How Substitutable Are Labor and Intermediates?*

Mons Chan†
Queen's University
Aarhus University

January 2021

Abstract

Empirical models of production often impose input complementarity and rule out an extensive margin in the decision to "make or buy" inputs. This paper develops a simple model of production which generalizes the standard Cobb-Douglas approach and allows labor and intermediates of similar types (or "tasks") to be complements, substitutes, or (importantly) outsourced entirely. Modeling this make-or-buy decision directly allows me to correct for the selection bias resulting from the endogenous outsourcing decision and to characterize the extensive margin of factor demand. I take the model to unique Danish data on task-level purchases of disaggregated labor (e.g., truck drivers), goods, and services (e.g., shipping) and find that labor and intermediates are gross substitutes. Estimated elasticities of substitution range from 1.5 to 4, with positive cross-price elasticities between 0 to 2 across inputs and industries. These results also hold in standard firm data using total labor and intermediate expenditure variables. Aggregating across firms, I show that demand for labor is becoming increasingly price elastic over time, driven by growing outsourcing and specialization. To illustrate the importance of allowing for flexible substitutability, I examine the effect of an increase in minimum wages in the Danish manufacturing industry, finding that ignoring outsourcing underestimates disemployment by 40%. This finding also has important implications for estimating productivity. I estimate the effect of recent decreases in Danish import tariffs on firm productivity and show that controlling for substitution triples the results relative to benchmark models which only control for price effects.

*I would like to thank Naoki Aizawa, Anmol Bhandari, Fatih Guvenen, Kyle Herkenhoff, Tom Holmes, Loukas Karabarbounis, Jeremy Lise, Ellen McGrattan, Amil Petrin, Joel Waldfogel, and Frederic Warzynski for their guidance, support, and suggestions. This paper has benefited from many useful comments and discussions with participants at the 2016 CEA, 2017 DIEW, 2018 IIOC, 2018 ETSG, and 2018 IZA World Labor Conferences, as well as multiple seminar presentations. I also thank the Department of Economics and Business at Aarhus University for support and making the data available. This paper was awarded the "Best Rising Star Paper" prize by the IIOC in 2018.

†Correspondence: mons.chan@queensu.ca

1 Introduction

Economists have been interested in estimating the technological relationship between inputs and outputs in production since at least the 19th century.[1] Ongoing improvements in the availability of firm-level data and estimation methodology have led to an explosion in papers that use these *production functions* as a building block for empirical studies of firm behavior. Despite these recent innovations, the existing methodological frontier still embodies restrictions on the data that may not hold. In particular, almost all empirical research done on production imposes that inputs are complements – cheaper intermediates increase demand for labor – and rules out extensive margin decisions to "make or buy" inputs.[2]

This paper uses a unique data set from Denmark to test and relax these restrictions. The data include the universe of labor inputs, occupations, and wages at the worker level, as well as detailed firm-level expenditure information for disaggregated intermediate goods and services. I group labor and intermediates by input "task" and show in the raw data that some firms are on the extensive margin and others are on the intensive margin. For example, to complete the task of transporting output to customers, some firms employ truck drivers, some outsource their shipping, and others do both. The choice of "make," "buy," or "both" differs across firms within very narrow industries. This suggests that different input tasks within a firm may be produced using different levels of technical efficiency, with comparatively low efficiency firms choosing to outsource. On aggregate, I show that firms are increasingly choosing to outsource these tasks in a manner consistent with a focus on core competencies.

I develop a simple model of task-based production which generalizes the standard Cobb-Douglas approach and rationalizes these patterns in the data. Firms employ a set of input tasks to produce a differentiated product, where each task can be provided in-house by labor or purchased on the market from another firm (or both). Each of these choices is associated with a fixed cost that drives the extensive margin decision. I allow labor and intermediates to be substitutes or complements, with substitutability varying by task. My model accounts for input mix heterogeneity by allowing each firm to have a vector of task-specific efficiency terms. Firms will outsource tasks if they face high wages relative to intermediate prices or if they are not very efficient at in-house production of that task.

[1] See Humphrey (1997) and notably Cobb and Douglas (1928) for early examples and a discussion of 19th-century precursors to the modern production function.

[2] Under profit-maximization, Cobb-Douglas functions impose both of these restrictions, while trans-log functions impose the latter.

To map the data to the model, I develop a simple clustering algorithm, which I use to assign goods, services, and occupations to tasks. Truck drivers and shipping services are both assigned to the transportation task because they are both disproportionately employed or produced by the transportation industry. The resulting mapping differs from standard labor aggregators in that occupations are not grouped by skill but rather by the input task they perform. Logistics professionals are grouped with freight handlers in the transportation task despite differences in skill.

A key contribution of this paper is to show how to estimate a generalized production function featuring flexible substitution and make-or-buy decisions over multiple inputs. I rely on a separability result to estimate the model in several stages. Since the elasticity of substitution between labor and intermediates can only be estimated using firms that employ both, the fact that firms may choose no labor or no intermediates introduces selection bias. I address this bias by modeling and estimating the discrete make-or-buy decision jointly with the intensive margin. This framework then allows me to estimate and characterize the intensive and extensive margins of factor demand, which I believe to be unique in the production estimation literature.

The estimated model then provides new estimates of labor demand and substitution elasticities. Relative to results in other papers, my estimates are much higher and more disaggregated. In particular, I estimate elasticities of substitution between labor and intermediates for a set of 13 different input types across 4 different industries. I find that labor and intermediates are gross substitutes, with elasticities ranging from 1.5 to 4,[3] strongly rejecting the Cobb-Douglas benchmark. Similarly, I find that cross-price elasticities of demand between labor and intermediates are positive and range from 0 to 2 at the firm-task level. Constructing a sequence of aggregate price elasticities, I show that demand for labor has become increasingly price elastic over time because of growth in outsourcing and specialization.

Since few firm-level data sets include information on disaggregated input use, I also show how to take the flexible framework to standard data with aggregate expenditures on labor and intermediates. Under some conditions, this aggregate input framework can be estimated easily in log-linear form. I demonstrate that allowing for flexible substitution in this simple way, even with standard data, makes a big difference when estimating firm productivity.

[3]This finding differs significantly from the literature that frequently either finds that labor and intermediates are complements or restricts these substitution elasticities to 1. For example, Doraszelski and Jaumandreu (2018) and Oberfield and Raval (2021) both find firm-level elasticities less than 1 using aggregated labor and intermediates.

Estimated elasticities of substitution between aggregate labor and intermediate indices are of similar sign and magnitude to the results for the disaggregated task model.

These results suggest that failing to allow for flexible input substitution and outsourcing may lead to biased results in empirical studies that rely on production function estimation. I illustrate this point by performing two empirical exercises, comparing results from my framework to benchmark models that do not include these features. First, I examine the effects of an increased minimum wage (wage floor) in the Danish manufacturing industry. The key contribution here is not only in being able to characterize flexible substitution patterns across intermediate and labor types, but also in being able to calculate the probability that any given firm outsources. I conduct a policy experiment where I raise the wage floor by 25 kroner ($4 USD) in 2011. The result is a total decrease in labor demand of 4.2%. The distribution of decreases across occupations ranges from 1.2% to 22.5%. I show that failing to account for substitution and extensive margin outsourcing underestimates these effects, sometimes by over 50%.

The second application is an extension of the literature that looks at the effects of competition and trade policy on productivity. I use a decrease in tariff protection for Denmark during the early to mid-2000s to estimate the effect of tariff reductions on productivity. De Loecker (2011) shows that failing to control for unobserved price effects when estimating productivity leads to an overestimate of the effects of trade protection on firm efficiency. I build upon this result and show that failing to control for input substitution actually biases results in the opposite direction, leading to an underestimate of the effects of trade protection on efficiency. When controlling for only the price effect, a removal of all tariffs leads to a productivity increase of 2.6%. When additionally controlling for input substitution, the estimated increase is almost three times higher at 6.9%. This exercise demonstrates that allowing for flexible substitution between labor and intermediates is important even in contexts where researchers do not have access to data on disaggregated inputs.

My paper draws from and contributes to several major strands of the literature. First, I make particular use of the proxy function methods developed by Olley and Pakes (1996), Levinsohn and Petrin (2003), and Wooldridge (2009). My extension to the production function estimation literature builds on recent work by Gandhi, Navarro and Rivers (2020), Doraszelski and Jaumandreu (2013), and Doraszelski and Jaumandreu (2018) in using the empirical content of the structural model to estimate model parameters. In relation to these papers, I develop a method for estimating a much more disaggregated model while controlling for outsourcing and selection bias. I draw on the extensive literature following Heckman

(1979) on correcting for selection bias. My approach differs from most in that I apply the selection correction to a structural model of production and estimate the selection problem jointly with the main equation. My paper deals with controlling for particular features of the data that may complicate estimates of productivity, and thus is closely related to an empirical strand of the trade and productivity literature typified by De Loecker (2011) and more recently Halpern, Koren and Szeidl (2015), Dhyne et al. (2017), and De Loecker et al. (2016). I am also methodologically akin to the recent literature on the extensive margin in trade, and in particular, to Helpman, Melitz and Rubinstein (2008), who use a similar approach to estimate the extensive margin of trade flows. My paper lies solidly amid the literatures on misallocation and multi-worker firms (Bagger, Christensen and Mortensen (2014)), estimating input elasticities (Senses (2010), Oberfield and Raval (2021), Raval (2019)) and the wage effects of outsourcing (Hummels et al. (2014), Goldschmidt and Schmieder (2017)), drawing significantly from each. Finally, this paper is tied to the vast literature on trade networks and firm-to-firm trade, especially those looking at task-level substitution between labor and intermediates (Chan and Xu (2017), Eaton et al. (2016)).

The paper proceeds as follows. In section 2, I introduce my data and provide evidence of heterogeneity and trends in input mix across firms. Section 3 discusses the difficulties faced by standard models of production. In section 4, I present an alternative model of production that resolves those difficulties. Section 5 describes the details of how I extend the estimation literature to take the alternative model to the data, with results presented in section 6. Sections 7 and 8 are the two primary applications of the framework: an estimation of the effects of increased wages in Danish manufacturing and the effects of tariff reductions on firm productivity. I conclude in section 9. Tables and figures are located at the end of the paper.

2 Heterogeneity and Trends in Input Composition

This section discusses the data that I use for this study and establishes a few facts that motivate and inform the subsequent model and analysis. In particular, I use data on disaggregated input use at the firm level to show that firms exhibit significant heterogeneity in input mix, at both the intensive and extensive margins. Most firms outsource at least one of their inputs. I show evidence that firms substitute between labor and intermediates of the same type, and that over a period of 11 years, Danish firms have been increasingly outsourcing their inputs, becoming more concentrated and focusing on their core competencies.

These empirical results imply that it is important to think of production at the disaggregated level so that these patterns across firms and over time can be accounted for and analyzed. In particular, I show that the labor share of input expenditure has been decreasing over time, but that this is at least partly due to outsourcing intermediate materials and services rather than substitution to capital.

2.1 Data Description

My primary data source is a register of Danish matched employer-employee (MEE) data collected by Statistics Denmark. I combine several registers for this analysis. The employee data is primarily from the Integrated Database for Labor Market Research (IDA). This individual-level panel contains information on employment status, occupation, wages, hours worked, education and employer for all individuals in Denmark aged 15 and above. My main use for this data is employment status, occupation, hourly/annual wages, hours worked and employer, all of which are recorded for the individual's primary job in November. I match this individual-level panel to a firm-level panel, the Firm Statistics Register, which covers the universe of firms in Denmark. The firm panel contains data on revenues, capital stock, aggregate intermediate expenditures and employment, as well as data on firm industry. By mapping the unique firm and individual identifiers together via the Firm Integrated Database for Labor Market Research (FIDA), I can match workers to firms, giving me the distribution of wages, hours and occupations at the firm-worker level. The high quality of the categorical data in the danish registers is noteworthy and gives me some confidence when applying my analysis to slices of the data defined by occupation and industry (measures which are sometimes noisy in other data)[4].

Given a matched panel of firm-individual level wages, occupations, industries, hours and firm accounting data, I merge in several additional data sets. First I match a subset of manufacturing firms to data on production, which includes quantity, revenue, type and other production data at the HS10 level. This gives me firm-product level measures of physical output and price. I also merge in trade data at the product level, allowing me to see trade flows for products at the HS6 level. Finally, and importantly, I use an additional dataset which includes detailed expenditure and input use data for a subset of large manufacturing firms. In particular, this data includes expenditures on a wide selection of services such

[4] See Hummels et al. (2014) and Bagger, Christensen and Mortensen (2014) for further details on the individual panels and data sets

as cleaning, law services, transportation, storage, ICT, management, etc. It also includes expenditure data on inputs such as energy and water, as well as detailed intermediate product expenditure at the HS6 level. This allows me to construct the input task expenditure measures I use throughout the paper.

2.2 Matching Labor and Intermediates

The main analysis of my paper involves measuring the degree to which labor and intermediates are substitutable. One fundamental idea upon which I build my theory (described in section 4) is that certain types of labor provide certain types of services or goods. For example, if a firm requires law services, they might hire lawyers, or contract the services of an external law firm. They would not, however, hire janitors to provide law services. This simple idea has several important implications. First, it is natural to think that the relationship between lawyers and law services in production is special, relative to, say, the relationship between lawyers and janitors, or lawyers and titanium hinges. While all of these may be fundamental inputs into production, it is much more likely that the firm treats lawyers and law services as strong substitutes (or complements) relative to other labor and purchased inputs.

In line with this idea, I propose a task-matched theory of production in which firms require a set of input tasks, each of which can be supplied by labor of a particular type, and/or purchased services and intermediates of that same type. This implies a mapping between occupations, goods/services, and industries. For example, the law services industry uses some set of occupations (lawyers, etc) to produce law services. Other firms which require law services may buy them from a law firm, or hire those same labor types (lawyers) directly to produce them in-house. Purchased law service intermediates and lawyers are mapped to the same "task".

The problem then is in matching occupations to services/intermediates so that we can be reasonably certain that we have correctly specified the inputs into this task-matched production theory. To do this, I develop a matching algorithm which determines the most likely mapping of occupations to intermediate services/products. The algorithm assigns detailed occupations (at the 4-digit DISCO level) to the industries in which they are most disproportionately employed relative to a measure of predicted employment which is based on overall industry and occupation employment shares. The idea is that the industry which uses a particular occupation to produce its primary output will employ that occupation disproportionately more than industries which use that occupation to produce intermediate tasks. Note that my ability to match inputs to tasks by what they do hinges critically on the fact that I have matched employer-employee data for every firm/worker in Denmark, thereby allowing me to observe the total distribution of labor inputs across industries, firms and occupations. I discuss this matching algorithm and the theory behind it in more detail in appendix A.

As an example, the algorithm determines that the primary occupations for the Transportation & Storage industry are: Transportation Managers, Transport Clerks, Heavy Truck and Lorry Drivers, Crane, Hoist and Related Operators, Messengers, Package Deliverers and Freight Handlers. These are the *primary* occupations which the transportation industry uses to produce its primary output, and so are the same occupations another firm would need to produce transportation services in-house. Of particular interest is the fact that my framework doesn't match labor to intermediates or capital based on skill, but rather on the task-composition of the occupation - what type of output an occupation produces. This gives me a set of occupations which vary in skill, but are all required to produce the aggregate Transportation input. This differs significantly from the existing literature on input substitution which typically disaggregates labor along the skill dimension rather than the task dimension. I follow a similar procedure to match services and product codes to industries, giving me a direct mapping between occupation, output type (intermediate) and industry.

I end up with a set of 13 different input types, meaning I group firms, occupations and intermediate services/goods into 13 categories, including an ”other services” category, as the intermediate expenditure data does not include detailed expenditure data for all service types. Appendix A provides a detailed list and summary statistics for the task grouping. This mapping then represents the empirical *network* of input links between industries which forms the basis of the model which I discuss in section 4.

2.3 Cross-Sectional Heterogeneity in Input Composition

One of the most striking features of the data input usage is the degree of heterogeneity in input mix across firms. Table 1 shows the input choice patterns for the Tools, Machinery and Goods industry across the matched tasks identified in the previous section. Consider the first row. A firm in this industry “Only Hires” if they employ some positive amount of transportation labor (as defined by the matching algorithm) but spend nothing on purchasing transportation services from other firms. In this sample, only 106 firms do all of their transportation services in-house. “Only Buys” is the opposite, with the majority of firms (3,220) choosing to entirely outsource their transportation services – hiring zero transportation labor in-house. Almost all of the remaining firms do some combination of both in-house production and outsourcing, with a few remaining firms appearing to do neither. The remaining input tasks display similar patterns. Most firms outsource some or all of their input production. This has several strong implications. First, thinking about production and productivity using aggregate input indices misses a lot of firm-level heterogeneity which likely affects firm behavior. As I will show later, failing to account for this input heterogeneity can lead to significant biases in estimates of productivity. Second, the fact that firms make extensive margin make-or-buy decisions implies that inputs are substitutes. I will use my structural model to show that labor and intermediates are in fact gross substitutes. However, when estimating elasticities of substitution, failing to account for the extensive margin which is so prominent in table 1 will lead to biased estimates. The key implication is that Cobb-Douglas, or even flexible trans-log production frameworks are inappropriate when wanting to model and estimate production with disaggregated inputs. See section 3 for a discussion of the problems which may occur when using these standard approaches.

To get an idea of how this heterogeneity in input mix is related to other firm characteristics, I briefly step away from the task-matched setting and plot in figure 1 the number of 2-digit occupations employed within a narrow industry (supermarkets) against firm revenue. The x-axis is the number of occupations employed by supermarket firms in 2002. Each circle represents the fraction of total industry establishments belonging to the firms in that category, while the red line plots the 90-10 revenue spread for firms in that category. For example, a large fraction of supermarkets employ only 5 2-digit occupations. This ranges from firms with a log revenue (in thousands of kroner) ranging from below 9.5, up to 11 and higher. The key takeaway from this graph is that larger firms do tend to employ more occupations and be more vertically integrated - a fact which I will exploit in my model and estimation. However, this is not a monotone relationship. There is considerable heterogeneity in the relationship between firm size/revenues and concentration. For example, if we consider supermarkets with a log revenue of roughly 11, this includes firms which employ only 5 or 6 occupations, as well as firms which employ 15 occupations. This suggests that there are other firm-level factors other than size/tfp which drive input mix choice and vertical integration/concentration.

2.4 Changes in Input Composition over Time

As shown above, firms exhibit an abundance of heterogeneity in input usage and internal structure in the cross section. Here I show that these distributions are changing over time as firms respond to changes in prices, productivity and market conditions by adjusting their optimal input mix. While changes over time differ by firm, the overall trend is an increase in concentration and a shift towards the firm's core competencies, which I define as an increasing employment share of *primary* occupations relative to total employment.

These trends can be clearly seen in figures 2 and 3. Figure 2 plots two variables. The first is the average number of occupations employed at the 2-digit level. This has been trending down between 1995 and 2009, from a high of around 15 occupations per firm down to just above 10 in 2009 – a 30% decrease over 15 years. The second is the average occupational Herfindahl index, which is a measure of occupational concentration within firms. A value of 1 means that a firm only employs a single 2-digit occupation. A value close to 0 implies that a firm employs equal shares of many occupations. Over this period, the herfindahl index has increased from about 0.265 to 0.39, a significant growth in within-firm occupational concentration. Note that both of these variables are calculated using a balanced panel of large firms, so that these trends are not due to firm exit or entry, but rather within-firm changes in composition. These calculations are also done with a time-consistent set of occupation codes, so this also does not include changes in occupational definition. Figure 3 shows the change in number of occupations employed in greater detail using the same balanced panel of large firms (defined as firms with > 50 employees in both 1995 and 2009). While a few firms became increasingly vertically integrated over this period and increased the set of occupations employed in-house, the vast majority became more concentrated, shedding two-digit occupations. As argued above and in the model, I propose that this is due to firms switching instead to purchased intermediates.

The final key fact that I describe is that firms are increasingly substituting away from intermediate labor in favor of purchased intermediates and their own primary labor. To show this, I perform a series of regressions of the intermediate-labor ratio (M/L)[5] on year and a set of controls and fixed effects. This results can be seen in table 2, where the M/L ratio for the average firm is increasing by 2% a year, controlling for firm level controls and fixed effects. The ratio of materials to capital and investment is also increasing. Figure 4 plots results of the same regression on year dummies, showing that the M/L ratio has increased

[5]The M/L ratio is deflated intermediate expenditure (M) divided by employed labor inputs (L).

by about 30% relative to 1992, and that it's also slightly pro-cyclical, in that it declined slightly in the great recession. The final two columns show that intermediate expenditure is also growing relative to capital and investment over this same period.

Table 3 shows the results of a similar regression, but this time looking at the ratio of primary labor (H) to total labor (L), where primary labor is the occupation set which is matched to the firm's industry, as described in section 2.2. Here, the ratio of primary labor (H) to total labor has also been increasing by about 3% a year (column 3). Column 6 also shows that the ratio of purchased intermediates (M) to intermediate labor (L-H) has been increasing at 3.8% a year on average. Figure 5 plots the regression in column 3 of table 3 on year dummies, showing that the share of primary labor in total labor has increased by about 60-70% since 1992.

3 Standard Models of Production

In this section I first lay out a general production framework and establish notation. I then review the standard approaches to modeling production and discuss why they may fail to account for the data patterns discussed in Section 2.

