Research and Development, Profits and Firm Value: A Structural Estimation

Missaka Warusawitharana*

Board of Governors of the Federal Reserve System

Abstract

The production function approach employed in the literature yields implausible estimates of the rate of return to R&D. This study presents an alternate structural approach. The underlying dynamic model assumes that firms invest in R&D in order to generate innovations that arrive stochastically. Optimality conditions pin down the rate of return. The estimated rate of return to R&D ranges from 2% to 5%, much lower than the values reported in the literature. The results also shed light on the obsolescence rate of R&D stocks, the impact of an innovation on profits, and the market value of the aggregate R&D stock.

December 10, 2009

*I thank Marco Cagetti, Taeyoung Doh, John Fernald, Simon Gilchrist, João Gomes, Jonathan Heathcote, Peter Howitt, Michael Palumbo, and seminar participants at the 2008 Econometric Society Summer meetings, 2008 SED meetings, 2009 Federal Reserve System Applied Microeconomics Conference, Bureau of Economic Analysis, University of Rochester, Wharton Macro Finance Lunch, and Federal Reserve Board for valuable comments and suggestions. The views expressed in this paper are mine and do not reflect the views of the Board of Governors of the Federal Reserve System or its staff. Contact: Division of Research and Statistics, Board of Governors of the Federal Reserve System, Mail Stop 97, 20th and Constitution Ave. NW, Washington, D.C., 20551. m1mnw00@frb.gov, (202) 452-3461.

1 Introduction

Estimating the rate of return to research and development (R&D) expenditures has value from an economic growth perspective (Romer (1990)) and a macro fluctuations perspective (Comin and Gertler (2006)).[1] A substantial literature employs the production function-based approach of Griliches (1979) and finds private rates of return to R&D in the range of 10 to 50%.[2] While the literature acknowledges these estimates may be implausibly high and discusses flaws in the production function approach (Griliches (2000, Chapter 4)), it continues to be the standard.

I argue that the production function approach has some major flaws. First, the estimation does not impose optimality, which would imply that the discounted value of the rate of return adjusted for the obsolescence of R&D would be less than or equal to the marginal cost. Many of the rate of return estimates in the literature violate this condition, unless one assumes a very high discount rate. Second, analysis reveals that the rates of return obtained using this approach are highly sensitive to the obsolescence rate of R&D stocks, which is typically fixed at 15% following Griliches and Mairesse (1984). However, we know little about the obsolescence rate of R&D stocks. Last, this approach ignores uncertainty in the outcome from R&D, a key feature of models with innovation. This suggests that estimates from production function methods may be unreliable due to model misspecification.

This study tackles these issues by estimating a dynamic optimizing model where firms invest in both physical capital and R&D. The model incorporates uncertainty by assuming that firms invest in R&D in order to increase the probability of generating innovations each period, where an innovation leads to an increase in profitability.[3] As profitability increases persist for many periods, firms optimally increase their capital stock following an innovation. The value of the firm includes both the value of the capital stock that generates output and profits, and the value of the R&D stock that generates future innovations. Solving the model yields explicit characterizations of the optimal R&D policy and the rate of return in terms of the underlying structural parameters.

I estimate the model using simulated method of moments.[4] The estimates imply a rate of return to R&D in the range of 2 to 5% for the entire sample. These estimates are much lower than the

[1]To clarify the terminology, I use 'rate of return' to describe the increase in next period payoff for a unit increase in current period R&D. This is also referred to in the literature as the 'gross rate of return'.

[2]See Table 3 of Okubo, Robbins, Moylan, Selzer, Schultz, and Mataloni (2006) for a survey of the literature.

[3]This setup follows the basic approach of the endogenous growth literature. See Romer (1990), Grossman and Helpman (1991), and Aghion and Howitt (1992).

[4]Other papers that use this method include: Cooper and Ejarque (2003); Hennessy and Whited (2005); Cooper and Haltiwanger (2006); Hennessy and Whited (2007); Lentz and Mortensen (2008); Eberly, Rebelo, and Vincent (2008); Akcigit (2009); and Bloom (2009).

values found in the literature. Optimality implies that the obsolescence rate of R&D limits the maximum rate of return. Furthermore, persistence in profitability reduces the rate of return as innovations benefit the firm not only in the next period, as measured by the rate of return, but also in subsequent periods. The rate of return measure also ignores the increase in firm value that arises from changes in the capital stock following an innovation. Thus, the model simultaneously generates a low rate of return and a high level of R&D investment as the rate of return focuses only on a narrow portion of the total benefit from R&D investment. Estimating the model for five selected R&D-intensive industries leads to similar estimates, with a maximum rate of return of 5% for the chips industry.

The structural estimation provides estimates of the increase in profitability from an innovation, the rate of innovation, and the obsolescence rate. The results indicate that an innovation leads to profitability increases of 15% to 20% for all R&D firms. This is equivalent to a permanent increase in profitability of 4% to 6%. The mean fraction of firms innovating each year ranges from 0.58 to 0.71, highlighting the uncertainty that firms face when undertaking R&D projects. This suggests that estimation methods that ignore uncertain outcomes from R&D may yield biased estimates due to model misspecification. The estimated obsolescence rates of 23% to 32% are higher than the values typically assumed in the literature. The higher estimates reflect the model's view of R&D stocks as measuring the potential for future innovations rather than the knowledge stock applicable for production, which would have a lower obsolescence rate.

The analysis also enables a quantification of the value of R&D investment, which existing empirical models of innovation such as Lentz and Mortensen (2008) do not provide. The model estimates imply that the market value of the aggregate R&D stock as a fraction of the market value of all firms equals 28%. This is similar to the fraction of aggregate R&D investment to all investment and suggests that the growth opportunities from potential innovations offsets the high obsolescence rate. The finding demonstrates that R&D stocks form an important component of the total capital stock in the U.S. economy, as argued by Hall (2001).

The paper is organized as follows. Section 2 details the model and derives the optimal policy function and the rate of return to R&D. Section 3 presents the data and discusses flaws in the production function approach. Section 4 presents the results from estimating the model for all R&D firms combined and for firms in selected R&D intensive industries. Section 5 evaluates policy experiments within the context of the model and Section 6 concludes.

2 Model

The model economy consists of a large number of heterogeneous firms. The firms can invest in a physical capital stock, K, and an R&D stock, S. The firms return any cash remaining after investment as a dividend to shareholders. The investment decisions are made with the objective of maximizing the value of the firm. The model does not incorporate leverage.

2.1 Physical investment and profits

Firms operate as standard Q-theory firms facing a downward sloping demand curve. In each period, the output of the i^{th} firm follows a constant returns to scale Cobb-Douglas specification with

$$Y(K_i, x_i) = x_i K_i^{\zeta} L_i^{1-\zeta},$$

where ζ denotes the elasticity of output with respect to capital (time subscripts omitted). The firm faces a downward sloping demand curve

$$P_i = d_i Y_i^{-\eta},$$

where d_i denotes a demand shift parameter and η equals a constant price elasticity term. These assumptions correspond to a monopolistic competition setting where each firm possesses a degree of pricing power. Assuming a deterministic wage process and a per period fixed cost of operations, c, the profits of the firm can be written as

$$\Pi(K_i, z_i) = z_i K_i^{\theta} - c,$$

where z_i inherits the properties of x_i and d_i, and the decreasing returns to scale parameter θ is a function of η and ζ. The above expression is derived by substituting in the optimal labor choice each period. The subsequent analysis employs the above profit function. The firm spends its profits on capital and R&D investments and returns any cash left as a dividend to shareholders, with negative dividends indicating a share issuance. Investment and disinvestment of physical capital incur a quadratic adjustment cost, $b\frac{I_i^2}{2K_i}$, and firms can freely disinvest.[5] The capital stock of the firm in the next period is given by

$$K_i' = K_i(1-\delta) + I_i,$$

where δ denotes the depreciation rate and I_i equals investment.

2.2 R&D investment

In additional to physical investment, firms also invest in R&D. The R&D stock of the firm does not directly impact the production function as in Griliches (1979) and others. Instead, the R&D stock stochastically affects the transition of profitability across periods. The model views R&D stocks as measuring the potential for future innovations rather than a measure of the stock of ideas applicable for production. When a firm's R&D activity is successful, the firm realizes a profitability jump in the next period.[6] If it was unsuccessful, the firm will not realize a jump in profitability. Thus innovations reflect discoveries by firms that lead to an increased profitability of the firm's capital stock.[7] A fraction of the R&D stock becomes obsolete each period reflecting the conclusion or abandonment of R&D projects. The model attempts to capture the inherently uncertain nature of the innovation process through this mechanism, as a firm would realize a negative return from its R&D investment in a period in which it failed to innovate.

Denote the accumulated R&D stock of the firm at the end of each period by S_i'. Let R_i equal the investment in R&D activity. The law of motion for S_i is given by

$$S_i' = S_i(1-\gamma) + R_i, \tag{1}$$

where γ denotes the rate at which R&D stocks become obsolete. Let j_i denote a binary variable that equals 1 if the firm successfully innovates, and 0 otherwise. The probability of a successful

[5]For simplicity, the model does not include any irreversibility of selling physical capital as in Abel and Eberly (1994), or any costs of external finance as in Gomes (2001) and Brown, Fazzari, and Petersen (2009). Including these frictions would likely have no significant impact on the parameter estimates of interest while substantially complicating the model solution and estimation.

[6]Other papers which provide a similar treatment of the innovation process include Thompson (2001), Klette and Kortum (2004), Aghion, Bloom, Blundell, Griffith, and Howitt (2005), and Li (2008).

[7]The vintage capital models of Greenwood and Jovanovic (1999) and Hobijn and Jovanovic (2001) emphasize macro level technological revolutions that have different impacts on the value of current and future capital.

innovation is given by a Bernoulli distribution with success probability

$$p(j_i = 1) = 1 - \exp(-a\frac{S_i'}{K_i}), \tag{2}$$

where a is a parameter that influences the success rate of innovations. Higher R&D stocks lead to a greater probability of a successful innovation. This particular parametrization implies that success probabilities are concave in S_i'. The success probability decreases as the current capital stock increases. The scaling by capital stock can be thought of as capturing an increase in R&D project size with firm size. As such, larger firms require a greater level of R&D investment to generate the same probability of success as a small firm. The scaling ensures that large firms do not benefit disproportionately from R&D activity.[8] In the event of success, log profitability jumps by a constant, λ, which measures the improvement in the firm's profitability from a successful innovation.[9] An alternate model where firm size affects the increase in profitability from an innovation and does not affect the success probability would be far more cumbersome to estimate. The transition equation for profitability includes a standard AR(1) component plus jumps from innovations:

$$\begin{aligned} \log(z_i') &= \mu + \rho\log(z_i) + \lambda j_i + \epsilon_i && (3) \\ j_i &\propto \mathsf{B}(p(\frac{S_i'}{K_i})) && (4) \\ \epsilon_i &\propto \mathsf{N}(0,\sigma^2), \end{aligned}$$

The distribution for j_i is independent of the distribution for ϵ_i. In this setup, the jump intensity varies endogenously with R&D stocks. Further, firms base their decisions upon realizations of z_i', and do not distinguish between changes in profitability due to exogenous shocks or innovations. Therefore, the impact of an innovation will decay at the same rate ρ as exogenous shocks. These assumptions yield the simplification that only the current level of profitability enters into the firm's policy functions.

The model is agnostic on the source of the jump in profitability from a successful innovation.

[8] The endogenous growth theory model of Romer (1990) implies that innovation increases in the level of R&D. Subsequent work by Jones (1995b) demonstrates that this relationship does not hold in the data. Jones (1995a), Young (1998), and Segerstrom (1998) introduce endogenous growth models without scale effects. In the current model, a lack of scaling leads to explosive value functions as large firms will continue to increase R&D spending and increase their probability of success.

