Discussion Paper No. 13-089

Estimating Dynamic R&D Demand: An Analysis of Costs and Long-Run Benefits

Bettina Peters, Mark J. Roberts, Van Anh Vuong, and Helmut Fryges

ZEW

Zentrum für Europäische
Wirtschaftsforschung GmbH

Centre for European
Economic Research

Discussion Paper No. 13-089

Estimating Dynamic R&D Demand: An Analysis of Costs and Long-Run Benefits

Bettina Peters, Mark J. Roberts, Van Anh Vuong, and Helmut Fryges

Download this ZEW Discussion Paper from our ftp server:

http://ftp.zew.de/pub/zew-docs/dp/dp13089.pdf

Die Discussion Papers dienen einer möglichst schnellen Verbreitung von neueren Forschungsarbeiten des ZEW. Die Beiträge liegen in alleiniger Verantwortung der Autoren und stellen nicht notwendigerweise die Meinung des ZEW dar.

Discussion Papers are intended to make results of ZEW research promptly available to other economists in order to encourage discussion and suggestions for revisions. The authors are solely responsible for the contents which do not necessarily represent the opinion of the ZEW.

Non-technical Summary

Estimating the private return to R&D investment has been a major goal for decades and most of the empirical literature has been built around the knowledge production function. In this framework, investment in R&D creates a stock of knowledge within the firm that enters into the firm's production function as an additional input factor. Estimates of the effect of this knowledge stock on output provides a measure of the ex post return to the firm's investment in R&D.

This paper provides a different approach to measuring the private payoff from R&D investment. We develop and estimate a dynamic, structural model of the firm's decision to invest in R&D and quantify the cost and long-run benefit of this investment. The dynamic model incorporates and quantifies linkages between the firm's R&D investment, product and process innovations, and future productivity and profits. And it provides a natural measure of the long-run payoff to R&D as the difference between the expected discounted sum of future profits if the firm undertakes R&D versus if it does not. The firm will choose to invest in R&D if this payoff is greater than the fixed or sunk cost they pay to invest in R&D.

We use firm-level data from the Mannheim Innovation Panel (MIP) for German manufacturing industries to estimate the dynamic structural model and to calculate the long-run payoffs to R&D. Comparing across industries for the firm with the median productivity level, we find that the expected long-run benefit of investing in R&D varies from a high of 43 million euros in the vehicle industry and 20 million euros in the chemical industry to a low of about 350 thousand euros in the plastic, non-metallic mineral products, and manufacturing nec industries. By combining estimates of the expected long-run benefit of R&D with the cost of R&D, we also estimate the distribution of net benefits across firms in each industry. We find that the expected net benefit varies substantially across industries and across firms that have already invested in R&D and those that are just starting R&D investments because of the substantial differences in the fixed versus sunk costs. Expressed as a proportion of firm value, our results show, for instance, that the net benefit for the median firm with prior R&D experience varies from 2.4 to 3.2 percent across five high-tech industries but varies from -4.6 to 0.6 percent for firms with no previous R&D experience. The negative value implies that the median inexperienced firm would not find it profitable to invest in R&D. Given unexperienced firms find R&D investment profitable and start performing R&D, we estimate a net benefit of 2.0 to 2.4 percent in the high tech industries. These net benefits are substantially smaller, around 0.2 percent in low-tech industries.

The estimated dynamic structural model of R&D demand can be used to simulate how a

change in the cost structure of R&D arising from, for example, a tax break or direct subsidy for R&D investment, affects the firm's investment choice and future productivity growth. We find that, in high-tech industries, a 20 percent reduction in the fixed cost of R&D leads after five years to an average increase of 7 percentage points in the probability a firm invests in R&D and a 4 percent increase in mean productivity. A similar reduction in the cost faced by firms just beginning to invest in R&D, however, has very little impact on the probability of investing or the level of productivity. That is, the two cost changes have very different impacts on firm incentives. Fixed cost reductions encourage all firms to invest. In contrast, the reduction in startup costs encourages new firms to begin investing but also reduces the option value of investing leading some firms to stop their R&D.

Das Wichtigste in Kürze

Die Schätzung der privaten Erträge aus Investitionen in Forschung und Entwicklung (FuE) steht seit langem im Fokus vieler empirischer Arbeiten. Als Ansatz wird zumeist eine Wissensproduktionsfunktion verwendet. In diesem Ansatz führen FuE-Investionen zu einer Erhöhung des firmeninternen Wissenskapitalstocks, der wiederum als ein Inputfaktor in die Produktionsfunktion eines Unternehmens eingeht. Die Schätzung des Effekts des Wissenskapitalstocks auf den Output eines Unternehmens stellt ein Maß für die Erträge aus FuE dar.

Dieses Papier stellt einen neuen Ansatz vor, die privaten Erträge aus Investitionen in FuE zu messen. Wir entwickeln und schätzen ein dynamisches strukturelles Modell der Entscheidung eines Unternehmens in FuE zu investieren, das konsistent ist mit einer Maximierung der langfristig erwarteten Nettoerträge aus dieser Investition. Das Modell berücksichtigt, dass die Entscheidung eines Unternehmens in FuE zu investieren, die Wahrscheinlichkeit für zukünftige Produkt- und Prozessinnovationen beeinflusst und sich die Einführung von Produkt- und Prozessinnovationen wiederum auf die zukünftige Produktivität und die Gewinne des Unternehmens auswirkt. Das Modell erlaubt es somit, den Einfluss der FuE-Entscheidung auf den Firmenwert (abdiskontierte Summe aller zukünftigen Gewinne) zu identifizieren. Die Differenz zwischen den Firmenwerten, wenn das Unternehmen in FuE investiert und wenn es nicht investiert, ist ein Maß für die langfristigen Erträge aus FuE. Ein Unternehmen wird sich für FuE-Aktivitäten entscheiden, wenn die langfristigen Erträge größer als die damit verbundenen Kosten sind. Dabei erlaubt das Modell, dass sich die Kosten im Falle einer Aufnahme von FuE-Aktivitäten (*Sunk Costs*) von denen bei Fortsetzung von FuE-Aktivitäten (*Fixkosten*) unterscheiden.

Wir verwenden Daten des Mannheimer Innovationspanels für das deutsche verarbeitende Gewerbe, um das Modell zu schätzen und die langfristigen Erträge zu berechnen. Unsere Ergebnisse zeigen eine hohe Variation in den erwarteten Erträgen aus FuE zwischen und innerhalb von Industrien. So reichen die erwarteten Erträge für ein Medianunternehmen (gemessen anhand seiner Produktivität) von 43 Millionen Euro in der Automobilindustries, über 20 Millionen Euro in der Chemischen Industrie bis zu rund 350 Tausend Euro in der Gummi-/Kunststoffverarbeitung oder in der Glas/Keramik/Steinwaren-Industrie. Berücksichtigt man neben den erwarteten langfristigen Erträgen aus FuE auch deren Kosten, dann lassen sich mittels des Modells auch die Nettoerträge aus FuE schätzen. Die Ergebnisse zeigen auch hier eine große Variation zwischen den Industrien sowie zwischen Firmen, die bereits in der Vorperiode FuE durchgeführt haben und solchen, die auf FuE in der Vorperiode verzichtet haben. Dies liegt auch darin begründet, dass die geschätzten Fixkosten deutlich geringer sind als die Sunk costs. So variieren die Nettoerträge, gemessen als Anteil am Firmenwert, für die Medianunternehmen in den fünf Hightech-Industrien zwischen 2.4 und 3.2 Prozent, wenn das Unternehmen FuE in der Vorperiode durchgefürt hat. Ohne FuE-Erfahrung liegen die Nettoerträge dagegen zwischen -4.6 to 0.6 Prozent. Ein negativer Wert bedeutet, dass in dieser Industrie das Medianunternehmen es als nicht profitabel erachtet in FuE zu investieren. Betrachtet man innerhalb der Gruppe der Unternehmen ohne FuE-Erfahrung nur solche, die FuE profitabel finden und daher FuE-Aktivitäten neu aufnehmen, dann erzielen sie Nettoerträge von 2.0 bis 2.4 Prozent. Diese Nettoerträge sind in den sieben Lowtech-Industrien deutlich niedriger und liegen bei rund 0.2 Prozent.

Das geschätzte dynamische strukturelle Modell kann genutzt werden, um kontrakfaktische Politiksimulationen durchzuführen. Zum Beispiel kann analysiert werden, wie sich eine Reduktion der FuE-Kosten als Folge einer FuE-Subvention auf die Entscheidung eines Unternehmens in FuE zu investieren und auf das zukünftige Produktivitätswachstum auswirkt. Unsere Ergebnisse zeigen, dass eine Reduktion der Fixkosten um 20 Prozent in den Hightech-Industrien nach 5 Jahren zu einem Anstieg der FuE-Beteiligung um etwa 7 Prozentpunkte und zu einer Zunahme der durchschnittlichen Produktivität um 4 Prozent führt. Dagegen hat eine Reduktion der Sunk Costs um 20 Prozent nur geringe Auswirkungen auf beide Größen. Beide Politikmaßnahmen haben somit sehr unterschiedliche Auswirkungen auf die FuE-Entscheidung eines Unternehmens. Während geringere Fixkosten für alle Unternehmen einen größeren Anreiz darstellt in FuE zu investieren, ist dies bei der Reduktion der Sunk Costs nicht der Fall. Einige Unternehmen werden zwar dadurch FuE-Aktivitäten aufnehmen, gleichzeitig sinkt der Optionswert für Unternehmen mit FuE-Erfahrung und ein Teil dieser Unternehmen stellt daher FuE-Aktivitäten ein.

Estimating Dynamic R&D Demand: An Analysis of Costs and Long-Run Benefits*

Bettina Peters†
Centre for European Economic Research (ZEW) and MaCCI

Mark J. Roberts
Pennsylvania State University and NBER

Van Anh Vuong
University of Cologne and Institute of Energy Economics

Helmut Fryges
Australian Innovation Research Centre, University of Tasmania

October 2013

Abstract

Using firm-level data from the German manufacturing sector, we estimate a dynamic, structural model of the firm's decision to invest in R&D and quantify the cost and long-run benefit of this investment. The model incorporates and quantifies linkages between the firm's R&D investment, product and process innovations, and future productivity and profits. The dynamic model provides a natural measure of the long-run payoff to R&D as the difference in expected firm value generated by the R&D investment. For the median productivity firm, investment in R&D raises firm value by 3.0 percent in a group of high-tech industries but only 0.2 percent in low-tech industries. Simulations of the model show that cost subsidies for R&D can significantly affect R&D investment rates and productivity changes in the high-tech industries.

Keywords: R&D demand, Innovation, Productivity, Dynamic structural model

JEL-Classification: L60, O31, O32

*We are grateful to Uli Doraszelski, Ken Judd, Jacques Mairesse, Joris Pinkse, Spiro Stefanou, Jim Tybout, and Hongsong Zhang for helpful comments and discussions. We thank the Center for European Economic Research (ZEW) for providing data access and research support.

†For further information on projects of the author see www.zew.de/staff_bpe as well as the ZEW annual report on www.zew.de/en

1 Introduction

Firm investment in R&D is a key mechanism generating improvements in firm performance over time. Estimating the ex post return to the firm's investment has been a major focus of empirical studies for decades and most of the empirical literature has been built around the knowledge production function developed by Griliches (1979). In this framework, firm investment in R&D creates a stock of knowledge within the firm that enters into the firm's production function as an additional input along with physical capital, labor, and materials. The marginal product of this knowledge input provides a measure of the return to the firm's investment in R&D and has been the focus of the empirical innovation literature.[1]

The goal of this paper is to estimate the payoff to R&D investment at the firm level. However, rather than focusing directly on how R&D impacts the production function, we focus on the firm's demand for R&D. This demand contains information on both the costs of R&D investment to the firm and, importantly for our purposes, the expected long-run payoff to the firm of undertaking R&D investment. We develop a dynamic structural model of the firm's demand for R&D, estimate it using micro data on German manufacturing firms, and summarize the implicit long-run payoff to R&D which rationalizes the firm's observed investment decision. Using this model we estimate the expected return which a profit maximizing firm faces when it makes its R&D investment decision.

Our model of the firm's dynamic demand for R&D captures several important features of the R&D investment process. First is the impact of R&D on the probability that the firm realizes a product or process innovation. Second is the effect of these realized innovations on the firm's productivity and short-run profitability. Third, these effects can be long-lived affecting the incentives of the firm to invest in the future and impacting the long-run value of the firm. Fourth, there is uncertainty about both the effect of R&D on innovation and the effect of innovation on productivity. Fifth, the cost of investing in R&D is likely to differ between firms that are spending to maintain ongoing R&D activities and firms that are just establishing new R&D activities. The structural parameters estimated in the model characterize the linkages

[1] See Hall, Mairesse, and Mohnen (2010) for a recent survey of the empirical studies using the knowledge production function framework.

between R&D, innovation, and productivity as well as the costs of starting or maintaining an R&D program.

