

NBER WORKING PAPER SERIES

THE EFFECT OF R&D ON QUALITY, PRODUCTIVITY, AND WELFARE

Mons Chan
Amil Petrin
Frederic Warzynski

Working Paper 30950
<http://www.nber.org/papers/w30950>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2023

We thank Mark Roberts and seminar participants at Toronto, Queen's, CAED 2019 and Yale. This Project has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement no. 822390. All errors are ours. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2023 by Mons Chan, Amil Petrin, and Frederic Warzynski. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Effect of R&D on Quality, Productivity, and Welfare
Mons Chan, Amil Petrin, and Frederic Warzynski
NBER Working Paper No. 30950
February 2023
JEL No. L1

ABSTRACT

In this paper we provide a methodology that jointly studies production and demand for multi-product firms using detailed firm-product level data from Denmark. We estimate marginal cost by combining production function estimation with a cost function that allows for quasi-fixed inputs. We use a discrete choice demand model that extends insights from Berry, Levinsohn and Pakes (1995) to obtain a measure of the demand shock (quality). We estimate the relationship between product (process) R&D and quality (efficiency), and find strong evidence that process innovation is related to higher efficiency, while product innovation is associated with higher product quality. We discuss the welfare implications of these two distinct innovation activities.

Mons Chan
Queen's University
Department of Economics
94 University Ave
Kingston K7L 3N6
Canada
mons.chan@queensu.ca

Frederic Warzynski
Aarhus University
fwa@asb.dk

Amil Petrin
Department of Economics
University of Minnesota
4-101 Hanson Hall
Minneapolis, MN 55455
and NBER
petrin@umn.edu

1 Introduction

Innovations in product quality and production methods are widely agreed to be an important source of long run growth in output per capita. Our understanding of how investment in research and development leads to these innovations has increased in recent years but the lack of detailed micro-level data has left many questions unanswered. A leading question is whether there is too much or too little R&D from the standpoint of maximizing per capita income.¹ Theoretically it can go either way and this has led to a debate about whether or how much to subsidize research and development.

Almost all research on this question uses plant-level revenue as the measure of output but aggregating prices and quantities over products at the plant confounds demand and supply, raising questions about the conclusions that can be drawn from such data. We show how to exploit new detailed firm-product level panel data to investigate how R&D investments affect product quality, production efficiency, and producer and consumer surplus. We estimate the spillover effects of R&D and we develop a methodology for answering whether there is too much or too little R&D from society's perspective and from the private sector's perspective.

The data include 10 digit product-specific revenues and quantities (in units) for all products sold in Denmark including all imports and all products produced domestically by manufacturing firms with 10 or more employees. The data also include measurements on the inputs used in production and R&D, including questions about total expenditure on R&D, percent of expenditures allocated to each of product and process innovations, whether or not any innovations have occurred recently, the number of patents granted, and other variables of potential interest in understanding R&D at the firm level.²

¹For example, see the discussion related to the Lisbon Agenda which was designed to encourage R&D in Europe to prepare for the reduction of European Union tariffs.

²This includes data from several European countries including Denmark, Belgium, France, Norway, Sweden, Spain and Slovenia (at least), and from Chile, Colombia, Japan, Korea, and to limited extent in the United States. In Europe the Lisbon Agenda led to a standardized methodology for collecting information about R&D levels, inputs and R&D subsidies for all countries in the Union (see the Oslo Manual, OECD (2005)). In the United States there are detailed annual surveys on R&D collected for many manufacturing firms that can be linked to the Census of Manufacturers data.

On the demand side we estimate standard logit and nested logit demand functions building on Berry (1994) and Berry et al. (1995) and develop a novel instrumental variable approach to recover estimates of quality conditional on price. To construct our instruments, we exploit a rare feature of our data. First, we use the fact that we can precisely identify inputs used by firms for the products that they make. Second, we also observe every import transaction of these inputs in Denmark, so that we can compute an average import price for every input used by Danish firms.

On the supply side we allow for multi-product production functions and multi-product cost functions by directly using quantities or total costs as the dependent variable respectively and inputs and outputs as explanatory variables (see Mundlak and Razin (1971) and Diewert (1973)). We allow for a productivity shock in the production function when thinking about estimation of the cost function. Because the productivity shock enters the cost function in a non-separable manner, the usual proxy techniques of Olley and Pakes (1996), Levinsohn and Petrin (2003), Wooldridge (2009), and Akerberg et al. (2015) are not consistent when applied to the cost function. Our insight is that we can use these techniques in estimating the production function and then recover the productivity shock, which is the missing variable that needs to be conditioned upon when estimating the cost function.

Parameter estimates of a cost function estimated only as a function of input prices and output quantities are not well-defined if any input is quasi-fixed. We follow Berndt and Morrison (1981) to allow for adjustment costs in quasi-fixed inputs (like capital). In this case the static variable cost function has as arguments variable input prices, output quantity, and the levels and changes of the quasi-fixed inputs. These are all observed in our data.

The estimated cost function at the product-level allows us to estimate marginal costs which, along with the unit price data, directly imply the markups.^{3,4} We then link the demand and supply estimates to the R&D production data to estimate the impact of R&D

³This generalization allows us to investigate whether the usual method of making sufficient assumptions to identify both demand and supply from supply are reasonable.

⁴See Dhyne et al. (2014) for related work using Belgian data.

on technical efficiency and product quality.

Our data come from a time when there is a stark increase in competition from China and other Eastern European countries as European Union tariffs and quotas start to decrease in 2003; from 1990 to 2009 the share of imported goods in Danish GDP rose from 28 to 40 percent. Thus our results also provide a look inside the relationship between productivity, innovation, competition, and welfare (see e.g. Arrow (1962), Syverson (2004), Aghion et al. (2005), Foster et al. (2008), and the discussion in Holmes and Schmitz Jr (2010)).

Our preliminary results are as follows. Price is positively correlated with quality, negatively correlated with technical efficiency, and positively correlated to the productivity residual based on the revenue production function (TFPR), marginal cost and markups. Technical efficiency (ω) and TFPR are also strongly correlated but less so with TFPC, the cost function technical efficiency term.⁵ Markups are positively correlated with technical efficiency and negatively correlated with marginal cost. Higher product quality is strongly negatively correlated with ω suggesting a tradeoff between product and process innovations. Changes in our estimated joint distribution of quality and technical efficiency are consistent with self-reported times of product and process innovation at the firm. Overall R&D expenditures predict increases in quality and technical efficiency. Expenditures on product innovations do a much better job of predicting quality improvements than technical efficiency improvements and similarly expenditures on process innovations predict increases in technical efficiency much better than quality improvements.

This paper builds on a long literature which seeks to estimate the returns to innovation (see Hall et al. (2010) for a survey). We differ from much of this literature in that we do not construct a stock of knowledge capital, but rather follow Doraszelski and Jaumandreu (2013) in allowing R&D investments to interact flexibly with the evolution of firm productivity. A more recent literature looks at the relationship between firm R&D investment, productivity, and exporting (see e.g. Peters et al. (2017) and Maican et al. (2020)). Many of these papers

⁵The negative correlation means they are "positively" correlated because one is production and one is cost.

estimate the effect of R&D on firm-level productivity and demand as an intermediate step. We build on these papers by showing how to do this at the product-level. Unlike most of these papers, we treat R&D spending as a continuous variable, rather than a discrete indicator, allowing us to estimate the marginal effect of policies such as R&D subsidies. Our paper also speaks to the literature on estimating markups (see De Loecker and Warzynski (2012), De Loecker et al. (2020), Raval (2020), and Flynn et al. (2019)) in that we are able to estimate and compare markups from both the production side using product-level cost functions, and the demand side using a product-level demand system.

Two papers in particular are closely related to ours. Van Den Bosch et al. (2019) also decompose the effect of R&D on prices and product-level production efficiency. Relative to their approach, our product-level production function makes fewer assumptions on input allocation, and we obtain measures of markups from both the supply (cost function) and demand side. Estimating a product-level demand model also allows us to examine the effect of R&D on product quality and perform welfare counterfactual exercises. The other is Jaumandreu and Mairesse (2010) who estimate the effect of process and product innovations on productivity and demand. However their unit of observation is at the firm level and their analysis restricts markups to be constant. In contrast, we allow markups and demand to vary across products, firms and time.

We start by discussing the various datasets that we use in section 2. We then turn to demand (section 3) and supply (section 4) estimation given this data. In Section 5 we discuss estimation of R&D production. Section 6 describes the econometric approach for the various pieces of the estimation. Section 7 shows and discusses results while Section 8 concludes.