[9] Kortum (1997) employs a search theoretic approach, in which the rate of arrival of ideas is exogenous and the efficiency of the improvement depends on the R&D stock.

This may arise from either improvements in the current products of the firm, the introduction of entirely new products, or productivity increases. More formally, a successful innovation may result in an increase in the productivity parameter x_i or the demand shifter d_i. The model does not take a stand on whether patent protection is necessary for generating an increase in profitability (see Boldrin and Levine (2008)). The above approach allows R&D investment to have a broad impact on the firm. Correspondingly, the endogenous growth literature highlights both quality improvements and new product introductions as the outcome of innovations.

The timing of the firm's decisions warrant clarification. Firms enter each period with an R&D stock, a capital stock, and a profitability level. The firms invest in R&D and capital during the period. At the end of the period, each firm discovers whether it successfully innovated or not. If a firm succeeds, its next period profitability will be higher than if it did not. The accumulated R&D stock carries over to the next period and a fraction of it becomes obsolete after the realization of the profitability level z'. Note that this timing sequence differs from the law of motion for capital: the current period expenditures adds to the R&D stock which impacts the realization of the next profitability level. This timing structure was chosen as it yields a value function that is separable in the R&D stock. It also captures the idea that R&D investment in the current period impacts the firm's profitability in the next period.

The dividends paid by the firm in each period is given by

$$(\Pi(K_i, z_i) - R_i)(1-\tau) - I_i - b\frac{I_i^2}{2K_i},$$

where τ denotes a linear tax rate. As in the tax code, R&D investment is treated as a tax deductible expense. The tax rate parameter, τ, will be calibrated using data on aggregate taxes and profits. The model thus incorporates factors that affect taxes payable by firms, such as debt financing and the research and experimentation tax credit, in a parsimonious manner. The actual research and experimentation tax credit is computed using a complex formula based on the growth rate of a subset of R&D expenditures.

Denote the value of the firm after the realization of z_i but prior to the obsolescence of the R&D stock as $V(K_i, S_i, z_i)$.[10] For notational convenience, define

$$D(K_i, z_i) = \Pi(K_i, z_i)(1-\tau) - I_i - b\frac{I_i^2}{2K_i}. \tag{5}$$

[10]This definition ensures notational symmetry between capital and R&D stocks.

The value of the firm can be expressed as a solution to the following Bellman equation:

$$\begin{aligned} V(K_i, S_i, z_i) &= \max_{I_i, K_i', R_i, S_i'} D(K_i, z_i) - R_i(1-\tau) + \beta E_z[V(K_i', S_i', z_i')], \\ K_i' &= K_i(1-\delta) + I_i, \\ S_i' &= S_i(1-\gamma) + R_i, \\ S_i' &\geq 0. \end{aligned} \tag{6}$$

The lack of a non-negativity constraint on R&D investment implies that firms can sell their R&D if necessary. While firms do so infrequently in the simulation, this assumption is necessary for the subsequent characterization of the R&D policy function. The reversibility assumption is also supported by anecdotal evidence of firms selling partially developed products to other firms, particularly in the pharmaceutical sector. The expectation in the Bellman equation is taken over the joint distribution for j_i, ϵ_i. The results in Bertsekas (2000, Chapter 7) yield the existence and uniqueness of the solution to the above problem. Substituting the expression for R_i into the maximization problem yields

$$\begin{aligned} V(K_i, S_i, z_i) &= \max_{I_i, K_i', S_i'} D(K_i, z_i) + S_i(1-\gamma)(1-\tau) - S_i'(1-\tau) + \beta E_z[V(K_i', S_i', z_i')], \\ K_i' &= K_i(1-\delta) + I_i, \\ S_i' &\geq 0. \end{aligned}$$

Observe that S_i does not impact the optimization problem for either K_i' or S_i'. This motivates the conjecture that the value of the firm can be simplified as follows:

$$V(K_i, S_i, z_i) = G(K_i, z_i) + S_i(1-\gamma)(1-\tau). \tag{7}$$

The value of the R&D stock equals $S_i(1-\gamma)(1-\tau)$ due to the model's timing convention. The effective value of S_i equals its value after obsolescence, which is reduced by the tax deductibility of

R&D expenditures. Substituting the above expression into the Bellman equation, one obtains:

$$\begin{aligned} G(K_i, z_i) &= \max_{I_i, K_i', S_i'} D(K_i, z_i) - S_i'(1-\tau) + \beta S_i'(1-\gamma)(1-\tau) + \beta E_z[G(K_i', z_i')], \qquad (8) \\ K_i' &= K_i(1-\delta) + I_i, \\ S_i' &\geq 0. \end{aligned}$$

This establishes our conjecture and demonstrates that the value function is separable in the R&D stock. Note that this is not a general result, and it arises from the particular assumptions made about the structure of the optimization problem. The separability ensures that the payoff from R&D investment is negative in periods where the firm does not innovate.

2.3 R&D policy

The above analysis simplifies the solution of the optimal R&D policy. The optimal choice of S_i' impacts the current period dividend payment, the level of the R&D stock carried over to the next period, and the transition function for profitability z. The first two pieces are linear in S_i'. Let $\tilde{S}_i'$ be the optimal policy in the interior region where the $S_i' \geq 0$ constraint does not bind. The following proposition characterizes the optimal R&D stock:

Proposition 1 *The optimal R&D stock of the firm when $S_i' > 0$ is given by*

$$\frac{\tilde{S}_i'}{K_i} = \frac{1}{a}\left[\log(a) - \log\left((1-\tau)(1-\beta(1-\gamma))\right) + \log\left(\frac{\beta(E_z[G(K_i', z_i')|j_i = 1] - E_z[G(K_i', z_i')|j_i = 0])}{K_i}\right)\right].$$

Proof. Appendix A. ■

Therefore, the optimal policy function for R&D stocks is given by

$$\frac{S_i'}{K_i} = \max\left(\frac{\tilde{S}_i'}{K_i}, 0\right).$$

Optimal R&D investment increases with the expected increase in firm value per unit of capital from an innovation. Decreasing returns to scale imply that $\frac{G(K_i, z_i)}{K_i}$ is decreasing in K_i. Therefore, the level of the optimal R&D stock per unit of capital decreases as firm size increases. This matches the negative relationship between firm size and R&D investment observed in the Compustat data set. The negative relationship between firm size and R&D intensity also arises in Akcigit (2009). This is in contrast to Klette and Kortum (2004), where firm size has no effect on R&D intensity.

The above expression indicates that one can decompose the total benefit from R&D investment into an increase in the residual R&D stock and an expected increase in firm value from an innovation. The increase in firm value from an innovation incorporates both increases in profitability in the next period and beyond, and the optimal rebalancing of the capital stock following an innovation. This decomposition indicates that the rate of return, which measures the expected marginal increase in next period profits, focuses only on a narrow portion of the total benefit from R&D investment. The next section examines the rate of return to R&D in this setting.

2.4 Rate of return to R&D

A key objective of estimating the above structural model is to identify the private rate of return to R&D, defined as the marginal impact of R&D investment on next period profits. This definition closely mirrors that employed in the production function literature, where the rate of return equals the marginal impact of R&D investment on next period value added. From an accounting perspective, the profits measure employed in the study equals the value added measure minus selling and administrative expenses. The impact of R&D on profits would be a more relevant variable for a firm's optimization decision than the impact of R&D on value added. Some algebra, detailed in Appendix B, yields the following expression for the expected rate of return to R&D:

$$\frac{\partial}{\partial {S_i}'} E_z\left[\Pi(K_i', z_i')\right] = \left(\frac{(K_i')^{\theta} z_i^{\rho} \exp\left(\mu + \sigma^2/2\right)\left[\exp(\lambda) - 1\right]}{E_z[G(K_i', z_i')|j_i = 1] - E_z[G(K_i', z_i')|j_i = 0]}\right) \frac{(1 - \beta(1-\gamma))}{\beta}(1 - \tau). \quad (9)$$

The above expression lends itself to a natural interpretation in the context of the model. The first term equals the expected increase in next period profits from an innovation divided by the expected increase in firm value. The $1 - \beta(1 - \gamma)$ term equals one minus the residual value of R&D investment in the next period after obsolescence. The tax rate enters the above expression as it affects the optimal R&D stock.

The intuition for the above formula arises from the fact that R&D policies are derived from optimality conditions. This implies that the discounted total marginal benefit to R&D investment equals its marginal cost, 1. Parameter changes that shift the total benefit from next period profits to an increase in firm value or increase the residual value of R&D result in a lower rate of return. The maximum return of $\frac{1-\tau}{\beta}$ is obtained when $\rho = 0$ and $\gamma = 1$. The rate of return increases with the obsolescence rate γ. Changes in the parameter that influences the success probability, a, impact the rate of return indirectly through the $E_z[G(K_i', z_i')|j_i]$ terms.

The rate of return to R&D does not directly affect the optimal R&D policy in this setting, as they are both jointly determined. As such, the private rate of return to R&D does not provide an appropriate statistic for forming R&D policies in the context of the model. The intuition for this result is that the R&D policy is determined by the expected total benefit to R&D, while the rate of return focuses only on the expected impact on next period profits, which is one component of the total benefit. For example, an increase in the persistence of profitability results in an increase in optimal R&D investment, while simultaneously lowering the rate of return as more of the benefit from an innovation accrues in later periods.

2.5 Statistics of economic interest

The above analysis focuses on the rate of return to R&D. The model also enables the computation of other statistics that may be of interest.

The optimal R&D policy increases with the expected increase in firm value from an innovation. The percentage increase in firm value equals

$$\frac{E_z[V(K_i', S_i', z_i')|j_i = 1] - E_z[V(K_i', S_i', z_i')|j_i = 0]}{E_z[V(K_i', S_i', z_i')]}.$$

While it is difficult to construct an empirical counterpart to the above measure, it is helpful for understanding R&D decisions.

The λ parameter estimates the impact of an innovation on next period profits. The total impact of an innovation also depends on the speed at which the increase in profitability mean reverts. The present value equivalent permanent increase in profitability from an innovation provides an alternate method of evaluating the impact of an innovation on profits. This is given by:

$$\frac{1-\beta}{1-\beta\rho}(\exp(\lambda) - 1).$$

The above statistic ignores changes in the optimal capital stock following an innovation. On the other hand, it provides a more direct comparison with quality ladder models of growth.

The model also provides an estimate of the market value of the aggregate capital stock, which includes both physical capital and R&D stocks. The expected market value of the R&D stock of each firm is given by

$$E_z[V(K_i', S_i', z_i')] - E_z[V^n(K_i', z_i')],$$

where $V^n(K_i', z_i')$ denotes the value of a non-R&D firm. The expression equals the difference in market value between a firm that engages in R&D and a firm with the same capital stock and profitability levels that does not. The above measure includes the growth opportunities associated with R&D investment and is greater than the replacement cost, which equals the book value. The ratio of the aggregate R&D stock to the value of all R&D firms equals

$$\frac{\sum_i \left(E_z[V(K_i', S_i', z_i')] - E_z[V^n(K_i', z_i')|z_i]\right)}{\sum_i E_z[V(K_i', S_i', z_i')|z_i]}.$$

The value of the aggregate R&D stock as a fraction of the market value of all firms equals the above expression times the market value of R&D firms as a fraction of the market value of all firms, which is obtained from the data and equals 0.723.

3 Estimation

The study estimates the above models using simulated methods of moments estimation (see Gourieroux, Monfort, and Renault (1993) for details). This method involves comparing a selected set of data moments with the same moments from an artificial data set obtained by simulating the model for a given set of parameters. The parameter estimates are obtained by minimizing a quadratic form of the difference between the data and simulated moments. Appendices C and D discuss the estimation in more detail.

