Empirical Modeling of R&D Demand in a Dynamic Framework*

Mark J.Roberts
Department of Economics
The Pennsylvania State University and NBER

Van Anh Vuong
Department of Economics
University of Cologne and Institute for Energy Economics (EWI)

February 2013

Abstract

Empirical analysis of firm-level investment in R&D and its effect on innovation patterns and productivity has advanced as a result of the collection of innovation surveys in many countries. The weak link in the analysis of the innovation surveys is the empirical model of firm R&D choice. In this paper we summarize how a dynamic, structural model of firm invesment can be used to estimate firm demand for R&D with the data collected in innovation surveys. The estimates provide a natural measure of the expected benefit to the firm of investing in R&D and allow the researcher to simulate how the firm's R&D investment will respond to cost or demand changes.

Key words: R&D demand, innovation, productivity, dynamic structural model
JEL codes: O31, O32, L60

1 Introduction

The study of a firm's investment in research and development (R&D) and its impact on productivity improvement has long been unified around the "knowledge production function" framework of Griliches (1979). In that framework firm's investments in R&D both cumulate and depreciate over time creating a stock of knowledge capital which enters as an input in the firm's production function. A large empirical literature, which is comprehensively surveyed in Hall, Mairesse, and Mohnen (2010), has focused on measuring the relationship between firm sales

*The authors are grateful to the editor, Spiro Stefanou, two anonymous referees, Florin Maican, Jacques Mairesse, and Bettina Peters for helpful comments and discussions. We thank The Center for European Economic Research (ZEW) for providing access to the data and research support.

and knowledge capital and using the estimates to construct a rate of return to R&D investment. The goal of measuring both the private and social rate of return to R&D has led to the development of more complex models describing the relationship between R&D, innovation outputs, and productivity, improved econometric methods for dealing with the endogeneity of R&D and measurement errors, and better data collection especially at the micro level. In this article we will focus on a small part of this literature: the application of dynamic structural models to measure the expected benefits of R&D investment and simulate investment responses to changes in R&D policy.

In the last decade a number of European countries have implemented innovation surveys that are designed to elicit firm-level information on innovation outcomes and the firms' expenditures on innovation activities, broadly defined, in order to help researchers and policy makers better understand a country's productivity performance. The innovation surveys have been used to develop scoring formulas providing a baseline for policy makers to evaluate the effectiveness and the progress of a country's or industry's innovation activities (EU (2011)). They are unique relative to most R&D surveys because they recognize that R&D expenditure may not be the only pathway through which firms' discover and implement new product or process innovations. The hope is that collecting data on both the incidence and magnitude of firm innovations will provide additional insights into the pathways that lead from firm investments on R&D and innovation to actual productivity or profit improvements by the firm.

A large empirical literature has developed analyzing the data in the innovation surveys. Most studies are built around the four-equation model of Crépon, Duguet, and Mairesse (1998) (hereafter, CDM) and the refinements in Mairesse, Mohnen, and Kremp (2005). This provides a framework specifying the linkages between R&D inputs, outcomes of the innovation process, and productivity. While very useful for organizing and identifying the correlations in these variables, this framework does not attempt to model the firm's demand for R&D in a way that exploits all the information that is available in the innovation and productivity data and cannot be used to conduct analysis of alternative policy options such as R&D subsidies.

In this article we will outline a more richly specified dynamic, structural model of the firm's demand for R&D that utilizes the information collected in the innovation surveys. The demand results from the firm choosing R&D to maximize the sum of discounted expected future profits while recognizing the impact of the R&D choice on future productivity and profits. By estimating this structural demand curve we are able to infer the expected benefit to the firm's investment in R&D. The structural demand model also allows us to change parameters in the firm's environment, for instance the degree of competition in the output market or introduction of R&D subsidies, and quantify how these changes affect the firm's decision to invest in R&D, its productivity, and the long-run impact on profitability. As such they provide a way to analyze the effects of policy decisions that affect the costs or benefits of R&D to the firms.

This survey will summarize alternative approaches to measuring the linkages between R&D, innovation, and productivity, how a detailed specification of the firm's choice of R&D affects the way the innovation survey data is utilized and the type of inference possible. We will illustrate the R&D demand model with an application to the Community Innovation Survey data from a set of German manufacturing industries.

2 Analysis of R&D Using Innovation Surveys

The Community Innovation Surveys (CIS) refer to a collection of internationally-harmonized firm-level surveys conducted separately in European countries. They rely on the Oslo Manual (OECD, 1992, 1996, 2005) to provide a common framework for defining innovation activities, both inputs and outputs of the innovation process, and so provide a basis for cross-country comparisons. These surveys started in 1993 and were initially conducted every four years with each wave covering a three-year time span. Since 2005 they are conducted every two years. In most countries each wave is a cross-sectional sample with partly overlapping firms across waves, although Germany, in particular, designed the survey as a panel of firms. The surveys cover firms from the manufacturing, mining, energy, and most service sectors and in many countries can be linked to other producer surveys. Mairesse and Mohnen (2010) describe the survey framework and discuss differences across countries and the limitations of the surveys in policy analysis and research work.

The surveyed firms are generally asked to report their revenue, number of employees, and investment in physical capital. The unique variables collected are input and output measures of the innovation process. In particular, each firm reports its R&D spending but also spending on a broader set of innovation activities including worker training in this area, acquistion of external knowledge, and capital, marketing, and design expenditures involved in developing a new product or introducing a new production process.

The firms also report whether they have realized any of four types of innovations, product, process, organizational, and marketing, during the time period under review. 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. Organizational innovation is defined as new business practices, workplace organization, or external relations and marketing innovation refers to changes in product design, packaging, product placement or promotion, and pricing methods. 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 variables indicating a product and process innovation are the most consistent measures of innovative outputs and have been collected in many waves of the CIS while the variables indicating organizational and marketing innovation have been included only since 2005. Some surveys also collect quantitative measures of innovation outputs such as the percentage reduction in cost due to a new production process or the percentage of firm revenue due to a newly-introduced product.

Based on the information the firms report on their characteristics (industry, region), spending on production inputs (labor, materials, and capital investment) and innovation inputs and outputs, the CIS data allows researchers to link the firm's invesment in R&D to its innovation outcomes and ultimately to revenue or profits. Crépon, Duguet and Mairesse (1998) introduce a conceptual framework that matches very naturally with the CIS data and has been widely used in its analysis. There are already several good surveys of this literature including Hall, Mairesse, and Mohnen (2010) and Hall (2011) and we only provide an overview so that we can contrast it with the structural demand model we present in the next section.[1]

The CDM model describes three linkages that relate innovation choice to firm performance.

[1]We will also not discuss estimation issues, econometric methods, or functional forms used in the estimating equations. Our focus is on how the basic relationship between the variables is specified.

