

**From the holdings of the Department of Special Collections , Princeton University Library,
Princeton, NJ**

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyright law. We do not charge any permission or usage fees for the publication of images of material in our collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff. For more information regarding copyright, credit, and citation please review the information on the Copyright, Credit and Citation Guidelines page of our website.

<https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines>. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. <https://library.princeton.edu/special-collections/ask-us>

Doctoral Dissertations, 2011-2026

Economics

2020

Essays in Industrial Organization and Trade

Hu, Kathleen

Princeton, NJ : Princeton University

<http://arks.princeton.edu/ark:/88435/dsp012227ms61f>

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. <https://www.copyright.gov/title17/title17.pdf>

ESSAYS IN INDUSTRIAL ORGANIZATION AND TRADE

KATHLEEN HU

A DISSERTATION
PRESENTED TO THE FACULTY
OF PRINCETON UNIVERSITY
IN CANDIDACY FOR THE DEGREE
OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE
BY THE DEPARTMENT OF
ECONOMICS
ADVISER: STEPHEN REDDING

JANUARY 2020

© Copyright by Kathleen Hu, 2019.

All rights reserved.

Abstract

In this dissertation, I explore productivity in multi-product establishments, complementarities in production, and tax harmonization in the EU. The first chapter develops a methodology for estimating within establishment, across product productivities. The key challenge in estimating productivity at the product-level within multi-product establishments is that even when there is data on the output of each good, inputs are recorded at the establishment-level. By leveraging optimality conditions in an establishment's cost minimization decisions, I infer the allocation of inputs across products in order to estimate establishment-product quantity-based total factor productivities. I find that within-establishment productivities exhibit nontrivial heterogeneity that is consistent with non-random product selection.

Then in chapter two, I examine the factors that determine the patterns of coproduction within multi-product establishments, and in particular, the relationship between an establishment's scope and the productivities of its goods. Using establishment-product data from India, I find that productivities are positively correlated with within-industry scope, which is consistent with both selection on productivity and with complementarities across products. To disentangle these two effects, I exploit variation in the timing of small scale industry dereservation in India and find only weak evidence of a causal effect of scope on productivity, suggesting that endogenous selection with correlated productivities may be the best explanation for the observed patterns of coproduction.

The third chapter, coauthored with Rowan Shi, addresses a different topic: the impact of tax dispersion on the allocation of firms, production, and resources across countries. We examine the potential impact of corporate tax harmonization within the European Union by building a general equilibrium model of multinational location and production with taxation and calibrating to recent aggregate data for the EU. If corporate taxation in the EU were harmonized to a common rate that would generate

the same amount of total tax revenue, our model predicts a significant reallocation of production from the countries that raise their tax rates to those that lower their rates.

Acknowledgements

I would like to thank my advisor Steve Redding, my committee members, and the participants of the trade lunch for invaluable feedback and support throughout my time at Princeton. I thank Rowan Shi for all her hard work on our joint project. Thanks to my friends and family, and especially my parents for always believing in me. Finally, I would also like to thank the International Economics Section for sponsoring me in 2016 and 2017 through the IES Summer Fellowship.

Contents

Abstract	iii
Acknowledgements	v
List of Tables	viii
List of Figures	ix
1 Measuring Productivity in Multi-Product Establishments	1
1.1 Introduction	1
1.2 Data: the Indian ASI	5
1.3 Theoretical Framework	7
1.3.1 A Model of Multi-Product Establishments	7
1.3.2 Introducing a Model of Demand	14
1.4 Methodology for Estimating Productivities	16
1.4.1 Production Function Estimation	18
1.4.2 Marginal Cost Estimation	21
1.4.3 Computing Productivities	23
1.5 Empirical Findings	24
1.5.1 Productivity in India	26
1.5.2 Within and Between Establishment Heterogeneity	28
1.5.3 The Core Competency Ladder	30
1.5.4 Comparison to Alternative Estimation Methods	32
1.6 Conclusions	35

2	Coproduction and Complementarities in India	38
2.1	Introduction	38
2.2	Coproduction in India	43
2.2.1	Characteristics of Multi-Product Establishments	44
2.2.2	Product Switching	46
2.2.3	Heterogeneity of Products Within Establishments	53
2.3	Productivity in Multi-Product Establishments	57
2.3.1	Theoretical Framework	57
2.3.2	Productivities of Coproduced Goods	61
2.3.3	Complementarities in Production	66
2.4	Conclusions	72
3	Evaluating Tax Harmonization in the European Union	
	<i>with Rowan Shi</i>	75
3.1	Introduction	75
3.2	Corporate Taxation in the European Union	78
3.3	A Model of Multinational Location Choice	80
3.3.1	Households	81
3.3.2	Capital Owners	82
3.3.3	The Nontradable Sector	82
3.3.4	The Tradable Sector	83
3.3.5	Equilibrium	88
3.4	Quantitative Methods	91
3.4.1	Solving for Equilibrium	92
3.4.2	Calibration	95
3.5	Counterfactual Tax Harmonization in th EU	100
3.5.1	The Spatial Distribution of Production	101
3.5.2	Employment and Wages	103

3.5.3	Tax Revenue and Worker Welfare	105
3.6	Conclusion	108
A	India’s Annual Survey of Industries	111
A.1	Data Variables	111
A.2	ASI Commodity Classification	112
A.3	Dereserved Products in the ASI	113
B	Productivity Estimation	114
B.1	Estimated Factor Intensities	114
B.2	Estimated Elasticities of Substitution by Product Group	116
B.3	Over-Identification Check of the Model’s Assumptions	116
C	Tax Dispersion and MP	119
C.1	Statutory Corporate Tax Rates in 2014	119
C.2	Tax Rates and Calibrated Productivity	119
C.3	Over-Identification Tests of the Calibrated Model	121
	Bibliography	124

List of Tables

1.1	Multi-product, industry, and sector establishments	5
1.2	Variability of productivity	29
1.3	Productivity under alternative estimation assumptions	33
2.1	Establishment characteristics	46
2.2	Frequency of product adding and dropping	47
2.3	Frequency of product and industry adding and dropping	48
2.4	Establishment characteristics and product activity	49
2.5	Establishment productivity and product adding	51
2.6	Product dropping	52
2.7	Output shares by product rank	54
2.8	Productivity and product switching	63
2.9	Productivity in multi-product establishments	64
2.10	Productivity and establishment scope	65
2.11	IV first stage regression	70
2.12	Complementarities regression	72
3.1	Calibration of parameters	99
A.1	Product dereservation by year	113
B.1	Estimates of factor intensities	115
B.2	Estimates of demand elasticities of substitution	116

C.1	Corporate tax rates	120
C.2	Calibrated core productivity and tax rate	121
C.3	Calibrated productivity shifter and tax rate	122
C.4	Over-identification moments	122
C.5	Over-identification of trade shares	123
C.6	Over-identification of MP shares	123

List of Figures

1.1	Productivity and price as a proxy for quality	26
1.2	Distributions of log productivity demeaned	27
1.3	Productivity for single product and multi-product establishments . .	28
1.4	Price for single product and multi-product establishments	28
1.5	Core competency ladder	32
1.6	Robustness against constant markups	33
1.7	Productivity using revenue shares	34
2.1	Within establishment product distribution	55
2.2	Coproduction matrix	56
3.1	Map of corporate tax rates across the EU	79
3.2	Tax rates and productivity	80
3.3	Counterfactual production shares	102
3.4	Counterfactual MP	102
3.5	Counterfactual firm entry	103
3.6	Counterfactual trade	104
3.7	Counterfactual employment	104
3.8	Counterfactual wages	105
3.9	Counterfactual tax revenue and welfare	106
B.1	Over-identification check	117

B.2 Over-identification check in bins	118
---	-----

Chapter 1

Measuring Productivity in Multi-Product Establishments

1.1 Introduction

Heterogeneity in the productivity of firms is an important element in addressing various economic questions, including the firm size distribution, aggregate productivity, and the reallocation of resources following a shock to the economy. Heterogeneity across firms is a topic that has been extensively explored in the literature, but there may be additional heterogeneity in the productivity and composition of products inside a firm. In this paper, I examine the within-establishment component of heterogeneity by estimating productivities that can vary across the products in an establishment's mix.¹ The importance of this channel could hint at selection effects and strategic linkages across products that explain how firms choose their product mix and how resources are reallocated both across and within firms when there are changes in the economic environment.

¹Here I begin discussion of multi-product establishments instead of multi-product firms because while much of the literature has studied these issues at the firm level, the dataset I use in this paper is establishment-level, and so my results speak to productivities within establishments. In the subsequent data section, I discuss the establishment distinction in more detail.

The growing research focusing on multi-product firms in recent years has demonstrated that multi-product firms are systematically different from single product ones across a range of observables: multi-product firms are larger, employ more people, and are more likely to export. They also tend to be more productive, as measured by firm-level productivity estimates such as TFPR and labor productivity. These facts as well as limitations of the available multi-product data imply that techniques used for single product firms need to first be adapted before they can be used for multi-product estimation.

In order to look at productivity heterogeneity within firms or establishments, it is necessary to first develop a model that incorporates product-specific productivities and then estimate these in the face of certain data limitations, namely the lack of product-level input data. As a theoretical framework for my estimation, I use a model of multi-product establishments that is both tractable and also offers richness to the interactions across products. Productivities are specific to the establishment as well as the product, allowing heterogeneity across and within establishments. The model imposes no restrictions on the correlations between productivities while also allowing endogeneity in productivity through a tractable channel to have theoretical mechanisms capable to generating the observed relationship between the establishment's endogenous product choices and productivities (see Bernard, Redding, and Schott (2010) [10], Goldberg et al (2010) [23]). I also incorporate a demand structure that allows cannibalization across an establishment's product lines, which generates additional strategic effects in product choices (the importance of which are discussed in Eckel and Neary (2010) [18] and Hottman, Redding, and Weinstein (2016) [26]).

For the estimation of product-specific productivities, the key empirical challenge is dealing with unobserved allocation of inputs at the product-level. I have data on multi-product establishments that includes quantity and price information for detailed products, but input expenditures and usage are recorded at the establishment level.

This is also true of all other multi-product datasets I am aware of. When estimating productivity for single product establishments or at the establishment level, we know the amount of each input used. With multi-product plants, only the total across all products is observed. Therefore, estimating at the product level requires inferring how the establishment allocates each input across its products. In previous literature, revenue shares have been used to allocate inputs across products, or the production technology has been assumed to be the same for all products. Both of these methods limit the heterogeneity across products by removing certain strategic interactions and reducing the scope of analysis to establishments producing within a narrower range of products. In this paper, I exploit first-order conditions for optimal input allocation to derive a model-consistent input allocation share rule that can be applied to establishments producing across a broad range of products. I then estimate the components of the input share equation in steps. Ultimately, I can infer the optimal quantities of inputs allocated to each product the establishment produces and use them to estimate total factor productivity.

I estimate product-specific productivities for a panel of Indian establishments and measure the relative importance of the within establishment component of productivity heterogeneity. I find that even after purging productivity measures of the establishment component, there remains significant heterogeneity in the relative productivity across products within establishments. Applying my estimates to the core competency idea, I examine the steepness of the competency ladder for establishments of different scope. Finally, I show how my estimates compare to ones using other assumptions and methods from the literature.

This paper is not the first to address the topic of productivity and heterogeneity in multi-product firms. This topic builds on earlier work exploring firm heterogeneity (Melitz 2003 [33], Bernard, Eaton, Jensen, and Kortum 2003 [9], Hottman, Redding, and Weinsten 2016 [26]). Since then, there has been a growing body of work on

multi-product firms. Bernard, Redding, and Schott 2010 [10] and Goldberg et al 2010 [23] established a set of empirical facts about multi-product firms and in particular how they differ systematically from single product firms. Many papers in the theoretical literature on multi-product firms, including Eckel and Neary 2010 [18] and Nocke and Yeaple 2014 [35], discuss endogenous product choices and reallocation following trade shocks. Among the more recent papers that estimate productivity in multi-product firms, Dhyne, Petrin, Smeets, and Warzynski 2017 [16] estimate firm-product technical efficiencies using a joint multi-product production function. Their method has the advantage of not needing to explicitly allocate observed firm-level inputs across products, but can only accommodate two unobserved technical efficiency terms by using investment and materials as proxies. The papers that are perhaps most similar to mine are De Loecker, Goldberg, Khandelwal, and Pavcnik 2016 [14] and Orr 2017 [37]. De Loecker et al use similar steps to estimate production functions for multi-product firms while also allowing greater flexibility in the functional forms, but their identification of input allocations across products relies on productivity being constant across all of a firm's products. Thus, their method is unable to speak to heterogeneity in productivity of products within firms. Orr estimates product-specific productivities but is limited to firms producing within a single industry in which factor intensities are assumed to be the same across all products. My paper builds on these earlier works to estimate establishment-product productivities across a broad range of industries.

The organization of the remainder of the paper is as follows: Section 1.2 presents details of the multi-product dataset, Section 1.3 sets up a theoretical framework of multi-product production and competition, Section 1.4 details the estimation procedure for product-specific productivities, Section 1.5 describes empirical findings, and Section 1.6 concludes.

1.2 Data: the Indian ASI

I use establishment-product panel data for India to measure productivity. This data, from India’s Annual Survey of Industries (ASI), collects plant-level information on establishment characteristics and input use as well detailed information about the set of output products and material inputs ². I use data for the years 2000 through 2010. The ASI is a census of all large establishments and a survey of smaller establishments (ones with fewer than 100 employees). This particular data is well-suited for research on multi-product firms because of the detailed information on the products produced within the plant. Additionally, the ASI captures many establishments that produce not only multi-products but across multiple industries. The procedure for estimating productivities I develop allows for differences in factor intensities and markups across an establishment’s products. Allowing these differences across the products that are produced by an establishment becomes relevant in the ASI data, which covers many establishments producing a wide scope of products, including across different industries. Table 1.1 shows the share of establishments in the data that produce across multiple products, industries, or sectors. Products, industries, and sectors are defined using the product coding system of the ASI as unique 5-, 2-, or 1-digit codes, respectively.

	Within category				
	Share firms	Share revenue	Mean	Median	Maximum
Multi-product	47.4%	79.9%	3.07	2	23
Multi-industry	19.1%	50.1%	2.37	2	9
Multi-sector	8.0%	32.4%	2.11	2	6

Table 1.1: Facts on multi-product establishments in the ASI 2000-2010

²Note that while the ASI includes information about the products that are used by an establishment as material inputs, it says nothing about which inputs are used in the production of which outputs.

Individual products in the ASI are identified by a 5 digit ASI Commodity Classification (ASICC) code. In my data sample, I observe 6,128 unique product codes in use. The ASICC, like other commodity classifications, is hierarchical. The first one, two, and three digits of the product codes group products into increasingly finer categories. When referring to broader categories of products, I define “industry” as the 2-digit ASICC and “sector” or “product group” as the 1-digit ASICC. An example of a 5 digit code is 63404 “Cotton pillow covers”, while 2 digit codes include 63 “Cotton, cotton yarn, and fabrics” and 1 digit codes consist of broad categories such as “Agriculture”, “Textiles”, and “Chemicals”. I observe 63 unique industries and 9 broad sectors in the data.

Table 1.1 also reports the share of revenue generated by the category in the second column. Multi-product establishments account for a nontrivial share (47.39%) of total establishments and account for a proportionally larger 79.93% of total revenue. This is consistent with the findings by Bernard, Redding, and Schott (2006) [10] for the US, where 39% of firms are multi-product and account for 87% of revenue. The last three columns give the mean, median, and maximum number of products, industries, or sectors for establishments that produce more than one of these categories. Out of over 300,000 establishment-years, there are many operating across multiple industries and sectors.

The primary variables of interest in the data are physical quantities and prices for the establishment’s set of output products. Physical quantities are generally reported in a standard unit of measure for each product. I also use the establishment level data on labor, material, and capital inputs. One advantageous feature of the ASI is that it records plant level prices for labor and intermediate materials. Material inputs are recorded as a set of product codes with quantities and prices much like the output products, but there is no information in the data on which output product lines use which inputs. For tractability in the modeling and estimation of production,

I aggregate the individual input products into a single measure of materials for the establishment using a Cobb Douglas aggregation function that controls for input prices. Details on the materials aggregator and other data variables are included in Appendix A.

Aside from the product-specific data within establishments, all other variables are at the establishment level. The data contains no identifiers to link establishments belonging to same firm. However, the available information does show that most establishments included in the estimation belong to single-establishment firms. Therefore, I conduct analysis at the establishment level, and for the purposes of this project, I make not distinctions between establishment and firms when connecting the estimates in this paper to discussions of multi-product firms in the literature.

1.3 Theoretical Framework

1.3.1 A Model of Multi-Product Establishments

I introduce a model of multi-product production in which the set of products chosen by a plant is an endogenous decision. Establishments thus self-select into being single product or multi-product. By imposing some structure on the production functions, I am able to infer the optimal input allocations that a plant would choose for each of the products in its set. This then allows me to estimate product-specific productivities for each plant.

I consider the production of each good within an establishment as a function:

$$Y_{ift} = L_{ift}^{\beta_L^i} M_{ift}^{\beta_M^i} K_{ift}^{\beta_K^i} \phi_{ift} \quad (1.1)$$

For establishment f to produce Y_{ift} units of product i at time t , f combines labor, materials, and capital inputs using a Cobb Douglas technology that is specific to

product i . Labor and materials are static inputs chosen freely each period, whereas capital is dynamic. The production is scaled by a Hicks neutral productivity term ϕ_{ift} that is specific to the product and establishment. X_{ift} is the quantity of input X from $\{L, M, K\}$ that the firm allocates to the production of product i . I assume each input is completely allocated across an establishment's set of products:

$$\sum_{i \in \mathcal{J}_{ft}} X_{ift} = X_{ft}, \quad X \in \{L, M, K\} \quad (1.2)$$

$\mathcal{J}_{ft}$ is the set of products the establishment makes in t . X_{ft} without an index for product represents the establishment total of input X . In other words, there are no joint or public use of inputs across product lines, and each input can be completely allocated across the product set. This assumption is necessary for identification, because input use is observable only at the establishment-year level in the data, and not by product.

Separability of the production function across goods provides a natural framework for understanding product-specific productivities and allows the adaptation of techniques from firm-level estimation in an analogous way. Thus, the production of a good i by a plant is determined by the inputs, technology, and productivity for product i only. However, interdependencies may still exist across a plant's products through the plant's input allocation decision, taking as given its capital stock at the beginning of the period, and through endogenous effects on productivities. I impose that the parameters of the production functions, the factor intensities β_X^i associated with each input X and output product i , may vary by product but not by establishment. This can be interpreted as there being a common technology for the production of each good, which any establishment uses to produce it. The factor intensities for a product do not vary with the establishment's size or product scope. This restricts interactions across an establishment's products through changes in factor intensity.

However, along with papers in the literature such as De Loecker et al (2016) and Orr (2017), I use this assumption for identification when estimating the production function using single-product establishments. The assumptions above are generally implicitly imposed in models without multiple products. Here, they allow for easy comparison to methods in the existing literature, give tractability to my estimation procedure, and allow me to focus on one channel for heterogeneity across products — through total factor productivity.

Within the economy, there are N different goods that establishments may produce. Each establishment chooses a set $\mathcal{J}_{ft}$ of products to produce in time t after drawing a vector of its potential productivities for every possible set $\vec{\phi}_{ft}(\mathcal{J}_{ft})$. This means that for any combination of the N possible products, the establishment sees what its productivities for each good would be if it were to choose a given set. In particular, I allow for an establishment's productivity at producing a good i to differ based on the entire set it produces. It is possible for $\phi_{ift}(\mathcal{J}) \neq \phi_{ift}(\mathcal{J}')$ for different product mixes $\mathcal{J}$ and $\mathcal{J}'$, where both sets include product i . The model accommodates endogenous spillover effects through Hicks neutral productivity that occur in coproduction.

It is important to note that while productivities may depend on what else is in the establishment's product mix, for tractability, I do not allow productivities to depend on the quantities of other goods produced or inputs used. For example, productivity at producing i can potentially depend on *whether* the establishment is concurrently producing another product j , but not on *how much* labor is put towards the production of j . This is important in order for the establishment's cost minimization problem to be solvable given productivities and desired output levels. Otherwise, the establishment's problem becomes a more complicated fixed point and would present a problem for identifying optimal input usage across products when productivities are also unobserved.

Very few assumptions on the distributions of productivity draws are needed, except for a few on the evolution of productivity conditional on the establishment remaining single product, which is used for the production function estimation step and discussed in greater detail in Section 1.4.1. By remaining agnostic on the form of the productivity distribution, I allow for a potentially rich set of correlations across products to influence the establishment's choice of product mix. In Chapter 2, I examine these correlations and potential interdependencies in productivity.

The establishment's problem is as follows: The establishment first observes its own productivity draws for each possible product set. It chooses its product mix for time t as well as its investment in capital, which comes online at time $t+1$. The establishment pays a fixed cost associated with its product mix. Next, it makes production decisions and sets its prices. During these decision steps, the establishment's information set includes economic aggregates and its own capital stock, productivity draws, and factor prices. The establishment's value function is:

$$V_t(\chi_{ft}) = \max_{\mathcal{J}_{ft}, I_{ft}} \Pi_f(\mathcal{J}_{ft}, K_{ft}) - F_{ft}(\mathcal{J}_{ft}) - d(K_{ft}, I_{ft}) \quad (1.3)$$

$$+ \rho \mathbb{E}[V_{t+1}(\chi_{ft+1}) | \chi_{ft}] \quad (1.4)$$

$$\text{s.t. } K_{f,t+1} = l(K_{ft}, I_{ft}) \quad (1.5)$$

where $\Pi_f(\cdot)$ is the optimized profit and $F_{ft}(\cdot)$ is the fixed cost of a product set. I_t is the investment in capital, $d(\cdot)$ is the investment and adjustment cost of capital, and $l(\cdot)$ is the law of motion of capital that governs how investment grows the capital stock and depreciation over time. I leave the dynamics and costs of capital adjustment as general functions because I do not need to know them in order to estimate productivity. All that matters is that capital is a dynamic input, and the current stock is taken as given at the beginning of each period. The state variables include the price vectors of other establishments, f 's own productivity draws, factor costs, and the current

capital stock:

$$\chi_{ft} = (P_{-ft}, \vec{\phi}_{ift}, \vec{W}_{ft}, K_{ft})$$

There are no costs of adjusting the product mix. Instead, the choice of product mix is a static decision that the establishment makes each period.