3.1 Basic Production Framework

The basic environment is an economy with a set $\mathcal{H}$ of industries and set $\mathcal{L}$ of labor types, where each industry $i \in \mathcal{H}$ consists of J_i firms. Each firm $j \in J_i$ produces some quantity of a differentiated product Y_{ji} using a vector of capital inputs $\{K_{kj}\}_{k \in \mathcal{K}_i}$, a vector of labor inputs $\{L_{\ell j}\}_{\ell \in \mathcal{L}_i}$ and a vector intermediate goods and services $\{Q_{hj}\}_{h \in \mathcal{H}_i}$ purchased from firms in the same or other industries. Here $\mathcal{K}_i$, $\mathcal{H}_i$ and $\mathcal{L}_i$ denote the input sets of capital goods, intermediate goods/services and labor types, respectively, required for production of output in industry i. To fix concepts, the input sets for coffee shops may include espresso machines, coffee beans and retail labor (baristas), but not titanium centrifuges, propellers or aerospace engineers, which are all integral inputs in the aerospace industry. This provides the following general production function:

$$Y_{ji} = F(\{K_{kj}\}_{k \in \mathcal{K}_i}, \{L_{\ell j}\}_{\ell \in \mathcal{L}_i}, \{Q_{hj}\}_{h \in \mathcal{H}_i}) \tag{1}$$

where in principle the function F can also differ across industry, firm and time, and may embody one or more dimensions of unobserved firm-level efficiency.

In the following sections I suppress the industry and time notation unless essential to the exposition. In general, all parameters vary at the industry level, while all inputs, outputs and efficiency terms vary across firm, industry and time. Also, for the duration of the paper, I will use the word "intermediate" to refer strictly to physical and service inputs purchased from external firms, in contrast to labor employed in-house which also acts as an input into production.

3.2 Cobb-Douglas Production Functions

In empirical applications, F_{jit} is often specified as a Cobb-Douglas function of aggregates,

$$F_j(\{K_{kj}\}_{k\in\mathcal{K}_i}, \{L_{\ell j}\}_{\ell\in\mathcal{L}_i}, \{Q_{hj}\}_{h\in\mathcal{H}_i}) = K_j^{\beta_K} L_j^{\alpha_L} Q_j^{\alpha_Q} e^{\omega_j} \tag{2}$$

where ω_j is a Hicks-neutral technical efficiency term. Disaggregated inputs are represented by "input indices" K_j, L_j, and Q_j, where the intermediate index (for example) is commonly represented as the sum of the disaggregated intermediates, weighted by their price (i.e.: total expenditure on all intermediates). This specification is convenient and easy to estimate, but it has several serious shortcomings.

The Cobb-Douglas functional form imposes the assumption that the elasticity of substitution between any two inputs is equal to one. Similarly, the cross-price elasticity of demand for each input relative to the others is strictly negative for profit maximizing firms, implying that all goods are gross complements[6]. For example, the short run change in demand for labor L_j in response to a change in the price of intermediates P^Q for a profit-maximizing price-taking firm,[7]

$$\epsilon_{P^Q}^{L_j} \equiv \frac{\partial L_j}{\partial P^Q}\frac{P^Q}{L_j} = -\frac{\alpha_Q}{1-\alpha_Q-\alpha_L} \tag{3}$$

[6] Of course, these inputs can still be net substitutes in the Cobb-Douglas setting

[7] This expression comes from the fact that labor demand for a price-taking profit-maximizing firm with the production function in 2, output price P and fixed capital is

$$L_j^* = \frac{\alpha_L}{W_j}\left[PK_j^{\beta_K}e^{\omega_j}\left(\frac{\alpha_L}{W_j}\right)^{\alpha_L}\left(\frac{\alpha_Q}{P^Q}\right)^{\alpha_Q}\right]^{\frac{1}{1-\alpha_Q-\alpha_L}}$$

Similar expressions implying gross substitution result in settings where firms set prices as well – see main model.

is negative and independent of relative prices or any firm heterogeneity. This restriction on substitution patterns may not hold in the data[8] and seems at odds with the outsourcing and make-or-buy literature which is predicated on the notion that labor and intermediates may be gross substitutes.

The Cobb-Douglas and other common functional forms such as the trans-log also are unable to rationalize the corner solutions seen in the disaggregated data, as both are undefined if any input is zero. Since the zeros are likely endogenous choices, discarding undefined observations may lead to biased parameter estimates or too few observations.[9] Assuming the elasticity of substitution is one for all inputs may also lead to misspecification bias and spurious measurements of productivity.[10]

These concerns regarding the aggregation and substitution assumptions inherent in the standard Cobb-Douglas framework are not new. The trade and networks literature is often concerned with flows of disaggregated inputs across countries and industries. In this setting, the most common empirical approach is to specify a disaggregated Cobb-Douglas in intermediates and (less often) labor.[11] The nascent literature on the transmission of shocks through network linkages has recently made some steps towards generalizing this framework,[12] but mostly along the lines of flexible substitution across intermediates, or relative to some labor aggregate. The macro labor literature has long been concerned about substitution patterns between labor and capital, leading to significant work in estimating the elasticity of substitution between capital and high/low skill labor. The standard framework here is a CES production function with some nesting to allow flexibility in substitution between aggregate capital and one or more types of labor.[13] Another common approach has been to estimate production parameters and input demand through the lens of a trans-log production or cost function. This allows for very flexible substitution patterns, but is both very demanding on the data, and as mentioned above is similarly undefined if any input is zero.

[8]The Cobb-Douglas assumption has been challenged many times both at the aggregate/macroeconomic level (see Antras, 2004) and at the firm or micro level (recently by Doraszelski and Jaumandreu, 2018)

[9]In the Danish data, zero firms employ both labor and intermediates for every task in table 1.

[10]See appendix D for a discussion and example.

[11]Long and Plosser (1983) is a canonical example. See also Acemoglu et al. (2012)

[12]See Carvalho et al. (2021) for a recent example.

[13]See Krusell et al. (2000) and the "canonical model" in Acemoglu and Autor (2011)

4 A Model of Production with Disaggregated Inputs and Outsourcing

This section develops a tractable framework which addresses the facts shown in section 2 and the issues with standard approaches discussed in section 3. In particular, a primary purpose of this paper is to develop a tractable framework for estimating production functions which accounts for 1) firm-level heterogeneity in input mix, 2) extensive-margin outsourcing, and 3) flexible substitution patterns between disaggregated inputs. With these goals in mind, along with the need to balance flexibility with tractability, this paper develops a "task-matched" CES framework which is a straightforward generalization of the disaggregated Cobb-Douglas framework common in the literature.

4.1 Production

The basic environment is the same as specified in section 3.1, with the key difference that each firm requires an industry-specific set of input "tasks" $\mathcal{H}_i$ (with cardinality H_i) which correspond to the output of different sectors $h \in \mathcal{H}$ in the economy. I specify the physical production function for firm j in industry i in period t in the following way:

$$Y_{jt} = K_{jt}^{\beta} \prod_{h \in \mathcal{H}_i} M_{hjt}^{\alpha_h} e^{\omega_{jt}} e^{\varepsilon_{jt}} \tag{4}$$

where each flexible input task M_{hjt} is a CES mix of intermediates Q_{hjt} purchased from industry h, and/or task-specific labor L_{hjt}[14]

$$M_{hjt} = \left(\gamma_h (e^{z_{hjt}} L_{hjt})^{\rho_h} + (1-\gamma_h) Q_{hjt}^{\rho_h}\right)^{1/\rho_h} \tag{5}$$

and K_{jt} is a standard measure of aggregate capital[15] which I assume is predetermined[16]. The parameter ρ_h determines the elasticity of substitution between purchased intermediates and

[14]For example, input M_{hjt} may be the quantity of transportation inputs required by the firm. These transportation inputs are a CES combination of services provided by in-house transportation labor L_{hjt} (such as truck drivers) and goods/services purchased from the transportation sector Q_{hjt} (such as long-distance shipping).

[15]My theory easily extends to the case where capital is also task-specific and flexibly substitutable with labor/intermediates. I provide details and discuss why I prefer the aggregate specification in appendix B.

[16]By *predetermined* I mean the level of capital in period t is fixed in period $t-1$. By *flexible* I mean that the input is chosen in period t and doesn't depend on past values of itself. See appendix E for details.

labor of type h. α_h and γ_h are the scale and distribution parameters for each task-matched set of labor and intermediates which, along with ρ_h, are allowed to vary by industry. z_{hjt} represents firm-task specific labor enhancing productivity relative to purchased intermediates[17]. I follow the standard assumptions for the hicks neutral productivity term, which is represented by a known component ω_{jt} (assumed to be first-order markov), and an ex-post i.i.d productivity shock ε_{jt}.

This production framework has several key implications and benefits. First, firms are differentiated by a vector of task-augmenting productivity terms $\{z_{hjt}\}_{h\in\mathcal{H}_i}$. Heterogeneity in task-efficiency (along with prices) will explain differences in input composition across firms and time. Second, this framework nests as a special case a disaggregated Cobb-Douglas production function[18] (as $\rho_h \to 0$),

$$Y_{jt} = K_{jt}^{\beta} \prod_{h\in\mathcal{H}_i} L_{hjt}^{\alpha_h\gamma_h} Q_{hjt}^{\alpha_h(1-\gamma_h)} e^{\tilde{\omega}_{jt}} e^{\varepsilon_{jt}} \tag{6}$$

as well as Leontief ($\rho_h \to -\infty$) and linear ($\rho_h \to 1$) production functions. This allows me to cleanly test the restrictions which are embodied in much of the existing empirical literature on production functions. In particular, I use this disaggregated Cobb-Douglas formulation as a "benchmark" against which to compare my flexible framework. Third, unlike Cobb-Douglas or trans-log functions, the matched CES production function is still defined for firms which either produce an entire task in-house ($Q_{hjt} = 0$) or outsource the entire task ($L_{hjt} = 0$). With the addition of fixed/adjustment costs (see section 4.3.1), the model is able to rationalize the endogenous corner solutions (zeros) observed in the data. Fourth, this nested structure implies that the production function is *weakly separable* in the different tasks, which implies that the firm's problem of input demand can be broken into multiple stages, which I discuss in section 4.3. This will allow me to estimate the parameters for each task separately. Finally, under some conditions, the CES aggregator in labor and intermediates can be seen as a first-order approximation of an arbitrary production function in those same factors[19], and can be built up from micro-foundations[20].

[17]This term encompasses any firm-level unobserved heterogeneity which leads firms to differ in their optimal labor mix for that input task given expected wages. This could include firm-level differences in labor/management productivity, differences in the relative importance of labor in the firm or industry production technology, or unobserved specification/measurement error.

[18]Here $\tilde{\omega}_{jt} = \omega_{jt} \sum_h z_{hjt}$ subsumes the task-enhancing productivity terms.

[19]See Doraszelski and Jaumandreu (2018) for a discussion.

[20]Suppose each aggregate task is produced using a continuum of sub-tasks. Each sub-task can be performed by a team of in-house task-specific labor types (transportation managers, logistics clerks, truck drivers) or outsourced to another firm. Rosen (1978) provides conditions on the structure of relative productivity over

Few data sets actually include input use at a disaggregated level. However, the implications and benefits of using this framework are not conditional on having such data. In the usual case where data is only available for aggregate labor and intermediate expenditure, I show[21] that the single-task flexible framework can be expressed as follows,

$$Y_{jt} = K_{jt}^{\beta} M_{jt}^{\alpha} e^{\tilde{\omega}_{jt}} e^{\varepsilon_{jt}} = K_{jt}^{\beta} \left(X_{jt}^{Q} \right)^{\alpha} (1 - S_{jt})^{-\frac{\alpha}{\rho}} a_t e^{\tilde{\omega}_{jt}} e^{\varepsilon_{jt}} \tag{7}$$

where X_{jt}^{Q} is total expenditure on intermediates, S_{jt} is the labor share in total variable expenditure, and a_t subsumes parameters which may vary over time. This specification has the benefit of being easily estimated in log-linear form, while still retaining the flexibility of the task-matched CES function. In the results and applications sections, I show that using this specification to allow for substitution over aggregate inputs results in substantially different estimates of firm efficiency.[22] It also, like the disaggregated specification, can flip the sign on estimated elasticity terms. To give a direct comparison to the strictly negative Cobb-Douglas result in 3, the short-run cross-price elasticity of demand in the aggregated case is

$$\epsilon_{P^Q}^{L_j} \equiv \frac{\partial L_j}{\partial P^Q} \frac{P^Q}{L_j} = \left(\frac{\rho}{1-\rho} - \frac{\alpha}{1-\alpha} \right) (1 - S_{jt}) \tag{8}$$

which can be either positive or negative, depending on the elasticity of substitution, scale parameter, and returns to scale[23].

While this generalized framework still embodies strong restrictions on the production technology – notably a unitary elasticity of substitution between aggregate tasks – it proves to be most convenient in terms of tractability, identification and parsimony. I argue that allowing for CES aggregation of task-matched labor and intermediates is sufficient for the task at hand, which is to estimate and control for the heterogeneous input demand, substitution and outsourcing patterns seen in the disaggregated data.

4.2 The Revenue Production Function

The main identification strategy used in this paper relies on firms' profit maximizing behavior. As such, I follow Klette and Griliches (1996), De Loecker (2011), and Halpern, Koren and Szeidl (2015) in assuming that firms face a simple downward sloping demand curve,

the task continuum under which aggregate task output is CES in labor and intermediates.

[21]See section 5.5 for a derivation in the H_i-task case and a discussion of estimation strategy.

[22]This is the specification I use when estimating the effects of tariffs on productivity in section 8

[23]See section 6.2 for the multi-input case, results and further discussion.

with $Y_{jt}(P^o_{jt}) = \psi_{jt}(P^o_{jt})^{-\eta^d}$. Here P^o_{jt} is the firm's output price relative to some industry price index, ψ_{jt} is a firm-specific demand shifter and η^d is the industry-level price elasticity of demand[24]. We can then write the firm's *revenue production function* as

$$R_{jt} \equiv P^o_{jt} * Y_{jt}(P^o_{jt}) = \psi_{jt}^{1-\theta} \left[K_{jt}^{\beta} \prod_{h \in \mathcal{H}_i} M_{hjt}^{\alpha_h} e^{\omega_{jt}} e^{\varepsilon_{jt}} \right]^{\theta} \tag{9}$$

where $\theta \equiv (\eta^d - 1)/\eta^d$. Note that $\theta = 1$ would imply perfect competition. Since I am working with revenues, this is one of the main equations I will take to the data. However, I differ from the cited papers in that I am able to use data on prices and sales to estimate η^d (and thus θ) directly rather than recovering them from the supply side.

4.3 Input Factor Demand

The task-matched production framework (along with some price and timing assumptions) provides an input demand system which forms the basis of my main estimation strategy. My approach relies on the property of weak separability in tasks embodied in the production function. Weak separability implies that the demand for any particular type of labor or intermediate depends only on the relative prices/productivity of the other sub-task inputs (L_{hjt} or Q_{hjt}) and overall demand for the task aggregate M_{hjt}[25]. Importantly, this means that the firm's input choice problem for each task can be broken up into several stages, across which their information set may evolve. Each of these stages provides a key relationship which I employ in estimating the model.

The three production stages and associated timing assumptions[26] are as follows: First, firms choose the level of each input M_{hjt} conditional on firm productivity ω_{jt} and expected prices/wages. This provides the expenditure share equations which I use to estimate the scale parameters $\alpha_h \theta$. Second, firms observe input prices P_{ht} and task productivity z_{hjt} and then choose how to acquire each input (the make-or-buy decision). Third, conditional on the

[24] This basic assumption on demand can be derived from both a CES or a Logit demand system. The former is more convenient when output prices are unobserved (as in De Loecker, 2011), while the latter is perhaps preferable when output prices and quantities are known, allowing η^d to be estimated directly from production data.

[25] See Nadiri (1982) and Varian (1992). I formally state the definition, proposition and proof of separability in appendix C

[26] While I call these "timing" assumptions, one could also think of this multi-stage decision process as reflecting decisions made at different levels of the firm. See appendix E for details and a formal statement of the assumptions on timing and prices.

make-or-buy choice, firms observe wages W_{hjt} and decide how much labor and intermediates of each type to use. This intensive-margin choice provides the input share equations which I use to estimate ρ_h and thus the elasticity of substitution.

4.3.1 Firm Scale and the Make-or-Buy Decision

In the first stage, I assume that firms choose optimal input levels M^*_{hjt} under price uncertainty, since the cost of input h depends on wages, prices and task-productivity which are not observed in the first stage. The solution to the firm's input choice is then the first order condition for M_{hjt}, or the *Expenditure Share Equation*,

$$M^*_{hjt} = \alpha_h \theta R_{jt} \mathbb{E}[P^I_{hjt}]^{-1} \tag{10}$$

which I will use to estimate $\alpha_h \theta$.

Given optimal input levels M^*_{hjt} from the first stage, the firm then observes input prices P_{ht} and task productivity z_{hjt} and chooses the cost minimizing arrangement {Buy, Buy and Make, Make} for each of the N tasks. I refer to the choice of buying and making as "Both" going forward.

It's important to note that the CES structure of this model does not by itself rationalize extensive margin make-or-buy decisions by the firm. The firm will always optimally require some positive amount of both L_{hjt} and Q_{hjt} regardless of the relative prices. Here I rationalize outsourcing by viewing the choice of input production technology as embodying some sort of fixed or adjustment costs. If a firm wishes to hire labor and run its own accounting department, there is a fixed cost f^L_{hjt} of doing so, which could include capital rental/setup, hiring costs, management costs, etc. Similarly, there is a fixed cost f^Q_{hjt} of purchasing intermediates, which could represent search or contracting costs. I allow these costs to differ by firm, task and year. Let $\mathcal{D}_{hjt} \in$ {Buy, Both, Make} represent the choice of procurement technology made by the firm. Given the assumptions thus far, and the solution to equation 10, the cost of each procurement technology is as follows:

$$\text{Cost}(\mathcal{D}_{hjt} \mid M^*_{hjt}) = \begin{cases} P_{ht} M^*_{hjt} (1-\gamma_h)^{-\frac{1}{\rho_h}} + f^{\text{Q}}_{hjt} & (\mathcal{D}_{hjt} = \text{Buy}) \\ P_{ht} M^*_{hjt} G_{hjt} + f^{\text{Q}}_{hjt} + f^{\text{L}}_{hjt} & (\mathcal{D}_{hjt} = \text{Both}) \\ \frac{\mathbb{E}[W_{hjt}]}{e^{z_{hjt}}} M^*_{hjt} {\gamma_h}^{-\frac{1}{\rho_h}} + f^{\text{L}}_{hjt} & (\mathcal{D}_{hjt} = \text{Make}) \end{cases} \tag{11}$$

where

$$G_{hjt} \equiv \left(\gamma_h^{\frac{1}{1-\rho_h}} \left(\frac{\mathbb{E}[W_{hjt}]}{e^{z_{hjt}} P_{ht}} \right)^{\frac{\rho_h}{\rho_h - 1}} + (1-\gamma_h)^{\frac{1}{1-\rho_h}} \right)^{\frac{\rho_h - 1}{\rho_h}} \tag{12}$$

can be seen as the cost discount or benefit from doing both (stemming from the shape of the CES production function). Note that $G_{hjt} \to 1$ if labor and intermediates approach perfect substitution ($\rho_h \to 1$). The choice of procurement technology then depends on (expected) relative prices, firm productivity z_{hjt}, as well as the shape of the task production function for that task-industry pair (i.e.: the parameters ρ_h and γ_h) and the fixed costs.

Figure 6 illustrates the relative cost curves as a function of firm productivity. As productivity goes to infinity, the marginal cost of an effective unit of labor goes to zero, in which case the cost of "make" asymptotes to the fixed cost of hiring labor f^L_{hjt} and the cost of "both" asymptotes to the sum of the fixed costs. The optimal choice of the firm as a function of productivity is the lower envelope of these three curves, as shown by the light blue line in the figure. The figure makes several mechanisms clear. First, if the fixed cost of hiring labor is too high, firms will always find it optimal to outsource the task. The location of the make & buy curve relative to the others depends on ρ_h. As $\rho_h \to 1$, the cost of doing both shifts up until the lower envelope curve only involves the extensive margin technologies (Perfect Substitution). As ρ_h decreases, the cost of doing both shifts down until it is always optimal to do both (Cobb-Douglas/Leontief).

This cost-minimization problem provides a set of productivity cutoffs which characterize the firm's optimal choice,

$$\text{Cost(Buy)} < \text{Cost(Both)} \iff z_{hjt} < \tilde{z}^1_{hjt} \tag{13}$$

$$\text{Cost(Both)} < \text{Cost(Make)} \iff z_{hjt} < \tilde{z}^2_{hjt} \tag{14}$$

$$\text{Cost(Buy)} < \text{Cost(Make)} \iff z_{hjt} < \tilde{z}^3_{hjt} \tag{15}$$

which I derive explicitly in section 5.3.1. For example, the firm chooses to "Buy" input h if and only if both conditions 13 and 15 hold. Given equation 12, these two conditions can be rearranged and inverted to give the following productivity cutoff rule: $z_{hjt} < \min\{\tilde{z}^1_{hjt}, \tilde{z}^3_{hjt}\}$, where the cutoff terms depend generally on expected wages, prices, firm scale and fixed costs. These three cutoffs correspond to the cost-curve intersections on figure 6. Similarly, the firm will only choose to do Both if $\tilde{z}^1_{hjt} < z_{hjt} < \tilde{z}^2_{hjt}$ and will only choose to Make if $z_{hjt} > \max\{\tilde{z}^3_{hjt}, \tilde{z}^2_{hjt}\}$. Modeling the extensive margin decision in this simple way will allow me to control for selection bias when estimating the substitution parameters (see section

5.3), and is the principle innovation provided by this paper.