We use the model to estimate the long-run payoff to R&D for a sample of German manufacturing firms in a range of high-tech and low-tech industries. The data source is the Mannheim Innovation Panel (MIP) collected by the Centre for European Economic Research (ZEW), which is the German contribution to the Community Innovation Survey (CIS) that is collected for most OECD countries. A unique aspect of the CIS data is that it includes survey questions on the product and process innovations realized by the firm as well as measures of R&D expenditure and variables to construct firm productivity. Because it distinguishes product and process innovation it will allow us to separate the effects of different innovation types on firms' performance.

The structural estimates can be briefly summarized. First, firms that invest in R&D have a substantially higher probability of realizing a product or process innovation but R&D investment is neither necessary nor sufficient for firm innovation. The group of high-tech manufacturing industries has a higher probability of innovation, given R&D, than the group of low-tech industries. Second, product innovation as well as process innovation lead to increases in future firm productivity but product innovations are more important in the high-tech industries while process innovations are more important for the low-tech industries. Third, firm productivity is highly persistent over time which implies that innovations that raise productivity will have long-run effects on firm performance. Fourth, fixed costs of maintaining ongoing R&D investment are significantly smaller than the sunk startup costs of beginning to invest in R&D. This means that firm R&D history is an important determinant of current R&D behavior.

Using the structural estimates we construct the expected payoff to firm R&D as the difference in the expected future value of the firm if it chooses to invest in R&D versus it does not. This varies with the productivity and size of the firm and can be constructed for all firms, not just firms that choose to invest. We find that the expected payoff, net of the cost of R&D, varies substantially across industries and across firms within the industry. For the five high-tech industries, the median firm with prior R&D experience has a net payoff equal to 3.0 percent of firm value. In contrast, the net payoff is often negative for firms in the low-tech industries or firms that must pay startup costs when they begin investing in R&D. In the seven low-tech industries the median firm with prior R&D experience has a negative expected net payoff, averaging -0.6 percent of firm value, which implies that it would not find it profitable to invest in R&D. We also compare the impact that R&D investment has on the short-run profits of the firm versus long-run firm value and find that the short-run payoff is a small fraction of the total benefit. The one-period payoff averages only 2 percent of the total increase in firm value for the median firm in the high-tech industries and 9 percent in the low-tech industries. This emphasizes the need to measure the return to R&D in a dynamic framework that captures the effect of the investment on long-run firm value.

The estimated dynamic structural model of R&D demand can be used to simulate how a change in the cost structure of R&D arising from, for example, a tax break or direct subsidy for R&D investment, affects the firm's investment choice and future productivity growth. We find that, in the high-tech industries, a 20 percent reduction in the fixed cost of R&D leads after five years to an average increase of 7.16 percentage points in the probability a firm invests in R&D and a 4.06 percent increase in mean productivity. A similar reduction in the cost faced by firms just beginning to invest in R&D has very little impact on the probability of investing or the level of productivity. The simulations also illustrate that the two cost changes have very different impacts on firm incentives. Fixed cost reductions encourage firms to continue investing if they already were or to begin investing if they were not. In contrast the reduction in startup costs encourages new firms to begin investing but also reduces the option value of investing leading some firms to stop their R&D.

In the next section of the paper we review some key ideas from the empirical literature estimating the private return to R&D. The third section develops the theoretical model of R&D investment. The fourth section discusses some important features of the data and section five develops the empirical model and estimation strategy. Sections six and seven discuss the empirical results and report counterfactual simulations of the model.

2 The Private Return to R&D Investment

The expected private return to a firm's R&D investment is one of the main factors driving the firm's decision to invest. Understanding the magnitude and determinants of the private return is key to explaining the observed patterns of R&D investment but also to predicting the likely response of firm investment to changes in the economic environment including policy changes that attempt to subsidize the cost of R&D activities. Estimating the private return to R&D has been been a major area of empirical research for decades.[2] Most of the empirical literature is built upon the knowledge production function developed by Griliches (1979). In this framework, firm investment in R&D, or innovation inputs more broadly defined, creates a stock of knowledge or expertise within the firm that enters into the firm's production function as an additional input along with physical capital, labor, and materials. In addition to being incremented positively by firm R&D expenditures, the knowledge stock can depreciate reflecting the fact that the firm's existing expertise can become irrelevant as new products, materials, and production processes are developed. The key concept of interest in this production function framework is the partial derivative of output with respect to the knowledge stock. This has been estimated as either the elasticity of output with respect to the knowledge stock or the marginal product of the knowledge stock. The marginal product can be interpreted as the gross rate of return to R&D while the net rate of return is defined as the marginal product minus the rate of depreciation.

This knowledge production function model has been extended in several ways, including incorporating R&D spillovers across firms or industries and using firm market value or Tobin's q as a long-run output measure.[3] Another extension incorporates more details on the innovation process that links R&D expenditure and subsequent productivity gains. The Community Innovation Surveys (CIS) have been developed to collect firm-level information on R&D expenditures and firm innovations including the development of new products and the adoption

[2] Surveys of the empirical literature are given by Mairesse and Sassonou (1991), Hall (1996), and Hall, Mairesse, and Mohnen (2010).

[3] See Griliches (1992) for discussion of spillovers and Hall, Mairesse and Mohnen (2010) for a recent review of the empirical evidence. Czarnitzki, Hall, and Oriani (2006) review the literature that measures the effect of the knowledge capital stock on firm market value.

of new or improved production processes. A large empirical literature built around the model of Crépon, Duguet, and Mairesse (1998) has extended the production function framework to incorporate innovation data collected in the CIS.[4] Even with these extensions the primary focus of the literature remains the estimation of either the output elasticity or marginal product of the knowledge capital stock.

An alternative approach to incorporating R&D in the firm's production process has been implemented by Aw, Roberts, and Xu (2011) and Doraszelski and Jaumandreu (forthcoming). Rather than trying to measure the firm's knowledge stock as a deterministic function of past R&D, they model the firm's productivity as a Markov process that is altered by the firm's endogenous choice of R&D. Changes in productivity resulting from both R&D investment and random shocks carry over into future productivity and the degree of persistence is determined by the parameters of the Markov process.[5] Doraszelski and Jaumandreu (forthcoming) implement hypothesis tests that allow them to discriminate between different variations of the knowledge capital and stochastic productivity specifications. Their results favor the endogenous stochastic productivity model. In this paper we adopt the endogenous stochastic productivity framework as one component of the dynamic model of R&D demand but instead of R&D we allow innovation outcomes to affect productivity using CIS data on firm product and process innovations.

Griliches (1979) identified several difficulties in applying the knowledge production framework including estimating the knowledge capital stock and its rate of depreciation from time-series data on R&D expenditures and clarifying the simultaneity between output and R&D expenditure. The simultaneity is particulary important. Current R&D expenditures increase the future knowledge stock which then increases future output through the production function. That is the mechanism of interest but, at the same time, current R&D expenditures are determined by past output and the firm's expectation of future output. Griliches warns that, without careful attention to model specification and formulation, estimates of the effect of R&D

[4]See Hall (2011) for a survey of the empirical studies and Mairesse and Mohnen (2011) and Mairesse, Mohnen and Kremp (2005) for discussion of the estimation issues that arise in using the CIS data. Roberts and Vuong (2013) provide a comparison of the structural model of R&D investment we develop in this paper and the framework from Crépon, Duguet, and Mairesse (1998).

[5]Griliches (1998) and Rogers (2010) estimate the knowledge capital model but also incorporate an exogenous stochastic process for productivity.

on output in this framework may largely reflect the effect of output on R&D (Griliches, 1979, p. 108). One limitation of the empirical studies of the return to R&D is that they tend to focus on the production function itself but have not tried to utilize the additional information contained in the firm's demand curve for R&D.[6] In this paper we develop an alternative approach to estimating the private return to R&D that focuses on the firm's dynamic demand curve for R&D rather than solely on the production function. In this way we model the simultaneous linkages between R&D and output that were identified by Griliches.

Our framework leads to a different formulation of the return to R&D. Rather than being the marginal product of the knowledge capital input, we estimate the benefit of the firm's R&D as the impact of the R&D choice on the expected future profits of the firm. This depends on how R&D affects output, the focus of the knowledge production function literature, but also on how the output change translates into the discounted sum of future firm profits. The stochastic nature of productivity in our framework generates a stochastic component in output and profits so that firms that invest in R&D have different output and profit distributions in future periods which leads to differences in expected long-run profits. This provides the basis for measuring the private return to the firm's R&D investment.

3 Theoretical Model

This section develops a theoretical model of a firm's dynamic decision to undertake R&D investment. In this framework the firm's current productivity is a key determinant of the decision to invest in R&D and future productivity evolves endogenously as a result of the firm's R&D choice. The framework recognizes both the uncertainty the firm faces about the ultimate impact of R&D spending on future productivity and the intertemporal tradeoff, with costs incurred up front but benefits likely delayed in time, that characterizes R&D investment. In the model the firm's choice to invest in R&D alters the probability that the firm will realize a product or

[6] The exceptions to this include Aw, Roberts, and Xu (2011) who estimate a dynamic demand curve for R&D using methods similar to the ones we apply here. Xu (2008) estimates a dynamic demand curve for R&D which includes both a private return to R&D but also an across firm spillover that generates potential social benefits from R&D. Bernstein and Nadiri (1989, 1991) also estimate a demand curve for R&D using a dynamic cost function model. Their model of R&D investment is analogous to an investment model for physical capital and they estimate an Euler equation for the knowledge capital stock.

process innovation in the future. If the firm realizes an innovation, this shifts the distribution of productivity and ultimately profits that they will be able to earn in future periods. The firm will choose to invest in R&D if the expected long-run payoff from this R&D-innovation-productivity process is greater than the current costs of investment.

The model contains four structural components that link R&D, innovations, productivity, and profits. The first is the firm's profit function $\pi(\omega_{it})$ where ω_{it} is firm i's productivity in year t. The second is the effect of the firm's R&D decision on the probability that the firm realizes either a product or process innovation in the future. This is represented by a cdf $F(d_{it+1}, z_{it+1}|rd_{it})$ where d, z, rd are measures of product innovation, process innovation, and R&D investment, respectively. The third component describes the process of productivity evolution, where process and product innovations affect the probability distribution of the firm's future productivity, $G(\omega_{it+1}|\omega_{it}, d_{it+1}, z_{it+1})$. The final structural component is the cost function for R&D investment, $C(rd_{it}|rd_{it-1})$. These costs will be either a sunk startup cost or a fixed maintenance cost depending on the firm's prior history of R&D investment. The next subsections discuss each of these components in turn.

3.1 Productivity and the Firm's Short-Run Profits

The firm's short-run marginal cost is given by

$$c_{it} = \beta_0 + \beta_k k_{it} + \beta_w w_t - \psi_{it}, \tag{1}$$

where c_{it} is the log of marginal cost, k_{it} is the log of firm capital stock, and w_t is a vector of market prices for variable inputs which every firm faces in period t. The firm-specific production efficiency ψ_{it} captures differences in technology or managerial ability and is known by the firm but not observable to the econometrician.[7] The capital stock is treated as a fixed factor in the short-run. Thus, there are two sources of cost heterogeneity across firms: the capital stock and the unobserved production efficiency.

Each firm is assumed to produce one product. The demand for firm i's product q_{it} is given

[7] Variation in input quality, which leads to variation in input prices, across firms will also be captured in ψ. We will model this source of quality variation as part of the unobserved firm efficiency.

by

$$q_{it} = Q_t \left(\frac{p_{it}}{P_t}\right)^{\eta} \exp(\phi_{it}) = \Phi_t (p_{it})^{\eta} \exp(\phi_{it}), \tag{2}$$

where Q_t is the aggregate industry output in period t and P_t is the industry price index which are combined into the industry aggregate Φ_t. The firm-specific variables are the firm's output price p_{it} and a demand shifter ϕ_{it} that reflects product desirability or quality. The demand shifter is known by the firm but not observed by the econometrician. The elasticity of demand η is negative and assumed to be constant for all firms in the industry.

Assuming the firm operates in a monopolistically competitive market, it maximizes its short-run profit by setting the price for its output equal to a constant markup over marginal cost: $p_{it} = \left(\frac{\eta}{1+\eta}\right) \exp(c_{it})$. Given this optimal price, the log of the firm's revenue r_{it} is:

$$r_{it} = (1+\eta)\ln\left(\frac{\eta}{1+\eta}\right) + \ln \Phi_t + (1+\eta)\left(\beta_0 + \beta_k k_{it} + \beta_w w_t - \omega_{it}\right). \tag{3}$$

Revenue productivity is denoted by ω_{it} and is defined as $\omega_{it} = \psi_{it} - \frac{1}{1+\eta}\phi_{it}$. Equation (3) implies that for a given capital stock, heterogeneity in the firm's revenue is driven by differences in production efficiency ψ and the demand shifter ϕ. We will refer to the unobserved revenue productivity ω_{it} simply as productivity.[8] Given the form of the firm's pricing rule there is a simple relationship between the firm's short-run profits and revenue:

$$\pi_{it} = \pi(\omega_{it}) = -\frac{1}{\eta}\exp(r_{it}). \tag{4}$$

The link between productivity ω and short-run profits will be an important determinant of the firm's demand for R&D.

