Technology-push, Demand-pull, and Strategic R&D Investment

by

Benjamin D. Leibowicz
Assistant Professor
Graduate Program in Operations Research and Industrial Engineering
Department of Mechanical Engineering
University of Texas at Austin

ETC 5.128D
204 E. Dean Keeton St.
Austin, TX 78712
(512) 537-5296
bleibowicz@utexas.edu

Abstract

In this study, a bilevel model is developed to determine the combination of technology-push and demand-pull policies that induces the socially optimal level of innovation for a given technology policy application. The model is bilevel in that it features inner agents (profit-maximizing firms) and an outer agent (welfare-maximizing policymaker). The inner problem is an oligopoly game in which each firm solves a two-stage stochastic decision problem. The firm chooses process and product R&D investments in the first stage and then chooses output levels in the second stage. The outcome of product R&D is uncertain. In the outer problem, the policymaker identifies the combination of technology-push and demand-pull policy interventions that induces the firms to reach the Nash equilibrium with the highest social welfare. This study goes beyond previous analyses of strategic innovation in oligopoly settings in that it explicitly incorporates uncertainty and includes a leader-follower interaction between the policymaker and the private sector. The model captures three critical market failures: incomplete appropriability of R&D, a negative production externality, and imperfect competition. Findings reveal that the optimal combination of technology-push and demand-pull policies varies depending on whether the policy motivation is to address a negative externality, reduce cost, or create demand. Stronger spillovers reduce product R&D expenditures but raise welfare because they make each dollar of R&D more effective. While welfare decreases with competition in the absence of technology policy, welfare increases with competition if optimal technology policies can be imposed.

1 Introduction

1.1 Market Failures

Without a policy remedy, firms generally engage in less than the socially optimal level of innovation because they cannot appropriate all the benefits of their own innovative effort. This is the case because the benefits of innovating spill over across firms, a phenomenon which has been documented empirically. Jaffe (1986) found that, controlling for its own R&D spending, a firm produces more patents when total other-firm R&D spending is higher. Bernstein and Nadiri (1988) estimated inter-industry spillovers and found that variable costs in each industry they examined decreased significantly in response to R&D in certain other industries. A number of studies have quantified the gap between private and social rates of return to R&D that these spillovers induce. In a pioneering empirical work, Mansfield et al. (1977) estimated private and social rates of return to innovation using data collected from firms in numerous industries in the northeastern U.S. Social rates exceeded private rates in most (but not all) industries, with estimated median social and private rates of 56% and 25%, respectively. Bernstein and Nadiri (1988) confirmed this general result and showed that the gap varies considerably across industries; the social rate of return exceeded the private rate by only 10–20% in electrical products and transportation equipment, but by a factor of roughly ten in scientific instruments. An implication of the gap between social and private rates of return is that the private sector will not conduct some innovative activities that would be socially beneficial. Jones and Williams (1998) derived the social rate of return to R&D analytically in an endogenous growth model and determined that optimal R&D spending in the economy is at least two to four times actual R&D investment.

Underinvestment in R&D is likely to be particularly severe in industries where the innovation market failure is accompanied by an externality market failure. A useful example which will motivate a set of numerical simulations presented later in this article is the energy sector, where many of the currently dominant technologies generate negative externalities such as air pollution, climate change, and energy security risks. While the social optimum might include significant R&D investment to develop cleaner energy supply technologies and more efficient end-use devices, in the absence of policy, the private sector will not account for these benefits of innovation and perform too little R&D. Although climate change has become a prominent issue, energy R&D spending has generally declined in the U.S. for several decades relative to total R&D spending and the size of the economy. This is particularly true in the private sector, where total private energy R&D investment is less than the R&D budgets of individual biotechnology companies. In response, some experts have recommended increasing energy R&D investment by an order of magnitude, which they claim is justified and feasible (Nemet and Kammen, 2007). Bosetti et al. (2009) included R&D decision variables in an integrated assessment model (IAM) and concluded that optimal energy R&D spending is approximately triple the current level if our goal is to stabilize greenhouse gas (GHG) concentrations.

1.2 Market Structure and Innovation

There is a substantial literature on the relationship between market structure and innovation. Although different studies have produced contrasting findings, it is clear that the level of innovative activity is strongly influenced by the degree of competition in the market. The traditional line of reasoning beginning with Schumpeter (1943) was that competition erodes post-innovation rents and therefore reduces the incentive to innovate. This Schumpeterian view was challenged theoretically by Arrow (1962) and empirically by a number of subsequent studies which found that competition has a positive effect on innovation. For example, Blundell et al. (1995) analyzed data on British manufacturing firms from 1972 to 1982 and econometrically determined that more competitive industries generated a greater number of technologically significant and commercially important innovations.

Some important studies in the literature offer a compromise by claiming that the relationship between competition and innovation is described by an inverted-U shape. At low levels of competition the relationship is positive, as firms look to escape competition by innovating to gain an advantage over their rivals. But at high levels of competition the relationship is negative, as rents in the market are insufficient to incentivize significant investment in costly innovation efforts. The theoretical possibility of an inverted-U relationship was suggested by Kamien and Schwartz (1976), who demonstrated analytically that under certain conditions the maximum amount of innovation occurs with some intermediate degree of rivalry between monopoly and perfect competition. Aghion et al. (2005) provided empirical evidence supporting the inverted-U hypothesis and developed an explanatory theory.

1.3 Process vs. Product Innovation

The literature distinguishes between process and product innovations. Process innovations make existing production processes more efficient and reduce costs. Product innovations lead to new products based on fundamentally different technologies than anything which previously existed, and that offer new benefits (Scherer, 1982). The distinction between process and product innovation is roughly analogous to related distinctions in the literature, such as incremental and radical innovation (Chandy and Tellis, 2000), or exploitation and exploration (Levinthal and March, 1993). In what Abernathy (1978) famously referred to as the “productivity dilemma,” firms must allocate scarce R&D resources between process and product innovation efforts. A common theme of the literature is that firms tend to overinvest in process innovation at the expense of product innovation. This occurs because process innovation can be accomplished through smaller projects, is less risky, has more immediate payoffs, and provides benefits which are usually easier to appropriate. Incumbent firms producing an existing good often have a vested interest in resisting new products that could threaten demand, and process innovations which improve and reduce the cost of the existing good make it more difficult for new technologies to compete and gain traction (Chandy and Tellis, 2000). As a result, underinvestment in R&D is typically more severe for product innovation.

Despite a glaring need for policies that support product innovation, policy analysis models designed to evaluate optimal R&D policies and investment levels often represent R&D only as a cost-reducing process consistent with process innovation. For example, the IAMs that incorporate energy R&D typically employ two-factor learning curves whereby R&D investment deterministically and smoothly reduces costs (Bosetti et al., 2009; Kypreos, 2007). This formulation neglects two defining characteristics of product R&D. First, product innovation is risky. Returns are highly uncertain, most projects fail to develop a profitable product, and a small number of major successes justify the much larger number of failures (Scherer, 1993). Second, early benefits stemming from product innovation typically arise by stimulating new demands rather than by meeting existing demands at lower cost. Cost reductions allowing the new technology to compete with existing ones typically only occur after a prolonged period of commercialization including moves toward standardization and mass production (Grubler, 1998). Early adoption of a new product may therefore depend critically on the presence of market niches where the unique performance advantages of the product are particularly desirable (Norberg-Bohm, 2000). The importance and unique characteristics of product innovation, combined with the dearth of analytical tools that properly represent it, mean that technology policy models must develop more appropriate formulations of product innovation.

1.4 Oligopolistic Models

The modeling framework developed in this study draws inspiration from and builds upon a series of previous studies that analyzed optimal innovation in an oligopoly setting. The studies summarized in this subsection, like the model developed in this analysis, all make use of game equilibrium concepts to determine the amount of innovation that the market undertakes.

In a pioneering study, Dasgupta and Stiglitz (1980) analyzed an oligopoly with endogenous process innovation. Their theoretical conclusions are essentially consistent with the inverted-U hypothesis. While a monopolist generally has insufficient incentives to undertake R&D expenditure, at high levels of rivalry R&D effort declines with competition. Levin and Reiss (1988) expanded the Dasgupta and Stiglitz (1980) framework by incorporating product R&D in addition to process R&D, and by including R&D spillovers. They showed analytically that process R&D and product R&D may be complements or substitutes depending on the parameterization of the model, in particular the relative magnitudes of spillovers. D'Aspremont and Jacquemin (1988) restricted their analysis to a duopoly with process R&D and spillovers. Their model is a two-stage game in which the duopolists first choose R&D investments then compete in the product market. The authors compare social welfare in three cases: a fully cooperative case, a fully non-cooperative case, and a hybrid case where the firms cooperate in the R&D stage but not in the product market stage. Results show that social welfare can be increased by permitting firms to engage in cooperative research where they share the costs and results of a research project. Suzumura (1992) conducted an analysis similar to that of D'Aspremont and Jacquemin (1988) but generalized the model to an oligopoly and considered a wider variety of welfare specifications. His findings demonstrate that in the presence of large spillovers, both the cooperative and non-cooperative equilibria include insufficient innovation. Lin and Saggi (2002) investigated the relationship between process and product R&D by constructing a duopoly model. They considered both Cournot and Bertrand competition and found that Cournot firms invest more in process R&D while Bertrand firms invest more in product R&D.

The analysis described in this article goes beyond these previous studies in two critical dimensions. First, it incorporates uncertainty in R&D outcomes to make the firm's optimization problem a two-stage stochastic one. All the aforementioned studies employed deterministic models, although some identified uncertainty as an important avenue for future research (Suzumura, 1992). Second, it includes an additional layer of decision variables that represent policy interventions designed to induce the socially optimal level of innovation. Unlike previous studies which served primarily to characterize market failures, this analysis considers policy remedies explicitly within its modeling framework. Table 1 summarizes the features of the oligopolistic models analyzed in the previous studies discussed above and clarifies how the present analysis relates to these prior efforts.

Table 1. Relationship of the present analysis to previous studies in the literature.