The profit function for a given product mix and capital stock is the maximized total variable profit that takes into account variable costs of production, but not fixed costs of selecting the product mix. The establishment solves the profit function as the set of prices for each product in its mix $\mathcal{J}_{ft}$ that maximizes total revenue net variable costs:

$$\Pi_f(\mathcal{J}_{ft}, K_{ft}) = \max_{\{P_{ift}\}_{i \in \mathcal{J}_{ft}}} \sum_{i \in \mathcal{J}_{ft}} P_{ift} Y_{ift}(P_{ift}) - C(\vec{Y}_{ft}, K_{ft}, \vec{W}_{ft}, \vec{\phi}_{ft}, \mathcal{J}_{ft}) \quad (1.6)$$

Given capital stock K_{ft} and desired levels of output $Y_{ift} \forall i \in \mathcal{J}_{ft}$ associated with its optimal pricing choices, the establishment chooses input allocations to minimize total static costs of production taking input costs as given. The cost function is:

$$C(\vec{Y}_{ft}, K_{ft}, \vec{W}_{ft}, \vec{\phi}_{ft}, \mathcal{J}_{ft}) = \min_{\vec{L}_{ft}, \vec{M}_{ft}, \vec{K}_{ft}} \sum_{i \in \mathcal{J}_{ft}} (W_{ft}^L L_{ift} + W_{ft}^M M_{ift}) \quad (1.7)$$

$$\text{s.t. } Y_{ift} \leq L_{ift}^{\beta_L^i} M_{ift}^{\beta_M^i} K_{ift}^{\beta_K^i} \phi_{ift} \quad (1.8)$$

$$\sum_{i \in \mathcal{J}_{ft}} K_{ift} \leq K_{ft} \quad (1.9)$$

Note that factor costs for labor and materials are indexed by establishment and year. Factor costs may vary across establishments and time, but not products. Also note that if all inputs were static, the input allocation problem could be solved product by product because establishments could freely choose the optimal amount of input for each product each period. However, with dynamic capital accumulation, the total quantity of capital that a establishment has to allocate across products is constrained

by its current capital stock, determined last period. This makes cost minimization non-separable across products because the model implies optimal ratios of input use for given input costs, and total capital cannot be adjusted for the current period.

Because the product mix decision is static and investments do not become available for use in production until the next period, there is no interaction between an establishment's investment decision and product mix choice in time t , so the two can be considered independently. Isolating the static component of the problem, the establishment chooses the product set that maximizes total profits minus fixed costs for the period:

$$\max_{\mathcal{J}_{ft}} \Pi_f(\mathcal{J}_{ft}, K_{ft}) - F_{ft}(\mathcal{J}_{ft}) \quad (1.10)$$

The establishment's problem can be solved backwards by solving for the cost function first for a given vector of desired outputs and the current capital, then the optimal pricing rule given a product set, and finally solving for the optimal product set. For the purposes of estimating productivities, it is not necessary to solve the entire establishment's problem. Likewise, the fixed cost function and capital adjustment and investment functions do not need to be known, and are left flexible. Identification primarily comes from first order conditions of the establishment's cost minimization problem. These first order conditions are:

$$\begin{aligned} \frac{\partial C}{\partial L_{ift}} &= W_{ft}^L = \frac{\lambda_{ift} \beta_L^i Y_{ift}}{L_{ift}} \\ \frac{\partial C}{\partial M_{ift}} &= W_{ft}^M = \frac{\lambda_{ift} \beta_M^i Y_{ift}}{M_{ift}} \\ \frac{\partial C}{\partial K_{ift}} &= \kappa_{ft} = \frac{\lambda_{ift} \beta_K^i Y_{ift}}{K_{ift}} \end{aligned}$$

The Lagrange multiplier on the capital constraint, κ_{ft} , is the shadow price of capital. The Lagrange multiplier on the production function constraint is λ_{ift} and, by the Envelope Theorem, the marginal cost of producing product i . Under a Cobb Douglas

production function, the first order conditions of cost minimization have this very tractable form. These first order conditions describe the optimal input choices for a given vector of desired outputs. By rearranging, I can express the input allocation shares:

$$S_{ift}^X = \frac{X_{ift}}{X_{ft}} = \frac{\lambda_{ift}\beta_X^i Y_{ift}}{\sum_{j \in \mathcal{J}_{ft}} \lambda_{jft}\beta_X^j Y_{jft}} \quad (1.11)$$

for each input X . This is the share of input X that establishment f optimally allocates to the production of product i . This derivation relies on the assumption that inputs can be completely allocated across products and that input costs may vary by establishment and over time, but not across products within an establishment. Solving the establishment's input allocation shares across products allows me to infer product-level input use, from which productivities can be backed out given known factor intensities β_X^i . This equation shows that input allocation shares are a function of marginal costs, factor intensities, and output quantities.

Leaving the demand system general for the moment, the optimal pricing rule derived from the solution to the variable profit maximization problem reveals additional interdependencies across products. First order conditions for each product i are:

$$\frac{\partial \Pi_{ft}}{\partial P_{ift}} = Y_{ift} + \sum_{j \in \mathcal{J}_{ft}} \frac{\partial Y_{jft}}{\partial P_{ift}} \left(P_{jft} - \frac{\partial C(\vec{Y}_{ft}, K_{ft}, \vec{W}_{ft}, \vec{\phi}_{ft}, \mathcal{J}_{ft})}{\partial Y_{jft}} \right) = 0$$

The price an establishment sets for product i will depend on the cross product demand elasticities with the other goods it produces. I can re-express the pricing rule as

$$P_{ift} = \sum_{j \in \mathcal{J}_{ft}} -\frac{\partial Y_{jft}/Y_{jft}}{\partial P_{ift}/P_{ift}} (P_{jft} - \lambda_{ft}^j)$$

For a general demand system in which multi-product establishments compete on prices, the optimal price is a function of the cross product demand elasticities and margins of the establishment's full set of products. Multi-product establishments are

subject to cannibalization effects across product lines and potentially rich patterns of substitution across products.

In the ASI and other establishment-product level datasets, there is information on prices and physical quantities but marginal costs are a complex object that cannot be observed in the data. However, if demand elasticities are known or estimated, then the pricing rule can be inverted to get marginal costs:

$$\vec{\lambda}_{ft} = \Delta_{ft}^{-1} \vec{Y}_{ft} + \vec{P}_{ft}$$

where Δ_{ft} is the N_{ft} by N_{ft} (number of products in $\mathcal{J}_{ft}$) matrix of demand derivatives $\frac{\partial Y_{ift}}{\partial P_{jft}}$.

1.3.2 Introducing a Model of Demand

I now impose additional structure on demand in order to make the estimation of marginal costs tractable across a broad range of industries. I adapt a demand model similar to that used by Hottman, Redding, and Weinstein (2016) [26] that features CES with non-atomistic establishments nested in a Cobb Douglas function over broad product groups. The Cobb-Douglas upper tier implies that there are no interdependencies in the establishment's pricing decision across product groups. Within a group of products, allowing non-atomistic establishments implies that establishments may be large and can affect the price index, leading to variable markups with cannibalization across product lines. In other words, pricing for an establishment's product i can affect quantity demanded for its product j when i and j are from the same product group, and thus pricing must be decided jointly for the product set rather than independently for each product. These are strategic trade-offs that distinguish multi-product establishments from single product ones.

For each of G broad product groups, consumers have Cobb Douglas preferences over the aggregate consumption index for each g , implying constant expenditure shares. Then within each product group, preferences are CES over varieties, which are establishment-product combinations. The demand function is:

$$U_t = \prod_{g=1}^G C_{gt}^{\alpha_{gt}}, \quad \sum_{g=1}^G \alpha_{gt} = 1$$

$$C_{gt} = \left[\sum_{(i,f) \in \Gamma_g} C_{ift}^{\frac{\sigma_g - 1}{\sigma_g}} \right]^{\frac{\sigma_g}{\sigma_g - 1}}, \quad \sigma_g > 1$$

where g is a 1-digit ASICC and Γ_g is the set of varieties in group g . This demand structure allows for variable markups and cannibalization effects within product groups.

Solving the model for the optimal pricing strategy reveals that the markup is the same for all of an establishment's products within a product group:

$$P_{ift} = \frac{\sigma_g - S_{gft}(\sigma_g - 1)}{(\sigma_g - 1)(1 - S_{gft})} \lambda_{ift}$$

where S_{gft} is f 's share of revenue in group g . The markup is variable, as it depends on the establishment's market share in a product group, and may differ across products from different groups produced by the same establishment, and for different establishments producing the same product.

Obviously the choice of how to divide these nests or groups is important. Here I will use broad industry groups defined by the 1 digit ASICC codes. Examples of these product groups are agricultural products, textiles, and manufactured articles. The advantage of this particular demand model is that it allows for strategic behavior in pricing and other interactions on the demand side between an establishment's products while still being very tractable because it requires estimating only

G elasticities of substitution. However, elasticities of substitution are assumed constant within groups, and interaction between groups is limited. In situations where researchers possess sufficiently rich data and in which more flexible substitution patterns are desirable, a more general, estimatable demand system such as BLP from the IO literature could be substituted here.

Finally, the choice of the product set $\mathcal{J}_{ft}$ is a combinatorial discrete choice problem that finds the subset of all possible products to maximize total profits. In considering the addition of a product to its set, the establishment trades off several effects. The new product is an additional line of profit. Also, there are potentially complementarity effects with its other products through Hicks neutral productivity. On the other hand, there is an added fixed cost to producing an extra product, and the new product could cannibalize the sales of existing products or cause diseconomies of scope. The optimal product mix choice is an interesting and computationally intensive problem, but unnecessary to solve for the purposes of estimating product-specific productivities.

1.4 Methodology for Estimating Productivities

The estimation of product-specific TFP within each establishment relies on the optimal share of each input that is allocated to a product, which can be expressed as:

$$S_{ift}^X = \frac{W_{ft}^X X_{ift}}{W_{ft}^X X_{ft}} = \frac{\lambda_{ift} \beta_X^i Y_{ift}}{\sum_{j \in \mathcal{J}_{ft}} \lambda_{jft} \beta_X^j Y_{jft}}$$

Input usage is observable at the establishment level and not the product level, which presents the primary challenge for estimating productivity for each good. However, using equation (1.11), I can back out the quantity of each input that the establishment would optimally put towards the production of each good. In the data, I can

observe variation in the cost of an input X across establishments and over time, but without product-level information about input use, there is no data about input costs varying by output product. I assume the cost of an input is the same across all of an establishment's products, which allows the W_{ft} term to cancel out in the input share equation. Then under Cobb Douglas production, the input quantity shares are a function of the marginal costs, factor intensities, and output quantities of each product in an establishment's mix. Output quantities Y_{ift} are observed in the data, so I just need to estimate factor intensities β_{ift}^X and marginal costs λ_{ift} to solve for input allocations. In the following subsections, I explain how to estimate these objects. Once the quantities of inputs used in the production of each good are inferred, it is straightforward to compute productivities as the residuals of the production function.

Earlier studies have used revenue shares to allocate observed establishment-level inputs across products. From the input share equation, I can show that my procedure would generate the same allocation under two additional assumptions on the model: (1) common factor intensities and (2) constant markups across all of the establishment's products. If output elasticities of each input and markups are the same for all products in the establishment's mix: $\beta_X^i = \beta_X^j$ and $\mu_{ift} = \mu_{jft} \forall i, j \in \mathcal{J}_{ft}$ where $\mu_{ift} = \frac{P_{ift}}{\lambda_{ift}}$, then the input share equation reduces to revenue shares

$$S_{ift}^X = \frac{W_{ft}^X X_{ift}}{W_{ft}^X X_{ft}} = \frac{P_{ift} Y_{ift}}{\sum_{j \in \mathcal{J}_{ft}} P_{jft} Y_{jft}} \quad (1.12)$$

Common factor intensities may be plausible for establishments that produce a narrow range of very similar products, but within the ASI data, I observe plants producing products across industries and sectors, for which it would be restrictive to assume common factor intensities. Table 1.1 showed that 19% of establishment observations produced in more than one industry, and these establishments have a significant impact on the economy, accounting for 50% of all revenue. Assuming common markups

across all products additionally restricts substitution patterns and demand interactions across products. I solve for input allocations without needing to impose common factor intensities or markups. I compare my results to productivities estimated using revenue share allocation in Section 1.5.4.

To proceed, the first steps in getting productivities are to estimate the factor intensities and marginal costs in order to infer the input allocations for each product.

1.4.1 Production Function Estimation

I estimate production function parameters using a proxy estimator following Olley-Pakes (1996) and Akerberg, Caves, and Frazer (2015). I limit the production function estimation to use only single-product establishments in order to use the observed input quantities. In single product establishments, all observed plant level inputs must also be the product-level inputs, so I avoid dealing with unobserved input allocation present in multi-product establishments. This assumes that the output elasticities are the same for any establishment producing a product, whether it is single product or multi-product.³

I pool single-product establishments within each industry (2 digit ASIC codes) to estimate the output elasticities $\{\beta_L^i, \beta_M^i, \beta_K^i\}$. In theory, I could estimate separate production functions for each product, but the dataset I have does not contain enough suitable observations of each product to accurately estimate the parameters at such a fine level. Pooling imposes that the output elasticities for each input are the same for products within the same industry, which is not unreasonable here because the

³The strategy of using single-product establishments to estimate production functions shared by multi-product establishments has been used by other authors, including De Loecker, Goldberg, Khandelwal, and Pavcnik in their 2016 paper on markups. They also use firm optimality conditions to infer input allocations across products, but the major difference between their work and mine is that they impose a constant firm-specific productivity that is needed for identification in their estimation method. In contrast, I estimate unobserved heterogeneous establishment-product productivities and achieve identification through different structural assumptions.

2 digit industries are fairly detailed. Ultimately, this corresponds to 36 industries where I have enough data to estimate factor intensities.

First, I consider the production function in logs:

$$y_{ift} = \beta_L^i l_{ift} + \beta_M^i m_{ift} + \beta_K^i k_{ift} + \log \phi_{ift} + \epsilon_{ift}$$

where lowercase letters indicate the logs of variables and ϵ_{ift} is measurement error or an unanticipated shock to production. The two major issues in estimating production functions with unobserved productivity, simultaneity and selection bias, have been discussed extensively in the literature (see Olley-Pakes [36], Levinsohn-Petrin [30], and Akerberg, Caves, Frazer [1]), and I will adapt the existing solutions here. I deal with the simultaneity problem using a static input demand control function for materials which includes the unobserved productivity and any other variables that determine demand for materials:

$$m_{ift} = m_t(\phi_{ift}, l_{ift}, k_{ift}, ms_{ift}, p_{ift})$$

Inverting this control function gives a proxy for unobserved productivity:

$$\log \phi_{ift} = m_t^{-1}(l_{ift}, m_{ift}, k_{ift}, ms_{ift}, p_{ift})$$

The market share ms_{ift} and output price p_{ift} enter the function because they influence profits and therefore the desired quantity of output. Note that the establishment's productivity draws for other goods have no effect on inputs put towards product i because only single product establishments are used in this step.

Following the literature, I form moments using the innovation term of log productivity each period. The law of motion for the productivity of a single-product establishment that continues to be single-product in the next period is a function of

previous log productivity with an additive innovation term:

$$\log\phi_{ift} = g(\log\phi_{ift-1}) + \xi_{ift}$$

Conditional on also producing only i next period, the establishment's productivity in time $t + 1$ is a function of last period's productivity plus an unanticipated innovation term ξ_{ift} . The use of this assumption on the evolution of productivity introduces some restrictiveness to the establishment's productivity draws. However, the assumption of this form is only imposed on the establishments producing only good i in both periods $t - 1$ and t , and so does not limit the correlations with products that are not chosen.

Identification comes from the moment

$$\mathbb{E}[\xi_{ift}Z_{ift}] = 0$$

where Z_{ift} includes current capital and lagged values of materials, labor, price, and market share. The production function parameters β are estimated using GMM.

An additional concern in estimating production functions is selection bias from the use of only single product establishments. Establishments consider their potential productivities when choosing to be single product, multi-product, or to exit the market, made clear in the establishment's problem described by equations (1.3) to (1.10). Thus within the single product dataset, estimates of production function parameters may be biased if the establishment's decision to produce the single product depends on unobserved productivity or input use. In particular, it is plausible under this framework that the choice to be single product depends on the current capital stock. Traditionally, the literature has dealt with selection bias through the use of an unbalanced panel and by estimating the probability of staying in the market as a control (see Olley Pakes [36]). In their work on markups in multi-product firms, De Loecker,

Goldberg, Khandelwal, and Pavcnik (2016) [14] addressed selection bias from the use of single product firm data with a selection control on the probability of being single product as a function of unobserved firm-level productivity. My estimation differs from theirs in not restricting productivity to be constant within the establishment, and thus the probability an establishment is single product in my model depends on the entire set of potential productivities for all goods. Once productivities are allowed to be product-specific, the choice to be single product depends on a multi-dimensional object of unobservables, which can't be controlled for using single product data without imposing restrictive assumptions on the productivity draws. To maintain the flexibility of the model regarding the productivity draws, I rely on the unbalanced panel to reduce selection bias in the estimates.

The estimated output elasticities are displayed in Appendix B, Table B.1. The average returns to scale across all product observations is 1.073.

1.4.2 Marginal Cost Estimation

To estimate marginal costs λ_{ift} , I impose structure in the form of a nested CES demand with non-atomistic establishment, detailed in the previous section. This demand function implies markups of the form

$$\frac{P_{ift}}{\lambda_{ift}} = \mu_{gft} = \frac{\sigma_g - S_{gft}(\sigma_g - 1)}{(\sigma_g - 1)(1 - S_{gft})} \quad (1.13)$$

where S_{gft} is establishment f 's share of revenue within product group g and observable in the data. Markups are constant within a product group for a given establishment. Thus, establishments may charge different markups across their product groups, and different establishments producing the same product may also charge different markups. I define the groups as sectors identified by the 1 digit ASIC codes.

Estimating marginal costs in this model comes down to estimating the elasticities of substitution for each product group σ_g .

I make use of the cost minimization first-order conditions from the supply side here as well to estimate the elasticities of substitution. Using $\lambda_{ift} = \frac{P_{ift}}{\mu_{g(i)ft}}$ in the establishment's cost minimization with respect to materials, I can factor out the markup for each group and express total establishment expenditure as:

$$W_{ft}^M M_{ft} = \sum_{g=1}^G \frac{(\sigma_g - 1)(1 - S_{gft})}{\sigma_g - S_{gft}(\sigma_g - 1)} \sum_{i,f \in \Gamma_g} \beta_M^i P_{ift} Y_{ift} \quad (1.14)$$

I use expenditure on materials because it has more complete data in the ASI, but the method would also work substituting labor expenditure and β_L^i . For each establishment-year, the expenditure on materials can be expressed as a function of markups, materials intensity, and sales for all its products. We have both materials input intensity and sales from production function estimation and the data, respectively, and the plant-product group specific markups are functions of the elasticities of substitution and revenue shares. The demand function I use imposes restrictions on the substitution patterns, and in particular limits markups to be common within an establishment-group-year, but it provides a great deal of tractability. It allows me to get all marginal costs for all establishment-product-years by estimating only $G = 9$ parameters, the σ_g s.

I estimate σ_g for each group g by nonlinear least squares using equation (1.14). The resulting estimates are reported in Appendix B, Table B.2. Then it's straightforward to compute marginal costs for each product-establishment-year using observed output prices and revenue shares:

$$\hat{\lambda}_{ift} = \frac{P_{ift}(\hat{\sigma}_g - 1)(1 - S_{gft})}{\hat{\sigma}_g - S_{gft}(\hat{\sigma}_g - 1)}$$

1.4.3 Computing Productivities

Finally, I compute the establishment-product level productivities using implied input allocations:

$$\hat{X}_{ift} = \frac{\hat{\lambda}_{ift} \hat{\beta}_X^i Y_{ift}}{\sum_{j \in \mathcal{J}_{ft}} \hat{\lambda}_{jft} \hat{\beta}_X^j Y_{jft}} X_{ft}, \quad X \in \{L, M, K\} \quad (1.15)$$

$$\hat{\phi}_{ift} = \frac{Y_{ift}}{\hat{L}_{ift}^{\hat{\beta}_L^i} \hat{M}_{ift}^{\hat{\beta}_M^i} \hat{K}_{ift}^{\hat{\beta}_K^i}} \quad (1.16)$$

Recall that the input allocation share equation is derived by summing the first order conditions of the cost minimization problem across products in the establishment's mix:

$$X_{ft} = \frac{\sum_{j \in \mathcal{J}_{ft}} \lambda_{jft} \beta_X^j Y_{jft}}{W_{ft}^X}$$

Thus, there is another expression for inferred product-specific inputs that could be used in estimating productivities:

$$\hat{X}_{ift} = \frac{\hat{\lambda}_{ift} \hat{\beta}_X^i Y_{ift}}{W_{ft}^X}, \quad X \in \{L, M, K\} \quad (1.17)$$

Under the model's assumptions, these two expressions should be equivalent. However, if the assumptions of a Cobb Douglas production function and no public inputs aren't completely true, or if there is simply estimation error in the output elasticities or marginal costs, there will be some difference in the productivity estimates using either the input share equation (1.15) or data on input costs with equation (1.17). I discuss the use of these two expressions as an over-identification check of the model's assumptions in Appendix B. There is a correlation of 0.9589 between productivities computed using the two methods, where the productivity is measured in logs and relative to the 5-digit product mean. For the results I will present in the main body of this paper, I use equation (1.15), which guarantees product-level inputs sum to the

observed establishment-level totals and provides a better comparison to alternative methods, detailed in Section 1.5.4.

Standard errors are not shown for the intermediate parameter estimates because each step of the procedure uses estimates from previous steps. For any statistics computed from the estimated productivities, standard errors need to be calculated by block bootstrap, where draws are of an establishment time series, over the entire estimation process.

1.5 Empirical Findings

Before getting into the analysis using the estimated productivities, I discuss a few issues related to quantity based productivities. Leveraging the availability of data on physical quantities manufactured for each product, I have estimated quantity based productivities, or TFPQ. This has an advantage over revenue based productivity of not being subject to the usual output price bias, which arises when revenues are deflated by an industry price index. Quantity based productivity also allows me to separate the physical productivity of an establishment from its efficiency at generating revenue, which may come from demand side factors like market power and product differentiation. However, there arises situations in which caution is needed when comparing quantity-based productivities across products or establishments.

First, the estimated productivities for different products are not directly comparable. Quantity based productivities are in terms of the unit of measurement, which often differs across products. In the data, we see wheat recorded in tons but cars recorded in numbers. This is an obstacle to comparisons of productivity across products, although for a given product, quantities in the ASI use a standard unit of measurement. In addition to differences in units, it is meaningless to directly compare the productivities of different products. The most clearcut comparisons of

productivity would be within a product, running any analysis product by product. However, with the large number of products in the data, doing every analysis product by product is not tractable and will make it difficult to discern broad trends and relationships across products. Instead, my approach for much of the following results will be to standardize productivity measures so that I can pool products and look at larger swaths of data at a time. How best to standardize depends on the application, but the main way I conduct cross-product analysis is to use the log demeaned productivity, which subtracts the log of the mean productivity within each 5 digit product code. Thus, results compare how productive an establishment is relative to the mean for the product.

Even when comparing the productivities of establishments producing the same product, it is important to note that these measurements are of physical productivity – in other words, how well a plant can turn inputs into numbers of units of output. Establishments that produce fewer, higher quality units will appear less productive by this measure. This is more problematic in industries with high vertical differentiation than those in which products are either homogeneous or horizontally differentiated. In applications where quality is an important margin of firm or establishment decisions, it is a good idea to complement the physical quantity productivity with a measure of quality using output prices. Figure 1.1, which plots the log demeaned productivity against log demeaned price, shows an inverse relationship between measured quantity-based productivity and output price in all sectors. For a full discussion on quantity and revenue based productivity measures, see the paper by Foster, Haltiwanger, and Syverson (2008) [21]. Concerns about unobserved differences in quality or subjective appeal could be alleviated by focusing on industries where physical quantities of products are homogeneous.