3.2 R&D Investment and Endogenous Productivity

A key component of our framework for endogenous productivity growth is that the firm can affect the evolution of productivity and profits over time by investing in R&D. By exploiting data on actual firm innovations we disaggregate this linkage into two components. First, the

[8] Data on firm sales revenue will contain information on ω_{it}. To estimate our model of R&D demand we do not need to separate ψ and ϕ but only need to quantify the effect of ω_{it} on firm profit. For studies identifying cost and demand shocks using quantity and price data see Foster, Haltiwanger, and Syverson (2008) and Roberts, Xu, Fan, and Zhang (2012).

firm's R&D affects the probability that the firm realizes a product or process innovation in the future. Innovations are denoted as z_{it+1} and d_{it+1} which are dummy variables equal to 1 if firm i realizes a process or product innovation in year $t+1$ and 0 if it does not. We recognize that innovations can take different forms, some may affect the production process and thus work through the shock in the marginal cost function ψ while others may represent new or improved products and work through the demand side shock ϕ. Throughout this paper we treat product and process innovations as distinct and allow them to impact the firm's productivity evolution in different ways. The linkage between R&D and innovation is represented by the cumulative joint distribution of product and process innovations conditional on whether or not the firm invests in R&D, $F(d_{it+1}, z_{it+1}|rd_{it})$. We expect that firms that invest in R&D will be more likely to realize product and process innovations in the next period. This specification captures the first aspect of the uncertainty that firms face when they invest in R&D, the technological uncertainty about the innovation process. The cdf must be general enough to recognize that the firm may have no product or process innovations when they invest in R&D and that they may realize one or both innovations even without R&D investment. The latter can result from luck, the effect of expenditures on R&D in the more distant past even if the firm is not currently investing, ideas that are brought to the firm by hiring experienced workers or other spillover channels, or changes in the production process that result from learning-by-doing without formal R&D investment.

We treat the firm as making a discrete decision $rd_{it} \in \{0, 1\}$ on whether or not to invest in R&D. This is driven by some aspects of the data set we will be using in the empirical application. In general, the measurement error in the continuous measures of R&D expenditure, and the product and process innovations, is more substantial than the error in the discrete variables.[9] In addition, in our data the probabilities of product and process innovation differ substantially between firms that invest in R&D and firms that do not (evidence is provided in Table 4) and we choose to develop the theoretical and empirical model with this in mind.

The second step of the R&D-productivity linkage models firm productivity as a stochastic

[9]See Mairesse, Mohnen, and Kremp (2005) for discussion and evidence on this point using firm data from the French innovation survey.

variable that is affected by the firm's past productivity and the current realizations of product and process innovations. The cdf $G(\omega_{it+1}|\omega_{it}, d_{it+1}, z_{it+1})$ captures the second aspect of uncertainty that firms face when they invest in R&D, the uncertainty about the economic value of an innovation. Even when they realize an innovation, the exact impact of that on future productivity and profits is unknown. It may also be the case that product and process innovations have different impacts on future productivity because they work through different channels on the demand and cost sides. Specifically, we assume that firm productivity evolves as:

$$\begin{aligned}\omega_{it+1} &= g(\omega_{it}, d_{it+1}, z_{it+1}) + \varepsilon_{it+1} \\ &= \alpha_0 + \alpha_1\omega_{it} + \alpha_2\omega_{it}^2 + \alpha_3\omega_{it}^3 \\ &\quad +\alpha_4 z_{it+1} + \alpha_5 d_{it+1} + \alpha_6 z_{it+1} d_{it+1} + \varepsilon_{it+1}.\end{aligned} \tag{5}$$

The function $g(\cdot)$ is the conditional expectation of future productivity and ε is a zero mean stochastic shock. This captures several important aspects of productivity evolution. First, the firm's productivity is assumed to persist over time, with the degree of persistence captured by the coefficients α_1, α_2, and α_3. This intertemporal persistence is an important feature of firm-level data on productivity. Second, innovations are allowed to systematically shift the mean of the distribution of future firm productivity and the magnitude of this effect is captured by the coefficients α_4, α_5, and α_6. The coefficient α_6 allows the possibility that the marginal effect of either a product and process innovation on future productivity will depend on whether the firm has the other type of innovation. Expected future productivity evolves only in the cases where the firm realizes a product or process innovation and this captures the fact that R&D expenditures alone are not sufficient to generate productivity improvements. Third, the specification recognizes that productivity change is affected by stochastic shocks ε_{it+1} that reflect the inherent randomness in the productivity process. We assume the productivity shocks ε_{it+1} are *iid* across time and firms and are drawn from a normal distribution with zero mean and variance σ_ε^2. Because of the persistence in productivity, the shocks in any period will be incorporated into future productivity levels, rather than have transitory effects.

Combining these two stages captures both the uncertainty and the endogeneity of the productivity process. By making investments in R&D the firm will alter the probability of getting a product or process innovation which in turn will alter the distribution of productivity that they face in future periods. We will refer to the first step as the innovation process and the second step as the productivity evolution process. By including the innovation process in the model, rather than linking R&D directly to productivity as in Aw, Roberts and Xu (2011) and Doraszelski and Jaumandreu (forthcoming), we can also gain some insight into whether R&D is working to improve productivity through the demand side or cost side of the firm's operations. In this framework, productivity improves with either cost reductions or revenue expansions. If we find that the overall linkage between R&D and productivity is primarily due to product innovations it suggests that R&D is working through the demand side while a finding of a more important role for process innovations suggests R&D is working through the cost side.

3.3 The Firm's Dynamic Decision to Invest in R&D

This section develops the firm's decision rule for whether or not to invest in R&D. The benefits of investing depend upon the effect of R&D on the firm's expected future productivity and the effect of productivity on future profits as developed in the last two subsections. The firm's decision will also depend on the costs of investing in R&D and these may differ for firms that are just beginning to invest in R&D activities and firms that are maintaining ongoing activities.

We assume that, at the start of period t, the firm observes its current productivity level ω_{it}, knows its short-run profit function and the processes for innovation and productivity evolution F and G. In addition, if the firm is maintaining an ongoing R&D investment then it observes a fixed cost γ_{it}^f of conducting R&D. Alternatively, if it is just beginning an R&D program, then it observes a sunk startup cost γ_{it}^s. Defining the discrete indicator variable rd_{it-1} which equals one if the firm invested in R&D in year $t-1$ and zero if it did not, the cost that firm i must pay in year t can be represented by the R&D cost function:

$$C(rd_{it}|rd_{it-1}) = \gamma_{it}^f rd_{it-1} + \gamma_{it}^s(1 - rd_{it-1}). \tag{6}$$

Both fixed and sunk costs are assumed to be *iid* draws from a known joint cost distribution C^γ. The combination of sunk startup costs together with the uncertainty about the profitability of

R&D investment will create an option value to the firm's investment decision.[10]

The firm's state variables $s_{it} = (\omega_{it}, rd_{it-1})$ evolve endogenously as the firm makes its R&D decision, $rd_{it} \in \{0, 1\}$. Given its state vector and discount factor β, the firm's value function $V(s_{it})$, before it observes the fixed and sunk cost, can be written as:

$$\begin{aligned} V(s_{it}) &= \pi(\omega_{it}) + \\ & \int_{\gamma^f, \gamma^s} \max_{rd \in \{0,1\}} \left(\beta E_t V(s_{it+1}|\omega_{it}, rd_{it} = 1) - C(rd_{it}|rd_{it-1}); \beta E_t V(s_{it+1}|\omega_{it}, rd_{it} = 0)\right) dC^\gamma \end{aligned} \tag{7}$$

where the expected future value of the firm is defined as an expectation over the future levels of productivity and innovation outcomes:

$$E_t V(s_{it+1}|\omega_{it}, rd_{it}) = \sum_{(d,z)} \int_\omega V(s_{it+1}) dG(\omega_{it+1}|\omega_{it}, d_{it+1}, z_{it+1}) dF(d_{it+1}, z_{it+1}|rd_{it}). \tag{8}$$

Equation (7) shows that the firm will choose to invest in R&D if the discounted expected future profits from doing R&D, $\beta EV(s_{it+1}|\omega_{it}, rd_{it} = 1)$, net of the relevant fixed or sunk cost, are greater than the expected future profits from not doing R&D, $\beta EV(s_{it+1}|\omega_{it}, rd_{it} = 0)$.[11] What makes these two expected future profits differ is the effect of R&D on the firm's future productivity. Using this specification we can define the marginal benefit of conducting R&D as:

$$\Delta EV(\omega_{it}) \equiv \beta E_t V(s_{it+1}|\omega_{it}, rd_{it} = 1) - \beta E_t V(s_{it+1}|\omega_{it}, rd_{it} = 0). \tag{9}$$

The firm will choose to invest in R&D if $\Delta EV(\omega_{it}) \geq C(rd_{it}|rd_{it-1})$. This will be the condition used in the empirical model to explain the firm's observed R&D choice.

Overall, this model endogenizes the firm's choice to undertake R&D investments as a comparison between the net expected future profits of the two alternatives. The optimal choice of

[10] See Dixit and Pindyck (1994) for models of uncertainty and sunk costs that generate hysteresis in investment patterns.

[11] The profit function $\pi(\omega_{it})$ and value function $V_{it}(s_{it})$ also depend on the exogenous variables in the firm's environment including the capital stock, variable input prices, aggregate demand shock, and industry demand elasticity. We have suppressed notation for these to highlight the role of R&D, process and product innovations, and productivity. In the empirical model we will define different firm types based on the exogenous variables and calculate the profit and value functions separately for each type.

whether or not to undertake R&D depends on whether the gains in expected future profits from conducting R&D outweigh the relevant startup or fixed cost. Using the empirical model we develop in section 5 we estimate the distribution of sunk and fixed costs faced by the firm and quantify ΔEV, the expected long-run payoff to investing in R&D.

4 Data

4.1 Firm Sample

The data we use to analyze the role of R&D in the productivity evolution of German firms are contained in the Mannheim Innovation Panel (MIP) survey collected by the Centre for European Economic Research (ZEW). The survey is conducted every year for firms in the manufacturing, mining, energy, water, construction and service sector. Firm samples are drawn from the Creditreform database according to the stratifying variables firm size, region, and industry.[12] These are representative for firms with German headquarters and at least 5 employees.

The survey started in 1993 for the manufacturing, mining, energy, water and construction sectors and added the service sector in 1995. The survey follows the form of the Community Innovation Surveys (CIS) that are administered in many OECD countries and adheres to the Oslo Manual which provides guidelines for the definition, classification, and measurement of innovation (OECD (1992, 1997, 2005)). Every year the same set of firms are asked to participate in the survey and to complete the questionnaire sent to them via mail. The sample is updated every two years to account for exiting firms, newly founded firms, and firms that developed to satisfy the selection criteria of the sample. Additionally a non-response analysis is performed via phone to check and correct for non-response bias. Every firm is in the panel, on average, for 2 to 3 years. Due to cost reasons, starting in 1998 the full questionnaire was only sent out every other year to all firms in the full-sample. However, information on variables of interest are asked retrospectively for the previous year to ensure the annual coverage. In odd years only short questionnaires with core questions are sent to a subset of firms. Therefore, the number of firms in odd years in the panel is significantly lower than in even years. This limits the ability

[12]The Creditreform database is the largest credit-rating agency in Germany and maintains comprehensive database of approximately 3.3 million German firms.

to follow individual firms over time. Participation in the survey is voluntary and the average response rate is about 25 percent, so each year there are approximately 5000 firms answering the questionnaires across all industries (see Rammer and Peters, 2013).

For the empirical analysis we focus on firms in the manufacturing sector, NACE industries 15 to 37, for a number of reasons. First, manufacturing has the best overall coverage in the survey. Second, prior to 2001 the firm questionnaires differ across manufacturing and service sectors and some of the necessary variables, such as the capital stock, are not always collected for other sectors. Finally, much of the reported R&D expenditure occurs in the manufacturing industries. We will focus on two groups of manufacturing industries. The high-tech (HT) industry group consists of five aggregates of two-digit manufacturing industries (NACE codes), chemicals (23, 24), non-electrical machinery (29), electrical machinery (30, 31, 32), instruments (33), and motor vehicles (34, 35). Based on OECD data these industries all have R&D-sales ratios that exceed .025. The low-tech (LT) industry group will include seven aggregated industries, food (15, 16), textiles (17, 18, 19), paper (20, 21, 22), plastic (25), non-metallic minerals (26), basic metals (27, 28) and manufacturing n.e.c. (36, 37), that all have much lower R&D-sales ratios. Our data covers the period 1993-2008. Due to the small number of observations in some industries we will have to combine manufacturing industries for some of the empirical analyses, particularly those requiring time-series data. In general, the parts of the model that can be estimated using the cross-sectional observations in the data will be estimated separately for each of the 12 industries. There are 18,655 cross-sectional observations that are used in estimation. The parts of the model that require time-series data because of the use of lagged variables will be estimated separately for the high-tech and low-tech groups of industries. We have 3,337 and 4,298 time-series observations for the high-tech and low-tech industries, respectively.

4.2 Variable Measurement

For the estimation we use data on firm revenue, variable cost, capital stock, innovation expenditures, and product and process innovations.[13] Firm revenue is the sum of domestic and export

[13] For 1999 and 2000 the panel does not contain information on the firms' capital stock and we impute these missing years using linear interpolation.

sales. Total variable cost is defined as the sum of expenditure on labor, materials, and energy and the firm's short-run profit is the difference between revenue and total variable cost. The firm's value is the discounted sum of the future short-run profits and thus measures the long-run resources that the firm has available to pay its capital expenses plus economic profits.