The first relationship describes the firm's choice of innovation inputs, generally R&D expenditure, as a function of firm and industry variables that are likely to shift the long-run return to R&D:

$$\begin{aligned} rd_i &= f_1(x_{di}) \\ R/Z_i &= f_2(x_{ki}). \end{aligned} \tag{1}$$

Here rd is the discrete decision variable for whether firm i invests in R&D and R/Z is a measure of the firm's research intensity, such as the R&D to sales ratio or R&D per worker. x_d and x_k are vectors of firm and industry characteristics explaining rd and R/Z, such as the number of workers, the market share of the firm, and the industry it belongs to. In empirical analyses researchers have employed numerous measures for innovation choice including actual spending on R&D, the broader measure of innovation spending, and R&D intensity. The choice can be modeled as a discrete decision or a mixture of discrete and continuous choices.

The second relationship, often referred to as the knowledge production function, links the firm's inputs of R&D to innovation outcomes:

$$n_i = h(r_i, x_{ni}). \tag{2}$$

Measures of innovation outcomes n are expressed as a function of R&D choice r_i, which is either the discrete rd_i or the continuous R/Z_i measure and a vector of factors x_n affecting innovating success such as firm or industry characteristics or other input choices of the firm. In addition to the discrete and continuous innovation outcome variables collected in the CIS and discussed in the last section, n could also be a measure of the number or value of patents held by the firm.

The third relationship links innovation success to firm performance measures ω_i :

$$\omega_i = g(n_i, x_{\omega i}). \tag{3}$$

Firm performance, generally sales, profits, labor productivity or total factor productivity (tfp), is modeled as a function of the innovation outcome n_i and other firm-specific factors such as capital stock, physical investment, or number of workers, x_ω.

This framework has provided the conceptual structure that many researchers have relied upon to analyze innovation survey data. Individual studies using this framework have extended it in various ways to incorporate institutional features, exploit variables collected in a specific survey, and deal with econometric issues including endogeneity of the R&D choice, unobserved firm characteristics that link the equations, and sample selection bias.[2] While the empirical studies are primarily cross-sectional, linking contemporaneous innovation inputs, outputs, and firm performance, there is a (complicated) panel structure to the data that has allowed some researchers to incorporate dynamic aspects of the innovation process by exploiting lags in the variables.[3]

Focusing on the studies using innovation data that are surveyed in Hall (2011) or summarized in OECD (2010), one of the first features characterizing the data is that R&D investments are neither necessary nor sufficient for a firm to realize an innovation. The probability of a firm introducing a new product to the market, if the firm invests in R&D, varies from 40 to 65 percent across countries, while the same probability for firms that do not invest in R&D is generally between 10 and 30 percent (OECD, 2010, chapter 1). While this suggests that R&D investment does matter in the innovation process it also indicates room for randomness or other systematic determinants in the linkage between R&D and innovation. This in turn has implications for the specification of equation (2). In particular, we want to allow firms to report innovations even when they do not have any formal R&D or innovation expenditures. A second finding is that R&D expenditure is more highly correlated with product innovations than process innovations and this suggests it will be important to distinguish the two types of innovations in equation (2). A third finding is that tfp is positively and significantly related to product innovations. In the studies using European data, the elasticity of tfp with respect to the share of the firm's revenue resulting from the new product introduction varies from .05 to .29 and is higher in the high-tech industries. This positive correlation holds up when product innovation is measured as a discrete variable. A fourth finding is that the correlation of process

[2] See Mairesse, Mohnen, and Kremp (2005) for analysis of the econometric issues that arise when applying the model to the CIS data.

[3] Peters (2009) estimates a dynamic random effect probit model of the innovation process using data in the Mannheim Innovation Panel, which is the source of the German CIS data. She discusses many aspects of the panel structure of the data.

innovation and tfp is much weaker, often zero and sometimes negative. The latter is possible if process innovations lead to revenue reductions. This could happen if the lower costs pass through into lower output prices and the firm faces inelastic demand. Alternatively, firms with low productivity levels may be forced to adopt process innovations in order to survive. One implication for the econometric model is that it is desirable to distinguish product and process innovations in the specification of equation (3). Fifth, the continuous measures of product and process innovation in the CIS are more likely to be impacted by measurement error than the discrete indicators of innovation or R&D variables and this can contribute to an attenuation of the estimated impact of innovation on productivity in equation (3).

Focusing on the broader set of micro data studies directly linking R&D and productivity, Hall, Mairesse, and Mohnen (2010) report a strong positive correlation between R&D investment and firm productivity. This reflects a combination of the equations (2) and (3) of the CDM model. The elasticity of tfp with respect to R&D intensity ranges from .01 to .25 and is centered on .08. The elasticities tend to be larger for firms in the manufacturing and high-tech sectors and lower for the service sector and low-tech manufacturing sectors. These estimates are used to compute a gross rate of return to R&D investment which is defined as the marginal product of knowledge capital input in the production function. Estimates in developed countries are high, ranging up to 75 percent, but are generally in the range of 20 to 30 percent (Hall, Mairesse, and Mohnen, 2010, Tables 2, 3). This literature focuses on a return to R&D that is based on the impact of knowledge capital in the production function. In the next section we will summarizere a dynamic model of R&D investment that leads to a different measure of the long run impact of R&D and can be interpreted as the expected payoff that a firm anticipates when it undertakes an R&D program.

3 A Structural Framework for R&D Investment

R&D investment is inherently a dynamic decision since the firm must incur costs in the present period for an anticipated gain in profits in future periods. There is likely to be a time lag between the firm's R&D expenditure, the new product or process innovations it produces, and measureable gains in productivity, sales, or profits. In addition, once innovations are realized, their impact on the firm may persist for a long time, changing the firm's environment and altering future investment decisions as well. R&D investment is unlikely to have a one-shot impact on the firm's performance. The magnitude of the gains from R&D investment may also be subject to a large random component that the firm cannot forsee. This randomness can arise at different stages. First, whether the R&D investment will produce any innovation at all and, if it does, which type of innovation is difficult to know at the time of investment. Second, the impact of a new product or a change in organization or production process on the firm's future profit is also difficult to predict. For both reasons the economic benefit of R&D investment to the firm is likely to be uncertain. In this section we outline a dynamic structural model that accounts for both the uncertainty and the intertemporal nature of the payoffs from R&D investment.[4] This model can be estimated with the type of micro data collected in the CIS and used to measure the expected costs and benefits for a firm choosing to invest in R&D.