I proceed now to discuss findings using the estimated productivities. In Section 1.5.1, I present the standardized productivity estimates and broad trends for India.




Figure 1.1: Productivity and prices

In Section 1.5.2, I examine the relative importance of within and between establishment heterogeneity in productivity. In 1.5.3, I explore my results in the context of core competencies. Finally in Section 1.5.4, I check the robustness of the estimation procedure to alternative assumptions on markups and input allocation.

1.5.1 Productivity in India

I estimate productivities for 3655 unique 5 digit products within 73,972 Indian establishments over the years 2000 to 2010. By using the demeaned log productivity, I center the productivity distribution of each good at zero, but there remains differences in the dispersion of productivity between products and even product groups (Figure 1.2).

Next, I compared the productivity of single product and multi-product establishments. To control for differences in productivity dispersion between products and because the single and multi-product establishments may produce different goods, I




Figure 1.2: Distributions of log productivity demeaned, by 1 digit ASIC

use the Z score of estimated productivity here, which divides the demeaned log productivity by the standard deviation within products. The literature has consistently found that multi-product firms tend to be more productive and have higher revenue. Within the ASI, I confirm that at the product level, goods produced within multi-product establishments tend to be more productive (Figure 1.3) and have higher prices (Figure 1.4) than their counterparts in single product plants. For both productivity and price, the distribution for multi-product establishments appears shifted to the right. One-sided tests of the differences in means between single-product and multi-product establishments finds that multi-product establishments have higher standardized log productivity, significant at the 0.1% level, and higher standardized log price, significant at the 1% level.




Figure 1.3: Single product and multi-product productivities




Figure 1.4: Single product and multi-product prices

1.5.2 Within and Between Establishment Heterogeneity

In estimating productivity at the product level, I am able to measure heterogeneity across products within establishments. While there have been many contributions in the literature to understanding heterogeneity across firms, there is less work on heterogeneity within firms. To address this topic, I decompose the heterogeneity in productivity within and across establishments. Similar to an exercise in Bernard, Eaton, Jensen, and Kortum (2003) [9] for heterogeneity across firms, I see how much of the variance in log productivity remains after purging product and establishment fixed effects.

Productivity measure	Variability
Unconditional	5.381
Within establishment	2.242
Within product	1.746
Within product, within establishment	0.870

Notes: Variability is the standard deviation of log productivity. The “within” measures subtract the mean of the log productivity in the given category. The “within product, within establishment” measure subtracts the establishment mean of the “within product” measure.

Table 1.2: Variability of log productivity with establishments and products

Comparing the unconditional variability, computed simply as the standard deviation of log productivity across all observations, to the within product variability in Table 1.2, there is a large decrease from 5.381 to 1.746. The unconditional measure faces the problems with comparing across products discussed at the start of this section, which suggests purging product fixed effects is important here to eliminate differences in units across product types. The “within product” removes that component and leaves the variability across establishments. The “within establishment” measure of 2.242 takes out the establishment fixed effect instead. Variability in scales of measured productivity is present here, so within establishment variability includes both heterogeneity in productivity within establishments as well composition effects of the different chosen product mixes. More of the total variability is attributable to differences in productivity values across products than differences across establishments.

The more telling comparison for understanding within and between establishment heterogeneity is the “within product” measure versus the “within product, within establishment” measure. The second purges the “within product” log productivity of its mean within each establishment. This takes out the variability between different products and between establishments so what’s left is the variability in productivities

relative to the product mean within an establishment. The “within product, within establishment” measures variability in how good an establishment is at producing the different products in its mix relative to the mean of each product. Compared to the within product measure of 1.746, nearly half of the variability is remaining after accounting for the common within establishment component, at 0.870. Differences in how productive an establishment is across the different products in its mix does appear to be a significant source of heterogeneity. This is consistent with non-random selection of products by establishments.

1.5.3 The Core Competency Ladder

Estimation of product-specific productivities within establishments has the potential to shed light on the product mix choices of firms. Evidence from empirical work on multi-product firms, particularly Bernard, Redding, and Schott 2010 [10], show that a firm’s product mix is not random. Instead, there seem to be systematic patterns to the products co-produced within a firm. Eckel and Neary (2010) [18] describe the idea of core competencies—that firms have a core competence product at which they are most efficient at producing. Additional products can be made with costly adjustment to the production process. Marginal costs increase for products further away from the core competency. Ultimately, firms produce a range of products moving away from its core. This idea of core competencies is consistent with my theoretical framework due to the potential linkages allowed across productivity draws, even though I have not imposed any assumptions of an exogenous core competency or how an establishment’s other products relate to it. Establishments that are productive at one certain good may tend to draw higher productivities for other similar products.

Defining an establishment’s most productive product, relative to each product’s mean, as its core competency, I examine the relative productivity of the other goods concurrently produced. Of interest is the steepness of the competency ladder, which

within the core competency framework, is a measure of how much adaptation of the production process is needed for a firm to add an additional product. I standardize log productivity values by demeaning log productivity and dividing by the standard deviation within each 5-digit product, which removes differences across products in measured units and dispersion. Each data point is the distance in standard deviations relative to the product mean of the observation's log productivity. On average, the difference in standardized log productivity between an establishment's core competency and its next most productive product is 0.748. This implies a productivity ratio of 2.11 when not in logs, meaning that on average, an establishment's top product is more than twice as productive as its next best. For context, Syverson (2004) [41] finds an average ratio of 1.92 between the 90th percentile and 10th percentile of TFP across plants in the US, and Hsieh and Klenow (2009) [27] find a ratio of over 5 for China. Here too it is important to note that these product-specific productivities are not directly comparable to their plant-level measures of revenue based productivity. However, it is still clear that the productivity differences across products produced by the same establishment are large.

Figure 1.5 plots the competency ladder: the mean standardized productivities as you move down the productivity ladder for establishments grouped by their product scope. Establishments that produce more products are on average more productive at their core competency. The competency ladder appears nonlinear and flatter for establishments producing more products. The flattening of the competency ladder may be explained by a fixed cost of adjustment or loss of efficiency in moving away from the core competency, potential sources of economies of scope. Unsurprisingly, the graph also shows that establishment's producing more products tend to have higher productivity in their best product. It also appears that establishments with greater scope are able to include relatively lower productivity goods in their mix.




Figure 1.5: Core competency ladder of standardized log productivity

1.5.4 Comparison to Alternative Estimation Methods

To check the robustness of some of my modeling choices, I estimate productivities under alternative assumptions and compare the results to the baseline estimation procedure described so far. I perform robustness checks against two alternatives: first, I assume constant markups within establishments, and second, I use revenue shares to allocate inputs across products.

Under the constant markup assumption where $\mu_{ift} = \frac{P_{ift}}{\lambda_{ift}} = \mu_{ft} \forall i \in \mathcal{J}_{ft}$, the input allocation share equation reduces to:

$$S_{ift}^X = \frac{W_{ft}^X X_{ift}}{W_{ft}^X X_{ft}} = \frac{\beta_X^i P_{ift} Y_{ift}}{\sum_{j \in \mathcal{J}_{ft}} \beta_X^j P_{jft} Y_{jft}}$$

With common markups, the share of an input allocated to a product is equal to the product's share of revenue weighted by the factor intensity. Under this assumption, it is not necessary to estimate marginal costs or demand parameters. Compared

to the nested CES demand with variable markups used for my baseline estimation, assuming common markups only produces different estimated productivities within establishments that operate in more than one product group. Figure 1.6 plots productivity estimated under common markups against the baseline estimates for only the multi-group establishments. The correlation between these two measures is high, at 0.963 even among only multi-group establishments (see Table 1.3).

Alternative assumption	Unconditional	Multi-industry	Multi-group
Constant markups	0.996	0.987	0.963
Revenue share allocation	0.987	0.991	0.977

Notes: Correlations are computed between productivities using the baseline estimation method and under alternative assumptions. The unconditional correlations include all observations. The remaining columns limit the sample to products produced by multi-industry (2 digit ASICC) and multi-group (1 digit ASICC) establishments.

Table 1.3: Correlations with productivities estimated under alternative assumptions




Figure 1.6: Productivity assuming constant markups




Figure 1.7: Productivity using revenue shares

As earlier works have often used revenue shares to allocate a firm or plant's observed inputs across its product mix, I also compare the baseline productivities to ones using revenue share input allocations. The use of revenue shares to allocate inputs is tractable because it doesn't require estimating any parameters first, such as the output elasticities or demand substitution elasticities that my baseline method requires. The downside, however, is that if the factor intensities or markups vary across products within the establishment's mix, revenue shares may bias the imputed inputs. The revenue share allocation method assigns the same share of all inputs to a given product, which fails to take into account different relative factor intensities within an establishment's product set. To keep the measure comparable, I use the same estimates of the factor intensities as in the baseline so that the only difference in productivity estimates arises from the different input allocation rule. The results are highly correlated with the baseline estimates (Table 1.3). Figure 1.7 also shows the relationship between the two estimates closely follows the 45 degree line, except in a few cases when revenue shares significantly over-estimate productivity. This suggests

that in cases where there is insufficient data for estimating production function and demand parameters, it may be a reasonable approximation to use revenue shares to allocate inputs across products.

1.6 Conclusions

In light of the growing multi-product literature, an important challenge in estimating productivity is that establishments that supply multiple products may produce each of them with different productivities. Firm and establishment-level measures of productivity ignore potential heterogeneity across the products. In this paper, I developed a flexible methodology for estimating within-establishments product-specific productivities. The main obstacle to adapting standard techniques to multiple products is that even in cases where product-level data is available, inputs are recorded at the establishment level. I address this by deriving a theoretically consistent expression for the shares of each input allocated to the products within the establishment's mix, using first-order conditions from the establishment's cost minimization and input allocation problem. The key insight here is that for the widely used Cobb Douglas production function, the share of an input optimally allocated to a product is proportional to a function of the product's factor intensity, marginal cost, and output quantity. I have tried to impose only the assumptions necessary for identification and tractability in order to allow for interactions between products in production, productivity draws, and demand which give richness to the model of multi-product establishments. These linkages across products may lead to the relationships between productivity, revenue, and product adoption that are observed in the endogenous product mixes of establishments. Compared to existing methods, I am able to estimate the set of unobserved productivities for each product within an establishment without restricting estab-

lishments to use the same technology or constant markups for every product they make.

I estimated the parameters of the production function and markups in steps in order infer input allocations across product and then computed productivity for each product-establishment observation. Next, I used the estimated productivities to revisit some classic issues in the literature on productivity measurement and heterogeneity. I confirmed the premium to multi-product establishments exists for measured productivity controlling for differences across products—multi-product establishments tend to be more productive and charge higher prices than single product ones. I show that controlling for systematic differences across product types, the within establishment, across products dimension of productivity contributes significant heterogeneity. Not only are establishments different, but the products chosen within an establishment are different from each other. The heterogeneity suggests that endogenous product selection is important for explaining observed differences across establishments. In the event of product, industry, or economy-wide shocks, this also implies different patterns of adjustment. The productivity estimates also speak to the steepness of the competency ladder as establishments move away from their most productive product. There is a sizable productivity drop from an establishment’s best product to it’s next best. The productivity ratio to subsequent products eventually flattens out and is also lower in establishments that produce more goods, which suggests non-linear costs of adjustment as establishments expand away from their core product and may be consistent with economies of scope or other linkages on the supply side.

The main findings reiterate that products within an establishment are not randomly chosen and are an important source of heterogeneity in production. These facts should guide how we think about measured aggregate productivity and reallocations within and across establishments, and highlight the need to better understand why

establishments choose to produce the products that they do. There is room for future work exploring linkages across products from both supply and demand sides to better understand the sources of product heterogeneity.

Chapter 2

Coproduction and Complementarities in India

2.1 Introduction

The importance of multi-product firms in terms of their share of total revenue and the way they differ from single product firms on a range of observables has been well established, as evidenced by the growing body of literature on them. The systematic differences between multi-product and single product establishments reveal a need to better understand how establishments choose not only how many products to produce, but also which products. Compared to models of single product establishments, endogenous selection of multiple products introduces the possibility of strategic interactions in the establishment's decisions across products. In particular, this paper aims to explore whether there are interdependencies and spillovers across products that help explain the pattern of coproduction.

The idea of complementarities in production has long existed in the theory (Baumol, Panzar, and Willig 1998 [6]) as motivating the existence of multi-product firms. Economies of scope resulting from complementarities between products are also often

cited as an incentive in mergers. It makes intuitive sense that one establishment¹ could coproduce certain similar, complementary goods more cheaply or productively than different establishments producing them separately. Complementarities may arise when similar processes, materials, knowledge, or experience go into the production of two goods. Allowing for complementarities between products makes the production decisions for individual goods interdependent. Complementarity effects have implications for how establishments react to shocks that affect either the entire market or individual goods. For example, policy directed at a given good would also have an effect on the production decisions for a good with which it has strong complementarity, and therefore the policy might have more far reaching effects than otherwise anticipated. Production complementarities can act through a variety of channels including input usage or costs, the fixed costs of producing or exporting, and productivities, not to mention there could be complementarities on the demand side as well (Bernard et al (2019)[8]).

This paper focuses on the relationship between productivity and coproduction. Productivity differences are a common way to think about heterogeneity across establishments or products in the international trade and macroeconomics literature, and inform the reallocation of resources in response to a shock. I examine the relationship between coproduction and the productivities of the goods within an establishment. Of particular interest is how the productivity of a good varies with the set of other products the establishment is producing and whether that provides evidence of complementarities in production.

¹Much of the literature on multi-product production has been at the firm-level. However, the Indian ASI which I use records data for establishments, with no available firm-level identifier. My analysis is conducted at the establishment level, and I use the terms “firm” and “establishment” interchangeably when discussing complementarities conceptually. Within my data, the majority of establishments report being in single-establishment firms. It is also reasonable to consider complementarities in production acting across products that are produced in the same location—within an establishment.

Using establishment-product data for India between 2000 and 2010, I begin by establishing a set of empirical facts about multi-product establishments. I compare my findings in the Indian data with facts about multi-product firms in the literature and discover patterns similar to what Bernard, Redding, and Schott (2010) found for US firms and Goldberg et al (2010) for Indian firms. However, I see more frequent product switching. I estimate the productivities of goods within establishments to answer how scope relates to productivity and test for one potential channel of interdependency across products — spillovers across product lines which modify the productivity of other goods. This paper’s main contribution is the modeling and estimation of production complementarities through the channel of establishment-product productivities, which is quantifiable across a broad range of industries. This allows the inclusion of establishments that produce a wide range of products across multiple industries, which are a small share of establishments but tend to account for a disproportionately large share of revenue². The theoretical framework I develop is flexible and does not attempt to explicitly model every form of complementarity. It allows me to apply the productivity estimation procedure developed in Chapter 1 to overcome the common issue that inputs are observed only at the establishment level. Finally, I address the challenge of identifying causation in the relationship between the product mix and productivities within an establishment. Because establishments endogenously choose their product mix, an observed correlation between coproduction and productivity could arise from causal complementarity effects or from selection of products with inherently correlated productivities. To identify the complementarity effect of interest, I exploit a policy change in India of the rollback of small scale industry promotion and find some suggestive evidence that producing more products in the same industry increases productivity.

²See table 1.1.

I build on the growing body of literature on multi-product firms and complementarities. The first set of empirical findings in this paper closely follow the work of Bernard, Redding, and Schott (2010), subsequently referred to as BRS, who study multi-product firms and product switching in the US. They find that multi-product firms are larger and more productive, and engage in product churning. They also construct a coproduction matrix, revealing a systematic pattern to the types of products that are frequently produced together. Heterogeneity across products in their model exists as product-specific consumer tastes, while productivity is firm-specific.

Goldberg et al (2010 Restat)[23] conduct a similar analysis on multi-product firms using the Prowess dataset for India from 1989 to 2003. Like Bernard, Redding, and Schott, they find multi-product firms to be larger and more productive controlling for industry, a similar product distribution within firms, and positive correlation between firms' extensive and intensive margins. However, they also found less product churning and rationalization, and significantly less frequent product dropping compared to BRS. In both empirical studies of multi-product firms, changes in product mix drive changes in overall output, which suggests that accounting for the firm's extensive margin is important for understanding reallocation and growth.

Also working with the Indian Prowess data, De Loecker, Goldberg, Khandelwal, and Pavcnik (2016)[14] estimate markups in multi-product firms. They allow for economies of scope in their multi-product model, but restrict them to entering through either input costs that differ across firms, the splitting of fixed costs of production, or through effects on firm-specific productivity. I adopt a similar model that also allows interdependencies in input and fixed costs, but I model establishment-product differences in productivities, which are able to capture product-specific effects of the establishment's product mix on productivity.

Another related area of literature is the theoretical work on multi-product firms and endogenous product choice in international trade that highlights the importance

of extensive margin reallocations in response to economic changes. Eckel and Neary (2010)[18] propose that multi-product firms have a core competency and are able to adapt the manufacturing process to produce additional goods, but at increasing marginal costs. Nocke and Yeaple (2014)[35] model firms choosing their optimal scope as a function of organizational capital and operational efficiency. Given a fixed endowment of organization capital, firms trade off offering more products and producing each product at a lower marginal cost. Mayer, Melitz, and Ottaviano (2014)[32] find that more competitive export markets induce firms to focus exports on their most productive products. Arkolakis, Ganapati, and Muendler (2019)[5] estimate fixed costs of exporting by firm, product, and destination and find economies of scope in market-access costs which counteract increasing production costs for products further from the establishment’s core competency. Dhyne, Petrin, Smeets, and Warzynski (2017)[16] estimate firm-product technical efficiencies in a multi-product production function. They find that import competition induces significant reallocation toward more productive firms as well as within firms toward more productive products. Bernard et al (2019)[8] document the prevalence of carry-along trade, when firms, especially productive ones, tend to export products they do not produce themselves. They find suggestive evidence of demand-scope complementarities to rationalize the positive correlations between firm productivity and the number of products produced and sourced for export. My work reiterates the common theme across these papers that endogenous product choices balance many counteracting forces such as increasing production costs, economies of scope in fixed costs, cannibalization effects, and demand complementarities acting on heterogeneous products. Selection under these forces generate the patterns of coproduction within firms and determine reallocation across and within firms in response to economic shocks.

Within the IO literature, there has long been discussion of economies of scope enjoyed by multi-product firms (Baumol, Panzar, and Willig 1998), but much of

the work has been either theoretical or empirical studies within specific industries (telecommunications in Beresteanu (2005)[7], banking in Iimi (2004)[28] and Pulley and Humphrey (1993)[38], and electric utilities in Gilsdorf (1994)[22]). A recent exception is Boehm, Dhingra, and Morrow (2019)[11], who find that establishments are more likely to expand into products with similar input mixes. In another related paper, Flagge and Chaurey (2014)[20] compare measures of distance between products and their relative importance for product adoption. They estimate that physical proximity to existing locations of production and input similarity have the largest effects. Following this area of empirical work, I estimate interdependencies using micro-data, but with a particular focus on the relationship between productivity and coproduction.

The remainder of the paper is organized as follows. First, Section 2.2 presents empirical facts on the patterns of coproduction in India, focusing on linking multi-production and product switching to establishment and product characteristics. In Section 2.3, I examine the relationship between coproduction and productivity within establishments. I begin with describing a theoretical framework that models productivity and product choice for multi-product establishments in Section 2.3.1. Then in 2.3.2, I establish a correlation between an establishment’s product mix and scope and its product-specific productivities. In Section 2.3.3, I use the dereservation policy in India to estimate the impact of an establishment’s product set on the productivities of its goods to test for the presence of complementarity effects. Section 2.4 concludes.

2.2 Coproduction in India

I explore coproduction in Indian establishments between 2000 and 2010 using data from the Annual Survey of Industries.³ Within establishments, the dataset provides

³Features of the dataset are described in Chapter 1 of this paper, Section 1.2, and in the Appendix, Section A.

detailed information on the set of products produced as well as sales and physical quantities for each product. At the establishment-level, the ASI records labor, material, and capital input usage and other characteristics. The information on products within establishments and the high rate of multi-product plants (see table 1.1) make this dataset well suited to examining coproduction and its relationship to the productivities of goods within an establishment’s product mix. I define products and industries by the 5- and 2-digit classification codes, respectively. I begin by comparing the frequency and patterns of coproduction in this data with empirical facts about multi-product firms recorded in the literature. I compute metrics from Bernard, Redding, and Schott (2010) [10], who studied coproduction in US firms between 1987 and 1997 to see if Indian establishments display similar patterns. Another natural comparison is to the work of Goldberg et al (2010), who also study multi-product producers in India, but use the Prowess dataset between 1989 and 2003. There are a few notable differences between the two Indian datasets that suggest the findings may not be identical. First, Prowess is firm-level while the ASI is establishment-level. The ASI also has better coverage of small producers. Also between 1989 and 2003, India was going through significant trade liberalization as well as reductions in industrial licensing and labor market deregulation. By the years of my data from the ASI, those policy changes had largely finished being implemented. Another important caveat is that because the product coding systems are different across data sources, the definitions of “products”, “industries”, and “sectors” differ, so rates of multi-production and product switching cannot be directly compared across the three papers.

2.2.1 Characteristics of Multi-Product Establishments

The literature on multi-product firms has consistently found that they differ from single product ones on a range of observables. Table 2.1 displays the results of OLS regressions of log firm characteristics (total output, employment, labor productiv-

ity, and establishment TFP) on a dummy for whether the establishment is multi-product, multi-industry, or multi-sector.⁴I include industry fixed effects to focus on within-industry differences in performance. Establishments producing multiple products, industries, and sectors tend to be larger, employ more workers, and are more productive than their single product counterparts, which is consistent with the literature. It is important to note that being multi-product is an endogenous choice, so differences in establishment characteristics should be interpreted as correlations rather than causal effects.

I employ two standard measures of establishment productivity widely used in the literature, revenue-based labor productivity and total factor productivity. Even among the literature on multi-product establishments, many empirical papers, including both BRS and Goldberg et al (2010), estimate productivity at the establishment-level. In the presence of multi-product establishments, these are not necessarily the theoretically appropriate measures of productivity. However, the challenge to estimating within-establishment product productivities is that data on inputs are generally recorded at the establishment-level only. In this Section of the paper, I use the standard measures of establishment productivity to examine how establishment characteristics vary with coproduction and compare my results to similar exercises in the literature. Later in Section 2.3, use the methodology developed in Chapter 1, which estimates theoretically consistent measures of productivities for multi-product establishments, to estimate within-establishment product productivities and use those to further examine how coproduction relates to productivity at the establishment-product-level.

⁴The establishment characteristics presented are computed as follows: Output is the total value of production across all of the establishment's products in a year. Employment is the quantity of labor in man-days. I use two measures of establishment-level productivity: labor productivity, which is value added per unit of labor (man-days), and revenue-based TFP measured using an index number calculation (see Caves, Christensen, and Diewert 1982 [13]).

