The Rise of Specialized Firms*

Lorenz K.F. Ekerdt
U.S. Census Bureau

Kai-Jie Wu
Penn State

October, 2023

[click for latest version]

Abstract

This paper studies firm diversification over 6-digit NAICS industries in U.S. manufacturing. We find that firms specializing in fewer industries now account for a substantially greater share of production than 40 years ago. This reallocation is a key driver of rising industry concentration. Specialized firms have displaced diversified firms among industry leaders—absent this reallocation concentration would have decreased. We then provide evidence that specialized firms produce higher-quality goods: specialized firms tend to charge higher unit prices and are more insulated against Chinese import competition. Based on our empirical findings, we propose a theory in which growth shifts demand toward specialized, high-quality firms, which eventually increases concentration. We conclude that one should expect rising industry concentration in a growing economy.

JEL codes: L25, O33, F14, L11, O47

*We thank our adviser Mark Bils for his invaluable guidance and good humor throughout this project. Special thanks to Nichole Szembrot at the Cornell FSRDC for her exceedingly helpful assistance with the Census data. We have also benefited greatly from discussions with George Alessandria, David Argente, Yan Bai, Lisa Kahn, Peter Klenow, Narayana Kocherlakota, Ronni Pavan, Xincheng Qiu, Nese Yildiz, and numerous seminar participants at the University of Rochester and various conferences. Any views expressed are those of the authors and not those of the U.S. Census Bureau. The Census Bureau has ensured appropriate access and use of confidential data and has reviewed these results for disclosure avoidance protection (project 2464). We acknowledge financial support from the National Science Foundation Award 2018048.

1 Introduction

Industry concentration—the market share of industry leaders—is traditionally used as an indicator of market competition. Thanks to the increased availability of administrative data covering the universe of U.S. firms and industries, recent work has documented a long-running upward trend in industry concentration, leading to a surge in interest in this development.[1] The proposed explanations vary considerably, ranging from changes in entry barriers, business dynamism, creative destruction, and globalization.[2]

We contribute to this debate by providing a new view on rising concentration. First, we show that rising concentration is not due to the expansion of gigantic firms operating in numerous industries. Instead, the Googles and Amazons of the world have become increasingly aberrant, and rising concentration is instead due to the faster growth of more specialized firms—firms that focus their production in fewer industries. Behind this lies a secular trend reallocating market share toward more specialized firms. This trend is particularly pronounced among industry leaders: our results suggest that industry concentration would have decreased absent the faster growth of specialized firms.

Second, we use proxies for product quality to show that specialized firms tend to produce higher-quality goods. Combined with our first finding, we therefore see rising industry concentration as stemming from a shift in the economy favoring firms who are particularly adept at producing higher-quality goods.

How do our empirical findings help us reinterpret rising industry concentration? We propose a theory in which quality is a luxury: as consumers become richer, they demand higher quality goods. Growth thus shifts demand toward firms producing higher-quality goods. Through the lens of our theory, we argue that one should expect concentration to increase in a growing economy, and so other inferences on this measure should be made relative to this trend.

Our measure of firm-level diversification relies on information on the universe of establishments from the U.S. Census Bureau. Comparing firm diversification over a long time period has only recently become possible due to the development of a harmonized NAICS code by Fort and Klimek (2018). The data enable us to observe firms' sales and employment over narrowly-

[1]See Autor, Dorn, Katz, Patterson, and Van Reenen (2017, 2020) for the trend between 1982-2012.

[2]See, for example, Furman and Orszag (2018), Gutierrez and Philippon (2017), Barkai (2020), Autor et al. (2017, 2020), Hsieh and Rossi-Hansberg (2019) Aghion, Bergeaud, Boppart, Klenow, and Li (2023), Akcigit and Ates (2021), Olmstead-Rumsey (2022), and Sui (2022), among others.

defined time-consistent industries, where an industry refers to a 6-digit code in the North American Industry Classification System (NAICS). We use a diversity index called effective number of industries (ENI) to measure firms' diversification based on how their sales and employment are distributed over industries. Our main results focus on the U.S. manufacturing sector, but we show that similar results hold in other sectors.