4.3.2 Labor and Intermediates

Given the choice of procurement technology $\mathcal{D}_{hjt}$, the firm then observes wages W_{hjt}[27]. The optimal choice of labor and intermediates for the two extensive margin cases are as follows:

$$(L^*_{hjt}, Q^*_{hjt}) = \begin{cases} (0, M^*_{hjt}(1-\gamma_h)^{-1/\rho_h}) & \text{if } \mathcal{D}_{hjt} = \text{Buy} \\ (M^*_{hjt}\gamma_h^{-1/\rho_h} e^{-z_{hjt}}, 0) & \text{if } \mathcal{D}_{hjt} = \text{Make} \end{cases}$$

If the firm decides to both hire labor in-house and purchase some amount of the input task on the market, then the firm's optimal choice of each is determined by the firm's cost minimization problem with respect to input task requirement M^*_{hjt}. Optimal labor and intermediate demand is then,

$$L^*_{hjt} = X_{hjt} W^{-1}_{hjt} S_{hjt} \tag{16}$$

$$Q^*_{hjt} = X_{hjt} P^{-1}_{ht} (1 - S_{hjt}) \tag{17}$$

where X_{hjt} is total expenditure on input h, $X^L_{hjt} \equiv W_{hjt} L_{hjt}$ is expenditure on labor of type h, and $S_{hjt} \equiv X^L_{hjt}/X_{hjt}$ is the labor share of total expenditure on h. Combining equations 16 and 17 provides the *Input Ratio Equation*[28],

$$\frac{L_{hjt}}{Q_{hjt}} = \left(\frac{P_{ht}}{W_{hjt}}\right)^{\frac{1}{1-\rho_h}} \left(\frac{\gamma_h}{1-\gamma_h}\right)^{\frac{1}{1-\rho_h}} (e^{z_{hjt}})^{\frac{\rho_h}{1-\rho_h}} \tag{18}$$

This will be the key estimating equation, along with 9 and 10.

---

[27] I assume that firm-task specific wages W_{hjt} are a function of some common market component W_{hit}, firm productivity and a firm-task component Θ_{hjt} which may represent compensating differentials or differences in labor market tightness across locations. See appendix E for details. While firms have some market power in setting wages, I assume wages are fixed in stage 3 when firms are deciding the input ratio for h.

[28] To see this, note that optimal expenditure shares S_{hjt} are functions of wages, prices and task productivity:

$$S_{hjt} \equiv \frac{\gamma_h^{\frac{1}{1-\rho_h}} \left(\frac{W_{hjt}}{e^{z_{hjt}}}\right)^{\frac{\rho_h}{\rho_h-1}}}{\left(\gamma_h^{\frac{1}{1-\rho_h}} \left(\frac{W_{hjt}}{e^{z_{hjt}}}\right)^{\frac{\rho_h}{\rho_h-1}} + (1-\gamma_h)^{\frac{1}{1-\rho_h}} P_{hit}^{\frac{\rho_h}{\rho_h-1}}\right)}$$

Plugging this into equations 16 and 17 and dividing one by the other provides 18

5 Estimation

Estimating the model proceeds in several steps. Firms in industry i have H_i input tasks which leads to the estimation of H_i expenditure share equations (10), H_i input ratio equations (18), and one revenue production function (9). I first recover the H_i scale parameters ($\{\alpha_h\theta\}_{h=1}^{H_i}$) using the set of expenditure share equations. Next I estimate the substitution parameters (notably $\{\rho_h\}_{h=1}^{H_i}$) using the input ratio equations. This involves addressing the selection bias which stems from the make-or-buy decision and is my main methodological contribution. Finally I recover the capital scale parameter (β) using the revenue production function.

5.1 Scale, Demand, and Wages

The assumption that inputs are flexible allows me to estimate all of the $\alpha_h\theta$ terms with a standard expenditure share approach. The main idea is that under some conditions, rearranging and taking expectations of equation 10 provides $\alpha_h\theta = \mathbb{E}\left[X_{hjt}/R_{jt}\right]$, which has a simple empirical analog that I use to estimate $\widehat{\alpha_h\theta}$ for each task-industry pair.

I estimate the industry-level demand elasticity η_i and firm-level demand shifters ψ_{jt} using a simple logit demand model and firm-level data on output quantities. This provides estimates of $\widehat{\theta}$ for each industry which I can use with $\widehat{\alpha_h\theta}$ to back out the physical production scale parameters $\widehat{\alpha}_h$. Finally, I specify a firm's expected wage for a particular task h as $\mathbb{E}[W_{hjt}] = g_w(W_{hit}, z_{hjt}, \omega_{jt}, \mathbb{E}[\Theta_{hjt}])$ for some unknown function g_w which I approximate with a polynomial $\hat{g}_w(h, i, t, W_{hjt-1}, R_{jt-1}, j)$. Lagged wages, revenues and firm fixed effects proxy for unobserved productivity and labor market heterogeneity, and industry-task-year effects capture the average industry-task-year wage component. I assume that this specification matches how the firm itself calculates expected wages and use the predicted values from the wage regression in the estimation of the structural model. See appendix F for further details on how I estimate the scale, demand, and wage terms.

5.2 Substitution Parameters

The basic strategy for estimating the substitution parameters is to use the input ratio equations (18), where ρ_h is identified off of variation in input prices relative to the labor-intermediate demand ratio. However, there are several difficulties.

First, since firm-task productivity may be correlated with prices, OLS estimates of ρ_h will be biased. To deal with this issue, I assume the task specific labor-enhancing productivity term z_{hjt} follows an AR(1) process: $z_{hjt} = z_h + \delta_h z_{hjt-1} + \zeta_{hjt}$ where the innovation term is i.i.d. normal $\zeta_{hjt} \sim N(0, \sigma_h)$. Given some substitution, we can then rewrite equation 18 in logs as

$$\ell_{hjt} - x^Q_{hjt} = a_{ht} - \frac{1}{1-\rho_h} w_{hjt} + \frac{\delta_h}{1-\rho_h} w_{hjt-1} + \delta_h(\ell_{hjt-1} - x^Q_{hjt-1}) + \frac{\rho_h}{1-\rho_h}\zeta_{hjt} \tag{19}$$

where lower case letters represent logged variables and a_{ht} is a time-specific effect which subsumes a set of fixed parameters. This assumption on z_{hjt} conveniently means that each of the H_i equations can be estimated independently. Also, while contemporaneous wages are still potentially correlated with ζ_{hjt}, the assumption on ζ_{hjt} provides a set of potential instruments for w_{hjt}. In theory, any function of lagged variables which is correlated with current wages may work as an instrument since the productivity innovation ζ_{hjt} is orthogonal to all lagged variables. In practice, I use lagged revenues, which are appropriate for two reasons. First, lagged revenues are uncorrelated with the error term the timing assumptions. Second, they are correlated with wages, since wages are correlated with firm tfp ω_{jt}, which is assumed persistent (first-order markov) and correlated with firm revenues.

The second problem is that the input ratio equation is only defined for firms which use both labor and intermediates. Since the choice to do both depends on prices and productivity, this introduces selection bias. To see this, note that we can express the innovation term in 19 as $\zeta_{hjt} = \mathbb{E}[\zeta_{hjt} \,|\, \mathcal{D}_{hjt} = \text{Both}] + \tilde{\zeta}_{hjt}$ which gives us

$$\begin{aligned} \ell_{hjt} - x^Q_{hjt} = a_{ht} - \frac{1}{1-\rho_h} w_{hjt} + \frac{\delta_h}{1-\rho_h} w_{hjt-1} + \delta_h(\ell_{hjt-1} - x^Q_{hjt-1}) \\ + \frac{\rho_h}{1-\rho_h}\mathbb{E}[\zeta_{hjt} \,|\, \mathcal{D}_{hjt} = \text{Both}] + \frac{\rho_h}{1-\rho_h}\tilde{\zeta}_{hjt} \end{aligned} \tag{20}$$

which will be the main estimating equation for ρ_h. Recall from section 4.3.1 that the firm will only do Both if $\tilde{z}^1_{hjt} < z_{hjt} < \tilde{z}^2_{hjt}$. Plugging in the AR(1) structure of z_{hjt}, we get that

$$\mathbb{E}[\zeta_{hjt} | \mathcal{D}_{hjt} = \text{Both}] = \mathbb{E}[\zeta_{hjt} | C^1_{hjt} < \zeta_{hjt} < C^2_{hjt}] \tag{21}$$

for some firm-specific cutoffs C^1_{hjt} and C^2_{hjt}. This will not in general be zero, introducing selection bias into the estimation.

5.3 The Selection Problem

The standard way to correct for this sort of selection bias would be to estimate a two sided multi-stage Heckman correction as in the literature following Heckman (1979). I instead approach the problem from a parametric maximum likelihood perspective for several reasons. First, as I will show, the selection condition is a function of the same parameters which characterize the input ratio equation. As such, estimating selection and input choice jointly increases the efficiency the estimation procedure. Second, this approach allows me to recover the distribution of fixed costs and perform counter factual experiments related to the firm's probability of outsourcing as a function of prices, demand, fixed costs and productivity.

5.3.1 Empirical Cutoffs

The core strategy to control for selection is to specify and estimate the firm's make-both-buy decision using maximum likelihood[29]. In order to do this, I need to fully specify the productivity cutoffs discussed in section 4.3.1. Since the main estimating equation is in terms of ζ_{hjt}, I can plug in the AR(1) structure of z_{hjt}, providing the following three cutoff terms: C^1_{hjt}, C^2_{hjt}, C^3_{hjt}, where for example, $C^1_{hjt} = \tilde{z}^1_{hjt} - z_{hjt-1} - \bar{z}_h$. It is important to note that each is a function of expected wages, total input expenditure and fixed costs. For example, $C^1_{hjt}(f^L_{hjt}, \mathbb{E}[W_{hjt}], P_{ht})$ is monotone increasing in expected wages and f^L_{hjt}, and monotone decreasing in intermediate price. This is intuitive – as the fixed and variable costs of hiring labor increase, a firm will need to be more productive in order for hiring both labor and intermediates to be cheaper than just buying. Similarly, $C^2_{hjt}(f^Q_{hjt}, \mathbb{E}[W_{hjt}], P_{ht})$ is monotone decreasing in prices and f^Q_{hjt} and increasing in expected wages. As the fixed and variable costs of purchasing intermediates increase relative to the cost of labor, even lower productivity firms will find it cost effective to do everything in house.

The fact that ζ_{hjt} is normally distributed provides the following choice probabilities. Note that here $\Phi()$ represents the CDF of the standard normal distribution, and that I am ignoring

[29] I will actually estimate the ML model jointly with the input ratio equation using the scores of the LLH function as moments in the GMM procedure.

for now the stochastic nature of the fixed costs.

$$\Pr(\mathcal{D}_{hjt} = \text{Buy}) = \Pr(\zeta_{hjt} < \min\{C^1_{hjt}, C^3_{hjt}\}) \equiv \Phi\left(\frac{\min\{C^1_{hjt}, C^3_{hjt}\}}{\sigma_h}\right)$$
$$\Pr(\mathcal{D}_{hjt} = \text{Both}) = \Pr(C^1_{hjt} < \zeta_{hjt} < C^2_{hjt}) \equiv \Phi\left(\frac{C^2_{hjt}}{\sigma_h}\right) - \Phi\left(\frac{C^1_{hjt}}{\sigma_h}\right)$$
$$\Pr(\mathcal{D}_{hjt} = \text{Make}) = \Pr(\zeta_{hjt} > \max\{C^3_{hjt}, C^2_{hjt}\}) \equiv 1 - \Phi\left(\frac{\max\{C^3_{hjt}, C^2_{hjt}\}}{\sigma_h}\right)$$

Note that since each cutoff term is a function of lagged task productivity z_{hjt-1}, the selection problem is identified off of the firms which choose Both in period $t-1$ and t, as well as the firms which chose Both in period $t-1$ but switched to Make or Buy in period t. In other words, the additional identification is coming from the "switchers" in the data – the firms which received productivity, wage or fixed cost shocks which pushed them over the cutoff into outsourcing or in-house production.

5.3.2 Unobserved Fixed Costs

While the firms observe their own productivity and fixed cost terms, both are unobserved by researcher. To deal with this issue, I assume that fixed costs f^L_{hjt} and f^Q_{hjt} follow an i.i.d. log-normal distribution, such that $\log(f^x_{hjt}) \sim N(\bar{f}^x, \sigma^x_h)$ for $x \in \{L, Q\}$. This implies that not only is productivity stochastic, but so are the cutoff terms. Thus, for example, the probability that the cost of outsourcing was cheaper than the cost of doing both for a given firm is

$$\Pr(\zeta_{hjt} < C^1_{hjt}) = \int \Pr(\zeta hjt < C^1_{hjt}(f)) \Pr(f^L = f) df$$

By the assumptions on ζ_{hjt} and fixed costs, this becomes

$$\Pr(\zeta_{hjt} < C^1_{hjt}) = \int \Phi\left(\frac{C^1_{hjt}(f)}{\sigma_h}\right) \phi\left(\frac{\log(f) - \bar{f}^L}{\sigma^L_h}\right) df$$

where as before, Φ is the CDF of the standard normal, and ϕ is the PDF of the standard normal. Similar terms can be derived for the other probabilities.

Given the structure of the cutoffs and fixed costs described above, we can re-derive the choice probabilities in section 5.3.1 and derive the selection correction term in equation 21:

$$\mathbb{E}[\zeta_{hjt} \,|\, \mathcal{D}_{hjt} = \text{Both}] = \mathbb{E}[\,\mathbb{E}[\zeta_{hjt} \,|\, C^1_{hjt}(f^L_{hjt}) < \zeta_{hjt} < C^2_{hjt}(f^Q_{hjt})]\,] \tag{22}$$

where the expectations are taken over ζ_{hjt} and the two fixed cost distributions. Equation 22 is the selection control term which will be included in the input ratio equation in order to control for the selection bias.

5.4 Joint Estimation of Input Ratio and Selection Problem

The goal is to get consistent estimates of parameters $\Omega_h = \{\rho_h, \delta_h, \sigma_h, \bar{f}_h^L, \sigma_h^L, \bar{f}_h^Q, \sigma_h^Q, a_{ht}\}$, which I obtain using two key expressions. The first is the input ratio equation (20) which provides the set of moment conditions

$$\mathbb{E}[\mathbb{Z}_{hjt}\tilde{\zeta}_{hjt}] = 0$$

where $\mathbb{Z}_{hjt}$ is a vector of functions of the exogenous variables. In practice, $\mathbb{Z}_{hjt}$ consists of lags of log wages, log revenue and log input ratios plus year indicators. I calculate the selection correction term using numerical integration over the distributions of ζ_{hjt} and the fixed costs, given a guess of Ω_h. It's important to remember that equation 20 is only defined for observations where firms choose both in two subsequent years. Let this set of observations be $\mathcal{J}_h^{IR} = \{(j,t) \in \mathcal{J} \,|\, \mathcal{D}_{hjt-1} = \text{Both}, \mathcal{D}_{hjt} = \text{Both}\}$, with J_h^{IR} the number of such observations and $\mathcal{J}$ the overall set of firm-year observations. The empirical moments are then constructed as

$$\mathcal{M}_h(\Omega_h; \mathcal{Y}_h^{IR}) = \frac{1}{J_h^{IR}} \sum_{\mathcal{J}_h^{IR}} \mathbb{Z}_{hjt}\tilde{\zeta}_{hjt}$$

where $\mathcal{Y}_h^{IR}$ is the data for observations $\mathcal{J}_h^{IR}$. The second expression is the likelihood function for the selection problem in section 5.3. Recall that the selection problem is defined over a different (super) set of observations relative to the input ratio equation. Let $\mathcal{J}_h^{Sel} = \{(j,t) \in \mathcal{J} \,|\, \mathcal{D}_{hjt-1} = \text{Both}\}$ be the set of firm-year observations in which the firm chose Both in the previous period and *anything* in the current period, with $J_h^{Sel} \geq J_h^{IR}$ the number of usable observations. Additionally define $\mathcal{A}_{hjt}$ as an indicator function which equals one if $\mathcal{D}_{hjt} = \mathcal{A}$ for $\mathcal{A} \in \{\text{Buy, Both, Make}\}$. Then the likelihood function is

$$\mathcal{L}(\Omega_h; \mathcal{Y}_h^{Sel}) = \prod_{\mathcal{J}_h^{Sel}} \Pr(\mathcal{D}_{hjt} = \text{Buy})^{\text{Buy}_{hjt}} * \Pr(\mathcal{D}_{hjt} = \text{Both})^{\text{Both}_{hjt}} * \Pr(\mathcal{D}_{hjt} = \text{Make})^{\text{Make}_{hjt}}$$

where the probabilities are over the distributions of productivity and fixed costs as defined as in section 5.3.2. The score of the log-likelihood, $\mathcal{S}_h$ (an N_h^S-element vector), is then defined as

$$\mathcal{S}_h(\Omega_h; \mathcal{Y}_h^{Sel}) \equiv \frac{\partial}{\partial \Omega_h} \log(\mathcal{L}(\Omega_h; \mathcal{Y}_h^{Sel}))$$

Let $N_h^{\mathbb{Z}}$ be the number of instruments for the input ratio equation, so that $N_h^{\mathcal{Q}} = N_h^{\mathbb{Z}} + N_h^S$ is the total number of moments. This provides the GMM objective function

$$\mathcal{Q}(\Omega_h) = g(\Omega_h)' \hat{W}_n g(\Omega_h)$$

where $g(\Omega_h) = [\mathcal{M}_h, \mathcal{S}_h]'$ is the $N_h^{\mathcal{Q}} \times 1$ vector of moments, and $\hat{W}_h$ is a $N_h^{\mathcal{Q}} \times N_h^{\mathcal{Q}}$ weighting matrix. I use a two-step estimator based on Hansen (1982), with the modification that the weighting matrix is block-diagonal, due to the two sets of moments being constructed with different sets of observations (though the observations used in the input ratio moments $\mathcal{J}_h^{IR}$ is a subset of those used in the selection problem $\mathcal{J}_h^{Sel}$).

5.5 Productivity

Given the preceding discussion on how to recover estimates of the scale and input substitution parameters, I lay out two methods for estimating the revenue production function itself, and critically, the unobserved productivity term ω_{jt}. The first method accounts for extensive margin selection by building upon first-stage estimates of the substitution and scale parameters from the previous section. The second method provides an alternate strategy for jointly estimating all of the model parameters when working with aggregate input data such that the extensive margin is not a concern.

5.5.1 Method 1

The basic strategy is to recover estimates of the production contribution of each input task, which is a function of data and parameters, using the procedure in sections 5.1 through 5.1. We can then plug these into the production function and use one of several approaches to estimate the remaining parameters (namely the scale coefficient on capital). Given the results of the joint estimation of the input ratio equation and the selection problem, we have a set of estimated parameters $\widehat{\Omega}_H = \cup_h \widehat{\Omega}_h$ where $\widehat{\Omega}_h = \{\widehat{\alpha_h \theta}, \widehat{\rho}_h, \widehat{\delta}_h, \widehat{\sigma}_h, \widehat{\bar{f}}_h^L, \widehat{\sigma}_h^L, \widehat{\bar{f}}_h^Q, \widehat{\sigma}_h^Q, \widehat{a}_{ht}\}$. After some rearranging, the input task term can be expressed as:

$$M_{hjt} = P_{ht}^{-1}(1-\gamma_h)^{\frac{1}{\rho_h}} X_{hjt}^Q (1 - S_{hjt})^{-\frac{1}{\rho_h}} \tag{23}$$

which then allows us to construct the production contribution of each input task $\widetilde{M}_{hjt}$, such that

$$\widetilde{M}_{hjt} = \begin{cases} X^Q_{hjt} & \text{if } \mathcal{D}_{hjt} = \text{Buy} \\ X^Q_{hjt}(1 - S_{hjt})^{-\frac{1}{\hat{\rho}_h}} & \text{if } \mathcal{D}_{hjt} = \text{Both} \\ L_{hjt} & \text{if } \mathcal{D}_{hjt} = \text{Make} \end{cases} \tag{24}$$

Define the *input technology* of firm j in period t as $c(jt)$, where $c(jt) = \{\mathcal{D}_{hjt}\}_{h \in \mathcal{H}_i} \in \mathcal{C}^{H_i}$ is the set of make-both-buy choices over all required input tasks. The revenue production function is then (adding the industry subscripts back in)

$$R_{jit} = \hat{\psi}^{1-\hat{\theta}_i}_{jit} \left[K^{\beta_i}_{jit} \prod_{h \in \mathcal{H}_i} \widetilde{M}^{\hat{\alpha}_{hi}}_{hjit} \Gamma_{c(jt)i} e^{\omega_{jit}} e^{\varepsilon_{jit}} \right]^{\hat{\theta}_i} \tag{25}$$

where $\Gamma_{c(jt)i}$ is an input technology specific term which subsumes the time-industry parameters particular to the firm's choice of input technology in period t. Following Griliches and Ringstad (1971), this can be restated (taking logs) as

$$\tilde{r}_{jit} = \beta_i k_{jit} + \gamma_{c(jt)i} + g_\omega(\omega_{jit-1}) + \eta^\omega_{jit} + \varepsilon_{jit} \tag{26}$$

where $\eta^\omega_{jit} \equiv \omega_{jit} - g_\omega(\omega_{jit-1})$ is the (mean zero) innovation to the firm's productivity, and

$$\widetilde{r}_{jit} \equiv \log \left(R^{\frac{1}{\hat{\theta}_i}}_{jit} \hat{\psi}^{\frac{\hat{\theta}_i - 1}{\hat{\theta}_i}}_{jit} \prod_{h \in \mathcal{H}_i} \widetilde{M}^{-\hat{\alpha}_{hi}}_{hjit} \right) \tag{27}$$

This just leaves one key parameter to estimate: β_i. Because capital is predetermined, the issue in estimating equation 26 is the correlation between k_{jit} and ω_{jit-1}. This can be tackled in two ways. One approach following Wooldridge (2009) and Levinsohn and Petrin (2003), which I will refer to as "WLP", is to provide conditions under which one of the flexible inputs is a strictly monotone in productivity, conditional on the other state variables. Optimal (log) input expenditure for some task h can be expressed as $x^*_{hjit} = g_{mh}(k_{jit}, \omega_{jit})$ for some unknown function g_{mh}. Given the assumptions above, this function can be inverted so that $\omega_{jit} = g^{-1}_{mh}(k_{jit}, x^*_{hjit})$. Taking expectations and plugging in provides

$$\mathbb{E}[\tilde{r}_{jit}] = \beta_i k_{jit} + \gamma_{c(jt)i} + g^{-1}_{mh}(k_{jit}, x^*_{hjit}) \equiv \phi_r(k_{jit}, x^*_{hjit}, \gamma_{c(jt)i}) \tag{28}$$

and so

$$\tilde{r}_{jit} = \beta_i k_{jit} + \gamma_{c(jt)i} + g_\omega(\phi_{rt-1} - \beta_i k_{jit-1} - \gamma_{c(jt-1)i}) + \eta^\omega_{jit} + \varepsilon_{jit} \tag{29}$$

with $\phi_{rt-1} \equiv \phi_r(k_{jit-1}, x^*_{hjit-1}, \gamma_{c(jt-1)i})$. The WLP approach is to estimate these two equations jointly, typically using some variety of polynomial to approximate ϕ_r and g_ω. I refer to any estimates using this approach as "M1-WLP".

