actors responding to varying incentives or due to market inefficiencies. As the preceding sections illustrate, there is a robust literature on the links between environmental policy and innovation, including both empirical studies and policy modeling, as well as representations of LBD. In contrast, diffusion of environmental technologies once they become cost-effective receives less attention. While diffusion is clearly a distinct area of empirical work, climate models diverge in their treatment of diffusion: top-down models without technology detail typically subsume diffusion into trends in sectoral productivity, while technology-rich bottom-up models typically include more ad hoc assumptions to constrain the penetration of new technologies.

Despite the mixed and relatively ad hoc treatment of diffusion in climate models, encouraging the diffusion of new technologies is clearly an important policy goal. Examples of policies that encourage diffusion of alternative technologies include tax credits for hybrid vehicles and California's Million Solar Roofs plan. In this section, we briefly review the empirical and simulation literature on technology diffusion. Here, there are a couple of issues to consider. One is the time lag between invention and adoption. Second is the flow of knowledge

across regions. In each case, we ask where each strand of research could usefully inform or be informed by the other.

The Timing of Diffusion

The diffusion of a new technology is a gradual, dynamic process; new technologies are not adopted en masse. Rather, adoption usually begins with a few early adopters, followed by a period of more rapid adoption, with the rate of adoption leveling off once most potential users have adopted the technology. This process generates the well-known S-shaped diffusion curve: the rate of adoption rises slowly at first, speeds up, and then levels off as market saturation approaches. The figure on page 30 illustrates such a curve, with market penetration plotted on the y-axis and the passage of time plotted on the x-axis. Early attempts to explain this process focused on the spread of information (e.g., epidemic models, such as Griliches [1957]) and differences among firms (e.g., probit models, such as David [1969]). More recently, researchers combined these explanations while adding potential strategic decisions of firms. These papers find that firm-specific differences explain most variation in adoption rates, suggesting that gradual diffusion is a rational process in response to varying incentives faced by individual actors.

Environmental technologies can be different, however. Incentives to adopt end-of-pipe technologies that serve only to reduce emissions must come from environmental regulation. Therefore, it is not surprising that the few studies that address adoption of environmental technologies find that regulations dominate all other firm-specific factors. Gray and Shadbegian (1998), Kemp (1998), Kerr and Newell (2003) and Popp (2005) all show that environmental regulations increase the probability of adopting environmentally friendly technologies. While the response varies by firm characteristics, regulations are the most important explanatory variable.

In contrast, energy-efficiency and fuel-saving technologies may be adopted more slowly, as it is cost savings, rather than a direct regulatory requirement, that matter. This can be influenced by policies that raise prices. Rose and Joskow (1990), for example, find that the adoption of fuel-saving technology by U.S. electric utilities is correlated positively with fuel prices. However, gradual diffusion is more likely and, even without climate policy, some adoption will occur, much like traditional technologies.

This suggests that modeling climate policy ideally would distinguish among different types of climate-mitigating technology, with particular emphasis on whether they might be subject to direct regulation. For example, mandating IGCC plants, hybrid cars, or even more

efficient building shells would imply more rapid technology diffusion of the sort noted for end-of-pipe technologies. Such policies might be motivated in part because of evidence of market failures in technology adoption beyond the pricing of greenhouse gas externalities. Jaffe and Stavins (1995) find that higher energy prices lead to greater use of insulation in new home construction but that the costs of installation are a more important consideration. This is part of a more general literature that finds inefficiencies in adoption of energy-saving technologies by consumers, with various explanations ranging from consumers using high discount rates, to credit-constrained consumers caring more about up-front costs than lifetime cost savings, to uncertainty over future costs (Jaffe et al., 2003). Hence, while the general result of the diffusion literature is that gradual diffusion is rational, market failures may contribute to slower adoption of energy-efficient technologies.