Our findings indicate a long-running expansion in the relative size of specialized firms over the past 40 years. While single-industry firms account for about 95% of U.S. manufacturing firms in both 1977 and 2017, their share of total employment has grown from 33% to 48%. The same pattern is found among multi-industry firms: the (unweighted) distribution of diversification has remained virtually unchanged, but more specialized firms now account for a substantially greater share of production. As a result, aggregate diversification, as measured by the size-weighted average effective number of industries, has fallen to about half of what it had been 40 years ago. We demonstrate that this reallocation toward specialized firms is not due to net entry; instead, it is fully driven by the reallocation between continuing firms, which in turn is largely driven by the faster growth of more specialized firms.

To show that said reallocation toward specialized firms is a key driver of rising industry concentration, we first show that the firm-level reallocation we document occurs within most industries and is not driven by a disproportionate expansion of industries with more specialized firms. We then measure specialized firms' contribution to industry concentration relative to that of diversified firms. We find that more specialized firms have dramatically increased their presence as industry leaders, defined as the top 4 largest firms in the industry, and consequently account for a greater market share among leaders. Using a simple counterfactual where we shut off the contribution of specialized firms' growth to industry concentration, we show that the expansion of specialized firms can account for more than the entire observed increase in manufacturing industry concentration.

We make two additional remarks to clarify our main findings. First, we make use of the interplant transfer variable in the Economic Census to separate horizontal and vertical diversification, where the latter refers to firms' expansion into industries in order to supply their own production in other industries with inputs. We find that horizontal diversification accounts for most of the variation in the observed ENI. Secondly, we demonstrate that the reallocation toward specialized firms is robust to including manufacturing firms' production in other sectors, and this trend is not unique to manufacturing, though it appears strongest there.

We argue that the secular trend of reallocation toward specialized firms reflects a demand shift favoring higher-quality products, where we define quality as a product-embedded characteristic that increases demand for the firm's product. Before formalizing this idea, we provide evidence favoring a link between specialization and product quality.

Because quality is hard to observe directly, we look at unit prices as a starting point. This approach is motivated by Rosen's (1974) theory of hedonic prices, and has been heavily applied in the trade literature. We find that specialized firms tend to charge higher unit prices within narrowly-defined product categories, suggesting that specialized firms produce higher-quality goods. To address alternative interpretations of this price differential we study the evolution of specialized firms' relative factor shares. The results do not support that the higher unit prices of specialized firms stem from distortions in input or output markets.

Our second approach is to assess firms' product quality by their responses to the dramatic expansion of Chinese imports beginning in the 1990s. This is motivated by a wealth of evidence that developing countries export relatively low-quality products to developed countries, and thus plausibly compete more directly with relatively low-quality domestic firms.[3]

We find that diversified firms' employment shrinks significantly more in response to the same change in exposure to Chinese import competition than the employment of specialized firms. The magnitude of the differential response is significant: our preferred specification suggests that Chinese import competition alone can account for around 40% of the difference in employment growth between specialized and diversified firms between 1997 and 2012. This differential response does not reflect differential responses by firm size, changes in the opportunity to export, or differential responses by firms' labor intensity. We thus view this result as favoring the view that specialized firms produce high-quality products, and consequently are less impacted by the increasing availability of cheaper and lower-quality Chinese imports.

To demonstrate that our findings can speak to the change in industry concentration, we build a general equilibrium model of heterogeneous firms featuring two key ingredients. First, we assume that quality is a luxury, and thus relative demand increasingly favors firms that produce higher-quality products as income grows (Bils and Klenow, 2001; Fajgelbaum, Grossman, and Helpman, 2011). We discipline this income effect using the empirical quality Engel curve, which captures the relationship between consumers' total expenditure and the unit prices of consumed products.

Secondly, we introduce a mechanism connecting a firm's level of diversification to the qual-

[3]See Hummels and Klenow (2005) and Hallak and Schott (2011) for developing countries in general, Schott (2008) for the specific case of Chinese exports.

ity of its products. Firms have access to two types of technologies with which to increase their profits: one reducing their marginal cost of production and the other increasing the quality of their products. The key distinction between these two types of technologies is that they differ in terms of their (dis)economies of scope, i.e. in how the average productivity of technology improvement changes with a firm's number of industries. We discipline their economies of scope using our estimated relationship between unit prices and firm diversification in conjunction with the employment-weighted distribution of firm diversification. Our calibration implies diseconomies of scope for quality improvement and hence specialized firms tend to produce higher-quality goods.