2 Data

We combine three main datasets for our analysis all covering the period 1999-2016. The first one is a production survey providing quarterly information about value and quantity of

each product made and sold by the firm by quarter. The second is an annual R&D survey detailing the innovation activities of firms, focusing on the amount of R&D spending and the share going to product and process innovation. The third one is a standard accounting dataset that contains standard information about firm revenue and input use (material, labor and capital).

2.1 Danish PRODCOM

The production survey (VARs) is the Danish version of PRODCOM, used in various countries (see e.g. Amity et al. (2019), Bernard et al. (Forthcoming), Dhyne et al. (2017) for Belgium, or Smeets and Warzynski (2019) for France). Firms are asked to describe the goods that are making (as a CN10 code), the quantity sold and the value. The VARs survey contains quarterly revenue (dkk) and sales (quantity) for goods at the 10-digit product level for a large set of manufacturing firms from 1995 through 2018. We aggregate this to the 8-digit product and annual level. Having revenue and quantity allows us to construct output prices (p_{gjt}) for each product (g), firm (j), year (t) observation as well as output market shares (s_{gjt}). This data forms the basis for the demand, productivity and cost function estimation. Table 1 shows the average number of products sold by firms and the number of firms interviewed for every quarter.

2.2 Danish CIS

The R&D survey is the Danish version of the Community Innovation Survey (CIS) which has been widely used in numerous papers covering different European countries (see e.g. Peters et al. (2018) for a recent example using German data). The survey has been run since 1995, but we only use the versions starting in 1999. The questions asked in the survey vary every year to better account for the complexity of firms' R&D activities, but the question about the amount of money spent on R&D is asked every year, as well as the number of R&D workers. Table 2 shows the number of firms reporting positive investment in R&D by year

and the weighted average of R&D investment. As we can see, the share of firms investing in R&D increases sharply, and also the amount of R&D investment.

In addition, every two years, firms are asked the share of R&D investment used for product innovation, process innovation and general innovation. As documented for other countries, the bulk of the investment is made for product innovation (see weighted shares in table 3), ranging from 68% in 2009 and 81% in 2005.

2.3 Accounting data

In order to estimate our production function, we need information about input use and we obtain this from the Accounting Statistics (Regnskab) provided by Statistics Denmark. This provides information about intermediate inputs use, number of employees, wage bill, and various components of fixed assets and investments that we use to build a standard perpetual inventory measure of capital. The information is available since 1995 until 2017, but we use it only since 1999. It covers more than 150,000 companies yearly and the quality of the match with the firms in the the production survey is very high.

2.4 Customs data

A big part of our strategy relies on using input prices to construct instruments for our demand and cost function estimation procedures. We rely on firm-product level customs data to construct these input price instruments. In particular, the data contains annual observations of import quantity and value at the firm and 8-digit product level. We construct price indices using firm-product import shares as weights (see below for details).

In addition, Tariff indices are obtained by combining WITS tariff data for Denmark⁶ from 1997 through 2013 with comtrade data on global trade flows in 1999. From the tariff data we obtain an indicator for every year-product-country observation which equals 1 if there is

⁶Tariffs are actually set at the EU level.

an effective (AHS) tariff on that 6-digit product from that country and 0 otherwise.⁷ From the trade data we obtain for each year-product-country observation the share of total global trade (exports) value in the year 1999.

2.5 Input survey

We additionally take advantage of an input use survey (VARK) that contains input use data for a subset of manufacturing firms from 2000 to 2018. The observations are expenditure at the firm/year/input level, where input is a 6-digit product code.⁸ In this paper, we only use this dataset as a key element to construct an instrument variable for prices, as is discussed later in the paper (see Chan, 2017 and Bernard et al., 2021 for other papers using this dataset).

We combine this data on input use with data on input prices and tariffs to construct instruments for output prices.

3 Demand

A critical aspect of our approach is that we want to be able to "back out" (or invert) from demand an estimate of product quality. This requires that there be a unique measure of product quality for every product given prices and market shares. If goods are weak substitutes in a random coefficients discrete choice demand setting then this uniqueness condition is satisfied for standard utility parameterizations (Berry et al., 1995).

We let utility depend on individual characteristics ν_i , price (p_j), observed product characteristics (x_j), and unobserved characteristics (ξ_j):

$$U(\nu_i, p_j, x_j, \xi_j; \vartheta)$$

⁷We also have the level of each tariff, but do not use levels in the construction of our input tariff index.

⁸The VARK data also includes data on other inputs such as services and water/electricity. We do not use this data in our analysis.

where ϑ is a vector of parameters to be estimated. A consumer will buy the good that provides the highest utility. Define A_j as the set of values of ν that induces the choice of j and let $P_0(d\nu)$ be the density of ν in the population. The market share of good j as a function of the characteristics of all goods competing in the market:

$$s_j(p, x, \xi; \vartheta) = \int_{\nu \in A_j} P_0(d\nu)$$

We start by considering a logit model. We write utility as:

$$u_{ij} = \beta x_j - \alpha p_j + \xi_j + \varepsilon_{ij} = \delta_j + \varepsilon_{ij}$$

with β the vector of taste parameters associated with characteristics x , α the demand elasticity parameter, and ε_{ij} assumed to be distributed type 1 extreme value. We collect all of the terms in utility that are constant for good j in δ_j . Berry et al. (1995) show for random coefficients models that there exists a unique $\delta(\theta) = (\delta(\theta)_1, \dots, \delta(\theta)_J)$ that matches observed to predicted markets shares. In the standard logit case the inversion is given by

$$\ln s_j - \ln s_0 = \delta_j = \beta x_j - \alpha p_j + \xi_j.$$

We will allow for price to be correlated with unobserved quality ξ , as in Berry et al. (1995).

A drawback of the logit IIA assumption is that the derivative of the inside share with respect to the number of goods is strictly positive. The nested logit specification, which follows Berry (1994), loosens the IIA assumption by allowing for the possibility that goods are closer substitutes for one another (relative to the simple logit). The utility function is now:

$$u_{ij} = \delta_j + \zeta_i + (1 - \sigma)\varepsilon_{ij}$$

where, following Berry (1994), for consumer i , ζ_i is common to all movies and has a distribution function that depends on σ such that if ε_{ij} is a random variable, then $\zeta_i + (1 - \sigma)\varepsilon_{ij}$ is

also extreme value. In this model the derivative of the inside share with respect to the number of goods approaches zero as σ approaches 1, so additional products simply cannibalize one another's share with no market expansion effect.

We implement this by including the product's share among inside goods - $\ln(\frac{s_j}{1-s_0})$ - as an explanatory variable. Suppressing time subscripts, we estimate:

$$\ln(s_j) - \ln(s_0) = \beta_0 + \alpha p_j + \sigma \ln(\frac{s_j}{1-s_0}) + \xi_j$$

The coefficient on the variable is the parameter σ , which indicates the degree of substitutability. When $\sigma = 0$, the model resolves to the simple logit; when $\sigma = 1$, inside products are perfect substitutes for another. $\frac{s_j}{1-s_0}$ is evidently endogenous and we instrument it in addition to prices.

$$\ln(s_j) - \ln(s_0) = \beta_0 + \alpha p_j + \sigma \ln(\frac{s_j}{1-s_0}) + \xi_j + \lambda p_j * \xi_j.$$

4 Supply

In this section we develop an approach to estimating the variable cost function with multiple outputs and adjustment costs for quasi-fixed inputs. We allow for the production function shock to affect the cost function. We also allow this productivity shock to be correlated with both output quantities and quasi-fixed inputs, which are arguments of the cost function. Allowing for correlation between quantities and productivity is important because more productive firms will in equilibrium - holding quasi-fixed inputs constant - produce more output. Allowing for correlation between quasi-fixed inputs and productivity is also important because as more productive firms are the firms that are likely to survive and accumulate quasi-fixed inputs over time.