Establishment characteristic	Multi-product _{<i>ft</i>}	Multi-industry _{<i>ft</i>}	Multi-sector _{<i>ft</i>}
Output _{<i>ft</i>}	1.15	1.15	1.15
Employment _{<i>ft</i>}	0.78	0.82	0.82
Labor productivity _{<i>ft</i>}	0.31	0.33	0.34
TFP _{<i>ft</i>}	-0.04	-0.01	-0.01

Results displayed for regressions of log establishment characteristics on an indicator for multi- product, industry, or sector. Regressions include an industry fixed effect for the establishment’s primary industry. Standard errors are clustered at the industry level. All coefficients are significant at the 1% level except for those on TFP.

Table 2.1: Establishment characteristics for multi-product, multi-industry, and multi-sector

2.2.2 Product Switching

Comparing the sets of products reported by establishments over time, I explore the product switching behavior in India. I find greater product churning in the ASI relative to what the literature has found in other product-level datasets such as the US Census by BRS. This difference is likely attributable, at least in part, to features of the ASI’s data collection method. Establishments are asked to self-report their products each year, which may result in some inconsistency in the product codes used year to year. This would exaggerate the amount of product switching.

Table 2.2 displays the rates of product adding and dropping between years in the ASI. All years are pooled for this analysis. The first panel shows the percentage, of all establishments and of only multi-product establishment, that added at least one product, dropped at least one, both, or neither. The ASI captures a high rate of product switching: 68% of establishments make changes to their product mix in a year. Therefore the implied five year transitions would be greater than observed for the US. The frequencies of both product adding and product dropping are higher among establishments that already produce more than one product, and in fact, 56% of them both add and drop over a given year. The high percentage of establishments

that add and drop products concurrently, also called product churning, suggests that the extensive margin of adjustment within establishments is not just on the scope. Establishments choose which products to produce, not just how many. In panel B of the table, where percentages for each activity are weighted by establishment output, I find that product switching also occurs more in larger establishments, although the difference is less stark than that found by BRS for US firms.⁵

Activity	All establishments	MP establishments
<i>A. Percent of establishments</i>		
None	32	18
Add only	6	12
Drop only	6	4
Both	56	66
<i>B. Output weighted percent of establishments</i>		
None	18	13
Add only	6	7
Drop only	6	5
Both	70	75
Percent of establishments that have added and dropped products since the previous year, with data from years 2000 to 2010 pooled.		

Table 2.2: Percent of establishments adding and dropping products in a year

Table 2.3 displays adding and dropping of 2-digit industries. The first column contains the rates for 5-digit products shown before, while the second column presents the percentage of establishments that add industries, drop industries, do both, or do neither over the course of a year. As expected, there is significantly less switching activity at the higher aggregation level of industries than for products. While 68% of

⁵Note that the ASI data is a census of larger establishments, but only collects data from a sample of small establishments. In order to examine product switching, I need to observe an establishment for at least two consecutive years, so small establishments will be under-represented. Additionally, it is not possible to distinguish between an establishment that does not appear in the sample for the rest of the panel and one that has exited. Therefore, statistics about product dropping do not include the extreme case of establishments dropping all of their products and exiting.

establishments switch products in a year, only 31% switch industries. Products that are added and dropped are not random. Establishments are more likely to stay within industries when adjusting their product mix. This suggests that there is perhaps some establishment-industry specific affinity, which could arise from correlated productivities within industries, complementarities in the productivities, or economies of scope in the fixed costs of adopting products.

Activity	Percent of establishments	
	Across products	Across industries
None	32	69
Add only	6	9
Drop only	6	7
Both	56	14

Percent of establishments that have added and dropped products or industries since the previous year, with data from years 2000 to 2010 pooled.

Table 2.3: Percent of establishments adding and dropping products or industries in a year

The previous exercises have established that multi-product establishments tend to be larger and more productive and also switch products more often. Next, I examine changes in establishment characteristics as the establishment either grows or shrinks its product mix. I run OLS regressions of log changes in establishment characteristics on dummies for net products added and net products dropped since the previous year. The product mix fixed effect controls for economic conditions specific to certain product mixes, while a year fixed effect is used control for time trends because I pool establishment observations across all years.

$$\Delta Z_{ft} = \beta_{add} \text{Net add}_{ft} + \beta_{drop} \text{Net drop}_{ft} + \alpha_{ft} + \mu_t + \epsilon_{ft} \quad (2.1)$$

Where Z_{ft} is one of four log establishment characteristics and α_{ft} is a product mix fixed effect. The results are presented in table 2.4. For establishments that begin with the same product mix, net product adding is associated with increases in total output, employment, and both productivity measures while net product dropping is associated with the opposite. Because the choice of products is an endogenous decision made by the establishment, changes in product mix are likely non-random, and these results can only be interpreted as correlations between changes in establishment characteristics and changes in scope. The correlations across a range of separately measured characteristics suggest that there are real changes happening to the organization of production when the establishment chooses to change its product mix. Therefore changes in product mix are associated with other changes in real economic activity.

	Log changes in establishment characteristics			
	Output _{ft}	Employment _{ft}	Labor productivity _{ft}	TFP _{ft}
Net add _{ft}	0.0652*** (0.00626)	0.0223*** (0.00298)	0.0281*** (0.00490)	0.0187*** (0.00541)
Net drop _{ft}	-0.0698*** (0.00523)	-0.0220*** (0.00304)	-0.0194*** (0.00406)	-0.0314*** (0.00568)
N	126108	126434	103437	103030
R^2	0.021	0.016	0.020	0.016

Regressions of log changes in firm characteristic between year $t - 1$ and t on dummies for net product adding and net product dropping over that period, with a product mix fixed effect and time fixed effect. All consecutive year establishment observations are pooled. Standard errors in parentheses are clustered by product mix. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2.4: Establishment characteristics and net product adding and dropping

Tables 2.5 and 2.6 further explore how the probability an establishment adds or drops a product varies with establishment and product characteristics. I regress the probability an establishment added a product since the previous period on estab-

lishment productivity in the starting year, controlling for employment and age, with product mix (α_{ft}) and year (μ_t) fixed effects.

$$Add_{ft} = \beta_p \log \text{Productivity}_{ft} + \beta_{age} \log \text{Age}_{ft} + \beta_{empl} \log \text{Employment}_{ft} + \alpha_{ft} + \mu_t + \epsilon_{ft} \quad (2.2)$$

Table 2.5 shows that for establishments with the same initial product mix, higher labor productivity is associated with a greater probability of adding products by the following year. The results on labor productivity are consistent with those found by BRS, but I find no significant effect for establishment TFP. Controlling for age, size (employment), and initial product mix, establishments that are more productive are more likely to add products. The correlation between higher labor productivity and more product adding suggests that the higher productivity of multi-product establishments found in table 2.1 may be a result of selection—high productivity establishments are more likely to add products and become multi-product.

Next I turn to the dropping of products from an establishment’s mmix. Table 2.6 presents regressions of the probability a product in the establishment’s mix is dropped next year on the product’s current log output and age.

$$Drop_{ift} = \beta_{out} \log \text{Output}_{ift} + \beta_{age} \log \text{Age}_{ift} + \alpha_i + \nu_f + \mu_t + \epsilon_{ift} \quad (2.3)$$

Note that the product adding regression in table 2.5 is at the establishment level while this regression of product dropping is for an establishment-product observation. I regress a dummy for whether a product in the establishment’s set will be dropped the next period on current characteristics of the product: its output and the number of years that it has been produced by the establishment. Both output and age are measured in logs and relative to the product-year mean, which controls for differences in output and tenure across products. An establishment fixed effect controls for unobserved establishment-level characteristics that affect the rate of product

	Add _{ft}	Add _{ft}	Add _{ft}	Add _{ft}
TFP _{ft}	-0.00668 (0.00396)	-0.00516 (0.00391)		
Labor productivity _{ft}			0.0226** (0.00739)	0.0213** (0.00809)
Employment _{ft}		0.00634* (0.00282)		0.00237 (0.00330)
Age _{ft}		-0.00958** (0.00352)		-0.00725* (0.00321)
<i>N</i>	106873	106873	106873	106873
<i>R</i> ²	0.150	0.150	0.152	0.152

Regressions of the probability of adding products between t and $t + 1$ on log productivity measures in year t , with controls for log employment and age. All specifications include product mix and year fixed effects. The regressions pool establishment observations over consecutive years. Standard errors in parentheses are clustered by product mix. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2.5: Establishment productivity and product adding

dropping. A product fixed effect controls for differences in drop rates across products, e.g. due to the broadness of a particular product definition. The regression pools establishment-product observations across all years, so I include a year fixed effect to control for time trends. I find scale and age dependence in product dropping decisions. Establishments are less likely to drop products in their mix that are larger and have been produced for longer, compared to other establishments producing the same product. The age and scale of a product are potential proxies for establishment-product-specific expertise, such as revenue-based productivity. This provides evidence of non-random selection of products within establishments and is consistent with theories emphasizing selection in explaining the patterns of coproduction. It suggests that establishments engages in product switching in order to reallocate resources towards product lines that are more productive or profitable.

	Drop _{<i>ift</i>}	Drop _{<i>ift</i>}	Drop _{<i>ift</i>}
Output _{<i>ift</i>}	-0.0133*** (0.000418)	-0.0261*** (0.000552)	-0.0277*** (0.000553)
Tenure _{<i>ift</i>}	-0.174*** (0.00177)	-0.0726*** (0.00206)	-0.0676*** (0.00201)
Fixed effects	None	Establishment	Establishment, product, year
<i>N</i>	276992	266149	265572
<i>R</i> ²	0.0393	0.2453	0.3550

Regressions of a dummy for product dropping between year t and $t + 1$ on log relative output and log relative tenure in t . Output and tenure are measured relative to the product-year average by log demeaning. Establishment-product observations are pooled across all years. Standard errors in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2.6: Regressions of product dropping by establishment-product

These empirical results have shown that product switching is more prevalent in multi-product and high revenue establishments. Adding products is associated with increases in output, employment, and revenue-based productivity, and more productive establishments are more likely to add products. The products that do get drop are more likely to be smaller and newer product lines, relative to other establishments producing the same product. The patterns of product switching suggest that the differences in establishment characteristics found in the previous subsection may be driven by endogenous selection effects: more productive establishments endogenously choose to add more products, leading to the observed systematic differences between single product and multi-product establishments. The positive correlation between establishment productivity and scope may result from selection alone, or there may be additional interdependencies that motivate multi-product establishments. The results provide evidence that product switching and scope choices are non-random. Establishments tend to add and drop products within industries, and product switch-

ing behaviors suggest reallocations across products towards the larger and possibly more productive ones.

2.2.3 Heterogeneity of Products Within Establishments

Having looked at the dynamics of product choice, I now turn to examining differences between products within an establishment's product mix. The evidence so far of non-random selection of products within establishments suggests that the differences across the products may be informative of the mechanisms that drive coproduction and reallocation patterns. The following tables examine heterogeneity across products and patterns in frequently coproduced goods.

The distribution of product shipments within establishments by the size of the product set, in table 2.7, shows significant heterogeneity. For each level of size of the product set, the table displays the average percent of total establishment output that is generated by each product, ranked from largest to smallest. Shipments are concentrated in the establishment's top products. Even for establishments producing ten different products, almost half of the revenue on average is generated by its largest product. In the context of the theory of core competencies and flexible manufacturing from Eckel and Neary (2010), establishments facing increasing marginal costs of producing additional products further away from their core competency. The heterogeneity in shipments at every level of scope is consistent with establishments facing increasingly greater costs, or lower productivity, as they add products further from their core. This also fits with the results from figure 1.5 in Chapter 1, which shows similar heterogeneity in productivities within establishments. The degree of heterogeneity in both productivities and output across products within establishments highlights the importance of thinking about productivity at the establishment-product level.

Next, I explore the shape of the product distribution. For reference, I compare the distribution of product shipments to the Pareto distribution, which has widely

Rank	Number of products produced by establishment									
	1	2	3	4	5	6	7	8	9	10
1	100	90	78	71	63	58	54	52	49	47
2		10	18	19	22	22	21	20	20	19
3			4	8	10	11	11	11	12	11
4				2	4	6	6	7	7	8
5					2	3	4	4	5	5
6						1	2	3	3	4
7							1	1	2	3
8								1	1	2
9									1	1
10										1

Table 2.7: Distribution of percent of establishment output by product rank

been found to fit the upper tail of the firm size distribution. If products within establishments follow a Pareto distribution the relationship between log rank and log size should be linear. I regress log product rank on log product share of establishment shipments, by total number of products. Figure 2.1 plots, for establishments producing 4, 6, 8, or 10 products, the fitted regression line as well as the actual values. The result appears very similar to the distributions found by BRS: the actual data falls above the fitted line in the middle of the distribution and below the line in the tails, implying thinner tails than the Pareto. The deviation from a log linear relationship, and thus from a Pareto distribution of products, appears greater as the number of products increases.

The coproduction matrix in figure 2.2 visualizes the frequency of coproduction across the range of products to show which ones tend to be produced together. The frequencies of coproduction are computed as the number of establishment-year observations that produce the combination of two products (at the 5-digit ASICC level) together. These frequencies are aggregated to the level of 2-digit industries to make the figure easier to read and interpret. For a given pair of industries, the data repre-




Figure 2.1: Within establishment product distribution, by total number of products

sented in the matrix is the number of establishment-year observations that produce at least one product from each industry.⁶ Darker squares indicate higher frequencies of coproduction. The darker areas along the diagonal show that there is more coproduction within an industry or sector, which is consistent with the findings that establishments are more likely to be only multi-product than multi-industry or multi-sector, and that product switching occurs more often within industries. Establishments tend to concentrate their product choices within industries, where the products may share certain characteristics. There are additional cells off the diagonal where the data show systematic patterns of coproduction. For example, there is relatively frequent

⁶ Industries are ordered by number in the graph. The ticks along the axes indicate the start of each 1 digit sector. See the data appendix for the categories. The resulting plot is symmetric along the diagonal. Thus, only the lower triangle is displayed here.

coproduction of products in the rubber (industry code 41) and plastic (42) industries with metal and machinery products (71 through 79).




Figure 2.2: Frequency of coproduction in India, summarized by 2 digit industries

These exercises using Indian data from 2000 to 2010 show similar trends for multi-product establishments as Bernard, Redding, and Schott (2010) which used US Census data from 1987 to 1997, and Goldberg et al (2010), who use the Prowess data on Indian firms from 1989 to 2003. Before drawing conclusions from differences in the frequency of product switching or establishment scope, it is important to keep in mind the differences in product coding and data collection across the datasets. The evidence so far, both in this paper and more generally in the multi-product literature, leave little doubt that coproduction is not random. Establishments choose their product mix endogenously, responding to differences in productivity, costs, technol-

ogy, and consumer tastes across products. The result is systematic patterns to which products are commonly produced within the bounds of the establishment and which tend to be produced by different establishments. These patterns across products depend on their identity and may be the result of correlations in an establishment’s ability to product certain goods, as well as interdependencies and spillover effects across products.

2.3 Productivity in Multi-Product Establishments

To better understand why certain products are produced together within establishments and the interdependencies that affect this choice, I examine the productivity specific to establishment-product observations. In part 2.3.1 of this section, I outline a theoretical framework for how establishments choose a product mix and the interdependencies in productivity that may arise. Part 2.3.2 presents evidence of a correlation between within-establishment productivities and scope. Then in 2.3.3 I use a policy change of small scale product dereservation in India as an instrument to disentangle causation and look for evidence of a complementarity effect of an establishment’s scope acting through productivity.

2.3.1 Theoretical Framework

As a theoretical framework for estimating within establishment product-level productivities and for thinking about product interdependencies, I build on the model of multi-product establishments developed in Chapter 1. Section 1.3 describes the endogenous product choice and production decisions of an establishment in the model. Establishments see their productivity draws ϕ_{ift} for all potential products $i = 1, \dots, N$ and all potential product mix combinations $\mathcal{J}$ before choosing a set of products $\mathcal{J}_{ft}$ to produce. This endogenous choice of the product set highlights the

non-random selection of products discussed above. The establishment pays a fixed cost of adopting a product set $F_{ft}(\mathcal{J})$ and then makes production decisions for all products in its set.⁷

The productivity term ϕ_{ift} is the Hicks neutral total factor productivity that enters the production function for good i . I make no assumptions on the distribution of productivity draws and allow it to vary by product, establishment, and time. Realized productivities may depend on the establishment's chosen product mix: $\phi_{ift}(\mathcal{J}_{ft})$. This implies that how productive the establishment is at making good i can potentially depend on which other products are in the current product mix $\mathcal{J}_{ft}$. This causal relationship between the product set and individual productivities is one form of production complementarity.

To the multi-product model of Chapter 1, I introduce some structure on productivity that is useful for thinking about production complementarities. Suppose there is a base productivity ω_{ift} that the firm draws for each product, whose realization is independent of the product set and any other establishment choices. The entire vector of base productivities is drawn from some joint distribution that might be establishment specific, but the realized draws are not affected by any of establishment's choices this period.⁸ One interpretation of these base productivities ω_{ift} is simply how productive the establishment would be if i were its sole product (i.e. $\mathcal{J} = \{i\}$). In the production stage, the productivity is then modified by a function of the product set. A simple parameterization of this is for the complementary component that depends on the set and the base productivity to be log additive:

$$\log \phi_{ift}(\mathcal{J}) = B_i(\mathcal{J}) + \log \omega_{ift} \quad (2.4)$$

⁷The production decisions for a given product mix are the desired output quantities Y_{ift} and the allocation of inputs across products $L_{ift}, M_{ift}, K_{ift}$.

⁸I specify this period because the distribution could potentially be affected dynamically by establishment choices, for example through research and development (R&D). However, I assume the base productivity draws are realized at the start of each period before the establishment makes any choices, and so are taken as given for that period.

where $B_i(\cdot)$ is a product-specific function that maps the set of products $\mathcal{J}$ onto a scalar that is additive to the product's log base productivity. I assume that the complementarity effect of the other products in the mix might depend on the identity of the product, but the effect of a given set is the same for all establishments that produce it.

I assume establishments have full knowledge of how product set choices influence productivity and take these into account when choosing a product mix.⁹ The productivity $\phi_{ift}(\mathcal{J})$ impacts both the extensive (whether to produce good i) and intensive (the output quantity of i if it is produced) margins of the establishment's decisions. However, complementarities depend only on the extensive margin. While productivity may be a function of the product mix, I do not allow it to depend on other endogenous establishment decisions such as output quantities or use of inputs. This assumption limits some types of interactions across products but is necessary for tractability of the establishment's problem and estimating productivities. It is important to note that although the model and proceeding estimation methods are consistent with the existence of complementarities, I do not assume them. This model nests cases where there are no complementarity effects and where there are diseconomies of scope. If $B_i(\mathcal{J})$ is zero for any product set, the establishment's productivity equals its base productivity and is unaffected by its product mix or scope, implying no complementarities to product i . There may also be cases when the product mix causes the productivity of a good to be lower than what it would be if the establishment were single product, $B_i(\mathcal{J}) < B_i(\{i\})$, $\mathcal{J} \neq \{i\}$. The complementarity function $B_i(\mathcal{J})$ is the object of interest that I estimate in the following sections.

The complementarity effect on productivity is the focus of this paper, but it is but one potential channel of interdependency across products in the model. Productivity draws may also be correlated. This feature is intuitive if establishments that are

⁹With the information assumptions, this formulation is isomorphic to the establishment drawing a vector of productivities for each possible set, presented in the first chapter.

productive in one good are also likely to be productive in other similar products.¹⁰ These correlations arise in the model from positive covariances between draws of the base productivities ω_{ift} . Additionally, multi-product establishments may experience economies in the fixed cost of adopting a product mix. For two similar products i and j , the fixed cost of producing them together may be less than the sum of costs of producing each alone: $F_{ft}(\{i, j\}) \leq F_{ft}(\{i\}) + F_{ft}(\{j\})$.

Within the scope of this paper, I focus specifically on potential complementarities in production that act through the productivity term. There are two primary challenges to identifying a complementarity effect of the product mix on productivity. The first is estimating establishment-product productivities when inputs are observed only at the establishment-level. I use the estimation procedure developed in Chapter 1, which is consistent with complementarities and correlations in productivity and with economies of scope in the fixed costs of adopting a product mix. The methodology uses output elasticities estimated from single product establishments and optimal cost minimization rules to infer input allocations by product and estimate establishment-product productivities. The second challenge is that the productivities within an establishment’s endogenously chosen product mix are jointly determined by non-random selection over potentially correlated products and any complementarity effects. To separately identify a causal effect of the product mix on productivity, I use exogenous variation in establishments’ product mixes generated by the rollback of small scale reservation policies in India.

Other types of interdependencies and complementarities may exist that aren’t captured by this model. For example, one could imagine spillover effects of coproduction that act through inputs, such as sharing lumpy units of capital between products or complementarities in sourcing inputs as in Boehm, Dhingra, and Morrow (2019).

¹⁰I use the term “similar” somewhat abstractly here, without formally defining how similarity or distance between products might be measured. Within the data, I consider sharing an industry as a reasonable proxy for similarity across products.

There could also be complementarities on the demand side from economies of scope in marketing costs or firms being able to increase demand for their products by bundling, as found by Bernard et al (2010) for Belgian exporters. However, in order to identify the channel of interest, a complementarity effect through productivity, the estimation procedure are not able to take into account these other forms of complementarities not included in the model.

2.3.2 Productivities of Coproduced Goods

To address the first challenge to estimating complementarity effects through productivity, I need to estimate within-establishment productivities for each product. Studies of coproduction have generally used firm or establishment-level measures of productivity, such as the labor productivity and TFP from the previous section and in BRS. However, the within-establishment heterogeneity in product shipments seen in figure 2.1 strongly suggests that there may be heterogeneity in the productivities of products within establishments as well. Then the measured establishment-level productivity would be a weighted average of the productivities of each good in the product mix. A major reason that establishment productivity is used in studies of multi-product establishments is likely the lack of product-specific input data that would allow the application of standard productivity estimation techniques such as Olley-Pakes [36] at the establishment-product level. Capital, materials, and labor inputs are recorded at the establishment or firm level in all datasets of which I am aware.

In Chapter 1, I developed a methodology for estimating product-specific productivities within an establishment's mix, which I implement here to examine the role of establishment-product level productivities in coproduction.¹¹ Refer to Chapter 1 for details of the estimation procedure and discussion of the estimated productivities. As hinted by the heterogeneity in product shipments, I found significant variabil-

ity in within-establishment productivity (table 1.2). I use the establishment-product productivities to take a closer look at how productivities are influenced by extensive margin product choices. In this subsection, I establish a correlation between an establishment's scope and the productivities of goods in its product mix and use these results to guide the estimation of complementarities in the following subsection.