Despite the interest in these types of policies, most top-down climate models without technology detail typically aggregate to the industry level. Such models have a difficult time modeling technology mandates, let alone making technology diffusion responsive to such policies. Bottom-up models and others with technology detail do have the capability to model such mandates but instead tend to focus on ad hoc and unresponsive assumptions constraining technology diffusion. The MERGE model, for example, assumes new technologies are constrained to 1 percent of production in the year they are introduced and can expand by a factor of three in each decade thereafter (Manne and Richels 2002). The NEMS model used by the U.S. Energy Information Administration uses both adjustment costs to increase the short-run costs of rapidly increasing capacity and absolute constraints based work by Mansfield and Blackman (DiPardo 2000). In contrast, the IGSM model introduces an additional, fixed factor that acts as an adjustment cost to mimic gradual penetration (McFarland et al. 2004). While this does make penetration responsive to price incentives, it does not consider the potential increase in diffusion rates created by a mandate or targeted regulation.

Absent this interest in modeling the effect of mandates on diffusion, it is worth noting that in many cases, simplistic representations that basically ignore diffusion may be sufficient. It is not unusual to see market penetration of a product in 10 years or less. For example, Popp (2001) finds that new energy patents have their largest effect on energy consumption within four years. Combined with the results of Popp (2002) that find that the median lag of energy-patenting activity with respect to prices is also about 4 years, this suggests that the main effect of price-induced technological change will occur within one decade. Thus, gradual diffusion will be less relevant in models where the time periods modeled are decades rather than individual years. Moreover, in highly stylized models that include highly aggregate representations of technology,

calibrating the model based on the observed effects of technological change will implicitly allow for diffusion, as the observed effects are not the full potential of technology but rather the realized gains from technology that has diffused. For example, Schwoon and Tol (2006) develop a top-down model that includes adjustment costs that effectively slow the penetration of new technology.

Diffusion across Regions

A further complication to modeling the generation and adoption of new technologies is the flow of these technologies across regions. This is particularly important when we consider technologies available to developing countries, such as China or India. The vast majority of R&D effort occurs in highly developed countries. In 2000, global R&D expenditures were at least \$729 billion—82 percent of this research was done in the Organisation for Economic Cooperation and Development countries and half was performed by the United States and Japan alone (National Science Board 2006). Moreover, in a global economy, policies in one nation may affect innovation in other nations. This is particularly likely when the regulating economy is an important market for other nations. For example, Lanjouw and Mody (1996) find that the majority of vehicle air-emissions patents granted in the United States are from foreign nations, even though the United States was the first country to adopt strict emissions standards.

There is a large body of literature on technological diffusion across countries (surveyed by Keller [2004]) but little work focusing specifically on environmental technologies. In general, there are two potential avenues through which foreign knowledge generally can have an influence in the domestic economy. First, as noted in our discussion of R&D spillovers, foreign knowledge may affect the productivity of domestic R&D. The “blueprints” represented by foreign patents may serve to inspire additional innovation by domestic inventors. Such productivity increases are knowledge spillovers, as the knowledge represented by the foreign patent creates a positive externality—knowledge from abroad that is borrowed by the domestic inventor (Griliches 1979). Second, new technologies owned by foreigners simply may be diffuse and adopted directly by domestic firms.

Buonanno, Carraro, and Galeotti (2003) provide perhaps the only example of foreign-knowledge spillovers in a simulation model, with a somewhat ad hoc calibration of spillover parameters to match R&D-output ratios. More generally, knowledge spillovers, and especially international technology transfer (after a technology is available for export), have not been modeled directly in simulations of climate policy. Given these two avenues for knowledge to enter the economy, representing foreign-knowledge influences is an important area for

development in integrated assessment models. At the same time, empirical research on the extent to which foreign knowledge influences domestic inventors will be necessary to calibrate these models, as little of the work on foreign-knowledge spillovers is at the level of aggregation necessary to calibrate such a model.