Study	Duopoly	Oligopoly	Process R&D	Product R&D	Spillovers	Uncertainty	Policy Intervention
Dasgupta and Stiglitz (1980)		✓	✓				
Levin and Reiss (1988)		✓	✓	✓	✓		
D'Aspremont and Jacquemin (1988)	✓		✓		✓		
Suzumura (1992)		✓	✓		✓		
Lin and Saggi (2002)	✓		✓	✓			
Present Analysis		✓	✓	✓	✓	✓	✓

1.5 Technology-push and Demand-pull

The previous subsections have reviewed the literature on market failures that justify technology policy intervention. As depicted in Table 1, a unique feature of the present analysis is that the model includes an additional layer of decision variables that represent the policy interventions intended to induce the socially optimal level of innovation.

Policy instruments employed to stimulate innovation are often divided into two groups: technology-push instruments that reduce the private cost of engaging in innovation and demand-pull instruments that create or expand markets to increase the private payoff to successful innovation (Nemet, 2009). Technology-push options include public R&D, government funding of private sector R&D, and support for higher education to enlarge the pool of innovators. Demand-pull options include subsidies for consumer purchases of new technologies, direct government procurement, and stronger intellectual property protection to increase appropriability. Each group of policy interventions has been favored at various times, and has resulted in mixed outcomes. The remainder of this subsection highlights the mixed records of success and failure achieved by past technology-push and demand-pull policy initiatives.

Innovation policy in the Big Science era during and after World War II emphasized technology-push instruments. Massive public R&D efforts such as the Manhattan Project and Apollo Program typify this approach to innovation, which was articulated by Bush (1945) in a report to President Franklin D. Roosevelt. This technology-push emphasis was consistent with the traditional linear model of innovation, describing sequential stages of research, development, demonstration, and diffusion (Gallagher et al., 2012). Arthur (2007) defines invention as the process of linking some purpose or need with an effect that can be exploited to satisfy it. Although the Big Science efforts were initiated based on perceived needs (e.g., to win World War II or to compete in the Space Race), the scientific advances they spawned suggested many uses and applications of these discoveries far outside the originally perceived needs. For example, advances in nuclear science achieved during the Manhattan Project later resulted in nuclear electric power, and the Apollo Program sparked advances in computing and integrated circuits.

In recent years, many innovation policies based on technology-push without simultaneous demand-pull measures have failed to achieve their goals. The U.S. launched the Mod program in the 1970s to develop a reliable, cost-competitive wind turbine. It cost over $200 million and was jointly administered by NASA and DOE. Despite accumulating useful design experience and experimental data, the program failed to realize a commercially successful design. Its focus on a 3 MW machine at a time when all installations were less than 100 kW was too radical a departure from the market, and 3 MW turbines would not be built for the next 20 years (Loiter and Norberg-Bohm, 1999). Between 1998 and 2002, the Netherlands subsidized research on second-generation biofuels through its GAVE program. Despite technological successes in the laboratory, the program failed to result in a commercial demonstration because, without policies to stimulate a market for these biofuels, the program's commercial partners did not view it as profitable (Suurs and Hekkert, 2009).

According to the demand-pull hypothesis, innovation is a function of market demand. Early support was provided by Schmookler (1966), who estimated a strong, positive relationship between patent applications from capital-goods-producing sectors and investment downstream in capital-goods-using industries. Dasgupta and Stiglitz (1980) showed that total R&D expenditure and R&D expenditure per firm increase with market size in their oligopolistic model.

Implemented by themselves, demand-pull policies are rarely consistent enough to stimulate the private sector to invest in risky R&D efforts. In the early 1980s, a 25% federal tax credit and 25% California tax credit for investment in windfarms resulted in substantial capacity installations. Although this experience caused a five-fold decrease in wind electricity costs, the technology was still not competitive with fossil generation. When the tax credits expired in 1987, investment quickly declined and most turbine manufacturers folded (Loiter and Norberg-Bohm, 1999). While the demand-pull incentives led to cost-reducing process innovations, they failed to stimulate product innovation. According to Nemet (2009), the rapid growth of the industry involved convergence on a single dominant design, and short-term profit opportunities favored process innovation at the expense of product R&D efforts. Demand-pull policies can also have unintended, adverse consequences. California's solar installation tax credits in the late 1970s and early 1980s were expensive and regressive while doing little to improve energy conservation (Taylor, 2008). Examining more recent solar rebate programs in California, Wiser et al. (2007) found that system purchasers do not get the full benefits of the incentives. Instead, system producers and installers reap higher profits by raising prices, with pre-rebate prices largely tracking the rebates themselves.

1.6 Balanced Innovation Policy

The recent empirical and case study literatures exhibit an overwhelming consensus that a balance of technology-push and demand-pull policies is generally required to successfully support emerging technologies. The overall notion is that the two groups of policy instruments are complements, not substitutes, and that an effective innovation policy portfolio should include both (Gallagher et al., 2012).

Kleinknecht and Verspagen (1990) examined Schmookler's (1966) earlier empirical work supporting the demand-pull hypothesis. They showed that after correcting a weakness in Schmookler's estimation approach, the relationship between demand and innovation is much weaker (although still significant). Rather than lend unique support to the demand-pull hypothesis, the data suggest a simultaneous relationship between demand and innovation in which the opposite direction of causality (technology-push) is just as likely. Watanabe et al. (2000) described what they refer to as a "virtuous cycle" involving positive feedbacks among R&D, adoption, and price reductions in the development of the Japanese solar PV industry. Pavitt (1984) showed that the relative importance of technology-push and demand-pull vary across industries, implying that careful consideration should be given to the appropriate allocation of resources between the two policy channels for a particular application.

Although experts agree that an innovation policy portfolio should consist of both technology-push and demand-pull instruments, there is a dearth of theoretical and modeling tools available to determine the proper balance of interventions for a particular application. This is important, since the right balance of technology-push and demand-pull likely varies across industries and technologies. The present study addresses this gap in the literature by developing a model to assess the performance of various combinations of R&D and price subsidies under different parameterizations of technologies and markets. As illustrated in Table 1, this model is unique in that it explicitly incorporates policy interventions and uncertainty in R&D outcomes. More generally, in a relatively compact modeling framework, the present analysis considers three critical market failures –incomplete appropriability of R&D, a negative production externality, and imperfect competition – as well as technology-push and demand-pull policy remedies.

2 Model

2.1 Overview

Figure 1 presents a visual overview of the bilevel modeling framework developed in this analysis and described in greater detail in the subsections that follow. The figure takes the form of a decision diagram from the field of decision analysis. As depicted, the inner agents are profit-maximizing firms that choose process and product R&D investments in the first stage and output levels in the second stage. In between the two stages the uncertainty related to the success or failure of product R&D is resolved. For simplicity, Figure 1 illustrates a monopoly case with only one firm, but in general there can be multiple firms competing in the R&D and product markets. The outer agent is a policymaker who seeks to identify the combination of technology-push and demand-pull policies that induces the firms to reach the Nash equilibrium with the highest social welfare.

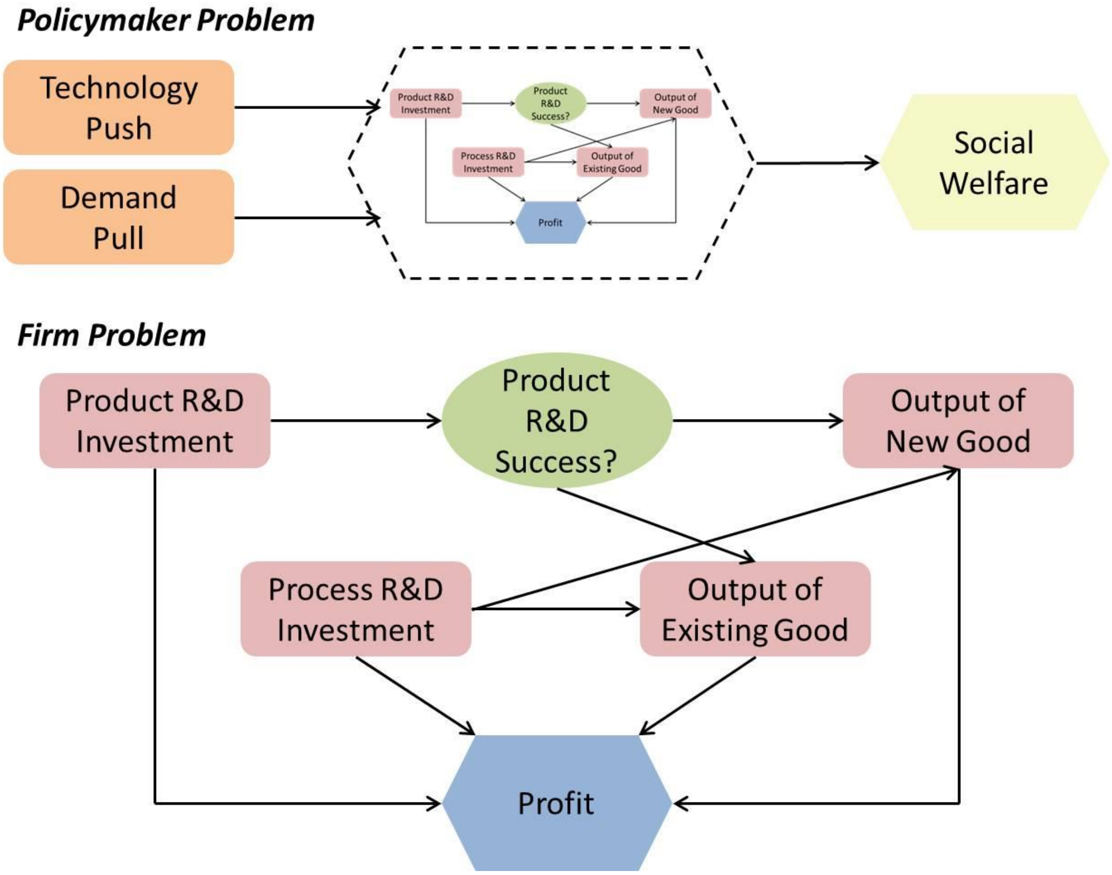

Figure 1. Visual overview of the bilevel modeling framework developed in this analysis.
In this decision diagram rectangles represent decisions, ovals represent uncertainties, and hexagons represent value calculations. The red rectangles are decisions of the profit-maximizing firms and the orange rectangles are decisions of the welfare-maximizing policymaker. For simplicity, only one firm is depicted in the diagram, but in general the modeling framework allows for multiple firms.