A special feature of the Community Innovation Surveys is that they provide measures of both innovation input and innovation outputs. Innovation input is measured by the firm's expenditure on a set of activities related to innovation. This measure includes R&D spending but also spending on worker training in this area, acquisition of external knowledge and capital, marketing, and design expenditures for producing a new product or introducing a new production process. The R&D variable we will analyze in the empirical model (rd_{it}) takes the value one if the firm reports a positive level of spending on innovation activities.

Innovation output captures the introduction of a new product or a new production process by the firm.[14] The Oslo Manual defines a product innovation as a new or significantly improved product or service. A process innovation refers to new or significant changes in the way products are produced, delivered, or supplied. The main purpose of a process innovation is to reduce production costs or to improve the quality of a product. For instance, the use of lasers to increase the quality of products in metal processing or the introduction of automation concepts are process innovations. The innovation does not have to be new to the market but only to the firm. A firm could report an innovation if it adopted a production technology or business practice from a competitor or expanded its product line even if the product was already offered by other firms.

The timing assumptions in the theoretical model about the relationship between R&D spending, innovation outcomes, and productivity are fairly general: R&D spending precedes innovation outcomes and innovations that are realized are assumed to affect productivity and profits in the period they are introduced. In the survey in year t, the firms are asked whether they introduced new or significantly improved products or services during the years $(t-2)$, $(t-1)$,

[14]Beginning in 2005 the survey also includes questions on organizational innovation, which is defined as new business practices, workplace organization, or external relations, and marketing innovation, which refers to changes in product design, packaging, product placement or promotion, and pricing methods. The time-series information on these variables is too short for them to be utilized in this study.

or t. The discrete variable *product innovation* d_{it} takes the value one if the firm reports yes to the question. The discrete variable for *process innovation* z_{it} equals one if the firm reports new or significantly improved internal processes during the years $(t-2)$ to t. In the empirical model this outcome will be related to R&D spending in the previous year $(t-1)$, so there is not a perfect match between the timing of the R&D and the realization of the innovations. This may lead us to overestimate the effect of R&D on innovation since the innovation variable could be capturing outcomes two years earlier. Attempting to use more distant lags of R&D spending exaggerates the problems caused by sample attrition and reduces the number of observations with the necessary current and lagged variables.

Table 1 summarizes the proportion of firms in the sample that report having positive innovation expenditures, the proportion of firms with successful product innovations, and the proportion with successful process innovations in each industry. The industries are aggregated into the high-tech and low-tech groups. In our sample the majority of firms report making expenditures on innovation activities but the proportions differ across industries. In the five high-tech industries the proportion varies from .731 to .818 while in the seven low-tech industries it varies from .514 to .642. The rate of product innovation is also higher in the high-tech industries. Between .650 and .771 of the firm/year observations report having a new product innovation while in the low-tech group the rate of product innovation varies from .392 to .592. This same difference exists for process innovation but the difference in magnitude between the high-tech and low-tech industries is not as large. The high-tech industries vary in a narrow band between .330 and .398 and all but one of the the low-tech industries vary between .245 and .327. The model developed in the last section will allow product and process innovations to occur at different rates given the firm's R&D expenditure and will allow them to each have a different impact on future productivity. This will lead to differences in the expected benefits of R&D across industries and help to explain differences in the proportion of firms that choose to invest in R&D.

By examining the data on R&D investment patterns we also see an important role for firm size. Table 2 reports the transition rates for firms' R&D activities between periods. The transition patterns in the data are important for estimating the sunk and fixed costs of conducting

R&D. There is a substantial pattern of movement of firms into and out of R&D activities over time. The rate at which firms begin conducting R&D varies from 17.63 to 33.71 percent depending on the firm's size class and this rate increases with the size class. The rate at which they leave varies from a high of 21.75 percent for the smallest size step class to 6.77 percent for the largest firms. The firm's capital stock will be an important dimension that is controlled for in the empirical work.

5 Empirical Model

5.1 Productivity Evolution

In this subsection we describe how we use the data in the MIP to estimate the R&D-innovation and innovation-productivity relationships. The first step is to estimate the joint probability distribution for innovations conditional on R&D, $F(d_{it+1}, z_{it+1}|rd_{it})$. Given that the three variables are discrete and observed in the data we estimate the joint probabilities as the fraction of observations reporting each of the four combinations of d_{it+1} and z_{it+1} conditioning on $rd_{it} = 0$ and $rd_{it} = 1$. The innovation probabilities are estimated separately for each industry.

Estimates of the process of productivity evolution, equation (5), are needed to construct the transition probabilities for productivity $G(\omega_{it+1}|\omega_{it}, d_{it+1}, z_{it+1})$. Unlike the innovation and R&D variables, the firm's productivity is not observable and these parameters will be estimated along with the firm's revenue function using the data on firm sales. The key parameters to be estimated are the cost elasticity of capital β_k, the parameters of the productivity process $\alpha_0, \ldots, \alpha_6$ and the elasticity of demand η.

The demand elasticity for each industry is estimated using the expression for short-run profit in the model, equation (4). The ratio of variable profit to firm revenue equals $-1/\eta$ and we use the mean profit-revenue ratio in each industry as an estimate of the inverse industry demand elasticity.

To estimate $\alpha_0, \ldots, \alpha_6$, and β_k we follow the proxy variable approach pioneered by Olley and Pakes (1996).[15] Their insight is that if the firm observes its own productivity level before

[15]The revenue function cannot be estimated consistently using OLS because the productivity level ω_{it}, which is contained in the error term, is likely to be correlated with the firm's capital stock k_{it}. The capital stock depends

choosing its variable input levels then input demands are functions of productivity and the fixed factors of production and we can infer information about productivity from observing the expenditure on variable inputs. Following Levinsohn and Petrin (2003) we focus on the choice of material spending and write the firm's demand for its intermediate input as $m_{it} = f_t(k_{it}, \omega_{it})$, where f_t is assumed to be strictly monotone in ω_{it} for a given k_{it}. Inverting the material demand function for ω_{it} and substituting it into equation (3) the revenue function can be written as:

$$\begin{aligned} r_{it} &= \delta_0 + \sum \delta_t D_t + (1+\eta)\left(\beta_k k_{it} - f_t^{-1}(k_{it}, m_{it})\right) + u_{it} \\ &= \delta_0 + \sum \delta_t D_t + h(k_{it}, m_{it}) + v_{it} \end{aligned} \tag{10}$$

where u_{it} and ν_{it} capture transitory shocks and measurement errors in firm revenue. The time dummies D_t control for the factor prices and aggregate demand shock and the intercept contains the demand elasticity. The function $h(k_{it}, m_{it}) = (1+\eta)\left[\beta_k k_{it} - \omega_{it}\right]$ controls for the joint effect of productivity and capital stock on the firm's revenue. By replacing $h(k_{it}, m_{it})$ with a cubic function of its arguments, equation (10) will be estimated separately for each industry using OLS.

Using the fitted value $\widehat{h_{it}}$ from equation (10) and substituting it into equation (5) we can recover the remaining structural parameters by estimating:

$$\begin{aligned} \widehat{h_{it}} &= \beta_k^* k_{it} - \alpha_0^* + \alpha_1(\widehat{h_{it-1}} - \beta_k^* k_{it-1}) - \alpha_2^*(\widehat{h_{it-1}} - \beta_k^* k_{it-1})^2 + \\ & \quad \alpha_3^*(\widehat{h_{it-1}} - \beta_k^* k_{it-1})^3 - \alpha_4^* z_{it} - \alpha_5^* d_{it} - \alpha_6^* d_{it} z_{it} - \varepsilon_{it}^* \end{aligned} \tag{11}$$

where $\alpha_2^* = \alpha_2 \frac{1}{(1+\hat{\eta})}$ and $\alpha_3^* = \alpha_3 \frac{1}{(1+\hat{\eta})^2}$. All other parameters with an asterisk denote the original parameter times $(1+\hat{\eta})$. Estimating this equation using NLLS yields the estimates $\hat{\alpha_0}, \ldots, \hat{\alpha_6}, \hat{\beta}_k$. An estimate of firm productivity can then be constructed as:

$$\hat{\omega}_{it} = -\frac{1}{1+\hat{\eta}}\widehat{h_{it}} + \hat{\beta}_k k_{it}.$$

This process differs slightly from the methodology developed by Olley and Pakes (1996) in two respects. First, productivity evolution is not an exogenous process but is affected by the

on the prior period investment which, in general, will be partly determined by the prior year's productivity ω_{it-1}. The assumption that productivity is serially correlated implies that current productivity and capital stock are correlated which causes the OLS estimates to be inconsistent.

firm's innovations and, as a result, the innovation variables enter into equation (11).[16] Second, because we are modeling productivity using the revenue function, we do not need to estimate the production function coefficients on the variable inputs of labor and materials. This simplifies equation (11) by removing the need to instrument variable input levels which would appear on the right hand side when using the production function as the starting point.[17]

5.2 Value Function and the Dynamic Choice of R&D

As described in section 2, the firm bases its R&D investment decision on a comparison of the long-run payoff from R&D, $\Delta EV(\omega_{it})$, with the realized fixed cost or startup cost, $C(rd_{it}|rd_{it-1})$. The probability that the firm chooses to invest in R&D is given by:

$$Pr\left(rd_{it} = 1|s_{it}\right) = Pr\left[C(rd_{it}|rd_{it-1}) \leq \Delta EV(\omega_{it})\right] \tag{12}$$

where the fixed costs and sunk startup costs of R&D investment are assumed to be distributed exponentially with mean γ^F and γ^S, respectively.

It is reasonable to assume that firms that perform R&D continuously might have different cost structures than firms that have to start the investment activity from scratch. It can be costly to set up and equip the research department or hire employees for the research unit. In this model, the firm's R&D cost is viewed as the expenditure the firm will need to spend to generate a process or product innovation. This cost includes the expenditure on employees and materials reported by the firm in the innovation surveys but should also include any adjustment cost that the firm incurs in starting or maintaining its operations. It should also include the capital costs of buildings and equipment used in the R&D process and these are unlikely to be reported in the innovation surveys. For this reason it is important to allow for some randomness or measurement error in the R&D expenditure. In the implementation, we also allow the cost

[16]Doraszelski and Jaumandreu (forthcoming) adopt an alternative approach to deriving an equation similar to (11). Like most of the recent production function literature, they specify a Cobb-Douglas production function but then invert the labor demand function that is derived from it. In this paper we are not interested in estimating the production function parameters but instead focus on the α parameters that describe the process of productivity evolution because these are the parameters that affect the future returns from R&D investment.

[17]Variable input prices are also arguments of the revenue function. The component of prices common to all firms will be captured by the time dummies. Firm-level variation in input prices will be one source of variation in the error term in equation (10).

distributions to differ across firms depending on firm size (measured by the value of the capital stock). This is reasonable since large firms will generally have larger expenditures on R&D if they choose to invest. The larger expenditures could reflect numerous projects being undertaken across different product lines or production processes within the firm or better access to credit which allows them to set up and maintain a larger research unit. One thing that the framework does rule out is persistence in the firm's R&D expenditure over time and we think that allowing the cost distribution to vary across firms of different sizes will control for the likely persistence that reflects differences in firm size.

The final piece of the empirical model is the construction of the value function and $\Delta EV(\omega_{it})$, the estimate of the payoff to R&D equations (7) and (9), respectively. Rust (1987) developed the nested fixed point algorithm for estimating dynamic discrete choice models and we use this methodology here. We discretize the state space $s_{it} = (\omega_{it}, rd_{it-1})$ into 100 grid points for productivity and two values for lagged R&D choice and use value function iteration to solve for the value function at each element of this discretized state space. In addition, firms are divided into discrete firm types based on the value of their capital stock, using 100 grid points, and 12 industries and the value function is estimated at each discrete state point for each of these firm types. We use a cubic spline to interpolate across the productivity and capital grid points for each industry and impute the firm's value function $V(s_{it})$ at each data point in the sample.

Assuming the firm's state variables s_{it} are independent of the cost draws and that the costs are *iid* across all periods and all firms, the likelihood function for the firms' R&D choice data can be expressed as

$$L(\gamma|rd, s) = \prod_{i}^{N} \prod_{t}^{T_i} \Pr(rd_{it}|s_{it}; \gamma), \tag{13}$$

where $\gamma = (\gamma^F, \gamma^S)$. The vectors rd and s contain every firm's R&D choice and state variables for each period, respectively. The total number of firms is denoted by N and T_i is the number of observations for firm i. We estimate the parameters of the cost distribution using the firms' discrete choices on R&D.