There also may be need to consider the interaction of international diffusion and domestic R&D. Lanjouw and Mody (1996) and Popp (2006c), for example, find evidence of innovation in countries that adopt regulations late, suggesting that these countries do not always take advantage of technologies “off the shelf” that have been developed elsewhere. Instead, adaptive R&D seems to be necessary to suit the technology to the local market, as these later patents are more likely to cite earlier foreign, rather than domestic, inventions. The need for adaptive R&D suggests that either: a) diffusion across countries will be slower than within countries, as an additional research step is needed to fit the technology to local conditions; or b) that foreign knowledge adopted directly (i.e., without adaptive R&D) will be an “imperfect fit” for local conditions and thus will be less effective than in the technology’s initial market.

Conclusion

Technological change is at once the most important and least understood feature driving the future cost of climate-change mitigation. Quantitative research has tended to follow two distinct and often isolated paths: empirical, microeconomic work on the nature of R&D and technological progress and simulation modeling of macroeconomic models with some technology detail and assumptions about how productivity and technology cost respond to various activities. While these simulation models often are used to provide quantitative guidance to policy debates, a review of recent work suggests that efforts to endogenize technological change often fail to calibrate themselves to important empirical benchmarks, at worst, and fail to clearly communicate their assumptions, at best. Meanwhile, microeconomic research often fails to focus on important questions asked by the simulation modelers.

What might a coherent research agenda look like, based on a review of both research areas? First, simulation models should communicate more clearly what benchmarks are being matched in terms of baseline productivity growth and energy-price elasticities, as well as the gross and net returns to energy R&D (both incremental and discrete increases from the baseline). To the extent that models involve multi-sector, or multi-region, R&D, assumptions about spillovers should be explained. When models include both R&D and LBD, efforts to avoid double counting the potential for R&D or experience to lower costs should be explained.

Second, microeconomic work should attempt to provide more information about spillovers across industries and regions. A high priority should be efforts to disentangle R&D and learning-by-doing effects to avoid the double counting just noted. Finally, given the prevalence of government R&D proposals, it would be valuable to explore the relative contribution of public versus private R&D. Simulation work also might usefully explore how both types of R&D could be usefully employed in a model.

The Intergovernmental Panel on Climate Change's *Third Assessment Report on Mitigation* noted the importance of exogenous technology assumptions, particularly the rate of improvement in energy efficiency, in estimating mitigation costs. Six years later, the *Fourth Assessment Report on Mitigation* describes a veritable explosion of research into how endogenizing technological change affects costs. Yet, in many ways, we are no closer to understanding the answer of how policies will affect the overall rate and direction of technological change. We have replaced previously transparent, exogenous assumptions with complex and often opaque assumptions about endogenous pathways. With a more coherent agenda combining empirical microeconomic analysis with a clear eye to the needs of the community of simulation modelers, and vice versa, perhaps we can improve our understanding of how technological change is likely to respond to current and future mitigation policies.