Dynamic structural models, beginning with Rust's (1987) model of a discrete investment decision for a firm, have been applied to a wide range of topics and used to evaluate policy options. Keane and Wolpin (1997) employ a dynamic model of discrete choice to study the effect of tuition reductions on school attendance and the occupation decisions of individuals. Rust and Phelan (1997) model the effect of social security and medicare on older workers' retirement decisions. In health economics, Gilleskie (1998) uses a dynamic structural model to study the impact of changes in health insurance coverage on doctor visits for medicare patients. More recent works involve the assessment of export decisions in international trade (Das, Roberts, and Tybout (2007)), entry and exit choices in industrial organization (Aguirregabiria and Mira (2007), and Collard-Wexler (forthcoming)) and firm investment (Ryan (2012)).[5]

In the area of firm decisions, the goal of a structural model is to estimate the primitives of technology and production costs, entry costs, and consumer preferences from data on the observed output, input, pricing, and investment choices of the firm. One of the advantages of a structural model is that it can be used to simulate the effects of changes in the firms' costs,

[4] There are many theoretical models of R&D investment in the literature that account for these factors. Aghion and Howitt (1992), Ericson and Pakes (1995), Klette and Griliches (2000), and Klette and Kortum (2004) are prominent examples. The focus of this review is on the empirical specification of a dynamic model of R&D.

[5] See Wolpin(1996), Aguirregabiria and Mira (2010) for general discussions of discrete dynamic models and Reiss and Wolak (2007) for a comprehensive overview of structural modeling in industrial organization.

or technology on the firms' decisions. For example, Nevo and Whinston (2010) discuss how structural estimates of firm cost and demand parameters can be used to simulate the effect of a merger on price and output decisions. The key requirement is that the parameters of the estimated model must not change when there are changes in the economic environment.

In the area of R&D investment there is a small number of papers that use micro data to estimate structural models. Aw, Roberts, and Xu (2011) use Taiwanese firm data to model the joint decision of a firm to export and to invest in R&D, where both choices are alternative paths to build expertise and improve productivity. Xu (2008) uses Korean firm data to estimate both a direct effect of the firm's own R&D and a spillover effect of their R&D onto the productivity of other producers. Peters, Roberts, Vuong, and Fryges (2013) (herafter PRVF) use micro data collected as part of the German CIS to estimate a model of the firm's discrete decision to undertake an R&D program. By explicitly modelling the firms' behavior these papers can then study, for example, how a change in the tax treatment of R&D expenditures, such as an investment tax credit changing the cost of R&D, could impact the firm's investment and innovation. In this section we will outline the model used in PRVF and discuss how it can be used to measure the long-run benefits of R&D investment to the firm and simulate the effect of a change in R&D costs.

3.1 The Basic Assumptions

Consider a single firm i making input and output choices at the beginning of time period t. The firm faces a logarithmic production function

$$y_{it} = \alpha_0 + \alpha_l l_{it} + \alpha_k k_{it} + \omega_{it} + \nu_{it}$$

y is log output, l is a set of variable inputs, such as labor, materials, and energy in logs, k is the log of capital or other fixed factors, and ω is a firm-specific productivity level which the firm observes and will be able to affect in the future through its choice of R&D. The random shock ν is a stochastic factor the firm cannot control and does not observe until after it has made its input decisions. Given its productivity level, fixed production factors, and exogenous input prices, the firm chooses its variable inputs to maximize the current period profit which we represent by the profit function $\pi(\omega_{it})$.[6] For simplicity we have written productivity as arising from just the production side. In general the output variable in most empirical models is the firm's sales deflated by a common output price index. Klette and Griliches (1996) show that if firms produce differentiated products and have different output prices then productivity measures based on sales production functions will also capture differences in demand characteristics across producers. The important point for this discussion is that the firm productivity level ω directly impacts firm profit and the firm will be able to impact its future productivity through its investments in R&D.[7]

Next, the firm makes a decision whether to invest in R&D, r_{it}, to improve the level of its future productivity. Firm i will incur an investment cost in the current period which is given by $c(r_{it})$. The benefits of R&D will be treated in two steps. First the firm's choice of r_{it} affects the probability of realizing an innovation in the next period n_{it+1}. Denote $F(n_{it+1}|r_{it})$ as the probability distribution of future innovation given the firm's choice of r_{it}. This captures the same part of the productivity process as equation (2) in the CDM model. Second, innovations can lead to improvements in the firm's future productivity and this will be represented by a distribution function $G(\omega_{it+1}|\omega_{it}, n_{it+1})$ that depends on both the firm's current productivity level and its realized innovations. This is analogous to equation (3) in the CDM model. Given this structure we can derive the firm's decision rule for r_{it} but before that it is useful to comment on features of this setup.

First, the cost $c(r_{it})$ includes the expenditure on employees and materials used in the R&D operations but, more generally, it also includes any adjustment costs which the firm faces in starting or maintaining its R&D investment and which are unlikely to be directly measured in the CIS data. These adjustment costs are likely to differ between firms that are maintaining one or more ongoing R&D projects and firms that are just starting to undertake formal R&D.

[6] The profit function will also be dependent on the fixed factor k_{it} and the exogenous input prices. We will ignore these factors in order to focus on the role of productivity and R&D in the explanation that follows. The more complete model in Peters, Roberts, Vuong, and Fryges (2013) incorporates exogenous state variables in the specification.

[7] PRVF begin with an explicit model of firm cost and demand where each component depends on an underlying firm-specific factor. They derive the firm's profit function $\pi(\omega_{it})$ and show that the productivity variable ω is a weighted sum of cost and demand characteristics. It is another reason why it is desirable to distinguish product and process innovations and allow each to have a different impact on firm productivity.

Other parts of the cost of R&D might not be directly measurable like the capital costs of buildings and equipment. From the modeling perspective, it will be useful to recognize that some of the costs of R&D might be unobserved by the reseacher and may need to be estimated within the model.

Second, the R&D-innovation relationship summarized by the distribution function $F(n_{it+1}|r_{it})$ is used to capture several features of the innovation process. It recognizes that there is a stochastic component to the innovation process. Conditional on the firm's choice of r_{it}, innovation outcomes are still random. We expect R&D investment to have a positive effect on the probability that the firm innovates (and this is consistent with the OECD figures cited above) however, the amount of the shift is an empirical question. The specification can also incorporate the idea that a firm can realize an innovation without any formal spending on R&D. This can reflect 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 through hiring of experienced workers, or changes in production processes that result from learning by doing without formal R&D investment. If the data are available, the innovation process can also be modeled so that n is multidimensional and there may be economies (or diseconomies) of scope in the innovation process. This could occur if a new product was developed with R&D spending but its manufacture required a change in the existing production processes. It is also possible to generalize this part of the model so that the innovation outcomes are also conditional on the firm's characteristics. In particular the firm's capital stock or size could have an impact on the probability they innovate. For example, large firms may have more opportunities to implement process innovations or have a more diverse product line so that product innovations are more likely.