4.1 Production

Estimation of the productivity shock is important for controlling for simultaneity in the cost function and for recovering an estimate of technical efficiency the changes of which can be related to expenditures on process innovations. In the case of a firm that produces a single-product the exercise of recovering the firm-level productivity shock is straightforward. With production given as

$$q_{ijt} = \beta_0 + \beta_l l_{it} + \beta_k k_{it} + \beta_m m_{it} + \omega_{it} + \eta_{it}$$

where in logs labor is l_{it} , k_{it} is capital, m_{it} is materials, the productivity term ω_{it} is assumed to be first-order Markov and may be correlated input choices, and η_{it} is an i.i.d. shock to production. $\beta = (\beta_0, \beta_l, \beta_k, \beta_m)$ are the elasticities of output of good j with respect to the inputs holding the other output quantities at the firm fixed. We experiment with the different proxy approaches, using the Wooldridge (2009) versions of Levinsohn and Petrin (2003) and Olley and Pakes (1996) estimators to allow for correlation between the technical efficiency error and inputs. Once the production function is estimated it is straightforward to estimate ω_{it} and add it as a regressor in a (potentially non-separable) cost function specification.

Many of the firms in our data are multi-product firms. Data on inputs used in production are aggregated across all the products at a firm. There are two approaches in the literature when inputs are aggregated, the first being a special case of the second. The first treats the multi-product firm as a collection of single product firms by dividing inputs among the products by (e.g.) the revenue share of each good. The second approach is to treat the production function as multi-product, and to estimate the multi-product transformation function described below (see Mundlak and Razin (1971) or Diewert (1973)). In the case of single product firms both approaches reduce to the standard production function setup.

Diewert (1973) shows that under mild regularity conditions there will exist a transformation function that relates the output of any good j to all other goods the firm produces and to aggregate input use. We add to that setup a productivity term that we call ω_{it} which we

assume follows a first-order Markov process and which may be correlated with both inputs and outputs. We write the production function for firm i producing good j as

$$q_{ijt} = \beta_0 + \beta_l l_{it} + \beta_k k_{it} + \beta_m m_{it} + \gamma' q_{it,-j} + \omega_{it} + \eta_{it} \quad (1)$$

where

$$q_{it,-j} = (q_{it1}, \dots, q_{it,j-1}, q_{it,j+1}, q_{itJ})$$

$q_{it,-j}$ is the vector of quantities produced of other goods. Holding overall input use constant γ_k is the additional amount of output j that would result from reducing output k by one unit holding input use constant. We modify the Wooldridge (2009) versions of Levinsohn and Petrin (2003) and Olley and Pakes (1996) estimators to allow for correlation between the technical efficiency error and inputs and correlation between technical efficiency and $q_{it,-j}$.

In practice, firms produce different sets of goods, so dealing with a vector is impractical. Indeed, it is difficult to identify precise environments where enough firms produce the same goods (see the discussion in Dhyne et al., 2022, 2023). To deal with this problem, we use an output aggregator for the production function of the other goods q_{-j} instead of treating the other goods as a vector.⁹

Our main aggregator is defined as follows:

$$q_{-j}^* = \log\left(\frac{\sum_{g \neq j} P_{m_g} Q_{m_g}}{P_{-j}}\right).$$

where P_{-j} is a firm-level price deflator constructed by using the observed prices of all the other goods produced by the firm, as in Eslava et al. (2004) and Smeets and Warzynski (2013).¹⁰

⁹Roberts and Supina (2000) use a similar quantity index when they estimate cost functions.

¹⁰If P_{igt} and $P_{ig(t-1)}$ are the unit values of good g produced by firm i at times t and $t - 1$, respectively, then the weighted average of the growth in output price is defined as

$$\Delta P_{i(-j)t} = \sum_{g \neq j} \bar{s}_{igt} \Delta \ln(P_{igt}),$$

4.2 Estimating Cost Function and Measuring Marginal Costs

We extend the settings of Lau (1976) and Berndt and Morrison (1981) who develop a variable cost function that accommodates both freely variable inputs ((m)aterials, (l)abor) and quasi-fixed inputs ((k)apital), where quasi-fixity is defined as having an adjustment cost function $c(\Delta k)$ such that

$$c(0) = 0, \quad c(\Delta k) > 0 \text{ if } \Delta k \neq 0$$

with Δ the first-difference operator. Under weak conditions they show that the variable cost (VC) function has as arguments the variable input prices ($P_I = (P_m, P_l)$), the quantity of output (q), and the levels and changes of the quasi-fixed inputs:

$$\ln VC_{it} = f(P_{It}, k_{it}, \Delta k_{it}, q_{it}).$$

Including the levels and changes is sufficient for (locally) controlling for the discontinuity introduced by the adjustment costs into the adjustment-cost-free cost function (i.e. the cost function written only with arguments input prices and output quantities). This new VC function gives the minimum variable costs necessary to achieve output q_{it} when facing prices P_{It} , given the level and change of the quasi-fixed input (s).

We let J_i be an index of the number of different products produced by i and we define the productivity shock from i 's production function as ω_{it} (see Section 4.1 for more detail). Our generalized variable cost function that accounts for multi-product production simultaneity induced by firm-specific productivity is written as:

$$\ln VC_{it} = f(P_{It}, k_{it}, \Delta k_{it}, q_{i1t}, \dots, q_{iJ_it}, \omega_{it}).$$

Allowing for multi-product production is straightforward in theory as we just extend the

where $\Delta \ln(P_{igt}) = \ln P_{igt} - \ln P_{ig(t-1)}$ and $\bar{s}_{igt} = \frac{s_{igt} + s_{ig(t-1)}}{2}$ (s_{igt} and $s_{ig(t-1)}$ are the revenue shares of product g at time t and $t-1$, respectively). The price index for each firm is then $\ln P_{i(-j)t} = \ln P_{i(-j)t-1} + \Delta P_{i(-j)t}$, where the price for the reference year is standardized ($P_{i'0} = 100$).

single-product variable cost function – which conditions on one output – to a setup that conditions on multiple outputs. The minimization problem then has firms solving for the minimum variable costs to produce the vector of outputs $(q_{i1t}, \dots, q_{iJ_t})$ conditional on quasi-fixed inputs and their adjustments.

One issue with this estimation approach is that - in principle - each permutation of production product-tuples should be considered a separate production technology. This works to reduce the number of observations of any production product-tuple. A second issue for estimation is that there is a curse of dimensionality if each quantity is allowed to interact flexibly with all other arguments in the cost function. We revisit both of these issues in the Estimation section.

The problem of simultaneity is more difficult when working with the cost function because the productivity shock will not enter the cost function in a separable way. This invalidates the use of the proxy methods (e.g. Olley and Pakes (1996), Levinsohn and Petrin (2003), Wooldridge (2009), Akerberg et al. (2015)) because they require that the estimated (production) function is separable in the productivity shock. In order to control for the simultaneity problem that is induced in the cost function by the simultaneity problem in the production setting we estimate the production function and recover the productivity shock and condition directly on it. In an environment with non-separable errors one must condition on both the realized values of all observed variables and the error in order to achieve consistency (i.e. instrumental variable methods are inconsistent for endogenous variables in non-linear settings).¹¹

5 Research & Development

We posit that R&D can impact product quality, and that product quality follows a process that is a function of past product quality, technical efficiency, and past R&D investments.

¹¹See e.g. Blundell and Powell (2003)

A simple starting point is

$$\xi_{it} = h(\xi_{i,t-1}, r_{\xi i, t-1})$$

and

$$\omega_{it} = h(\omega_{i,t-1}, r_{\omega i, t-1}),$$

It is straightforward to accommodate spillovers if we define for every product $R_{\xi j, t-1}^i = \sum_j r_{\xi j, t-1}$, where the sum is taken of the set of products whose R&D may potentially spill over. We then add this new variable to the conditioning set to get

$$\xi_{it} = h(\xi_{i,t-1}, r_{\xi i, t-1}, R_{\xi j, t-1}^i)$$

and similarly

$$\omega_{it} = h(\omega_{i,t-1}, r_{\omega i, t-1}, R_{\omega j, t-1}^i).$$

We consider several specifications of these equations in the empirical section. We also run these specifications with instruments for R&D that use the previous year's estimated profits - as given by our supply side estimates - as an instrument.

6 Estimation

6.1 Demand

For now, we estimate a very basic logit demand model. Markets (and thus market shares) are defined by 6-digit product code and year. We further subdivide markets by quantity type. So, those goods which have a 6-digit code of "840505" and are measured in "kg" are considered separate from goods measured in "units". Goods themselves are aggregated at the 8-digit level. The regression is thus:

$$\ln(s_{gj}) - \ln(s_O) = \delta_{gj} = \alpha p_{gj} + \beta_j + \beta_g + \xi_{gj}$$

where s_{gj} is the physical quantity share of 8-digit good g at the 6-digit level, s_O is the outside share, p_{gj} is the unit price of good g produced by firm j , ξ_{gj} is unobserved demand/quality, and β_j, β_g are firm and product specific intercepts. We suppress the notation for year, t , for ease of reading.