6 Empirical Results

6.1 Estimates of the Productivity Process

Estimates of the probability of an innovation conditional on the firm's prior period investment in R&D, $\Pr(d_{it+1}, z_{it+1}|rd_{it})$, are reported for each industry in Table 3. There is a strong but not perfect relationship between R&D investment and innovation outcomes. Columns (2) through (5) show the probability of realizing each combination of product and process innovation given that the firm does not engage in R&D. Columns (6) through (9) report these probability for firms that conduct R&D. Focusing first on the firms that did not engage in R&D, column (2) shows that, on average, they have approximately a 78 percent chance of not having either a product or process innovation in the next year. This estimate is very similar across industries varying only from a low of .716 in electronics to .822 in basic metals. It does not differ significantly between the low-tech and high-tech industry groups. What is more important to note is that approximately 22 percent of the firms still realize innovations even if their R&D spending is zero and the most common outcome among the three combinations is that they have both product and process innovations ($d = 1, z = 1$). This indicates that prior period R&D is neither necessary or sufficient for the firm to report realizing an innovation. Our model recognizes this possibility in the link between R&D and future productivity.

For the firms that invest in R&D we observe that they are much less likely to report no innovation. Column (6) shows that between 9.0 and 27.1 percent of the firms that conduct R&D report no innovations in the next year. This probability does vary between the industry groups, being significantly higher for the low-tech industry group. This reflects a combination of lower R&D effort in these industries, even when the firm reports conducting R&D, and fewer technological opportunities for innovations. Among the three possible combinations of innovation outcomes, the most common is that the firm reports both a product and process innovation ($d = 1, z = 1$) with between 44.8 and 63.8 percent of the R&D firms reporting both innovations. Among these firms the success rate for introducing a new product innovation is in general much higher than the rate for a new process. The two exceptions are the paper and basic metals industries where these two probabilities are similar. For both of these industries large scale production is important and this could give a strong incentive for firms to invest to improve their production efficiency.[18]

Table 4 reports the estimates for the demand elasticities for each industry in the high-tech and low-tech sectors. For instance, the estimate of $(1 + 1/\hat{\eta})$ in the chemicals industry is .708. This implies a demand elasticity $\hat{\eta}$ of -3.425 which is reported in the third column. The demand elasticity is important in converting productivity into profit as seen in equations (3) and (4). The estimates vary substantially across industries ranging from -2.994 in the food industry to -7.937 in vehicles.

Table 5 reports the estimates of the productivity evolution process for the high-tech and low-tech sectors using equation (11). The double and single asteriks denote parameter estimates different from zero at the .01 and .05 significance levels, respectively. The cost elasticity of capital in the high-tech sector is estimated to be $\hat{\beta}_k = -0.056$ and in the low-tech sector is -0.060. Negative values of β_k imply firms with a higher capital stock have lower production costs because they use less variable inputs.

The positive coefficient estimates for z and d indicate that firms that realize innovations have, on average, higher future productivity levels compared to those that do not have any kind of innovation. The marginal effects of adopting a new process or developing a new product is nearly identical for high-tech firms. A new process innovation z contributes on average 1.4 percent to productivity gain and a new product innovation d contributes 1.3 percent. There is no additional effect from having both product and process innovations. The coefficient on the interaction term $d * z$ is -.014 which just outweighs the direct effect of the second innovation. Basically, firms with either or both types of innovation have, on average, 1.4 percent higher productivity in the next year.

The difference in the effect of the two types of innovations is more pronounced in the low-tech industries. Firms that introduced a new product have on average 0.2 percent higher future

[18] If we construct Table 3 using rd_{it-1} as the conditioning variable, so there is a two-year lag between R&D and innovation, we get a very similar pattern of innovation rates. Among the firms with $rd_{it-1} = 0$, 74.3 percent report no innovation. There is no difference between the high-tech and low-tech industries. Among the firms with $rd_{it-1} = 1$, 18.7 percent report no innovation and the average is twice as large in the low-tech sectors compared with the high-tech sectors. The estimates of innovation probabilities by industry are not sensitive to the use of one or two period lags in R&D.

productivity while a new process innovation raises productivity by 1.0 percent. One reason for the weaker impact of product innovation on future productivity is that new or improved products may represent less substantial changes over existing products in these industries.[19] If a firm realizes both product and process innovation the estimated interaction term, which is -0.002, partially offsets the marginal effect of the second innovation type. The three coefficients together imply that firms with a process innovation have 1 percent higher future productivity regardless of whether or not they also have a product innovation and firms with just a product innovation do not have significantly higher productivity in the next period.[20]

The effect of past productivity on the current productivity level is measured by the coefficients of ω_{t-1}, its squared and cubic terms. Past productivity is highly persistent. There is a non-linear relationship between current and lagged productivity for high-tech firms as seen by the statistically significant effect of ω_{t-1}^2. These higher-order terms are not significant in the low-tech industries implying a linear relationship between the current and lagged productivity level. The persistence of the productivity process has a substantial impact on the long-run payoff from R&D because it determines how quickly the productivity gains from an innovation depreciate. Lower values of α_1 imply more rapid depreciation of the productivity and profit gains from an innovation d or z and, because it depends on the discounted stream of future profits from the innovation, the long-run payoff to R&D will fall. Overall, larger coefficients on the innovation variables and higher persistence of the productivity process both raise ΔEV, the expected long-run payoff to R&D.

The empirical literature measuring the return to R&D has often constructed the elasticity of output, usually measured as firm revenue, with respect to R&D expenditure. Hall, Mairesse, and Mohnen (2010) review this literature and report that the elasticity estimates based on production function models vary from .01 to .25 and are centered around .08. Doraszelski and Jaumendreu (forthcoming, Table 5) report estimates of the distribution of firm-level estimates

[19] This interpretation is supported by data on products that are new to the market. In the MIP the proportion of firms introducing products that are new to the market varies from 39 to 51 percent in the high-tech industries but 16 to 31 percent in the low-tech industries.

[20] We conducted a sensitivity check on the specification of the productivity process by allowing separate industry intercepts in the productivity evolution equation (5). They were never statistically significant and this is not surprising given the strong effect of past productivity.

for ten Spanish manufacturing industries. The average value over all firms is .015 and the average at the industry level varies from -.006 to .046 across the ten industries with half of the industries falling between .013 and .022. Using the results reported in Tables 3, 4, and 5 we construct an analogous measure using the discrete R&D variable: the proportional gain in firm revenue resulting if the firm moves from not investing in R&D ($rd_t = 0$) to investing in R&D ($rd_t = 1$).[21] Table 6 provides estimates of this shift on the log of future revenue for each industry. For the five high-tech industries, the elasticity of revenue with respect to discrete shift in R&D varies from .021 to .058 while they are generally smaller, ranging from .008 to .026 for the low-tech industries. In our dynamic framework this is one component of the expected benefit of R&D investment ΔEV but it is not the sole focus of our estimation.

6.2 Estimates of the Cost of an R&D Program

The final set of parameters we estimate is the startup and fixed costs of establishing and maintaining an R&D program. To account for size differences across firms that will be reflected in the magnitude of their R&D expenditures, we allow the startup cost and fixed cost distribution C^γ to vary for three size groups of firms, small, medium and large.[22] We estimate the fixed and startup cost parameters $(\gamma_s^F, \gamma_m^F, \gamma_l^F, \gamma_s^S, \gamma_m^S, \gamma_l^S)$ where the subscript denotes the size category by maximizing the likelihood function in equation (13). Table 7 reports the estimated means of the distribution of startup (γ^S) and fixed costs (γ^F). The first three rows report the results for the high-tech group distinguishing between firm sizes. The average costs for firms in the low-tech group are reported in the last three rows.

A number of general patterns stand out across all specifications. First, fixed costs are smaller than startup costs for all firm sizes. This means that firms that were previously engaged in R&D

[21]The revenue increase resulting from R&D depends on how R&D affects innovation, how innovation affects productivity, and how productivity translates into revenue. The difference in log revenue when rd=1 and rd=0 is constructed as:

$$\Delta r = (1+\eta) \sum_{(d,z)} [g(\omega, d, z) - g(\omega, 0, 0)] [\Pr(d, z|rd = 1) - \Pr(d, z|rd = 0)]$$

for all $(d, z) \in \{(1, 0), (0, 1), (1, 1)\}$.

[22]Firms were divided into size categories based on their capital stock. Firms with a capital stock up to the 33rd percentile of the firm distribution are considered to be small. Large firms have capital stock exceeding the 66th percentile.

have to incur a smaller cost if they want to continue their R&D activities while firms that did not have any previous R&D activities will have to pay a higher amount to start their R&D activities.

In Table 7 we estimate an average startup cost for doing R&D for small firms in high-tech industries of EUR 3.98 mln, more than six times higher than the fixed cost of EUR 0.65 mln. In the high-tech sector, the ratio of startup costs to fixed costs is approximately 6 for small, medium, and large firms. In the low-tech sector the ratio is between 4 and 5. The difference between fixed and startup costs is crucial in explaining the pattern of R&D choice in the data. If the fixed cost is low relative to the startup cost, continuing to do R&D is more attractive because it allows firms to avoid paying the higher cost if it restarts its investment. Even facing negative shocks that lower the expected return of R&D would have less of an effect on the firm quitting R&D. A high startup cost prevents firms from starting to do R&D which can contribute to the high persistence for non-R&D firms seen in Table 2. On the other hand, reducing the gap between fixed and startup costs would imply more switching between starting and quitting R&D. The magnitude of the cost estimates in the low-tech sector ranges between half and one-third of the estimates in high-tech.

A second pattern that stands out is that both fixed and startup costs increase with the firm's capital stock. There is a positive correlation in the data between capital stock and productivity and the payoff to conducting R&D is increasing with the capital stock. Despite this higher payoff not all large firms conduct R&D and this reflects the higher costs that they face.

We can assess the goodness of fit of the dynamic model by simulating the firms' investment choices, given their capital stock and productivity level, and computing the percentage of correct predictions. Using the cost estimates in Table 7, the model fits the data well. In the high-tech industries, the overall percent of correct prediction is 77.84 and the model does a better job of correctly predicting that firm will chose to invest (84.46 percent) than not invest (54.16 percent). For the low-tech sector the model correctly predicts 65.73 percent of all cases, with 66.93 percent correct predictions for positive investment and 64.21 percent correct predictions of no investment.

6.3 Expected Benefits and Costs of R&D

Using these parameter estimates and equation (9), we construct $\Delta EV(\omega)$ the expected long-run payoff to investing in R&D. This measures the difference in the present value of expected future profits that accrue to the firm if it engages in R&D in a year versus if it does not engage in R&D. This benefit depends on the industry-level measures (profit function, demand elasticity, and innovation probabilities) and the firm-level variables (productivity and the capital stock) and varies across firms in an industry as a result. It captures the randomness that arises in the relationship between R&D investment and a product or process innovation, captured in the model by $F(d', z'|rd)$, as well as the randomness between innovation outcomes and productivity, captured in the model by $G(\omega'|\omega, d', z')$.

Tables 8 and 9 provide estimates of $\Delta EV(\omega)$ using innovation outcomes at five different percentiles (5, 25, 50, 75, and 95) of the productivity distribution within each industry. The values in all cells in these tables are averaged over the capital stocks and years. Table 8 covers the five high-tech industries and Table 9 reports the values for the seven low-tech industries. The first five rows of Table 8 show that, as the productivity of a firm in the chemical industry increases from the 5th (-.299) to 95th (2.053) percentile, $\Delta EV(\omega)$ rises from 0.965 million to 87.131 million euros. This reflects the impact of the higher productivity resulting from R&D on the firm's expected future profits. Every industry shows the benefit of R&D increasing with firm productivity but the level of the benefit differs across industries. Comparing the group of industries in Tables 8 and 9 we see that the marginal benefits of R&D are much larger in the high-tech industry group. At the upper end, in the electronics industry the high productivity firms have benefits from an R&D program averaging over 111 million euros. In contrast, the benefits of an R&D program in the low-tech industries is always less than four million euros and generally only exceeds one million euros for the highest productivity firms. This illustrates that the payoff to R&D is very specific to an industry reflecting differences in profit functions. If we rank industries by the expected marginal benefit at the median of the productivity distribution, the vehicle, chemical, and electronics industries have the highest expected payoffs to R&D, followed by machinery and instruments. Minerals and manufacturing nec products have the lowest expected benefits.

In the model developed above, firm i will choose to do R&D if $C(rd_{it}|rd_{it-1}) \leq \Delta EV(\omega_{it})$. The realized costs of firms that choose to do R&D will be described by a truncated cost distribution where $\Delta EV(\omega_{it})$ is the truncation point. For example in the chemical industry, low-productivity firms have a marginal benefit of R&D of 0.965 million euros, so only firms that have R&D costs less than this will choose to invest. Because firms with the same observable productivity and capital stock spend different amounts on R&D to realize the same expected gross benefit $\Delta EV(\omega_{it})$ they have different net benefits from their R&D investment. In section 6.4 we report the distribution of net benefits to R&D across firms.

The fourth and fifth columns of Table 8 report the mean fixed and startup costs among the firms in the five high-tech industries that choose to conduct R&D ($rd = 1$). For example, in the first row of the table, the low productivity firms in the chemical industry that invest in R&D will have an average R&D expenditure of 0.437 million euros if they had previously conducted R&D, and so were paying a fixed cost to maintain it, or 0.475 million euros if they were paying a startup cost to begin an R&D program. The mean truncated expenditure on R&D rises with the level of productivity because the truncation point ΔEV rises with productivity and thus high productivity firms are willing, on average, to invest more money in R&D programs than low productivity firms. The R&D expenditure differs across industries, reflecting differences in the distribution of productivity, capital stocks, and profit function parameters but the differences are fairly small for fixed costs ($rd_{t-1} = 1$) and larger for startup costs ($rd_{t-1} = 0$). The fixed costs for the median productivity firm are almost always less than 3 million euros while the expenditure by the median firm starting up an R&D program can range as high as 12.0 million euros in the case of the vehicle industry. Examining the patterns for the low-tech industry group in Table 9, we observe the same increase with productivity and higher costs for firms that were inexperienced ($rd_{t-1} = 0$). Not surprisingly, the mean firm R&D costs are lower in these industries reflecting the fact that the benefit of R&D investment is also lower.