Third, the stochastic evolution of productivity, represented by the cumulative distribution function $G(\omega_{it+1}|\omega_{it}, n_{it+1})$ is an important component because it captures how a firm's innovation impacts future profits and thus the economic value of the innovation. Olley and Pakes (1995) develop the idea of treating productivity as an exogenous first-order Markov process and writing it as $\omega_{it+1} = g(\omega_{it}) + \varepsilon_{it+1}$ where $g(\omega_{it})$ is the expectation of future productivity by the firm given its current productivity level and ε_{it+1} is a zero mean, *iid* shock to future productivity. They show how to incorporate this assumption on productivity evolution into a production function model and estimate both the parameters of the production function and the process of productivity evolution. More recently, Doraszelski and Jaumandreu (forthcoming) and Aw, Roberts, and Xu (2011) have generalized the assumption to make productivity growth endogenous. The productivity process can be affected by the firm's R&D investment: $\omega_{it+1} = g(\omega_{it}, r_{it}) + \varepsilon_{it+1}$.[8] In the specification used in PRVF, the firm's realized innovations are assumed to shift the distribution of future productivity: $\omega_{it+1} = g(\omega_{it}, n_{it+1}) + \varepsilon_{it+1}$. In this case the firm's expectation of period $t+1$ productivity depends on its period t productivity and its realized innovations. The firm will base its expected return to R&D on this part of the productivity evolution process, recognizing that innovation outcomes themselves are not fully predictable. This specification recognizes firm-level productivity is persistent and implies that shocks to productivity get incorporated into future productivity levels and thus their effect, along with the effect of the firm's innovations, persist into the future.

This specification explicitly requires panel data in order to estimate these intertemporal linkages in the productivity process. Because many of the CIS data sets are primarily cross-sectional, the CDM applications generally model differences in the *level* of productivity across firms as a function of innovations as in equation (3), although some papers do model productivity *growth* as a function of innovations and this is closer to the specification in PRVF and used here.

3.2 The Dynamic Choice of R&D

Given this setup it is possible to derive an empirical model of the firm's optimal choice of R&D. If the firm is forward looking it will incorporate information about the likely effects of its R&D choice on its future innovations, productivity, and profits. This will lead to an explicit formulation of the firm's dynamic demand for R&D, a feature that is absent from equation (1).

Assume that firm i chooses its sequence of R&D expenditures $\{r_{it}\}$ to maximize the dis-

[8] Doraszelski and Jaumandreu (forthcoming) show that the assumption of stochastic productivity with R&D investment can be viewed as a generalization of the knowledge capital model of Griliches (1979). This specification can also be the basis for incorporating an R&D spillover across firms. Xu (2008) models the firm's productivity to depend on its own R&D investment but also on the investment choices of other firms within the same industry.

counted sum of expected future profits net of the cost of R&D. Its value function can be written as:

$$V(\omega_{it}) = \pi(\omega_{it}) + \max_{r_{it}}[\beta EV(\omega_{it+1}|\omega_{it}, r_{it}) - c(r_{it})] \tag{4}$$

where β is the discount factor. The firm's value function V is the sum of its current period profit π and maximized discounted future expected value net of the cost for investment. The expected value term $EV(\omega_{it+1}|\omega_{it}, r_{it})$ is important in the R&D decision because it captures all future payoffs to the firm from investing in R&D. R&D investment in year t will lead to a higher probability of innovation and a higher path for productivity in year $t+1$ which in turn raises future profits. The expected future value can be written more precisely in terms of the stochastic assumptions about the innovation and productivity processes:

$$EV(\omega_{it+1}|\omega_{it}, r_{it}) = \int_{n,\omega} V(\omega_{it+1})dG(\omega_{it+1}|\omega_{it}, n_{it+1})dF(n_{it+1}|r_{it}) \tag{5}$$

The three terms on the right hand side of the equation identify the steps from R&D to innovation, innovation to future productivity, and future productivity to future long-run profits. Because this is an expectation over the future values of innovation and productivity, the specification captures the fact that future innovations and productivity are not known with certainty when the firm chooses its R&D.

In the discrete case modeled by Aw, Roberts, and Xu (2011) and PRVF, R&D is treated as a discrete choice by the firm and the R&D variable is given as $rd_{it} = 1$ if the firm chooses to invest in R&D and $rd_{it} = 0$ if it does not. In this case the marginal benefit to the firm of investing in R&D is the difference in the expected future value between the two choices. This gives us a natural way to define the long-run payoff to the firm from R&D investment:

$$\Delta EV(\omega_{it}) = EV(\omega_{it+1}|\omega_{it}, rd_{it} = 1) - EV(\omega_{it+1}|\omega_{it}, rd_{it} = 0). \tag{6}$$

The term ΔEV is simply the increment to the expected future value of the firm if they choose to invest in R&D. Since productivity is persistent, a gain in productivity does not only show up as additional profit in the next period but also in all subsequent periods and $\Delta EV(\omega_{it})$ captures this long-run payoff. Furthermore, a higher level of productivity implies a higher return on the investment and makes it more likely for firms to conduct R&D in the future which again will positively influence the firms' productivity path. This long term gain adds to the benefit of investing in R&D.

When deciding the level of R&D investment the firm must weigh this expected gain in future profits against the current cost of R&D. The final element needed to complete the empirical model of R&D demand is a specification of the cost function for R&D. For example, in Aw, Roberts, and Xu (2011) and PRVF, the authors treat firm costs as independent draws γ_{it} from an underlying cost distribution $C(\gamma)$ and estimate parameters describing this distribution as part of the model. In this discrete framework, the firm's demand for R&D is simply the probability that they choose $rd_{it} = 1$, which is the probability that $\Delta EV(\omega_{it}) \geq \gamma_{it}$.[9] This model of R&D demand puts substantial structure on equation (1) by tying the firm's choice to the subsequent likely effects of R&D on the firm's future profits. In addition to providing a summary measure of the payoff to R&D, $\Delta EV(\omega_{it})$, it also provides a basis for simulating how the firm's optimal choice of R&D will change with changes in either the costs or benefits of R&D.

3.3 Estimates for the German Manufacturing Industries

PRVF estimate this dynamic model using firm-level data from the Mannheim Innovation Panel which is collected by the ZEW and is the source of the CIS data for Germany. The key components of the data are a measure of firm profits, the firm's capital stock, the discrete pattern of firm R&D decisions, and the discrete pattern of process and product innovations realized by the firm. To illustrate how the model works we provide estimates for a simplified version of their model. We divide the manufacturing industries into two broad groups based on OECD (2011) definitions. The high-tech industry group includes the chemical, non-electrical machinery, electrical machinery, instruments and motor vehicle industries. These industries

[9] The firm's choice of R&D can also be treated as a continuous variable and the firm will choose the R&D expenditure to set the marginal benefit of an additional euro of R&D equal to the marginal cost:

$$\frac{\partial EV(\omega_{it+1}|\omega_{it}, r_{it})}{\partial r_{it}} = \frac{\partial c(r_{it})}{\partial r_{it}}.$$

In this case we would interpret the marginal benefit as the effect of an extra euro of expenditure on R&D on the firm's value.

have a ratio of R&D to sales that is greater than .025. The low-tech industry group includes food, textiles, paper, plastic, basic metals, non metallic minerals, and wood products industries which have much lower R&D to sales ratios. We estimate the model separately for the high-tech and low-tech groups because we believe that the opportunities for technological advances or new product introductions are likely to differ for the two groups and this will affect the relationship between R&D and productivity.