From table 2.4, net product adding is associated with an increase in establishment-level productivity while net product dropping is associated with a decrease. At the product-level within establishments, this relationship could arise from changes to the productivity of goods remaining in the mix or from a compositional change to the product mix through selection. Table 2.8 shows the changes in productivities of incumbent goods when there is either net product adding or dropping. The coefficients are from a regression of log change in productivity for a product that remains in the mix on a dummy for net product adding and a dummy for net product dropping. The regression also includes product (γ_i) and initial product mix (α_{ft}) fixed effects, so it measures how productivity changes vary with net adding and dropping within a product and for establishments producing the same initial product mix.

$$\Delta \log \phi_{ift} = \beta_{add} \text{Net add}_{ft} + \beta_{drop} \text{Net drop}_{ft} + \alpha_{ft} + \gamma_i + \mu_t + \epsilon_{ift} \quad (2.5)$$

The observations included in the regression are limited to incumbent products that are produced in both years $t-1$ and t , and so are neither added nor dropped in the period. Thus, the regression estimates how productivity varies with net adding and dropping of other products. Net product adding, which equates to an increase in scope, is associated with an increase in productivity of incumbent products. Table 2.8 shows

¹¹Caution is required when comparing quantity-based TFP across products because the productivity measure is sensitive to the units of output. All productivity estimates used in this paper are measured relative to the 5-digit product mean. When comparing results in this section to earlier ones using establishment-level measures of productivity, it is important to be aware of fundamental differences in the measures. The establishment-product level productivities here are quantity-based. In contrast, establishment labor productivity and TFP are revenue based.

that the estimated increase in establishment productivity found earlier cannot only be the result of adding high productivity products that raise the establishment average. The productivities of incumbent goods also tend to increase when the establishment adds products in net. Again, it is necessary to stress that the positive correlation between increases in scope and increases in the productivities of incumbent products may be the result of either selection or complementarities. If productivity draws are positively correlated across some products, then establishments may produce the incumbent products more productively at the same time that they add other related products, giving rise to the pattern found here. Also possible is a complementarity effect where increasing scope causes increases in productivity for incumbent products.

	Change in productivity $_{ift}$
Net Add $_{ft}$	0.0839*** (0.0304)
Net Drop $_{ft}$	-0.0374 (0.0288)
N	56989
R^2	0.077

Regression of log change in productivity between $t - 1$ and t on indicators for whether the establishment added or dropped products in net between years, with fixed effects for year, product, and product mix. Establishment-product observations are pooled across all years. Bootstrap standard errors in parentheses are clustered by establishment. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2.8: Changes in productivity with net product adding or dropping

Earlier in this paper, I showed that establishment-level labor productivity¹² is higher in multi-product and multi-industry establishments (see table 2.1). With estimates of productivity at the product level, one would expect to find

¹²The effect on establishment TFP was not significant.

	Multi-product	Multi-industry
Log ϕ_{ift}	0.171*** (0.0346)	0.0458 (0.066)
N	244190	244190
R^2	0.0094	0.0082

Table reports regressions of log demeaned productivity on indicators for whether the establishment is multi-product or multi-industry, with year and product fixed effects. Bootstrapped standard errors in parentheses are clustered by establishment. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2.9: Productivity in multi-product establishments

that establishment-product productivities are also higher for multi-product and multi-industry establishments. Table 2.9 compares the productivities for establishments that are multi-product or multi-industry to those that produce only one product or industry. I find that indeed productivity is higher for products made by multi-product establishments, relative to same good when produced by single product establishments. The coefficient on producing across multiple industries, while positive, is not statistically significant. I am unable to conclude that for a give product, multi-industry establishments are more productive than single-industry establishments, which include establishments producing multiple goods in one industry. This suggests that the relationship between coproduction and productivity may be stronger within industries, and is consistent with the relatively high frequency of within-industry coproduction.

Next I establish a correlation between an establishment's scope and the productivities of its product mix. I regress log productivity on the establishment's scope, the total number of products it produces, along with product and year fixed effects in the first column of table 2.10. The fixed effects control for different effects of scope across products as well as trends over time. Establishments that produce a larger number

of products tend to be more productive at what they make compared to other establishments that produce the same product, although the statistical significance is weak. Further breaking down this finding in columns 2 and 3, I find that the positive coefficient on scope is approximately twice as large when counting only the number products in the same industry, whereas the number of products produced in other industries has no significant effect on productivity. This pattern of results show that the connection between the number of products and productivity is stronger within the same industry than across different industries. The correlation between coproduction and productivity depends not only on the number of products, but also the identity of the other products in the set.

	Log ϕ_{ift}	Log ϕ_{ift}	Log ϕ_{ift}
Number of products $_{ft}$	0.0381~ (0.0218)		
Products in same industry $_{ift}$		0.0764*** (0.0188)	0.0770*** (0.0179)
Products in other industries $_{ift}$			-0.00649 (0.0312)
N	244190	244190	244190
R^2	0.009	0.009	0.010

Regressions of log productivity, demeaned by the average for the product, on measures of the establishment's scope, with fixed effects for year and product. Bootstrapped standard errors in parentheses are clustered by establishment. ~ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2.10: Productivity and establishment scope

That product switching behavior is correlated with changes in productivity of other products in the mix indicates there is non-random selection of products. Producing a larger number of products, especially within the same industry, is correlated with higher productivity. This pattern could arise from positive correlation of an

establishment’s product-specific productivities, especially within the same industry. Then high productivity draws for incumbent products are correlated with high productivities for new varieties that are endogenously added to the set. Another possibility is that there are spillovers across product lines that produce a complementarity effect when producing similar products together. In the subsection, I use an exogenous policy change as an instrument to try to disentangle correlation from causation in the relationship between coproduction and productivity.

2.3.3 Complementarities in Production

I use the Indian data to examine one particular channel of interdependencies: effects on the productivity of a good by the establishment’s product mix. This interaction is represented in the model as $B_i(\mathcal{J})$, the component of product i ’s productivity that is a function of the establishment’s product mix. In the absence of a theory-guided functional form, I choose a simple and tractable form for $B_i(\cdot)$: I estimate the effect of the product mix as linear in the number of products the establishment produces in the same industry. Intuitively, we expect greater spillover effects across products that are more similar, and the results in table 2.10 show that within-industry scope is more strongly correlated with productivity than the total number of products or products in different industries. The baseline specification I use is:

$$\log(\phi_{ift}) = \beta \log(\text{Number of products}_{ind(i)ft}) + \mu_{if} + \nu_t + \epsilon_{ift}$$

The parameter of interest is β , which measures the percentage change in a product’s productivity for each percent increase in the total number of products the establishment produces within the industry of i . The regression includes fixed effects for the variety¹³, to control for idiosyncratic characteristics of the establishment-product

¹³“Variety” refers to an establishment-product pair (i, f) .

combination, and for the year, to control for time trends in productivity. Therefore all identification comes from time-series changes within establishment-product pairs. In theory, the complementarity effect may depend on exactly which products are being coproduced. In this linear specification, then, the coefficient β would be the average spillover effect of all other same-industry products in the mix. A positive estimate of β would be evidence of complementarity effects while a negative estimate would suggest diseconomies of scope.

An endogeneity problem is present when trying to estimate complementarity effects because the establishment chooses the product set after seeing how productive it would be at each good. From the previous exercises, I find a correlation between high productivity and producing more products concurrently, especially when producing more within the same industry, but this relationship in the data could result from either complementarities or from correlated productivities. To separate out production complementarity through productivity from other channels of interdependencies across products, I exploit the exogenous variation created by a policy change that occurred during the span of my data in India: the rollback of small-scale reservations.

Small-scale industry (SSI) promotion was a policy beginning in 1967 that aimed to boost employment and growth in small businesses. It limited the production of certain products, typically in labor-intensive manufacturing industries, to small and medium enterprises. By 1996, over 1,000 products fell under the policy. Studies of the SSI policy (Mohan 2002 [34]; Martin, Nataraj, and Harrison 2017 [31]) suggest that the choice of products to include was largely arbitrary. A small or medium enterprise was defined as one with fixed assets under a certain threshold which, by 1999, was 10 million rupees. In 1997 following trade liberalization and other deregulation, India began the gradual dereservation of products covered by the policy, removing the restrictions on their production. The bulk of the dereservations occurred between 2003 and 2008, and in 2015 the last of the reservation policies were abolished. In

their 2017 paper, Martin, Nataraj, and Harrison study the removal of India’s SSI promotion and find that dereservation actually encouraged growth in output and employment. I use the concordance created by Martin, Nataraj, and Harrison to match dereserved products from the SSI list to ASICC product codes. Between 2000 and 2010, I am able to match the dereservation of 155 unique products in the ASI, spanning 35 2-digit industries.¹⁴

Variation in the timing of dereservation across products is a source of exogenous variation in the probability an establishment chooses to produce these products. The removal of products from the SSI list was gradual, with matched products being dereserved in every year of the data except 2009. The differences in the timing of dereservation generate variation across products. For establishments that were previously too large to produce a product, the dereservation increases the probability that they will begin producing the good. One way to interpret the policy’s effect is that establishments that were restricted had an infinite fixed cost for producing a product mix that includes the good. Dereservation represents a decrease in the fixed cost. For small establishments that were already allowed to produce the product, dereservation gave them the opportunity to expand. For both small and large establishments, I would expect the dereservation of a product to affect the probability of producing it, although effects may be heterogeneous. To measure the effect of within-industry scope on a product’s productivity, I aggregate to the industry level: I instrument measures of establishment scope in an industry at time t with the number of products in the industry that have been dereserved by t . I expect dereservation to be associated with scope, but it is unlikely to be directly related to establishment-product productivities. Then degree of dereservation in an industry is a source of exogenous variation across industries and time that can help separately identify a causal effect of the set of products on productivity from non-random selection into the set of products.

¹⁴See appendix A.3 for more detail on dereserved products in the data.

The first stage regression of the number of products in the same industry ($\text{Products}_{ind(i)ft}$) on the number dereserved in the industry ($\text{Dereserved}_{ind(i)t}$) is summarized in table 2.11.

$$\text{Products}_{ind(i)ft} = \alpha \text{Dereserved}_{ind(i)t} + u_{if} + v_t + e_{ift} \quad (2.6)$$

The first column pools all establishment-products, regardless of establishment size. The number of dereserved products in an industry is associated with an increase in the number of products that establishments produce in that industry. Dereservation has no significant effect when the regression in column 2 is limited to small enterprises, which were already allowed to produce the products under SSI promotion. In the third column which includes only non-SME establishments, the effect is what is expected when the dereservation policy allows previously constrained large establishments to begin producing the products.¹⁵ While first stage regression suggests that dereservation has the expected effect on the number of products in an industry, the dereservation variable does not vary across establishments within an industry-year, and only a small share of products (155 out of 3655 total observed in the data) are dereserved during the panel. This has resulted in low power of the instrument, seen by the F statistics of the first stage regressions. The estimated effects of the log number of dereserved products on log number of products establishments produce is the expected sign and statistically significant in the pooled and non-SME regressions, but the magnitude of the effect is small.

The IV estimates of equation 2.3.3 are presented in table 2.12. I regress log productivity on the log number of products produced in the same industry, instrumented

¹⁵The data does not directly indicate which enterprises were considered small for the purposes of the SSI reservations. I define SMEs in the data as establishments with less than 10 million rupees in net capital. This measure is an approximation of whether the establishment was affected by the reservation policy because official SME designations use different data on fixed assets, and because some exceptions were made for larger establishments that produced the products prior to their reservation.

	<u>All</u>	<u>SME</u>	<u>Not SME</u>
	Products _{ind(i)ft}	Products _{ind(i)ft}	Products _{ind(i)ft}
Dereserved _{ind(i)t}	0.00255** (0.0008545)	0.00107 (0.000949)	0.00319* (0.00134)
N	95483	40048	52574
R^2	0.874	0.865	0.877
F	6.44	0.72	4.12

Regressions of log number of products in the industry of i on log number of products dereserved in that industry. All regressions include variety and year fixed effects. Bootstrapped standard errors in parentheses are clustered by establishment. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 2.11: Effect of dereservation on within-industry scope

with the number of dereserved products in the industry, with variety and year fixed effects. The structural regression is:

$$\log(\phi_{ift}) = \beta \text{Products}_{ind(i)ft} + \mu_{if} + \nu_t + \epsilon_{ift} \quad (2.7)$$

The first column of the table reports OLS estimates while the remaining columns report IV estimates over different samples. In the second column, which includes all establishments, I estimate that for a 1% increase in the number of same-industry products, the productivity of the good increases on average 7.6%. This coefficient is an order of magnitude larger than the OLS estimate in the table's first column, but is not statistically significant. The third and fourth columns break down establishments by whether they are small and medium enterprises (SME) or not. There was no significant relationship found between dereservation and scope for SMEs in the first stage, and thus the coefficient in the IV regression for SME is likewise not significantly different from zero. The regression in column 4 includes only non SMEs and finds a similar estimate to the regression using all establishments.

An additional concern is that for products that were dereserved during the panel, there may be additional interactions with productivity or product choice that are not accounted for in the model. For example, some large establishments were allowed to produce a reserved product through a grandfather clause, but had to limit their production volume. Constraints on output quantities would violate the assumptions of the productivity estimation procedure and lead to bias in the estimated productivities. Dereservation could also lead to general equilibrium effects on the product’s profitability. To address these concerns, I use a restricted sample for the regression in column 5 that removes establishment-product observations with products that are observed being dereserved in the data. The regression on the remaining observations estimates the complementarity effect in products that were not themselves affected by the dereservation policy, and so may better satisfy the exclusion condition for instrumental variables. The estimated complementarity effect is similar in magnitude to the coefficients for all establishments and for only non-SMEs and has weak statistical significance (between 5% and 10%). The results fail to find strong evidence of complementarity effects in productivity, suggesting that selection is a better explanation for the observed correlation between within-industry scope and productivity. However, I find a very weakly significant positive effect on never-reserve products in non-SME establishments, which suggests that perhaps complementarity effects shouldn’t be ruled out completely.

When interpreting these results, there are a few limitations to consider. Estimating complementarity effects from only non SME establishments may be subject to selection bias if complementarity effects (whether positive, negative, or zero) vary across establishments—in other words, if the function $B_i(\mathcal{J})$ is not only product specific but also establishment specific. For example, some establishments may be organized such that they are better able to share knowledge or experience across product lines. Then the coefficient may overestimate the magnitude of complementar-

	OLS, All	IV, All	IV, SME	IV, Not SME	IV, Not SSI
	Log ϕ_{ift}	Log ϕ_{ift}	Log ϕ_{ift}	Log ϕ_{ift}	Log ϕ_{ift}
Products _{ind(i)ft}	0.152*** (0.0227)	7.614 (10.36)	10.626 (21.66)	7.849 (5.379)	7.791~ (4.232)
N	95483	95483	40048	52574	50712

OLS and IV regressions of log demeaned productivity on log number of products in same industry. All regressions include variety and year fixed effects. “SME” includes only establishments with less than Rs. 10 million in net capital. “Not SME” includes only establishments with more. “Not SSI” includes only observations of products that were never reserved and are produced by non SME establishments. Bootstrapped standard errors in parentheses are clustered by establishment. $\sim p < 0.10$, $* p < 0.05$, $** p < 0.01$

Table 2.12: Effect of within-industry scope on productivity

ities if higher establishment complementarity potential is correlated with producing more products and accumulating a capital stock above the SME threshold. Also, the functional form of productivity here assumes that any complementarity effects are linear in only the number of products in the same industry, but in reality one might expect that complementarities may be nonlinear in scope, or vary by industry or even for specific product combinations. In that case, the effect I estimated would be an average across industries.¹⁶

2.4 Conclusions

In this paper, I aimed to explore the factors that determine the patterns of coproduction within multi-product establishments and why they appear systematically different from single product establishments. In particular, I was interested in the relationship between coproduction and productivity. Using detailed data on products within establishments in India, I examined the relationship between coproduction

¹⁶Because there is no variation in the dereservation instrument within an industry-year, I am not able to identify industry-specific complementarities. There is also not enough power in the instrument to test more flexible functional forms than the linear specification considered here.

and establishments or product characteristics, which were consistent with theories of multi-product establishments in the literature and previous empirical work. Guided by these findings, I used estimates of within-establishment product productivities to establish a correlation between the number of products in an establishment and the productivities of those products. I then looked for evidence of complementarity effects in productivity by exploiting exogenous variation from a policy change to separately identify a causal effect from correlation.

First I presented a set of empirical facts about multi-product establishments and product switching behavior in India. Establishments that produce multiple products or industries tend to be larger, have more employees, and are more productive. The rates of product adding and dropping are higher in the Indian ASI data across one year than in the US Census data over five year periods, which may be a result of different aggregation for what is defined as a product or of the data collection method of the ASI. Multi-product and multi-industry establishments and establishments with higher revenue, employment, and productivity engage in more product adding and dropping. I also find evidence of scale and age dependence in product dropping. Looking within establishments, there is significant heterogeneity in the revenue share and productivity of products. There are systematic patterns to which products tend to be produced together by establishments that suggest frequent coproduction within industries as well as across certain combinations of industries. The empirical facts match what has been recorded in the literature for US firms, even though the economies are quite different. That coproduction in Indian establishments between 2000 and 2010 follows the same patterns as in US firms between 1987 and 1997 suggests that the findings in BRS may be broadly generalizable. These empirical findings pointed strongly to non-random product choices within establishments, which could result from selection or additional interdependencies across products.

Next, I looked at the productivities of goods within establishments to help explain the non-random pattern of product choices. I observed increases in productivities during years when establishments added products on net. Products in establishments with greater scope, especially within the same industry, tended to be more productive. The results established a positive correlation between an establishment's scope and its productivities. To test for any complementarity effects, however, it was necessary to separately identify a causal effect of coproduction on productivity from non-random selection in the presence of correlations in the productivity draws. Using the dereservation of products in India as an instrument, I estimated any causal effects on productivity of same-industry scope and found only weak evidence complementarity. The results suggest that endogenous selection with correlated productivities across products may be a better explanation for the observed correlation between scope and productivity, but complementarities are potentially economically important and may merit further examination.

There are many ways that product choice and production decisions could be interdependent across products, which requires a fundamentally different way of thinking about the establishment's choices from models of single product establishments, especially when considering reallocation effects across products. Figuring out which types of interdependencies exist, disentangling their effects on product choice, and estimating their importance are needed for a full understanding of how multi-product establishments operate. Otherwise, we won't be able to predict how establishments choose their products and production in the face of shocks or policy changes that affect certain portions of the economy. The findings here, especially the strong correlation between productivity and coproduction, indicate that the reallocation across products within an establishment is an important margin of adjustment. If there are complementarities affecting productivity, the pattern of reallocation will differ and so will measured establishment-level or aggregate productivities.

Chapter 3

Evaluating Tax Harmonization in the European Union

with Rowan Shi

3.1 Introduction

One of the goals of the European Union is to reduce economic frictions between its member states. Through free trade, a monetary union, and policy to encourage the free movement of labor, the EU takes steps towards integrating its member states into one economic unit. However, they still remain considerably differentiated along other dimensions, such as fiscal policy and taxation. This paper investigates the impact of one particular dimension of variation across member countries: their corporate tax rates.

In particular, corporate taxes are levied on economic activity within a country's borders. Dispersion in tax rates across the EU could therefore distort where production is carried out, especially in the context of multinational production. Multinational firms choose where to locate their affiliates and concentrate production to serve global markets, both decisions that are affected by the dispersion in tax rates

across countries. Multinational firms, defined as firms that own plants in multiple countries, account for a disproportionately large share of international trade activity and employment. For example, Bernard, Redding, and Schott (2009)[24] found that in 2000, multinationals employed 27.4% of the U.S. civilian work force.

When multinationals take tax rates into account, there are two margins of note. Firstly, when deciding where to set up their affiliates, tax rate dispersion may result in location decisions that would otherwise be suboptimal. Secondly, conditional on the location of their affiliates, production may be distributed across them in a similarly distorted way. Both scenarios result in spatial misallocation. The goal of this paper is to quantify this misallocation and to consider the welfare effects of an EU-wide tax harmonization.

In order to address these questions, we develop a general equilibrium model of multinational location choice and production that features export platforms and corporate taxation by country. Firms can serve foreign markets through direct FDI, exports from home, or by exporting from a third country in which it has an arms-length production facility, often referred to as an export platform or bridge FDI in the literature. Allowing export platforms introduces additional richness to the location decisions of multi-national firms through interactions across production locations, but also increases the computational burden of solving the firm’s problem. Thus, we introduce a novel algorithm to solve the model quickly and robustly, described in Section 3.4.

We calibrate the model to recent aggregate data on the EU and conduct a counterfactual of tax harmonization. This imposes a common tax rate on all countries, which we calibrate to raise the same total tax revenue as before harmonization. We find that the previously higher-tax countries gain in employment and firms, enjoy higher wages, and see an increase in production reallocated away from the previously lower-tax countries. At the same time, the low-tax countries collect higher tax rev-

enue and see welfare gains in the short run, although there is little change in equalized long run worker utility.

Our paper builds on the multinational production literature and on the literature examining tax dispersion. Helpman, Melitz, and Yeaple (2004)[25] characterized the selection over firm productivity into exporting or FDI, a key component of many models of MP including ours. However, by not allowing export platforms, they limit the interactions between a firm’s decisions of how to serve each market. Ramondo and Rodríguez-Clare (2013)[39] estimated the gains from trade and MP in a Ricardian model based on Eaton Kortum (2002)[17]. Their model has no fixed cost of multinational production, so production occurs wherever variable costs are lowest, with no discrete choice of firm locations. Most similar to our work is the paper on multinational production and location choice with export platforms by Tintelnot (2017)[42], which looks at the affiliate location and production choices of German multinationals within a set of twelve countries. We use elements from Tintelnot’s model to construct our own, placing it within a general equilibrium framework and incorporating taxation. In order to address the combinatorial discrete choice (CDC) problem of the firm’s affiliate location set, we build on the insights from Jia (2008)[29]; Antràs, Fort, Tintelnot (2017)[3]; and Arkolakis Eckert (2017)[4] to develop a new algorithm, which we use to solve the CDC problem for a set of 27 countries. Finally, our work relates to the literature on tax incentives of multinational location choice, including Devreux and Griffith (2003)[15], and Buetter and Ruf (2007)[12]. Fajgelbaum, Morales, Serrato, Zidar (2018)[19] studied state and federal taxes in the US as a source of spatial misallocation and found welfare gains from tax harmonization. Our paper aims to address similar issues for multinational firms in the EU. Whereas Fajgelbaum et al. explored the effects of many forms of taxation at the state and federal level and modeled firms that choose a single state in which to locate and produce, we focus

on the country-specific corporate tax and allow multinational firms to place foreign affiliates and produce in a set of countries.


The remainder of the paper is organized as follows. Section 3.2 provides background on the corporate taxation system of the EU. Section 3.3 sets down our theoretical framework of multinational production and exporting with country-specific tax rates. Section 3.4 describes how we solve and calibrate the model. Section 3.5 presents a counterfactual of tax harmonization within the EU, and Section 3.6 concludes.

3.2 Corporate Taxation in the European Union

The European Union has an integrated market and currency union, but corporate taxation is individually controlled by the member states. As a result, there exists significant dispersion in corporate tax rates across countries, with top marginal tax rates ranging from 10% in Bulgaria to 33.99% in Belgium.

The base for corporate taxation can be either territorial, on profits made within the country, or residential, on the worldwide profits of firms registered in a country. In practice, however, there is an extensive set of deferral policies and double taxation treaties between EU countries that limit taxation on profits made in another EU country. Additionally, the EU Interest Royalty Directive and the EU Parent-Subsidiary Directive eliminate withholding tax on interest, royalties, and repatriated profits across member states. Thus, in our model, we assume that EU firms are taxed on their profits in the country of production. In most countries, the corporate tax is assessed at a flat rate, with some providing reduced rates for firms with profits below a certain threshold. To capture the incentives for multinational enterprises, which tend to be large firms, we use the top marginal statutory tax rate of each country. Figure 3.1 maps these for each EU country considered.