2.2 The Firm's Problem

Let F denote the set of firms in the market. Each firm $f \in F$ faces a two-stage stochastic expected profit maximization problem. Here, $\theta \in \Theta$ represents a possible state of the world, a joint realization of all random variables. The firm's problem is Equation 1, where ρ_θ is the probability that state of the world θ occurs and $\Pi_{f\theta}$ is the firm's profit in that state. The decision variables are introduced in the paragraph that follows.

Equation 1

$$maximize\ E\left[\Pi_{f\theta}\right] = \sum_{\theta \in \Theta} \rho_\theta\, \Pi_{f\theta}$$

In the first stage, the firm chooses investment in process R&D (S_f) and investment in product R&D (T_f). Both decision variables are continuous. Investing in process R&D deterministically reduces the cost of producing an existing good. Investing in product R&D raises the probability of successfully developing a new good that can then be produced in addition to the existing good. In the second stage, the firm chooses output levels for the existing good ($X_{f\theta}$) and the new good ($Y_{f\theta}$). The θ subscripts reflect the fact that the firm learns which state of the world has occurred before choosing output levels. The state of the world is defined by the outcomes of all firms' efforts to develop the new good. Each state θ can be thought of as a vector of binary 0 and 1 elements indicating which firms fail and succeed, respectively; as such, the number of states is given by $|\Theta| = 2^{|F|}$. The profit function is Equation 2, which will be explained in greater detail below.

Equation 2

$$\Pi_{f\theta} = \left(P_{X\theta} - C_{Xf}\right)X_{f\theta} + (P_{Y\theta} - c_Y)Y_{f\theta} - S_f - \varphi\left(1 - \eta_{tp}\right)T_f$$

The first term of Equation 2 is revenue minus production cost for the existing good. The second term is analogous, but for the new good. The final two terms are the costs of the firm's R&D investments, where φ is the cost of product R&D relative to process R&D. $P_{X\theta}$ and $P_{Y\theta}$ are the prices of the existing and new goods, respectively, determined according to the inverse demand functions in Equation 3.

Equation 3

$$P_{X\theta} = a_X - b_X \sum_{f \in F} \left(X_{f\theta} + \sigma Y_{f\theta}\right)$$

$$P_{Y\theta} = a_Y + \eta_{dp} - b_Y \sum_{f \in F} \left(Y_{f\theta} + \sigma X_{f\theta}\right)$$

In Equation 3, the parameter σ $(0 \leq \sigma \leq 1)$ captures the substitutability between the existing and new goods. From the firm's perspective, a higher σ means that producing the new good cannibalizes more demand for the existing good.

The cost of producing the existing good (C_{Xf}) is a deterministic, decreasing function of S_f, as specified in Equation 4. The numerator c_{X0} is the initial unit cost before any process R&D is performed, and the exponent α controls the returns to process R&D. The parameter ν_S $(0 \leq \nu_S \leq 1)$ enables spillovers where one firm's process R&D investment reduces production costs for other firms as well as for itself.

Equation 4

$$C_{Xf} = \frac{c_{X0}}{\left(1 + S_f + \nu_S \sum_{g \neq f} S_g\right)^{\alpha}}$$

Note the parameters η_{tp} and η_{dp} that appear in Equation 2 and Equation 3, respectively. These parameters represent technology-push and demand-pull innovation policies. η_{tp} is a subsidy that covers a fraction of product R&D spending. η_{dp} is an absolute dollar amount subsidy for consumer purchases of the new good that shifts its demand curve.

Equation 5 defines the probability of the firm successfully developing the new good (p_f) as a function of its product R&D investment (T_f). This relationship is assumed to take on the logistic (S-shaped) functional form. At first, returns to product R&D are increasing because it requires some critical threshold of resource commitment to establish an R&D program; but ultimately, decreasing returns set in as there is a natural upper limit to the probability of success. The parameter ν_T $(0 \leq \nu_T \leq 1)$ enables spillovers where one firm's product R&D investment increases the probability of product R&D success for other firms as well as for itself.

Equation 5

$$p_f = \frac{\lambda}{1 + exp\left[-\kappa\left(T_f + \nu_T \sum_{g \neq f} T_g - \mu\right)\right]}$$

The logistic function has three important parameters. λ is the saturation level, the maximum probability of success the firm approaches as it continues to invest more in product R&D. κ controls the steepness of the function. If κ is low, then the probability of success increases gradually with product R&D investment; if κ is high, then there is some critical level of R&D investment near which all improvement in probability of success is achieved. μ determines the position of the probability function and is more specifically its inflection point, where probability of success reaches half its maximum value λ. A higher μ thus means the firm must invest more in product R&D to achieve a given probability of success.

Each state of the world θ corresponds to a realization of all product R&D efforts in the market. Let the parameter $\beta_{f\theta} \in \{0,1\}$ indicate whether firm f successfully develops the new good in state of the world θ (1 if success, 0 if failure). The probability that state θ occurs is then given by Equation 6.

Equation 6

$$\rho_\theta = \sum_{f \in F} \beta_{f\theta} p_f + (1 - \beta_{f\theta})(1 - p_f)$$

At this point the formal mathematical description of the firm's problem is complete. Before proceeding, it is useful to clarify the firm's motivations for investing in process and product R&D. Process R&D reduces the cost of producing the existing good, which can lead to greater profits from that market. Product R&D improves the probability of ending up in states of the world where the firm successfully develops the new good, which offer the opportunity to earn more profit.

2.3 Nash Equilibrium

The firms in the market behave as an oligopoly, competing in a game where they simultaneously choose R&D investments and output levels (in all possible states of the world) to maximize expected profit subject to the decisions of all other firms. A Nash equilibrium of the game exists wherever no firm could increase its expected profit by unilaterally modifying its decisions. Since this is a multi-agent game with nonlinear functions, in general it is possible that there are multiple equilibria. Following Dasgupta and Stiglitz (1980), Lin and Saggi (2002), and Suzumura (1992), two simplifying assumptions are made that greatly simplify the solution strategy. First, all firms are assumed to be identical. Second, it is assumed that the outcome of interest is the symmetric Nash equilibrium in which all firms make the same decisions.

With these simplifying assumptions, the equilibrium is determined by solving the expected profit maximization problem of a representative firm. It is implemented in the General Algebraic Modeling System (GAMS) as a nonlinear program (NLP) and computed using the PATHNLP optimization solver. To help ensure that the solution is a global rather than local maximum, the solver is run using a multi-start routine in which the solver is executed many times for different random initial points.

2.4 Welfare

For any combination of technology-push (η_{tp}) and demand-pull (η_{dp}) policies, the firms in the market reach a Nash equilibrium as outlined in the two preceding subsections. This is the inner problem of the bilevel framework. In the outer problem, the policymaker seeks to identify the combination of subsidy levels η_{tp} and η_{dp} that induces the Nash equilibrium which is most desirable for society as a whole. To develop this outer problem, it is first necessary to define a welfare metric that measures the desirability of an equilibrium outcome.

Welfare (W) consists of four components: producer surplus (PS), consumer surplus (CS), subsidy cost (SC), and externality damage (ED). These components are depicted graphically in Figure 2. Consumer surplus is the green triangle below the demand curve and above the equilibrium price. Producer surplus is the red rectangle below the equilibrium price, above the supply curve, and to the left of the equilibrium quantity. The subsidy cost is the amount of money the policymaker must provide to fund the technology policies. The externality damage component captures the possibility that the existing good has an associated negative production externality ω. This possibility is included because the policymaker might be particularly interested in stimulating the development of a new good if the existing one has undesirable characteristics which the market does not internalize (e.g. it pollutes). Mathematical expressions for welfare and its components are provided in Equation 7.

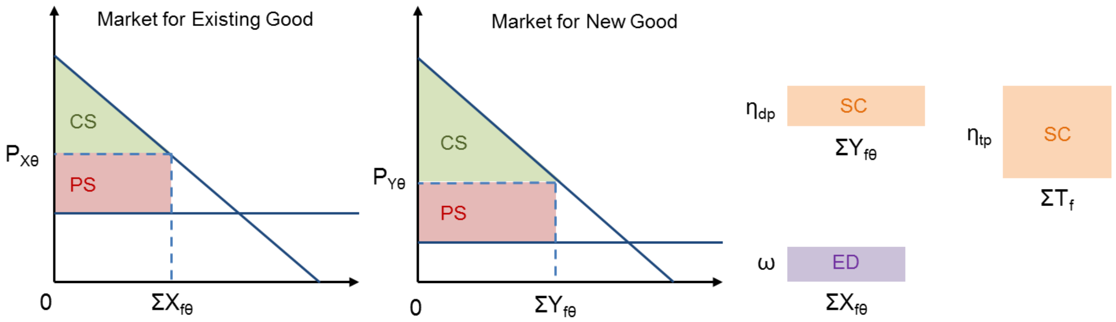

Figure 2. Graphical representation of the four components of welfare.
PS = producer surplus, CS = consumer surplus, SC = subsidy cost, and ED = externality damage.