Instruments

A key issue in estimating this specification is dealing with possible correlation between prices and unobserved quality. In particular, there may be both cross-sectional correlation (goods with higher prices may have higher unobserved quality) and also correlation over time (changes in quality may be correlated with changes in price). We also do not observe product characteristics for any goods other than firm, product code, and how the good is measured (for example, by kg or by liter). We pursue two strategies to deal with these issues.

The panel structure of the data allows us to include both firm and product-specific intercepts. This controls for much of the cross-sectional correlation between price and quality, leaving only components of ξ_{gj} which vary by market (time) to be potentially correlated with prices.

To control for remaining correlation, we construct and use several different instruments. The first is a weighted input price index (PI) which combines price information on the set of physical inputs (h) used by each firm in the production of their output (g). The assumption here is that changes in input market prices are not correlated with changes in the unobserved quality of output of any particular firm. The second is a weighted input tariff index (TI) where again we assume that changes in input tariffs are uncorrelated with changes in unobserved output product quality. The third is a most-important input price instrument (SP). While the first two instruments are firm-year level variables, here we instrument for the output price of product g using only the market price of the input h which is most correlated with the price of g . Thus this instrument is at the product-year level.

Input Sets and Input Weights

All of the instruments we employ are based on the set of intermediate goods purchased by firms to use in the production of their final goods. The main problem with doing so is that firms may respond to changes in input prices by adjusting their input set, which may affect unobserved product quality. For example, suppose a firm makes water bottles and we define z_{jt} as the mean unit price of all inputs purchased in period t by firm j . If between $t-1$ and t , the firm switches from using plastic to using steel to make its water bottles, the price index instrument may increase, as steel may be more expensive than plastic. This will certainly be reflected in the output price. However, the unobserved quality and demand for the water bottle will also likely increase, violating the conditions for a valid IV.

To avoid this problem, we construct input-set weights which are constant across all markets/years, and thus invariant to changes in input set. This should strip the price index instruments of correlation with changes in unobserved quality. The input set weights are constructed as follows.

Define firm j 's input set for period t as $\mathcal{H}_{jt} \equiv \{h | q_{hjt} > 0\}$, where q_{hjt} is the quantity of intermediate h purchased by j in period t . Firm j 's total input set is then $\mathcal{H}_j \equiv \bigcup_t \mathcal{H}_{jt}$. We further define the per-unit *input requirement set* for product g as $\overline{\mathcal{H}}_{gj} \equiv \{\bar{q}_{hgj} | h \in \mathcal{H}_j\}$ where $\bar{q}_{hgj} \equiv \frac{\sum_t q_{hjt}}{\sum_t q_{gjt}}$ is the total amount of h purchased per unit of g produced over the sample period.

Input Price Indices

The first type of instrument we construct are weighted indices of the set of input prices faced by the firm. The motivation for this instrument is that changes in input prices lead to changes in marginal costs of production, and thus changes in output prices. Let p_{ht} be a measure of the market price of input h . Then the price index instrument z_{gjt}^1 is constructed as $z_{gjt}^1 \equiv \sum_{h \in \mathcal{H}_j} \bar{q}_{hgj} p_{ht}$, where the time-invariant weights $\bar{q}_{hgj}$ are as described above. Given the time invariance of the input set weights, we can be reasonably sure that changes in input

prices are orthogonal to changes in unobserved quality.

We obtain measures for p_{ht} either by looking at the weighted average price paid for h by importers, or the weighted average price charged by domestic producers. In both cases, we construct $p_{ht} \equiv \sum_{jt} w_{hjt} p_{hjt}$, where w_{hjt} is a firm-input-year weight. The import data records quantity imported of h by firm and year in both units and kg at the 8-digit level, so we construct two import price indices, one using kg share weights and the other using unit share weights. Import prices are the implied price (per unit import expenditure) by the importing firm j in that year. In the output data, quantity is recorded in various measures, but these measures are consistent within 8-digit product codes, so we use quantity shares for weights. Prices here are the output price (per-unit sales revenue) for good $g = h$. This gives us three potential product-level instruments for g - import prices weighted by kg share, import prices weighted by unit share, and output prices weighted by quantity share, each covering a different set of 8-digit products. Since the kg-weighted import prices have the most coverage, they are our preferred price index instrument.

Input Tariff Indices

The second type of instrument we construct are weighted indices of the set of input tariffs faced by the firm. The motivation here is similar to the price index instrument, in that holding the input set fixed, changes in input tariffs will affect output prices but not unobserved quality. Let τ_{ht} be a measure of import tariffs on input h . Then the tariff index instrument z_{gjt}^2 is constructed as $z_{gjt}^2 \equiv \sum_{h \in \mathcal{H}_j} \bar{q}_{hgj} \tau_{ht}$.

The tariff measure for input h is obtained using WITS data on EU import tariffs and Comtrade data on global trade flows. In general the tariff measure for h is $\tau_{ht} \equiv \sum_c \lambda_{hc} \tau_{hct}$ where λ_{hc} is a weight equal to the share of country c in the global trade of h in 1999. We obtain these weights from the Comtrade data. The country specific tariff measure τ_{hct} is either the effectively applied (AHS) import tariff for good h from country c in year t , or an indicator which equals 1 if the AHS tariff is positive, and 0 otherwise. The weighted tariff

indicator is our preferred instrument.

Most-Correlated Input Price

The indices above have a significant amount of noise which potentially limit their power as instruments. Our third instrument attempts to get around this problem by selecting, for each output good price p_{gjt} , the input price p_{ht} which is most correlated over time. The instrument is $z_{gjt}^3 \equiv \bar{q}_{hgj} p_{ht}^g$ where

$$p_{ht}^g \equiv (p_{ht} | r(p_{ht}, p_{gjt}) > r(p_{kt}, p_{gjt}) \forall k \neq h, h \in \mathcal{H}_j, k \in \mathcal{H}_j)$$

where $r(x, y)$ refers to the Pearson correlation coefficient. Since inputs are recorded at the 6-digit level in VARK, the actual procedure is to select the 8-digit input price which is most correlated with the output price from within all of the 6-digit input categories used by the firm.

This method is similar in spirit to the instrument selection method proposed by Belloni et al. (2012) and Chernozhukov et al. (2015), with the main difference that our procedure is run at the *firm-product* level rather than at the *variable* level.

6.2 Production function

We adopt the standard algorithm suggested by Wooldridge (2009) to deal with the Akerberg et al. (2015) criticism of Olley and Pakes (1996) and Levinsohn and Petrin (2003). We use our data to estimate physical productivity at the firm-product level (TFPQ) following Dhyne et al. (2017).¹² As shown in Diewert (1973), the production function of any good j produced by the firm can be written as a function of inputs and the production of the other goods

¹²As discussed previously, this estimation is not really central to our analysis and we mostly need to recover a measure of TFP to use it and deal with the endogeneity in our cost function estimation, so we simplify the discussion and refer the reader to DPSW for more details on the method.

$(-j)$ produced by the firm:

$$q_{jt} = \beta_0 + \beta_l^b l_t + \beta_k^b k_t + \beta_m^b m_t + \gamma_{(-j)} q_{(-j)t} + \varepsilon_{jt} \quad (2)$$

with the production parameters $\beta^b = (\beta_l^b, \beta_k^b, \beta_m^b)$ denoting the percentage change in output of product j due to a one percent change in any one input holding other inputs and production of the other outputs constant. γ_{-j} is the percent change in the output of good j that results from increasing the output of the other goods by one percent holding overall input use constant. The shock can be decomposed in two components, i.e. $\varepsilon_{jt} = \omega_{jt} + \eta_{jt}$. η_{jt} is assumed to be i.i.d. error upon which the firm does not act (like measurement error or specification error). ω_{jt} is the technical efficiency shock, a state variable observed by the firm but unobserved to the econometrician. ω_{jt} is assumed to be first-order Markov and is the source of the simultaneity problem as firms potentially observe their shock before choosing their freely variable inputs l_t and m_t . k_t also responds to ω_{jt} but with a lag as investments made in period $t - 1$ come online in period t . This assumption allows k_t to be correlated with expected value of ω_{jt} given $\omega_{j(t-1)}$. but the innovation in the productivity shock $\xi_{jt} = \omega_{jt} - E[\omega_{jt}|\omega_{j(t-1)}]$ is unknown at the time the investment decision was made in $t - 1$ and is therefore uncorrelated with current k_t .