3.1 Data

The data for the estimation are obtained from the Compustat Annual data set. The data set includes information on profits, capital expenditures, and balance sheet items for listed U.S. corporations. The market value of equity is obtained from the linked CRSP data set. The sample period extends from 1987 to 2006, and was chosen to provide a stable tax environment. The sample excludes financial firms and regulated utilities.

Figure 1 plots aggregate R&D investment as a fraction of the sum of R&D and physical investment over time. The figure demonstrates the substantial level of R&D investment conducted by U.S. corporations. This provides suggestive evidence that the rate of return to R&D is lower than reported, as it would be difficult to reconcile the high level of R&D investment with substantial private returns under standard concavity assumptions. The high level of R&D investment also suggests that calibrations and estimations of RBC models may be misspecified when they do not take into account R&D investment.

The data on research and development expenditures (Compustat Annual data item 46) is available for more than half the observations. This series measures company funded R&D and excludes those funded by the government. As such, the results in this study measure the private rate of return to R&D. The baseline definition of the R&D firms sample includes all firm-year observations with positive R&D values. The non-R&D firms consist of all other observations. The study estimates the model first using data on R&D firms and then using data on all firms assuming that non-R&D firms have $a = 0$, which forces the success rate of innovations to be zero. As a robustness check, the analysis is repeated using an alternate definition of R&D firms that includes only firms that report R&D for all years, with positive values in at least one. This reduces the sample size by about 10%. In general, the results would be biased if some of the non-R&D firms engage in substantial R&D activity. Accounting rules provide some comfort in this regard, as they clearly specify the classification and reporting of R&D expenditures. Furthermore, most firms that report positive values for R&D expenditures do so for all years, while many firms do not engage in R&D expenditures at all. The fraction of observations with firms either reporting R&D data for all years or none at all equals 87.75%. These facts indicate that a clear set of R&D firms can be extracted from the data.

The production function regressions require the construction of the R&D stock for each firm. The R&D stock for an assumed obsolescence rate is constructed using a perpetual inventory method, with historical R&D expenditures from 1975 onwards. The initial R&D stock is obtained assuming a real growth rate of 5% for R&D expenditures. The initial assumptions have little impact, as the regressions use observations from 1987 onwards, matching the sample selection for the structural estimation. The log value added measure equals net sales (data item 12) minus cost of goods sold (data item 41).

Panel A of Table 1 presents summary statistics using the two definitions of R&D firms and for non-R&D firms under the baseline definition. Size equals the book value of assets (data item

6). Profitability equals operating income before depreciation (data item 13) scaled by book assets. Investment is defined as capital expenditure on plant, property, and equipment (data item 30) minus retirements (data item 184). Tobin's Q is computed as the market value of assets scaled by the book value of assets.[11] An alternative would be to define Q as the market value of capital scaled by the replacement value of capital. While this definition may be preferred where measurement error is a concern (Erickson and Whited (2000)), it also leads to large outliers, particularly in R&D intensive firms, that may bias the estimation. The variables are Winsorized at the 99% level to eliminate outliers. The data used to compute the second moments and autocorrelation terms adjust for firm and year fixed effects within the R&D and non-R&D firms, respectively.

The level of R&D spending is quite substantial when compared to total assets and is larger than the rate of investment. However, aggregate expenditures on R&D are lower than total capital expenditures within R&D firms, as small firms do a disproportionately larger level of R&D spending. This is consistent with the policy function given in Proposition 1. R&D firms have noticeably different firm values and profitability levels. While these firms have a much higher average Tobin's Q value they exhibit lower average profitability levels after accounting for their R&D expenditures.[12]

Panel B of Table 1 presents summary statistics for R&D firms in five selected R&D intensive industries. These industry groups mostly follow the 49 industry classifications constructed by Kenneth French using 4-digit SIC code data.[13] The one exception being the medical equipment category where I combined the medical and laboratory equipment industries. Appendix E details the SIC codes used in each industry category. These industries were chosen as they have the highest concentration of firms engaging in R&D as well as the highest R&D intensities.

Firms in these industries are on average smaller than all firms that engage in R&D. They tend to have higher valuations and higher R&D intensities. The mean level of R&D investment is particularly high in the software industry and the pharmaceutical industry. The chips industry has the highest level of profitability and the pharmaceutical industry has a negative mean profitability level. This arises due to the presence of small firms that have a very high R&D intensity, resulting in large negative profitability levels.

[11] Market value of assets equals the book value of assets + market value of equity - book value of equity - deferred taxation.

[12] Chan, Lakonishok, and Sougiannis (2001) find that the market value of firms reflect their R&D activity.

[13] The industry classifications are available at http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

3.2 Identification

The results obtained from indirect inference methods are sensitive to the matched moments. In particular, it is important to select moments that are informative about the parameters of interest. Intuitively, if a given simulated moment varies strongly with a parameter, matching this moment to the data will be informative about the underlying parameter.

The matched moments include the first moment of profitability for the R&D firms. This will be informative about the parameters of the profitability process, as well as the parameters influencing the success rate and impact of R&D spending. The second moments of profits help identify σ. The first and second moments of R&D spending provide information on the R&D parameters $\lambda, a,$ and γ. The first moment of Tobin's Q helps inform the returns to scale parameter θ and the R&D parameters. The second moment of Q also help identify θ and the volatility of profits. The autocorrelation coefficient for profits and R&D spending helps inform the ρ parameter. Finally, the second moment of investment identifies the level of adjustment cost parameter b. It should be noted that each moment provides information on almost all of the parameters.

One potential concerns is whether the estimation can separately identify the R&D related parameters. The identification of these parameters arises from matching the profitability and valuation moments, in addition to the R&D moments. Changes in the R&D related parameters effect the R&D policies in the model, thereby changing the steady state profitability and valuation levels.

Figure 2 plots the change in the mean R&D to assets ratio and mean profitability level with the R&D related parameters. Panels A, B, and C plot the variation in these moments with changes in λ, a, and γ, respectively. The figures are constructed using the point estimates presented in Section 4.1. An increase in λ leads to a sharp increase in R&D investment as firms respond to the increased benefits from an innovation. Over these parameter values, an increase in a lowers R&D investment by reducing the marginal impact of R&D investment on firm value. Although an increase in γ reduces the optimal R&D stock, it does so less than one-to-one, leading an increase in R&D investment to counteract the increased obsolescence.

In each case, changes in R&D related parameters also have a clear, though potentially counterintuitive, effect on the mean profitability level. Increases in R&D investment lead to more innovations and higher realizations for z. This also increases the steady state capital stock, which leads to lower average profitability levels due to the presence of decreasing returns to scale in the profit function. In addition, average profitability levels decrease as they are computed net of R&D investment.

The average valuation levels are also informative about the R&D parameters in the model. For instance, an increase in λ will lead to an increase in Tobin's Q as firms face larger potential profit increases from innovations. Thus, the estimation of the R&D parameters uses the moments for R&D investment, profitability, and firm value.

Some of the auxiliary parameters in the model are calibrated to simplify the estimation of the parameters of interest. The calibrated parameters are the discount factor β, the per period fixed cost of operations c and the depreciation rate δ. The calibrated value of $\beta = 1/1.04$ implies a return to capital of 4%.[14] The fixed cost of operation is set at 15% of the profits prior to R&D expenditures for the median firm in each sector.[15] The depreciation rate $\delta = 0.063$, a value which equals the mean investment rate for R&D firms. This ensures that the steady state investment rate in the simulated data, which equals the depreciation rate, matches the actual data. The μ parameter is not identified with the chosen moments as it acts as a scaling parameter. Fix $\mu = -0.3$.

3.3 Estimating the rate of return using production functions

This section presents the rate of return estimates obtained using the production function approach applied to the Compustat data set. While there are many well known econometric issues with this approach, the analysis highlights two that, to the best of my knowledge, have not been highlighted in the literature. Specifically: the estimated rates of return to R&D violate an optimality condition and they vary sharply with the assumed obsolescence rate.

The production function approach pioneered by Griliches (1979) specifies value-added by the firm as a multiplicative function of factor inputs, including the R&D stock:

$$Y_{i,t} = a_t K_{i,t}^{\alpha} L_{i,t}^{\beta} S_{i,t}^{\nu}. \tag{10}$$

A linear regression of log value-added on log inputs yields the elasticity of value-added with respect to the R&D stock ν. Doraszelski and Jaumandreu (2009) build on this approach and estimates a model where R&D investment influences the Markov transition function for productivity. The

[14] An alternative calibration that is employed by Bloom (2009) sets $\beta = 1/1.065$. I estimate the model using this value and find broadly similar results. Recent evidence also points to a lower equity premium for the sample period. Fama and French (2002) find an equity premium in the range of 2.55% to 4.32% using dividend and earnings growth rates. Cogley and Sargent (2008) argue that the declining equity premium arises due to a pessimistic prior following the Great Depression.

[15] Eberly, Rebelo, and Vincent (2008) estimate fixed costs to be 22% of revenue net of variable costs, including R&D expenditures.

construction of the R&D stock $S_{i,t}$ in this regression follows a perpetual inventory method using data on R&D expenditures applied to equation (1). The obsolescence rate of R&D investment necessary for the construction of the R&D stock is typically assumed to be 15%, following Griliches and Mairesse (1984). One common finding in this literature is that the estimated elasticity, ν, varies little with the assumed obsolescence rate, γ.

The rate of return to R&D in this framework equals the dollar payoff next period for a marginal dollar investment and is given by:

$$\frac{\partial Y_{i,t}}{\partial S_{i,t}} = \nu \frac{Y_{i,t}}{S_{i,t}}.$$

The above expression highlights a potential difficulty with this approach. The rate of return varies with both the elasticity parameter ν and the measured R&D stocks. However, changes in the obsolescence parameter, γ, change the measured R&D stocks while having little effect on the elasticity estimate, ν (see Hall and Mairesse (1995)). This implies that different assumptions on the obsolescence rate will lead to different rates of return, a difficulty compounded by the lack of evidence on the appropriate obsolescence rate for R&D stocks.

Define the total payoff to R&D investment in this framework to be

$$\frac{\partial Y_{i,t}}{\partial S_{i,t}} + 1 - \gamma.$$

This equals the rate of return plus the residual value of R&D investment. In a frictionless setting, optimal R&D investment by the firm requires that

$$\beta E\left[\frac{\partial Y_{i,t}}{\partial S_{i,t}} + 1 - \gamma\right] \leq 1,$$

i.e., the marginal cost of R&D investment should be less than or equal to the total payoff. The production function approach, however, does not impose the above optimality condition, and many of the estimates reported in the literate violate it, unless one assumes a very high discount rate. Uncertainty in the outcome from R&D, by itself, would not justify a high discount rate as only the component of uncertainty that correlates with aggregate conditions affect the discount rate.

Table 2 presents the results of the production function regression for different assumptions of γ. The sample ranges from 1987 to 2006, and all variables except employment are adjusted for inflation using the GDP deflator for non-residential investment. The capital stock is constructed using the perpetual inventory method of Salinger and Summers (1983). The R&D stock is constructed using a perpetual inventory method with assumed obsolescence rates ranging from 0 to 0.25. Panel A reports the results of pooled regressions that includes year and industry dummies, while Panel B reports the results of fixed effects regressions with year dummies. The standard errors are heteroskedasticity robust and cluster at the firm level. The rates of return and the total payoff are computed using the sample median of value-added to R&D stocks.

The results from the pooled regression (Panel A) match those found in the literature. The elasticity of value added with respect to R&D is significant, both economically and statistically. The rates of return estimates are quite high, particularly under the standard assumption of $\gamma = 0.15$. Note that the implied total payoff to R&D are substantially higher than one, indicating violation of the above optimality condition unless one assumes a very high discount rate. Furthermore, different assumptions on the obsolescence rate result in sharply different rates of return estimates.