Equation 7

$$W = PS + CS - SC - ED$$

$$CS = \frac{1}{2}\left[\left(a_X - b_X\sigma\sum_{f\in F} Y_{f\theta}\right)\sum_{f\in F} X_{f\theta} + \left(a_Y + \eta_{dp} - b_Y\sigma\sum_{f\in F} X_{f\theta}\right)\sum_{f\in F} Y_{f\theta}\right]$$

$$PS = (P_{X\theta} - C_X) \sum_{f \in F} X_{f\theta} + (P_{Y\theta} - c_Y) \sum_{f \in F} Y_{f\theta}$$

$$SC = \eta_{dp} \sum_{f \in F} Y_{f\theta} + \eta_{tp} \sum_{f \in F} T_f$$

$$ED = \omega \sum_{f \in F} X_{f\theta}$$

2.5 The Policymaker's Problem

In the outer problem of the bilevel framework, the policymaker seeks to identify the combination of technology-push (η_{tp}) and demand-pull (η_{dp}) subsidies that induces the firms to reach the inner problem Nash equilibrium with the highest expected social welfare. As implemented in this analysis, the policymaker considers a total of 121 discrete policy combinations formed as pairs of 11 discrete choices for the technology-push subsidy and 11 discrete choices for the demand-pull subsidy. Considering a broad range of potential policy combinations rather than formulating the problem solely to identify the welfare-maximizing combination serves to elucidate the full behavior of the model. In any technology policy application there are numerous and deep uncertainties, so it is crucial to consider the full space of potential outcomes to ensure that a small error in estimating parameters does not translate into a significant loss of welfare. Ultimately, this approach allows the policymaker to plot welfare as a function of the considered η_{tp} and η_{dp} values.

3 Numerical Simulations and Results

This section describes two sets of numerical simulations and presents their results. The first set of simulations distinguishes three policy motivations for stimulating innovation and shows how the optimal combination of technology-push and demand-pull policies depends on the primary motivation in a particular application. The simulations in this set are parameterized in a stylized manner but reflect examples from the energy industry. These examples are used to investigate the full behavior of the model. The second set of simulations constitutes a sensitivity analysis performed on the parameters that control the three market failures captured in the modeling framework: incomplete appropriability of R&D, a negative production externality, and imperfect competition. Results from this set of simulations reveal how these market failures influence outcomes such as R&D investment, expected profit, and expected welfare. The findings also indicate how the optimal technology policy, as well as the ease of enhancing welfare through technology policy, varies with the strengths of the market failures.

3.1 Policy Motivations for Stimulating Innovation

The first set of numerical simulations consists of three cases that correspond to three distinct motivations for enacting policies to stimulate innovation. To add some concreteness to these simulations, the three cases are parameterized in a stylized fashion based on examples from the energy industry. The goal of these simulations is not to create parametrically accurate representations of real world policy applications, but rather to use the model to investigate how the best approach to innovation policy depends on the motivation for policy intervention.

Case 1 (Address Negative Externality) depicts a scenario in which product R&D is very expensive but provides an opportunity to replace an existing good that generates a negative production externality. The new good would be slightly more costly than the existing good and a close substitute for it, but would address the externality problem. A concrete example to have in mind could be replacing nuclear fission with nuclear fusion, which would generate little to no waste and be easier to control safely (Bednyagin and Gnansounou, 2011).

Case 2 (Reduce Cost) describes an application in which the primary motivation for investing in product R&D is the promise of developing a new good which is a close substitute for the existing good, but can be produced at lower cost. In this case, product R&D does not require excessive investment and is reasonably likely to succeed, but process R&D for the existing good is still relatively effective and there are no differences between the two goods in terms of production externalities. This case is consistent with the introduction of a new generation of a good, such as replacing crystalline silicon solar PV cells with organic solar PV alternatives that offer the potential for lower costs (Kalowekamo and Baker, 2009).

Case 3 (Create Demand) corresponds to a scenario in which the new good has a large potential market and is not a close substitute for the existing good. In this case the welfare benefit of product R&D would be the creation of a new market with substantial demand that was previously unmet. The electricity generation sector generally does not offer examples that fall under this case because electricity is more or less a homogeneous good and generation options are therefore close substitutes. The demand side, however, offers many examples of this class of technology policy application. While it is difficult to predict the radically different demand-side technologies that will be available in the future, innovations such as autonomous vehicles or personal robots could satisfy demands that are partially unmet by current technologies and have an impact on energy consumption patterns[1].

Table 2 specifies all parameter value assumptions adopted for the three cases.

Table 2. Parameterizations of the cases in the Policy Motivations for Stimulating Innovation numerical simulation set.
Parameters whose values vary across the case studies are shaded in yellow.

Parameter	Description	Case 1 Combat Negative Externality	Case 2 Reduce Cost	Case 3 Create Demand
a_X	Demand intercept, X	20	20	20
b_X	Demand slope, X	1	1	1
a_Y	Demand intercept, Y	20	20	40
b_Y	Demand slope, Y	1	1	1
σ	Substitution parameter	0.8	0.8	0.2
n	Number of firms	1	3	3
ν_S	Process R&D spillover strength	0.3	0.3	0.3
ν_T	Product R&D spillover strength	0.3	0.3	0.3
c_{X0}	Initial unit cost, X	5	5	5
c_Y	Unit cost, Y	6	3	6
α	Effectiveness of process R&D	0.05	0.2	0.2
φ	Relative cost of product R&D	10	5	10
κ	Probability function steepness	0.4	0.4	0.4
λ	Probability function maximum	0.3	0.8	0.3
μ	Probability function inflection point	10	10	10
ω	Negative production externality, X	2	0	0

Results for Case 1 (Address Negative Externality) are illustrated in Figure 3. In all four plots, each grid cell corresponds to a combination of product R&D subsidy η_{tp} (horizontal axis)

---

[1] An industry with many applications that would fall under Case 3 (Create Demand) is the pharmaceutical industry, where any new drug that treats a disease or condition that was previously untreated would qualify.

and new good price subsidy η_{dp} (vertical axis). The heat plots in Figure 3c and Figure 3d show product and process R&D investment per firm, respectively. In the absence of sufficiently strong technology policy, the firms entirely forego product R&D. The sudden switch to substantial investment in product R&D once the subsidies are high enough reflects the initially increasing returns to product R&D in the probability of success function. Below some critical level of investment, it is not sensible to undertake any product R&D. Where the policies suddenly induce product R&D investment, there is a clear decline in process R&D investment. The model endogenously captures the tradeoff between product and process R&D, and based on these results the two activities are clearly substitutes in this case. Unsurprisingly, expected profit increases with both subsidy levels (Figure 3b). The most important result from a policy standpoint is the expected welfare heat plot in Figure 3a. It suggests that it is difficult to achieve a better outcome than the no-policy baseline in Case 1 (Address Negative Externality). Welfare increases for only a very narrow range of policy combinations that primarily utilize the technology-push channel to just barely provide a strong enough incentive for the firms to invest in product R&D. Beyond this, stronger policy interventions induce slightly more product R&D but result in lower welfare than the no-policy baseline. The results of Figure 3 suggest that it can be very difficult to improve social welfare by relying on technology policy to combat a negative production externality.

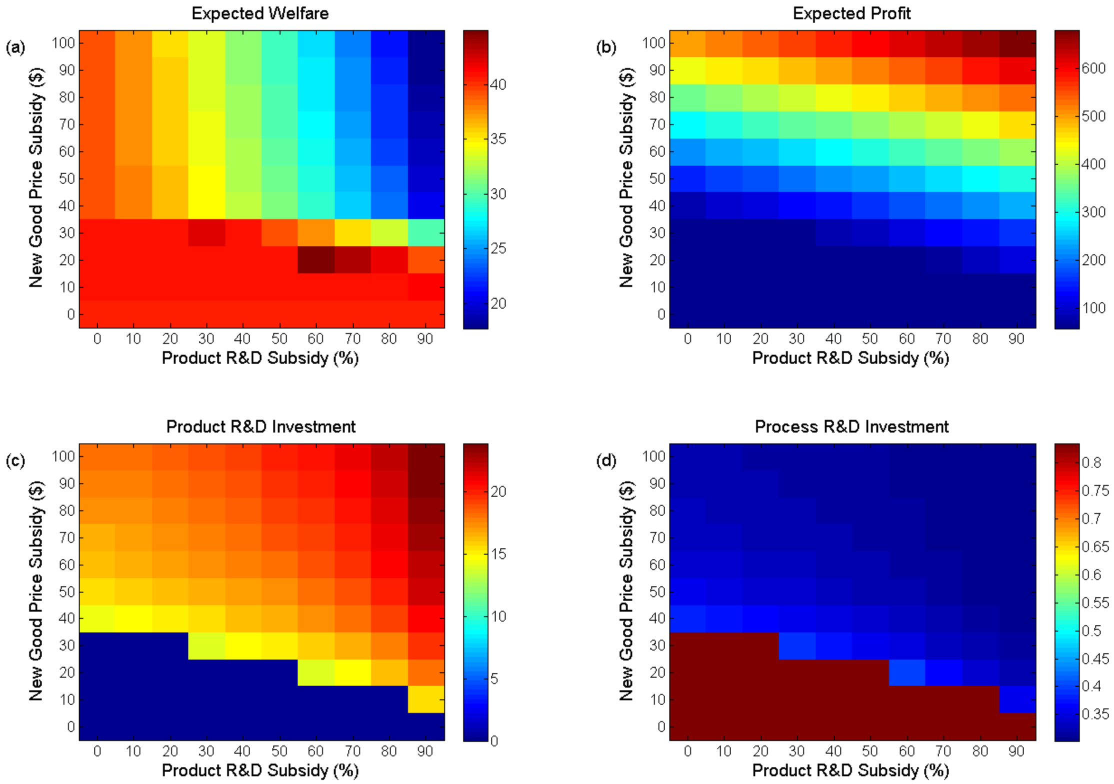

Figure 3. Results for Case 1 (Address Negative Externality).
In the absence of technology policy, the private sector does not perform product R&D. When sufficient incentives are in place to stimulate product R&D investment, process R&D expenditure drops, implying that the two types of R&D are substitutes. Expected welfare is higher than that in the no-policy baseline under only a few combinations of technology-push and demand-pull interventions, and not by a significant amount. These combinations emphasize technology-push more than demand-pull.