Figure 3.1: Corporate tax rates vary widely across the EU




Map of EU member states shaded by their top marginal statutory corporate tax rate in 2014. Lighter shading denotes higher rates. See Appendix C for a table of tax rates by country.

In addition, a member state’s corporate tax rate is strongly correlated with empirical measures of productivity. Figure 3.2 graphs the relationship between statutory corporate tax rates and aggregate value added per worker. There is a strong positive relationship, perhaps forming the foundation for the common narrative that low-tax countries rely on their low rates to “poach” production and jobs from more productive countries.¹ From the perspective of a firm choosing where to locate production, it must trade off production capability with tax burden. This tension is the key driving force we capture in our model and analysis. In response, large high-tax countries, notably Germany and France, have pushed for tax policy harmonization within the

¹The notable exception is Ireland, which features both high value added per worker and low tax rates. A likely explanation is that Ireland is widely regarded as a “tax haven”, into which large corporations shift income to avoid taxation. This would inflate measured value added per worker for Ireland and other tax haven countries. Due to low data availability, we do not address income shifting in this paper.

Figure 3.2: Low-tax countries tend to be less productive, as measured by VA per worker.



Source: Firm-level data from AMADEUS in 2014. We compute country-level aggregates of value added per worker among manufacturing firms.

EU. These proposals have predictably been met with strong resistance from low-tax countries.

3.3 A Model of Multinational Location Choice

We model multinationals in a tradable goods sector with monopolistic competition similar to Tintelnot (2017), where firms self-select into multinational production, but we also build in corporate taxation which impacts production decisions on both the extensive and intensive margins. Firms place affiliates in an endogenously chosen set of countries, make production and exporting decisions, and are taxed on their profits in the country of production. Then this sector is embedded in a general equilibrium model with a perfectly competitive non-tradable sector and consumers who are either workers or capitalists. We can use our model to evaluate both the short run effects

of a policy, holding populations and firms constant in each country, as well as long run effects in the presence of endogenous firm entry and labor mobility.

3.3.1 Households

Consumers live in a set J of countries. There are $|J| = C$ countries in total. They consume housing, a homogeneous nontradable good, and a final tradable good. Preferences across the three consumption categories are Cobb Douglas. A representative consumer in country k has utility

$$V_k = u_k \left((c_k^T)^\beta (c_k^N)^{1-\beta} \right)^\eta h_k^{1-\eta}$$

where u_k is the utility shifter in country k , c_k^T and c_k^N denote quantities of tradables and nontradables consumed, and h_k is the quantity of housing consumed. The utility shifter u_k can be interpreted as a fixed level of amenities in country k . Each country is endowed with a fixed quantity H_k of housing, with each unit of housing renting for s_k . Given the Cobb-Douglas structure of demand, households spend a share η of their income on consumer goods and $1 - \eta$ on housing. Similarly, among consumption goods, households spend a share β on tradable goods and $1 - \beta$ on nontradables. In turn, the tradable baskets are composed of differentiated varieties ω of goods. These goods are aggregated with a constant elasticity of substitution ρ as follows:

$$c_k^T = \left[\int_{\omega} c_k(\omega)^{\frac{\rho-1}{\rho}} d\omega \right]^{\frac{\rho}{\rho-1}}, \quad \rho > 1.$$

Households each inelastically supply a unit of labor, which is used in the production of goods and the payment of fixed costs. In addition to wage income, workers also receive the rebated tax revenue collected by their home country, T_k .² For country k

²We model government spending of tax revenue as a direct transfer to workers for simplicity. As a result, worker welfare in our model is sensitive to this assumption on tax revenue.

with L_k measure of labor, total income is $I_k = w_k L_k + T_k$, where w_k is the equilibrium wage. Through utility maximization, the indirect utility of a household in country $k \in J$ is

$$V_k = u_k \left(\frac{1}{\eta} \left(\frac{P_k^T}{\beta} \right)^\beta \left(\frac{P_k^N}{1-\beta} \right)^{1-\beta} \right)^{-\eta} H_k^{1-\eta} \frac{(\eta I_k)^\eta}{L_k}$$

where P_k^T is the CES-accurate price index over all tradable consumption goods available, and P_k^N is the price of the nontradable good.

3.3.2 Capital Owners

In addition to workers, there are absentee capital owners in each country that own a share b_k of a global portfolio consisting of housing and firms' profits from all countries. We will define profits in subsequent sections that describe firms, but briefly, the profits in a country consist of after-tax variable profits minus the fixed costs of setting up affiliates. Profits may be positive in the short run, but long run profits, which also take out payments of entry costs, equal zero on aggregate due to free entry in each country. Capitalists consume the nontradable and tradable goods with the same relative preferences as households, but do not consume housing. We assume capitalists in each country are massless so that they do not contribute to congestion. Their income and expenditure is $b_k \sum_j (s_j H_j + \Pi_j) = b_k \sum_j ((1-\eta)I_k + \Pi_j)$, where Π_j represents total net profits of firms originating in country j . The presence of capitalists who consume in k but have income from holdings in all countries allows for trade imbalances to occur.

3.3.3 The Nontradable Sector

In the perfectly competitive nontradable sector, firms produce a homogeneous good using a country-specific, linear technology. If L_k^N units of labor work in the nontradable sector, $\alpha_k L_k^N$ units of nontradable good is produced. Households can supply their

labor to either the nontradable or tradable sector, resulting in a single equilibrium wage for each country. For simplicity, there is no tax on nontradable goods.

3.3.4 The Tradable Sector

The focus of our analysis is on firm location and production in the tradable goods sector. We model the tradable sector to allow for multinational firms to emerge endogenously. In particular, each firm has a country of birth, indexed i , but is not restricted to production there. Consequently, when serving a foreign market, the firm has a choice between exporting from its origin country, FDI in the destination market, or sourcing from an export platform in a third country. Our modelling of tradable firms adapts elements from Tintelnot (2017) and Eaton Kortum (2002).

The traded sector is populated by a continuous measure M_i of firms in each country $i \in J$, determined by free entry. Firms may enter in each origin country by paying a fixed cost f_i^e in units of labor to draw a core productivity and operate. Each firm is identified by its country of origin i , a single good variety ω , and a core productivity $\varphi_i(\omega)$.³ The core productivity is drawn from a distribution specific to its birth country, G_i . The single variety firms engage in monopolistic competition in each export market they serve. At birth, each firm is endowed with a production facility in its origin country. However, it may also choose to establish production facilities in any foreign country $j \in J$ at a fixed cost of f_j units of labor. These can be interpreted as the costs of forming a subsidiary in a foreign country as well as setting up the plant and distribution network. These costs for the home country are included in f_i^e , so there is no additional fixed cost associated with the plant in the origin. Once all production locations are established, a production productivity

³Our model differs here from Tintelnot (2017), in which each firm produces a continuum of products. This modeling choice, along with an assumption that the elasticity of substitution is the same within and across firms, allowed Tintelnot to have an interior solution for the share of a firm's varieties sourced from each of its production locations. However, for simplicity we omit this feature as it is unnecessary for our calibration and counterfactual.

$\varphi_{ij}(\omega)$ associated with each plant is revealed. Then, the firm chooses how much to produce in each production location, as well as where to sell the goods produced by each affiliate. Each country j taxes profits of plants located within its borders at a flat rate of τ_j , which is known to firms at the time they make their affiliate location decisions.

As with nontradables, firms produce tradable goods using only local labor as an input. Producing outside of its home country carries an additional iceberg cost of arms-length production. Finally, firms can sell to any country $k \in J$ using any of its plants, subject to an iceberg cost of transport. Thus, when a firm born in i produces at an affiliate located in j to sell to consumers in k , it encounters marginal cost

$$\frac{w_j \gamma_{ij} d_{jk}}{\varphi_{ij}(\omega)}$$

where $\gamma_{ij} \geq 1$ is the arms-length cost of production for a firm originating in i when it produces in j , and $d_{jk} \geq 1$ is the iceberg trade cost of shipping from j to market k .

Given the timing of the firm problem, it can be solved backwards. First, conditional on an already-established set of affiliates and their plant productivities, we find the optimal price and quantity if the firm were to serve destination k from an affiliate in j . We call this step the intensive margin. Then, using these optimal decisions, we can determine the best sourcing strategy for each destination market. This choice is the conditional extensive margin, since it is made dependent on a set of affiliates and plant productivities. Finally, given the solutions to the first two problems, we find the set of affiliates maximizing the expected overall firm profits (over all possible plant productivity draws). This final affiliate location choice is the extensive margin of production.

We begin with the intensive margin. Since entering the export markets bears no fixed costs, producing firms will serve all destinations. Sales to one destination do not

affect sales to other destinations and production costs are linear, so each destination can be considered independently. As such, the firm chooses a price for each destination to maximize profits from that destination. Given CES demand, tradable good prices have a common markup of $\frac{\rho}{\rho-1}$ over marginal costs. The quantity of production is determined by CES demand at the price the firm sets. Tax rate dispersion does not directly affect this step of the firm decision.

Next, we consider the conditional extensive margin. If the firm were to serve market k from an affiliate in j , it would derive variable profits

$$\frac{1}{\rho} (1 - \tau_j) X_k^T (P_k^T)^{\rho-1} \left(\frac{\rho}{\rho-1} \frac{w_j \gamma_{ij} d_{jk}}{\varphi_{ij}} \right)^{1-\rho}.$$

where X_k^T is total expenditure on tradables in country k (including both households and capitalists).⁴ Then, the firm simply chooses the production location $j \in \mathcal{J}$ that maximizes variable profits after taxes.⁵ The firm's sourcing decision for a market k is therefore influenced not only by its marginal costs of serving k , but also the corporate tax rate τ_j in each j . Because profits are taxed in the country of production, differences in tax rates across countries may lead to distortions in the firm's sourcing decisions, and thus its distribution of production across affiliates. More specifically, firms may optimally choose a sourcing location that does not have the lowest marginal cost for serving a given market k , but instead maximizes net variable profit for the market by having a sufficiently low tax rate.

Finally, total firm profits given a set of locations $\mathcal{J}$ is the sum of these profits by market less the fixed costs of setting up production in each country in $\mathcal{J}$. For the extensive margin, the firm therefore chooses its set of production locations to maximize the expected value of total profits. This choice is also distorted by the variation

⁴Each firm produces only one variety, so we drop indexing by variety ω when discussing the decisions for a given firm.

⁵Throughout this paper, we use $\mathcal{J}$ to denote a firm's affiliate location set. This is not to be confused with J , the complete set of countries in the economy. Firms choose their location set as a subset of all possible countries: $\mathcal{J} \subseteq J$.

in tax rates through their effects on the intensive and conditional extensive margins. To complete the characterization of the firm problem, we assume the idiosyncratic location productivity $\varphi_{ij}(\omega)$ is drawn from a distribution depending on the firm's core productivity. In particular, we assume the distribution is Fréchet with scale $\varphi_i(\omega)T_j$ and a common shape θ across all firms and countries. T_j is a productivity shifter specific to country j that affects all production located there. Then, given a location set $\mathcal{J}$, we can solve for closed form expressions describing the probabilities of sourcing from each country and expected total profits using standard properties of the Fréchet.

The probability that destination k is served by $j \in \mathcal{J}$ is

$$\pi_{ijk}(\mathcal{J}) = \frac{(1 - \tau_j)^{\frac{\theta}{\rho-1}} (w_j \gamma_{ij} d_{jk} / T_j)^{-\theta}}{\Theta_{ik}(\mathcal{J})}$$

where

$$\Theta_{ik}(\mathcal{J}) = \sum_{j \in \mathcal{J}} (1 - \tau_j)^{\frac{\theta}{\rho-1}} (w_j \gamma_{ij} d_{jk} / T_j)^{-\theta}$$

summarizes location set $\mathcal{J}$'s profit potential for market k . A location set has higher profit potential to serve market k when the countries in the set have lower tax rates, lower labor costs, higher productivity shifters, lower arms-length costs from the origin i , or lower trade costs to market k . Putting it all together, a firm with core productivity summarized by $z_i(\omega) \equiv \varphi_i(\omega)^{\rho-1}$ ⁶ chooses the location set that maximizes total expected profits

$$\max_{\mathcal{J} \subseteq J} \left[\frac{1}{\rho} A z_i(\omega) \sum_k X_k^T (P_k^T)^{\rho-1} \Theta_{ik}(\mathcal{J})^{\frac{\rho-1}{\theta}} - \sum_{j \in \mathcal{J}} w_j f_{ij} \right]$$

⁶For the remainder of this paper, we summarize core productivity using $z_i(\omega) \equiv \varphi_i(\omega)^{\rho-1}$ for notational ease. In all of our aggregate equations of interest, core productivity φ_i appears to the power $\rho - 1$.

where $A \equiv \Gamma\left(\frac{\theta-\rho+1}{\theta}\right)\left(\frac{\rho}{\rho-1}\right)^{1-\rho}$ is a constant and $f_{ij} \equiv \mathbb{1}[j \neq i]f_j$ for notational simplicity. In order for the expectation to exist, we impose the technical assumption that $\theta > \rho - 1$. The solution $\mathcal{J}^*$ summarizes the firm's extensive margin decision.

Examining the expected profit function, expected variable profits are proportional to the sum of profit potentials $\Theta_{ik}(\mathcal{J})$, weighted by market size $X_k^T(P_k^T)^{\rho-1}$. Since profit potential mechanically increases as another element is added to $\mathcal{J}$, adding a production location increases expected variable profits. However, firms must pay a fixed cost to add this plant. As a consequence, firms trade off the higher expected profits of a new plant with the flat fixed cost of its establishment. Because taxes enter in the expression for profit potential, tax rate dispersion directly affects the firm decision at this stage. A location j 's contribution to profit potential takes into account the marginal cost components associated with production in j for k , as well as the prevailing tax rate. Even though marginal costs may be low in a particular location j , a high tax rate could dampen its profit potential.

Since the location set $\mathcal{J}$ is chosen before plant-specific productivities are realized, all firms with the same birthplace and core productivity z_i will choose the same optimal set $\mathcal{J}^*$. We call this policy function $\mathcal{J}_i^*(z_i)$. However, firms with identical core productivity will have different production patterns within this set based on the actual idiosyncratic draws of plant-specific productivities.

To summarize, corporate tax rates directly distort firm decisions in two areas: first, differences in tax rates across countries affect the set of production locations chosen by the firm; second, given a chosen set of production locations, they also affect the firm's sourcing decisions. These distortions are caused by the differences in tax rates across countries, and not by the mere presence of corporate taxation. In fact, if the tax rate were uniform, firm decisions would be identical to the case of no taxation given aggregates. The differences in tax rates across countries introduce an additional

consideration for location and sourcing decisions, which firms trade off against input costs, trade costs, arm's length production costs, and market conditions.

3.3.5 Equilibrium

While each firm in the economy has varying production and trade patterns stemming from idiosyncratic draws of φ_{ij} , the law of large numbers implies the probabilities and expectations derived above are exactly realized over the measure of firms. For example, the sourcing probability $\pi_{ijk}(z_i)$ is also the fraction of i -firms with affiliates in $\mathcal{J}$ that serve k from their j plant.⁷ We can also define λ_{ijk} , the share of sales in k produced in j by firms born in i :

$$\lambda_{ijk} = \frac{M_i \int_{z_i} z_i \Theta_{ik}(z_i)^{\frac{\rho-1}{\theta}} \frac{\pi_{ijk}(z_i)}{1-\tau_j} dG_i(z_i)}{\sum_{i'} \sum_{j'} M_{i'} \int_{z_{i'}} z_{i'} \Theta_{i'k}(z_{i'})^{\frac{\rho-1}{\theta}} \frac{\pi_{i'j'k}(z_{i'})}{1-\tau_{j'}} dG_{i'}(z_{i'})}.$$

This expression is the key share of the model, on which many of the later equilibrium conditions and variables depend. For example, the (j, k) trade share is simply $\sum_i \lambda_{ijk}$. We define μ_{ij} , the share of firms born in i with a production location in j , by aggregating over all firm whose optimal location set includes a plant in j :

$$\mu_{ij} = \int_{z_i} \mathbb{1}[j \in \mathcal{J}_i^*(z_i)] dG_i(z_i).$$

Equilibrium in the tradable sector is characterized by the aggregate price index, the labor market clearing condition, balance of payments, and the long run endogenous measures of workers and firms in each country. In the long run, households are freely mobile across countries, and there is free entry of firms into each origin

⁷We take a notational shortcut by writing the sourcing probability of the optimal location set for a firm with core productivity z_i as $\pi_{ijk}(\mathcal{J}_i^*(z_i)) = \pi_{ijk}(z_i)$, and similarly for $\Theta_{ik}(z_i)$.

country.⁸ Consequently, the population L_i and measure of headquartered firms M_i in each country are determined endogenously by these conditions. Free mobility allows households to move to countries where they would have higher utility, but there are counteracting congestion forces from limited housing stock and labor market equilibrium. Ultimately, utility is equalized across all countries in the long run equilibrium:

$$V_k = u_k \left(\frac{1}{\eta} \left(\frac{P_k^T}{\beta} \right)^\beta \left(\frac{P_k^N}{1-\beta} \right)^{1-\beta} \right)^{-\eta} H_k^{1-\eta} \frac{(\eta I_k)^\eta}{L_k} = \bar{V}, \quad \forall k. \quad (3.1)$$

In the long run, firms enter freely in each origin country up to the point where the total expected value of entry is zero. As a result, the aggregate variable profits less fixed costs of setting up plants in each country equals the cost of entry:

$$M_i w_i f_i^e = \frac{1}{\rho} \sum_k X_k^T \sum_j (1 - \tau_j) \lambda_{ijk} - M_i \sum_j \mu_{ij} w_j f_{ij} \quad (3.2)$$

Thus, total profits in each origin i are zero in long run equilibrium: $\Pi_i = 0, \forall i$. This also implies that capital owners receive positive income only from housing rents in the long run.

Integrating over all firms, the price index in destination country k is

$$(P_k^T)^{1-\rho} = A \sum_{i,j} M_i \int_{z_i} z_i \Theta_{ik}(z_i)^{\frac{\rho-1}{\theta}} \frac{\pi_{ijk}(z_i)}{1-\tau_j} dG_i(z_i). \quad (3.3)$$

The price index definition integrates over the distribution of core productivities and includes the terms $\Theta_{ik}(z_i)$ and $\pi_{ijk}(z_i)$, which depend on the optimal location set for each productivity. Thus, the equilibrium condition incorporates the firm's discrete location choice decision, which precludes an analytical solution. Developing a com-

⁸Even though the EU has the single market program that allows movement of workers across its members, in practice, labor may not be perfectly mobile due to frictions such as language differences. Imperfectly mobile labor would imply that even in the long run equilibrium, utility may not be equalized across countries. It is important to note this when considering the real-world welfare and wage impacts of tax harmonization, the counterfactual we run in section 3.5.

putational method for solving the firm's problem tractably even for a large number of potential locations is a key contribution of this paper and is detailed in the following section.

The labor market clearing condition in each country of production j is

$$w_j L_j = X_j^N + w_j f_j^e M_j + \sum_i M_i \mu_{ij} w_j f_{ij} + \frac{\rho - 1}{\rho} \sum_k X_k^T \sum_i \lambda_{ijk} \quad (3.4)$$

where the first term is payments to labor in the nontraded sector, while the rest are payments to labor in the traded sector. Total payments to labor in tradables can be decomposed into the fixed costs of entry in j , fixed costs paid by foreign multinationals setting up affiliates in j , and variable production costs from all firms producing in j . On the other hand, payments to labor in nontradables are simply equal to sales because of perfect competition.

The presence of capitalists in each country who consume goods locally but have income from a global portfolio leads to potential trade imbalances. The balance of payments condition gathers all expenditure on traded goods as follows

$$X_k^T = \beta \eta I_k + \beta b_k \sum_j ((1 - \eta) I_j + \Pi_j). \quad (3.5)$$

The first term captures expenditure on tradables by households while the second term corresponds to tradables spending by capitalists.

Given parameters $\{\beta, \eta, u_j, H_j, b_j, \alpha_j, \rho, \tau_j, \theta, G_j, T_j, f_j, f_j^e\}_{j=1}^C$ and $\{d_{jk}, \gamma_{ij}\}_{i,j,k=1,\dots,C}$, an equilibrium is defined as a set of $\{P_j^T, X_j^T, w_j, L_j, M_j\}_{j=1}^C$ that satisfies equations (3.1) to (3.5) above.

3.4 Quantitative Methods

The main quantitative challenge of the model is solving the firm’s problem and embedding the solution in general equilibrium. More specifically, there is no analytic solution to the firm’s extensive margin decision of where to set up plants. However, the entire distribution of the firm’s optimal location set (across all core productivities) appears in many of the equilibrium conditions, including in the aggregate shares λ_{ijk} and μ_{ij} . As a result, solving for equilibrium and calibrating the model requires solving the firm’s problem for many values of core productivity. It is therefore important to solve it robustly and quickly.

To this end, we use a novel algorithm developed in Shi (2019) [40] to numerically solve for the optimal policy function $\mathcal{J}_i^*(\cdot)$ along the range of core firm productivity. This feature is desirable in our context because, as noted above, many key variables in the model require integration over the entire distribution of firm choices. By solving for the policy function directly, we avoid discretizing the distribution of core productivities into a grid $\{z_i^n\}_n$ and solving for each optimal set $\mathcal{J}_i^*(z_i^n)$ separately. In addition, it precludes any numerical error resulting from interpolating firm choices between gridpoints.

Once the firm problem can be solved, we next turn our attention to numerically computing the equilibrium as well as performing a calibration in anticipation of a quantitative counterfactual. We describe these methods in this section as follows: subsection 3.4.1 sketches the intuition behind the algorithm and describes how we use it to solve for equilibrium; and subsection 3.4.2 shows how we calibrate the model by simulated method of moments using data from the European Union.

3.4.1 Solving for Equilibrium

The main challenge in computing the numerical equilibrium is solving the firm's problem, which is a combinatorial discrete choice over a large set of possible countries. In other words, each firm chooses an optimal set of production locations among 2^C possibilities to maximize expected total net profits. Without any further structure, this problem grows exponentially in the number of countries. Additionally, it must be solved for many possible values of firm productivity, multiplying the potential computation time considerably. However, the expected profit function has several key properties, of which the algorithm takes advantage to solve for the policy function $\mathcal{J}_i^*(\cdot)$. We replicate the expected profit function below for reference.⁹

$$\begin{aligned}\Pi_i(\mathcal{J}, z) &= \frac{1}{\rho} A z \sum_k X_k (P_k^T)^{\rho-1} \Theta_{ik}(\mathcal{J})^{\frac{\rho-1}{\theta}} - \sum_j \mathbb{1}[j \in \mathcal{J}] w_j f_{ij} \\ \Theta_{ik}(\mathcal{J}) &= \sum_j \mathbb{1}[j \in \mathcal{J}] (1 - \tau_j)^{\frac{\theta}{\rho-1}} (w_j \gamma_{ij} d_{jk} / T_j)^{-\theta}\end{aligned}$$

To begin, the profit function is continuously increasing in firm productivity z , but discrete if $\mathcal{J}$ changes. Intuitively, this property implies that if a particular set $\mathcal{J}^*$ is optimal for a firm with productivity z , it will also be the optimal choice for a firm with a similar productivity $z + \varepsilon$. As a result, the optimal policy function $\mathcal{J}_i^*(\cdot)$ will be constant over some range of productivity, before jumping to a new set discontinuously. In turn, this new set will be the best choice for some range of productivity, before the optimal set switches again. Then, the productivity range can be partitioned into a finite list of intervals, each associated with a different optimal set. The problem now becomes finding the productivity values where the optimal sets change, as well as the optimal sets in question. Note the optimal choice depends both on firm productivity z and location of birth i , so we must solve for the policy function associated with

⁹Note that the profit function here is for a given firm with affiliate location set $\mathcal{J}$ and core productivity z , which is not the same as the aggregate profit in i , previously denoted Π_i .

each country of birth. In what follows below, we omit the i subscript for notational lightness.