As the economy becomes richer and increasingly favors higher-quality products, specialized firms thus grow disproportionately. However, the magnitude of this reallocation is unclear, and its effect on concentration is generally ambiguous. We hence use our calibrated model to show that growth can replicate an economically significant reallocation toward specialized firms and a rising trend in industry concentration. To do so, we let aggregate technology improve so that the real consumption in the model replicates the path of U.S. real consumption per capita over the past 40 years.

This exercise produces two key insights. First, our proposed mechanism is quantitatively relevant: growth alone can generate a reallocation toward specialized firms half as large as that observed in the data. Secondly, and more importantly, we show that our calibrated model generates an increasing trend in industry concentration without targeting any trends other than consumption. We conclude that one should expect industry concentration to increase in a growing economy, and so other changes in concentration should be viewed relative to this trend.

The paper proceeds as follows. Section 2 describes the data used and defines our measure of diversification. Section 3 investigates the firm-level patterns underlying the reallocation of production toward specialized firms. Section 4 establishes a connection between specialization and product quality by examining cross-sectional variation in unit prices and responses to Chinese import competition. Section 5 introduces a theory of heterogeneous firms in which specialized firms tend to produce higher-quality products. Finally, section 6 concludes.

Related literature. As a side note to rising industry concentration, Autor et al. (2020) and Hsieh and Rossi-Hansberg (2019) remark that industry leaders operate in fewer industries than they used to, which they attribute to the largest firms increasingly focusing on their core competencies.[4] In contrast, our paper emphasizes an alternative channel: the reallocation of market share from diversified to specialized firms, which is mainly driven by the faster growth of more specialized firms. In Section 3, we demonstrate that our new channel contributes to most of the overall increase in specialization. Moreover, we show that the reallocation toward specialized firms is the main driver of rising industry concentration.

Ma (2022) documents the increasing (unweighted) average specialization of innovative firms and connects this to pro-patent reform.[5] We do not view this theory as being in competition with ours because we focus on a reallocation toward more specialized firms rather than the increasing specialization of large firms. Moreover, Section 3 shows that the change in the unweighted average only accounts for around 1% of the total increase in aggregate specialization.

Our result that more specialized firms are less sensitive to changes in Chinese import competition is parallel to that of Argente, Moreira, Oberfield, and Venkateswaran (2021). Though our empirical design is similar to theirs, there are two main distinctions. First, they focus on differential responses by firm scope, measured as the number of products or establishments within industries, whereas we focus on differential responses by diversification across industries. Second, Argente et al. (2021) treat changes in Chinese import competition as a common demand shock, and rationalize the stronger contraction of more diversified firms via a theory of joint production across a firm's products. In contrast, we view the differential response as occurring because changes in Chinese import competition are not a common shock: Chinese imports are more substitutable with the products of diversified firms, and thus the same change in Chinese imports leads to greater competitive pressure on more diversified firms.

Our theory of quality is conceptually related to but distinct from that proposed by Holmes and Stevens (2014) in that we view industries as being partitioned into segments consisting of firms that directly compete with each other. In contrast to their paper, which defines segments based on trade costs, we model segmentation based on product quality and thus by consumers' willingness-to-pay and products' income elasticities of demand. Accordingly, we add non-homothetic preferences over segments and endogenize firms' sorting into said segments.

[4]For instance, some studies argue that large diversified firms expand across space and specialize across industries in response to spatial aggregation of the economy, such as trade liberalization (Bernard, Redding, and Schott, 2011) or information technology improvement (Hsieh and Rossi-Hansberg, 2019).

[5]Specifically, Ma (2022) documents a falling unweighted average number of industries among innovating firms, defined as firms that have ever been granted a patent.

2 Data and measurement for diversification

Measuring firms' diversification across industries requires data on production activities (e.g., sales and employment) at the firm-by-industry level. In this section, we first describe said data, and the definition of firms and industries used. We then introduce our measure for diversification and map its evolution over time.