The Case 2 (Reduce Cost) results depicted in Figure 4 bear some similarities to the Case 1 (Address Negative Externality) results, but also some important differences. As in the prior case, the firms do not invest in product R&D in the absence of technology policy. Once the policy incentives are sufficiently strong, there is a sudden increase in product R&D expenditure accompanied by a reduction in process R&D expenditure. Again, the two types of R&D are substitutes. The policy threshold at which the private sector begins to invest in product R&D is lower in this case, and almost every policy combination featuring a positive price subsidy induces some product R&D investment (with only one exception). The expected welfare results in Figure 4a are the most revealing for clarifying how Case 2 (Reduce Cost) differs from Case 1 (Address Negative Externality). In this case, the optimal technology-push and demand-pull combination leads to a significant social welfare improvement relative to the no-policy baseline, and the number of policy combinations that enhance welfare is far greater. A close examination of Figure 4a suggests that there is little rationale for enacting technology-push policies if the primary motivation for innovating is to reduce cost. Expected welfare generally declines from left to right in the figure and it is possible to achieve substantial welfare gains by relying solely on a demand-pull intervention. Nevertheless, there could be a useful role for technology-push policy if there is some upper limit on the demand-pull policy. For example, if the new good price subsidy is constrained to a maximum of $10, expected welfare can be increased by providing a 10% subsidy for product R&D.

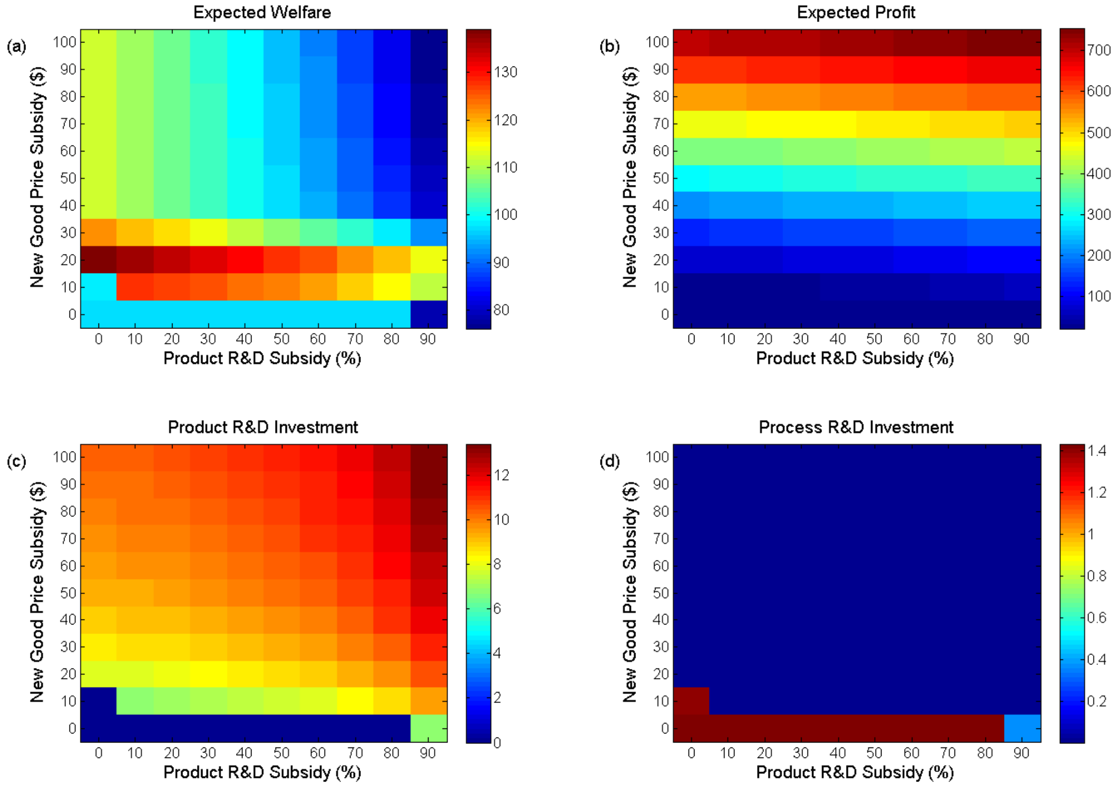

Figure 4. Results for Case 2 (Reduce Cost).
The technology policy threshold at which the private sector begins to invest in product R&D is lower in this case than in Case 1 (Combat Negative Externality). Again, product and process R&D are substitutes. Compared to the prior case, there are far more policy combinations that enhance welfare, and some of them lead to significantly higher welfare than the no-policy baseline. If the primary motivation to stimulate innovation is to reduce cost, policy intervention should emphasize demand-pull more than technology-push.

Case 3 (Create Demand) results are plotted in Figure 5. It is possible to induce product R&D investment with fairly weak policies. Once again, process and product R&D are evidently substitutes, although in this case the shift from process to product R&D occurs more gradually as the policies become stronger compared to the two previous cases. This is a sensible outcome given the lower value of the substitution parameter σ in this case. The expected welfare results in Figure 5a demonstrate that a very wide range of policy combinations lead to significantly higher expected welfare than the no-policy baseline. From the perspective of the policymaker, improving social welfare through technology policy appears to be an easier task if the application involves innovation that would create new markets with previously unmet demand. Given that the vertical pattern in Figure 5a is more pronounced than the horizontal pattern, expected welfare depends more on the demand-pull policy than on the technology-push policy. The highest expected welfare results are achieved under policy combinations that emphasize demand-pull more than technology-push, but considering the robustness of welfare improvements across many different policy combinations, the exact balance between the two is less important in this case.

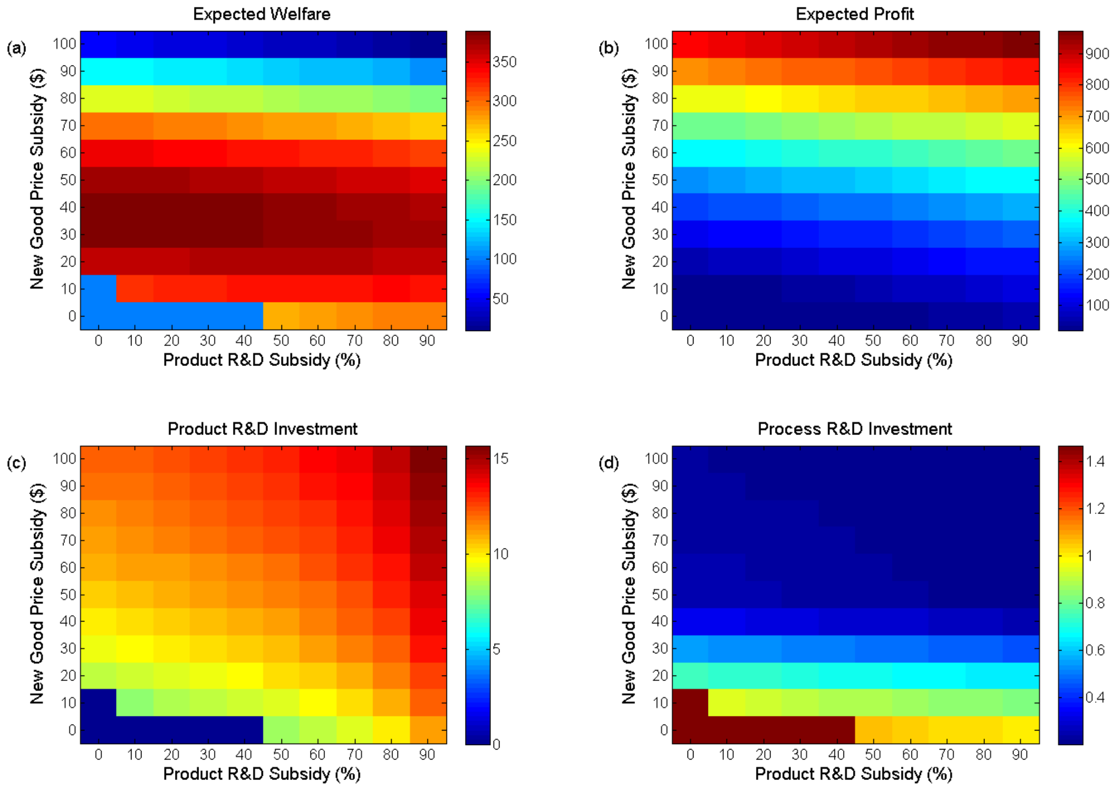

Figure 5. Results for Case 3 (Create Demand).
Once again, product and process R&D are substitutes, although the switch from process to product R&D occurs more gradually than in the previous cases. A very wide range of policy combinations enhance social welfare significantly relative to the no-policy baseline, suggesting that it is easier to improve welfare through technology policy alone if the application involves innovation that would create large new markets to satisfy previously unmet demand. The exact balance between technology-push and demand-pull is not critical, but expected welfare is highest under combinations that emphasize demand-pull.

Taken together, the results of the numerical simulations in this set suggest that enhancing social welfare through technology policy is very difficult if the innovation would primarily address a negative externality, less difficult if the innovation would primarily reduce cost, and substantially easier if the innovation would primarily create new market demand. The policy implications of these findings for energy technology innovation are clear. Rather than focus on technologies that provide the same service as existing technologies but at a lower cost or with less of a negative externality, it would be prudent to apply technology policies to innovations that could also create new market demands or augment existing ones. Perhaps the balance of R&D expenditures should be shifted away from large-scale power generation projects in the direction of novel end-use goods and distributed generation technologies capable of providing power in places that currently lack reliable electric grids (e.g. developing countries).

3.2 Market Failure Sensitivity Analysis

The bilevel modeling framework of this analysis captures three critical market failures: imperfect appropriability of R&D, a negative production externality associated with the existing good, and imperfect competition. The numerical simulations of this subsection shed light on how these market failures influence outcomes such as product R&D expenditures and expected welfare, and affect the optimal combination of technology-push and demand-pull policy interventions. Beginning with a reference parameterization, sensitivity analysis is performed on the model parameters that control the strengths of the three market failures. The incomplete appropriability market failure is related to the values of the spillover parameters ν_S and ν_T; if their values are higher, R&D is less appropriable and the innovation market failure is stronger. The negative production externality associated with the existing good is simply the parameter ω. Imperfect competition is tied to the number of firms n, with more of a market failure if the market consists of fewer firms. This approach is clarified in Table 3.