We use as our two proxies investment and materials, and we write these input demands as $i_t(k_t, \omega_{jt}, \omega_{(-j)t})$ and $m_t = m(k_t, \omega_{jt}, \omega_{(-j)t})$. If the bivariate function (i_t, m_t) is one-to-one and onto with $(\omega_{jt}, \omega_{(-j)t})$ then this bivariate bijection can be inverted and there exist functions $g_j(\cdot)$ and $g_{(-j)}(\cdot)$ such that $\omega_{jt} = g_j(k_t, i_t, m_t)$ and $\omega_{(-j)t} = g_{(-j)}(k_t, i_t, m_t)$. For either j we approximate

$$\omega_j = g_j(k_t, i_t, m_t) = \mathbf{c}_j(k_t, i_t, m_t)' \beta_{\omega_j}$$

where $\mathbf{c}_j(k_t, i_t, m_t)$ is a known vector function of (k_t, i_t, m_t) chosen by researchers. The

nonparametric conditional mean function for either j is given as

$$E[\omega_{jt}|\omega_{t-1}] = f_j(\mathbf{c}_j(k_{t-1}, i_{t-1}, m_{t-1})' \beta_{\omega_j}) \quad j = 1, 2$$

for some unknown functions $f_1(\cdot)$ and $f_2(\cdot)$. The error now becomes

$$[\xi_{jt} + \epsilon_{jt}](\theta) = q_{jt} - \beta_l^j l_t - \beta_k^j k_t - \beta_m^j m_t - \gamma_{-j}^j q_{-jt} - f_j(\mathbf{c}_j(k_{t-1}, i_{t-1}, m_{t-1})' \beta_{\omega}) \quad j = 1, 2.$$

Let the set of conditioning variables be given as (e.g.) $x_{jt} = (q_{-j,t-1}, k_t, k_{t-1}, i_{t-1}, m_{t-1}, m_{t-2})$.

Let θ_0 denote the true parameter value. The conditional moment restrictions for each equation are given as

$$s(x_{jt}; \theta) \equiv E[[\xi_{jt} + \epsilon_{jt}](\theta)|x_{jt}] \text{ and } s(x_{jt}; \theta_0) = 0 \quad j = 1, 2.$$

6.3 Cost

Our main goal in estimating the cost function is to obtain product-level marginal costs and markups which we can then use to determine how investments in R&D affect prices and demand. In the absence of input prices one can, in principle, recover marginal costs from the production function alone following the method popularized by De Loecker and Warzynski (2012). Since we do have measures of input prices, we can estimate the cost function and recover more direct measures of marginal costs and markups. One issue we face is how to condition the variable cost function on the quantities of each product produced by the firm. This is a problem since firms may produce many different products (raising a dimensionality problem since we want to interact quantity of each good with other inputs in the cost function) and this set of products differs by firm. We overcome this problem by estimating the firm's variable cost function at the product-year level, where for each product we condition on its own quantity and a quantity index of all the other products produced by the firm (similarly to how we specify the product-level production function). Our empirical

specification for the variable cost function is thus

$$\log VC_{jt} = f^{vc}(k_{jt}, \Delta k_{jt}, \bar{w}_{jt}, z_{gjt}^1, q_{gjt}, q_{(-g)jt}, \omega_{gjt}) \quad (3)$$

where f^{vc} is a flexible (translog) function of logs in capital (k_{jt}), changes in capital (to account for adjustment costs), mean wage ($\bar{w}_{jt}$), an input price index (z_{gjt}^1)¹³, quantity of product g (q_{gjt}), the sum of all other products ($q_{(-g)jt}$), and product-level production efficiency (ω_{gjt}). Product-level marginal cost is calculated as $mc_{gjt} = \frac{VC_{jt}}{q_{gjt}} \frac{\partial \log VC_{jt}}{\partial \log q_{gjt}}$. Markups are then $\mu_{gjt}^c = \frac{p_{gjt}}{mc_{gjt}}$ where the superscript on μ indicates that this markup was recovered from the cost side rather than the demand side.

7 Results

7.1 Prices and Quantities

We begin by establishing the raw relationships between firm investments in R&D and product-level prices and output. Table 4 shows the results of three separate OLS regressions, where we regress (for example) log product price on lagged log spending in the three different types of R&D as well as a set of controls including firm/product characteristics, year and dummy variables indicating whether the firm made positive investments in each type of R&D. The dependent variables for each regression are listed at the top of each column. The results confirm our intuition regarding the nature and effect of product and process R&D on product characteristics. Specifically, column 1 shows that log prices increase when firms invest in product R&D (possibly because of an increase in product quality) while prices decrease when firms invest in process R&D (possibly due to improvements in manufacturing efficiency).

Column 2 shows the results for the same regression where now the dependent variable is

¹³We use the kg-weighted import price index, deflated by the product input requirement set, described in section 6.1.

the log quantity of the product sold. Intuitively, investments in product R&D are associated with decreases in quantity sold, while process R&D investments are associated with increases in output. In sum, a 10% increase in spending on product R&D increases prices by 3.9% and decreases quantity sold by 4.8%. The same increase in process R&D decreases prices by 1.3% on average, while increases quantity sold by 2.3%. Neither product nor process R&D has a significant relationship with the number of products sold by a firm, but general R&D does seem to be positively correlated with firm-level product scope. These results give us confidence that our measures of R&D make sense and do have significant relationships to the endogenous firm outcomes in which we are interested. But in order to determine the welfare effects of research and development, we need to understand the underlying mechanisms driving these relationships between R&D and market outcomes. To do this, we turn to our model of demand and supply.

7.2 Demand

Our main goal in estimating the demand model is to recover estimates of unobserved product quality ξ_{gj} that we can then use to measure how investments in R&D affect market outcomes and ultimately welfare. As discussed in section 6.1, we estimate separate demand models for each 6-digit market in our data, defined as the set of 8-digit product codes within each 6-digit nest. Table 5 shows two sets of results.

Since we estimate the model separately for each 6-digit product code, we end up with too many parameters to report. Instead, in the top panel, we show the results for one of our 6-digit product markets (Dining/Living Room Wooden Furniture, 940360) to give an idea of how our instruments perform. The first column shows the results for an OLS regression where we do not instrument for product price. We find a very small median price elasticity of -0.31. Since product quality is likely positively correlated with price, the estimates of the demand elasticity will be biased towards zero. Column 2 shows that we are unable to get sensible results using Hausman-style instruments. Columns 3-5 show the results for our

three individual instrumental variables. IV1 is the firm-level input price index, IV2 is the firm-level tariff index, and IV3 is the firm-product level best-input-price instrument. All three give similar results, with much more reasonable median elasticities of around -3.0. Note that basic producer theory implies that firms are unlikely to operate on the inelastic part of the demand curve, so we should expect (median) elasticities to be greater than -1 in absolute value. Columns 6 and 7 show the results when we use different combinations of the instruments. This allows us to test our instruments (we fail to reject the validity of our instruments) and we get similar results for our demand model parameters. The bottom panel shows summary stats for the entire pooled sample of results across all the 6-digit markets in our sample. Specifically, the median estimated elasticity across all of the 6-digit products when we use IV1 is -2.65. IV3 provides a median elasticity of -2.71. The fact that all three instruments provide relatively consistent results across all of the goods is reassuring and indicates that these may be useful instruments for estimating demand in other contexts. Our estimated demand model thus provides us with demand elasticities $\hat{\epsilon}_{gjt}$ and product quality $\hat{\xi}_{gjt}$ for each product g sold by firm j in period t . Given the structure of the demand model, we can also recover the implied markups (μ_{gjt}^d) and marginal costs (mc_{gjt}^d). In particular, a profit maximizing firm facing logit demand will set product prices such that $p_{gjt} = \frac{\epsilon_{gjt}}{1+\epsilon_{gjt}} mc_{gjt}$. Since we have $\hat{\epsilon}_{gjt} \equiv -\hat{\alpha} p_{gjt} (1 - s_{gjt})$, this provides us with markups $\hat{\mu}_{gjt}^d = \hat{\epsilon}_{gjt} / (1 + \hat{\epsilon}_{gjt})$ and marginal costs $\hat{mc}_{gjt}^d = p_{gjt} / \hat{\mu}_{gjt}^d$, where the superscripts indicate that these measures are recovered from the demand model.