References

- Bazon, Coleman, and Kent Smetters. 1999. Discounting Inside the Beltway. *Journal of Economic Perspectives* 13(4): 213–228.
- Brunneimer, S., and M. Cohen. 2003. Determinants of Environmental Innovation in US Manufacturing Industries. *Journal of Environmental Economics and Management* 45: 278–293.
- Buonanno, Paolo, Carlo Carraro, and Marzio Galeotti. 2003. Endogenous Induced Technical Change and the Costs of Kyoto. *Resource and Energy Economics* 25(1): 11–34.
- Byatt, Ian, Ian Castles, Indur Goklany, David Henderson, Nigel Lawson, Ross McKittrick, Julian Morris, Alan Peacock, Colin Robinson, and Robert Skidelsky. 2006. The Stern Review: A Dual Critique, Part II: Economic Aspects. *World Economics* 7(4): 199–224.
- Clarke, Leon, James Edmonds, Henry Jacoby, Hugh Pitcher, John Reilly, and Richard Richels. 2006. *CCSP Synthesis and Assessment Product 2.1, Part A: Scenarios of Greenhouse Gas Emissions and Atmospheric Concentrations*. Draft report. Washington, DC: U.S. Climate Change Science Program.
- Dai, Yixin, David Popp, and Stuart Bretschneider. 2005. Institutions and Intellectual Property: The Influence of Institutional Forces on University Patenting. *Journal of Policy Analysis and Management* 24: 579–598.
- David, Paul A. 1969. *Contribution to the Theory of Diffusion*. Stanford Center for Research in Economic Growth, Memorandum #71. Stanford, CA: Stanford University.
- DiPardo, Joseph. 2000. Outlook for Biomass Ethanol Production and Demand. Washington, DC: U.S. Energy Information Administration.
- Energy Information Administration. 2003. The National Energy Modeling System: An Overview. Washington, DC: Energy Information Administration.
- Fischer, Carolyn, and Richard G. Newell. 2004. Environmental and Technology Policies for Climate Change and Renewable Energy. Discussion Paper 04-05. Washington, DC: Resources for the Future.
- Gillingham, Kenneth, Richard Newell, and William Pizer. 2007. Modeling Endogenous Technological Change for Climate Policy Analysis. RFF Working Paper. Washington, DC: Resources for the Future.

- Goolsbee, Austan. 1998. Does Government R&D Policy Mainly Benefit Scientists and Engineers? *American Economic Review* 88(2): 298–302.
- Goulder, Lawrence, and Koshy Mathai. 2000. Optimal CO₂ Abatement in the Presence of Induced Technological Change. *Journal of Environmental Economics and Management* 39: 1–38.
- Goulder, Lawrence, and Stephen Schneider. 1999. Induced Technological Change and the Attractiveness of CO₂ Abatement Policies. *Resource and Energy Economics* 21: 211–253.
- Gray, W.B., and R.J. Shadbegian. 1998. Environmental Regulation, Investment Timing, and Technology Choice. *Journal of Industrial Economics* 46: 235–256.
- Griliches, Z. 1957. Hybrid Corn: An Exploration in the Economics of Technological Change. *Econometrica* 25: 501–522.
- Griliches, Z. 1979. Issues in Assessing the Contribution of Research and Development to Productivity Growth. *Bell Journal of Economics* 10(1): 92–116.
- Hall, Bronwyn H. 1996. The Private and Social Returns to Research and Development. In *Technology, R&D and the Economy*, edited by B.L.R. Smith and C.E. Barfield. Washington, DC: The Brookings Institution and American Enterprise Institute, 140–162.
- Helpman, E. (ed.). 1998. *General Purpose Technologies and Economic Growth*. Cambridge, MA: MIT Press.
- Intergovernmental Panel on Climate Change. 2007. *Climate Change 2007: Mitigation of Climate Change*. Geneva: Intergovernmental Panel on Climate Change.
- Jaffe, Adam B. 1986. Technological Opportunity and Spillovers of R&D: Evidence from Firms' Patents, Profits, and Market Value. *American Economic Review* 76(5): 984–1001.
- Jaffe, A.B., and K. Palmer. 1997. Environmental Regulation and Innovation: A Panel Data Study. *Review of Economics and Statistics* 79(4): 610–619.
- Jaffe, A.B., R.G. Newell, and R.N. Stavins. 2003. Technological Change and the Environment. In *Handbook of Environmental Economics*, edited by K.G. Mäler and J. Vincent. Amsterdam: Elsevier, 461–516.
- Jaffe, A.B., and R.N. Stavins. 1995. Dynamic Incentives of Environmental Regulation: The Effects of Alternative Policy Instruments on Technology Diffusion. *Journal of Environmental Economics and Management* 29: S43–S63.