The final two columns of Tables 8 and 9 report the probability a firm conducts R&D based on its productivity, industry, and prior experience. Several patterns are evident. First, for the high-tech industries the probability of maintaining an R&D program is generally above

.90 for firms that have prior experience. This reflects the high benefits of conducting R&D seen in column (3) and relatively low cost in column (4). For firms that do not have prior R&D experience the probability is substantially lower, resulting from startup costs that are higher than the maintenance fixed costs. The gap between the probabilities in the last two columns is a measure of the effect of the higher startup costs, because the expected benefit of conducting R&D faced by the firm is the same independent of their experience. Focusing on the R&D probabilities in the low-tech industries in Table 9 we observe the same pattern of higher probabilities with higher productivity, reflecting the higher marginal benefits seen in column (3), and with experience, caused by the lower fixed costs relative to sunk startup costs. The primary difference between these industries and the ones in Table 8 are that the magnitudes of the estimates are substantially lower for the low-tech industries. This reflects a lower payoff to R&D in these industries. While the cost distributions are lower, as seen in Table 7, they are not enough to compensate for the lower benefits accruing to R&D.

6.4 The Return to R&D

The patterns of benefits and costs reported in Tables 8 and 9 are predictions from the estimated model across different values of productivity and prior R&D experience and emphasize how variations in the state variables impact the expected long-run benefits of R&D. In this section we turn to the actual data and calculate the long-run expected benefit of R&D, $\Delta EV(\omega_{it})$, for each firm observation given its observed productivity, capital stock, and industry. We can then compare this with the cost of R&D to calculate an expected net benefit for each observation. Let γ_{it} be the fixed or startup cost draw that the firm gets, then the expected net benefit, prior to observing the cost draw, is $\Delta EV(\omega_{it}) - E(\gamma_{it})$ where $E(\gamma_{it})$ is the mean of the distribution of γ_{it}. The $E(\gamma_{it})$ will depend upon the firm's prior R&D experience and size category as seen in Table 7. We will normalize this expected net benefit by the value of the firm $V(s_{it})$ given by equation (7) and define the summary measure:

$$NB_{it} = \frac{\Delta EV(\omega_{it}) - E(\gamma_{it})}{V(s_{it})}$$

This normalized expected net benefit of R&D will vary across firms and time depending on the productivity, capital stock, and R&D history. Normalizing by the value of the firm corrects for differences in firm size that will be reflected in $\Delta EV(\omega_{it})$ and provides a more useful metric for interpreting proportional differences in the benefit of R&D across firms and industries. For many firms NB will be negative reflecting the fact that, given the expected costs and benefits faced by the firm, it is optimal to not invest in R&D. This measure provides a useful description of how the expected net payoff to R&D varies across the whole distribution of firms in operation, even firms that do not invest in R&D.

Alternatively, the return to R&D can be calculated just for the firms that actually invest in R&D. Firms will only choose to do R&D when $\Delta EV(\omega_{it}) - \gamma_{it} > 0$, and we can characterize the net benefit of R&D for firms that choose to do R&D using the truncated mean of the distribution of γ. The expected net benefit of R&D when firms choose to do R&D is defined as:

$$TNB_{it} = \frac{\Delta EV(\omega_{it}) - E(\gamma_{it}|\Delta EV(\omega_{it}) - \gamma_{it} > 0)}{V(s_{it})}$$

Unlike NB, TNB will always be positive. This measure is the model equivalent of calculating the returns to R&D using a sample of firms that all invested in R&D.

Both NB and TNB capture the fact that current R&D expenditure affects the future path of productivity and R&D choices. They are useful measures of the firm-level, long-run net expected benefit of investing in R&D. Given estimates of NB and TNB for all observations in the data we can then summarize the entire distribution of expected R&D returns.

In Table 10 we present the 25th, 50th, and 75th percentiles of the distributions of NB_{it} and TNB_{it} across observations in the data. We divide the firms by industry and by their prior R&D status because this affects whether they pay a fixed cost (continuing firms columns 3-5) or a startup cost (startup firm columns 6-8) when they invest in R&D. For example, in the case of the chemical industry the three percentiles of the distribution of NB for continuing R&D firms are .021, .031, and .037 indicating that the expected net benefit of R&D varies in a fairly narrow range relative to firm value. The median firm will have an expected long-run net payoff to an R&D program of 3.1 percent of firm value. When we truncate the payoffs to recognize that firms will choose not to do R&D in high cost situations, the distribution of net payoffs, TNB, is not greatly affected. The truncated percentiles are .023, .033, and .037 indicating just a slight rightward shift in the lower tail of the distribution. This occurs because the fraction of firms with negative NB is small and so there is little difference between the truncated and untruncated cost means for continuing firms.

The picture is different when we look at firms that were not previously conducting R&D. In the chemical industry, the 25th, 50th and 75th percentilles of the distribution of NB are -.045, .001, and .016. In particular, a large percentage of the firms would have negative expected net benefits because the startup costs they would pay exceed ΔEV and so would not choose to invest in R&D. The median firm would choose to do R&D but would have an expected net payoff equal to one-tenth of one percent of the value of the firm. When we focus on the expected returns of the firms that would actually choose to invest we observe that the percentiles for this industry vary from .019 to .027. Each percentile of the distribution is less than the corresponding percentile for the distribution of continuing firms because of the higher startup costs these firms must pay.

Across the other four high-tech industries a similar pattern is observed. There are fairly small differences in the distribution of NB and TNB for the continuing firms because the fixed costs they would pay tend to be small relative to the benefits and so most firms would choose to continue to invest. There are more substantial differences in the two distributions for starting firms because many of these firms would not choose to do R&D. Even the median firm has a negative expected return in the machinery and instruments industries. Comparing TNB across industries for the continuing firms we find very similar distributions. The median firm in each industry has an expected net benefit of R&D that is between 2.6 and 3.2 percent of the firm's value. For the firms beginning to do R&D, the median varies from 2.0 to 2.4 percent across industries.

In the low-tech industry group we see a different picture reflecting much lower returns to R&D investment. For all of the industries the median of the distribution of NB is negative and the 75th percentile is at its highest .001. The percentiles of the distribution of TNB lie between .001 and .003, indicating small net benefits of R&D relative to firm value even for the firms that find R&D profitable. The pattern is even stronger for the starting firms. The 75th percentile of NB is negative across all industries and the 75th percentile of TNB does not exceed .003. Overall, evaluating these returns at the mean cost levels, less than one quarter of the low-tech firms would choose to invest in R&D.

The measures in Table 10 summarize the long-run payoff to R&D, including the effect on the incentive to invest in R&D in future periods. We can also use the model to calculate the short-run benefit of R&D. We define this as the increment to next period expected profits if the firm chooses R&D in the current period versus if it does not. We will express this as a ratio to the long-run benefit :

$$SNB_{it} = (E\pi(\omega_{it+1}|rd_{it} = 1) - E\pi(\omega_{it+1}|rd_{it} = 0))/\Delta EV(\omega_{it})$$

This measure recognizes that current R&D spending will affect productivity and profit in the next period. Table 11 reports percentiles of the distribution of SNB_{it} across observations. In the case of the high-tech industry group, the table shows that the one-period payoff from R&D is small relative to the long-run payoff. For the median firm, the next period profit accounts for only between 0.9 and 2.7 percent of the total long-run payoff to R&D. Even at the 75th percentile the short-run payoff is at most 6.6 percent of the long-run payoff. In the low-tech sector, we observe that the short-run profit increase accounts for a larger fraction of the total long-term benefit. At the 75th percentile, this fraction varies from 8.1 to 21.1 percent but the short-run benefit is still small relative to the long-term payoff. This is not surprising because this short-run effect only captures the immediate effect of R&D on productivity (a 1.4 percent increase as seen in Table 4) and firm profit. It does not capture the payoff resulting from a permanently higher level of productivity in future periods or the increase this will have on the probability the firm continues to invest in R&D in the future. For at least 75 percent of the firms in these industries, these long-run impacts account for over 79 percent, and often over 90 percent, of the payoff to current period R&D. This emphasizes the need to examine R&D choice and its benefit in a dynamic framework.

One reason that the long-run benefit is much more substantial than the short-run benefit is that the productivity process is highly persistent. Given the coefficients on lagged ω reported in Table 5 we see that productivity gains resulting from innovations (or random shocks) will depreciate very slowly which implies that the gain in profits from an innovation and subsequent productivity improvement will last over many periods. As the degree of productivity persistence falls, the long-run gains, relative to the one-period gain, will decline. In the limit, if there was no productivity persistence, so that innovations only affected the current period profits, the firm would only invest in R&D if it could cover the entire fixed or sunk cost with the one-period profit increase. Everything else equal, a reduction in the α_1 coefficient in equation 5 has a large impact on the probability of conducting R&D. To illustrate this we simulate the model using the first observation for each firm in the high-tech industries and construct the mean probability of R&D across the firms after 5, 10, and 20 years. We repeat the simulation for different values of α_1. The results show that when $\alpha_1 = .961$, as reported in Table 5, the mean R&D probability stabilizes at .77. When α_1 is reduced to .941, the mean probability falls to .49, at $\alpha_1 = .923$ it falls to .28, and at $\alpha_1 = .903$ it falls to .17. Any value of $\alpha_1 < .7$ generates a probability of R&D that is virtually 0. More rapid depreciation of the productivity effect of the innovations leads to substantial reductions in the incentive to invest in R&D.

7 Counterfactual Analysis: R&D Subsidies

In the previous section we showed that the expected net benefits of investing in R&D vary across firms with differences in their productivity, capital stock, industry, and R&D history. In addition, productivity is an endogenous determinant of the R&D decision. Changes in the underlying economic environment in which the firms operate will shift the distribution of net benefits, cause some firms to make a different choice about R&D investment, and thus lead to different long-run changes in firm productivity and profits. Policy instruments such as direct subsidies to firms or tax treatment of R&D expenditures can affect the cost of R&D. Using the estimated model we simulate reductions in both the fixed and startup costs of investing in R&D.

We can simulate more favorable treatment of R&D costs for all firms by reducing the fixed cost of R&D. The top panel of Table 12 shows the effect of reducing the fixed cost by 20 percent, relative to the estimates in Table 7, for the high-tech industries. The second through fourth columns summarize the change in the cross-sectional distribution for the probability of investing in R&D after 5, 10, and 20 years experience with the lower cost environment. After 5 years we observe that the mean probability has increased by .0716 percentage points (from .7400 in the initial regime to .8116 after the cost reduction). However, the changes in the percentiles indicate that not all firms are more likely to invest in R&D. For at least 5 percent of the firms there is no increase in the probability of doing R&D. This is due to the fact that their probability of doing R&D is virtually one before the cost change so there is no impact of the lower fixed cost. The median indicates that half of the firms have an increase of at least .064 percentage points. For some firms the increase in the probability is substantial, the 75th percentile is .116. The increased rate of investment in R&D then leads to a shift in the distribution of productivity. In column (5), the mean of the firm productivity distribution increases by 4.06 percent after 5 years. The percentage improvement in productivity is not uniform across all firms. The lowest five percent of the firms have no increase in productivity, the median firm increase is .38 percent, and the top 5 percent of the firms have an increase of at least 8.77 percent. The reduction in cost is most likely to alter the investment decision of firms that have costs that are near the threshold of ΔEV while it will leave unchanged the decision of firms with very low or very high initial costs. As a result the proportional improvements in productivity are not equal across firms.

Table 12 also reports the R&D and productivity responses after 10 and 20 years of the lower cost environment. These show that the heterogeneity in responses remains even after many years. The R&D distibution does not shift substantially over the longer time period, relative to the change after five years, while the productivity distribution continues to shift to the right, especially between the 25th and 75th percentiles. This reflects the fact that once firms begin investing in R&D they can continue to realize innovations and gain further productivity improvements over time. The increase in R&D investment occurs fairly quickly and this can generate a continual series of productivity improvements.

We also examine the effect of a 20 percent reduction in the startup costs of conducting R&D in the lower panel of Table 12. The first row of the panel shows that there is very little change in the mean probability of investing in R&D. Correspondingly, the mean change in productivity across all firms is also close to zero, rising .28 percent after 5 years and .51 percent after 20 years. The percentiles of the distribution indicate that for at least half the firms the reduction in startup cost actually results in a decrease in the probability of investing in R&D. In contrast to the fixed cost reduction, which always increases the firm's probability of investing, a sunk cost reduction has two effects. It lowers the entry cost for firms that are not investing and this raises their probability of investment but it also lowers the option value of continuing to invest for firms that are conducting R&D. This can lead some firms to stop investing. Overall, the sunk cost reduction has a less powerful effect on investment incentives than the reduction in fixed cost. However, it is important to point out that the two cost changes are not equivalent in terms of the overall cost of the subsidy. The fixed cost reduction is applicable to all investing firms while the sunk cost reduction only applies to the firms that begin to invest in R&D. The latter is a much smaller number of firms in our sample.