2.1 Data

Our primary data come from the U.S. Census Bureau's Longitudinal Business Database (LBD) for every year from 1977 to 2017. The LBD tracks the universe of U.S. private, non-farm establishments based on administrative records from business tax filings. In the LBD, we observe each establishment's number of employees and payroll. Moreover, the LBD tracks establishments longitudinally, enabling us to observe life cycle information such as entry, exit, and age.

Information beyond employment and payroll comes from the Census of Manufactures (CMF) every five years from 1977 to 2017, which covers the universe of U.S. manufacturing establishments. The CMF collects a variety of information, including, but not limited to, the total value of shipments, value-added, employment, wage payroll, and total value of assets.

The LBD and CMF are particularly suitable for our analysis because each establishment is assigned a firm identifier and a time-consistent industry code, which ensures that our measurement of diversification is consistent across years. Conceptually, a firm is a collection of establishments under common ownership. The firm identifier is constructed based on a combination of payroll tax filing information and questions about ownership on the Economic Census and other surveys. Industries are defined by the 6-digit North American Industry Classification System (NAICS).[6] Applying the method developed by Fort and Klimek (2018), the Census assigns a 2012 vintage 6-digit NAICS code to each establishment in every year based on the primary activity performed at the establishment.[7] Harmonized industry codes over a long period are unique in the U.S. Census data and have only become available recently.[8]

[6] Within manufacturing 6-digit NAICS is often treated as a product market.

[7] For readers who are interested in the construction of firm identifiers and industry codes in the U.S. Census data, we recommend Sections 8 and 10 in Chow, Fort, Goetz, Goldschlag, Lawrence, Perlman, Stinson, and White (2021).

[8] Most public-use data sets, such as Compustat and Orbis, are at the firm level and hence cannot be used to observe firms' cross-industry production. Other establishment-level data sets covering the U.S. do not have detailed industry codes that are consistent over a long period.

We focus our analysis on U.S. manufacturing for three main reasons. First, the manufacturing sector has better data coverage than other sectors (e.g., retail trade, wholesale trade, and services). The Census did not survey all non-manufacturing establishments until a major expansion of industry coverage in 1992. Second, the richness of the data allows us to form a more comprehensive understanding of the patterns we study. In Section 4 for example, the analysis of firms' responses to Chinese import competition is not feasible for sectors without data on trade flows; the unit-price analysis is also infeasible because we do not observe quantities from the Census of other sectors. Third, the economic meaning of NAICS industries is not compatible across broad sectors, and so neither is the resulting diversification measure. For example, for firms producing non-tradable goods, such as retail or dining services, the relevant definition of a market may not be a 6-digit NAICS industry, but an industry-and-location pair (see, e.g., Rossi-Hansberg, Sarte, and Trachter, 2020). In Section 3.7.3, we show that our main findings regarding firm diversification are (i) robust to counting manufacturing firms' production in other sectors and (ii) not unique to manufacturing.

To form our main sample, we take establishments with positive payroll and employment (and sales for the CMF), and keep only establishments in the manufacturing sector, namely, establishments whose first two digits of the NAICS code are 31, 32, or 33. For a given firm identifier, we thus only use the firm's manufacturing establishments for the benchmark analysis. We will show that our main findings regarding firm diversification are robust to counting those firms' production in other sectors. We then aggregate production variables for all establishments with the same firm identifier and 6-digit NAICS code to form our benchmark firm-by-industry-level database. Other details about sample selection are reported in Appendix A.1.

Notation. For the remainder of this paper, we index industries with $j \in J$, where J is the set of all industries, and firms with $i \in I$, where I is the set of all firms. We call firm i active in industry j if its size (sales or employment) is positive in that industry. The set I_j collects the active firms in industry j, and J_i collects firm i's active industries. For each firm-by-industry-level variable x_{ij}, we let $x_i = \sum_{j \in I_j} x_{ij}$ be the firm-level value, $x_j = \sum_{i \in I_j} x_{ij}$ be the industry-level value, and $x = \sum_{i \in I} x_i = \sum_{j \in J} x_j$ be the total value of x for each year. The operator "$|\cdot|$' denotes the number of elements of a set. For example, $|J_i|$ denotes the number of active industries for firm i, and $|I_j|$ denotes the number of active firms in industry j.