- Jones, C., and J. Williams. 1998. Measuring the Social Rate of Return to R&D. *Quarterly Journal of Economics* 113(4): 119–135.
- J. O. Lanjouw, and A. Mody. 1996. Innovation and the International Diffusion of Environmentally Responsive Technology. *Research Policy* 25: 549–571.
- Kerr, Suzi, and Richard Newell. 2003. Policy-Induced Technology Adoption: Evidence from the US Lead Phasedown. *Journal of Industrial Economics* 51(3): 317–343.
- Kneese, Allen V., and Charles L. Schultz. 1978. *Pollution, Prices and Public Policy*. Washington, DC: Brookings Institute.
- Klaassen, Ger, Asami Miketa, Katarina Larsen, and Thomas Sundqvist. 2003. Public R&D and Innovation: The Case of Wind Energy in Denmark, Germany, and the United Kingdom, IR-03-011. Laxenburg, Austria: International Institute for Applied Systems Analysis.
- Keller, Wolfgang. 2004. International Technology Diffusion. *Journal of Economic Literature* 42:3: 752–782.
- Kemp, Rene. 1998. The Diffusion of Biological Waste-Water Treatment Plants in the Dutch Food and Beverage Industry. *Environmental and Resource Economics* 12: 113–136.
- Löschel, A. 2002. Technological Change in Economic Models of Environmental Policy: A Survey. *Ecological Economics* 43: 105–126.
- Manne, Alan, and Richard Richels. 1992. *Buying Greenhouse Insurance: The Economic Costs of CO₂ Emission Limits*. Cambridge, MA: MIT Press.
- Mansfield, Edwin. 1977. Social and Private Rates of Return from Industrial Innovations. *Quarterly Journal of Economics* 91: 221–240.
- Mansfield, Edwin. 1996. *Estimating Social and Private Returns from Innovations Based on the Advanced Technology Program: Problems and Opportunities*. NIST GCR 99-780. Gaithersburg, MD: National Institute of Standards and Technology.
- Mowery, D., and N. Rosenberg. 1979. The Influence of Market Demand Upon Innovation: A Critical Review of Some Recent Empirical Studies. *Research Policy* 8: 102–153.
- Pakes, A. 1985. On Patents, R&D, and the Stock Market Rate of Return. *Journal of Political Economy* 93(2): 390–409.
- Edenhofer, O., K. Lessman, C. K. Emfert, M. Grubb, and J. Köhler. 2006. Induced Technological Changes: Exploring its Implications for the Economics of Atmospheric

- Stabilization. Synthesis Report from the Innovation Modeling Comparison Project. *Energy Journal* (special issue): 1–51.
- Manne, A., and R. Richels. 2004. The Impact of Learning-By-Doing on the Timing and Costs of CO₂ Abatement. *Energy Economics* 26(4): 603–619.
- McFarland, J.R., J.M. Reilly, and H.J. Herzog. 2004. Representing Energy Technologies in Top-Down Economic Models Using Bottom-Up Information. *Energy Economics* 26(4): 685–670.
- National Science Board. 2006. *Science and Engineering Indicators 2006*. Arlington, VA: National Science Foundation.
- Newell, Richard, Adam Jaffe, and Robert Stavins. 1999. The Induced Innovation Hypothesis and Energy-Saving Technological Change. *The Quarterly Journal of Economics* 114(3): 941–975.
- Nordhaus, William. 1994. *Managing the Global Commons: The Economics of Climate Change*. Cambridge, MA: MIT Press.
- Nordhaus, William. 2002. Modeling Induced Innovation in Climate Change Policy. In *Technological Change and the Environment*, edited by A. Grubler, N. Nakicenovic and W. Nordhaus. Washington, DC: Resources for the Future Press, 259–290
- Otto, Vincent, Andreas Löschel, and John Reilly. 2006. *Directed Technical Change and Climate Policy*. MIT Joint Program on the Science and Policy of Global Change Report No. 134. Cambridge, MA: MIT Global Change Joint Program.
- National Science Foundation. 2006. *Research and Development in Industry: 2003*, NSF 07-314. Arlington, VA: National Science Foundation, Division of Science Resources Statistics.
- Nemet, Gregory F. 2006. Beyond the Learning Curve: Factors Influencing Cost Reductions in Photovoltaics. *Energy Policy* 34(17): 3218–3232.
- Pizer, William. 2002. Combining Price and Quantity Controls to Mitigate Global Climate Change. *Journal of Public Economics* 85(3): 409–434.
- Popp, David. 2001. The Effect of New Technology on Energy Consumption. *Resource and Energy Economics* 23(3): 215–239.
- Popp, David. 2002. Induced Innovation and Energy Prices. *American Economic Review* 92(1): 160–180.