Table 3. Reference parameter values and minimum and maximum values in the sensitivity analysis on parameters that control the three market failures.
The process and product R&D spillover strengths (ν_S and ν_T) control the imperfect appropriability market failure. The negative production externality associated with the existing good (ω) controls the negative production externality market failure. The number of firms (n) controls the imperfect competition market failure.

Parameter	**Description**	**Reference Value**	**Minimum Value**	**Maximum Value**
a_X	Demand intercept, X	20		
b_X	Demand slope, X	1		
a_Y	Demand intercept, Y	20		
b_Y	Demand slope, Y	1		
σ	Substitution parameter	0.5		
n	Number of firms	2	1	4
ν_S	Process R&D spillover strength	0.3	0	1
ν_T	Product R&D spillover strength	0.3	0	1
c_{X0}	Initial unit cost, X	5		
c_Y	Unit cost, Y	4		
α	Effectiveness of process R&D	0.1		
φ	Relative cost of product R&D	10		
κ	Probability function steepness	0.4		
λ	Probability function maximum	0.5		
μ	Probability function inflection point	10		
ω	Negative production externality, X	4	0	20

Sensitivity analysis results for product R&D expenditures under optimal policy interventions[2] are reported in Figure 6. As shown in Figure 6a, stronger spillovers reduce investment in product R&D, consistent with the logic of the incomplete appropriability market failure. The marginal reduction in product R&D decreases with the spillover strength. As the negative production externality associated with the existing good rises in Figure 6b, the product R&D investment eventually increases. This occurs because the optimal policy becomes more forceful to induce a stronger transition from the existing good to the new good if the existing good generates a larger negative externality. Since the number of firms in the market varies in the imperfect competition sensitivity analysis, the solid and dotted lines in Figure 6c respectively correspond to total product R&D and per-firm product R&D. In line with the traditional Schumpeterian argument, which holds that greater competition erodes post-innovation rents and reduces incentives to invest in R&D, per-firm product R&D declines as competition intensifies. However, even though per-firm product R&D declines, this is more than offset by the presence of more firms in the industry doing R&D. As a result, total product R&D expenditures in the industry actually increase with the number of firms.

---

[2] Without technology policy, product R&D expenditures are always $0.

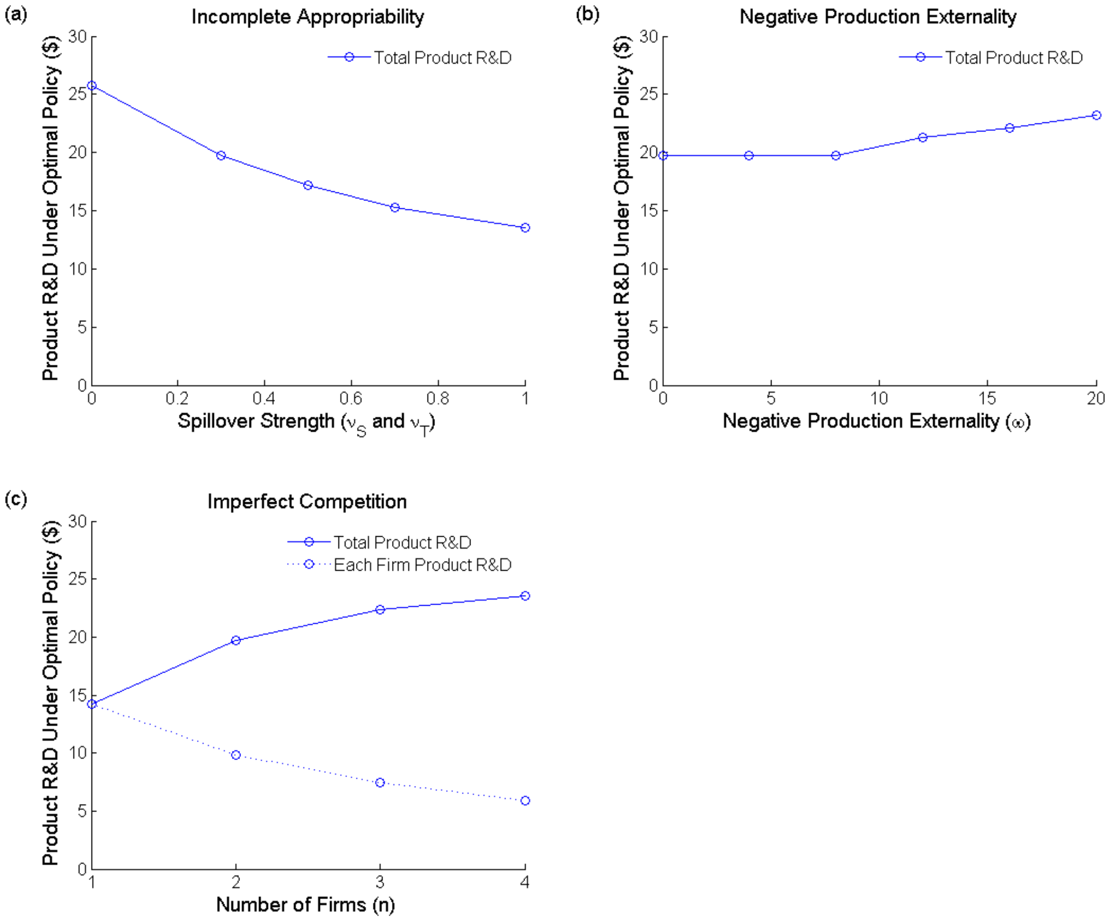

Figure 6. Sensitivity analysis results for product R&D investment.

Sensitivity analysis results for expected profit per firm are plotted in Figure 7. In the absence of technology policy, the only market failure parameter that has a significant impact on profit is the number of firms. As one would expect, per-firm profit is highest in the monopoly case and declines with greater competition following a convex profile (Figure 7c). The negative production externality does not factor into the firms' profit maximization problem and therefore has no effect on profit. Spillovers could theoretically influence profits but their effect is barely discernible in Figure 7a. When optimal technology policy intervention is allowed, firms generally earn greater profits because they benefit from the subsidies. This increase in profit due to the policy is larger with strong spillovers (Figure 7a) or a monopoly (Figure 7c). Firms profit most from the introduction of optimal technology policy if the negative externality associated with the existing good is substantial (Figure 7b). In such circumstances the optimal policy features a generous subsidy for product R&D to intensify the transition to the new good. As a result, the firms earn greater profits.

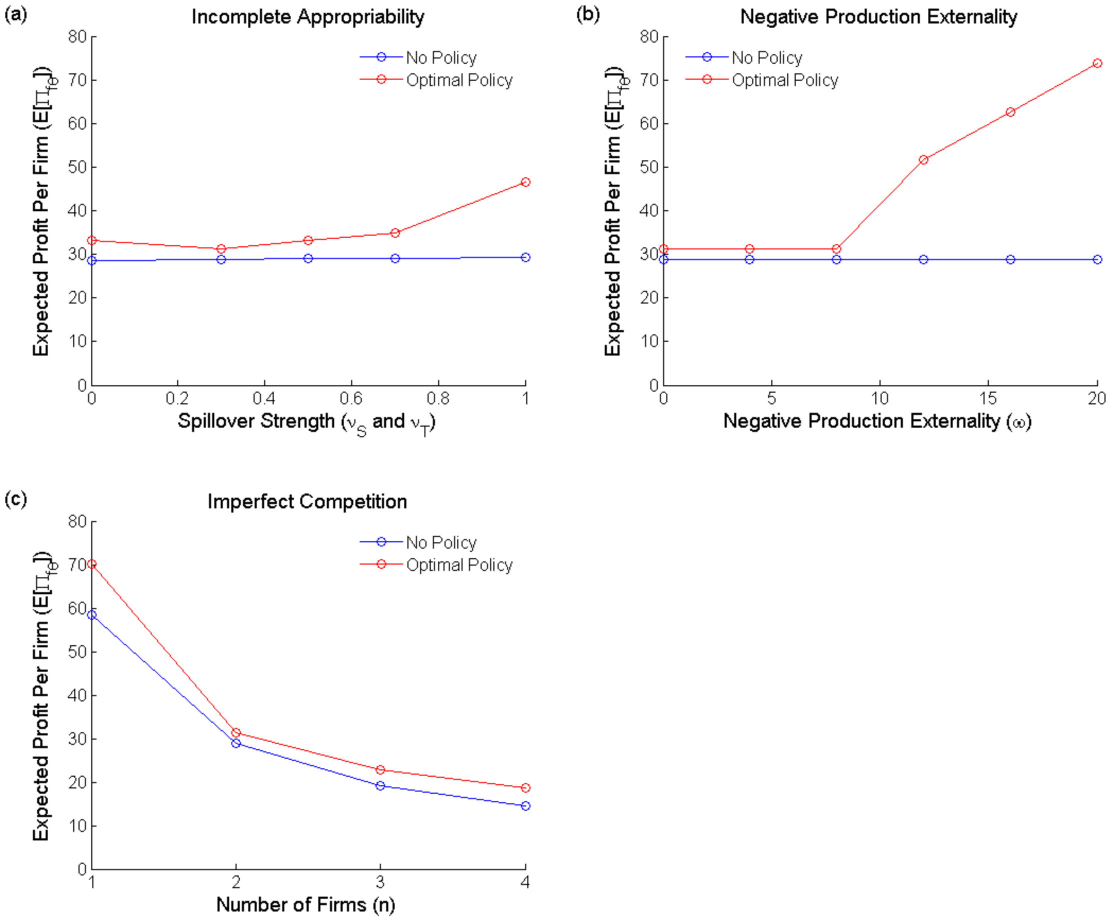

Figure 7. Sensitivity analysis results for expected profit per firm.