Before discussing the results, there are a few special features of the German CIS data and the model specification that are needed to interpret the results. We treat the R&D variable rd as a discrete choice where it equals 1 if the firm reports any expenditure on innovation inputs. This is a broader definition that also captures firms that spend money on acquiring external knowledge or on marketing and design features but do not report direct R&D expenditures. We use two discrete measures of firm innovation: product innovations will be denoted as nd and process innovations denoted as nz. Each will be set equal to one if the firm reports that type of innovation and zero if it does not.

In modeling the cost of R&D, $c(r_{it})$ in the theory model, we distinguish two types of firms. The first are firms that had invested in R&D in the previous period. In this case the cost of conducting R&D is the cost of maintaining their ongoing innovation activities and we will refer to this as a fixed per period cost. This could include all expenses that are incurred for inventing, introducing or implementing a new product or process such as operating costs related to R&D staff, materials, software and equipment. It can also include expenditures spent on intangibles bought from others such as licenses. The second type of firm did not invest in R&D in the previous period and is likely to face additional costs involved with starting up an R&D program. These can include costs such as the capital costs of a new research facility or hiring cost for R&D workers. These costs are only paid when the firm begins R&D. In each of these cases the firm's cost will be modeled as a random draw from either an underlying fixed cost distribution or a startup cost distribution, depending on their previous R&D experience. By modeling the cost of R&D in this way we recognize that firms that are starting up an R&D program may face different, presumably higher, costs than experienced firms. We also recognize that we cannot observe all the underlying costs relevant to the firm's decision.

The first feature to be estimated from the data is the effect of R&D on the probability of a future innovation, the analogue of $F(n_{it+1}|r_{it})$ in the theory model. Table 1 summarizes the probability that a firm reports adopting a new product or process innovation given that it did ($rd_{it} = 1$) or did not ($rd_{it} = 0$) invest in R&D. The first column of the table shows that, among the firms that did not report any R&D expenditure ($rd_{it} = 0$), the probability of not having a subsequent product or process innovation is .769 or .786 in the two industry groups. Perhaps more surprising is that approximately 22 percent of the firms report adopting an innovation with the most common outcome being both a product and a process innovation. This occurs with a probability of .131 and .115 in the two groups. R&D investment in the previous period is not a necessary condition for an innovation. The difference in innovation probabilities between the low tech and high tech industry groups is small when the firms do not invest in R&D. When the firms do engage in R&D ($rd_{it} = 1$) the innovation probabilities are very different. The probability of no innovation is much lower, .108 and .206 in the two industry groups, although it is not zero implying that R&D investment is not sufficient to guarantee future innovations. In addition, the probability of a product innovation is much higher than the probability of a process innovation and the most likely outcome is that the firm reports both types of innovations. Clearly, firms that invest in R&D report higher rates of innovation. Among this group, the innovation probabilities are also higher for the high-tech industry group probably reflecting more opportunities for product or process improvements. Overall the data are consistent with R&D expenditures leading to future innovations by the firm.

Table 1 The Probability of an Innovation Conditional on R&D investment

R&D Choice	$rd_{it} = 0$				$rd_{it} = 1$			
Product Innovation nd_{it+1}	No	Yes	No	Yes	No	Yes	No	Yes
Process Innovation nz_{it+1}	No	No	Yes	Yes	No	No	Yes	Yes
High-tech Industries	0.769	0.060	0.040	0.131	0.108	0.243	0.034	0.615
Low-tech Industries	0.786	0.059	0.040	0.115	0.206	0.180	0.073	0.541

The second part of the model is the productivity evolution process, the analogue of $\omega_{it+1} = g(\omega_{it}, n_{it+1}) + \varepsilon_{it+1}$ in the theory model. PRVF modeled the systematic part of the productivity process as $g(\omega_{it}, \omega_{it}^2, \omega_{it}^3, nd_{it+1}, nz_{it+1}, (nd_{it+1}nz_{it+1}))$, that is as a third-order polynomial in ω_{it} and complete set of dummies for the different combinations of process and product innovations.[10] Table 2 reports the estimates of the coefficients on the discrete innovation variables. Focusing on the firms in the high-tech industries, a firm that reports a product innovation has productivity that is 0.7 percent higher than a firm with no innovations. Firms with process innovations have productivity that is 1.0 percent higher and firms that have both product and process innovations have productivity that is 0.9 percent higher (the sum of all three coefficients). A similar pattern is seen in the low-tech industries but the contribution of product innovations is smaller, 0.1 percent addition to productivity. In both cases the interraction term $nd * nz$ is negative and approximately equal to the coefficient on nd, implying that firm's that report both types of innovations do not have any additional productivity improvement relative to firms that just report a process innovation. Process innovations are found to have a larger impact on productivity than product innovations.

Table 2 Effect of Innovations on Productivity Growth (standard errors)

	Intercept	$nd = 1$	$nz = 1$	$nd * nz = 1$
High-tech Industries	.010 (.003)	0.007(0.004)	0.010(0.009)	-0.008(0.009)
Low-tech Industries	.005 (.002)	0.001(0.003)	0.010(0.004)	-0.002(0.005)

Combining the estimates from Table 1 and Table 2, we can measure the contribution of R&D to firm revenue.[11] This elasticity is equal to .030 for the high-tech industry group, firms that conduct R&D have 3.0 percent higher revenue, holding capital fixed, than firms that do not do R&D. The estimate for the low-tech group is .026. The lower estimate for this industry reflects a smaller effect of product innovation on productivity in Table 2 but also a lower innovation success rate as reported in Table 1. This measure is conceptually similar to the elasticity of sales with respect to R&D capital that is estimated by many authors using

[10] The productivity level ω is not directly observable in the data but is estimated as part of the model jointly with the parameters of the $g(\omega_{it}, n_{it+1})$ function using the approach developed by Olley and Pakes (1996). See the full version of the PRVF paper for the details.

[11] We calculate the revenue elasticity of R&D based on the specification of PRVF. Given explicit assumptions about firm's cost and demand functions we can write the firm revenue as a function in its productivity level ω and the demand elasticity in the output market η, $\ln S_{it}(x_{it}, \omega_{it}, \eta) = (1+\eta)(\beta x_{it} - \omega_{it})$ where x_{it} is a vector of variable input prices and fixed capital stocks. Table 1 and Table 2 provide estimates to contruct the average effect of R&D investment on productivity. Combining this with the demand elastictiy estimates yields the effect of R&D on firm revenue. We estimate η to be -4.98 for high-tech and -5.59 for low-tech industries.

the CDM framework. The main difference is that these estimates correspond to a difference between firms that do R&D and those that do not, while the CDM estimates generally treat R&D as a continuous variable. Estimates of the sales elasticity reported by Hall, Mairesse, and Mohnen (2010, Table 2a) vary from .01 to .25 and are centered on .08.