A second possible approach is to follow Doraszelski and Jaumandreu (2013) in noting that the first order conditions of the firm provide me with an exact expression for g_{mh}. With the additional assumption that $g_\omega(\omega_{jit-1}) = \delta_\omega \omega_{jit-1}$, the estimating equation becomes

$$\tilde{r}_{jit} = \beta_i k_{jit} + \gamma_{c(jt)i} + \delta_\omega\Big(\frac{1}{\theta_i}(x^*_{hjit-1} - \log(\widehat{\alpha_{hi}\theta_i\widehat{\eta}_{hi}}) - \hat{\eta}_{hjit-1} - r_{jit-1}) - \beta_i k_{jit-1} + \tilde{r}_{jit-1}\Big) + \eta^\omega_{jit} + \varepsilon_{jit} \tag{30}$$

I refer to estimates with this second approach as "M1-DP" for "direct proxy". Both of these equations can be estimated with GMM relying on just the variables present on the right hand side, as they are all exogenous by the timing assumptions on η^ω_{jit} and ε_{jit}. For example, given the second equation, let $\mathbb{Z}^2_{jit} = \{k_{jit}, x^*_{hjit-1}, \hat{\eta}_{hjit-1}, r_{jit-1}, \tilde{r}_{jit-1}, k_{jit-1}, c(jt)\}$ be the set of instruments. Consistent estimates of $(\beta_i, \delta_\omega, \gamma_{c(ji)t})$ can then be obtained with moments of the following form:

$$\mathbb{E}[\mathbb{Z}^2_{jit}(\eta^\omega_{jit} + \varepsilon_{jit})] = 0$$

5.5.2 Method 2

The second method is much simpler, but only applicable when the inputs are aggregated to the point where firms are not making (observable) extensive margin input decisions. The key difference is that here all of the model parameters are estimated jointly (excepting of course the selection parameters, which are not relevant if there is no selection in the model). The fact that every firm uses the same acquisition technology (Both) means the production function can be written (in logs using equation 23) as

$$r_{jit} = (1-\theta_i)\log\psi_{jit} + \beta_i\theta_i k_{jit} + \sum_{h\in\mathcal{H}_i}(\alpha_{hi}\theta_i x^Q_{hjit} - \frac{\alpha_{hi}\theta_i}{\rho_{hi}}\log(1-S_{hjit})) + b_{ti} + g_\omega(\omega_{jit-1}) + \eta^\omega_{jit} + \varepsilon_{jit} \tag{31}$$

Here k_{jit}, x^Q_{hjit} and S_{hjt} are correlated with unobserved productivity. Just as with method 1, this equation can be estimated using either WLP or the direct proxy approach described above, with the added twist that since intermediate expenditure x^Q_{hjit} and input shares s_{hjt} are potentially correlated with the innovation to productivity η^ω_{jit}, additional instruments must be included for each term. In practice, I use lagged expenditures and input shares of each type, as these are orthogonal to the productivity innovation by assumption. The procedure otherwise proceeds exactly as in method 1. Note that with a single task (aggregate labor and intermediate indices), this method corresponds to the model in equation 7. I refer to estimates obtained with these two methods as M2-WLP and M2-DP.

6 Results

This section presents the results from the model estimation. First I discuss the input substitution results, followed by a discussion of the resulting elasticities and how they differ from previous estimates in the literature. I follow with a discussion of the productivity estimates, fixed cost distributions and other aspects/implications of the estimated model.

6.1 Input Substitution

I estimate the model described in the previous sections for 4 different industries: Food Products, Wood & Paper Products, Heavy Industry and Extraction, and Tools, Machinery and Consumer Goods (which I also refer to as "Manufacturing"). As discussed in the estimation section, the parameters for each industry-task pair are estimated separately. Since there are 12 different inputs and 4 industries, this means that I am presenting the results from roughly 50 different estimation procedures, each with 19 parameters to estimate, for a total of about 900 parameter estimates. Here I focus on the key objects of interest - the scale and substitution parameters $\alpha_h\theta$ and ρ_h, as well as the associated elasticities of substitution. These results are spread across several tables. Table 4 presents estimates of ρ_h and the elasticities of substitution ϵ_h for the manufacturing industry using the full selection-correction model. Table 5 shows estimates of ρ_h for the same industry, where ρ_h is estimated using several methods to demonstrate the bias from not correcting for endogeneity and selection. Tables 6 and 7 present the scale and substitution parameter estimates for all four industries. Table 8 presents the elasticities of substitution for all the tasks and industries. Table 9 has the results from an aggregate single input task model.

Table 4 shows some of the key results from this paper. The first column lists the input type for which the estimation was done (i.e.: h). The second column is the estimate of ρ_h for that task using the full selection-correction model developed in the previous section. The third column is the elasticity of substitution implied by the point estimate of ρ_h. The estimates of ρ_h are of particular interest, because they allow me to test whether or not the assumptions built into the standard Cobb-Douglas framework hold in the data. Recall that $\rho_h \to 0$ corresponds to Cobb-Douglas, while $\rho_h = 1$ is perfect substitution. The estimates I get in the full model are solidly between the two, ranging from roughly 0.4 to 0.6, which corresponds to elasticities of substitution between 1.6 and 2.5. The results mean that labor and intermediates are gross substitutes, and that substitutability differs significantly across input tasks. Marketing labor is much more substitutable with purchased marketing services than ICT labor is relative to purchased ICT services. The estimates are all very precise and significantly different from both 0 and 1.[30] These results hold across all four industries, as shown in table 7.

Table 5 presents estimates of ρ_h for the same industry using three different estimators. The OLS estimates (column 2) are from a simple regression of the log input ratio ($\log(L_h/X_h^Q)$) on log wages and time dummies (equation 19 without the lagged variables and transformed error term). The IR column refers to estimates using the GMM procedure in section 5.2, where ρ_h is estimated controlling for the endogeneity of wages, but not controlling for selection bias. The last column (IR+SEL) are the results from the full structural model (same as table 4). Clearly the OLS estimates are suffering from significant biases. For the disaggregated inputs, the estimates range from -1.279 to 6.424, most of which are very imprecise. This is due to the correlation between wages and unobserved productivity z_{hjt}. Controlling for endogeneity by transforming the error and instrumenting for wages leads to much more precise and reasonable estimates of ρ_h, with most between 0.8 and 1, which implies that labor and intermediates are highly substitutable. A few inputs have estimates greater than 1, which is difficult to interpret, as $\rho_h = 1$ implies perfect substitutability. These estimates are all relatively imprecise, however, and so we cannot reject the null that they are equal to or less than one. Moving to the selection corrected model as discussed above provides much more reasonable estimates.

The estimates of $\alpha_h\theta$ in table 6 are primarily useful for getting an idea of the relative importance of each task in each industry. For example, it's obvious that manufacturing goods and labor are the biggest input into manufacturing firm's production process, followed by heavy industry inputs (which includes labor and products related to raw materials, chemicals, etc). Of the service inputs, Engineering is the largest, followed by Transportation and ICT. Note that all of these estimates are deflated by the demand parameter θ_i, which for the manufacturing industry I estimate to be 0.804 (from a demand elasticity of -5.124). Thus

[30] I do not report significance stars, since it's not clear which null hypothesis is the most informative ($\rho_h = 0$ vs. $\rho_h = 1$).

(for example) the physical production contribution of heavy industry inputs in manufacturing is actually $0.076/0.804 = 0.095$.

Table 8 shows the estimated elasticities of substitution for all four industries including manufacturing. As hinted by the ρ_h estimates, most of the elasticities lie between 1.5 and 4, which is significantly larger than the Cobb-Douglas case which restricts these elasticities to equal 1. Note that these elasticities are significantly different from what has previously been estimated in the literature at the aggregate level (typically studies find elasticities of substitution < 1 for aggregate labor and aggregate intermediates). I obtain different estimates for several reasons. First, my data is much more disaggregated, exposing intra-task substitution which may be difficult to detect at the aggregate level. Second, I employ a unique strategy for estimating the elasticity parameters, and third, previous studies have not explicitly accounted for extensive margin input selection and outsourcing. The heterogeneity in elasticities suggest that the effects of changes in prices, wages or productivities will have very different effects across different tasks and industries. It's difficult to see any strong patterns, other than that relative substitution elasticities appear to be consistent across industries. ICT is uniformly less substitutable than Transportation in all industries, while Transportation is in turn less substitutable than Marketing across the board. The Food industry is the only one which employs or purchases a significant number of food-related inputs, which is why the other industries lack elasticities for food inputs (i.e.: Food $\notin \mathcal{H}_{\text{Wood}}$).

I also estimate an aggregate input or single-task model, where I allow the aggregate labor and aggregate intermediate input indices to be flexibly substitutable. I estimate this model the same way as the disaggregated model, with the exception that I do not control for selection bias since every firm is at the intensive margin when looking at aggregate labor and intermediates. Notably, the results (table 9) for each industry lie in the same range as the selection-corrected estimates for the disaggregated model and solidly reject the Cobb-Douglas null hypotheses ($\rho = 0$). This implies that controlling for flexible substitution may be important even with commonly available data on aggregate input use, and also that the results from such estimates are not too different from the substitutability results when looking at disaggregated input data. I demonstrate this idea in section 8 where I show that allowing for flexible substitution when estimating the effect of tariffs on productivity provides significantly different results than the Cobb-Douglas benchmark. This is true for both a model with aggregate inputs (the same single-task model I estimate here) and in the disaggregated model.

6.2 Demand Elasticities

Calculating demand elasticities is the focus of a large segment of the literature, since they are an important input into evaluating the effects of various policies which may affect wages or prices. One benefit of my framework is that I can not only recover estimates of own and cross-price elasticities which differ significantly from standard estimates using benchmark cobb-douglas or trans-log models, but also construct aggregate elasticities of demand up from the firm-task level elasticities. In this sense my exercise is similar to Oberfield and Raval (2021) and Raval (2019), except that I am able to estimate elasticities at the disaggregated input level, giving me aggregate elasticities for particular labor and intermediate types across different industries.

Most papers which aim to calculate demand elasticities are interested in the own-price demand elasticity of labor. For example, papers which look at the effects of minimum wages (Kreiner, Reck and Skov (2020)) or the effects of trade and offshoring on labor demand (Senses (2010)). My estimates differ from most of the literature in that I am allowing specific types of labor to be flexibly substitutable with particular types of intermediate goods. This allows tariffs and policies which affect some types of goods or occupations to have heterogeneous effects across firms and demand for other inputs, conditional on firm-level exposure and elasticity terms. To see this, consider the own-price elasticity of demand for labor of type h:

$$\epsilon^{L_{hjt}}_{W_{hjt}} = \frac{\partial L_{hjt}}{\partial W_{hjt}} \frac{W_{hjt}}{L_{hjt}} = \underbrace{-1 - \frac{\alpha_h \theta}{1 - \alpha\theta} S_{hjt}}_{\text{Direct and Scale Effect}} - \underbrace{\frac{\rho_h}{1-\rho_h}(1 - S_{hjt})}_{\text{Substitution Effect}} \tag{32}$$

where $\alpha\theta \equiv \sum_h \alpha_h \theta$. In the Cobb-Douglas case, the own price elasticity is restricted to the direct and scale effect, with $S_{hjt} = 1$. Thus we will always obtain estimates very close to 1, since the scale term is generally quite small. When we additionally allow for substitution (the second term) and firm-level *exposure* to price changes, these elasticities can differ significantly from 1. By *exposure* I mean that the degree to which a firm is sensitive to changes in the wage depends on their input mix. For $\rho_h > 0$, the elasticity is decreasing in the expenditure share of intermediates, meaning that the more a firm has already outsourced, the more sensitive they will be to changes in the wage. The aggregate elasticity of demand is then calculated as the weighted sum of firm-level elasticities across all firms in the industry:

$$\epsilon^{L_{ht}}_{W_{ht}} = \frac{\partial L_{ht}}{\partial \overline{W}_{ht}} \frac{\overline{W}_{ht}}{L_{ht}} = -\frac{1}{1-\rho_h} + \left(\frac{\rho_h}{1-\rho_h} - \frac{\alpha_h \theta}{1-\alpha\theta} \right) \sum_j \frac{L_{hjt}}{L_h t} S_{hjt} \tag{33}$$

where $\overline{W}_h t$ is the mean industry wage for labor h. Note that this term is strictly negative for $\rho_h \in (0,1)$ and the contribution of an individual firm to the aggregate elasticity depends on their industry labor share and own labor input share S_{hjt}. Similarly, a look at the aggregate cross-price elasticity of demand for labor of type h w.r.t. the price of the same-type intermediate,

$$\epsilon^{L_{ht}}_{P_{ht}} = \frac{\partial L_{ht}}{\partial P_{ht}} \frac{P_{ht}}{L_{ht}} = \left(\frac{\rho_h}{1-\rho_h} - \frac{\alpha_h \theta}{1-\alpha\theta} \right) \sum_j \frac{L_{hjt}}{L_h t} (1 - S_{hjt}) \tag{34}$$

demonstrates that firms which only hire labor have no *exposure* to changes in input prices, so the effect of a change in input prices from, say, a change in trade policy on employment will depend on the existing distribution of outsourcing patterns in the economy. Importantly, this formulation allows for a positive response of labor demand to an increase in input prices. In the Cobb-Douglas benchmark, an increase in any price causes an increase in the total input price index for the firm, leading to a decrease in demand for all inputs from the scale effect. Here, an increase in the cost of outsourcing ICT services may increase domestic demand for ICT workers, depending on the value of ρ_{ICT}.

I report the aggregate demand elasticities in table 10. There are several results worth noting. First, there is considerable heterogeneity in aggregate price sensitivity across different input types. Marketing and Transportation labor is much more sensitive to changes in wages than Security and ICT labor. A 10% increase in the wage for transportation labor leads to a 17.1% decrease in demand in 2000. Also, these elasticities are increasing in absolute value over time. Between 2000 and 2011, the elasticity for Cleaning labor increased from -1.38 to -1.84. This is primarily due to a shift away from hiring labor in-house and towards outsourcing – increasing price sensitivity for labor. Looking at the last two columns, it's clear that the cross price elasticities are positive and significantly higher than the Cobb-Douglas benchmark (which would be strictly negative and lie at around the value of $-\alpha_h \rho$ from table 4). Indeed, a 10% increase in the cost of transportation services would increase demand for transportation labor by 6.7%. In the benchmark case, demand for transportation labor would decrease by 0.5%.

Table 11 reports the distribution of own-wage demand elasticities across firms in the middle of my sample (2006). The results show significant differences across firms. While some firms are essentially at the benchmark case of the direct + scale effect close to -1 (see equation 32), others are very sensitive to changes in wages, with elasticity terms ranging as low as -3.11. These estimates are much lower than the standard estimates in the literature (for example, Senses (2010)), which stems from the fact that I estimate that matched intermediates are substitutes for labor, and so the effect of a wage increase is magnified by the tendency of the firm to mitigate the cost increase by substituting away from labor towards intermediates. Note that this also means that the scale effects on other inputs are actually somewhat mitigated by the firm's ability to "cushion the blow" so to speak. In a purely Cobb-Douglas world, the full effect of the price increase would be passed on to the other inputs via the scale effect.

6.3 Demand Results

Table 12 shows the results from the simple logit demand system I estimate from the production data, as described in section 5.1. The key regression is relative market shares at the cn2 level on log price and firm/product dummies. With the exception of the Food industry, all the demand elasticity terms are fairly reasonable and highly significant. Note that I estimate these demand elasticities with the full set of firms for which I observe production data, which is a superset of the firms for which I have detailed input data. A demand elasticity of -5 implies that $\theta \approx 0.8$. These results are in line with standard parameters used in the trade literature, though I estimate them directly from the output quantity and price data.

7 The Effects of an Increase in the Wage Floor in Danish Manufacturing

A key result from my framework is that I'm able to recover not only estimates of flexible substitution and demand elasticities between labor and intermediates, but also extensive-margin outsourcing probabilities for each firm-task pair. This is, as far as I know, unique in the production literature. There's a big literature on outsourcing and offshoring which suggests that firms respond to changes in the relative cost of labor by shifting production of intermediates outside the boundaries of the firm and instead contracting that work out to other firms who may be able to do the work at lower cost (due to higher labor productivity for that task or perhaps access to cheaper labor). Accounting for this mechanism is thus important when studying the employment effects of changes in wages or intermediate prices.

This is particularly salient if one wishes to investigate the effect of a policy such as a minimum wage or union-bargained wage floor. While a number of mechanisms play a role in determining the ultimate equilibrium effect of an increase in minimum wages, a first-order mechanism is the potential dis-employment effect stemming from a decrease in demand for labor. In this section, I use my framework to study the effects of a change in the wage floor in the Danish manufacturing sector on labor demand, prices and output. While a lot of studies have examined this sort of policy from a number of angles, I innovate in several directions. First, by structurally estimating firm-level demand for different types of workers and intermediates, I can estimate changes in demand for disaggregated types of labor. This will matter because low-wage occupations will certainly be affected differently than high wage occupations. Second, The degree to which they are affected also differs across occupations and industries due to differences in substitutability and differences in the distribution of workers across firms, both of which I estimated. Third, and importantly, while most previous studies have only been able to look at intensive margin adjustments based on demand elasticities, I am able to calculate the probability that any give firm outsources all of a given occupation in response to the change in the wage floor. Failing to account for these extensive margin adjustments will lead to an under-estimate of the dis-employment effect of increased wages. Note, however, that my exercise is done in industry equilibrium, in that firms are optimizing their behavior relative to other firms in the industry, but not in general equilibrium. As such, my focus is on just the first-order effect of wage policies on labor demand by firms.

My specific experiment relates to the structure of wage bargaining in Denmark. Denmark does not have an official minimum wage. However, most industries do have binding wage floors which have been negotiated via collective bargaining agreements with one of the major labor unions which operate in Denmark. These wage floors are renegotiated on a regular schedule, and effective wage floors can differ significantly across industries, ranging from 110 to 138 DKK in 2015. My counterfactual will be to simulate an increase in the wage floor for the manufacturing industry (which has one of the lower wage floors) from 110 DKK to 135 DKK, which is roughly the effective wage floor for banks, cinemas and discotheques (I use the basic income figures from Kreiner, Reck and Skov (2020)[31]). I do this both at the

[31] See table A.1 in that paper for a detailed summary of the various collective bargaining agreements across

industry level, and for just a particular occupation group to illustrate the interconnections between input types.

7.1 Construction of the counter-factual wages

One feature of the data is that I see wages for every worker at all of these firms. Thus I can simulate a change in minimum wages at the worker level and aggregate up, or directly change the average wage at the firm level. A look at the wage distribution does indicate that some workers do get paid less than the negotiated basic income wage floors. This could be for a number of reasons including contractual exceptions, temporary workers, young workers (who face a lower wage floor – see Kreiner, Reck and Skov (2020)) or measurement issues. I assume that whatever reasons exist for these wages to be below the wage floor will still exist should the wage floor increase. Thus I create two sets of counter-factual wages. The first attempts to preserve these features of the data by increasing all wages which are lower than the new wage floor by a maximum of 25 kroner (the difference between the old and new wage floors), with a new maximum wage of 135 dkk for any wage which was previously below the new floor. Let an individual n's wage be denoted W_{nt}. Counterfactual wages under the first scheme are:

$$\widehat{W}_{nt} = \begin{cases} W_{nt} & \text{if } W_{nt} > 135 \\ 135 & \text{if } W_{nt} \in [110, 135] \\ W_{nt} + 25 & \text{if } W_{nt} < 110 \end{cases} \tag{35}$$

I then aggregate up to the firm-task level as before, giving me the counterfactual firm-task wages $\widehat{W}_{hjt}$. The second strategy simply sets $\widehat{W}_{hjt} = \max\{W_{hjt}, 135\}$. As an example, figure 7 shows the actual and counterfactual distributions of firm-task wages for transportation labor in the manufacturing industry in 2011, which is the year for which I simulate the increase in the wage floor. Note that the wage floor doesn't bind for a significant subset of firms who were already paying higher wages. I assume that firms take the new wage floors into account when calculating their expected wages.