The fixed effect regression (Panel B) leads to uniformly lower elasticity estimates. The difference between fixed effect estimation and pooled estimation was highlighted by Hall and Mairesse (1995). The fixed effect estimates lead to lower total payoffs to R&D that do not violate optimality. However, this regression does not clearly identify the rate of return to R&D as different obsolescence rates lead to sharply different rates of return. This arises again due to the fact that the elasticity of value-added does not vary with the assumed obsolescence rates. The five estimates of ν are, in fact, statistically indistinguishable from one another. These findings demonstrate that the production function approach fails to provide reliable estimates of the rate of return to R&D due to the lack of information on γ and the lack of optimality restrictions.

4 Results

This section presents the results from estimating the structural model. First, I estimate the model using data on R&D firms and then estimate it using data on all firms assuming that non-R&D firms do not engage in R&D for some exogenous reason. Then I present results obtained by estimating the model for R&D firms in five selected R&D intensive industries. As a robustness check, the last subsection presents results obtained from estimating the model using an alternate definition of R&D firms.

The treatment of firms that engage in R&D and those who do not in the estimation warrants some discussion. Most of the firms in the data can be split in to two categories: one group engages in R&D investment in all periods, and another group that engages in no R&D in any period. In fact, less than 15% of the sample reports R&D investment in some periods, but not others. This indicates that the decision to engage in R&D is very much a firm specific fixed choice that does not fit well into a dynamic framework. As such, this study considers the decisions of firms not to engage in R&D as reflecting some exogenous factor.[16]

4.1 R&D firms

The estimation of the model on R&D firms identifies the impact of R&D on profits, the success rate from R&D, and the obsolescence rate of R&D stocks. The rate of return to R&D and other statistics of interest can be inferred from the structural estimates and the simulated distribution of firms. Table 3 presents the results of the simulated method of moments estimation. Panel A reports the matched moments for the data and simulations, respectively. Panel B reports the estimated parameter values. Panel C reports some statistics of economic interest from the simulated data set.

Examination of the matched moments indicates that the model succeeds in matching the first and second moments of R&D to assets. The potential for increases in profitability captured in the model is sufficient to generate the high level of R&D investment observed in the data. The model comes close to matching the Tobin's Q moments, which incorporate the value of the R&D stock and reflect growth options due to potential innovations. The model fares less well at matching the first and second moments of profitability. The estimation yields a ρ value that is much higher than regression estimates, resulting in higher autocorrelations for both profitability and R&D expenditures.

The model has difficulty generating the high level of Tobin's Q given the observed levels of profitability. A related difficulty arises from matching the variation in Tobin's Q given the variation in profitability. A high autocorrelation value helps with both of these challenges. For a given distribution of profits, an increase in persistence will increase Tobin's Q for high profitability firms and lower it for low profitability firms. This impact will be greater for high Q firms, resulting in higher first and second moments for Tobin's Q. A high autocorrelation value also increases R&D investment by extending the benefit of an innovation further into the future.

The results indicate that R&D firms generate an economically and statistically significant in-

[16]For instance, consider two pharmaceutical firms, one specializing in drug development, and the other specializing in manufacturing drugs developed by other firms. The decision on whether a firm engages in a business line that requires R&D investment would be influenced by a variety of exogenous factors, including the expertise of the firm's founders.

crease in profitability from a successful innovation. The point estimate for λ of 0.152 implies that a successful innovation increases the profitability of the firm by 16.4% ($= \exp(\lambda) - 1$). Ignoring any changes in the capital stock, this is equivalent in present value terms to a 4.3% permanent increase in profitability. This corresponds to the increase in productivity from an innovation in a quality ladder model. In the steady state, 66.9% of firms in the R&D sector innovate each period. However, there is sharp cross-sectional variation in the probability of innovation, reflecting the sensitivity of the R&D policy function to the current profitability and size of the firm.

An application of equation (9) yields the distribution of the expected rate of return to R&D in the simulations. The estimates imply a mean rate of return of only 3.3%, with a 95th percentile value of 5.4%. These values are substantially lower than the rates of return reported in the literature. The residual value of R&D investment at the end of the next period equals $0.655 (= \beta * (1 - \gamma))$, implying a maximum expected rate return of 35.9%. The rate of return is driven lower by the high ρ parameter, which shifts the benefits of an innovation further into the future. This decreases the ratio of expected increase in next period profits from an innovation to the expected increase in firm value, which also reflect changes in the optimal capital stock following an innovation. While the rate of return is much lower than the values reported in the literature, it does not deter firms from maintaining a high level of R&D investment.

The mean expected increase in firm value from an innovation equals 16.4%. A large impact of innovations on firm value induces firms to incur the high level of R&D expenditures observed in the data and the simulations. While the appropriate sample counterpart is unclear, the estimate is consistent with anecdotal evidence of sudden jumps in firm value upon the release of news about potential innovations.

A regression of log value-added, which equals zk^{θ} in the model, on the logs of the capital stock and the lagged R&D stock replicates the production function regression on the simulated sample. This regression is clearly misspecified in this setting and yields a median rate of return of 54.4%, much higher than the true value. This finding demonstrates that the production function approach can lead to misleading results due to model misspecification.

The estimated obsolescence rate of $\gamma = 0.319$ is not directly comparable to the values employed in the production function approach. The R&D stock in the model measures the potential for generating future innovations. Prior innovations from R&D would be reflected in a higher level of profitability, z. In contrast, the R&D stock in the production function approach reflects all the accumulated knowledge application for production. The estimated high obsolescence rate indicates that a substantial proportion of a firm's R&D expenditures ceases to have any value if the firm fails to innovate. Hall (2007) obtains obtains similar estimates for γ using a valuation based approach, thus providing support for this study's view of R&D stocks as measuring the stock of experiments towards the next innovation. It is less plausible to view the stock of ideas applicable for production becoming obsolete at these rates. The γ estimate implies a half-life of R&D expenditures of 1.80 years, indicating a moderate lifespan for R&D projects.

The point estimate for the adjustment cost parameter b is somewhat higher than those found by Cooper and Haltiwanger (2006), and Eberly, Rebelo, and Vincent (2008). The goodness-of-fit statistic reported in Panel B, $\hat{\Phi}$, has values in the same range as those reported in Cooper and Haltiwanger (2006).[17] The precision of the estimates reflects the sensitivity of the chosen moments to changes in the estimated structural moments.

4.2 All firms

The previous estimates employed only data on firms that engage in R&D. This section presents results of an alternative specification that uses data on all firms to estimate the model. Estimating the model on all firms has the benefit of yielding the market value of the R&D stock as a fraction of the market value of all firms. Extending the model to non-R&D firms requires assuming that these firms do not engage in R&D for some exogenous reason. The exogeneity of whether to engage in R&D can be justified as most firms in the sample either report R&D for all years, or report no R&D.

Table 4 presents the results from estimating the model on all firms. The matched moments now include additional moments on profitability and Tobin's Q for non-R&D firms, enabling the model to use information on these firms for estimating the non-R&D related structural parameters. The second moment of investment and the autocorrelation of profitability are computed using data on all firms. Panel A reports the matched moments for the data and simulations, respectively. Panel B reports the estimated parameter values. Panel C reports some statistics of economic interest from the simulated data set.

The estimates result in a slightly worse fit for the first and second moments of R&D than before. The valuation moments from the simulated sample are lower than the data for R&D firms but not for the non-R&D firms. This suggests that there may be additional growth options associated with R&D that is not captured by the model, or that there may be differences in the underlying

[17] See Appendix C for details.

non-R&D structural parameters. The simulated profitability moments of R&D firms are similar to those reported above. The model also overshoots the profitability moments for non-R&D firms. The estimates continue to overstate the persistence of profitability. The goodness-of-fit statistic, $\hat{\Phi}$, is sharply higher than before, indicating that the model fits the data less well for all firms.

The point estimate for λ implies a 20.0% increase in profitability from a successful innovation. This is larger than the corresponding value obtained in the baseline estimation. The lower value of ρ requires an innovation to have a greater initial impact on profits. In present value terms, the estimate corresponds to a 3.64% permanent increase in profitability, lower than the comparable estimate for the baseline model. The steady state rate of innovation is lower, with a mean probability of success of 0.578. This arises due to the lower R&D intensity level.

The mean rate of return to R&D expenditures equals 4.6%, higher than the estimated value from the baseline model. Nonetheless, this estimate is much lower than the values obtained using the production function approach. Intuitively, the lower estimate for ρ implies that more of the benefit of an innovation is realized in the next period, thereby increasing the ratio of expected increase in profits from an innovation to the expected increase in firm value. The similar estimate for γ results in an unchanged residual value of R&D expenditures.

The expected increase in firm value from an innovation has a direct impact on R&D policy in the model. The mean expected increase in firm value of 12.7% is lower than the corresponding value in the baseline model. There is also less cross-sectional variation in the increase in firm value. The higher impact of a successful innovation on profits does not translate to firm values, as the benefit dissipates faster.

Replicating the production function regression with the simulated sample yields a median rate of return of 48.8%, similar to the results obtained above and much higher than the actual rate of return. This indicates that the production function method can substantially overstate the rate of return due to model misspecification.

The estimation yields a market value of the aggregate R&D stock equal to 28.2% of the market value of all capital.[18] In comparison, the aggregate value of R&D investment to all investment in the sample period equals 24.6%. The similar values of the R&D stock and the R&D intensity demonstrate that the higher growth options associated with R&D investment offsets the higher obsolescence rate. The result indicates that a substantial portion of the aggregate capital stock consists of R&D stocks and highlight the need for incorporating R&D investment into models of the

[18] The R&D Satellite Accounts recently constructed by the BEA measures the book value of R&D stocks to be 7.3% of the sum of R&D and physical capital as of 2003.

macroeconomy (Comin and Gertler (2006) and Barlevy (2007)). As R&D stocks comprise a large fraction of all intangible capital, the finding supports the argument of Hall (2001) that intangible capital accumulation forms an important component of U.S. investment.

4.3 Industry level estimation

The results from the previous sections provide estimates of the model at a macroeconomic level. Estimating the model at an industry level sheds light on the variation in R&D-related parameters across industries.

Table 5 presents the results from estimating the model on R&D firms within five selected R&D-intensive industries. Panel A reports the first moments of R&D to assets, Tobin's Q, and profitability for the data and the model. The model is able to generate the level of R&D and valuations for all industries except pharmaceuticals. As before, the mean level of profitability is somewhat higher in the data, with the model coming closest with the chips industry. Overall, the results indicate that while the model matches features of the data for most of the R&D intensive industries, it fails to replicate the data for pharmaceuticals. This may be due to the model's failure to capture the long gestation period of R&D projects in the pharmaceutical industry, where a typical project takes 8 to 12 years before becoming a marketable drug.

Panel B presents the R&D-related structural parameters obtained from estimating the model for the five selected industries. The estimated λ values indicate that the impact of an innovation on profits is highest for the pharmaceutical industry and reflects the high level of R&D investment by these firms. The economic and statistical significance of the point estimates indicate that innovations have a clear impact on the profitability of firms in each of these industries. The obsolescence rates vary across the industries, with lower values for industries with higher profitability levels. The model's inability to reconcile the high level of R&D investment and the low level of profitability in pharmaceutical firms is reflected in a substantially higher goodness-of-fit measure compared to the other industries.

Panel C reports the statistics of economic interest from the estimation. The results highlight the uncertainty that firms face in generating innovations. The rate of return to R&D varies little across the industries, and is similar to the values obtained from the estimation on all R&D firms. On the other hand, the equivalent permanent increase in profitability and the increase in firm value from an innovation tend to be higher than the previous estimates, particularly for industries with very high R&D activity. The results indicate that the variation in R&D activity across industries reflects the effect of innovations on profitability and firm value, as highlighted in the model.