Figure 8 illustrates how expected welfare varies with the sensitivity parameters in the no-policy case and in the optimal policy case. Based on the substantial gaps between the no-policy and optimal policy lines in all three subfigures, the results clearly indicate that technology policy intervention can generate significant welfare gains across a wide range of parameter settings. Figure 8a shows that expected welfare rises with the spillover strength in both the no-policy and optimal policy cases. Although stronger R&D spillovers reduce total expenditures on product R&D (see Figure 6a), they also make each dollar spent on product R&D more productive for the industry, and the net effect is that stronger spillovers are a positive from a social welfare point of view. While these spillovers are the mechanism behind the innovation market failure, they can evidently be desirable for society as a whole. The results in Figure 8b show the intuitive result that expected welfare declines with the negative production externality in both the no-policy and optimal policy cases. The interesting feature of the results is that expected welfare declines far less steeply with the negative production externality in the optimal policy case, suggesting that the gains from technology policy can be large if the targeted externality is severe. As the externality becomes larger, the optimal policy reduces the expected welfare loss by inducing a stronger transition from the existing good to the new one. As previously discussed in subsection 1.2, the economic literature contains opposing claims about the relationships among competition, innovation, and welfare. The findings presented in Figure 8c are striking in that expected welfare decreases with greater competition in the no-policy case, but increases with greater competition in the optimal policy case. In other words, the Schumpeterian argument that competition erodes post-innovation rents and leads to less than the socially desirable level of innovation appears to hold in the absence of technology policy. But if optimal technology-push and demand-pull policies can be imposed, it is possible to achieve better welfare outcomes with more intense competition. Observing the gap between the no-policy and optimal policy expected welfare lines in Figure 8c, technology policy is a more powerful means of improving social welfare if the industry in question is competitive.

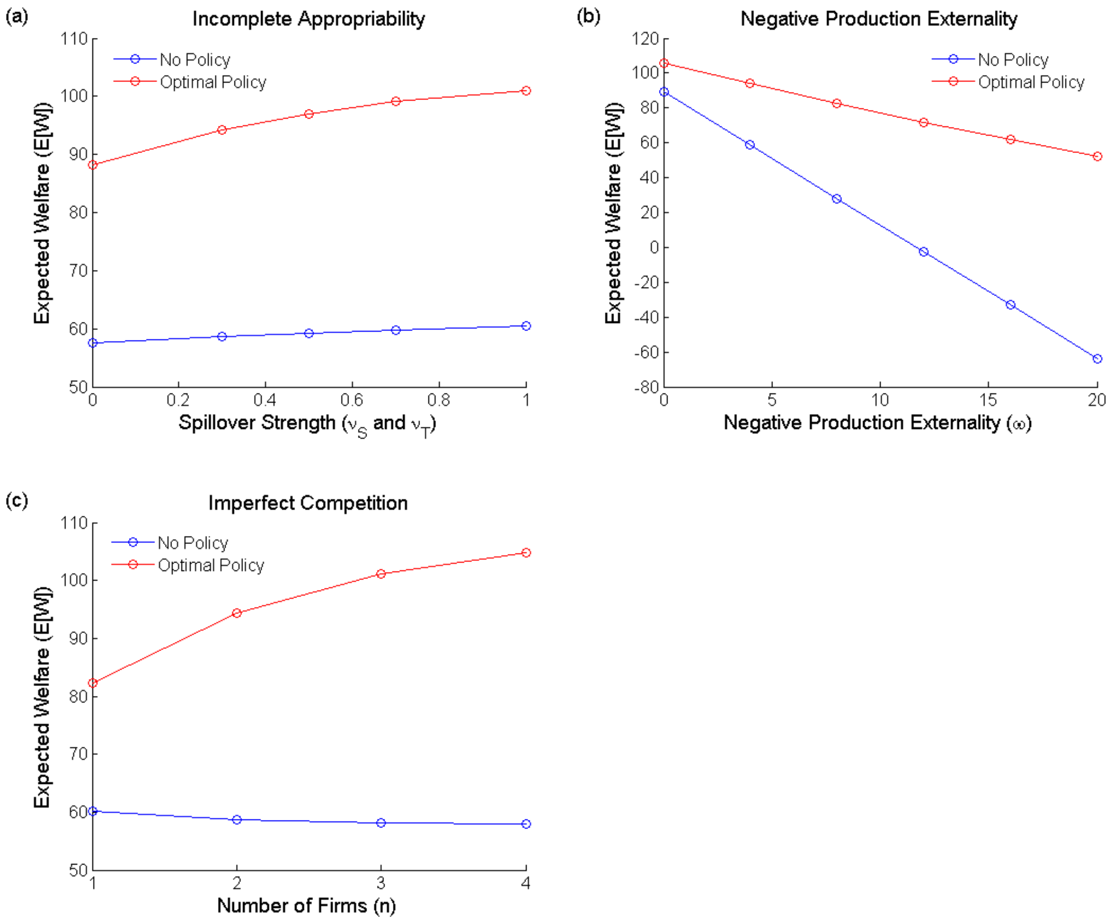

Figure 8. Sensitivity analysis results for expected welfare.

In revealing how the optimal technology policy combination varies with the three market failures, it should first be noted that the optimal demand-pull policy (price subsidy for the new good) remains the same across all parameterizations considered in this sensitivity analysis. This is quite striking and suggests that the optimal demand-pull policy is largely determined by parameters other than those varied in this sensitivity analysis. Evidence for this interpretation can be seen in Figure 5, in which the optimal policy intervention for Case 3 (Create Demand) features a higher new good price subsidy than the level which is always identified as optimal in this sensitivity analysis ($20). In other words, the optimal demand-pull policy in this sensitivity analysis is more a product of the reference parameterization – particularly its fixed assumption about the demand for the new good – than the parameters which are varied.

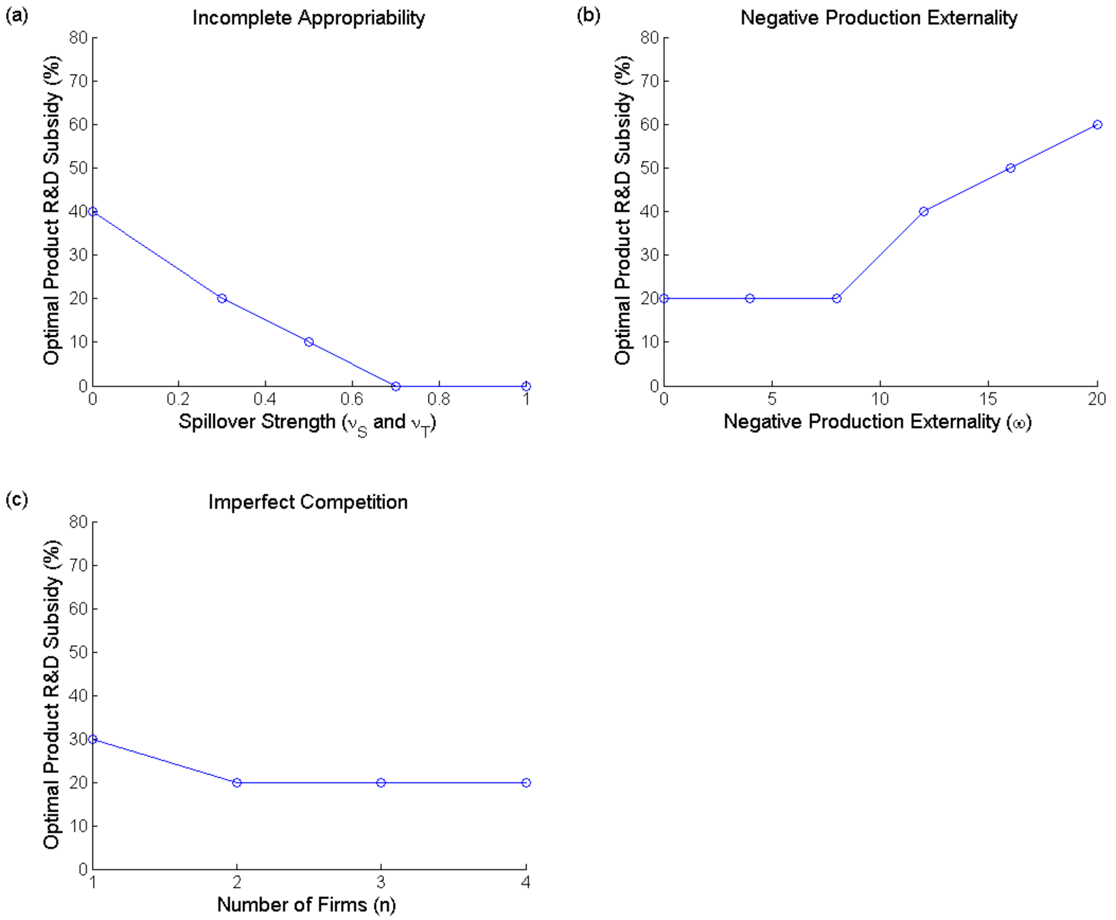

Figure 9. Sensitivity analysis results for the optimal technology-push policy.

Unlike the optimal demand-pull policy, the optimal technology-push policy (product R&D subsidy) is sensitive to the parameters that control the three market failures. Figure 9 reveals that the optimal product R&D subsidy generally decreases with the strength of R&D spillovers (Figure 9a), increases with the negative production externality associated with the existing good (Figure 9b), and decreases with the number of firms in the market (Figure 9c). The first of these findings is not obvious ex ante because stronger R&D spillovers should imply a more severe innovation market failure, and total product R&D investment does, in fact, decline (see Figure 6a). But with stronger spillovers the benefit of each dollar spent on product R&D is greater in terms of its effect on the probability of product R&D success. This latter effect evidently dominates and stronger spillovers mean that the optimal policy intervention features a weaker technology-push policy. The negative production externality finding is intuitive, since a more severe negative externality associated with the existing good calls for a stronger policy signal to intensify the transition from the existing good to the new good. Monopolistic competition leads to a higher optimal product R&D subsidy than market configurations with multiple firms. Even though the monopolist invests more in product R&D on a per-firm basis (see Figure 6c), in the monopoly case there is only one firm to invest in R&D and thus no spillovers to benefit from. A stronger technology-push policy is needed to raise the monopolist's product R&D expenditure closer to the total amount that would be spent by an industry with more participants performing R&D. From a policy perspective, the distinction between pure monopoly and any higher degree of competition appears to be more meaningful than differences at higher levels of competition.