7.2 Calculation of the outsourcing probabilities

Given the assumptions on the model, calculating the probability that any given firm outsources in response to a change in expected wages is straightforward. Due to the nature of

Danish industries

the experiment (an increase in expected wages) my exposition focuses on the key case of a firm shifting from choosing Both to choosing Buy (which I call "outsourcing"). Thus we are interested in the probability that a firm crosses the cutoff C^1_{hjt} from above. Note that the probability of outsourcing and the intensive margin change in labor demand, both depend on the entire *vector* of wages faced by the firm. The probability of outsourcing task h depends directly on the price/wage ratio for h, but also indirectly on the costs of all other inputs via the optimal scale of the firm. An increase in the cost of input k may lead a firm to reduce its scale in the industry equilibrium. Since the outsourcing decision depends in part on the ratio of fixed costs to input expenditure, a decrease in optimal expenditure can cause a firm to outsource an input, even if the price of that input didn't change. Using the assumptions on task productivity and the fixed costs, I can explicitly calculate the probability of outsourcing any given task for any firm in response to a change in prices. I explain exactly how I do this calculation in appendix G.

The expected change in labor from a given change in own-type wages can then be calculated as:

$$\mathbb{E}[\Delta L_{hjt}] = \Pr(\text{Outsource})(-L_{hjt}) + (1 - \Pr(\text{Outsource}))(\%\Delta W_{hjt} \times \epsilon^{L_{hjt}}_{W_{hjt}}) \tag{36}$$

Note that while this equation is in terms of change in demand for L_{hjt} from a change in W_{hjt}, as mentioned above the demand depends on the entire vector of wages and prices. When doing the full-industry counterfactual, I take these total changes into account. Similarly, the intensive margin change in labor demand is also calculated to take into account the optimal response to the wage changes for all of the labor types employed by the firm.

7.3 Experiment and Results

I perform two different counterfactual exercises, focusing on the increase in the wage floor for all labor in the manufacturing industry. Given the counterfactual wages, I use the above procedure to calculate expected change in labor demand at the firm level, then aggregate up to get the expected percent change in aggregate demand for labor in the manufacturing industry for each labor type. The primary purpose of this exercise is to quantify the importance of accounting for flexible substitution and outsourcing when evaluating changes in wage policy such as minimum wages and wage floors set by collective bargaining. As such, I calculate the change in labor demand for three cases. Table 13 shows the results for a subset of the occupation types. The first case ("C-D"), which is analogous to the disaggregated Cobb-Douglas specification in section 4 assumes that $rho_h = 0$ such that elasticities of substitution are constrained to the scale effects discussed in section 6.2, own-price demand elasticities are 1, and there is no outsourcing. I consider this the benchmark case. The second case ("Subst.") shows the change in labor demand if we take the flexible demand elasticities into account ($\rho_h \neq 0$) but ignore the possibility of outsourcing. The third case uses the estimated substitution patterns along with the outsourcing probabilities.

The results show significant heterogeneity across labor types in response to the increased wage floor. Because the Cobb-Douglas specification shuts down the firm's ability to substitute away from labor, the results in the C-D column represent two effects – the direct own-wage demand elasticity, which equals 1, plus the scale effect from the total change in the input price index for the firm, which depends on the scale and demand parameters $\alpha_h \theta$ as well as the distribution of employment and wage changes. For example, demand for cleaning labor declines by -10.4%, which is much higher than other occupations, because cleaning labor is generally lower wage than other labor so the firm-task price or cleaning labor increases more due to the wage floor increase than for other labor types (a mean increase of 12.9%). The other labor types have smaller changes since fewer firms experience increases in their labor costs due to the increased wage floor. ICT is a high-wage labor type, and so the mean increase in firm-level wage is only 0.8%, with a decrease in labor demand of 1.1%.

Including the substitution effect (the "Subst" column) has two effects. First, since firms can now substitute away from labor towards intermediates, the own-price demand elasticity is much higher, increasing the dis-employment effect of the wage floor. However, this ability for firms to mitigate the increased cost of labor means that an increase in the wage for transportation occupations has a smaller effect on the firm's overall input price index, and thus a smaller effect on demand for all other types of labor relative to the Cobb-Douglas benchmark. The effect depends on the substitutability of the task. ICT has a low elasticity of substitution (1.59 in the manufacturing industry) and so allowing for substitution only decreases aggregate demand slightly relative to the benchmark. Marketing and Transportation, however, are much more substitutable, with elasticities of 2.09 and 2.59 respectively. Thus moving to the flexible substitution framework nearly doubles the effect of the minimum wage relative to the benchmark.

Moving to the total specification, we can see that for some labor types, the probability of outsourcing is very low, and so there's little difference between the substitution case and the full framework. ICT and Marketing are barely affected. Cleaning, Transportation and Heavy Industry occupations are more heavily affected. Failing to account for outsourcing would underestimate the dis-employment effects of the wage floor relative to the flexible substitution case by 20%, 13% and 29% respectively. Compared to the benchmark case, these differences are 116%, 98% and 130%. When looking at total change in labor demand, accounting for flexible substitution and outsourcing increases the effect by 40%. Thus it's clear that accounting for both flexible substitution and outsourcing are vital for estimating the effects of a wage policy on employment.

These numbers may overestimate the actual dis-employment effect of a 25 kroner increase in the wage floor, since this exercise does not account for equilibrium changes in wages or prices in other industries. A decrease in labor demand in manufacturing due to increased wages will increase the labor supply in other industries, pushing equilibrium wages back down. On the other hand, if firms substitute away from labor towards intermediates, as is the case in my framework, then output demand will grow in other industries, increasing labor demand and pushing equilibrium wages up. However, my framework does account for several important effects. First, the direct change in labor demand by each firm in the industry due to the increase in wages. Second, because I model and estimate the industry equilibrium, the firm's input adjustments do take into account changes in optimal firm scale which is determined by the shape of output demand in the industry (θ_i). Also, while my framework doesn't account for equilibrium changes in wages or labor supply, it will be a vital component of any such general equilibrium approach, as failing to account for input substitution and outsourcing will lead to significantly under-estimated employment effects relative to the benchmark model.

My current work is on extending this experiment to calculate the effect of a wage floor on the other industries in my data via intermediate price linkages. As mentioned above, an increase in the wage floor will increase the output price index for the industry, as well as increase demand for intermediate output from other industries. Both of these effects feed into the optimal behavior of firms in the other industries, affecting labor demand and output. My framework allows the calculation of these effects, which will be included in the next draft of this paper.

8 Estimating the Effect of Tariff Protection on Productivity with Flexibly Substitutable Inputs

This paper proposes that failing to account for flexible substitution and outsourcing leads to imprecise or misleading estimates of firm productivity. This is of particular importance for empirical studies where the goal is to estimate, for example, the effects of trade policy, market competition or R&D on the evolution of firm efficiency. In this section I focus on one of these questions: the effect of tariff protection on productivity. There's a significant recent literature looking at this and related questions, where a large part of the focus is on estimating models of production while controlling for the difficulties which arise due to unobserved prices and markups, or when firms produce multiple distinct products. De Loecker (2011) examines the effects of quota protection on firm productivity while controlling for unobserved variation in output prices. Dhyne et al. (2017) examines the effects of import competition on productivity using production data which allows them to avoid the problem of unobserved prices while separately identifying product-level production functions. Similarly De Loecker et al. (2016) uses data on Indian manufacturing to look at the effect of trade liberalizations and tariff reductions on prices and markups.

To investigate the importance of controlling for input substitution when estimating productivity, I conduct a study similar to De Loecker (2011). In particular, I use data on Danish manufacturing firms and product-level tariffs to investigate the effect of tariff reductions between 2000 and 2006 on firm-level productivity. I first do a very similar exercise to De Loecker (2011) where I use revenue shares and product mix as a method of stripping price variation out of the productivity term. My results are surprisingly quite similar to his, despite some small methodological differences and the completely different data. I then move to my setting where I additionally control for input substitution at both the aggregate input level and with disaggregated inputs, finding that input variation tends to bias the results in the opposite direction from price variation. In the rest of this section I describe how I construct the data, then compare my estimation strategy to the benchmark in the literature, and finally discuss the results.

8.1 Data

For this exercise I focus on just the tools, machinery and goods industry, which I refer to as "manufacturing". The goal is to estimate the effect of changes in firm-specific tariff protection on firm productivity. To do this, I start with a database of tariff lines for Denmark obtained from the World Bank's WITS database[32]. This data has product-level tariffs at the HS6 level for each year and country with which Denmark has trade agreements. I construct the firm-level tariff exposure term as follows. Let $\tau_{gct} \in 0, 1$ indicate whether or not there exists an effective (AHS) tariff on imports of good g (at the hs6 level) from country c in year t. Define λ_{gc} as the value share of country c in total world trade of product g in 1999[33], which I define as the pre-sample period. I then construct product-level tariffs for product g in year t as $\tau_{gt} = \sum_c \lambda_{gc}\tau_{gct}$. The firm-level tariff protection is then $\tau_{jt} = \sum_{g(j)} \lambda_{gj}\tau_{gt}$, where λ_{gj} is the product/revenue share of good g for firm j, and each firm sums over the set of goods which it produces. The product share weights are constructed using production data for the Danish manufacturing sector, where for each firm I observe output and revenues for each good at the hs6 level[34].

I follow De Loecker (2011) in constructing sector-level demand shifters as a market share weighted average of product-level revenue: $q_{st} = \sum_j ms_{jst}r_{jst}$ where ms_{jst} is the firm's market share for sector s aggregated up to the 2-digit (hs) level, and r_{jst} are firm level sales for goods in that sector. Unlike De Loecker, I do observe these variables at the firm-product level, and thus can construct these demand shifters directly from the data. The firm-specific total demand shifter is then a revenue-share weighted sum of the total demand shifters across segments $\sum_s \beta_s rs_{sjt}q_{st}$, where the β_s coefficients are to be estimated. Note that by following De Loecker in the construction of these demand terms, I am implicitly making the same assumptions as him in regards to input proportionality across products.

8.2 Estimation

In this section I describe the "benchmark" specification, which will closely follow the strategy developed by De Loecker (2011), and then outline how I apply the new framework developed in this paper. The purpose of this exercise is not a full replication of De Loecker's strategy, but rather to conduct a similar exercise to establish baseline estimates, and then investigate how those estimates change when additionally taking input variation into account. The

[32] Available at https://wits.worldbank.org/

[33] Obtained from the United Nations Comtrade Database, available at https://comtrade.un.org

[34] I weight the tariff protection terms using both revenue shares and simple averages. The results are essentially the same

baseline model is an aggregate-input Cobb Douglas revenue production function:

$$R_{jt} = \psi_{jt}^{1-\theta} K_{jt}^{\beta\theta} L_{jt}^{\alpha_L \theta} Q_{jt}^{\alpha_Q \theta} e^{\omega_{jt}\theta} \tag{37}$$

where following the standard procedure, L_{ht} is labor hours and Q_{jt} is deflated expenditure on intermediates. I allow tariffs to potentially affect revenues through both demand and productivity as follows. I assume the demand shock takes the form $\log \psi_{jt} \equiv q_{jt} + a_1 \tau_{jt} + \xi_j + \tilde{\xi}_{jt}$ where $q_{jt} \equiv \sum_s \beta_s rs_{sjt} q_{st}$. Productivity follows the same assumptions as in the body of the paper, with the additional assumption that lagged tariff protection may affect the evolution of productivity, so $\omega_{jt} = g_\omega(\omega_{jt-1}, \tau_{jt-1}) + \eta_{jt}^{\omega}$. The procedure is obtain estimates of ω_{jt} while controlling for demand variation and contemporaneous tariffs using WLP, and then estimating the effects of tariffs on productivity with a simple regression of $\hat{\omega}_{jt}$ on lagged productivity and lagged tariffs, i.e.: $\omega_{jt} = \delta_\omega \omega_{jt-1} + a_2 \tau_{jt-1} + \eta_{jt}^{\omega}$.

In addition to this benchmark model, I estimate productivity using both methods described in 5.5. In particular, I estimate productivity in a model with aggregate labor and aggregate intermediates using method 2 (joint estimation of all model parameters), and the full disaggregated model from the main section with method 1, where I first estimate the input scale and substitution parameters and then estimate the remainder of the model using WLP[35]. The key difference from the strategy outlined in those sections is that now tariffs and demand terms are included in the WLP control function.

8.3 Results: The Effect of tariffs on Productivity

This exercise is done over a period in which tariff protection (and this import competition) fell for Danish firms. The mean value of the firm-specific tariff exposure term dropped from 0.77 in 2000 to 0.22 in 2006. Similarly, the mean drop in tariff protection over this period was 0.52. Recall that a value of $\tau_{jt} = 1$ indicates that every good produced by the firm has an associated effective tariff applied to foreign imports of that good from every country. A value of zero means none of the firm's products enjoy shelter from import tariffs.

I estimate several models. The first model (WLP) estimates productivity in an aggregate input model using WLP while omitting the demand terms and parameters ψ_{jt} and θ. The second model (DL) is the benchmark model discussed in the previous section, where I control

[35] I also estimate the results using the direct proxy (DP) approach and get similar results. I omit them from the results section for ease of exposition.

for demand, but not input substitution. The third model is the aggregate-input matched ces (MC) model with the demand shifters omitted. The MC model is estimated with the M2-WLP method. This provides an idea of how input substitution biases productivity separate from the demand effect. The fourth model (DL-MC) is the aggregate input matched ces with demand shifters included, estimated using M2-WLP. The fifth model (DL-MC2) is the disaggregated input matched ces with demand shifters, estimated using the two-step M1-WLP method.

The results are shown in table 14. The first two rows roughly mirror the results from De Loecker (2011). Ignoring the effects of demand and output prices, the effect of tariffs on firm productivity is negative, statistically significant and equal to -0.058. The interpretation is that eliminating the tariffs on all products would raise productivity by about 6 percent. When controlling for unobserved price variation (DL), this effect drops in magnitude to -0.026, which is in line with De Loecker's results. Thus failing to control for prices will lead to overestimates of the effect of tariffs on productivity. The third row (MC) moves to a specification without price controls, but where I do control for input substitution. Here the effect of tariffs is much larger in magnitude, at -0.103. This suggests that failing to control for substitution will lead to *underestimates* of the effects of tariffs on productivity. i.e.: the bias from input substitution moves in the opposite direction as the price effect. This is born out in the combined model (DL-MC) where I control for both the price effect and the substitution effect. The estimate here lies between the price-effect-only estimate and the substitution-only estimate, suggesting that moving from full tariffs to no tariffs would raise productivity by 7 percent (a parameter estimate of -0.069). This is similar in magnitude to the naive WLP estimation, but only by coincidence, as the price and substitution effects move in opposite directions. Finally, the last row shows the estimate when controlling for price and substitution in the full model, where I use the full disaggregated model and the two-step estimator (M1-WLP). While these results are not directly comparable to the other 4, since it requires much more detailed data and a significantly different estimation strategy, it also provides estimates of the tariff effect on productivity which are significantly higher than the DL estimate, at -0.046.

This exercise has made it clear that the bias which results from ignoring input substitution can be significant. Depending on the estimation method and the data available, the estimated tariff effect is as much as double or triple the magnitude of estimates obtained when only controlling for unobserved prices. This stresses the need to control for both effects when estimating productivity.

9 Conclusion

I develop a new method for modeling and estimating production with disaggregated inputs, flexible substitution patterns between labor and intermediates, and extensive-margin outsourcing. I motivate this framework by using detailed input use data from Denmark to show that firms substitute along the intensive and extensive margins between labor and intermediate goods/services – facts which cannot be rationalized using standard models of production. Applying my framework to this data, I estimate that labor and intermediate goods are gross substitutes, with the elasticity of substitution ranging from 1.5 to 4 across tasks and industries. My framework also generates positive cross-price elasticities of demand between matched labor and intermediates, ranging from 0 to 2 at the firm-task level. I aggregate these up and show that trends in outsourcing have caused an increase in demand elasticities and price sensitivity across all input types between 2000 and 2011.

I demonstrate the importance of accounting for substitution and outsourcing by applying my framework to several empirical applications. First, I examine the effect of a 25 kroner (about $4 USD) increase in minimum wage in the Danish manufacturing industry. I estimate that demand for labor drops by about 4.2%. Shutting down the outsourcing and substitution channels results in an estimate that is 40% lower. I also estimate the effect of a change in tariffs in Denmark during this period to estimate the effects of trade protection on technical efficiency. Ignoring input substitution biases estimated effects of tariffs on productivity downward. When controlling for both price effects and substitution, I estimate that removing all tariffs would result in a 6.9% increase in productivity, which is almost triple the estimate obtained when controlling for only price effects.

The main contributions of this paper are methodological and empirical. Demonstrating how firms adjust on the intensive and extensive margin in response to changes in firm productivity and prices, and that disaggregated labor and intermediates are imperfect substitutes, has important implications for questions related to trade and labor policy. Similarly, researchers interested in the evolution of firm productivity and its relationship with policies and market conditions such as tariffs or competition may wish to account for unobserved productivity. Estimating production functions at the disaggregated level also allows researchers to link and quantify the effects of a policy related to a particular industry or occupation, to the heterogeneous effects on any other industry or occupation.

References

Acemoglu, Daron, and David H. Autor. 2011. "Handbook of Labor Economics." Chapter Skills, tasks and technologies, 1043–1171. Amsterdam [u.a.]:North-Holland.

Acemoglu, Daron, Vasco M Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. 2012. "The network origins of aggregate fluctuations." *Econometrica*, 80(5): 1977–2016.

Antras, Pol. 2004. "Is the US aggregate production function Cobb-Douglas? New estimates of the elasticity of substitution." *Contributions to Macroeconomics*, 4(1): 1161.

Bagger, Jesper, Bent Jesper Christensen, and Dale T Mortensen. 2014. "Wage and labor productivity dispersion: The roles of total factor productivity, labor quality, capital intensity, and rent sharing."

Carvalho, Vasco M, Makoto Nirei, Yukiko U Saito, and Alireza Tahbaz-Salehi. 2021. "Supply chain disruptions: Evidence from the great east japan earthquake." *The Quarterly Journal of Economics*, 136(2): 1255–1321.

Chan, Mons. 2017. "Firm Organization and the Transmission of Shocks."

Chan, Mons, and Ming Xu. 2017. "Trade, Occupation Sorting and Inequality." *University of Minnesota, Working Paper.*

Cobb, Charles W, and Paul H Douglas. 1928. "A theory of production." *The American Economic Review*, 18(1): 139–165.

De Loecker, Jan. 2011. "Product differentiation, multiproduct firms, and estimating the impact of trade liberalization on productivity." *Econometrica*, 79(5): 1407–1451.

De Loecker, Jan, Pinelopi K Goldberg, Amit K Khandelwal, and Nina Pavcnik. 2016. "Prices, markups, and trade reform." *Econometrica*, 84(2): 445–510.

Dhyne, Emmanuel, Amil Petrin, Valerie Smeets, and Frederic Warzynski. 2017. "Multi Product Firms, Import Competition, and the Evolution of Firm-product Technical Efficiencies." National Bureau of Economic Research.

Doraszelski, Ulrich, and Jordi Jaumandreu. 2013. "R&D and productivity: Estimating endogenous productivity." *Review of Economic Studies*, 80(4): 1338–1383.

Doraszelski, Ulrich, and Jordi Jaumandreu. 2018. "Measuring the bias of technological change." *Journal of Political Economy*, 126(3): 1027–1084.

Eaton, Jonathan, Samuel Kortum, Francis Kramarz, et al. 2016. "Firm-to-Firm Trade: Imports, exports, and the labor market."

Gandhi, Amit, Salvador Navarro, and David A Rivers. 2020. "On the identification of gross output production functions."

Goldschmidt, Deborah, and Johannes F Schmieder. 2017. "The rise of domestic outsourcing and the evolution of the German wage structure." *The Quarterly Journal of Economics*, qjx008.

Griliches, Zvi, and Vidar Ringstad. 1971. *Economies of Scale and the Form of the Production Function.* North-Holland Publishing Company.

Halpern, László, Miklós Koren, and Adam Szeidl. 2015. "Imported inputs and productivity." *The American Economic Review*, 105(12): 3660–3703.

Hansen, Lars Peter. 1982. "Large sample properties of generalized method of moments estimators." *Econometrica*, 1029–1054.

Hausman, Jerry A. 1996. "Valuation of new goods under perfect and imperfect competition." In *The economics of new goods*. 207–248. University of Chicago Press.

Heckman, James. 1979. "Sample specification bias as a selection error." *Econometrica*, 47(1): 153–162.

Helpman, Elhanan, Marc Melitz, and Yona Rubinstein. 2008. "Estimating trade flows: Trading partners and trading volumes." *The Quarterly Journal of Economics*, 123(2): 441–487.

Hummels, David, Rasmus Jørgensen, Jakob Munch, and Chong Xiang. 2014. "The wage effects of offshoring: evidence from danish matched worker-firm data." *American Economic Review*, 104(6): 1597–1629.

Humphrey, Thomas M. 1997. "Algebraic production functions and their uses before Cobb-Douglas." *Economic Quarterly*, 83(1): 51–83.

Klette, Tor Jakob, and Zvi Griliches. 1996. "The inconsistency of common scale estimators when output prices are unobserved and endogenous." *Journal of applied econometrics*, 4: 343–361.

Kreiner, Claus Thustrup, Daniel Reck, and Peer Ebbesen Skov. 2020. "Do lower minimum wages for young workers raise their employment? Evidence from a Danish discontinuity." *Review of Economics and Statistics*, 102(2): 339–354.

Krusell, Per, Lee E Ohanian, José-Víctor Ríos-Rull, and Giovanni L Violante. 2000. "Capital-skill complementarity and inequality: A macroeconomic analysis." *Econometrica*, 68(5): 1029–1053.

Levinsohn, James, and Amil Petrin. 2003. "Estimating production functions using inputs to control for unobservables." *The Review of Economic Studies*, 70(2): 317–341.

Long, John B, and Charles I Plosser. 1983. "Real business cycles." *Journal of political Economy*, 91(1): 39–69.