4.4 Estimates using an alternate definition of R&D firms

The baseline estimation includes all firm-year observations with positive R&D data in the construction of the sample. This section presents the results of estimating the model only on firms that report R&D for all of the years they report data. This excludes firms that have positive R&D in some years but missing values in others and includes several observations where a firm that typically does positive R&D reports a zero value. Table 6 presents the results of the estimation. Panel A reports the matched moments for the data and simulations, respectively. Panel B reports the estimated parameter values. Panel C reports some statistics of economic interest from the simulated data set.

The results indicate a better fit for these data compared with that obtained in Section 4.1. The model delivers a slightly lower first moment of R&D compared to the data. The model has difficulty reconciling the observed levels of profitability and valuations as before, resulting in a good fit for Tobin's Q and a higher level of profitability than in the data. The autocorrelation terms are also higher than the corresponding data values.

The estimates imply a statistically and economically significant increase in profitability from an innovation of 15.5%. This translates to a higher permanent increase in profitability than in Section 4.1, due to a higher autocorrelation term. The lower a and γ parameters generate a higher probability of innovation, with significant cross-sectional variation.[19] Nonetheless, firms continue to face significant uncertainty in their R&D investment with substantial losses if they fail to innovate in a given period. The high level of R&D investment is driven partly by the large increase in firm value from an innovation.

The rate of return to R&D continues to be far below the estimates reported in the production function literature. Optimality conditions imply that the rate of return is bounded by the residual value of R&D, implying a maximum rate of return of 0.273. The rate of return is further reduced by the high degree of persistence in profitability, which implies that much of the benefit from an innovation accrues in the periods following the innovation. On the other hand, replicating the production function regression using the simulated sample yields a rate of return of 36.3%. This indicates that the production function method can generate rates of return substantially different

[19]The 5th percentile of the probability of success equals 0, as a small fraction of firm-year observations in the steady state have an R&D stock of 0.

from the true value due to model misspecification.

5 Policy experiments

This section presents the results of counterfactual experiments on the following: a subsidy on R&D expenditures and an increase in the impact of a successful innovation on profits. The policy experiments reflect the model and are partial equilibrium in nature. The analysis provides a method to quantify the impact of changes in key parameters on the level of R&D expenditures, profits, firm value, and the rate of innovation within the context of the model.

5.1 A subsidy on R&D expenditures

An active policy debate exists regarding the use of the tax code to subsidize R&D investment. The research and experimentation tax credit uses a complex formula based on growth rates in these expenditures, which comprise only a subset of total R&D. The estimation presented in the previous section accounts for the tax credit through the calibration of the gross tax rate on operating income, τ. The tax credit amounted to about 2.5% of all spending on R&D by firms from 1990 to 2001 (Moris (2005)). Currently, the tax credit is not permanent and requires occasional Congressional reauthorization. The counterfactual experiment studies the impact of a further 5% R&D tax credit within the context of the model (see Bloom, Griffith, and van Reenen (2002) for a cross-country study on the impact of tax subsidies on innovation). Table 7 reports the moments of interest using the actual estimates, and with a 5% subsidy on R&D expenditures. Panel A reports the values for the baseline model with R&D firms, Panel B reports the values using the estimates from all firms, and Panel C reports the results with the alternate definition of R&D firms. While the analysis does not impose any other taxes to offset the drop in tax revenue, it can provide a basis for a cost-benefit analysis of R&D tax credits.

The subsidy for R&D expenditures leads to higher spending and a greater level of innovation under all models. The tax subsidy increases the rate of innovation by 2.8, 2.2, and 5.8 percentage points for the models estimated in Sections 4.1, 4.2, and 4.4, respectively. This is driven by a 0.4% to 0.9% increase in the rate of R&D expenditures in response to the additional tax credit. The findings imply a mean increase in R&D expenditures of 0.8 to 1.5 dollars for a dollar increase in the tax credit (computed as the change in mean R&D expenditures divided by the cost of the subsidy). In comparison, Hall and van Reenen (2000) estimate a 1 to 1 impact of R&D tax credits on R&D

expenditures. The mean profitability level decreases due to the curvature of the profit function as the increase in the capital stock offsets the profit increase. The effect of the tax subsidy on R&D expenditures and the rate of innovation does not increase with the estimated rate of return across the models. This arises due to the joint determination of the private rate of return and the optimal R&D expenditure level.

This analysis quantifies the impact of an R&D tax subsidy on firm's policies. One clear limitation of this analysis is that it does not shed light on the optimal level of R&D subsidies. Tackling this question would require expanding the model and the empirical analysis to incorporate spillovers from R&D investment.

5.2 An increase in the impact of a successful innovation

An increase in the impact of a successful innovation provides a method of incorporating into the model changes in patent policy that favor innovators. Gilbert and Shapiro (1990) argue for viewing patent policy in terms of the degree of profits accruing to the innovator. Eaton and Kortum (1999) estimate a model of international research and patenting and find a strong effect of strengthening patent rights on productivity growth. In the context of this study, increased patent protection can be viewed as an exogenous increase in λ. While the following analysis is uninformative on the optimal level of patent protection, it quantifies the effect of a change in patent policy on R&D expenditures and the rate of innovation. Table 8 reports the moments of interest using the actual estimates and the new values obtained with a 5% increase in λ.[20] Panel A reports the values for the baseline model with R&D firms, Panel B reports the values using the estimates from all firms, and Panel C reports the results with the alternate definition of R&D firms.

The increase in the impact of a successful innovation on profits leads to higher investment in R&D for all models. The rate of innovation increases by 3.3, 2.5, and 7.6 percentage points, respectively. The increase in λ has a larger impact on the rate of innovation compared to the R&D tax subsidy. It also has a bigger impact on firm value. The tax subsidy benefits all R&D firms, while the increase in λ benefits firms only if they innovate. As firms with higher valuations invest more in R&D and are more likely to innovate, the change in λ has a bigger impact on these firms, as reflected in the larger increase in the mean value of Tobin's Q.

[20] This calculation keeps the fixed cost unchanged while λ varies. In the estimation, the fixed cost would vary with the parameter vector, specifically the size of the median firm.

6 Conclusion

The R&D estimation literature has focused more on establishing statistical significance of R&D investment on productivity, rather than a precise quantification (Griliches (2000)). However, policy prescriptions require precise estimates. One approach would be to improve upon available data sources (Griliches (1994)). Another approach would be to improve upon our methods. I argue that the production function method currently employed fails to provide reliable estimates of the rate of return to R&D. Using an optimality framework, I obtain rates of return that are much lower than the values reported in the literature. The approach also yields estimates of other statistics of interest, such as the rate of innovation by firms, the impact of an innovation on profits, and the market value of the R&D stock.

Although this study obtains a low private rate of return to R&D using a specific model, the main argument applies more broadly. Optimality conditions imply that the private return to R&D can not be as high as reported in the literature unless both of the following conditions apply: most of the payoff to R&D investment is realized in next period profits, and R&D investment has a high obsolescence rate. However, it is unlikely that either one of these conditions hold.

The failure of the production function method to estimate the private rate of return to R&D indicates that one should treat the estimates of the social returns to R&D obtained using this approach with equal or more skepticism. This motivates the need for different approaches, such as the model based approach of Comin (2004), or the estimation of the spillovers from R&D carried out by Bloom, Schankerman, and van Reenen (2007). Further research into this area may prove fruitful.

Appendix

A Proofs

Proposition 1 *The optimal R&D stock of the firm when $S_i' > 0$ is given by*

$$\frac{\tilde{S}_i'}{K_i} = \frac{1}{a}\left[\log(a) - \log\left((1-\tau)(1-\beta(1-\gamma))\right) + \log\left(\frac{\beta(E_z[G(K_i', z_i')|j_i = 1] - E_z[G(K_i', z_i')|j_i = 0])}{K_i}\right)\right].$$

Proof. The first order condition for the optimal R&D stock yields:

$$-(1-\tau) + (1-\tau)\beta(1-\gamma) + \beta\frac{\partial E_z[G(K_i', z_i')]}{\partial S_i'} = 0. \tag{A.1}$$

The impact of R&D spending on the expected value of the firm in the next period can be clarified by an application of the law of iterated expectations,

$$E_z[G(K_i', z_i')] = E_z[G(K_i', z_i')|j_i = 1]p(j_i = 1) + E_z[G(K_i', z_i')|j_i = 0]p(j_i = 0).$$

Substituting the expression for $p(j_i)$ given in (2) one obtains:

$$E_z[G(K_i', z_i')] = E_z[G(K_i', z_i')|j_i = 1](1 - \exp(-a\frac{S_i'}{K_i})) + E_z[G(K_i', z_i')|j_i = 0]\exp(-a\frac{S_i'}{K_i}). \tag{A.2}$$

Recall that the R&D stock has no effect on the conditional expectation of $G(K_i', z_i')$ given j_i. The derivative of the above expression with respect to S_i' yields:

$$\frac{\partial E_z[G(K_i', z_i')]}{\partial S_i'} = \frac{a}{K_i}\exp(-a\frac{S_i'}{K_i})\left(E_z[G(K_i', z_i')|j_i = 1] - E_z[G(K_i', z_i')|j_i = 0]\right).$$

Substituting the above expression into the first order condition given in (A.1) yields the optimal policy function for the firm's R&D stock,

$$(1-\tau)(1-\beta(1-\gamma)) = \beta\frac{a}{K_i}\exp(-a\frac{\tilde{S}_i'}{K_i})\left(E_z[G(K_i', z_i')|j_i = 1] - E_z[G(K_i', z_i')|j_i = 0]\right) \tag{A.3}$$

$$\Rightarrow \frac{\tilde{S}_i'}{K_i} = \frac{1}{a}\left[\log(a) - \log\left((1-\tau)(1-\beta(1-\gamma))\right) + \log\left(\frac{\beta(E_z[G(K_i', z_i')|j_i = 1] - E_z[G(K_i', z_i')|j_i = 0])}{K_i}\right)\right].$$

Some algebra reveals that the second order condition with respect to S_i' is negative, ensuring that the F.O.C.s yield the optimal policy in the interior region. ■

B Expected rate of return to R&D

The expected rate of return to R&D expenditures is given by:

$$\frac{\partial}{\partial {S_i}'} E_z \left[\Pi(K_i', z_i')\right] = \frac{\partial}{\partial {S_i}'} E_z \left[z_i'(K_i')^{\theta}\right]. \tag{B.1}$$

The conditional expectation of z_i' can be written as,

$$E_z[z_i'] = E_z[z_i'|j_i = 1]p(j_i = 1) + E_z[z_i'|j_i = 0]p(j_i = 0). \tag{B.2}$$

Using the transition equation for z_i' in (3) and the log-normality assumption, the conditional expectations can be written as

$$\begin{aligned} E_z\left[z_i'|j_i = 1\right] &= \exp\left(\mu + \lambda + \rho\log(z_i) + \sigma^2/2\right) \\ &= z_i^{\rho}\exp\left(\mu + \sigma^2/2 + \lambda\right). \\ E_z\left[z_i'|j_i = 0\right] &= z_i^{\rho}\exp\left(\mu + \sigma^2/2\right). \end{aligned}$$

The partial derivative of the probability of an innovation equals:

$$\frac{\partial}{\partial {S_i}'} p(j_i = 1) = \frac{a}{K_i}\exp(-a\frac{S_i'}{K_i}). \tag{B.3}$$

Substituting in the above expressions yield the following simplification:

$$\begin{aligned} \frac{\partial}{\partial {S_i}'} E_z\left[\Pi(K_i', z_i')\right] &= (K_i')^{\theta}\frac{a}{K_i}\exp(-a\frac{S_i'}{K_i})\left[z_i^{\rho}\exp\left(\mu + \sigma^2/2 + \lambda\right) - z_i^{\rho}\exp\left(\mu + \sigma^2/2\right)\right] \\ &= (K_i')^{\theta}\frac{a}{K_i}\exp(-a\frac{S_i'}{K_i})z_i^{\rho}\exp\left(\mu + \sigma^2/2\right)[\exp(\lambda) - 1]. \end{aligned}$$

Substituting the expression for $\frac{a}{K_i}\exp(-a\frac{S_i'}{K_i})$ given in (A.3) and collecting terms yields the following expression for the rate of return to R&D.:

$$\frac{\partial}{\partial S_i{}'} E_z\left[\Pi(K_i', z_i')\right] = \left(\frac{(K_i')^\theta z_i^\rho \exp\left(\mu + \sigma^2/2\right)\left[\exp(\lambda) - 1\right]}{E_z[G(K_i', z_i')|j_i = 1] - E_z[G(K_i', z_i')|j_i = 0]}\right)\frac{(1 - \beta(1-\gamma))}{\beta}(1-\tau).$$

C Simulated method of moments

The indirect inference method of Gourieroux, Monfort, and Renault (1993) obtains parameter estimates by matching a set of selected moments from the data to those obtained by simulation. Denote the true values of the structural parameters by Ψ^*. The matched moments can be written as a solution to a minimization problem $Q(Y, M)$, where Y denotes the data and M the moments to be matched. The data moments are then given by

$$\hat{M} = \arg\min_M Q(Y_N, M), \tag{C.1}$$

where Y_N denotes a data matrix with N observations. The corresponding moments for the simulated data set with parameter vector Ψ and $n = N \times S$ observations are given by

$$\hat{m}(\Psi) = \arg\min_M Q(Y_n, M). \tag{C.2}$$

The study picks $S = 8$, which is within the recommended range.