7.3 Supply

7.3.1 Production function

We estimate the multi product production function described in subsection 4.1 using all firm-product observations within the same 2-digit CN code. Table 6 shows the results from the physical production function estimation for a subset of 6 industries. As we see, the input coefficients are all positive and in the range of other studies in the literature, except for labor

in two industries (coefficients are negative but not significant). The coefficients of $\ln Q_{-j}$ on the other end are all negative and significant.

7.3.2 Cost function

We then use our estimates of product-level production efficiency $\hat{\omega}_{gjt}$ to estimate the variable cost function as described in section 6.3. In particular we estimate a pooled cost function where each observation is a 6-digit product-firm-year tuple. Table 7 shows selected parameter estimates from the cost function. As expected, variable costs are increasing in quantity, mean wages and input prices, and decreasing in capital. Adjustment costs do seem to play a role, as variable costs are increasing in adjustments to the capital stock. Using the estimated variable cost function, we can calculate the marginal cost at the product level as $\widehat{mc}_{gjt}^c = \frac{VC_{jt}}{q_{gjt}} \frac{\partial \log \widehat{VC}_{jt}}{\partial \log q_{gjt}}$ where the translog form of the cost function gives us an analytical expression for $\frac{\partial \log \widehat{VC}_{jt}}{\partial \log q_{gjt}}$. Estimated cost-side markups are then $\hat{\mu}_{gjt}^c = \frac{p_{gjt}}{\widehat{mc}_{gjt}^c}$, where the superscripts indicate that these measures are recovered from the variable cost model.

7.4 The effect of R&D on Product Quality and Production Efficiency

So far we have estimates of product quality, marginal costs and markups from the demand model, estimates of production efficiency from the production function, and estimates of marginal costs and markups from the variable cost function. We can use these estimates to answer our questions of interest: How does R&D affect producer and consumer surplus? Is it mostly through the demand channel, or cost/supply channel? Do different types of R&D work through different channels, and are firms investing too much or too little in R&D?

The first step is to establish the relationship between investments in R&D and fundamentals of our models of supply and demand. In particular, we want to estimate how investments

in R&D drive the evolution of product quality:

$$\xi_{gjt} = f(\xi_{gjt-1}, r_{jt-1}^{\xi}, r_{jt-1}^{\omega}, r_{jt-1})$$

production efficiency:

$$\omega_{gjt} = f(\omega_{gjt-1}, r_{jt-1}^{\xi}, r_{jt-1}^{\omega}, r_{jt-1})$$

and marginal cost:

$$mc_{gjt} = f(r_{jt-1}^{\xi}, r_{jt-1}^{\omega}, r_{jt-1})$$

where r_{jt}^{ξ} is investment (in log dkk) in “Product R&D”, r_{jt}^{ω} is investment in “Process R&D”, and r_{jt} is investment in “General R&D”.

We begin by examining the relationship between product quality and R&D. We run a series of regressions of the form

$$\xi_{gjt} = \beta_0 + \mathbf{RD}_{t-1}\beta + \rho\xi_{gjt-1} + \mathbf{X}_{gjt}\gamma + \varepsilon_{gjt}$$

where $\mathbf{RD}_{t-1}$ and $\mathbf{X}_{gjt}$ are vectors of firm-level R&D variables and product-firm controls. We run this regression separately for each of the measures of product quality ξ_{gjt} obtained from our different demand model specifications.

The results are in Table 8. The first column shows that the measure of product quality obtained from the demand model estimated with OLS seems to be mostly uncorrelated with the different measures of R&D. Once we move to columns 3 through 7 where we are using product quality obtained by instrumenting for price in the demand model we obtain much larger and significant effects of RD on quality. Overall, we find that lagged investments in product R&D significantly increase product quality in the subsequent year, with similar coefficients across all three single-instrument specifications. On the other hand, investments in process R&D seems to have little to no significant effect on product quality across most of the specifications. General R&D seems to have a negative though not statistically significant

effect on product quality. This confirms our ex-ante understanding that investments in product R&D are aimed at improving product quality and demand. It also reassures us that the demand model is picking up true variation in quality across products and firms.

Given our demand model, we can also decompose the effect of R&D investment on product prices into two components: the marginal cost and markups. Recall that $\log p_{gjt} = \log mc_{gjt} + \log \mu_{gjt}$. Given this log linear identity, we can run the following regression

$$\log y = \beta_0^d + \mathbf{RD}_{t-1}\beta^{\mathbf{d}} + \mathbf{X}_{gjt}\gamma^{\mathbf{d}} + \varepsilon_{gjt}^d$$

where y is price, marginal cost or markups. The estimated coefficients from the marginal cost and markup regressions will add up to the coefficients from the price regression, allowing us to see whether R&D investments affect prices primarily by changing marginal costs (the supply channel), or by changing markups (the demand channel). The results from these regressions are shown in Table 9. The first column is the same regression as column 1 from table 4 with the exception that here we are restricted to a smaller sample due to the limited intersection of the R&D and price-instrument data. The results are essentially the same despite this much smaller sample, with product R&D investments increasing prices and process R&D decreasing prices. The next two columns decompose those marginal effects on price into the effect through marginal costs and the effect through markups. We find that almost all (87%) of the increase in prices that comes from product R&D is coming from increases in the marginal cost of production. Of the 5.5% increase in price that comes from a 10% increase in investment, 4.8 percentage points are due to increasing marginal costs while only 0.7 percentage points are from an increase in markups. The last column shows a regression of $\log(-\epsilon_{gjt})$ on the same variables. The negative sign on the product R&D coefficient means that investments in product R&D tend to move products and firms onto the less elastic portion of the demand curve (note that since $\epsilon_{gjt} < 0$ a decrease in $\log(-\epsilon_{gjt})$ implies the elasticity decreases in absolute value). This is what allows the firm to increase

its markup for that product.

Investments in process R&D also seem to act almost entirely through *reductions* in marginal costs and have no effect on markups. This is not entirely surprising, since we saw that in Table 8 that process R&D has little to no effect on product quality. This implies that through the lens of the demand model, these firms tend to pass all of the cost savings derived from increased process efficiency and decreased marginal costs on to consumers through reductions in price. Interestingly, general R&D investments seem to have an even larger effect on the demand elasticity and markups, with a 10% increase in general R&D investments leading to 1.2% increases in markups.

Finally, we can do exactly the same decomposition but using the measures of marginal cost and markups derived from the supply side of the model (the production and variable cost functions). Note that these two approaches use very different data and very different assumptions, and so we might not expect ex-ante that our estimates from the demand and supply side will coincide. The results are shown in Table 10. As before, the first column is the regression of log prices on product-firm controls and our set of (lagged) R&D investment measures. The point estimates are slightly different from the results in Table 9 due to the differing sample size, but not statistically different. Perhaps surprisingly, despite using totally different data and methods to estimate the supply-side marginal cost and markups, we find remarkably similar qualitative results. In particular, almost all of the effect of both product and process R&D investments on prices is due to changes in marginal costs. About 2/3rds of decrease in price from increased process R&D investment is due to decreases in marginal costs, with the other 1/3rd coming from decreases in the markup. The supply-side estimates actually suggest that increases in product R&D increase marginal costs *more* than the increase in prices, with the relative difference coming from a *decrease* in markups, which is unexpected. The fourth column shows the effect of R&D investments on product-level production efficiency ω_{gjt} . As expected, investments in product R&D decrease production efficiency (which explains the rise in marginal costs), while process R&D increases

productivity (lowering marginal costs).

A significant recent literature has use supply-side data to estimate firm and product markups in an attempt to estimate how markups and thus firm market power has been changing over time. There's a related literature which debates the degree to which these estimates of markups are reasonable, given the assumptions that go into the estimation process. We can get a sense of how reasonable these markups are by comparing them to the markups we obtain from the demand model. In our sample, the median markup obtained from the demand model is 1.07, while the median markup obtain from the cost function is 1.18. That's not an insignificant difference, but considering that they come from somewhat different samples, it's reassuring that neither is particularly unrealistic.

8 Conclusion

We examine the joint decisions of pricing and R&D expenditures in newly disaggregated firm-level panel data sets. We show how to use these data to estimate the primitive parameters of demand, output production, and R&D production in the face of simultaneity concerns. Our results document two intuitive effects of R&D on firm's profitability. More process innovation is shown to be associated with lower marginal costs and higher productivity, while product innovation leads to increase in product quality. Welfare is therefore increasing through these two distinct channels.

References

Akerberg, Daniel A, Kevin Caves, and Garth Frazer, "Identification properties of recent production function estimators," *Econometrica*, 2015, 83 (6), 2411–2451.