Nadiri, M Ishaq. 1982. "Producers theory." *Handbook of mathematical economics*, 2: 431–490.

Nevo, Aviv. 2001. "Measuring market power in the ready-to-eat cereal industry." *Econometrica*, 69(2): 307–342.

Oberfield, Ezra, and Devesh Raval. 2021. "Micro data and macro technology." 2, Wiley Online Library.

Olley, G. Steven, and Ariel Pakes. 1996. "The dynamics of productivity in the telecommunications equipment industry." *Econometrica*, 64(6): 1263–1297.

Raval, Devesh R. 2019. "The micro elasticity of substitution and non-neutral technology." 1, Wiley Online Library.

Rosen, Sherwin. 1978. "Substitution and division of labour." *Economica*, 45(179): 235–250.

Senses, Mine Zeynep. 2010. "The effects of offshoring on the elasticity of labor demand." *Journal of International Economics*, 81(1): 89–98.

Stole, Lars A, and Jeffrey Zwiebel. 1996. "Intra-firm bargaining under non-binding contracts." *Review of Economic Studies*, 63(3): 375–410.

Varian, Hal R. 1992. *Microeconomic analysis.* . 3. ed ed., New York [u.a.]:Norton.

Wooldridge, Jeffrey M. 2009. "On estimating firm-level production functions using proxy variables to control for unobservables." *Economics Letters*, 104(3): 112–114.

Table 1: Input Usage by Input Type for the **Tools, Machinery and Goods** Industry.

Input Type	Only Hires	Only Buys	Both	Neither
Transportation & Storage	106	3,220	2,502	137
Information Communications Tech.	88	3,015	2,747	115
Legal & Accounting	103	2,595	3,173	94
Architecture & Engineering	1,531	549	3,405	480
Marketing & Sales	68	4,722	855	320
Training & Employment	110	4,339	823	693
Cleaning & Maintenance	374	2,040	3,381	170
Wood & Related	88	4,061	1,252	564
Heavy Industry & Extraction	60	1,713	3,982	210
Tools, Machinery, Goods	1,561	-	4,381	-
Total Observations: 5,965. Observations with all inputs: 0				

Note: Each cell of this table contains the number of firm-year observations where the firm either hires some positive quantity of labor of a particular type, or spends some positive amount of money on an intermediate of the same type, or both/neither. Labor and Intermediate Goods/Services are matched to "Input Type" or task using the matching algorithm I describe in section 2.2 and appendix A.

Table 2: **Materials to Labor Ratio**

VARIABLES	(1) M/L	(2) M/L	(3) M/L	(4) M/L (Srv.)	(5) M/L (Mfr.)	(6) M/K	(7) M/I
Year	0.0146***	0.0141***	0.0204***	0.0166***	0.0221***	0.0298***	0.0696***
	(0.000168)	(0.000169)	(0.000387)	(0.000602)	(0.000283)	(0.000354)	(0.00125)
Observations	293,069	278,373	47,919	33,514	64,641	276,377	160,043
Firm FE	YES	YES	YES	YES	YES	YES	YES
Addtl. Controls	NO	YES	YES	YES	YES	YES	YES
Firm Size	> 10	> 10	> 50	> 10	> 10	> 10	> 10
Industries	All	All	All	Services	Manufacturing	All	All

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Firm size refers to total employment. Additional controls include revenues capital stock and firm size.

The dependent variables (M/L) are all expressed in logs, so the coefficient estimates represent percentage change.

Table 3: **Primary Labor to Total Labor Ratio**

VARIABLES	(1) H/L	(2) H/L	(3) H/L	(4) H/L (Srv.)	(5) H/L (Mfr.)	(6) M/(L-H)	(7) M/L
year	0.0233***	0.0261***	0.0278***	0.0456***	0.0189***	0.0380***	0.0213***
	(0.000429)	(0.000441)	(0.000970)	(0.000944)	(0.000482)	(0.000354)	(0.000264)
Observations	108,358	101,835	26,446	35,560	65,966	101,816	102,029
Firm FE	YES	YES	YES	YES	YES	YES	YES
Addtl. Controls	NO	YES	YES	YES	YES	YES	YES
Firm Size	> 10	> 10	> 50	> 10	> 10	> 10	> 10
Industries	Matched	Matched	Matched	Services	Manufacturing	Matched	Matched

Standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Notes: Firm size refers to total employment. Additional controls include revenues capital stock and firm size. Here H is industry-matched primary labor and L is total labor. Matched Services are: Transportation Storage ICT Legal Accounting Architecture Engineering Marketing Training Security Cleaning. Manufacturing includes Food Textiles Clothing Wood Paper Extraction Tools Furniture Machinery Consumer Goods.

The dependent variables (H/L) are all expressed in logs, so the coefficient estimates represent percentage change.

Table 4: Substitution Parameters and Elasticities for **Tools, Machinery and Goods** Industry.

Input Type	Substitution Parameter (ρ_h)	Elasticity of Substitution (ϵ_h)
Transportation	0.521	2.09
	(0.011)	
ICT	0.371	1.59
	(0.019)	
Legal & Accounting	0.545	2.20
	(0.040)	
Engineering	0.584	2.40
	(0.003)	
Marketing	0.614	2.59
	(0.007)	
Employment and Training	0.507	2.03
	(0.025)	
Security	0.433	1.76
	(0.013)	
Cleaning and Maint.	0.478	1.92
	(0.005)	
Other Services	0.495	1.98
	(0.001)	
Wood and Paper	0.472	1.89
	(0.006)	
Heavy Industry	0.424	1.74
	(0.014)	
Tools, Machinery, Goods	0.590	2.44
	(0.053)	

Note: This table provides estimates of the substitution parameter (ρ_h) from the full model controlling for endogeneity and selection (equation 20). $\rho_h \to 0$ corresponds to the Cobb-Douglas benchmark while $\rho_h \to 1$ implies perfect substitution and $\rho_h \to -\infty$ implies Leontief complementarity. Standard Errors are reported in parentheses. Elasticities of substitution (ϵ_h) are simple transformations of ρ_h. Specifically, $\epsilon_h = 1/(1-\rho_h)$. See section 6.1 for further discussion of these results.

Table 5: Comparison of Estimated Substitution Parameters for **Tools, Machinery and Goods** Industry under different estimation assumptions.

	(OLS)	(IR)	(IR+SEL)
Input Type	ρ_h	ρ_h	ρ_h
Transportation	2.642	0.934	0.521
	(0.460)	(0.049)	(0.011)
ICT	2.983	0.815	0.371
	(0.442)	(0.063)	(0.019)
Legal & Accounting	-0.205	0.884	0.545
	(0.138)	(0.031)	(0.040)
Engineering	3.556	0.871	0.584
	(1.098)	(0.086)	(0.003)
Marketing	6.683	0.869	0.614
	(5.714)	(0.224)	(0.007)
Employment and Training	-1.279	1.015	0.507
	(1.399)	(0.029)	(0.025)
Security	2.084	1.055	0.433
	(0.222)	(0.043)	(0.013)
Cleaning and Maint.	-0.744	0.830	0.478
	(0.293)	(0.053)	(0.005)
Other Services	0.165	0.916	0.495
	(0.061)	(0.028)	(0.001)
Wood and Paper	2.442	0.821	0.472
	(0.653)	(0.094)	(0.006)
Heavy Industry	6.424	0.689	0.424
	(5.260)	(0.226)	(0.014)
Tools, Machinery, Goods	0.466	0.827	0.590
	(0.059)	(0.126)	(0.053)

Note: This table provides estimates of the substitution (ρ_h) parameter under different estimation assumptions. Results in the OLS column are estimated using equation 18 (in logs) without controlling for endogeneity or selection. Column (IR) provides GMM results from equation 19 controlling for endogeneity but not selection. Column (IR+SEL) are results from the full model controlling for endogeneity and selection (equation 20). $\rho_h \to 0$ corresponds to the Cobb-Douglas benchmark. Standard Errors are reported in parentheses. See section 6.1 for further discussion of these results.

Table 6: Scale Parameters ($\alpha_{hi}\theta_i$) for all Four Industries

	Industry			
Input Type	Food	Wood	Heavy	Machinery
Transportation	0.045	0.041	0.045	0.024
	(0.002)	(0.002)	(0.002)	(0.001)
ICT	0.008	0.008	0.013	0.022
	(0.000)	(0.001)	(0.000)	(0.001)
Legal & Accounting	0.007	0.007	0.014	0.013
	(0.000)	(0.001)	(0.001)	(0.001)
Engineering	0.011	0.017	0.038	0.043
	(0.001)	(0.002)	(0.002)	(0.001)
Marketing	0.030	0.009	0.022	0.019
	(0.001)	(0.000)	(0.002)	(0.001)
Employment & Training	0.001	0.001	0.002	0.003
	(0.000)	(0.000)	(0.000)	(0.001)
Security	0.003	0.001	0.004	0.009
	(0.000)	(0.000)	(0.001)	(0.001)
Cleaning	0.010	0.004	0.005	0.006
	(0.001)	(0.000)	(0.000)	(0.001)
Other Services	0.132	0.153	0.145	0.137
	(0.004)	0.003	(0.002)	(0.001)
Food	0.121			
	(0.005)			
Wood & Paper	0.008	0.107	0.009	0.048
	(0.001)	(0.005)	(0.000)	(0.003)
Heavy Industry	0.041	0.067	0.128	0.076
	(0.002)	(0.004)	(0.003)	(0.002)
Tools, Machinery, Goods	0.033	0.071	0.072	0.164
	(0.001)	(0.003)	(0.003)	(0.002)

Note: Standard Errors are reported in parentheses. The columns for Wood, Heavy Industry and Tools, Machinery, Goods do not have estimates for the Food Products input task as these industries do not typically buy or employ that input. See section 6.1 for further discussion.

Table 7: Substitution Parameters (ρ_{hi}) for all Four Industries

Input Type	Food	Wood	Heavy	Machinery
	Industry			
Transportation	0.596	0.555	0.573	0.521
	(0.000)	(0.011)	(0.023)	(0.011)
ICT	0.541	0.421	0.544	0.371
	(0.002)	(0.032)	(0.010)	(0.019)
Legal & Accounting	0.488	0.563	0.510	0.545
	(0.008)	(0.027)	(0.033)	(0.040)
Engineering	0.531	0.577	0.550	0.584
	(0.024)	(0.008)	(0.025)	(0.003)
Marketing	0.693	0.719	0.681	0.614
	(0.065)	(0.091)	(0.066)	(0.007)
Employment & Training	0.528	0.621	0.588	0.507
	(0.051)	(0.063)	(0.015)	(0.025)
Security	0.422	0.452	0.505	0.433
	(0.026)	(0.019)	(0.017)	(0.013)
Cleaning	0.581	0.457	0.490	0.478
	(0.016)	(0.133)	(0.121)	(0.005)
Other Services	0.402	0.896	0.572	0.495
	(0.000)	(0.383)	(0.000)	(0.001)
Food	0.741			
	(0.097)			
Wood & Paper	0.446	0.495	0.415	0.472
	(0.156)	(0.010)	(0.031)	(0.006)
Heavy Industry	0.469	0.393	0.516	0.424
	(0.043)	(0.007)	(0.013)	(0.014)
Tools, Machinery, Goods	0.614	0.665	0.543	0.590
	(0.092)	(0.039)	(0.070)	(0.053)

Note: These estimates of ρ_h are all from the full specification which corrects for selection bias (equation 20). Standard Errors are reported in parentheses. The columns for Wood, Heavy Industry and Tools, Machinery, Goods do not have estimates for the Food Products input task as these industries do not typically buy or employ that input. The Cobb-Douglas benchmark is $\rho_h \to 0$ while $\rho_h \to 1$ implies perfect substitution and $\rho_h \to -\infty$ implies Leontief complementarity. See section 6.1 for further discussion of these results.

Table 8: Elasticities of Substitution between labor and intermediate inputs for each input type, in each industry.

Input Type	Industry: Food	Wood	Heavy	Machinery
Transportation	2.48	2.25	2.34	2.09
ICT	2.18	1.73	2.19	1.59
Legal & Accounting	1.95	2.29	2.04	2.20
Engineering	2.13	2.36	2.22	2.40
Marketing	3.26	3.56	3.13	2.59
Employment & Training	2.12	2.64	2.43	2.03
Security	1.73	1.82	2.02	1.76
Cleaning	2.39	1.84	1.96	1.92
Other Services	1.67	9.62	2.34	1.98
Food	3.86			
Wood & Paper	1.81	1.98	1.71	1.89
Heavy Industry	1.88	1.65	2.07	1.74
Tools, Machinery, Goods	2.59	2.99	2.19	2.44

Note: Elasticities of substitution (ϵ_{hi}) are simple transformations of estimated elasticity parameters from table 7. In particular, $\epsilon_{hi} = 1/(1-\rho_{hi})$. See section 6.1 for further discussion of these results.

Table 9: Estimates for Aggregate Input (Single-Task) Model.

Parameter	Food	Wood	Heavy	Machinery
$\alpha\theta$	0.450	0.486	0.497	0.564
	(0.001)	(0.000)	(0.001)	(0.000)
ρ	0.417	0.653	0.568	0.620
	(0.360)	(0.078)	(0.087)	(0.055)
Elasticity of Subs. (ϵ)	1.715	2.882	2.315	2.632
Observations	1,525	806	2,417	5,963

Note: This table contains estimates of scale ($\alpha\theta$) and substitution (ρ) parameters for the model with a single aggregate input task. As in other tables, the elasticity of substitution term is calculated as $\epsilon = 1/(1-\rho)$. Standard errors are in parenthesis. Total observations is over all firms in the industry which I see in the detailed data. See section 6.1 for further discussion.

Table 10: Aggregate Price Elasticities of Demand for Labor in the **Tools, Machinery and Goods** Industry.

	Own-Wage ($\epsilon^{L_{ht}}_{W_{ht}}$)		**Cross-Price** ($\epsilon^{L_{ht}}_{P_{ht}}$)	
Input Type	2000	2011	2000	2011
Transportation	-1.71	-1.82	0.67	0.79
ICT	-1.19	-1.23	0.17	0.21
Legal & Accounting	-1.58	-1.67	0.57	0.66
Engineering	-1.38	-1.45	0.34	0.41
Marketing	-2.27	-2.35	1.24	1.34
Employment & Training	-1.55	-1.54	0.55	0.54
Security	-1.02	-1.07	0.02	0.06
Cleaning	-1.38	-1.84	0.37	0.84
Other Services	-1.44	-1.5	0.28	0.34
Wood & Paper	-1.16	-1.22	0.1	0.17
Heavy Industry	-1.36	-1.36	0.27	0.28
Tools, Machinery, Goods	-1.33	-1.37	0.14	0.17

Note: Elasticity terms are with respect to the own-type wage and own-type intermediate price. For example, the first row shows price elasticities of demand for transportation labor with respect to wages for transportation labor as well as the market price P_{ht} of transportation services. Aggregate elasticities are calculated as the labor-share weighted sum of firm-level elasticity terms. See section 6.2 for discussion and derivations.

Table 11: Distribution of Own-Wage Elasticities of Demand across Firms in the **Tools, Machinery and Goods** Industry in 2006.

Input Type	**Mean**	**St. Dev.**	**Min**	**Max**
Transportation	-2.01	0.20	-2.08	-1.02
ICT	-1.48	0.15	-1.59	-1.03
Legal & Accounting	-1.89	0.35	-2.19	-1.01
Engineering	-1.57	0.51	-2.41	-1.05
Marketing	-3.03	0.39	-3.16	-1.02
Employment & Training	-1.93	0.25	-2.03	-1.01
Security	-1.52	0.33	-1.76	-1.01
Cleaning	-2.43	0.71	-3.11	-1.01
Other Services	-1.49	0.15	-1.92	-1.16
Wood & Paper	-1.78	0.25	-1.89	-1.06
Heavy Industry	-1.64	0.17	-1.74	-1.09
Tools, Machinery, Goods	-1.42	0.33	-2.43	-1.19

Note: Elasticity terms are with respect to the own-type wage.

Table 12: Estimates of Demand Elasticities from Firm Production Data.

Variable	Food	Wood	Heavy	Machinery
Log Price	-8.560	-3.847***	-5.360***	-5.124***
	(5.880)	(0.295)	(0.764)	(0.899)
Observations	2,321	2,430	4,554	16,361

Note: Demand elasticity terms are estimated using firm-product level data on output and prices as described in section 5.1 and appendix F.

Table 13: Change in labor demand from a 25 kroner increase in the wage floor in the manufacturing industry.

Labor Type	$\overline{\Delta W_{hjt}}$	Change in Labor Demand		
		C-D	Subst.	Subst. + Outsrc.
Other Services	2.4%	-3.0%	-3.9%	-4.3%
Transportation	5.9%	-4.0%	-7.0%	-7.9%
ICT	0.8%	-1.1%	-1.2%	-1.2%
Marketing	3.2%	-1.7%	-3.2%	-3.6%
Cleaning	12.9%	-10.4%	-18.7%	-22.5%
Heavy Industry	3.4%	-3.3%	-5.9%	-7.6%
Manufacturing	3.2%	-3.6%	-4.5%	-4.8%
Total	3.2%	-3.1%	-3.8%	-4.2%

Note: Counterfactual aggregate labor demand if wage floor increased in 2011. $\overline{\Delta W_{hjt}}$ is the mean percent change in firm-task wage for each input type. Column 3 shows the results for the Cobb-Douglas benchmark ($\rho_h \to 0$). Column 4 shows the results when $\rho_h \neq 0$. Column 5 is when $\rho_h \neq 0$ and the probability of outsourcing is allowed to be nonzero as well.

Table 14: The Impact of Tariff Protection on Productivity

Approach	Corrections	Estimate (β_2)
WLP	Productivity	**-0.058**
		(0.012)
DL	Productivity & Price Variation	**-0.026**
		(0.005)
MC	Productivity & Substitution	**-0.103**
		(0.028)
DL-MC	Productivity, Prices, Substitution	**-0.069**
		(0.019)
DL-MC2	Productivity, Prices, Substitution	**-0.046**
		(0.007)

Note: The first four rows show the results of productivity regressed on lagged productivity and lagged tariffs, estimated using aggregate labor and intermediates. The fifth row is estimated with the two-step "M1-WLP" estimation strategy, using disaggregated data on input usage. As such, results from the fifth row is not directly comparable to the first four, which all use a similar estimation methodology.

Figure 1: **Number of occupations employed by firms in the supermarkets industry relative to their log revenue.**

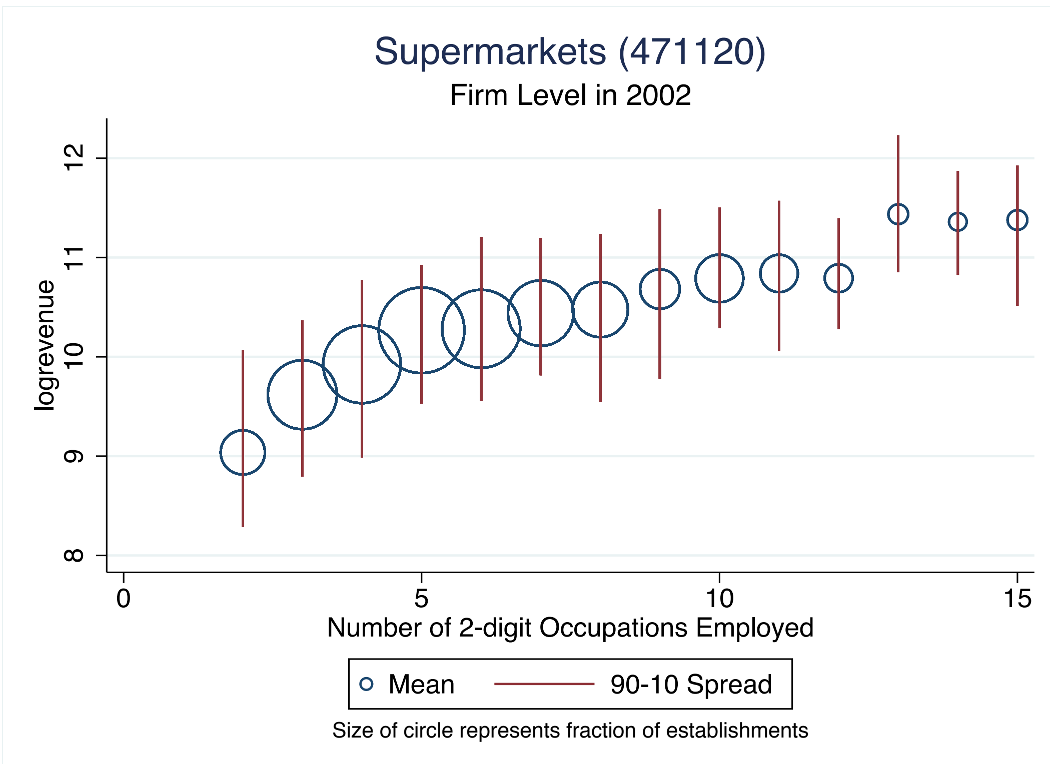

Figure 2: **Changes in occupational concentration over time**

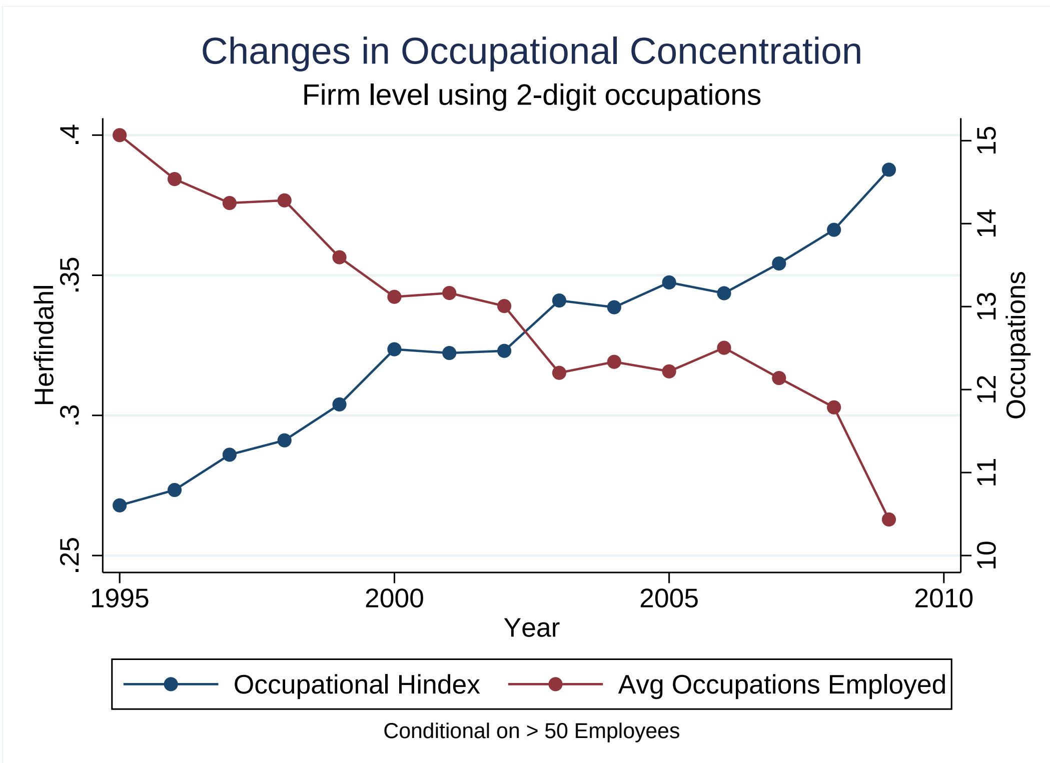

Figure 3: **Changes in number of occupations employed by a balanced panel of firms from 1995 to 2009.**