2.2 Measuring diversification

A natural measure of diversification is simply the firm's number of industries, $|J_i|$. However, this measure does not take into account how unevenly the firm's production (e.g., sales or employment) is distributed across industries. For example, one would consider a firm that produces evenly in two industries to be more diversified than another that concentrates most of its production in one of its two industries, but these two firms are assigned the same degree of diversification if measured by their number of industries. Ignoring non-uniformity in firms' distribution of production across industries thus potentially discards significant variation in diversification.

To circumvent such issues, we instead focus on the *effective number of industries* (ENI), a diversity index that reflects not only a firm's number of industries but also the evenness of how its production is distributed over industries. The ENI is the reciprocal of the Herfindahl-Hirschman Index (HHI), an index that is commonly used in the economic literature to measure market concentration. One can interpret the ENI as the hypothetical number of equally-sized industries that would exhibit the same HHI as the firm's actual employment distribution. We prefer the ENI because its units can be given a clear meaning.

In this paper, we measure production as either sales or employment. For simplicity, we will refer to the production variable as employment for the remainder of this subsection, but it should be understood that all arguments apply equally to sales. Let x_{ij} denote firm i's employment in industry j and x_i its total employment. The firm's effective number of industries is defined as:

$$\mathrm{ENI}_i = \frac{1}{\mathrm{HHI}_i} = \left[\sum_{j \in J_i} \left(\frac{x_{ij}}{x_i} \right)^2 \right]^{-1}. \tag{1}$$

Concretely, the ENI discounts a firm's number of industries by the unevenness of its distribution of employment across industries. To see this, note that the ENI is bounded between one and the firm's number of industries, $1 \leq \mathrm{ENI}_i \leq |J_i|$. For a multi-industry firm, the ENI equals the number of industries if and only if the firm's employment is uniformly distributed across its active industries. Fixing the set of active industries J_i, the ENI decreases when the distribution becomes more uneven.[9] As more and more of a firm's employment is concentrated in one industry, the

[9]Let $s_{ij} = \frac{x_{ij}}{x_i}$ be firm i's share of production in industry j. The ENI can be expressed as:

$$\mathrm{ENI}_i = \frac{|J_i|}{1 + |J_i|^2 \,\mathrm{Var}(s_{ij})},$$

where the variance is taken across industries. Fixing the set of active industries J_i, the ENI decreases with $\mathrm{Var}(s_{ij})$, ENI converges to 1. This index is widely used in political science to measure the effective number of parties, and in ecology to measure the effective number of species.[10]

2.3 Initial data description

Figure 1 plots the trend of aggregate firm diversification from 1977 to 2017 for U.S. manufacturing. Panel (a) plots the average number of industries weighted by firm employment from the LBD, showing a substantial decline over the sample period. The average U.S. manufacturing firm has about 11 industries in 1977, but the number is cut to half after 40 years. Panel (b) plots the average effective number of industries (ENI) weighted by firm size. The dark solid line is calculated from the LBD, where employment is used to measure size. When using data from the LBD, we only calculate the ENI using employment because the LBD does not include establishment-level sales information. To test if the trend is sensitive to the choice of firm size variables, we calculate the ENI based on both employment and sales using the CMF data and compare the two ENI measures. Accordingly, the light dashed line is calculated from the CMF, where sales are used to measure size. As expected, the ENI is smaller than the number of industries because most multi-industry firms' production is not evenly distributed across industries.

There are two main takeaways from the figure. First, average diversification, no matter whether measured by the actual or effective number of industries, declines by around 50% over the sample period. Although the two measures yield different values, they reveal the same trend of diversification in terms of the ratio scale. Moreover, the cross-firm correlation coefficient between the number of industries and employment-based ENI in our sample is 0.88.[11] The two measures capture similar cross-firm variation in diversification.

Second, the trends of employment-based ENI and sales-based ENI are closely aligned, decreasing from about 4 to 2 over the sample period. Indeed, the cross-firm correlation coefficient

which captures the unevenness of the employment distribution.