To define the two additional properties of the expected profit function precisely, we introduce the “marginal value” of location j as:

$$D_j(\mathcal{J}, z) = \Pi(\mathcal{J} \cup \{j\}, z) - \Pi(\mathcal{J} \setminus \{j\}, z) .$$

This difference describes the additional expected profits of including j in $\mathcal{J}$ for a firm of productivity z . Note that the firm problem is difficult to solve precisely because this marginal value depends on $\mathcal{J}$. In other words, the firm problem comprises of a collection of *interdependent* decisions: knowing the extra expected profits attributable to j is impossible without also knowing which other locations are included. However, this marginal value (and, indirectly, the expected profit function) has two important regularities. First of all, D_j is “decreasing in $\mathcal{J}$ ”: given $\mathcal{J} \subset \mathcal{J}'$, then $D_j(\mathcal{J}, z) > D_j(\mathcal{J}', z)$. The marginal value of including j in $\mathcal{J}$ decreases if other additional locations are added to $\mathcal{J}$. Broadly, the profit function can be interpreted as concave in its first argument. This shape arises from substitution effects across plant locations. All else equal, a potential affiliate is less beneficial for a firm with a large set of plants compared to a firm with fewer, since it is relatively less likely that the new location can serve any market at lower cost than the other affiliates. Secondly, D_j is increasing in z : all else equal, additional locations generate higher marginal benefit for more productive firms. Economically, this property implies complementarity between firm productivity and production location sets, which we loosely call “complementarity”.

The algorithm relies on these two properties to guide the search for marginal productivities where the optimal set changes. The broad idea is to squeeze the optimal

set between two sets.¹⁰ The intuition is as follows. Suppose over an interval of productivity $[z_l, z_r]$, we know a subset of items $\underline{\mathcal{J}} \subseteq \mathcal{J}^*$ that are definitely included in the optimal set. Likewise, suppose we also know a superset containing the best set, $\mathcal{J}^* \subseteq \overline{\mathcal{J}}$.¹¹ These sets can be interpreted as “lower” and “upper” bounds on the optimal set. Then, $\overline{\mathcal{J}} \setminus \underline{\mathcal{J}}$ contains the locations that may or may not be in the optimal set.

Consider one of these locations, j . For all $z \in [z_l, z_r]$, it is the case that $D_j(\overline{\mathcal{J}}, z) \leq D_j(\mathcal{J}^*, z) \leq D_j(\underline{\mathcal{J}}, z)$: the marginal value of including j is bounded from above and below. It is bounded from below by the marginal value of including j in the upper bound set, while it is bounded from above by the marginal value of inclusion into the minimal set. This conclusion is a direct consequence of the expected profit function’s concavity in its first argument. For any z with $D_j(\overline{\mathcal{J}}, z) > 0$, it is therefore optimal to include j : the marginal value of its inclusion into $\mathcal{J}^*$ is definitely positive. In particular, let $\tilde{z}$ be the productivity where this marginal value is 0, so that $D_j(\overline{\mathcal{J}}, \tilde{z}) = 0$. From the complementarity property, we conclude that any $z \in [\tilde{z}, z_r]$ will have a positive marginal value to j ’s inclusion. Then, firms of this productivity will place a plant in j .

Likewise, for any $z \in [z_l, z_r]$ with $D_j(\underline{\mathcal{J}}, z) < 0$, it is better to exclude j : here, we compute the maximal marginal value associated with j ’s inclusion, by adding it to the lower bound set. If this maximum possible marginal value is negative, then firm would be better off not opening a plant in country j . If now $\tilde{z}$ denotes the productivity for which j ’s maximum possible marginal value $D_j(\underline{\mathcal{J}}, \tilde{z}) = 0$, then any $z \in [z_l, \tilde{z}]$ will exclude j from its optimal set. Repeating this logic for all uncertain locations j , the lower and upper bound sets can be updated along the productivity distribution. In particular, for all productivity values, the lower bound set will grow while the

¹⁰Papers featuring combinatorial discrete choice problems in the literature (Jia 2008; Antràs, Fort, Tintelnot 2017; Arkolakis Eckert 2017) use a very similar idea, but for a discrete value of z .

¹¹It is always possible to find such a pair of subsets – the empty set and the entire set satisfy these conditions.

upper bound set will shrink. The algorithm iterates this procedure, progressively squeezing the optimal set $\mathcal{J}^*$ from above and below. Ultimately, it returns a list of intervals partitioning the range of core firm productivity, along with the optimal set for each interval.

This type of solution is desirable because the relevant integrals over the firm distribution can be more simply written. For example, the price index (3.3) aggregates individual firm prices over the entire productivity distribution. However, the firm-specific components are simply $\Theta_{ik}(\mathcal{J}_i^*(z))$, $\pi_{ijk}(\mathcal{J}_i^*(z))$, and z . The first two only require the firm's optimal set, and are therefore constant within the intervals returned by the algorithm's solution. If these intervals are indexed by e , let the associated endpoints be $[\underline{z}^e, \bar{z}^e]$, corresponding $\Theta_{ik}^e \equiv \Theta_{ik}(\mathcal{J}^{*e})$, and similarly for π_{ijk}^e . Then, the price index can be re-expressed as

$$(P_k^T)^{1-\rho} = A \sum_{i,j} M_i \sum_{e=1}^{E_i} (\Theta_{ik}^e)^{\frac{\rho-1}{\theta}} \frac{\pi_{ijk}^e}{1-\tau_j} \int_{\underline{z}^e}^{\bar{z}^e} z \, dG_i(z).$$

For many distributions G , this resulting integral is easily and reliably computed.

Finally, the algorithm to solve the firm problem is embedded into a larger procedure to calculate equilibrium. We guess a set of aggregates $(w_j, X_k^T P_k^{T\rho-1}, M_i)$, solve the firm problem for this vector, then check the aggregate equations (3.1) to (3.5). The equilibrium solution of the model is found by solving for the fixed point of the aggregates.

3.4.2 Calibration

The goal of this paper is to examine the general equilibrium implications of tax harmonization in the European Union. With this goal in mind, we calibrate the

model to aggregate outcomes from the EU.¹² First, we must take a stand on the form of G_i , the core productivity distribution. We assume it is Pareto with common shape ξ and country-specific scale $z_j^{\min}$. For each member country, we estimate this scale as well as the productivity shifter for producing there (T_j), the fixed costs of entry (f_j^e) and setting up affiliates (f_j), the utility shifters (u_j), and the shares of the global portfolio owned by capitalists (b_j). The main estimation procedure uses simulated method of moments to match the bilateral affiliate location shares. We use data on location shares, production shares, consumption of tradable goods, and number of firms in each country, while solving for endogenous wages and prices.

The remaining parameters are calibrated outside of the SMM estimation by using common literature values or computing them directly from data. We set the elasticity of substitution between varieties of tradable goods ρ to 6, which implies a markup of $\frac{\rho}{\rho-1} = 1.2$, or 20% over marginal cost. For the shape parameter of the Pareto distribution, which governs the dispersion of core productivity draws, we use $\theta = 7$. This value is within the range estimated by Eaton and Kortum (2002), and also used by Tintelnot (2017). Finally, we use Tintelnot’s estimate of the Pareto shape parameter governing the distribution of core productivity draws for plant, $\xi = 6.436$. Tintelnot estimates this parameter using maximum likelihood and micro-data on German multinationals.¹³ We estimate our own values of the Pareto productivity scale parameter $z_j^{\min}$, which is country-specific.

The model also includes a set of bilateral iceberg trade costs and bilateral iceberg arms-length production costs. Following the gravity literature, we parameterize these as functions of distance, shared language, and shared border between each pair of

¹²Due to its small size and data limitations, we omit Malta from our analysis. We calibrate our model to the remaining $C = 27$ EU countries in 2014.

¹³Tintelnot identifies his set of parameters including the productivity dispersion with data on the affiliate location sets and output of each affiliate of German multinational firms. Although he does not study the EU, his estimation includes many of the same countries as ours. In both models, the dispersion parameter is assumed to be common across countries. Thus, we believe this is a reasonable value to use in our calibration.

countries.

$$d_{jk} = \beta_{cons}^d \times \text{dist}_{jk}^{\beta_{dist}^d} \times \beta_{contig}^d \text{contig}_{jk} \times \beta_{lang}^d \text{lang}_{jk}$$

$$\gamma_{ij} = \beta_{cons}^\gamma \times \text{dist}_{ij}^{\beta_{dist}^\gamma} \times \beta_{contig}^\gamma \text{contig}_{ij} \times \beta_{lang}^\gamma \text{lang}_{ij}$$

We normalize $d_{jj} = 1$ and $\gamma_{ii} = 1$ for own-country pairs.¹⁴ We use coefficients estimated by Tintelnot and data on bilateral distances, shared language, and contiguity from CEPII to construct these bilateral costs.

Determining the remaining parameters relies on EU member states data in 2014 obtained from Eurostat. The aggregate data includes national accounts, consumption prices and expenditures, and notably the foreign affiliates statistics (FATS). Some parameters of the model can be computed directly from this data. Specifically, the labor force in each country is used for L_i and the number of firms originating from each country is used for M_i .¹⁵ The expenditure shares of goods and tradables from the representative household's utility function, $\eta = 0.75$ and $\beta = 0.7$, are computed from aggregate expenditure data.

We estimate the remaining set of parameters $\{z_j^{\min}, T_j, f_j, f_j^e, u_j, b_j\}_{j=1}^C$ using simulated method of moments. Recall that $z_j^{\min}$ is the scale of the core productivity distribution for each origin country, T_j shifts the productivity of plants producing in j regardless of the firm's origin, f_j is the fixed cost of setting up an affiliate, f_j^e is the fixed cost of entering, u_j shifts the level of amenities to households, and b_j is the share of the global portfolio owned by the country's capitalists. These country-specific parameters, along with the bilateral trade and arms-length production costs, govern how firms, affiliates, and production is located across the EU member states. Our calibration strategy guesses values for a small set of composite parameters $\{w_j f_j, w_j / T_j, (P_j^T)^{1-\rho}\}_{j=1}^C$ which are needed to solve the firm's problem and

¹⁴Setting own-costs as 1 is a normalization given the presence of T_j and $z_j^{\min}$.

¹⁵While M_i is an equilibrium outcome, the free entry condition (3.2) can be inverted for f_i^e that exactly rationalizes this mass of firms.

determine an equilibrium of the model. We then calibrate the model by matching the implied matrix of affiliate location shares $\mu_{ij}, \forall i, j = 1 \dots C$ to the EU data, subject to equilibrium constraints on the price index and shares of tradables production in each country.

For a set of guesses for $\Phi = \{w_j f_j, w_j/T_j, (P_j^T)^{1-\rho}\}_{j=1}^C$, we first solve the firm's problem over the core productivity distribution for each country. The firm's problem uses aggregate values w_j/T_j , $X_j^T/(P_j^T)^{1-\rho}$, and $w_j f_j$ for all j and returns the optimal location sets along the distribution of z . We use data on expenditures on tradable goods X_k^T from Eurostat, purged of shipments from outside the EU countries of our model. Next, we estimate the scale z_i^{min} of the core productivity distribution in each country $G_i(z)$ by exactly matching the sum of foreign affiliate location shares in i , $\sum_{j \neq i} \mu_{ij}$.

Then we can numerically compute the aggregate conditions. In particular, we compute the affiliate locations shares μ_{ij} , the price indices P_k^T , and shares of production in each country $\frac{Y_j^T}{\sum_i Y_i^T}$. We match the model's implied location shares to the data:

$$M_{ij}(\Phi) = (z_i^{min})^\xi \sum_{e=1}^{E_i} \mathbb{1}[j \in \mathcal{J}_e^*] ((z_i^e)^{-\xi} - (\bar{z}_i^e)^{-\xi}) - \mu_{ij}^{data} \quad (3.6)$$

The model is calibrated by minimizing the sum of squares of the objective moment conditions M_{ij} for all countries i and j , subject to matching the price indices and production shares to data.

For any guess of the composite parameters Φ , it is possible to find an equilibrium by setting a free parameter in each of the equilibrium conditions for labor market clearing, zero profits, balance of payments, and utility equalization. The rest of these parameters can be calibrated outside of the simulated method of moments optimization using the calibrated composite parameters Φ^* . We compute the equilibrium wages w_j using the labor market clearing condition. With wages w_j , we can back out the parameters f_j and T_j from the guesses of $w_j f_j$ and w_j/T_j in Φ^* . The rest of the

parameters are chosen as the residuals that satisfy equilibrium conditions: for each country, we set the fixed costs of entry f_j^e to satisfy the zero profit condition, b_j to satisfy the balance of payments, and u_j to equalize household utilities.

Parameter	Calibration strategy
$\rho = 6$	Tintelnot (2017)
$\theta = 7$	Tintelnot (2017), Eaton Kortum (2002)
$\xi = 6.436$	Tintelnot (2017)
d_{jk}	Tintelnot (2017) and CEPII
γ_{ij}	Tintelnot (2017) and CEPII
L_i	labor force by country in Eurostat
M_i	number of enterprises by country in Eurostat FATS
$\eta = 0.75$	consumption share of household expenditure in Eurostat
$\beta = 0.7$	tradables share of consumption expenditure in Eurostat
f_j	SMM, target affiliate shares μ_{ij} from FATS
T_j	SMM, target share of production in j
z_i^{min}	SMM, target sum of affiliate shares $\sum_{j \neq i} \mu_{ij}$
f_j^e	SMM, solve zero profit condition in equilibrium
u_j	SMM, solve utility equalization in equilibrium
b_j	SMM, solve balance of payments in equilibrium

Table 3.1: Calibration of parameters

Table 3.1 summarizes the calibration strategy for each of the model's parameters. The top panel lists the parameters taken from the literature or computed directly from data. The bottom panel summarizes those parameters obtained through our simulated method of moments estimation and the equilibrium conditions or data moments used to target them. In Appendix C, we examine how well the calibrated model fits the untargeted moments wages, trade shares, and multinational production shares, as an over-identification test.

3.5 Counterfactual Tax Harmonization in the EU

Harmonization of corporate tax rates within the EU has been subject of discussion in recent years. Imposing a common corporate tax rate across all member states would eliminate incentives from tax rate differentials, which could result in different equilibrium patterns of multinational location and production, and subsequent general equilibrium effects on other aggregates. In our counterfactual exercise, we are interested in how tax harmonization would affect firms' location decisions and production patterns, and how these may influence employment, wages, tax revenue, and welfare in general equilibrium. Optimal choice of a harmonized rate is outside the scope of this paper. Instead, we consider a counterfactual with a common tax rate that generates the same long run total tax revenue as the calibrated equilibrium for the EU in 2014.

We find that a harmonized rate of 27.26% in each country generates the same total tax revenue. This counterfactual policy implies 22 countries would increase their tax rates while 5 countries would decrease them. We examine both the short run and long run effects of tax harmonization, exploring its effects on the spatial reallocation of employment and production, and the resulting impact on tax revenue, wages, and welfare in each country. In the short run, we assume that consumption, production, and affiliate locations can be adjusted, but there is no entry or exit of firms, and households cannot relocate across countries. We find the new short run equilibrium with the harmonized tax rate while holding the measures of labor and firms born in each country constant to the calibrated values. Then in the long run equilibrium, free entry and free mobility allow the measures of workers and firms in each country to adjust endogenously.¹⁶

¹⁶The long run counterfactual equilibrium satisfies the full set of equilibrium conditions in equations (3.1) to (3.5). The short run equilibrium solves equations (3.3) to (3.5) while holding employment L_i and firms M_i constant at the levels from the calibrated equilibrium.

In our counterfactual, harmonization leads to an increase in the total aggregate level of tradables production in the EU. Holding firms and workers constant in the short run, total production increases by 0.01%. Once firms and workers are allowed to move in the long run, we find a more sizeable increase in tradable production of 0.32%. While modest, the increase in equilibrium production after harmonization suggests that the original calibrated equilibrium had some misallocation due to tax dispersion. In the following sections, we discuss how harmonization differentially affects countries depending on the change in their tax rate.

3.5.1 The Spatial Distribution of Production

The harmonization of tax rates to 27.26% results in a reallocation of production in the tradables sector from previously low-tax countries that raised their tax rate during harmonization toward previously high-tax countries that now have lower rates. Figure 3.3 shows the percent change in tradables production shares by country in both the short run (in blue) and long run (in red) equilibria against the change in the tax rate after harmonization.¹⁷

In the short run, there is a modest change in the shares of tradable production that is decreasing in the change in tax rates. Firms cannot enter or exit, and instead, we see a reallocation of production in the short run through MP in both the extensive and intensive margins. Figure 3.4a shows the percent change in the number of foreign affiliates from i in host country j over the change in tax rate in j , while figure 3.4b shows the percent change in multinational production flows.¹⁸ Firms place more foreign affiliates and increase multinational production in the countries that lowered

¹⁷For all graphs that display both the short run and long run variables, the short run is plotted in blue while the long run is plotted in red. Values are percent changes relative to the calibrated equilibrium with tax dispersion.

¹⁸All changes are percentage changes in the relevant variable relative to the calibrated equilibrium. The figure leaves out the own-values $PLANTS_{ii}$, the number of domestic plants in each country, and MP_{ii} , domestic production. Both panels are plotted against tax rate changes in the host country where production occurs.





Figure 3.3: Tradable production shares after harmonization

their tax rate. Despite short run changes in MP flows of up to 500%, the overall reallocation of production across countries is small because multinational production makes up only a small share of the tradables production in the calibrated model.



(a) Foreign affiliate plants

(b) Multinational production flows

Figure 3.4: Changes in MP extensive and intensive margins after harmonization

In the long run firm entry adjusts, and we see a sizeable increase in the number of firms born in the countries that lower their rate because the decrease in the tax rate increases expected value of entering for a given core productivity. The effect is apparent in figure 3.5, which plots the long run percentage change in the firms entering in a country against the change in tax rate there. Reallocation of production

across countries occurs through increased domestic production in previously high-tax countries. In the long run, the change in domestic production seems to crowd out changes in MP, perhaps because foreign multinationals face increased competition from the expanded domestic market in previously high-tax countries. Figure 3.4 shows little effect of harmonization on multinational production in the long run.




Figure 3.5: Long run firm entry

The reallocation of production across countries after tax harmonization also has implications for the patterns of trade. Figure 3.6a shows changes in domestic absorption while figure 3.6b shows the changes in international trade flows, both displayed against the change in tax rate in the country of production. The increased long run production in previously high tax countries is partially absorbed by the domestic market and partially exported, while previously low-tax countries see a decrease in their exports and domestic absorption.

3.5.2 Employment and Wages

Tax harmonization induces changes to the pattern of production across countries, which also has general equilibrium effects on employment and wages, especially in




Figure 3.6: Changes in trade patterns after harmonization

the long run. Figure 3.7 shows the reallocation of long run employment in percentage changes, which declines in the country's change in tax rate. The largest magnitude effect is in Cyprus, where employment drops by more than 3%. With free labor mobility in the long run, employment is shifted from previously low-tax countries to previously high-tax countries, although the magnitude of the effect is smaller than that on firm entry.




Figure 3.7: Long run employment

As a result, wages in previously high-tax countries increase. Both the nominal wage (figure 3.8a) and the real wage (figure 3.8b) are plotted as percentage changes

from the calibrated equilibrium. In the short run when there is low reallocation in tradables production and no change in employment or firm entry across countries, there is little change in wages. In the long run, however, both real and nominal wages are decreasing in the change in tax rate after harmonization. The increase in firms and production in previously high-tax countries results in increased competition in those labor markets, and thus a rise in wages. The same factors moving in the opposite direction for previously low-tax countries lead to a decrease in real and nominal wages there. At the extreme, real wage drops by 3.7% in the previously lowest tax country, Bulgaria.




Figure 3.8: Changes in wages after harmonization

3.5.3 Tax Revenue and Worker Welfare

In both the short and long run, tax revenue collected in each country closely tracks the change in tax rate (see figure 3.9a for the percentage change in tax revenue collected in each country relative to the calibrated equilibrium). Intuitively, countries that decreased their rate see lower tax revenue, while those that increased their rate see higher tax revenue. The changes in tax revenue are larger in magnitude in the short run, when firms are unable to enter or exit. The long run effect is somewhat offset by the reallocation of firms and production toward the countries that decreased their tax rate.




Figure 3.9: Changes in tax revenue and welfare after harmonization

Figure 3.9b shows changes in short and long run worker utility in each country after harmonization. Utility is equal across all countries in the long run due to labor mobility, but in the short run, there are differential effects on worker welfare across countries. The dominant short run effect driving changes welfare is the significant change in tax revenue, which we assume are rebated to workers in the country they are collected. This results in welfare changes in the same direction as the changes in taxes: short run welfare gains for previously low-tax countries and welfare losses for previously high-tax countries. At the upper extreme, workers in Bulgaria see a 1.25% increase in welfare in the short run after harmonization.

We find equalized utility is virtually unchanged in the long run harmonized counterfactual compared to the calibrated equilibrium. The low magnitude of the welfare effect may be due to a few reasons. First, the calibrated equilibrium has low multinational production relative to total production in the tradables sector.¹⁹ Because much of the production in the economy is domestic, the location distorting effects of tax dispersion are likely small to begin with. If there is little spatial misallocation from tax dispersion in the first place, then harmonization will not produce large welfare effects. Second, the model includes multiple distortions: the size of the tax wedge in each country, tax dispersion across countries, and the markup on tradable

¹⁹See Appendix C for a comparison of MP shares in the calibrated model and data.

goods all potentially distort the allocation of production. In general equilibrium, it is difficult to separate the effects of the other potential sources of misallocation on utility from the outcome of eliminating tax dispersion. Additionally, our measure of utility is sensitive to the functional form chosen in the model, and in particular, to our simplifying assumption that tax revenue is distributed to workers in the country of collection.

In the context of the political discussions regarding EU tax harmonization, the idea has generally been supported by the higher tax countries, such as France and Germany, but opposed by others that have set lower taxes, such as Ireland and Hungary. While our welfare results show an increase in short run worker utility for the low-tax countries, the utility measure is subject to the limitations discussed above. More importantly, previously low-tax countries would lose on employment, wages, and the ability to attract firms in the long run as production is reallocated to those previously high-tax countries that lower their tax rate during harmonization. These are perhaps the more relevant metrics that low-tax countries consider in the tax harmonization discussion.