- Popp, David. 2003. Innovation in Climate Policy Models: Implementing Lessons from the Economics of R&D. In *Proceedings of the 2003 Energy Modeling Forum Workshop on Climate Change Impacts and Integrated Assessment*, August 5, Snowmass, CO.
- Popp, David. 2004. ENTICE: Endogenous Technological Change in the DICE Model of Global Warming. *Journal of Environmental Economics and Management* 48: 742–768.
- Popp, David. 2005. Lessons From Patents: Using Patents to Measure Technological Change in Environmental Models. *Ecological Economics* 54(2-3): 209–226.
- Popp, David. 2006a. R&D Subsidies and Climate Policy: Is there a “Free Lunch”? *Climatic Change* 77(3-4): 311–341.
- Popp, David. 2006b. They Don't Invent Them Like They Used To: An Examination of Energy Patent Citations Over Time. *Economics of Innovation and New Technology* 15(8): 753–776.
- Popp, David. 2006c. International Innovation and Diffusion of Air Pollution Control Technologies: The Effects of NO_x and SO₂ Regulation in the U.S., Japan, and Germany. *Journal of Environmental Economics and Management* 51(1): 46–71.
- Rose, N. L., and P. L. Joskow. 1990. The Diffusion of New Technology: Evidence from the Electric Utility Industry. *RAND Journal of Economics* 21: 354–373.
- Rosenberg, N. 1974. Science, Invention and Economic Growth. *Economic Journal* 84: 90–108.
- Rosenberg, N. 1982. *Inside the Black Box: Technology and Economics*. Cambridge, UK: Cambridge University Press.
- Sampat, Bhaven N. 2006. Patenting and US Academic Research in the 20th Century: The World Before and After Bayh-Dole. *Research Policy* 35: 772–789.
- Scherer, F.M. 1965. Firm Size, Market Structure, Opportunity and the Output of Patented Inventions. *American Economic Review* 55: 1097–1123.
- Schmookler, Jacob. 1966. *Invention and Economic Growth*. Cambridge: Harvard University Press.
- Schneider, Stephen, and Lawrence Goulder. 1997. Commentary: Achieving Low-Cost Emissions Targets. *Nature* 389(4): 13–14.
- Schwoon, Malte, and Richard Tol. 2006. Optimal CO₂-Abatement with Socio-Economic Inertia and Induced Technological Change. *The Energy Journal* 27(4): 25–59.

- Söderholm Patrik, and Ger Klaassen. 2007. Wind Power in Europe: A Simultaneous Innovation-Diffusion Model. *Environmental and Resource Economics* 36: 163–190.
- Söderholm, P., and T. Sundqvist. 2003. Pricing Environmental Externalities in the Power Sector: Ethical Limits and Implications for Social Choice. *Ecological Economics* 46(3): 333–350.
- Stern, Nicholas. 2006. *The Stern Review of the Economics of Climate Change*. Cambridge: Cambridge University Press.
- Stern, Nicholas. 2007. Testimony before the U.S. Senate, Energy and Natural Resources Committee, February 13.
- U.S. Department of Energy. 2004. *US Department of Energy's FY2005 Budget Request to Congress*. Washington, DC: U.S. Department of Energy.

Figure

The S-Shaped Diffusion Curve