The structural parameters are then obtained by minimizing a quadratic form of the distance between the data and simulated moments.

$$\hat{\Psi} = \arg\min_\Psi N\left[\hat{M} - \hat{m}(\Psi)\right]' \hat{W}\left[\hat{M} - \hat{m}(\Psi)\right], \tag{C.3}$$

where $\hat{W}$ denotes a positive definite weighting matrix. The value of the above function at the minimum, denoted by $\hat{\Phi}$, provides a goodness-of-fit measure. The optimal weighting matrix is given by

$$\hat{W} = \left[N\text{var}(\hat{M})\right]^{-1}. \tag{C.4}$$

The above covariance matrix is calculated with the actual data set using the influence function method of Erickson and Whited (2000). The estimator is asymptotically normal for fixed S with covariance matrix given by

$$\begin{aligned} \sqrt{N}(\hat{\Psi} - \Psi^*) &\sim \mathsf{N}(0, \Sigma) \\ \Sigma &= (1 + \frac{1}{S}) \left[\frac{\partial^2 Q}{\partial \Psi \partial M'} \left(\frac{\partial Q}{\partial M} \frac{\partial Q}{\partial M}' \right)^{-1} \frac{\partial^2 Q}{\partial M \partial \Psi'} \right]^{-1}. \end{aligned} \tag{C.5}$$

While $\frac{\partial Q}{\partial M}$ can be evaluated analytically, numerical methods are required to obtain $\frac{\partial^2 Q}{\partial \Psi \partial M}$. Both partial derivatives are computed using simulated data evaluated at the data moments.

D Numerical solution

The simulations require a numerical solution of the value function for R&D firms. The capital grid has 101 points and the profitability grid has 21 points. The capital grid is centered around an approximation of the median size of the firm, given the parameters. Simulations which result in steady state firm sizes near the boundaries of the grid are discarded in the estimation. The profit grid is formed using the quadrature method of Tauchen and Hussey (1991), with a mean value obtained by guessing the success rate. The endogenous jumps in z from an innovation are handled by interpolating firm value over two more grids constructed using the transition equation for profitability conditional on whether the firm innovates. The expected value of the firm is obtained using the law of iterated expectations.

The simulated sample is generated using the value and policy functions for R&D firms. The law of motion for profitability is generated directly using the transition equations (3). The firm's decisions are obtained using linear interpolation of the policy functions. The estimation with all firms solves the model assuming $a = 0$ for non-R&D firms. The simulated panel data set for this estimation maintains the same fraction of firms in the R&D sector as in the data. The simulation is run for 100 years, with the initial 50 discarded as a burn-in sample. The value of the quadratic form of the distance between the data moments and simulated moments is computed for each simulation. The program searches for the parameters that minimize this distance using the simulated annealing algorithm. Each estimation involved evaluating more than 50,000 candidate parameter sets and took 2 to 4 weeks of computing time.

E Industry classification

The selected industry groups are constructed based on 4-digit SIC codes. The list of SIC codes included in each industry category is as follows:

- Chips: 3622, 3661-3666, 3669-3679, 3810, 3812
- Hardware: 3570-3579, 3680-3689, 3695
- Medical equipment: 3693, 3811, 3820-3827, 3829-3851
- Pharmaceuticals: 2830, 2831, 2833-2836
- Software: 7370-7373, 7375.

References

Abel, Andrew, and Janice C. Eberly, 1994, A unified model of investment under uncertainty, *American Economic Review* 84, 1369–1384.

Aghion, Philippe, Nick Bloom, Richard Blundell, Rachel Griffith, and Peter Howitt, 2005, Competition and innovation: An inverted-U relationship, *Quarterly Journal of Economics* 120, 701–728.

Aghion, Philippe, and Peter Howitt, 1992, A model of growth through creative destruction, *Econometrica* 60, 323–351.

Akcigit, Ufuk, 2009, Firm size, innovation dynamics, and growth, Working paper, University of Pennsylvania.

Barlevy, Gadi, 2007, On the cyclicality of research and development, *American Economic Review* 97, 1131–1164.

Bertsekas, Dimitri P., 2000, *Dynamic Programming and Optimal Control* vol. 1. (Athena Scientific Belmont, MA).

Bloom, Nick, 2009, The impact of uncertainty shocks, *Econometrica* 77, 623–685.

Bloom, Nick, Rachel Griffith, and John van Reenen, 2002, Do R&D tax credits work? Evidence from a panel of countries 1979-1997, *Journal of Public Economics* 85, 1–31.

Bloom, Nick, Mark Schankerman, and John van Reenen, 2007, Identifying technology spillovers and product market rivalry, NBER working paper no. 13060.

Boldrin, Michele, and David K. Levine, 2008, Perfectly competitive innovation, *Journal of Monetary Economics* 55, 435–453.

Brown, James R., Steven M. Fazzari, and Bruce C. Petersen, 2009, Financing innovation and growth: Cash flow, external equity and the 1990s R&D boom, *Journal of Finance* 64, 151–185.

Chan, Louis K. C., Josef Lakonishok, and Theodore Sougiannis, 2001, The stock market valuation of research and development expenditures, *Journal of Finance* 56, 2431–2456.

Cogley, Timothy, and Thomas J. Sargent, 2008, The market price of risk and the equity premium: A legacy of the Great Depression?, *Journal of Monetary Economics* 55, 454–476.

Comin, Diego, 2004, R&D: A small contribution to productivity growth, *Journal of Economic Growth* 9, 391–421.

Comin, Diego, and Mark Gertler, 2006, Medium-term business cycles, *Americal Economic Review* 96, 523–551.

Cooper, Russell W., and Joao Ejarque, 2003, Financial frictions and investment: Requiem in Q, *Review of Economic Dynamics* 6, 710–728.

Cooper, Russell W., and John Haltiwanger, 2006, On the nature of capital adjustment costs, *Review of Economic Studies* 73, 611–633.

Doraszelski, Ulrich, and Jordi Jaumandreu, 2009, R&D and productivity: Estimating endogenous productivity, Working paper, Harvard University.

Eaton, Jonathan, and Samuel Kortum, 1999, International technology diffusion: Theory and measurement, *International Economic Review* 40, 537–570.

Eberly, Janice, Sergio Rebelo, and Nicolas Vincent, 2008, Investment and value: A neoclassical benchmark, NBER working paper no. 13866.

Erickson, Timothy, and Toni M. Whited, 2000, Measurement error and the relationship between investment and q, *Journal of Political Economy* 108, 1027–1057.

Fama, Eugene F., and Kenneth R. French, 2002, The equity premium, *Journal of Finance* 57, 637–659.

Gilbert, Richard, and Carl Shapiro, 1990, Optimal patent length and breadth, *RAND Journal of Economics* 26, 106–112.

Gomes, Joao F., 2001, Financing investment, *American Economic Review* 90, 1263–1285.

Gourieroux, Christian, Alain Monfort, and Eric Renault, 1993, Indirect inference, *Journal of Applied Econometrics* 8, S85–S118.

Greenwood, Jeremy, and Boyan Jovanovic, 1999, The information-technology revolution and the stock market, *American Economic Association (Papers and Proceedings)* 89, 116–122.

Griliches, Zvi, 1979, Issues in assessing the contribution of research and development to productivity growth, *Bell Journal of Economics* 10, 92–116.

Griliches, Zvi, 1994, Productivity, R&D, and the data constraint, *American Economic Review* 84, 1–23.

Griliches, Zvi, 2000, *R&D, Education and Productivity: A Retrospective.* (Harvard University Press Boston, MA).

Griliches, Zvi, and Jacques Mairesse, 1984, Productivity and R&D at the firm level, in Zvi Griliches, eds.: *R & D, Patents, and Productivity* (University of Chicago Press, Chicago, IL).

Grossman, Gene M., and Elhanan Helpman, 1991, *Innovation and Growth in the Global Economy.* (MIT Press Boston, MA).

Hall, Bronwyn H., 2007, Measuring the returns to R&D: The depreciation problem, NBER working paper no. 13473.

Hall, Bronwyn H., and Jacques Mairesse, 1995, Exploring the relationship between R&D and productivity in French manufacturing firms, *Journal of Econometrics* 65, 263–293.

Hall, Bronwyn H., and John van Reenen, 2000, How effective are fiscal incentives for R&D? A review of the evidence, *Research Policy* 29, 449–469.

Hall, Robert E., 2001, The stock market and capital accumulation, *American Economic Review* 91, 1185–1202.

Hennessy, Christopher A., and Toni M. Whited, 2005, Debt dynamics, *Journal of Finance* 60, 1129–1165.

Hennessy, Christopher A., and Toni M. Whited, 2007, How costly is external financing? Evidence from a structural estimation, *Journal of Finance* 62, 1705–1745.

Hobijn, Bart, and Boyan Jovanovic, 2001, The information-technology revolution and the stock market: Evidence, *American Economic Review* 91, 1203–1220.

Jones, Charles I., 1995a, R&D-based models of economic growth, *Journal of Political Economy* 103, 759–784.

Jones, Charles I., 1995b, Time series tests of endogenous growth models, *Quarterly Journal of Economics* 110, 495–525.

Klette, Tor Jacob, and Samuel Kortum, 2004, Innovating firms and aggregate innovation, *Journal of Political Economy* 112, 986–1018.

Kortum, Samuel, 1997, Research, patenting and technological change, *Econometrica* 65, 1389–1419.

Lentz, Rasmus, and Dale T. Mortensen, 2008, An empirical model of growth through innovation, *Econometrica* 76, 1317–1373.

Li, Dongmei, 2008, Financial constraints, R&D investment, and stock returns: Theory and evidence, Working paper, University of California, San Diego.

Moris, Francisco, 2005, The U.S. research and experimentation tax credit in the 1990s, *National Science Foundation*, InfoBrief NSF05-316, Arlington, VA.

Okubo, Sumiye, Carol A. Robbins, Carol E. Moylan, Brian K. Selzer, Laura I. Schultz, and Lisa S. Mataloni, 2006, R&D satellite account: Preliminary estimates, Working paper, Bureau of Economic Analysis.

Romer, Paul M., 1990, Endogenous technological change, *Journal of Political Economy* 98, S71–S102.

Salinger, Michael A., and Lawrence H. Summers, 1983, Tax reform and corporate investment: A microeconometric simulation study, in Martin Feldstein, eds.: *Behavioral simulation methods in tax policy analysis* (University of Chicago Press, Chicago, IL).