One of the primary gains from the dynamic structural model is a measure of the long-run expected benefit of performing R&D, $\Delta EV(\omega_{it})$.[12] This captures the immediate gain in profit from having a higher productivity level but also the gain to a firm from being on a more favorable future productivity path. Table 3 reports estimates of this benefit. As described above this will vary with the productivity level of the firm and we report it for the 5th, 25th, 50th, 75th and 95th percentiles of the productivity distribution for firms in high-tech and low-tech sectors. A high-tech firm with productivity level of -0.174 (the 5th percentile of the productivity distribution) has an expected gain from investing in R&D of 1.869 million euros. This gain goes up to 24.251 million for a firm at the 95th percentile of the productivity distribution. The increase in ΔEV as productivity rises is also present within the low-tech group, however the difference in magnitude is pronounced. A firm in the high-tech group with a productivity level at the 95th percentile faces a gain from R&D investment of more than 24 million Euro, whereas it is only 3.5 million Euro for a comparably productive firm in the low-tech industries. This result reflects overall lower productivity and profit levels in the low-tech industries. The positive relationship between firm productivity and the benefit from performing R&D causes the high-productivity firms to self select into R&D investment at a higher rate than low-productivity firms.

[12] The model also provides estimates of the distribution of fixed costs by firms maintaining an R&D program and startup costs for firms that are beginning to invest. The startup costs are consistently higher than the fixed costs which leads to hysteresis in the pattern of R&D participation. Both costs are estimated to be higher in the high-tech industries. The mean fixed cost and startup costs for the high-tech industries are 1.55 and 7.63 million euros. For the low-tech industries they are 0.88 and 3.07 million euros.

Table 3 The Expected Benefit of R&D

Industry Group	percentile of ω distribution	ω	$\Delta EV(\omega)$ (million EUR)	$NBV(\omega)$ fixed cost	$NBV(\omega)$ startup cost	$\Delta\pi/\Delta EV$
High-tech	0.05	-0.174	1.869	-0.006	-0.060	0.011
	0.25	0.184	6.775	0.011	-0.019	0.010
	0.50	0.460	11.907	0.016	-0.001	0.015
	0.75	0.689	16.660	0.015	0.005	0.027
	0.95	1.111	24.251	0.010	0.006	0.095
Low-tech	0.05	-0.241	0.191	-0.033	-0.130	0.026
	0.25	-0.033	0.546	-0.015	-0.079	0.023
	0.50	0.223	1.446	0.002	-0.030	0.027
	0.75	0.504	3.053	0.008	-0.005	0.043
	0.95	0.811	3.538	0.005	-0.001	0.146

An alternative way to characterize the gains from R&D is to express them net of the cost of investment. In the PRVF model, the marginal benefit of R&D is determined by the firm's productivity level while the cost that the firm must spend to achieve this benefit is treated as random with a different distribution depending on whether it is a fixed or startup cost. Let $E(\gamma)$ be the mean of the distribution of costs, either fixed or startup, across the firms, then the net benefit of R&D at the mean cost level is:

$$NBV(\omega_{it}, \gamma) = \frac{\Delta EV(\omega_{it}) - E(\gamma)}{V(\omega_{it})} \tag{7}$$

where we normalized by the value of the firm. NBV is a measure of the expected net benefit of R&D investment expressed as a share of the value of the firm. This will vary across firms depending on their productivity level thus generating a distribution of returns to R&D across firms. In addition, it can be negative. If a firm has low productivity so that the benefits of R&D are small, but the costs of conducting R&D are high then the numerator of NBV can be negative and the firm will not choose to invest in R&D.

The fourth and fifth columns of Table 3 report the value of NBV for different productivity levels in the two industries. The measures also depends on whether the firm is paying a fixed cost to maintain its R&D or a higher startup cost to begin. For firms that pay a fixed cost (column 4) we see that NBV is negative for the lowest productivity firms in the high-tech industries but is positive for firms above the 25th percentile of productivity. For a median productivity firm, the net payoff is equal to 1.6 percent of firm value. This is very different for a firm that is contemplating entering into R&D investment. Since startup costs are estimated to be higher than fixed costs, only the highest productivity firms will find it profitable to begin investing in R&D. In the fifth column of the table we see that for a startup firm with the median productivity level in the industry NBV is negative 0.1 percent of firm value, so the firm would not choose to invest in R&D. For the low-tech industries, the major difference is that the negative values of NBV extend over a larger range of the productivity distribution. For startup firms, even ones at the 95th percentile of the productivity distribution the expected net benefits are negative. This will lead to much lower rates of R&D participation in the low-tech industries and for firms without previous R&D experience.

We have emphasized that an advantage of the dynamic model is that it estimates the long-run payoff to R&D as the increment to the firm's value. Because the productivity gains are long-lived this will be larger than the one-period gain in profit. Within the model we can also calculate the one-period or short-run profit gain that would accrue as the difference in future expected profits if the firm undertakes R&D versus does not:

$$\Delta\pi(\omega_{it}) = \pi(\omega_{it+1}|\omega_{it}, rd_{it} = 1) - \pi(\omega_{it+1}|\omega_{it}, rd_{it} = 0).$$

This captures the fact that the firm undertaking R&D in year t will likely have higher productivity and profits in $t + 1$ but does not capture any of 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.

The final column of Table 3 reports the ratio of the short-run gain to the long-run gain: $\Delta\pi(\omega_{it})$ / $\Delta EV(\omega_{it})$. It is apparent from the reported numbers that the short-term gain from R&D is only a small fraction of the long-term investment gain. In the high-tech group this ratio ranges from 1.0 percent to 9.5 percent depending on the firm's productivity level. Low-tech firms realize a larger share of the benefits in the short run. This is consistent with less persistence in the productivity levels for low-tech firms. Overall, the fact that the long-run expected benefit of R&D is substantially larger than the short-run benefit emphasizes the need to examine R&D choice and calculate the benefits to R&D in a dynamic framework.

4 Policy analysis

A second advantage of estimating a structural model is to provide insights into the potential effect of policy changes. Using the estimated model we can change conditions in the firm's environment and ask how the firm will respond in terms of its optimal choice of R&D and how this response will affect the level of innovation and productivity going forward. A variety of possible policy experiments can be conducted but in the case of R&D investment the most interesting ones are likely to focus on the effects of cost subsidies. These can take the form of tax credits for firm investments in R&D or changes in the expensing of R&D costs for tax purposes. They could also arise from the establishment of R&D consortium that share development costs among firms or research institutions. These cost subsidies could be designed to increase the investment levels of existing firms, such as by reducing fixed costs of investment, or to increase the number of firms investing in R&D, such as by subsidizing startup costs of new investors. Although we did not develop the role of demand side factors in this review article, a more complete model of R&D investment would incorporate demand differences and competition among firms. This could be used to assess the impact of policies that increase output market competition or alter the protection of intellectual property rights on the firm's incentive to invest in R&D.