- Aghion, Philippe, Nick Bloom, Richard Blundell, Rachel Griffith, and Peter Howitt**, “Competition and innovation: An inverted-U relationship,” *The Quarterly Journal of Economics*, 2005, *120* (2), 701–728.
- Amiti, Mary, Oleg Itskhoki, and Jozef Konings**, “International shocks, variable markups, and domestic prices,” *The Review of Economic Studies*, 2019, *86* (6), 2356–2402.
- Arrow, Kenneth**, “Economic Welfare and The Allocation of Resources for Invention,” *The Rate and Direction of Inventive Activity*, Princeton University Press and NBER, 1962.
- Belloni, Alexandre, Daniel Chen, Victor Chernozhukov, and Christian Hansen**, “Sparse models and methods for optimal instruments with an application to eminent domain,” *Econometrica*, 2012, *80* (6), 2369–2429.
- Bernard, Andrew B., Ilke Van Beveren, and Hylke Vandenbussche**, “Multi-product exporters, carry-along trade and the margins of trade,” *The Review of Economic Studies*, Forthcoming.
- Berndt, Ernst R and Catherine J Morrison**, “Capacity utilization measures: underlying economic theory and an alternative approach,” *American Economic Review*, 1981, *71* (2), 48–52.
- Berry, Steven, James Levinsohn, and Ariel Pakes**, “Automobile prices in market equilibrium,” *Econometrica*, 1995, pp. 841–890.
- Berry, Steven T**, “Estimating discrete-choice models of product differentiation,” *The RAND Journal of Economics*, 1994, pp. 242–262.
- Blundell, Richard and James L Powell**, “Endogeneity in nonparametric and semiparametric regression models,” *Econometric society monographs*, 2003, *36*, 312–357.

- Bosch, J Van Den, B Cassiman, Konings J, and S Vanormelingen**, “The Impact of Innovation on Productivity: Profitability and Technical Efficiency,” Technical Report 2019.
- Chernozhukov, Victor, Christian Hansen, and Martin Spindler**, “Post-selection and post-regularization inference in linear models with many controls and instruments,” *American Economic Review*, 2015, *105* (5), 486–90.
- Dhyne, E., A. Petrin, and F. Warzynski**, “Deregulation and Spillovers in Multi-Product Production Settings,” *Working Paper*, 2014.
- Dhyne, Emmanuel, Amil Petrin, Valerie Smeets, and Frederic Warzynski**, “Multi product firms, import competition, and the evolution of firm-product technical efficiencies,” Technical Report, National Bureau of Economic Research 2017.
- Diewert, W Erwin**, “Functional forms for profit and transformation functions,” *Journal of Economic Theory*, 1973, *6* (3), 284–316.
- Doraszelski, Ulrich and Jordi Jaumandreu**, “R&D and productivity: Estimating endogenous productivity,” *The Review of Economic Studies*, 2013, *80* (4), 1338–1383.
- Eslava, Marcela, John Haltiwanger, Adriana Kugler, and Maurice Kugler**, “The effects of structural reforms on productivity and profitability enhancing reallocation: evidence from Colombia,” *Journal of Development Economics*, 2004, *75* (2), 333–371.
- Flynn, Zach, Amit Gandhi, and James Traina**, “Measuring markups with production data,” *Available at SSRN 3358472*, 2019.
- Foster, Lucia, John Haltiwanger, and Chad Syverson**, “Reallocation, firm turnover, and efficiency: selection on productivity or profitability?,” *American Economic Review*, 2008, *98* (1), 394–425.

- Hall, Bronwyn H, Jacques Mairesse, and Pierre Mohnen**, “Measuring the Returns to R&D,” in “Handbook of the Economics of Innovation,” Vol. 2, Elsevier, 2010, pp. 1033–1082.
- Holmes, Thomas J and James A Schmitz Jr**, “Competition and productivity: a review of evidence,” *Annu. Rev. Econ.*, 2010, *2* (1), 619–642.
- Jaumandreu, Jordi and Jacques Mairesse**, “Innovation and Welfare: Results from Joint Estimation of Production and Demand,” Technical Report, National Bureau of Economic Research 2010.
- Lau, Lawrence J**, “A characterization of the normalized restricted profit function,” in “The Hamiltonian Approach to Dynamic Economics,” Elsevier, 1976, pp. 131–163.
- Levinsohn, James and Amil Petrin**, “Estimating production functions using inputs to control for unobservables,” *The Review of Economic Studies*, 2003, *70* (2), 317–341.
- Loecker, Jan De and Frederic Warzynski**, “Markups and firm-level export status,” *American Economic Review*, 2012, *102* (6), 2437–71.
- , **Jan Eeckhout, and Gabriel Unger**, “The rise of market power and the macroeconomic implications,” *The Quarterly Journal of Economics*, 2020, *135* (2), 561–644.
- Maican, Florin G, Matilda Orth, Mark J Roberts et al.**, “The Dynamic Impact of Exporting on Firm R&D Investment,” Technical Report, National Bureau of Economic Research 2020.
- Mundlak, Yair and Assaf Razin**, “On multistage multiproduct production functions,” *American Journal of Agricultural Economics*, 1971, *53* (3), 491–499.
- Olley, G Steven and Ariel Pakes**, “The Dynamics of Productivity in the Telecommunications Equipment Industry,” *Econometrica*, 1996, pp. 1263–1297.

- Peters, Bettina, Mark J Roberts, and Van Anh Vuong**, “Firm R&D investment and export market exposure,” Technical Report, National Bureau of Economic Research 2018.
- , —, —, and **Helmut Fryges**, “Estimating dynamic R&D choice: an analysis of costs and long-run benefits,” *The RAND Journal of Economics*, 2017, 48 (2), 409–437.
- Raval, Devesh**, “Testing the production approach to markup estimation,” *Available at SSRN 3324849*, 2020.
- Roberts, Mark J and Dylan Supina**, “Output price and markup dispersion in micro data: the roles of producer heterogeneity and noise,” in “Industrial Organization,” Emerald Group Publishing Limited, 2000.
- Smeets, Valerie and Frederic Warzynski**, “Estimating productivity with multi-product firms, pricing heterogeneity and the role of international trade,” *Journal of International Economics*, 2013, 90 (2), 237–244.
- and —, “Multi-Product Firms in French Manufacturing: the Rise of Subcontracting and Global Sourcing,” *Working Paper*, 2019.
- Syverson, Chad**, “Market structure and productivity: A concrete example,” *Journal of Political Economy*, 2004, 112 (6), 1181–1222.
- Wooldridge, Jeffrey M**, “On estimating firm-level production functions using proxy variables to control for unobservables,” *Economics letters*, 2009, 104 (3), 112–114.

9 Tables

Table 1: Summary statistics on the production survey

year	# products		# firms	# single product firms	# firms producing btw 2 and 5 products		# firms producing more than 5 products	#	
	mean	median				%			%
1997	3.12	2	4314	1,768	40.98%	2,054	47.61%	492	11.40%
1998	3.07	2	4407	1,920	43.57%	2,015	45.72%	472	10.71%
1999	3.12	2	4204	1,850	44.01%	1,907	45.36%	447	10.63%
2000	3.18	2	4155	1,824	43.90%	1,875	45.13%	456	10.97%
2001	3.24	2	4089	1,813	44.34%	1,833	44.83%	443	10.83%
2002	3.28	2	4038	1,816	44.97%	1,773	43.91%	449	11.12%
2003	3.30	2	3936	1,766	44.87%	1,712	43.50%	458	11.64%
2004	3.29	2	3787	1,709	45.13%	1,646	43.46%	432	11.41%
2005	3.32	2	3676	1,659	45.13%	1,592	43.31%	425	11.56%
2006	3.36	2	3606	1,632	45.26%	1,575	43.68%	399	11.06%
2007	3.29	2	3569	1,623	45.47%	1,551	43.46%	395	11.07%
2008	3.18	2	3712	1,744	46.98%	1,561	42.05%	407	10.96%
2009	3.21	2	3457	1,659	47.99%	1,408	40.73%	390	11.28%
2010	3.22	2	3238	1,562	48.24%	1,316	40.64%	360	11.12%
2011	3.33	2	2935	1,389	47.33%	1,195	40.72%	351	11.96%
2012	3.53	2	2582	1,209	46.82%	1,045	40.47%	328	12.70%
2013	3.48	2	2643	1,245	47.11%	1,077	40.75%	321	12.15%
2014	3.59	2	2612	1,246	47.70%	1,059	40.54%	307	11.75%
2015	3.61	2	2587	1,271	49.13%	1,001	38.69%	315	12.18%
2016	3.64	2	2603	1,282	49.25%	1,014	38.96%	307	11.79%
2017	3.73	2	2592	1,293	49.88%	988	38.12%	311	12.00%
2018	3.69	1	2637	1,325	50.25%	990	37.54%	322	12.21%