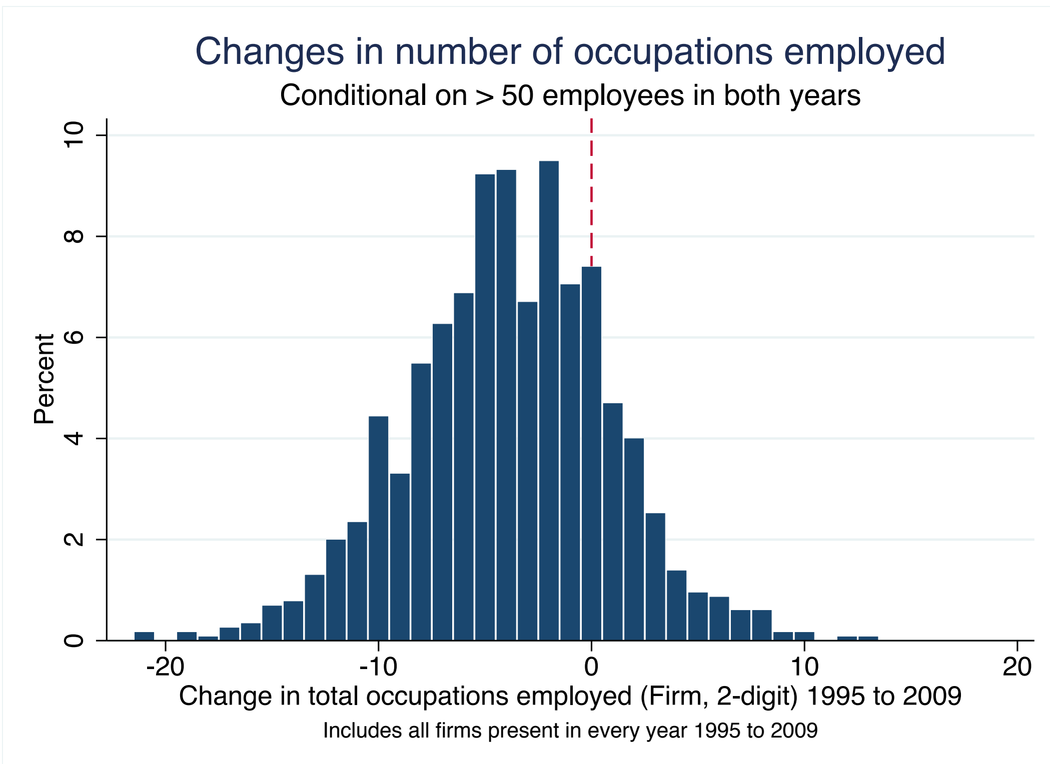

Figure 4: **Change in Materials to Total Labor Ratio over time**

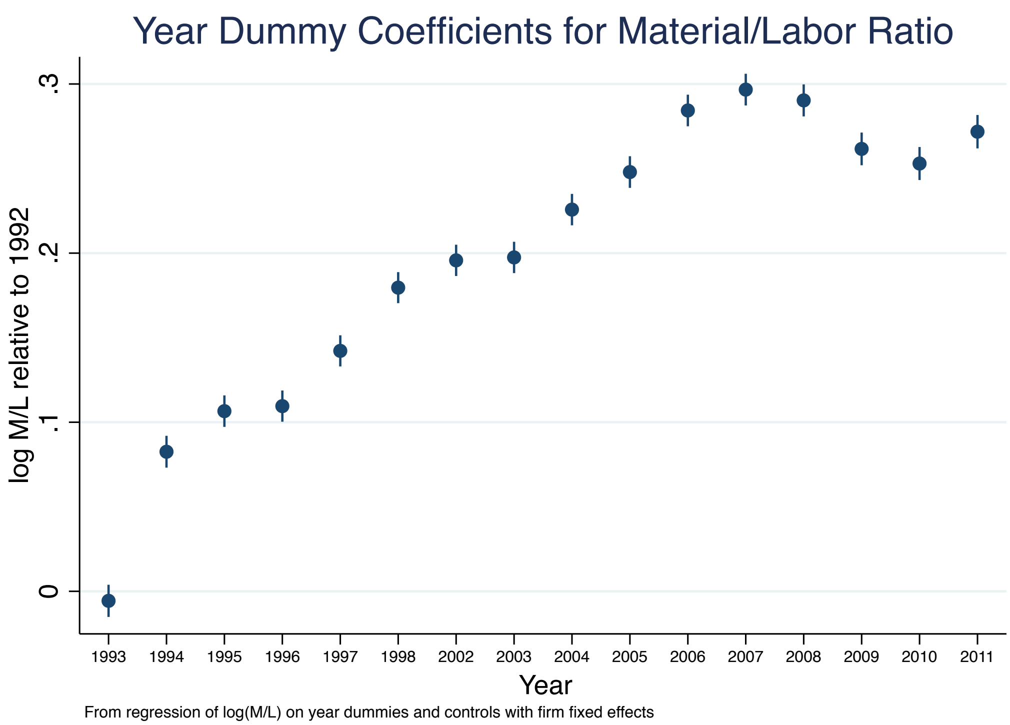

Figure 5: **Change in Primary Labor Share over time**

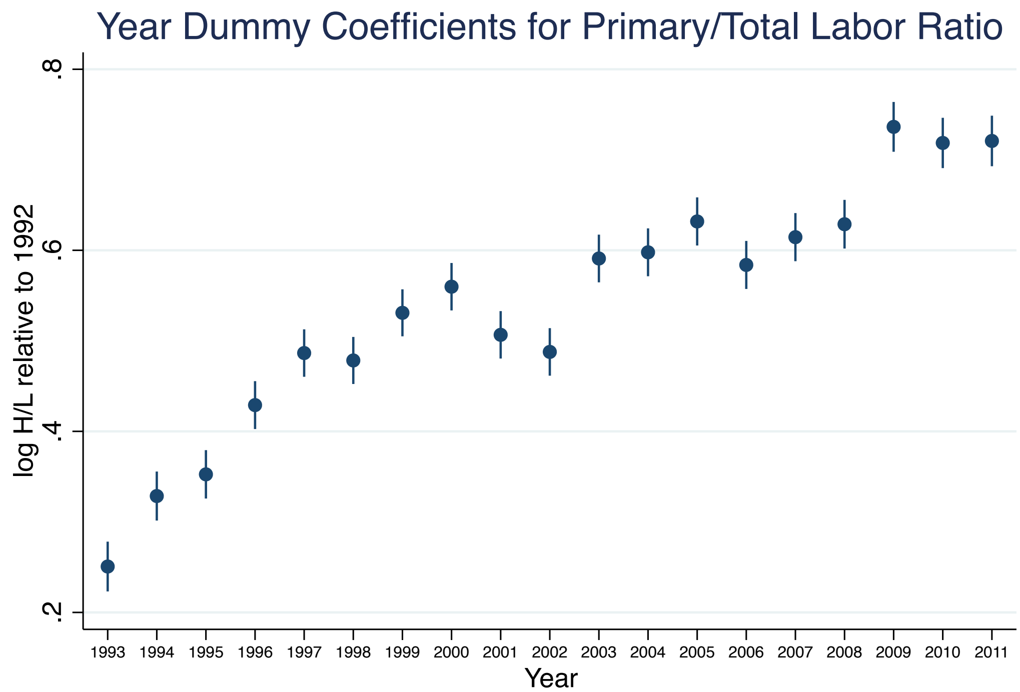

From regression of log(H/L) on year dummies and controls with firm fixed effects where H is industry matched labor and L is total labor.

Figure 6: **Illustration of the Make-Both-Buy decision**

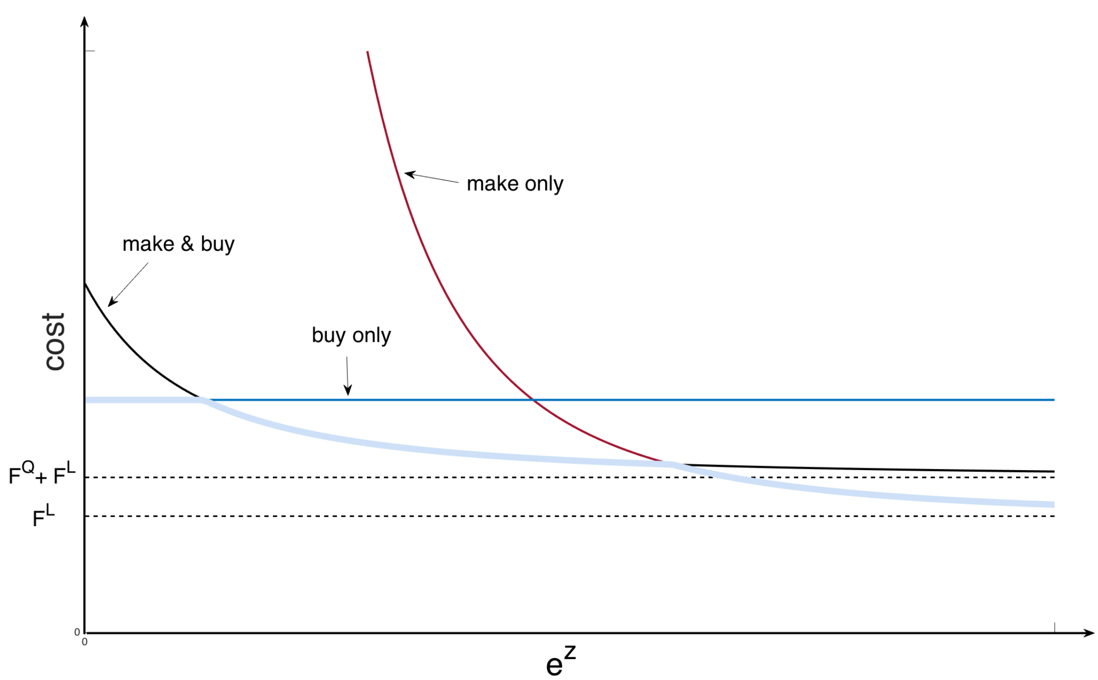

Figure 7: The distribution of real and counterfactual firm-task wages for transportation labor in the manufacturing industry in 2011.

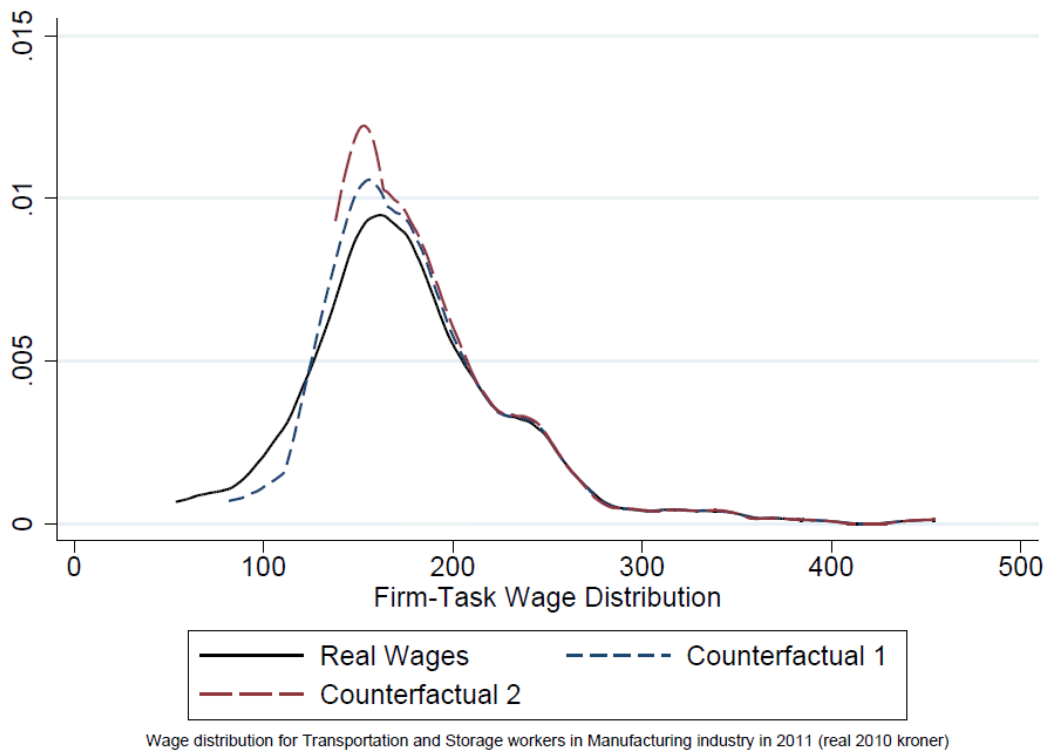

A Labor-Task Matching Algorithm and Results

In order to map the "task-matched" model of production to the data, I first define a set of input tasks which are used by firms in production. This is based on the set of intermediate input purchases which I observe in the data. Next, I need to determine which occupations a firm would employ in order to produce these tasks in-house. The obvious way to do this is to look at the firms which sell those intermediates and see which occupations they use in production. For example, I see that most firms spend money on transportation services. To figure out what occupations would be needed in order to produce transportation services in-house, I look at firms which sell transportation services and determine what occupations they use to do so. This is complicated by the fact that transportation firms themselves require non-transportation inputs and thus hire lawyers and janitors and web developers. To deal with this, I use a simple weighting algorithm to match occupations to the industry where they are most likely to act as *primary* labor – i.e. labor which is directly involved in producing the firm's primary output. Here I outline the steps of this matching algorithm.

A.1 Defining Input Tasks

The core of this paper is built around a unique set of data on purchased intermediate goods and services[36]. I take these highly disaggregated inputs and aggregate them up to a level which facilitates tractability while still representing very different and discrete tasks. This process is somewhat arbitrary, but intuitive. My data is split into two. Data on purchases of services, and data on purchases of physical intermediates at the HS6 level. I aggregate the purchased goods to the 2-digit HS2 level and manually define an aggregation mapping so that I end up with 13 tasks, including one "other" category. I then map NACE industry codes to these 13 task categories. The mapping of intermediates and industries into tasks is shown on table 15.

A.2 Mapping Occupations to Tasks

To match occupations to tasks I use data on the universe of firms and workers in the Danish matched employer-employee data. Let L_{oi} be the actual employment of occupation o

[36] The form which is used to collect this data, along with the full list of included services is available (in Danish) at http://www.dst.dk/pukora/epub/upload/17114/form.pdf

Table 15: Mapping of Industries and Intermediates into Tasks

Input Type	**Service/Product (HS) Codes**	**Industry (NACE) Codes**
Transportation	992*, 998006	494, 521, 5224, 532
ICT	993*	62, 63
Legal & Accounting	997001-997003	69, 70
Engineering	991001-991003	71
Marketing	994001	73
Employment & Training	996001	78
Security	995002	80
Cleaning	998005	812
Food	04,11,15-24	10, 11, 12
Wood & Paper	44,48	02, 16, 17
Heavy Industry	25, 27-35, 37-40, 54, 55, 68-79	06, 09, 19-24
Tools, Machinery, Goods	82-95	25-32

in industry i. Similarly L_i is total employment in industry i, L_o is total employment of occupation o, and $L = \sum_{o,i} L_{oi}$ is overall employment. Define $\widehat{L_{oi}} \equiv s_i L_o$ as the *predicted* employment of occupation o in industry i, where $s_i = L_i/L$ is industry i's share of total employment. I call this predicted employment because it is the level of employment of o in i you would expect if occupations were distributed across industries in equal proportion to the size of the industry.

The mapping of occupation to task is then based on determining which industry employs each occupation most disproportionately relative to its predicted employment. Formally, define the overall set of industries and occupations as I and O respectively. I assume there is a unique many-to-one mapping from the set of occupations to the set of tasks such that O can be partitioned into task subsets O_h. Under a similar assumption, let $I_h \subset I$ be the set of industries mapped to task h. The set of occupations o mapped to task h is then

$$O_h = \left\{ o \in O \,|\, \arg\max_{i \in I} \frac{L_{oi}}{\widehat{L_{oi}}} \in I_h \right\}$$

I perform this operation at the 4-digit occupation level, mapping to 2-digit NACE industries. I get similar results if I first aggregate industries and then perform the same mapping. I do not report the full mapping here (it will be available in an online appendix), but an example is the transportation services task. The set of 4-digit ISCO occupation codes mapped to this task are: Production and operations department managers in transport, storage and communications (1226), General managers in transport, storage and communications (1316), Transport clerks (4133), Drivers of motor vehicles (8320), Heavy truck and lorry drivers (8324), Crane, hoist and related plant operators (8333), Messengers, package and luggage porters and deliverers (9151) and Freight handlers (9333).

B Model with Task-Specific Capital

The basic theory of production written down in section 4 can easily be extended to accommodate task-specific capital. This may be a desirable extension, since it makes sense to think of input tasks being produced in-house jointly by labor and capital (say trucks and truck drivers), or outsourced to a trucking firm (in which case the firm may not need its own trucks). Suppose the physical production function is written as

$$Y_{jit} = \prod_{h \in \mathcal{H}_i} M_{hjit}^{\alpha_{hi}} e^{\omega_{jit}} e^{\varepsilon_{jit}} \tag{38}$$

where each input task M_{hjit} is a CES mix of intermediates Q_{hjit} purchased from industry h, and/or a task-specific capital-labor composite $\tilde{L}_{hjt} \equiv L_{hjt} K_{hjt}^{\beta_h} e^{z_{hjt}}$

$$M_{hjit} = \left(\gamma_{hi} \tilde{L}_{hjit}^{\rho_{hi}} + (1 - \gamma_{hi}) Q_{hjit}^{\rho_{hi}} \right)^{1/\rho_{hi}} \tag{39}$$

K_{hjt} is the task-specific capital which, combined with labor, produces in-house input task of type h. Solving this model is straightforward and very similar to the baseline model with aggregate capital. The difficulty comes in estimating the model when task-specific capital is not observed. One approach would be to assume that aggregate capital is a dynamic input, in that firms must decide the total value of the capital stock one period ahead of time, but that capital can be allocated across tasks flexibly. So, suppose $K_{hjt} \equiv \mu_{hjt} K_{jt}$, where $\sum_h \mu_{hjt} = 1$ and μ_{hjt} can be chosen every period without cost. Given this assumption, conditional on a firm choosing the BOTH technology, the marginal rate of substitution between labor and intermediates is

$$\mathrm{MRS}_h = \frac{\gamma_h}{1 - \gamma_h} \left((\mu_{hjt} K_{jt})^{\beta_h} e^{z_{hjt}} \right)^{\rho_h} L_{hjt}^{\rho_h - 1} Q_{hjt}^{1 - \rho_h}$$

and the optimal allocation of capital across tasks is defined by the following system of equations

$$\mu_{hjt}^* = \frac{\alpha_h \theta S_{hjt} \beta_h}{\sum_k \alpha_k \theta S_{kjt} \beta_k}$$

In principle, this allows the researcher to back out the optimal allocation of capital across tasks, as well as estimate the β_h parameters, even if capital is only observed in aggregate. However, this also implies that the production function is no longer weakly separable with respect to $\mathcal{R}_i$ and the ratio of labor and intermediates for each task h depends on the input mix for all other tasks $k \neq h$. Similarly, the choice of input technology for task h also then depends on the choice of technology for all other tasks k, as the relative cost of producing in-house depends on the available capital stock, which depends on whether or not the other tasks are outsourced. Since the problem can no longer be split by task, all parameters in the model would have to be estimated jointly, which quickly becomes intractable.