Segerstrom, Paul S., 1998, Engogenous growth without scale effects, *American Economic Review* 88, 1290–1310.

Tauchen, George, and Robert Hussey, 1991, Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models, *Econometrica* 59, 371–396.

Thompson, Peter, 2001, The microeconomics of an R&D-based model of endogenous growth, *Journal of Economics Growth* 5, 263–283.

Young, Alywn, 1998, Growth without scale effects, *Journal of Political Economy* 106, 41–63.

Table 1: Summary statistics

Panel A reports summary statistics for firms classified by their R&D investment into R&D firms and non-R&D firms. Panel B reports summary statistics for firms in 5 selected industries, categorized by the 4-digit SIC codes detailed in Appendix E. 'Medical eq.' and 'Pharma.' denote the medical equipment and pharmaceutical industries, respectively. The baseline definition of R&D firms in Panel A includes all firm-year observations with positive R&D expenditures. The non-R&D firms consists of all firm-year observations with missing or zero R&D expenditures. The alternate definition of the R&D firms includes all firms that report non-missing R&D expenditures in each year with positive values in some. The data are at an annual frequency. The source data are from the Compustat Industrial files from 1987 to 2006. All variables are scaled by the book value of assets of the firm and Winsorized at the 1% level. Tobin's Q is defined as the market value of assets scaled by the book value of assets. Profitability equals operating income before depreciation scaled by book assets. R&D expenditure is obtained from Compustat data item 46, and treated as an expense in calculating profitability. Investment equals capital expenditures on plant, property, and equipment net of retirements scaled by book assets. The standard deviations adjust for firm and year fixed effects within each group of firms.

Panel A: Economy-wide data

	R&D firms		Non-R&D firms		R&D firms (alternate)	
Variable	Mean	Std.	Mean	Std.	Mean	Std.
Log assets	5.034	2.131	5.193	1.860	4.993	2.109
Tobin's Q	2.174	1.224	1.528	0.680	2.244	1.267
Profitability	0.072	0.146	0.144	0.100	0.067	0.150
R&D to assets	0.112	0.082			0.120	0.086
Investment to assets	0.063	0.062	0.094	0.089	0.062	0.076
Number of observations	35409		38207		31344	

Panel B: Selected industry data

	Chips		Hardware		Medical eq.		Pharma.		Software	
Variable	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.
Log assets	4.765	1.894	4.605	1.956	4.161	1.607	4.566	1.967	4.356	1.554
Tobin's Q	2.078	1.212	2.169	1.300	2.383	1.261	3.417	1.669	2.836	1.585
Profitability	0.101	0.141	0.074	0.166	0.078	0.148	−0.127	0.187	0.047	0.190
R&D to assets	0.119	0.073	0.145	0.080	0.116	0.073	0.261	0.144	0.183	0.111
Investment to assets	0.074	0.071	0.058	0.054	0.054	0.053	0.059	0.069	0.058	0.063
Number of observations	4832		2617		4321		3790		4567	

Table 2: Production function regressions

The table reports the results of a regression of the log of value-added on logs of labor, capital, and the stock of R&D. Panel A reports the results of pooled regressions, and Panel B reports the results of fixed effects regressions. The regressors include year dummies. The pooled regressions also include industry dummies. The point estimates correspond to the estimated elasticities of value-added with respect to the inputs. The columns vary the assumed obsolescence rate in the construction of the R&D stock. The standard errors are heteroskedasticity robust and adjust for clustering at the firm level. The R&D stock variable is constructed using a perpetual inventory method. The data are from the Compustat Industrial file from 1987 to 2006. All variables, except labor, are adjusted for inflation using the GDP deflator for business investment. Computation of the estimated rate of return to R&D, which equals $\nu\frac{Y}{S}$, uses the estimated elasticity ν and the sample median of value added scaled by the R&D stock, which varies with the assumed obsolescence rate γ.

Panel A: Pooled regressions

	Obsolescence rate				
Regressor	$\gamma = 0$	$\gamma = 0.05$	$\gamma = 0.10$	$\gamma = 0.15$	$\gamma = 0.25$
Labor	0.682	0.678	0.675	0.673	0.669
	(0.017)	(0.017)	(0.017)	(0.017)	(0.017)
Capital stock	0.173	0.158	0.151	0.148	0.147
	(0.014)	(0.014)	(0.014)	(0.014)	(0.014)
R&D stock	0.183	0.204	0.213	0.219	0.223
	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)
Number of observations	39523	39523	39523	39523	39523
Adjusted R-squared	0.892	0.893	0.894	0.895	0.896
Estimated rate of return ($= \nu\frac{Y}{S}$)	11.8%	22.8%	33.1%	43.1%	62.1%
Total payoff to R&D ($= \nu\frac{Y}{S} + 1 - \gamma$)	1.12	1.18	1.23	1.28	1.37

Panel B: Fixed effects regressions

	Obsolescence rate				
Regressor	$\gamma = 0$	$\gamma = 0.05$	$\gamma = 0.10$	$\gamma = 0.15$	$\gamma = 0.25$
Labor	0.773	0.773	0.773	0.773	0.771
	(0.016)	(0.016)	(0.016)	(0.016)	(0.016)
Capital stock	0.133	0.129	0.127	0.124	0.118
	(0.015)	(0.016)	(0.016)	(0.016)	(0.016)
R&D stock	0.063	0.057	0.055	0.056	0.061
	(0.019)	(0.017)	(0.016)	(0.016)	(0.015)
Number of observations	39523	39523	39523	39523	39523
Adjusted R-squared	0.874	0.874	0.874	0.874	0.876
Estimated rate of return ($= \nu\frac{Y}{S}$)	4.1%	6.3%	8.5%	11.0%	17.1%
Total payoff to R&D ($= \nu\frac{Y}{S} + 1 - \gamma$)	1.04	1.01	0.98	0.96	0.92

Table 3: R&D firms

Panel A reports the matched moments from the actual and simulated data sets for the model estimated on R&D firms. Panel B reports the estimated structural parameters and a goodness-of-fit statistic. Panel C presents some key statistics from the simulation detailed in Section 2.5. The estimation is carried out using simulated method of moments. Tobin's Q is defined as the market value of assets scaled by the book value of assets. Profitability equals operating income before depreciation scaled by book assets. R&D expenditure is obtained from Compustat data item 46 and treated as an expense in calculating profits. Investment equals capital expenditures on plant, property, and equipment net of retirements. Tobin's Q for the simulated data includes the value of the R&D stock in the numerator.

Panel A: Moments

Moment	Data	Model
First moment of R&D to assets	0.112	0.115
Second moment of R&D to assets	0.019	0.021
First moment of Tobin's Q	2.174	2.153
Second moment of Tobin's Q	6.224	5.842
First moment of profitability	0.072	0.094
Second moment of profitability	0.027	0.040
Autocorrelation of profitability	0.435	0.917
Autocorrelation of R&D to assets	0.289	0.481
Second moment of investment to assets	0.008	0.013

Panel B: Parameter estimates

Parameter	θ	ρ	σ	λ	a	b	γ	$\hat{\Phi}$
Estimate	0.785	0.887	0.285	0.152	3.761	4.185	0.319	8494
Standard error	(0.012)	(0.001)	(0.003)	(0.003)	(0.045)	(0.154)	(0.009)	

Panel C: Economic impact

Variable of interest	5th Percentile	Mean	95th Percentile
Expected rate of return to R&D	2.0%	3.3%	5.4%
Probability of innovation	0.121	0.669	0.884
Expected increase in firm value from an innovation	10.4%	16.4%	26.8%
Equivalent permanent profitability increase	-	4.28%	-
Production function method rate of return	-	54.4%	-

Table 4: All firms

Panel A reports the matched moments from the actual and simulated data sets for the model estimated on R&D and non-R&D firms. Panel B reports the estimated structural parameters and a goodness-of-fit statistic. Panel C presents some key statistics from the simulation detailed in Section 2.5. The estimation is carried out using simulated method of moments. Tobin's Q is defined as the market value of assets scaled by the book value of assets. Profitability equals operating income before depreciation scaled by book assets. R&D expenditure is obtained from Compustat data item 46 and treated as an expense in calculating profits. Investment equals capital expenditures on plant, property, and equipment net of retirements. Tobin's Q for the simulated data includes the value of the R&D stock in the numerator. The autocorrelation coefficient for profitability accounts for differences in mean z across the R&D and non-R&D firms in the simulated data.

Panel A: Moments

Moment	Data	Model
First moment of R&D to assets	0.112	0.104
Second moment of R&D to assets	0.019	0.021
First moment of Tobin's Q (R&D firms)	2.174	1.895
First moment of Tobin's Q (non-R&D firms)	1.528	1.569
Second moment of Tobin's Q (R&D firms)	6.224	4.147
Second moment of Tobin's Q (non-R&D firms)	2.799	2.556
First moment of profitability (R&D firms)	0.072	0.106
First moment of profitability (non-R&D firms)	0.144	0.165
Second moment of profitability (R&D firms)	0.027	0.041
Second moment of profitability (non-R&D firms)	0.031	0.040
Autocorrelation of profitability	0.469	0.859
Autocorrelation of R&D to assets	0.289	0.423
Second moment of investment to assets	0.012	0.017

Panel B: Parameter estimates

Parameter	θ	ρ	σ	λ	a	b	γ	$\hat{\Phi}$
Estimate	0.859	0.820	0.333	0.182	3.450	2.863	0.317	22650
Standard error	(0.003)	(0.001)	(0.002)	(0.001)	(0.032)	(0.014)	(0.003)	

Panel C: Economic impact

Variable of interest	5th Percentile	Mean	95th Percentile
Expected rate of return to R&D	2.4%	4.6%	7.1%
Probability of innovation	0.025	0.578	0.855
Expected increase in firm value from an innovation	9.6%	12.7%	16.0%
Equivalent permanent profitability increase	-	3.64%	-
Production function method rate of return	-	48.8%	-
Ratio of aggregate R&D value to the value of all capital	-	28.2%	-

Table 5: Selected industries

Panel A reports the matched first moments of R&D to assets, Tobin's Q, and profitability from the actual and simulated data for 5 selected industries; the other matched moments are not reported. Panel B reports the estimated R&D related structural parameters and a goodness-of-fit statistic. Panel C presents some key statistics from the simulation detailed in Section 2.5. The estimation is carried out using simulated method of moments. Tobin's Q is defined as the market value of assets scaled by the book value of assets. Profitability equals operating income before depreciation scaled by book assets. R&D expenditure is obtained from Compustat data item 46 and treated as an expense in calculating profits. Investment equals capital expenditures on plant, property, and equipment net of retirements. Tobin's Q for the simulated data includes the value of the R&D stock in the numerator. 'Medical eq.' and 'pharma.' denote medical equipment and pharmaceuticals, respectively.