Another consideration for low-tax countries might be potential endogeneity of productivity. As we saw in figure 3.2, countries with lower productivity as measured by value added per worker also tend to have lower statutory tax rates. The same pattern holds between statutory tax rates and our calibrated country-specific productivity parameters.²⁰ One story that could explain the correlation between productivity and tax rates endogenously is that low productivity countries may use a low tax rate to attract firms in the hopes of improving their productivity, through learning by doing or technological transfers from FDI. Tax harmonization at the rate used in our coun-

²⁰The country-specific productivity parameters in our model are the core productivity scale z_i^{min} of the Pareto distribution and the productivity shifter T_j . See Appendix C for more details on how they vary with tax rates.

terfactual would be counter to this purpose, as it reallocates firms and production away from the originally low-tax, low productivity countries.

3.6 Conclusion

In the presence of taxation, a firm’s decision of where to produce depends on the costs and productivity in each country as well as the tax rate. In both the data and our calibrated model, we find a positive correlation between a country’s productivity and its statutory tax rate.²¹ This relationship could potentially result from an endogenous choice by low productivity countries to attract firms, through either domestic entry or placement of foreign affiliates. This can benefit countries by increasing employment and wages or by generating endogenous increases in productivity through learning by doing or technology and productivity spillovers from FDI. With harmonized taxes, these countries would lose the ability to use their tax rate differential to attract firms and production.

To better understand these effects of tax dispersion on the spatial allocation of economic activity, in this paper we compare a general equilibrium calibrated to the EU in 2014 to a counterfactual of tax harmonization. We make three contributions: First, we build a model of multinational production with corporate taxation where firms endogenously choose a set of countries in which to set up foreign affiliates. Second, this endogenous location choice is a computationally intensive combinatorial discrete choice problem, which we solve using a novel algorithm. Compared to previous work, we are able to solve the firm’s problem for a large set (27 in our model) of potential locations. Finally, we calibrate our model to aggregate EU data and run a quantitative counterfactual to gain insight into the potential effects of tax harmonization.

²¹Appendix C shows the correlation between calibrated country-specific productivity parameters and 2014 tax rate.

We conduct a counterfactual exercise of harmonized taxes within the EU to a common rate that would generate the same total tax revenue as in the current calibrated equilibrium. We examine two versions of equilibrium: a short run case holding the distribution of employment and firms constant across countries and a long run case with free labor mobility and free entry of firms. In both, harmonization leads to an increase in overall production as well as a reallocation of production from the countries that have raised their tax rate to those that have lowered it.

Holding firms and workers constant, production flows toward previously high-tax countries through multinational production both in terms of the number of affiliates placed there and the volume of production. The magnitude of the reallocation of production is limited by the immobility of workers and firms in the short run, so the biggest response is in tax revenue, which tracks closely with the change in tax rate in each country. Because we assume tax revenue is rebated to workers in the country where it is collected, the change in tax revenue is also reflected by changes in worker utility. Previously low-tax countries see welfare gains for its workers, while previously high-tax countries see welfare losses.

In the long run equilibrium, we find greater entry of firms into the previously high-tax countries, accompanied by a reallocation of production in the same direction. The distribution of workers also shifts towards previously higher tax countries, although to a lesser extent than firm entry. As a result wages changes, both real and nominal, are decreasing in the tax rate change of a country. Multinational production returns to levels closer to the calibrated equilibrium. In the long run, the ability to relocate production natively to previously high-tax countries through firm entry supplants the need for increased MP. The reallocation of production offsets some of the effect on tax revenue of the rate changes after harmonization. Interestingly, we find the long run equalized welfare is virtually unchanged from the calibrated equilibrium.

Overall, our results suggest large distributional effects of harmonization in the long run. Currently low-tax countries stand to lose in terms of employment, firm entry, wages, and share of global production. However, they would enjoy increased tax revenue which may contribute to welfare gains in the short run. Because countries with high tax rates also tend to be more productive, the removal of the effect of tax dispersion on the spatial distribution of production leads to an increase in total production.

Appendix A

India's Annual Survey of Industries

A.1 Data Variables

The variables in the ASI that are used to estimate productivities include employment in mandays, net capital stock, quantities and prices of material inputs, and total expenditure on each of these inputs, which are recorded at the establishment-level, and the sale prices and physical quantities manufactured for each product listed in the establishment's set. All financial values are deflated to constant 2005 Rupees using a national GDP deflator.

I aggregate product-level price and quantity data on materials to a single establishment-level measure. The Cobb-Douglas materials aggregator computes an index for materials as:

$$M_{ft} \equiv \frac{\sum_h W_{ft}^h M_{hft}}{\prod_h \left(\frac{W_{ft}^h}{\gamma_{hft}} \right)^{\gamma_{hft}}}$$

where h indexes an individual input product, $W_{ft}^h M_{hft}$ is the expenditure on input h , W_{ft}^h is the price of input h , and $\gamma_{hft} = \frac{W_{ft}^h M_{hft}}{\sum_{h'} W_{ft}^{h'} M_{h'ft}}$. The index aggregates the individual input products, weighting by their prices and share of total expenditure. The capital stock K_{ft} is measured as the net closing value of fixed assets and is deflated using the annual capital deflators from Allcott et. al (2005) [2]. The quantity of labor

is total mandays worked for all employees, and labor expenditures are inclusive of wages, bonuses, and firm contributions to worker funds.

From 2000 to 2002, plants provide information on their top ten products. After 2002, the ASI survey allows establishments to list additional products. Some establishments list multiple products using the same product code. For the purposes of this study, I aggregate variables within a single 5-digit code. The scope, or number of products, of an establishment is the number of unique 5-digit codes produced.

For the estimation, I drop entire establishment-years with missing or negatives values for inputs or any of its products. It is necessary to remove an establishment from the estimation even if only some of its products have incomplete data because otherwise the input allocation across all of the products would be affected. The same is done for establishments in which the product-level material input data is missing crucial values. For a few 2-digit industries where there were not enough observations to estimate the production function, I dropped any establishments that produced at least one product from the industries. I also drop establishment-year observations that were in operation for less than the full year (12 months).

A.2 ASI Commodity Classification

An example of the aggregation at the 2-, 3-, and 5-digit levels:

78: Electronics equipment & parts, excluding bio-chemical

781: Telephone/telecommunications/transmission equipment

782: Audio/video/sound apparatus and parts

783: Computers and computer equipment, peripherals, & parts

78304: Computers

78313: Hard disc drives

78314: Keyboards

78332: Software

Descriptions of the 1-digit and 2-digit codes are included the Estimation Appendix. Table B.2 includes the 1-digit sectors and table B.1 identifies the 2-digit industries for which I had sufficient observations to estimate production function parameters.

A.3 Dereserved Products in the ASI

Year	Number dereserved in year	Cumulative dereserved
2001	8	8
2002	1	9
2003	28	37
2004	15	52
2005	15	67
2006	14	81
2007	46	127
2008	27	154
2009	0	154
2010	1	155

Table A.1: Number of products dereserved each year relative to year 2000

The number of products that have been dereserved each year are counted out of the number of products that had reservation status in 2010. I do not distinguish between products that were once reserved but were removed from the list prior to 2000 and products that were never subject to reservation because the change in reservation status is what identifies within-industry scope. The instrumental variable used is the log of the cumulative number of dereserved products in the industry of the left-hand side product.

Appendix B

Productivity Estimation

B.1 Estimated Factor Intensities

Table B.1 contains the estimated production function parameters for the 36 industries used in the analysis. Note that I needed to observe a plant being single product and producing the same thing for at least 2 consecutive years to use it in the estimation because one year lagged variables are used as instrument. For this reason, as well as observing missing values for some essential variables, much fewer observations were used in the production function estimation than are available in the full dataset and than I ultimately estimate productivities for. Some industries had very few usable observations and had to be dropped from the analysis.

Industry	Observations	$\hat{\beta}_L$	$\hat{\beta}_M$	$\hat{\beta}_K$	RTS
11 Animal Products	327	0.544	0.244	0.371	1.159
12 Vegetable Products	566	0.290	0.311	0.508	1.109
13 Other Food	611	0.423	0.494	0.116	1.034
15 Beverages and Tobacco	2214	0.588	0.243	0.163	0.995
21 Ores	602	0.646	0.119	0.220	0.985
23 Mineral Fuels	246	0.330	0.478	0.050	0.859
29 Non-Metallic Minerals	303	0.581	0.076	0.118	0.774
31 Inorganic Chemicals	299	0.902	0.009	0.453	1.365
34 Fertilizers and Pesticides	154	1.062	0.194	0.194	1.450
35 Dyes, Inks, and Paint	197	0.921	0.305	0.350	1.576
36 Cosmetics and Detergents	283	0.556	0.240	0.022	0.818
37 Explosives and Combustibles	1506	0.504	0.277	0.088	0.870
41 Rubber	413	0.752	0.134	0.333	1.217
42 Plastic	538	0.516	0.261	0.241	1.019
43 Hides and Leather	129	0.308	0.621	0.738	1.668
44 Leather Products	318	0.502	0.492	0.149	1.142
51 Wood Products	354	0.656	0.254	0.093	1.003
55 Paper	533	0.506	0.319	0.191	1.016
56 Printing	514	0.460	0.342	0.007	0.809
57 Packaging Materials	399	0.568	0.190	0.113	0.871
62 Silk and Wool	110	0.437	0.258	0.110	0.805
63 Cotton	1990	0.270	0.406	0.217	0.893
64 Synthetic Textiles	439	0.140	0.356	0.426	0.922
65 Other Fabrics	490	0.497	0.217	0.361	1.075
71 Iron and Steel	1588	0.326	0.572	0.248	1.145
73 Aluminum, Tin, and Lead	306	0.473	0.487	-0.248	0.712
74 Other Base Metals	336	0.543	0.295	-0.184	0.653
75 Tools and General Purpose Machines	420	0.596	0.113	0.355	1.064
76 Industry-Specific Machinery	297	0.193	0.476	0.103	0.773
77 Electrical Machinery	894	0.602	0.211	0.473	1.286
78 Electronics	178	0.396	0.486	0.282	1.164
82 Vehicles, Aircraft, and Ships	250	0.175	0.335	0.747	1.257
92 Jewelry and Decor	479	0.060	0.547	0.150	0.757
93 Athletic Goods and Toys	142	0.454	0.170	0.294	0.917
94 Glass, Ceramics, and Stone	1178	0.592	0.263	0.161	1.016
95 Other Manufactured Items	42	0.522	0.139	0.347	1.008

Table B.1: Estimates of β

B.2 Estimated Elasticities of Substitution by Product Group

ASICC 1	$\hat{\sigma}_g$
1 Agriculture	2.461
2 Ores and Minerals	3.418
3 Chemicals	1.732
4 Rubber, Plastic, and Leather	13.862
5 Wood, Cork, and Paper	3.190
6 Textiles	19.387
7 Metals and Machinery	2.002
8 Transportation Equipment	1.340
9 Other Manufacturing	1.563

Table B.2: Estimates of σ_g

B.3 Over-Identification Check of the Model's Assumptions

As briefly discussed in Section 1.4.3 of the main body of the paper, summing across products within an establishment of the first-order conditions of the cost minimization problem provides two ways of computing the final estimates of productivity, which can be used as an over-identification check of some of the model's assumptions.

$$W_{ft}^X X_{ft} = \sum_{j \in \mathcal{J}_{ft}} \hat{\lambda}_{jft} \hat{\beta}_X^j Y_{jft} \quad (\text{B.1})$$

The left-hand side of equation (B.1) is simply the establishment's total expenditure on input X , which is observable in the data. On the right-hand side, the expression is a weighted sum of product-specific marginal cost times output, weighted by the output

elasticity of the product with respect to input X . Only the right-hand side uses estimated values β_X^i s and λ_{ift} s. If the assumptions of the Cobb Douglas functional form of production, no shared use of inputs, and those supporting the production function and marginal cost estimation steps are all satisfied, then the two sides should be exactly equal. Of course, it is unlikely for the equation to hold exactly even if all assumptions are true due to sampling error in the estimation of the output elasticities and marginal costs. However, it can still be used as the basis for an over-identification test.




Figure B.1: Data vs model expenditures

Figure B.1 plots establishment input expenditure in the data against the model implied right-hand side expression in a log scale for labor and materials. Both inputs show strong positive relationship between expenditures in the data and implied by the model, but the fit appears better for labor expenditure. The weaker fit for materials may be due to error introduced by the materials aggregator or differences in measure-

ment error. From the binned plot in figure B.2, the model-implied expression using the estimated parameters has a tendency to overestimate establishment-level labor expenditure and underestimate establishment-level materials expenditure relative to the data.




Figure B.2: Data vs model expenditures, binned

The correlation between productivities computed in the two ways, using either input prices or model implied expression, is 0.894. The correlation using log demeaned productivities is 0.9589.

Appendix C

Tax Dispersion and MP

C.1 Statutory Corporate Tax Rates in 2014

Table C.1 shows the top marginal statutory tax rate in each EU country included in our analysis. Note that we omit Malta because of its small size and limited data.

C.2 Tax Rates and Calibrated Productivity

Figure C.2 shows calibrated values of the Pareto core productivity scale parameter by 2014 tax rate. Figure C.3 shows the calibrated values of the country-specific productivity shifters.

Consistent with what we found for value added per worker in the data, both calibrated productivity parameters (the productivity of firms entering in a country z_i^{min} as well as the productivity of firms foreign or domestic that produce in a country T_j) are positively correlated with the statutory tax rate.

Member state	Tax rate
Austria	25%
Belgium	33.99%
Bulgaria	10%
Croatia	20%
Cyprus	12.5%
Czech Republic	19%
Denmark	23.6%
Estonia	20%
Finland	20%
France	33.33%
Germany	30%
Greece	26%
Hungary	19%
Ireland	12.5%
Italy	31.4%
Latvia	15%
Lithuania	15%
Luxembourg	22.47%
Netherlands	25%
Poland	19%
Portugal	21%
Slovakia	22%
Slovenia	17%
Spain	28%
Sweden	22%
United Kingdom	20%

Table C.1: Top marginal corporate tax rates by country




Table C.2: Calibrated core productivity and tax rate

C.3 Over-Identification Tests of the Calibrated Model

We check the calibration by examining the correlation between the model and data for three untargeted moments: equilibrium wages, trade shares, and multinational production shares. The correlations are presented in table C.4.

The equilibrium wages solved in the calibrated model are highly correlated with country-level wage data from Eurostat. Note that this correlation does not include Croatia, for which wage data was missing. The correlation between model and data trade shares is also high. However, our calibrated model does not multinational production shares well. From figures C.5 and C.6, it is apparent that the model predicts too much trade and not enough MP. For most country pairs, our calibrated model predicts very little MP as a share of the country's production. This likely limits the impact we find of tax dispersion on the spatial allocation of production




Table C.3: Calibrated productivity shifter and tax rate

Moment	Correlation
Wages	0.838
Trade shares	0.796
MP shares	0.149

Table C.4: Matching of untargeted moments in the calibrated model

because we expect multinational production to be the most affected. It may explain why we find virtually no long run welfare effect of tax harmonization.




Table C.5: Over-identification of trade shares




Table C.6: Over-identification of multinational production shares

Bibliography

- [1] Daniel A. Akerberg, Kevin Caves, and Garth Frazer. Identification Properties of Recent Production Function Estimators. *Econometrica*, pages 2411–2451, November 2015.
- [2] Hunt Allcott, Allan Collard-Wexler, and Stephen D. O’Connell. How Do Electricity Shortages Affect Industry? Evidence from India. *The American Economic Review*, pages 587–624, March 2016.
- [3] Pol Antràs, Teresa C. Fort, and Felix Tintelnot. The Margins of Global Sourcing: Theory and Evidence from U.S. Firms. *American Economic Review*, pages 107 (9): 2514–64, 2017.
- [4] Costas Arkolakis and Fabian Eckert. Combinatorial Discrete Choice. *Working paper*, January 2017.
- [5] Costas Arkolakis, Sharat Ganapati, and Marc-Andreas Muendler. The Extensive Margin of Exporting Products: A Firm-level Analysis. *NBER Working Paper Series 16641*, January 2019.
- [6] William Baumol, John Panzar, and Robert Willig. *Contestable Markets and the Theory of Industry Structure*. Harcourt Brace Jovanovich, 1988.
- [7] Arie Beresteanu. Nonparametric Analysis of Cost Complementarities in the Telecommunications Industry. *The RAND Journal of Economics*, pages 36(4): 870–889, 2005.
- [8] Andrew B. Bernard, Emily J. Blanchard, Ilke Van Beveren, and Hylke Vandenbussche. Carry-Along Trade. *Review of Economic Studies*, pages 86: 526–563, 2019.
- [9] Andrew B. Bernard, Jonathan Eaton, J. Bradford Jensen, and Samuel Kortum. Plants and Productivity in International Trade. *The American Economic Review*, pages 1268–1290, September 2003.
- [10] Andrew B. Bernard, Stephen J. Redding, and Peter K. Schott. Multiple-Product Firms and Product Switching. *The American Economic Review*, pages 70–97, March 2010.

- [11] Johannes Boehm, Swati Dhingra, and John Morrow. The Comparative Advantage of Firms. *Working paper*, 2019.
- [12] Thiess Buettner and Martin Ruf. Tax Incentives and the Location of FDI: Evidence from a Panel of German Multinationals. *International Tax and Public Finance*, pages 14: 151–164, 2007.
- [13] Douglas W. Caves, Laurits R. Christensen, and W. Erwin Diewert. The Economic Theory of Index Numbers and the Measurement of Input, Output, and Productivity. *Econometrica*, pages 50(6): 1393–1414, November 1982.
- [14] Jan De Loeker, Pinelopi K. Goldberg, Amit K. Khandelwal, and Nina Pavcnik. Prices, Markups, and Trade Reform. *Econometrica*, pages 445–510, March 2016.
- [15] Michael P. Devereux and Rachel Griffith. Evaluating Tax Policy for Location Decisions. *International Tax and Public Finance*, pages 10: 107–126, 2003.
- [16] Emmanuel Dhyne, Amil Petrin, Valerie Smeets, and Frederic Warzynski. Multi Product Firms, Import Competition, and the Evolution of Firm-Product Technical Efficiencies. *NBER Working Paper Series 23637*, July 2017.
- [17] Jonathan Eaton and Samuel Kortum. Technology, Geography, and Trade. *Econometrica*, pages 70(5): 1741–1779, 2002.
- [18] Carsten Eckel and J. Peter Neary. Multi-Product Firms and Flexible Manufacturing in the Global Economy. *The Review of Economic Studies*, pages 188–217, January 2010.
- [19] Pablo D. Fajgelbaum, Eduardo Morales, Juan Carlos Suárez Serrato, and Owen Zidar. State Taxes and Spatial Misallocation. *Econometrica*, pages 86(1): 333–376, 2018.
- [20] Matthew Flagge and Ritam Chaurey. *Firm-Product Linkages and the Evolution of Product Scope*. PhD thesis, Columbia University; SUNY Binghamton, 2014.
- [21] Lucia Foster, John Haliwanger, and Chad Syverson. Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability? *The American Economic Review*, pages 394–425, March 2008.
- [22] Keith Gilsdorg. Vertical Integration Efficiencies and Electric Utilities: A Cost Complementarity Perspective. *The Quarterly Review of Economics and Finance*, pages 34(3): 261–282, 1994.
- [23] Pinelopi K. Goldberg, Amit K. Khandelwal, Nina Pavcnik, and Petia Topalova. Multiproduct Firms and Product Turnover in the Developing World: Evidence from India. *The Review of Economics and Statistics*, pages 92(4): 1042–1049, 2010.

- [24] Elhanan Helpman, Marc J. Melitz, and Stephen R. Yeaple. Importers, Exporters and Multinationals: A Portrait of Firms in the U.S. That Trade Goods. *The American Economic Review*, pages 300–316, March 2004.
- [25] Elhanan Helpman, Marc J. Melitz, and Stephen R. Yeaple. Export Versus FDI with Heterogeneous Firms. *Producer Dynamics: New Evidence from Micro Data, NBER Chapter*, pages 513–552, 2009.
- [26] Colin J. Hottman, Stephen J. Redding, and David E. Weinstein. Quantifying the Sources of Firm Heterogeneity. *Quarterly Journal of Economics*, pages 1291–1364, July 2016.
- [27] Chang-Tai Hsieh and Peter J. Klenow. Misallocation and Manufacturing TFP in China and India. *The Quarterly Journal of Economics*, pages 124(4): 1403–1448, 2009.
- [28] Atsushi Iimi. Banking Sector Reforms in Pakistan: Economies of Scale and Scope, and Cost Complementarities. *Journal of Asian Economics*, pages 15: 507–528, 2004.
- [29] Panle Jia. What Happens When Wal-Mart Comes to Town: An Empirical Analysis of the Discount Retail Industry. *Econometrica*, pages 76(6): 1263–1316, 2008.
- [30] James Levinsohn and Amil Petrin. Estimating Production Functions Using Inputs to Control for Unobservables. *The Review of Economic Studies*, pages 317–341, April 2003.
- [31] Leslie A. Martin, Shanthi Nataraj, and Ann E. Harrison. In with the Big, Out with the Small: Removing Small-Scale Reservations in India. *The American Economic Review*, pages 107(2): 354–386, 2017.
- [32] Thierry Mayer, Marc J. Melitz, and Gianmarco I. P. Ottaviano. Market Size, Competition, and the Product Mix of Exporters. *American Economic Review*, pages 104(2):495–536, 2014.
- [33] Marc Melitz. The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity. *Econometrica*, pages 1695–1725, November 2003.
- [34] Rakesh Mohan. *Small Scale Industry Policy in India : A Critical Evaluation*. New Delhi : National Council of Applied Economic Research, 2001.
- [35] Volker Nocke and Stephen Yeaple. Globalization and Multiproduct Firms. *International Economic Review*, pages 993–1018, November 2014.
- [36] Steven G. Olley and Ariel Pakes. The Dynamics of Productivity in the Telecommunications Equipment Industry. *Econometrica*, pages 1263–1297, November 1996.

- [37] Scott Orr. *Productivity Dispersion, Import Competition, and Specialization in Multi-product Plants*. PhD thesis, University of Toronto, 2017.
- [38] Lawrence Pully and David Humphrey. The Role of Fixed Costs and Cost Complementarities in Determining Scope Economies and the Cost of Narrow Banking Proposals. *The Journal of Business*, pages 66(3): 437–462, 1993.
- [39] Natalie Ramondo and Andrés Rodríguez-Clare. Trade, Multinational Production, and the Gains from Openness. *Journal of Political Economy*, pages 121(2): 273–322, 2013.
- [40] Rowan Shi. Quickly optimizing over a lattice. https://github.com/rowanxshi/cdc_solver/blob/master/algorithm.pdf, 2019.
- [41] Chad Syverson. Product Substitutability and Productivity Dispersion. *Review of Economics and Statistics*, pages 86(2): 534–550, 2004.
- [42] Felix Tintelnot. Global Production with Export Platforms. *The Quarterly Journal of Economics*, pages 157–209, 2017.