In this case we simulate the effect of a permanent cost reduction for R&D investment. This reduction can be the consequence of a direct government support such as a grant or the consequence of an indirect government support such as a tax credit.[13] We implement this change by having the firms receive their R&D cost draws from a cost distribution with a lower mean. We then simulate the firm's optimal investment decision under the subsidized cost regimes over a time horizon of 20 years and construct the resulting productivity path for an

[13] One qualifier must be added to this exercise. We do not argue that subsidies are justified because the current R&D choices we observe in the data are inefficiently low as they would be if firms do not realize the full benefits of their investments due to positive across-firm spillovers. Instead we view this exercise as a way to assess how sensitive R&D decisions are to expected costs.

average firm and the industry R&D investment rate. We consider two different cost subsidy programs, one that aims at supporting incumbent firms by reducing the fixed cost they pay to maintain an R&D program and the second that aims to encourage entry by subsidizing startup costs. For each exercise we calculate the effect of a 20 percent cost reduction on the firm's R&D choice.

We will simulate the effects of cost reductions on the path of average firm productivity in the industry and the proportion of firms that choose to invest in R&D and do this separately for the high-tech and low-tech industries. Facing a lower R&D cost, the firm will be more likely to conduct R&D (using the discrete dynamic demand curve for R&D), this will lead to a higher probability of innovation and higher productivity. In turn this leads to a higher expected future value for the firm. The impact of the subsidy can be estimated by simulating how each firm responds under the new cost regime.

Table 4 reports the results of the simulations. The first two columns simulate how mean industry productivity and the proportion of firms choosing to invest in R&D evolves over time under the estimated parameters. In the high-tech industries mean productivity rises from .317 in the initial year to .591 after 20 years. The proportion of firms investing in R&D rises from .722 after 5 years to .815 after 20 years. The remaining columns report the effect of the reduction in fixed and startup costs. In both cases there is very little impact on the mean productivity in the industry but there is an increase in the proportion of firms conducting R&D under the fixed cost subsidy. The proportion of firms rises to .854 over 20 years in this case. There is virtually no impact on the participation rate of firms under the entry cost subsidy.

We find that subsidies to fixed cost have a greater effect on investment participation rates than subsidies for startups. This response results from how each cost change affects ΔEV. The larger is the difference between the fixed cost and the startup cost the more important is the firm's prior experience in determining its participation. A decrease in entry cost makes it more likely that new firms will begin investing but at the same time reduces the benefit of doing R&D today which will lower the participation rate of existing R&D firms. The entry and exit patterns offset each other and this results in an ambiguous effect on the overall participation rate. A decrease in fixed cost makes it more likely that existing firms will continue to invest in R&D but it also raises the value of investing in R&D to others and thus encourages entry. The entry and exit patterns reinforce each other and lead to an increase in the proportion of firms investing in R&D.

In the low-tech industries there is a more substantial response to the cost subsidies but again the response is in the participation rate, not the mean productivity, and is more substantial under the fixed cost subsidy. In that case the proportion of firms rises to .626 after 20 years as opposed to .566 under the base case without a subsidy.

What these results suggest is that in industries where the private incentives to invest are already quite strong, cost subsidies do little to encourage additional firms to undertake R&D investment. The heterogeneity in returns across firms, which is driven by productivity differences in this simplified framework, implies that the subsidies act only on those firms that have expected benefits just below the expected costs, and in the high-tech industries there are relatively few firms that are not already investing. With the low-tech industries, the expected benefits are lower and, while the fixed and startup costs are lower as well, more firms are on the margin where the expected benefits are close to the expected costs. In this case cost subsidies can help to encourage firms to continue investing in R&D even if they do little to encourage new firms to undertake R&D. Another implication is that since these policy changes act only on the marginal firms that are near the investment threshold, it may be very difficult to find any impact on the average measures of productivity or profits or revenue across all firms in the industry. We believe that these simple simulations indicate that cost reduction matters, though it only impacts a subset of firms and thus its effect might be quite limited.

Table 4 Simulations of Productivity and R&D Investment with Cost Reductions

	Base Case		20% Reduction in Fixed Cost		20% Reduction in Entry Cost	
year	ω	$Pr(rd=1)$	ω	$Pr(rd=1)$	ω	$Pr(rd=1)$
			High-tech Industries: $\omega_1 = 0.317$			
5	0.382	0.722	0.383	0.766	0.382	0.725
10	0.458	0.755	0.461	0.799	0.458	0.759
15	0.527	0.786	0.532	0.829	0.527	0.790
20	0.591	0.815	0.597	0.854	0.591	0.817
			Low-tech Industries: $\omega_1 = 0.198$			
5	0.229	0.506	0.230	0.556	0.229	0.515
10	0.264	0.528	0.266	0.582	0.264	0.537
15	0.297	0.548	0.300	0.606	0.297	0.556
20	0.327	0.566	0.331	0.626	0.327	0.574

This experiment illustrates one use of the structural model of R&D demand to quantify the effect of a permanent cost reduction on firm investment behavior and the industry producitivity development. Other possible interesting policy options include a temporary reduction in fixed or entry cost. Policy makers may also consider a subsidy program that is designed for specific group of firms. The model used by PRVF can also link the cost reduction to observable firm characteristics such as the size or age of the firm or the industry in which it operates. The investment response of the affected group of firms, large versus small for example, could then be quantified.

5 Conclusion

The study of the firm-level relationship between R&D investment and productivity has been a major area of research since the pioneering work of Griliches (1979) and early empirical study by Griliches and Mairesse (1984). In the last decade this has been spurred by the collection of firm-level innovation surveys in a number of countries. Empirical studies using these surveys have identified a set of robust facts about the importance of R&D in the innovation process and the contribution of innovations, especially product innovations, to the firm's productivity or revenue. The framework developed by Crépon, Duguet, and Mairesse has proven very useful in organizing the empirical studies that cross many countries. It is flexible, allowing the researcher to exploit unique features of each country's data in a comparable setting, and has identified some robust patterns in the relationship between R&D, innovation, and productivity.

In this article we highlight how the firm's choice of R&D investment can be empirically modeled in a dynamic setting. The model recognizes that the firm is forward-looking, so that its demand for R&D depends on how R&D investment affects future innovation, productivity, and profits. The firm is making its choices under uncertainty about the exact outcomes of the innovation process. The dynamic decision model also provides a useful measure of the expected benefit to the firm of investing in R&D: the full impact the investment has on the expected future value of the firm. This measure, which can be quantified using the firm-level innovation data, recognizes the dynamic linkages present in the innovation process. In this way the structural framework provides a methodology to measure the long-run expected payoff to R&D in a way that is consistent with the stochastic and dynamic nature of the innovation and productivity process. Estimating the model provides a basis for counterfactuals and policy simulations.