[10] In general, the ENI belongs to a broader class of diversity indices:

$$D_i^{\rho} = \left[\sum_{j \in J_i} \left(\frac{x_{ij}}{x_i} \right)^{\rho} \right]^{\frac{1}{1-\rho}},$$

where the parameter $\rho \leq 0$ governs the extent to which the number of industries is discounted by unevenness. This index equals the number of industries when $\rho = 0$. We choose $\rho = 2$ as our benchmark measure because the economic literature is already familiar with the concept of HHI. Another commonly used value is $\rho = 1$, where the index equals the exponential of the entropy.

[11] The correlation is weighted by employment; the unweighted correlation coefficient is 0.86.

Figure 1: Aggregate firm diversification in U.S. manufacturing, 1977-2017

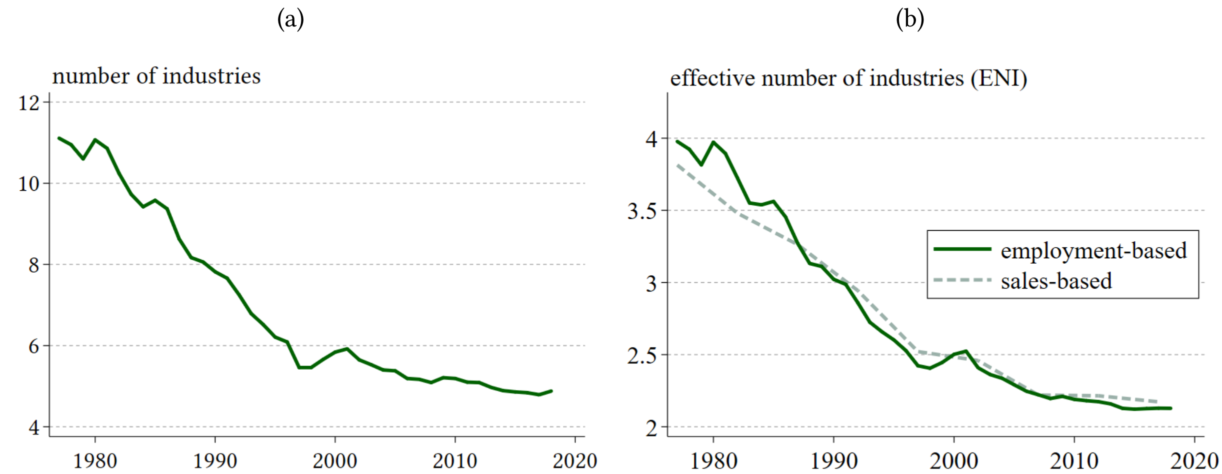

Notes: Panel (a) depicts the employment-weighted average number of industries. Panel (b) depicted the sales-weighted average sales-based ENI and the employment-weighted average employment-based. Data for number of industries and employment-based ENI are from the LBD. Data for sales-based ENI are from the CMF.

between the two ENIs is 0.92 in our sample.[12] Both the aggregate pattern and the cross-firm variation in ENI are thus robust to using employment or sales to measure size.

Because employment is available in both the LBD and the CMF, and because the change in diversification across the employment- and sales-based indices are similar in ratio scale, we use the logarithm of employment-based ENI as our benchmark measure of diversification for the remainder of this paper. For brevity's sake, we will often refer to the employment-weighted average log ENI as aggregate diversification.

3 Empirical findings

This section investigates the firm-level patterns underlying the dramatic decline in aggregate diversification in U.S. manufacturing. The findings help us direct the scope of our study and support our proposed mechanism.

[12] The correlation is weighted by employment; it is 0.89 if weighted by sales and 0.94 if unweighted.

Figure 2: Share of single-industry firms (%), 1977 vs. 2017

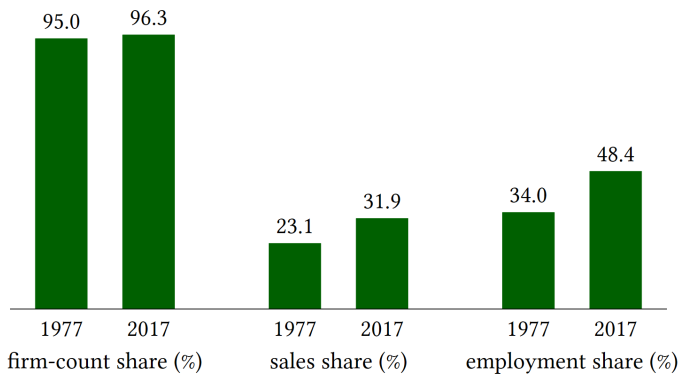

Notes: Figure depicts the firm count share, sales share, and employment share of single industry firms in U.S. manufacturing for the years 1977 and 2017 respectively. Data source: CMF.