If we disaggregate the fixed cost and sunk counterfactual results by the size of the firm we observe that the change in R&D and productivity is not uniform across the size distribution. Table 13 reports the mean change in the proportion of high-tech firms investing in R&D. The top panel reports the change for the 20 percent reduction in fixed cost. Small, medium, and large firms are distinguished where size is defined by the the capital stock and the categories each contain one-third of the firms in the industry. We observe that after 5 years the small and medium firms have a larger increase in their probability of investing than the larger firms. Initially the increase in R&D investment observed in Table 12 is more concentrated among the small and medium firms. The difference across size groups gradually diminishes over time. The lower panel of the table reports the change in mean investment when the startup cost is reduced by 20 percent. We observe that small firms increase their rate of investment when the startup cost falls while the medium and large firms reduce theirs. The negative pattern observed in Table 12 arises because the entry of small firms is outweighed by the exit of the medium and large firms. While the overall impact of the startup cost reduction is small it does not uniformly affect firms of different sizes.

8 Conclusions

Measuring the private return to R&D investment has been a major goal of the productivity literature and the knowledge production function model has been the primary empirical tool. In this framework, firm investments in R&D accumulate and depreciate over time creating a stock of knowledge within the firm that enters as an additional input in the production function. Estimates of the effect of this knowledge stock on output provide a measure of the return to R&D.

In this paper we take a different approach to measuring the private payoff from R&D investment. We develop an empirical model of the firm's dynamic demand for R&D and use it to measure the long-run expected payoff to a firm's R&D. In this model a firm's investment in R&D raises the probability that it develops a product or process innovation. The firm's productivity is an endogenous state variable whose transition path is altered by these product and process innovations. Realizations of productivity then determine the firm's profits. The expected long-run payoff from investing in R&D is the difference between the expected discounted sum of future profits if the firm undertakes R&D versus if it does not. The firm will choose to invest in R&D if this payoff is greater than the fixed or sunk cost they pay to invest.

This dynamic demand model captures several important components of the R&D and innovation process. It views R&D as an investment that raises the probability that the firm will be on a higher productivity and profit path in the future. The difference in the expected value of the firm between these two paths provides a natural measure of the expected payoff to the R&D investment. There is uncertainty at two stages of the process; whether the firm successfully introduces an innovation and the realization of the productivity level. The model also recognizes that payoffs are realized in the future and expectations of these future payoffs are critical to the firm's R&D decision. Finally, since we estimate the firm's decision rule for R&D investment we can conduct counterfactual experiments that change the economic environment in which the firm operates and simulate how this affects the firm's R&D choice and future productivity.

The empirical model is designed to exploit the micro data that is collected in the Community Innovation Surveys. For Germany, this includes firm panel data on R&D expenditures, product and process innovations realized by the firm, and variables to construct productivity and short-run profits. The four key structural components of the model are: the firm's profit function which relates productivity to profit, the evolution of firm productivity which depends upon product and process innovations realized by the firm, the probability of a product or process innovation given the firm's choice of R&D, and the fixed and startup costs of investing in R&D.

The structural parameter estimates show, first, firms that invest in R&D have a higher probability of realizing a product or process innovation but R&D investment is neither necessary nor sufficient for firm innovation. The group of high-tech manufacturing industries has a higher probability of innovation, given R&D, than the group of low-tech industries. Second, product innovation as well as process innovation lead to increases in future firm productivity but product innovations are more important in the high-tech industries while process innovations are more important for the low-tech industries. Third, firm productivity is highly persistent over time which implies that innovations that raise productivity will have long-lived effects on firm performance. Fourth, fixed costs of maintaining ongoing R&D investment are significantly smaller than the sunk startup costs of beginning to invest in R&D. This means that firm R&D history is an important determinant of current R&D choice.

Comparing across industries for the firm with the median productivity level we see that the expected benefit of investing in R&D varies from a high of 43 million euros in the vehicle industry and 20 million euros in the chemical industry to a low of about 350 thousand euros in the plastic, non-metallic mineral products, and manufacturing nec industries. This difference in the benefit of R&D will lead to very different rates of R&D investment across industries.

Combining estimates of the expected benefit of R&D with the cost of R&D we summarize the distribution of net benefits across firms in each industry. This net benefit differs substantially between firms that have already invested in R&D and those that are just starting R&D investments because of the substantial differences in the fixed versus sunk cost of investment. Expressed as a proportion of firm value, this net benefit for the median experienced firm in an industry varies from .024 to .032 across five high-tech industries but varies from -.046 to .006 for firms with no previous R&D experience. The negative value implies that the median inexperienced firm would not find it profitable to invest in R&D. We also examine the distribution of net benefits for firms that find R&D investment profitable. This distribution indicates a net benefit of .020 to .024 for startup firms in the high tech industries. These net benefits are substantially smaller, around .002 for the median firm, in the group of low-tech industries. Finally, we compare the expected short-run (one period) benefit that the firm gets from R&D with the expected long-run benefit and find that the long-run benefit is between five and 30 times larger depending on the industry. This emphasizes the need to examine R&D choice and its benefit in a dynamic framework.

The empirical model is used to conduct counterfactual experiments simulating whether an R&D subsidy leads to an increase in productivity. This question is at the heart of many discussions regarding the costs and benefits of public subsidies and we simulate different subsidy policies by changing the cost of R&D. The results show that in the high-tech industries, a 20 percent reduction in the fixed cost of R&D leads after five years to an average increase of 7.16 percentage points in the probability a firm invests in R&D and a 4.02 percent increase in mean productivity. A similar reduction in the cost faced by firms just beginning to invest in R&D has very little impact on the probability of investing or the level of productivity. The simulations illustrate that the two cost changes have very different impacts on firm incentives. Fixed cost reductions encourage all firms to invest. In contrast, the reduction in startup costs can lead some firms to stop their R&D because the cost of restarting is now lower.

References

[1] Aw, B.Y., M.J. Roberts, and D.Y. Xu (2011), R&D Investment, Exporting and Productivity Dynamics, *The American Economic Review*, Vol. 101, No. 4 (June), pp. 1312-1344.

[2] Bernstein, J.I. and M.I. Nadiri (1989), Research and Development and Inter-Industry Spillovers: An Empirical Application of Dynamic Duality, *The Review of Economic Studies*, Vol. 56, pp. 249-269.

[3] Bernstein, J.I. and M.I. Nadiri (1991), Product Demand, Cost of Production, Spillovers, and the Social Rate of Return to R&D, NBER Working Paper w3625, Cambridge, MA.

[4] Crépon, B., E. Duguet, and J. Mairesse (1998), Research Innovation and Productivity: An Econometric Analysis at the Firm Level, *Economics of Innovation and New Technology*, Vol. 7, No. 2, pp. 115-158.

[5] Czarnitzki, D., B.H. Hall, and R. Oriani (2006), The Market Valuation of Knowledge Assets in US and European Firms, in *The Management of Intellectual Property*, D. Bosworth and E. Webster (eds.), Cheltenham, UK: Edgar Elgar, pp. 111-131.

[6] Dixit, A.K. and R.S. Pindyck (1994), *Investment Under Uncertainty*, Princeton: Princeton Unversity Press.

[7] Doraszelski, U. and J. Jaumandreu (forthcoming), R&D and Productivity: Estimating Endogenous Productivity, *Review of Economic Studies.*

[8] Foster, L., J. Haltiwanger, and C. Syverson (2008), Reallocation, Firm Turnover and Efficiency: Selection on Productivity or Profitability, *American Economic Review*, Vol. 98, No. 1, pp. 394-425.

[9] Griliches, Z. (1979) Issues in Assessing the Contribution of Research and Development to Productivity Growth, *Bell Journal of Economics*, Vol. 10, No. 1 (Spring), pp. 92-116.

[10] Griliches, Z. (1992), The Search for R&D Spillovers, *The Scandinavian Journal of Economics*, Vol 94, pp. 29-47.

[11] Griliches, Z. (1998), R&D and Productivity: The Unfinished Business, in *R&D and Productivity: The Econometric Evidence*, Chapter 12. Chicago: The University of Chicago Press.

[12] Hall, B.H. (1996), The Private and Social Returns to Research and Development, in *Technology, R&D, and the Economy*, B.L.R. Smith and C.E. Barfield (eds.), The Brookings Institution, Washington D.C.

[13] Hall, B.H. (2011), Innovation and Productivity, NBER Working Paper No. 17178, Cambridge, MA.

[14] Hall, B.H., J. Mairesse, and P. Mohnen (2010), Measuring the Returns to R&D, in *Handbook of the Economics of Innovation,* B.H. Hall and N. Rosenberg (eds.), Vol. 2, Chapter 22, Elsevier, pp. 1033-1082.

[15] Levinsohn, J. and A. Petrin (2003), Estimating Production Function Using Inputs to Control for Unobservables, *Review of Economic Studies*, Vol. 70, No. 2 (April), pp. 317-342.

[16] Mairesse, J. and M. Sassenou (1991), R&D and Productivity: A Survey of Econometric Studies at the Firm Level, *STI Review*, OECD, Vol. 8, pp. 9-46.

[17] Mairesse, J. and P. Mohnen (2010), Using Innovation Surveys for Econometric Analysis, in *Handbook of the Economics of Innovation*, B.H. Hall and N. Rosenberg (eds.), Vol. 2, Chapter 26, Elsevier, pp. 1129-1155.

[18] OECD (1992, 1997, 2005), *Oslo Manual: Proposed Guidelines for Collecting and Interpreting Technological Innovation Data, 1st, 2nd and 3rd edn.*, Paris.

[19] Olley, G.S. and A. Pakes (1996), The Dynamics of Productivity in the Telecommunications Equipment Industry, *Econometrica*, Vol. 64, No. 6 (November), pp. 1263-1297.

[20] Rammer, C. and B. Peters (2013), Innovation Panel Surveys in Germany, in *Handbook of Innovation Indicators and Measurement*, Fred Gault (ed.), Edward Elgar: Cheltenham, UK and Northampton, MA, USA, pp. 135-177.

[21] Roberts, M.J. and V.A. Vuong (2013), Empirical Modeling of R&D Demand in a Dynamic Framework, *Applied Economic Perspectives and Policy*, Vol. 35, No. 2 (June), pp. 185-205.

[22] Roberts, M.J., D.Y. Xu, X. Fan, and S. Zhang (2012), A Structural Model of Demand, Cost and Export Market Selection for Chinese Footwear Producers, NBER Working Paper 17725, Cambridge, MA.

[23] Rogers, M. (2010), R&D and Productivity: Using UK Firm-Level Data to Inform Policy, *Empirica*, Vol 37, pp. 329-359.

[24] Rust, J. (1987), Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher, *Econometrica*, Vol. 55, No. 5 (September), pp. 999-1033.

[25] Xu, D.Y. (2008), A Structural Empirical Model of R&D Investment, Firm Heterogeneity and Industry Evolution, Working Paper, Duke University.

Table 1: Innovation Rates by Industry - pooled over firms and years

Industries	Proportion of Innovating Firms rd	Proportion with Product Innovation d	Proportion with Process Innovation z
High-Tech			
Chemicals	0.7866	0.7081	0.3633
Machinery	0.7702	0.7147	0.3609
Electronics	0.8053	0.7449	0.3977
Instruments	0.8176	0.7706	0.3300
Vehicles	0.7309	0.6504	0.3955
Low-Tech			
Food	0.5425	0.4732	0.2580
Textiles	0.5135	0.4643	0.2027
Paper	0.5174	0.3919	0.2453
Plastic	0.6422	0.5915	0.3266
Mineral	0.5887	0.5257	0.3113
Basic Metals	0.5938	0.4785	0.3164
Manuf. nec	0.6060	0.5283	0.2697
Average	0.6596	0.5868	0.3148