C Separability

This section establishes the weak separability result more formally. I start with the assumption on production:

Assumption 1. *The physical production function for firm j in industry i in period t is*

$$Y_{jit} = K_{jt}^{\beta} \prod_{h \in \mathcal{H}_i} M_{hjt}^{\alpha_h} e^{\omega_{jt}} e^{\varepsilon_{jt}} \tag{40}$$

where each input task M_{hjt} is a CES mix of intermediates Q_{hjt} purchased from industry h, and/or task-specific labor L_{hjt}

$$M_{hjt} = \left(\gamma_{hi}(e^{z_{hjt}} L_{hjt})^{\rho_h} + (1-\gamma_h) Q_{hjt}^{\rho_h}\right)^{1/\rho_h} \tag{41}$$

I also make the following assumptions about the capital, labor and intermediate inputs:

Assumption 2. *Capital (K_{jit}) is predetermined, while all labor (L_{hjit}) and intermediate (Q_{hjit}) inputs are flexible.*

I restate Nadiri (1982) in making the following definition:

Definition C.1. Let $\mathcal{X}_i = \{L_h, Q_h \,|\, h \in \mathcal{H}_i\}$ be the set of labor and intermediate inputs required by a particular industry, and let production function $F(\mathcal{X}_i)$ be twice differentiable and strictly quasi-concave. Define $\mathcal{R}_i$ as the partition of $\mathcal{X}_i$ into mutually exclusive subsets $\mathcal{X}_{hi} = \{L_h, Q_h\} \,\forall\, h \in \mathcal{H}_i$. $F(\mathcal{X}_i)$ is *weakly separable* with respect to $\mathcal{R}_i$ if the marginal rate of substitution between L_h and Q_h is independent of $\mathcal{X}_{ki} \,\forall\, k \neq h$.

Proposition C.1. *Given Assumption 2, the production function stated in Assumption 1 is weakly separable with respect to partition $\mathcal{R}_i$.*

Proof. By inspection of the first order conditions of the firm's cost minimization problem (see equations 16 and 17). $\square$

D The Productivity Bias from Ignoring Substitution

One key difference between the task-matched CES and a more common Cobb-Douglas approach is that the CES allows for flexible and heterogeneous substitution patterns across inputs, while the Cobb-Douglas does not. This is obviously of interest when the researcher explicitly desires to measure demand or substitution elasticities. However, it should also be of interest to anyone interested in estimating productivity using a production function. Klette and Griliches (1996) and De Loecker (2011) show that failing to account for price variation when estimating productivity with deflated revenues leads to unobserved prices "polluting" estimates of productivity. I briefly provide a similar argument here for why input substitution may pollute productivity estimates. See section 8 for a demonstration of the idea via estimating the effects of tariff protection on different estimates of productivity.

Suppose the true production function and dgp is as described in the preceding sections, but that the researcher mistakenly believes that firms use a disaggregated Cobb-Douglas technology in the same set of inputs. To what degree with they get the wrong answer when estimating productivity? To see this, note that the task matched CES production function in equation 4 can be rewritten in Cobb-Douglas form:

$$Y_j = e^{\tilde{\omega}_j} K_j^{\beta} \prod_{h \in \mathcal{H}} L_{hj}^{\alpha_{\ell h}} Q_{hj}^{\alpha_h} \tag{42}$$

where the modified productivity term $\tilde{\omega_j}$ is

$$\tilde{\omega}_j = \omega_j + \sum_{h \in \mathcal{H}} \log \mathcal{G}_{hj} \tag{43}$$

and

$$\mathcal{G}_{hj} \equiv ((1 - \gamma_h)(1 - S_{hj})^{-1})^{\frac{\alpha_h}{\rho_h}} L_{hj}^{-\alpha_{\ell h}} \tag{44}$$

It's immediately clear that any estimates of productivity using the Cobb-Douglas formulation will contain not only true productivity ω_j but also the input variation term $\mathcal{G}_{hj}$. This is a problem, since $\mathcal{G}_{hj}$ is correlated with ω_j through L_{hj}, and also potentially correlated with variables which the researcher may want to regress on tfp, such as changes in tariffs or competitive pressure. In particular, input shares S_{hj} will shift in response to unobserved changes in input prices, with the degree of pollution depending on the magnitude of ρ_h. Input shares may also be correlated with unobserved productivity through wages. Since the input share and labor terms are correlated in different directions, overall direction of bias is a question of relative magnitude. Estimating the model with the task-matched framework allows for flexible estimation of firm productivity while controlling for and skimming out the input variation from the productivity term. Since it cleanly nests the Cobb-Douglas, this allows researchers to test the degree to which estimates of productivity gained using the Cobb-Douglas framework may be tainted by unmeasured variation in inputs.

E Assumptions on Productivity, Timing and Prices

This appendix provides the formal definitions, assumptions and technical details which I employ in estimating the model.

Since identification of the model requires several timing assumptions, it is convenient to define the information set of the firm in period t as $\mathcal{I}_t$. This information set contains all of the information with which the firm enters period-t and thus uses to make period-t choices such as input or outsourcing decisions. $\mathcal{I}_t$ contains information relevant to the firm such as (lagged) prices, inputs and productivity. Let X_t denote a generic input in period t. Following Gandhi, Navarro and Rivers (2020), I define any input such that $X_t \in \mathcal{I}_t$ as *predetermined*, implying that $X_t(\mathcal{I}_{t-1})$ is a function of the previous period's information set. Capital is commonly treated as a predetermined input. Define any input which is not predetermined as *variable*. Additionally, an input which is variable and where the optimal choice X_t^* is a function of lagged values of itself is defined as being *dynamic*, whereas an input which is variable but not dynamic is *flexible*. An input with adjustment costs might be dynamic. Labor is frequently treated as dynamic, while materials are typically treated as flexible.

My assumptions on total factor productivity are standard in the literature. Nevertheless it is useful to state these assumptions formally.

Assumption 3. *The hicks neutral productivity term $\omega_{jit} \in \mathcal{I}_t$ is observed by the firm prior to making period-t decisions and is Markovian, such that $\mathbb{E}[\omega_{jit}|\mathcal{I}_{t-1}] = \mathbb{E}[\omega_{jit}|\omega_{jit-1}] =$*

$g_\omega(\omega_{jit-1})$ *for some continuous function* $g(.)$*. Also,* $\varepsilon_{jit} \notin \mathcal{I}_t$ *and is i.i.d. across firms and time.*

I also make the following assumptions about the capital, labor and intermediate inputs:

Assumption 4. *Capital* (K_{jit}) *is predetermined, while all labor* (L_{hjit}) *and intermediate* (Q_{hjit}) *inputs are flexible.*

Since the production technology is weakly separable (see appendix C), I can make the following timing assumption. Define the information sets $\mathcal{I}_t^{'}$ and $\mathcal{I}_t^{''}$ such that $\mathcal{I}_t \subset \mathcal{I}_t^{'} \subset \mathcal{I}_t^{''} \subset \mathcal{I}_{t+1}$.

Assumption 5. (Timing) *Upon entering period* t*, firms choose their optimal scale conditional on* $\mathcal{I}_t$*. This provides a vector of optimal input requirement terms* $\{M^*_{hjt}(\mathcal{I}_t)\}_{h\in\mathcal{H}_t}$*. Firms then observe* $\mathcal{I}_t^{'}$ *(where* $M^*_{hjt} \in \mathcal{I}_t^{'}$*) and choose whether to fulfill their input requirement by hiring labor to produce it in-house (make), purchasing it on the market (buy), or doing both. Given their make and/or buy decision, firms finally observe* $\mathcal{I}_t^{''}$ *and choose levels of* L_{hjt} *and* Q_{hjt}*.*

While I call these "timing" assumptions, one could also think of this multi-stage decision process as reflecting decisions made at different levels of the firm. Perhaps firm scale is determined by top management, with total input requirements passed on to subdivisions of the firm. These subdivisions, which are responsible for providing the input tasks, then have the autonomy to decide how these tasks are provided, be it via in-house labor and/or outsourcing. The differences in information sets then reflect not sequential realizations over time, but differences in information sets across organizational layers in the firm.

The identification strategy used in this paper relies on the optimizing behavior of the firm, especially in response to changes in the costs of inputs. As such, I need to be explicit about the assumptions on factor prices. This assumption also makes explicit how information differs over time or division.

Assumption 6. (Prices) *The firm-task specific marginal cost of labor* (L_h) *is* W_{hjt}*, a function of common (industry) market wage component* ($W_{ht} \in \mathcal{I}_t$)*, firm productivity* ($z_{hjt} \in \mathcal{I}_t^{'}, \omega_{jt} \in \mathcal{I}_t$) *and a firm-task component* $\Theta_{hjt} \in \mathcal{I}_t^{''}$*. Firms in industry* i *face a common market price* $P_{ht} \in \mathcal{I}_t^{'}$ *for intermediate* Q_h*.*

This implies first that firms face price uncertainty when making their scale and make-or-buy decisions. Second, Θ_{hjt} may contain compensating differentials or differences in labor market tightness across locations, implying that firms face imperfect labor markets. Third, since wages are allowed to depend on firm productivity and other unobserved firm-task components of Θ_{hjt}, firms may possess some measure of market power in the setting of wages. While I do not model the evolution of wages or prices directly in this paper, I allow and control for these different components in my empirical strategy. In particular, the latter implication, which is supported by a significant literature on wage setting, is an important part of my strategy for identifying the elasticity of substitution between labor and intermediates.

F Estimation Details

F.1 Scale Parameters

My strategy for recovering $\alpha_h\theta$ closely follows Gandhi, Navarro and Rivers (2020). Recall that firms choose optimal input levels M^*_{hjt} under price uncertainty. I assume this uncertainty takes the following form,

Assumption 7. (Price Uncertainty) *Let $P^I_{hjt} \in \mathcal{I}''_t$ be the ex-post marginal cost of aggregate input M_{hjt}. While $P^I_{hjt} \notin \mathcal{I}_t$, firms do observe a noisy signal $\tilde{P}^I_{hjt} \in \mathcal{I}_t$ of marginal costs, where $\tilde{P}^I_{hjt} = P^I_{hjt}e^{-\eta_{hjt}}$, η_{hjt} is i.i.d., and $\mathbb{E}[\eta_{hjt}] = 0$.*

Given Assumption 7, we have $\mathbb{E}[P^I_{hjt}] = \tilde{P}^I_{hjt}\bar{\eta}_h$ where $\bar{\eta}_h \equiv \mathbb{E}[e^{\eta_{hjt}}]$ and thus can rewrite equation 10 as $X_{hjt} = \alpha_h\theta R_{jt} + \alpha_h\theta R_{jt}(\bar{\eta}_h^{-1}e^{\eta_{hjt}} - 1)$. Dividing through by R_{jt} and taking expectations of both sides provides $\alpha_h\theta = \mathbb{E}\left[X_{hjt}/R_{jt}\right]$. I take the empirical analog of the expectation so that my estimate of $\alpha_h\theta$ is

$$\widehat{\alpha_h\theta} = \frac{1}{J_h}\sum_{\mathcal{J}_h}\left[\frac{X_{hjt}}{R_{jt}}\right] \tag{45}$$

Note that in principle, this approach allows for the scale parameters on the flexible inputs to vary over time, though I hold them fixed for my application.

F.2 Demand Estimation

As mentioned in section 4.2, I can estimate the the demand parameters in several ways. In section 8.2 I follow De Loecker (2011) and Halpern, Koren and Szeidl (2015) in using a

CES specification where demand is proxied by revenue shares. This is so that I can directly compare my estimates of the effects of tariffs on productivity to the existing literature. See 8.2 for details. Alternately, one can estimate the demand parameters directly off of firm-product output data. This is the approach I take in order to get estimates of the demand elasticity θ and the demand shifter ψ_{jt}. The empirical specification is a simple logit in log prices and a firm effect. I let the indirect utility for individual ℓ for purchasing product i from firm j in period t be equal to $V_{\ell jit} = D_j + \eta_i^d p_{\ell jit} + \xi_{jit} + \epsilon_{\ell jit}$. This gives the standard estimating equation

$$ms_{jit} - ms_{oit} = D_j + \eta_i^d p_{jit} + \xi_{jit} \tag{46}$$

with $ms_{jit} = \log(MS_{jit})$ the log market share in industry i for firm j, ms_{oit} represents the log outside share, p_{jit} is log price, and D_j is a firm-specific fixed effect. To control for the endogeneity of prices, I instrument with a basic Hausman-style strategy, where I use the average price of all other goods within the firm's narrow industry as a proxy for a demand shock[37]. The identifying assumption is that, controlling for firm/product-specific means, the industry demand shocks are uncorrelated with unobserved variation in product quality. This then provides estimates of $\hat{\eta}_i^d$ from which I construct $\hat{\theta}_i$, and $\hat{\psi}_{jt} = MS_{oit}Q_{it}e^{D_j+\xi_{jit}}$.

F.3 Expected Wages

The timing assumptions in the body of the paper imply that firms base their make-both-buy decision on expected wages, conditional on their information set $\mathcal{I}_t'$ at the time of making the decision. The expected wage can be written

$$\mathbb{E}[W_{hjt} \,|\, \mathcal{I}_t'] = \mathbb{E}[W_{hjt} \,|\, W_{hit}, z_{hjt}, \omega_{jt}, \mathbb{E}[\Theta_{hjt}]] = g_w(W_{hit}, z_{hjt}, \omega_{jt}, \mathbb{E}[\Theta_{hjt}]) \tag{47}$$

for some unknown function g_w. I approximate g_w with $\hat{g}_w(h, i, t, W_{hjt-1}, R_{jt-1}, j)$, where lagged wages, revenues and firm fixed effects proxy for unobserved productivity and labor market heterogeneity, and industry-task-year effects proxy for the average industry-task-year wage component. In particular, I run the following regression for each industry-task pair:

$$W_{hjt} = \hat{h}_{hj}(W_{hjt-1}, R_{jt-1}) + b_t + b_{hj} + \varepsilon_{hjt}^w \tag{48}$$

[37] See Hausman (1996) and Nevo (2001) for further discussion.

where $\hat{h}_{hj}$ is an industry-specific polynomial in lagged wages and revenues[38], b_{ht} is a task-time effect capturing average market wages for labor type h in year t, and b_{hj} is a firm fixed effect for task h, capturing persistent heterogeneity in compensating differentials or labor market conditions. I assume that this specification matches how the firm itself calculates expected wages and use the predicted values from specification 48 in the estimation of the structural model.

G Calculation of the outsourcing probabilities

There are a few simple results, which translate directly to the other cutoffs, that are needed to calculate the probability of outsourcing. First I formally restate several assumptions from the main paper, then the results.

Assumption 8. (Task Productivity) *The task specific labor-enhancing productivity term* $z_{hjt} \in \mathcal{I}_t^{'}$ *is Markovian. More specifically it follows an AR(1) process:* $z_{hjt} = z_h + \delta_h z_{hjt-1} + \zeta_{hjt}$ *where the innovation term is i.i.d. and* $\zeta_{hjt} \sim N(0, \sigma_h) \notin \mathcal{I}_t$.

Assumption 9. (Distribution of Fixed Costs) *Fixed costs* f_{hjt}^L *and* f_{hjt}^Q *follow an i.i.d. log-normal distribution, such that* $\log(f_{hjt}^A) \sim N(\bar{f}^A, \sigma_h^A)$ *for* $A \in \{L, Q\}$.

Lemma G.1. *Cutoff* $C_{hjt}^1(f_{hjt}^L, \mathbb{E}[W_{hjt} \,|\, \mathcal{I}_t^{'}], P_{ht})$ *is a continuous function which is monotone increasing in* $\mathbb{E}[W_{hjt} \,|\, \mathcal{I}_t^{'}]$ *and* f_{hjt}^L*, and monotone decreasing in* P_{ht}.

Proof. By the definition of C_{hjt}^1 and assumption 5. □

Lemma G.2. *Given any particular expected wage* $\mathcal{W}$ *and price* $\mathcal{P}$*, then for any realization of firm-task productivity* ζ_{hjt}*,* $\exists! \, \tilde{f}_{hjt}^L$ *s.t.* $\zeta_{hjt} = C_{hjt}^1(\tilde{f}_{hjt}^L, \mathcal{W}, \mathcal{P})$*. In addition,* C_{hjt}^1 *is invertible in* f_{hjt}^L *such that we can write* $\tilde{f}_{hjt}^L = g_c^1(\mathcal{W}, \mathcal{P}, \zeta_{hjt})$.

Proof. Follows from assumptions 8, 9 and lemma G.1 □

Lemma G.3. *Define* $\hat{f}_{hjt}^L$ *and* $\tilde{f}_{hjt}^L$ *s.t.* $\zeta_{hjt} = C_{hjt}^1(\hat{f}_{hjt}^L, \widehat{\mathcal{W}}, \mathcal{P})$ *and* $\zeta_{hjt} = C_{hjt}^1(\tilde{f}_{hjt}^L, \mathcal{W}, \mathcal{P})$*. Then* $\widehat{\mathcal{W}} > \mathcal{W} \implies \hat{f}_{hjt}^L < \tilde{f}_{hjt}^L$.

Proof. Follows from lemmas G.2 and G.1 □

[38] In practice I experiment with first, second and third degree polynomials. The results are nearly identical across all specifications, so I use the results from the linear approximation.

The intuition behind these lemmas is that as the cost of labor increases, the productivity cutoff at which the firm is indifferent between outsourcing and choosing Both increases, since only more productive firms will still find it profitable to employ labor. While we do not observe the firm's fixed costs, we can also characterize a fixed cost cutoff. Thus as wages increase, the fixed cost at which a firm with a given productivity would remain indifferent between choosing Buy and Both decreases. This allows me to characterize the counterfactual choice probability as follows. Let counterfactual wages be denoted with the hat symbol, and define the counterfactual and realized fixed cost cutoffs as $\hat{f}^L_{hjt} = g^1_c(\mathbb{E}[\widehat{W}_{hjt} \,|\, \mathcal{I}'_t], P_{ht}, \zeta_{hjt})$ and $\tilde{f}^L_{hjt} = g^1_c(\mathbb{E}[W_{hjt} \,|\, \mathcal{I}'_t], P_{ht}, \zeta_{hjt})$. The goal is to get the probability that a firm outsources in period t under the counterfactual wage given that they chose ($\mathcal{D}_{hjt} =$ Both) when faced with realized wages. This probability can be expressed as

$$\Pr\left[\zeta_{hjt} < C^1_{hjt}(f^L_{hjt}, \mathbb{E}[\widehat{W}_{hjt} \,|\, \mathcal{I}'_t]) \,|\, \zeta_{hjt} > C^1_{hjt}(f^L_{hjt}, \mathbb{E}[W_{hjt} \,|\, \mathcal{I}'_t])\right] = \frac{F(\tilde{f}^L_{hjt}) - F(\hat{f}^L_{hjt})}{F(\tilde{f}^L_{hjt})} \tag{49}$$

Where the equality follows from lemmas G.1 to G.3 and F represents the distribution of f^L_{hjt}. Given the estimated model, the firm's productivity innovation ζ_{hjt} is known but fixed costs remain unobserved. However, since I have estimates of the distribution of fixed costs, I can calculate the two fixed cost cutoff terms and then use equation 49 to calculate the probability of outsourcing from a given change in wages. In essence this calculation says: given the increase in wages, the fixed-cost cutoff must have decreased. What is the probability that the firm's unobserved fixed-cost draw was above the counterfactual cutoff $\hat{f}^L_{hjt}$ given that we know it must have been below the realized cutoff $\tilde{f}^L_{hjt}$.

The procedure is as follows. First, calculate $\tilde{f}^L_{hjt}$ by inverting the cutoff equations which provide

$$\tilde{f}^L_{hjt} = \left(1 - \left[\left(\frac{e^{z_{hjt}} P_{ht}}{\mathbb{E}[W_{hjt} \,|\, \mathcal{I}'_t]}\left(\frac{\gamma_h}{1-\gamma_h}\right)^{\frac{1}{\rho_h}}\right)^{\frac{\rho_h}{1-\rho_h}} + 1\right]^{\frac{\rho_h - 1}{\rho_h}}\right)(1 - S_{hjt})^{\frac{\rho_h - 1}{\rho_h}} X_{hjt} \tag{50}$$

where using the model estimates and first order conditions we can obtain

$$e^{z_{hjt}} P_{ht}\left(\frac{\gamma_h}{1-\gamma_h}\right)^{\frac{1}{\hat{\rho}_h}} = \left(\frac{L_{hjt}}{X^Q_{hjt}}\right)^{\frac{1-\hat{\rho}_h}{\hat{\rho}_h}} W_{hjt}^{\frac{1}{\hat{\rho}_h}} \tag{51}$$

Calculating the counterfactual fixed cost cutoff $\tilde{f}^L_{hjt}$ involves calculating the counterfactual optimal expenditure level X_{hjt} and share S_{hjt} (both of which are endogenous functions of the wage), which can then be plugged into equation (50) along with the counterfactual wages. In particular,

$$\widehat{X}_{hjt} = X_{hjt}(1 + \%\Delta W_{hjt} * \epsilon^{X_{hjt}}_{W_{hjt}})$$
$$\widehat{S}_{hjt} = S_{hjt}(1 + \%\Delta W_{hjt} * \epsilon^{S_{hjt}}_{W_{hjt}})$$

This then provides estimates of the probability of outsourcing, Pr(Outsource). The expected change in labor from a given change in own-type wages can then be calculated as:

$$\mathbb{E}[\Delta L_{hjt}] = \Pr(\text{Outsource})(-L_{hjt}) + (1 - \Pr(\text{Outsource}))(\%\Delta W_{hjt} \times \epsilon^{L_{hjt}}_{W_{hjt}}) \tag{52}$$

Note that while the above derivations have been in terms of change in demand for L_{hjt} from a change in W_{hjt}, as mentioned above the demand depends on the entire vector of wages and prices. This changes the calculation of the counterfactual $\widehat{X}_{hjt}$ and $\widehat{S}_{hjt}$ since changes in other prices affect demand for h via firm scale. When doing the full-industry counterfactual, I take these total changes into account. Similarly, the intensive margin change in labor demand is also calculated to take into account the optimal response to the wage changes for all of the labor types employed by the firm.