3.1 Reallocation toward specialized firms

We begin by documenting a significant shift in production toward specialized firms between 1977 and 2017. Figure 2 focuses on the most specialized firms—firms that only produce in one industry. The figure compares the share of single-industry firms in 1977 versus 2017, where the left two bars display the firm-count shares, the middle two display the sales shares, and the right two display the employment shares. Perhaps not surprisingly, most firms are single-industry firms and most single-industry firms are relatively small. In 1977, 95% of manufacturing firms operate in a single industry, but they only account for 23% of sales and 34% of employment in the sector. Forty years later, the firm-count share of single-industry firms has only increased by 1 percentage point, but they now account for a 9pp higher share of sales and a 15pp higher share of employment.

A similar pattern holds among multi-industry firms. In Figure 3, we group multi-industry firms by their log effective number of industries (ENI). Panel (a) shows the distributions of the number of firms against log ENI, where the height of each bar represents the share of multi-industry firms in that group. The filled bars depict the distribution in 1977, and the hollow bars depict the analogous distribution in 2017. The figure shows that these distributions are nearly identical. Panel (b) and (c) show the distribution of sales and employment against log ENI respectively, where the height of each bar represents the share of sales or employment relative to the total among multi-industry firms. Compared with the 1977 distribution, more specialized firms in 2017 account for a significantly larger market share, whereas diversified firms have lost market

Figure 3: Distribution of diversification among multi-industry firms, 1977 vs. 2017