Table 2: Transition Rates for R&D Investment

	$rd_{it} = 0$	$rd_{it} = 1$	Capital Range[a]
$rd_{it-1} = 0$	.813	.188	[0, .15]
	.824	.176	(.15, .42]
	.772	.228	(.42, .92]
	.745	.255	(.92, 1.75]
	.787	.214	(1.75, 3.04]
	.717	.283	(3.04, 5.49]
	.663	.337	(5.49, 10.83]
	.669	.331	> 10.83
$rd_{it-1} = 1$	.218	.783	[0, .15]
	.194	.806	(.15, .42]
	.215	.785	(.42, .92]
	.169	.832	(.92, 1.75]
	.172	.828	(1.75, 3.04]
	.143	.857	(3.04, 5.49]
	.081	.919	(5.49, 10.83]
	.068	.932	> 10.83

a millions of Euros

Table 3: Probability of Innovation Conditional on Past R&D: $\Pr(d_{t+1}, z_{t+1} \mid rd_t)$								
	$rd_t = 0$				$rd_t = 1$			
Product Innovation	$d = 0$	$d = 1$	$d = 0$	$d = 1$	$d = 0$	$d = 1$	$d = 0$	$d = 1$
Process Innovation	$z = 0$	$z = 0$	$z = 1$	$z = 1$	$z = 0$	$z = 0$	$z = 1$	$z = 1$
High-Tech Industries								
Chemicals	0.779	0.048	0.048	0.124	0.112	0.214	0.045	0.629
Machinery	0.786	0.055	0.039	0.120	0.100	0.249	0.035	0.616
Electronics	0.716	0.092	0.028	0.163	0.100	0.262	0.029	0.609
Instruments	0.779	0.044	0.035	0.142	0.090	0.317	0.010	0.582
Vehicles	0.783	0.058	0.050	0.108	0.139	0.172	0.052	0.638
Low-Tech Industries								
Food	0.767	0.047	0.043	0.142	0.243	0.170	0.047	0.540
Textiles	0.791	0.072	0.038	0.099	0.251	0.247	0.054	0.448
Paper	0.782	0.038	0.082	0.097	0.271	0.136	0.141	0.453
Plastic	0.786	0.079	0.020	0.115	0.148	0.171	0.045	0.636
Mineral	0.776	0.068	0.021	0.135	0.188	0.156	0.043	0.613
Metals	0.822	0.023	0.041	0.113	0.171	0.123	0.115	0.590
Manuf. nec	0.775	0.085	0.035	0.106	0.170	0.258	0.066	0.507
Average	0.779	0.059	0.040	0.122	0.165	0.206	0.057	0.572

Table 4: Demand Elasticity Estimates (standard error)			
Industry	$1{+}1/\eta$	η	sample size
High-Tech			
Chemicals	0.708 (0.005)**	-3.425	1361
Machinery	0.803 (0.002)**	-5.076	2644
Electronics	0.753 (0.005)**	-4.049	1413
MPO	0.763 (0.006)**	-4.219	1429
Vehicles	0.874 (0.003)**	-7.937	911
Low-Tech			
Food	0.666 (0.008)**	-2.994	1162
Textiles	0.697 (0.003)**	-3.300	990
Paper	0.697 (0.003)**	-3.300	1669
Plastic	0.798 (0.003)**	-4.950	1396
Mineral	0.675 (0.005)**	-3.077	959
Metals	0.822 (0.001)**	-5.618	2773
Manuf. nec	0.765 (0.004)**	-4.255	872

** significant at the .01 level, * significant at the .05 level

Table 5: Productivity Evolution Parameters (standard error)		
	High-Tech Industries	Low-Tech Industries
k	-0.056 (0.002)**	-0.060 (0.002)**
ω_{t-1}	0.961 (0.008)**	0.978 (0.005)**
ω_{t-1}^2	0.030 (0.012)*	0.006 (0.008)
ω_{t-1}^3	-0.008 (0.005)	0.001 (0.004)
d	0.013 (0.005)**	0.002 (0.004)
z	0.014 (0.008)	0.010 (0.005)*
$d * z$	-0.014 (0.009)	-0.002 (0.007)
$intercept$	0.010 (0.003)**	0.010 (0.002)**
$SE(\varepsilon)$	0.1010	0.1088
sample size	3337	4298

** significant at the .01 level, * significant at the .05 level

Table 6: Elasticity of Revenue w.r.t. R&D

High-Tech Industries	
Chemicals	0.021
Machinery	0.036
Electronics	0.024
Instruments	0.029
Vehicles	0.058
Low-Tech Industries	
Food	0.008
Textiles	0.009
Paper	0.010
Plastic	0.022
Mineral	0.011
Metals	0.026
Manuf. nec	0.015

Table 7: Dynamic Parameter Estimates: Fixed Cost and Startup Cost

	Fixed Cost	Startup Cost
High-Tech Industries		
Small Firms	0.655 (0.025) **	3.980 (0.216) **
Medium Firms	1.933 (0.055) **	12.215 (1.367) **
Large Firms	4.544 (0.154) **	26.840 (1.009) **
Low-Tech Industries		
Small Firms	0.368 (0.018) **	1.540 (0.300) **
Medium Firms	0.907 (0.037) **	3.986 (0.372) **
Large Firms	1.675 (0.016) **	8.262 (0.066) **

** significant at the .01 level, * significant at the .05 level

Table 8: Benefits and Costs of Conducting R&D for High-Tech Industries						
	ω	$\Delta EV(\omega)$	$E(\gamma\|\gamma < \Delta EV)$		$Pr(rd_t = 1)$	
Industry			$rd_{t-1} = 1$	$rd_{t-1} = 0$	$rd_{t-1} = 1$	$rd_{t-1} = 0$
Chemicals	-0.299	0.965	0.437	0.475	0.378	0.088
	0.316	6.328	1.949	2.921	0.877	0.377
	0.849	20.506	2.921	8.136	0.998	0.731
	1.251	37.870	2.971	12.282	1.000	0.896
	2.053	87.131	2.972	16.702	1.000	0.990
Machinery	-0.227	1.685	0.709	0.819	0.587	0.154
	0.072	5.456	1.654	2.520	0.907	0.382
	0.301	9.572	2.112	4.191	0.980	0.548
	0.563	15.289	2.311	6.196	0.997	0.705
	0.886	19.616	2.347	7.476	0.999	0.785
Electronics	-0.296	3.367	1.273	1.610	0.724	0.231
	0.048	9.028	2.319	4.055	0.946	0.466
	0.332	16.882	2.766	6.950	0.994	0.669
	0.765	37.807	2.868	12.119	1.000	0.901
	1.445	111.846	2.869	16.709	1.000	0.997
Instruments	-0.458	0.396	0.187	0.196	0.267	0.053
	-0.078	1.620	0.640	0.780	0.648	0.193
	0.204	3.499	1.135	1.630	0.849	0.343
	0.565	7.653	1.754	3.351	0.972	0.553
	0.944	10.773	1.948	4.500	0.992	0.660
Vehicles	-0.071	14.532	2.568	5.919	0.896	0.426
	0.090	28.995	2.959	9.607	0.983	0.632
	0.242	43.650	3.060	11.995	0.997	0.756
	0.391	56.940	3.081	13.433	0.999	0.824
	0.581	61.831	3.084	13.831	1.000	0.842

Table 9: Benefits and Costs of Conducting R&D for Low-Tech Industries

Industry	ω	$\Delta EV(\omega)$	$E(\gamma \mid \gamma < \Delta EV)$		$Pr(rd_t = 1)$	
			$rd_{t-1} = 1$	$rd_{t-1} = 0$	$rd_{t-1} = 1$	$rd_{t-1} = 0$
Food	-0.600	0.047	0.023	0.024	0.058	0.014
	-0.065	0.120	0.059	0.060	0.141	0.035
	0.590	0.409	0.187	0.200	0.399	0.115
	1.296	1.884	0.631	0.864	0.857	0.404
	2.031	4.009	0.913	1.687	0.978	0.632
Textiles	-0.569	0.040	0.020	0.020	0.058	0.014
	-0.190	0.096	0.047	0.048	0.132	0.032
	0.532	0.414	0.187	0.202	0.451	0.136
	0.957	1.074	0.407	0.506	0.759	0.307
	1.360	2.198	0.641	0.977	0.929	0.504
Paper	-0.554	0.045	0.022	0.022	0.059	0.014
	-0.108	0.122	0.059	0.061	0.152	0.038
	0.492	0.414	0.188	0.203	0.424	0.124
	0.983	1.250	0.467	0.587	0.778	0.320
	1.492	2.826	0.758	1.236	0.954	0.551
Plastic	-0.266	0.060	0.030	0.030	0.070	0.016
	-0.010	0.202	0.097	0.100	0.216	0.054
	0.204	0.358	0.166	0.176	0.349	0.094
	0.496	0.624	0.273	0.303	0.520	0.158
	0.761	0.594	0.262	0.289	0.503	0.151
Mineral	-0.653	0.048	0.024	0.024	0.059	0.014
	-0.129	0.111	0.054	0.055	0.132	0.032
	0.457	0.324	0.151	0.160	0.337	0.092
	0.906	0.824	0.343	0.396	0.624	0.217
	1.706	2.304	0.709	1.038	0.908	0.465
Metals	-0.250	0.130	0.063	0.065	0.115	0.027
	-0.007	0.476	0.216	0.233	0.352	0.094
	0.178	0.872	0.363	0.420	0.535	0.163
	0.397	1.504	0.547	0.707	0.719	0.262
	0.684	2.061	0.668	0.946	0.817	0.339
Manuf. nec	-0.369	0.054	0.027	0.027	0.071	0.017
	-0.035	0.175	0.084	0.087	0.210	0.053
	0.277	0.356	0.164	0.175	0.377	0.106
	0.576	0.628	0.272	0.304	0.558	0.178
	0.954	0.631	0.273	0.306	0.559	0.179

Table 10: Long-Run Return to R&D Given R&D History							
		Continuing Firms			Startup Firms		
		25th	50th	75th	25th	50th	75th
High-Tech Industries							
Chemicals	NB	0.021	0.032	0.037	-0.045	0.001	0.016
	TNB	0.023	0.033	0.037	0.019	0.024	0.027
Machinery	NB	0.024	0.031	0.035	-0.042	-0.013	0.002
	TNB	0.026	0.032	0.035	0.020	0.023	0.026
Electronics	NB	0.027	0.032	0.037	-0.014	0.006	0.020
	TNB	0.027	0.032	0.037	0.020	0.024	0.027
Instruments	NB	0.008	0.024	0.029	-0.095	-0.046	-0.019
	TNB	0.018	0.026	0.030	0.015	0.020	0.023
Vehicles	NB	0.020	0.028	0.031	-0.015	0.003	0.011
	TNB	0.020	0.029	0.031	0.015	0.020	0.022
Low-Tech Industries							
Food	NB	-0.018	-0.007	0.001	-0.087	-0.044	-0.012
	TNB	0.002	0.002	0.003	0.001	0.002	0.003
Textiles	NB	-0.014	-0.006	0.001	-0.074	-0.040	-0.015
	TNB	0.002	0.002	0.003	0.002	0.002	0.003
Paper	NB	-0.014	-0.005	0.000	-0.074	-0.039	-0.019
	TNB	0.002	0.002	0.003	0.002	0.002	0.003
Plastic	NB	-0.011	-0.006	-0.003	-0.061	-0.041	-0.029
	TNB	0.002	0.002	0.002	0.002	0.002	0.002
Mineral	NB	-0.016	-0.008	-0.001	-0.081	-0.048	-0.022
	TNB	0.002	0.002	0.003	0.002	0.002	0.002
Metals	NB	-0.006	-0.002	0.000	-0.042	-0.029	-0.018
	TNB	0.002	0.003	0.003	0.002	0.003	0.003
Manuf. nec	NB	-0.011	-0.006	-0.003	-0.063	-0.044	-0.028
	TNB	0.002	0.002	0.002	0.002	0.002	0.002

Table 11: Ratio of Short-Run to Long-Run Return on R&D			
	25th percentile	50th percentile	75th percentile
High-Tech Industries			
Chemicals	0.015	0.017	0.023
Machinery	0.012	0.018	0.033
Electronics	0.007	0.009	0.014
Instruments	0.015	0.020	0.030
Vehicles	0.014	0.027	0.066
Low-Tech Industries			
Food	0.075	0.083	0.089
Textiles	0.059	0.069	0.081
Paper	0.065	0.077	0.087
Plastic	0.087	0.110	0.211
Mineral	0.087	0.103	0.113
Metals	0.063	0.079	0.129
Manuf. nec	0.070	0.096	0.150

Table 12: Counterfactual Reductions in Cost for High-Tech Industries

	Change in R&D Proportion			Proportional Change in Productivity		
Percentile	5 years	10 years	20 years	5 years	10 years	20 years
			20 Percent Reduction in Fixed Cost			
Mean	0.0716	0.0709	0.0658	0.0406	0.0559	0.0494
5th	0.0000	0.0000	0.0020	0.0000	0.0000	0.0002
25th	0.0180	0.0300	0.0340	0.0005	0.0023	0.0056
Median	0.0640	0.0740	0.0660	0.0038	0.0110	0.0185
75th	0.1160	0.1080	0.0980	0.0122	0.0297	0.0383
95th	0.1660	0.1440	0.1280	0.0877	0.1593	0.1128
			20 Percent Reduction in Sunk Startup Cost			
Mean	-0.0007	0.0002	-0.0009	0.0028	0.0018	0.0051
5th	-0.0220	-0.0160	-0.0160	-0.0065	-0.0046	-0.0032
25th	-0.0100	-0.0060	-0.0060	-0.0010	-0.0013	-0.0016
Median	-0.0020	-0.0020	-0.0020	-0.0003	-0.0003	-0.0004
75th	0.0060	0.0060	0.0040	0.0002	0.0015	0.0015
95th	0.0240	0.0220	0.0180	0.0086	0.0148	0.0141

Table 13: Mean Change in R&D by Firm Size

Firm Size	Change in R&D Proportion 5 years	10 years	20 years
	20% Reduction in Fixed Cost		
Small	0.0777	0.0745	0.0621
Medium	0.0782	0.0747	0.0704
Large	0.0589	0.0635	0.0651
	20% reduction in Sunk Startup Cost		
Small	0.0027	0.0021	0.0001
Medium	-0.0009	0.0003	-0.0005
Large	-0.0040	-0.0017	-0.0022