4 Conclusions

The bilevel modeling framework developed in this study has been designed to determine the optimal combination of technology-push and demand-pull interventions for a particular technology policy application. The model goes beyond previous studies of strategic innovation in an oligopoly context by incorporating uncertainty in R&D outcomes and by including an outer layer of decision variables that represent the optimal policy intervention. In a relatively compact framework, the model captures three market failures which can have an important influence on innovation and optimal technology policy: incomplete appropriability of R&D, a negative production externality associated with the existing good, and imperfect competition.

The first set of numerical simulations used three parameter cases corresponding to three distinct policy motivations for stimulating innovating to explore the behavior of the model. These motivations are to address a negative externality (Case 1), to reduce cost (Case 2), and to create demand (Case 3). The results reveal that the firms do not perform risky product R&D in the absence of technology policy. Process and product R&D are substitutes in the model, consistent with the productivity dilemma concept (Abernathy, 1978) and the tradeoff between exploitation and exploration. The optimal combination of technology-push and demand-pull measures, as well as the ease of enhancing welfare through technology policy, varies across the cases. If the innovation would primarily serve to address a negative externality, the optimal policy includes a strong technology-push subsidy but it is quite difficult to enhance welfare through technology policy relative to the no-policy baseline. If the innovation would primarily reduce cost, it is slightly easier to achieve higher welfare through technology policy and the optimal intervention emphasizes demand-pull. If the innovation would primarily create demand, a wide range of technology policy combinations can substantially increase welfare and these tend to rely mostly on demand-pull.

The second set of simulations was a sensitivity analysis performed on the model parameters that control the three market failures. The results show that firms perform less product R&D under stronger spillovers, but that expected welfare is higher because the R&D that is performed more powerfully translates into a higher probability of product R&D success. Technology policy is an effective tool for mitigating the welfare loss stemming from a severe negative production externality by intensifying the transition from the dirty existing good to a cleaner new good. Firms earn lower profits and perform less product R&D on a per-firm basis as competition becomes stronger, consistent with the traditional Schumpeterian argument (Schumpeter, 1943). However, total product R&D in the industry rises with the number of firms. Interestingly, expected welfare decreases with competition in the no-policy case, but increases with competition in the optimal policy case. Claims that some amount of market power is beneficial for innovation therefore seem to hold less weight if technology policies can be well designed and effectively imposed. While the optimal demand-pull policy is not sensitive to the market failure sensitivity analysis parameters, the optimal technology-push policy exhibits sensitivity. The optimal product R&D subsidy decreases with spillover strength, increases with the negative production externality, and is higher for a monopoly than for more competitive market structures.

The analysis presented in this article has offered a methodology for assessing strategic innovation and optimal technology policy, and has produced a number of valuable economic and policy insights. In the future, practical application of this methodological approach will require integration into broader modeling frameworks and estimation of parameters for each particular application. For example, integrating strategic energy R&D and technology policy into an IAM could be accomplished by adding the necessary decision variables to a multi-agent framework such as a computable general equilibrium IAM, or through an iterative model coupling approach like that employed by Leibowicz (2015). The practical value of this model for optimal technology policy design would be greatly enhanced by more extensive empirical research to estimate or at least constrain the values of critical parameters. For example, there is a pressing need for empirical research to deepen our understanding of R&D spillovers and the relationship between R&D investment and the likelihoods of various innovation outcomes.

5 Acknowledgments

The author began this study at the International Institute for Applied Systems Analysis (IIASA), where he was supported by the institute's Peccei Award. Arnulf Grubler contributed valuably to the conceptual development of this project. Volker Krey provided helpful advice on implementing and solving the model. Funding for completion of this research was obtained through the Department of Energy, Office of Science PIAMDDI grant (DE-SC005171) awarded to the Stanford University Energy Modeling Forum. While at Stanford, the author received useful suggestions from Lawrence Goulder, Charles Kolstad, James Sweeney, John Weyant, and participants in a Stanford Environmental and Energy Policy Analysis Center (SEEPAC) seminar where he presented an earlier version of this study.

6 References

Abernathy, W.J., 1978. The Productivity Dilemma: Roadblock to Innovation in the Automobile Industry. Johns Hopkins University Press, Baltimore.

Aghion, P., Bloom, N., Blundell, R., Griffith, R., Howitt, P., 2005. Competition and innovation: an inverted-u relationship. Quarterly Journal of Economics 120, 701–728.

Arrow, K., 1962. Economic welfare and the allocation of resources for invention, in: The Rate and Direction of Inventive Activity: Economic and Social Factors. Princeton University Press, Princeton, NJ, pp. 609–626.

Arthur, W.B., 2007. The structure of invention. Research Policy 36, 274–287.

Bednyagin, D., Gnansounou, E., 2011. Real options valuation of fusion energy R&D programme. Energy Policy 39, 116–130.

Bernstein, J.I., Nadiri, M.I., 1988. Interindustry R&D spillovers, rates of return, and production in high-tech industries. American Economic Review 78, 429–434.

Blundell, R., Griffith, R., Reenen, J. Van, 1995. Dynamic count data models of technological innovation. Economic Journal 105, 333–344.

Bosetti, V., Carraro, C., Massetti, E., Sgobbi, A., Tavoni, M., 2009. Optimal energy investment and R&D strategies to stabilize atmospheric greenhouse gas concentrations. Resource and Energy Economics 31, 123–137.

Bush, V., 1945. Science the endless frontier: A report to the president on a program for postwar scientific research. Washington, D.C.

Chandy, R.K., Tellis, G.J., 2000. The incumbent's curse? Incumbency, size, and radical product innovation. Journal of Marketing 64, 1–17.

D'Aspremont, C., Jacquemin, A., 1988. Cooperative and noncooperative R&D in duopoly with spillovers. American Economic Review 78, 1133–1137.

Dasgupta, P., Stiglitz, J., 1980. Industrial structure and the nature of innovative activity. Economic Journal 90, 266–293.

Gallagher, K.S., Grübler, A., Kuhl, L., Nemet, G., Wilson, C., 2012. The energy technology innovation system. Annual Review of Environment and Resources 37, 137–162.

Grubler, A., 1998. Technology and Global Change. Cambridge University Press, Cambridge, UK.

Jaffe, A.B., 1986. Technological opportunity and spillovers of R&D: Evidence from firms' patents, profits, and market value. American Economic Review 76, 984–1001.

Jones, C.I., Williams, J.C., 1998. Measuring the social return to R&D. Quarterly Journal of Economics 113, 1119–1135.

Kalowekamo, J., Baker, E., 2009. Estimating the manufacturing cost of purely organic solar cells. Solar Energy 83, 1224–1231.

Kamien, M.I., Schwartz, N.L., 1976. On the degree of rivalry for maximum innovative activity. Quarterly Journal of Economics 90, 245–260.

Kleinknecht, A., Verspagen, B., 1990. Demand and innovation: Schmookler re-examined. Research Policy 19, 387–394.

Kypreos, S., 2007. A MERGE model with endogenous technological change and the cost of carbon stabilization. Energy Policy 35, 5327–5336.

Leibowicz, B.D., 2015. Growth and competition in renewable energy industries: Insights from an integrated assessment model with strategic firms. Energy Economics 52, 13–25.

Levin, R., Reiss, P., 1988. Cost-reducing and demand-creating R&D with spillovers. RAND Journal of Economics 19, 538–556.

Levinthal, D.A., March, J.G., 1993. The myopia of learning. Strategic Management Journal 14, 95–112.

Lin, P., Saggi, K., 2002. Product differentiation, process R&D, and the nature of market competition. European Economic Review 46, 201–211.

Loiter, J.M., Norberg-Bohm, V., 1999. Technology policy and renewable energy: public roles in the development of new energy technologies. Energy Policy 27, 85–97.

Mansfield, E., Rapoport, J., Romeo, A., Wagner, S., Beardsley, G., 1977. Social and private rates of return from industrial innovations. Quarterly Journal of Economics 91, 221–240.

Nemet, G.F., 2009. Demand-pull, technology-push, and government-led incentives for non-incremental technical change. Research Policy 38, 700–709.

Nemet, G.F., Kammen, D.M., 2007. U.S. energy research and development: Declining investment, increasing need, and the feasibility of expansion. Energy Policy 35, 746–755.

Norberg-Bohm, V., 2000. Creating incentives for environmentally enhancing technological change: Lessons from 30 years of U.S. energy technology policy. Technological Forecasting and Social Change 65, 125–148.

Pavitt, K., 1984. Sectoral patterns of technical change: Towards a taxonomy and a theory. Research Policy 13, 343–373.

Scherer, F.M., 1982. Inter-industry technology flows and productivity growth. Review of Economics and Statistics 64, 627–634.

Scherer, F.M., 1993. Pricing, profits, and technological progress in the pharmaceutical industry. Journal of Economic Perspectives 7, 97–115.

Schmookler, J., 1966. Invention and Economic Growth. Harvard University Press, Cambridge, MA.

Schumpeter, J.A., 1943. Capitalism, Socialism, and Democracy. Allen and Unwin, London.

Suurs, R.A.A., Hekkert, M.P., 2009. Competition between first and second generation technologies: Lessons from the formation of a biofuels innovation system in the Netherlands. Energy 34, 669–679.

Suzumura, K., 1992. Cooperative and noncooperative R&D in an oligopoly with spillovers. American Economic Review 82, 1307–1320.

Taylor, M., 2008. Beyond technology-push and demand-pull: Lessons from California’s solar policy. Energy Economics 30, 2829–2854.

Watanabe, C., Wakabayashi, K., Miyazawa, T., 2000. Industrial dynamism and the creation of a virtuous cycle between R&D, market growth, and price reduction: The case of photovoltaic power (PV) generation development in Japan. Technovation 20, 299–312.

Wiser, R., Bolinger, M., Cappers, P., Margolis, R., 2007. Analyzing historical cost trends in California's market for customer-sited photovoltaics. Progress in Photovoltaics: Research and Applications 15, 69–85.