We summarize an empirical application, which is a simplified version of the model estimated in Peters, Roberts, Vuong, and Fryges (2013), to data collected in the German CIS. We find that R&D investment has substantial effects on the probability a firm realizes a product or process innovation and that these innovations do have positive effects on firm productivity. The expected benefit of R&D investment varies positively with the firm's productivity and is substantially larger in a group of high-tech industries than in a group of less R&D intensive industries. As a share of firm value, the expected benefit of R&D net of R&D costs varies from –0.6 to 1.6 percent across firms with different productivity levels in the high-tech industries and from -3.3 to 0.8 percent in the low-tech industries. Firms with negative net benefits would choose not to invest in R&D. The results provide a picture of the firms that find it profitable to invest in R&D.

While this article has focused on a dynamic model of discrete R&D choice, there are a number of ways the basic framework can be extended. With some modifications, R&D or R&D intensity can be treated as a continuous choice variable for the firm. We would then be able to quantify how an additional euro of R&D expenditure would be translated into future profits and simulate the effects of policy changes on the amount of R&D or innovation spending by the firm. A more substantial extension would develop the output demand curve faced by the firm and recognize that the firm operates in a market environment. This could then be used to address two long-standing questions in the R&D literature: Are there spillover effects of one firm's R&D investment onto the long-run profits of its competitors and what is the effect of output market competition on investment choices?

References

[1] Aghion, Phillippe and Peter Howitt (1992), "A Model of Growth Through Creative Destruction," *Econometrica*, Vol. 60, No. 2, pp. 323-353.

[2] Aguirregabiria, Victor and Pedro Mira (2007), "Sequential Estimation of Dynamic Discrete Games," *Econometrica,* Vol. 75, No. 1 (January) pp. 1-53.

[3] Aguirregabiria, Victor and Pedro Mira (2010), "Dynamic Discrete Choice Structural Models: A Survey," *Journal of Econometrics*, Vol. 156, No. 1, pp. 38-67.

[4] Aw, Bee Yan, Mark J. Roberts, and Daniel Yi Xu (2011), "R&D Investment, Exporting and Productivity Dynamics," *The American Economic Review*, Vol. 101, No. 4 (June), pp. 1312-1344.

[5] Collard-Wexler, Allan (forthcoming), "Demand Fluctuations in the Ready-Mix Concrete Industry," *Econometrica*.

[6] Crépon, Bruno, Emmanuel Duguet, and Jacques Mairesse (1998), "Research Innovation and Productivity: An Econometric Analysis at the Firm Level," in *Economics of Innovation and New Technology*, Vol. 7, No. 2, pp.115-158.

[7] Das, Sangamitra, Mark J. Roberts, and James R. Tybout (2007), "Market Entry Costs, Producer Heterogeneity, and Export Dynamics," *Econometrica,*Vol. 75, No. 3 (May), pp.837-873.

[8] Doraszelski, Ulrich and Jordi Jaumandreu (forthcoming), "R&D and Productivity: Estimating endogenous Productivity," *Review of Economic Studies*.

[9] Ericson, Richard and Ariel Pakes (1995), "Markov Perfect Industry Dynamics: A Framework for Empirical Work," *Review of Economic Studies*, Vol. 62 (January), pp. 53-82.

[10] European Union (2011), *Innovation Union Scoreboard,* Brussels.

[11] Gilleskie, Donna (1998), "A Dynamic Stochastic Model of Medical Care Use and Work Absense," *Econometrica*, Vol. 66, No. 1, pp. 1- 45.

[12] Griliches, Zvi (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.

[13] Griliches, Zvi and Jacque Mairesse (1984), “Productivity and R&D at the Firm Level,” in *R&D, Patents, and Productivity*, Zvi Griliches (ed.), Chicago: University of Chicago Press, pp. 339-374.

[14] Hall, Bronwyn H. (2011), “Innovation and Productivity,” NBER Working Paper No. 17178.

[15] Hall, Bronwyn H., Jacques Mairesse, and Pierre 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.

[16] Keane, Michael P. and Kenneth I. Wolpin (1997): "The Career Decisions of Young Men", Journal of Political Economy, Vol. 105, No. 3, pp. 473-522.

[17] Klette, Tor J. and Zvi Griliches (1996), "The Inconsistency of Common Scale Estimators When Output Prices are Unobserved and Endogenous," *Journal of Applied Econometrics,* Vol. 11, No. 4 (July-August), pp. 343-361.

[18] Klette, Tor J. and Zvi Griliches (2000), "Empirical Patterns of Firm Growth and R&D Investment: A Quality Ladder Model Interpretation," *The Economic Journal*, Vol. 110 (April), pp. 363-387.

[19] Klette, Tor J. and Samuel Kortum (2004), "Innovating Firms and Aggregate Innovation," *Journal of Political Economy*, Vol 112, No. 5, pp. 986-1018.

[20] Mairesse, Jacques and Pierre 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.

[21] Mairesse, Jacques, Pierre Mohnen, and Elisabeth Kremp (2005), "The Importance of R&D and Innovation for Productivity: A Reexamination in Light of the French Innovation Survey," *Annales D’Economie Et De Statistique,* No. 79/80, pp. 489-527.

[22] Nevo, Aviv and Michael Whinston (2010), "Taking the Dogma out of Econometrics: Structural Modeling and Credible Inference," *Journal of Economic Perspectives*, Vol. 24, No. 2, pp. 69-82.

[23] OECD (1992, 1996, 2005), *Oslo Manual: Proposed Guidelines for Collecting and Interpreting Technological Innovation Data, 1st, 2nd and 3rd edn.*, Paris.

[24] OECD (2010), *Measuring Innovation: A New Perspective,* Paris.

[25] OECD (2011), ISIC Rev. 3 Technology Industry Definition, Paris.

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

[27] Peters, Bettina (2009), "Persistence of Innovation: Stylized Facts and Panel Data Evidence," *Journal of Technology Transfer,* Vol. 34, pp. 226-243.

[28] Peters, Bettina, Mark J. Roberts, Van Anh Vuong, and Helmut Fryges (2013), "Firm R&D, Innovation, and Productivity in German Industries," Working Paper.

[29] Reiss, Peter and Frank Wolak (2007), "Structural Econometric Modeling: rationales and Examples from Industrial Organization," in *Handbook of Econometrics, Vol. 6A*, James J. Heckman and Edward Leamer (eds.), North-Holland.

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

[31] Rust, John and Christopher Phelan (1997), "How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets," *Econometrica*, Vol. 65, No. 4, pp. 781-832.

[32] Ryan, Stephen (2012), "The Costs of Environmental Regulation in a Concentrated Industry," *Econometrica*, Vol. 80, No.3 (May), pp. 1019–1062.

[33] Wolpin, Kenneth I. (1996), "Public-Policy Uses of Discrete Dynamic Programming Models," *American Economic Review Papers and Proceedings*, pp. 427 -432.

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