Table 2: Summary statistics on R&D spending

	# firms surveyed	# firms with:			% firms with:		
		own R&D division	pos R&D spending	pos R&D personnel	own R&D division	pos R&D spending	pos R&D personnel
1999	4,168	588	1,034	1,034	14.11%	24.81%	24.81%
2000	4,077		765	689		18.76%	16.90%
2001	3,536	303	566	560	8.57%	16.01%	15.84%
2002	3,736		697	696		18.66%	18.63%
2003	3,637	282	723	722	7.75%	19.88%	19.85%
2004	3,490		618	616		17.71%	17.65%
2005	3,464	457	692	692	13.19%	19.98%	19.98%
2006	3,908		599	598		15.33%	15.30%
2007	4,534	626	841	841	13.81%	18.55%	18.55%
2008	4,438	889	1,488	1,436	20.03%	33.53%	32.36%
2009	4,545	850	1,064	1,045	18.70%	23.41%	22.99%
2010	4,322	815	1,027	1,021	18.86%	23.76%	23.62%
2011	4,424	782	983	981	17.68%	22.22%	22.17%
2012	4,698	848	1,081	1,081	18.05%	23.01%	23.01%
2013	4,787	849	985	979	17.74%	20.58%	20.45%
2014	4,901	854	1,006	1,005	17.43%	20.53%	20.51%
2015	5,044	870	1,044	1,045	17.25%	20.70%	20.72%
2016	4,522	854	1,095	1,097	18.89%	24.21%	24.26%
2017	3,321	813	1,034	1,034	24.48%	31.14%	31.14%

Table 3: Share of product, process and general R&D spending

	(1)	(2)	(3)
	Product	Process	General
2001	74.79%	15.44%	9.77%
2003	76.09%	16.65%	7.26%
2005	80.76%	13.52%	5.72%
2007	74.32%	19.29%	6.39%
2009	69.26%	19.75%	10.99%
2011	75.07%	15.96%	8.96%

Note: Share of total R&D spending on product innovation, process innovation, and general development by year from the Danish R&D survey.

Table 4: Elasticities of Price and Quantity w.r.t. R&D

	(1)	(2)	(3)
	Log Price	Log Quantity	Number of Products
Product R&D	0.388 (0.030)	-0.484 (0.035)	0.088 (0.110)
Process R&D	-0.127 (0.036)	0.225 (0.042)	-0.084 (0.143)
General R&D	0.014 (0.034)	0.054 (0.040)	0.294 (0.140)
Observations	13,069	13,069	1,778

Note: R&D measures are in 1-period lagged logs of spending (dkk). Controls include dummies for all types of R&D, year, product rank, number of products, and firm size. Standard errors are in parentheses.

Table 5: Demand Estimation Results

	OLS	Hausman	IV1	IV2	IV3	1+3	1+2+3
Detailed Demand Estimation							
Example: Dining/Living Room Wooden Furniture (940360)							
Price (α)	0.40 (0.02)	-5.8 (7.11)	4.0 (1.12)	3.9 (0.51)	4.1 (0.61)	4.7 (0.62)	4.2 (0.55)
Obs.	1,328	1,328	551	551	351	351	351
F-Stat	-	0.78	74.4	56.2	45.3	29.5	20.7
Med. Elast.	-0.31	4.43	-3.07	-2.99	-3.10	-3.58	-3.20
Overall (all cn6 codes with ≥ 50 obs)							
Obs.	66,868	66,868	34,484	34,240	20,759	19,293	19,192
Med. Elast.	-0.35	-0.01	-2.65	-2.77	-2.71	-1.88	-1.53

Note: Standard errors are in parentheses.

Table 6: Production Function Estimation Results

Production Function Estimates						
	Meat Products 16	Plastics 39	Fab. Metal 73	Machinery & Equip. 84	Electrical Machinery 85	Furniture 94
$\ln Q_{-j}$	-0.171 (0.012)	-0.095 (0.006)	-0.100 (0.005)	-0.112 (0.004)	-0.082 (0.008)	-0.135 (0.004)
ℓ	0.320 (0.065)	-0.095 (0.074)	-0.042 (0.052)	0.180 (0.049)	0.216 (0.097)	0.371 (0.048)
k	0.113 (0.037)	0.317 (0.036)	0.187 (0.024)	0.055 (0.018)	0.191 (0.037)	0.145 (0.019)
m	0.540 (0.047)	0.762 (0.059)	0.723 (0.038)	0.731 (0.042)	0.419 (0.075)	0.651 (0.037)
Obs.	1,412	2,539	3,942	7,411	2,032	3,957

Note: Standard errors are in parentheses.

Table 7: Cost Function Estimation Results

Cost Function Estimates		
	Variable Cost	(s.e.)
Quantity	0.28	(0.07)
Quantity ²	0.02	(0.001)
Capital	-1.30	(0.10)
Δ Capital	0.08	(0.01)
Mean Wage	2.56	(0.36)
Input Prices	0.13	(0.06)
Obs.	15,565	

Note: Standard errors are in parentheses. The independent variables in the translog cost function regression are firm capital, change in capital, mean wage, an input price index, output quantity of good g , output quantity of all other goods produced by the firm, and product-level production efficiency.

Table 8: The Effect of R&D on Product Quality (ξ_{gjt})

	(1) OLS	(2) Hausman	(3) IV1	(4) IV2	(5) IV3	(6) IV1+3	(7) IV1+2+3
Prod. R&D	-0.055 (0.031)	-0.011 (0.076)	0.463 (0.188)	0.474 (0.191)	0.552 (0.211)	0.363 (0.155)	0.229 (0.140)
Proc. R&D	-0.009 (0.038)	-0.048 (0.093)	-0.456 (0.238)	0.078 (0.244)	0.049 (0.269)	0.156 (0.197)	0.145 (0.179)
Gen. R&D	0.115 (0.038)	0.065 (0.094)	-0.254 (0.240)	-0.508 (0.246)	-0.446 (0.271)	-0.357 (0.198)	-0.205 (0.180)
Observations	2,956	2,954	2,323	2,361	1,498	1,503	1,499

Note: R&D measures are in 1-period lagged logs of spending (dkk). Controls include dummies for all types of R&D, year, product rank, number of products and firm size. S.E. (in brackets) clustered at firm level. Columns represent regressions on the ξ terms generated by different first-stage demand estimations.

Table 9: Effect of R&D on Prices, MC, Markups & Elasticities

	(1)	(2)	(3)	(4)
	$\log p_{gjt}$	$\log mc_{gjt}$	$\log \mu_{gjt}$	$\log (-\epsilon_{gjt})$
Prod. R&D	0.550 (0.098)	0.477 (0.094)	0.073 (0.029)	-0.207 (0.074)
Proc. R&D	-0.343 (0.124)	-0.317 (0.119)	-0.026 (0.038)	0.065 (0.100)
Gen. R&D	0.066 (0.125)	-0.056 (0.120)	0.122 (0.038)	-0.353 (0.101)
Observations	1,241	1,241	1,241	1,234

Note: R&D measures are in 1-period lagged logs of spending (dkk). Controls include dummies for all types of R&D, year, product rank, number of products and firm size. S.E. (in brackets) clustered at firm level. Elasticities calculated using IV3 specification.

Table 10: The Effect of R&D on MC, Markups & Productivity

	(1)	(2)	(3)	(4)
	$\log p_{gjt}$	$\log mc_{gjt}$	$\log \mu_{gjt}$	ω_{gjt}
Prod. R&D	0.678 (0.067)	0.769 (0.068)	-0.089 (0.025)	-0.227 (0.079)
Proc. R&D	-0.377 (0.084)	-0.263 (0.085)	-0.119 (0.031)	0.372 (0.098)
Gen. R&D	0.081 (0.061)	0.072 (0.081)	0.010 (0.022)	-0.102 (0.069)
Observations	3,589	3,589	3,589	3,292

Note: R&D measures are in 1-period lagged logs of spending (dkk). Controls include dummies for all types of R&D, year, product rank, number of products and firm size. S.E. (in brackets) clustered at firm level. Markups and Marginal Costs calculated using estimated cost function.