Panel A: First Moments

Industry	R&D to assets		Tobin's Q		Profitability	
	Data	Model	Data	Model	Data	Model
Chips	0.119	0.121	2.07	2.03	0.101	0.118
Hardware	0.145	0.145	2.17	2.19	0.074	0.106
Medical equipment	0.116	0.116	2.38	2.46	0.070	0.107
Pharmaceutical	0.261	0.238	3.42	3.23	−0.127	−0.026
Software	0.183	0.189	2.84	2.82	0.047	0.096

Panel B: Parameter estimates

Industry	λ	α	γ	$\hat{\Phi}$
Chips	0.152	3.788	0.344	2183
	(0.012)	(0.159)	(0.041)	
Hardware	0.192	3.090	0.277	640
	(0.009)	(0.092)	(0.008)	
Medical equipment	0.110	4.178	0.369	1055
	(0.005)	(0.114)	(0.034)	
Pharmaceutical	0.254	3.044	0.409	5293
	(0.006)	(0.067)	(0.015)	
Software	0.146	2.196	0.366	1877
	(0.003)	(0.091)	(0.024)	

Panel C: Economic impact

Variable of interest	Chips	Hardware	Medical eq.	Pharma.	Software
Probability of innovation	0.687	0.735	0.669	0.730	0.594
Expected rate of return to R&D	4.5%	3.1%	2.9%	2.7%	2.8%
Increase in firm value from an innovation	19.7%	23.4%	16.1%	34.3%	20.5%
Equivalent permanent profitability increase	3.9%	6.2%	4.3%	8.3%	5.0%
Production function method rate of return	79.5%	51.3%	56.9%	27.7%	31.0%

Table 6: R&D firms (alternate)

Panel A reports the matched moments from the actual and simulated data sets for the model estimated on R&D firms under the alternate definition. Panel B reports the estimated structural parameters and a goodness-of-fit statistic. Panel C presents some key statistics from the simulation detailed in Section 2.5. The estimation is carried out using simulated method of moments. Tobin's Q is defined as the market value of assets scaled by the book value of assets. Profitability equals operating income before depreciation scaled by book assets. R&D expenditure is obtained from Compustat data item 46 and treated as an expense in calculating profits. Investment equals capital expenditures on plant, property, and equipment net of retirements. Tobin's Q for the simulated data includes the value of the R&D stock in the numerator.

Panel A: Moments

Moment	Data	Model
First moment of R&D to assets	0.120	0.115
Second moment of R&D to assets	0.022	0.022
First moment of Q	2.244	2.285
Second moment of Q	6.640	7.265
First moment of profitability	0.067	0.095
Second moment of profitability	0.027	0.035
Autocorrelation of profitability	0.439	0.897
Autocorrelation of R&D to assets	0.285	0.532
Second moment of investment to assets	0.008	0.011

Panel B: Parameter estimates

Parameter	θ	ρ	σ	λ	a	b	γ	$\hat{\Phi}$
Estimate	0.791	0.936	0.192	0.144	3.390	4.589	0.233	6027
Standard error	(0.018)	(0.001)	(0.002)	(0.007)	(0.062)	(0.363)	(0.005)	

Panel C: Economic impact

Variable of interest	5th Percentile	Mean	95th Percentile
Expected rate of return to R&D	1.2%	2.1%	3.6%
Probability of innovation	0.000	0.712	0.923
Expected increase in firm value from an innovation	11.1%	18.6%	35.7%
Equivalent permanent profitability increase	-	5.94%	-
Production function method rate of return	-	36.2%	-

Table 7: Counterfactual experiment on a subsidy to R&D expenditures

The table reports selected statistics from a counterfactual experiment on an additional tax subsidy to R&D expenditures. Panel A reports the results based on the estimation using R&D firms, Panel B reports the results based on the estimation using all firms, and Panel C reports the results based on the estimation using the alternate definition of R&D firms. The estimate row reports the moments with no additional subsidy on R&D expenditures, and the experiment row reports the results with an additional 5% tax subsidy. s denotes the size of the additional tax subsidy. The statistics reported consist of the steady state mean of R&D expenditures to assets, profitability, Tobin's Q, and the probability of innovation. The structural parameters for each model are reported in Panel B of Tables 3, 4, and 6, respectively.

Panel A: R&D firms

	s	R&D to assets	Profitability	Tobin's Q	$\bar{j}$
Estimate	0.000	0.115	0.094	2.153	0.669
Experiment	0.050	0.121	0.089	2.301	0.697

Panel B: All firms

	s	R&D to assets	Profitability	Tobin's Q	$\bar{j}$
Estimate	0.000	0.104	0.106	1.895	0.578
Experiment	0.050	0.108	0.101	2.053	0.600

Panel C: R&D firms (alternate)

	s	R&D to assets	Profitability	Tobin's Q	$\bar{j}$
Estimate	0.000	0.115	0.095	2.285	0.712
Experiment	0.050	0.124	0.074	2.431	0.770

Table 8: Counterfactual experiment on an increase in the impact of an innovation

The table reports selected statistics from a counterfactual experiment on an increase in the impact of an innovation. Panel A reports the results based on the estimation using R&D firms, Panel B reports the results based on the estimation using all firms, and Panel C reports the results based on the estimation using the alternate definition of R&D firms. The estimate row reports the moments with the estimated value for λ, and the experiment row reports the results with a λ value equal to 1.05 times the estimate. The statistics reported consist of the steady state mean of R&D expenditures to assets, profitability, Tobin's Q, and the probability of innovation. The structural parameters for each model are reported in Panel B of Tables 3, 4, and 6, respectively.

Panel A: R&D firms

	λ	R&D to assets	Profitability	Tobin's Q	$\bar{j}$
Estimate	0.152	0.115	0.094	2.153	0.669
Experiment	0.159	0.122	0.089	2.507	0.702

Panel B: All firms

	λ	R&D to assets	Profitability	Tobin's Q	$\bar{j}$
Estimate	0.182	0.104	0.106	1.895	0.578
Experiment	0.191	0.110	0.096	2.261	0.603

Panel C: R&D firms (alternate)

	λ	R&D to assets	Profitability	Tobin's Q	$\bar{j}$
Estimate	0.144	0.115	0.095	2.285	0.712
Experiment	0.151	0.125	0.067	2.540	0.788

Figure 1: Aggregate R&D investment as a fraction of total investment

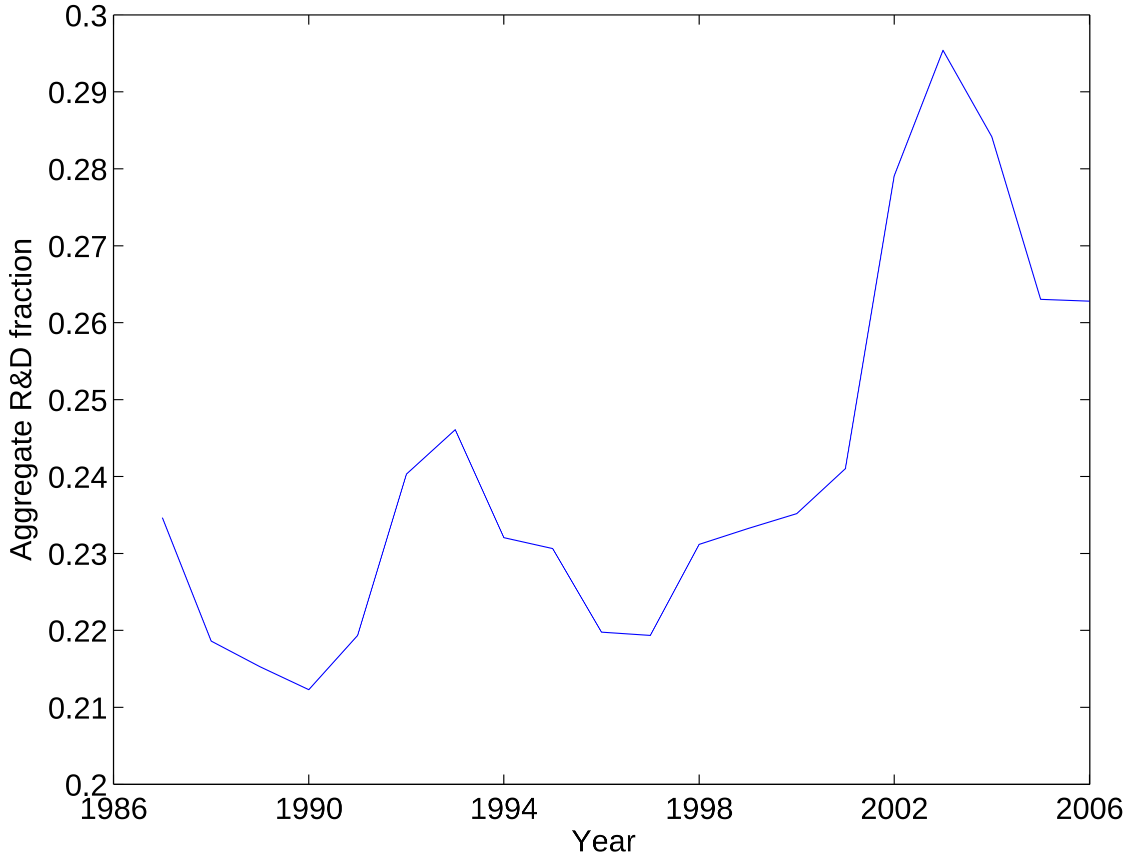

The figure plots the time series of aggregate R&D investment as a fraction of the sum of aggregate R&D and physical investment. These aggregates are obtained from the sample used to estimate the model detailed in Section 3.1. The physical investment aggregate includes investment by both R&D firms and non-R&D firms.

Figure 2: Mean R&D investment and profitability

Panel A: Variation with λ

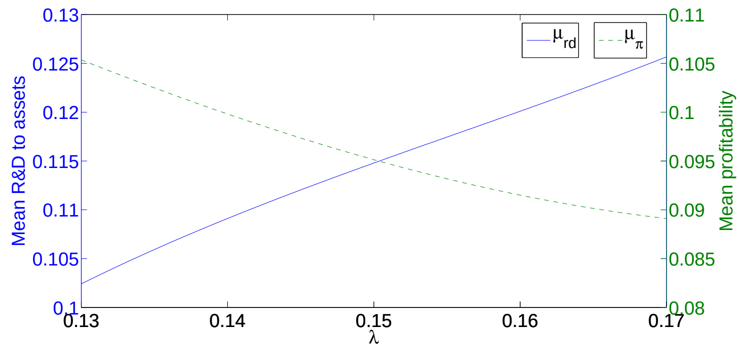

Panel B: Variation with a

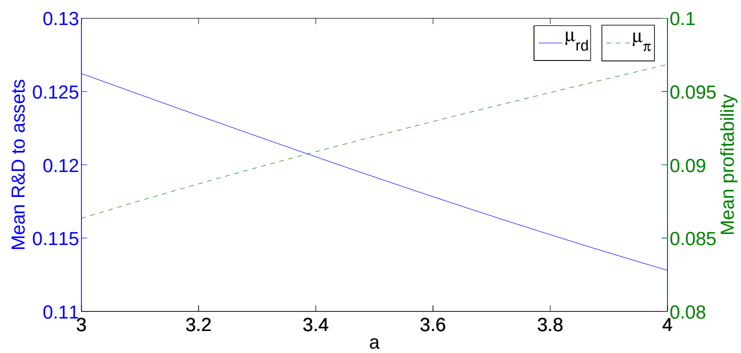

Panel C: Variation with γ

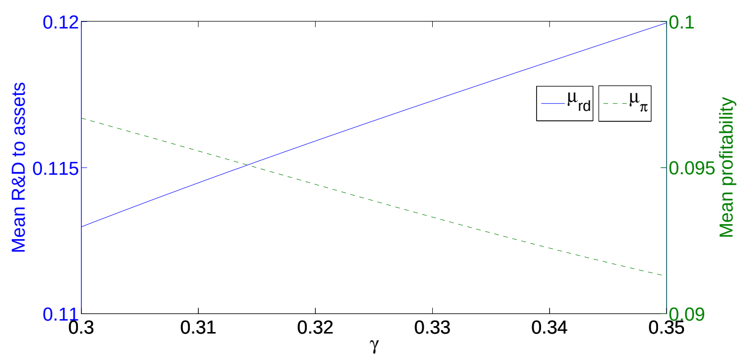

Notes: The figures plot the first moment of R&D and profitability as each of the R&D related structural parameters vary while keeping the others fixed at their point estimates reported in Panel B of Table 3. Panels A, B, and C plot mean R&D to assets (left axis) and mean profitability (right axis) as λ, a, and γ varies, respectively. The solid line denotes mean R&D to assets (the left axis variable) and the dashed line denotes mean profitability (the right axis variable). The plotted values are obtained by solving the model for 100 evenly spaced values of λ, a, and γ, respectively.