(a) unweighted

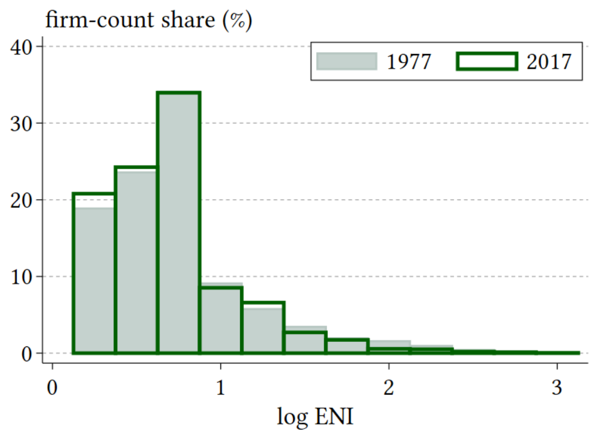

(b) sales-weighted

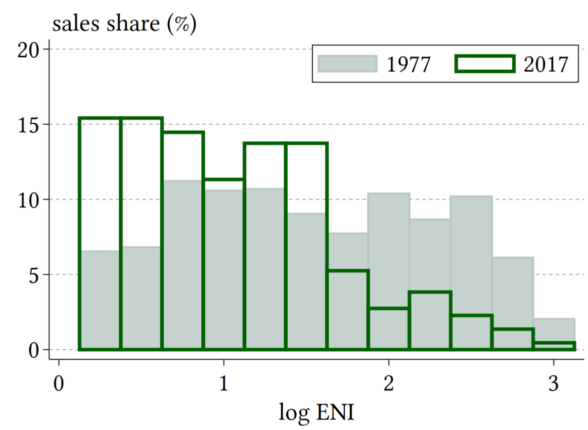

(c) employment-weighted

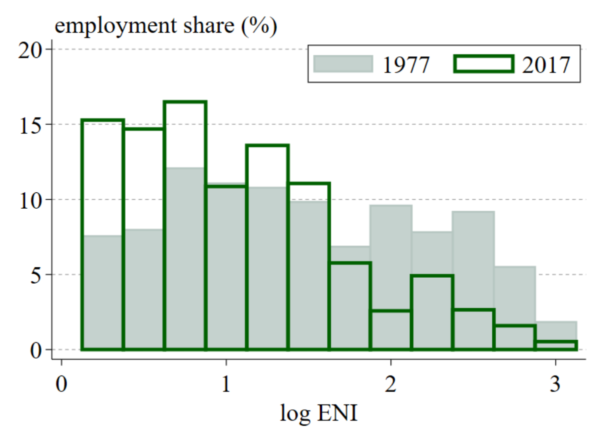

Notes: Firms are grouped by their log ENI. The grouping cutoffs are 0.25, 0.5, ..., and 3, where each group collects firms whose log ENI is weakly smaller than the cutoff and is strictly greater than the previous cutoff. To comply with Census disclosure rules, we distort the right tail by (i) aggregating firms whose log ENI is greater than 2.5 and (ii) reassigning the shares to group 2.5, 2.75, and 3 by 50%, 30%, and 20% of the aggregated value. The same treatment is applied to all four distributions. Data source: CMF.

share. To illustrate the magnitude, take $\ln \text{ENI}_i = 1.5$ as a reference point separating specialized and diversified firms. About 30% of sales and 20% of employment among multi-industry firms has been reallocated from diversified firms to specialized firms.

The key takeaway is that specialized firms now account for a significantly greater share of production than they did 40 years ago, and this is not because there are relatively more of them, but because they have become larger than before.

3.2 Reallocation is a secular trend among continuing firms

We proceed by decomposing the pattern observed in Figure 3 using the full LBD panel from 1977 to 2017. This exercise produces two insights: first, that the reallocation occurs at a similar rate throughout our sample period, and second, that the trend is among continuing firms rather than among entering or exiting firms.

We implement this analysis by using the accounting method developed by Melitz and Polanec (2015). This framework decomposes the change of an aggregate variable (size-weighted average) into the contributions of net entry (entry minus exit) and those of continuing firms. In our case, the aggregate variable of interest is the employment-weighted average of log ENI. To simplify notation, we let $n_{i,t} = \ln \text{ENI}_{i,t}$, and let n_t denote its size-weighted average in year t,

$$n_t = \sum_{i \in I} \left(\frac{x_{i,t}}{x_t} \right) n_{i,t}, \tag{2}$$

where $x_{i,t}$ here denotes firm i's employment in year t, and x_t the total manufacturing employment in year t. For any two consecutive years, $t = 1$ and 2, we categorize firms that appear in at least one year ($I_1 \cup I_2$) into three groups: continuers ($C = I_1 \cap I_2$), entrants ($E = I_2 \setminus I_1$), and exiters ($X = I_1 \setminus I_2$). As shown in Melitz and Polanec (2015), the annual change in aggregate diversification, $\Delta n = n_2 - n_1$, can be decomposed into the following four terms:

$$\Delta n = \underbrace{\left(\frac{x_2^E}{x_2} \right) \left(n_2^E - n_2^C \right) - \left(\frac{x_1^X}{x_1} \right) \left(n_1^X - n_1^C \right)}_{\text{net entry}} + \underbrace{\Delta \, \overline{n}^C}_{\text{within } C} + \underbrace{\Delta \, \text{cov}_C \left(\frac{x_i}{\overline{x}^C}, \, n_i \right)}_{\text{between } C}. \tag{3}$$

The first two terms capture the contribution of net entry, where x_t^g denotes the total employment and n_t^g denotes the conditional average log ENI (weighted by size) for each group $g \in \{E, X, C\}$ in each year t. Supposing entering firms are, on average, less diversified than continuing firms

($n_2^E < n_2^C$), entry would drive down aggregate diversification, and the magnitude of this decline increases with the employment share of entering firms. Likewise, supposing exiting firms tend to be less diversified than continuing firms ($n_1^X < n_1^C$), exit would raise aggregate diversification, and the magnitude of this rise increases with the employment share of exiting firms.

The third and fourth terms are the Olley and Pakes (1996) decomposition among continuing firms. The third term, $\Delta \overline{n}^C$, is the change of the unweighted average of log ENI among continuers. This term captures the *within-firm* change of diversification. If this term is the main driver of the aggregate change, one would not lose much information viewing the trend as a representative firm changing diversification over time.

The last and most important term, $\Delta \mathrm{cov}_C \left(x_i / \overline{x}^C, n_i \right)$, is the change in the covariance between firm size and diversification among continuers, where size is normalized by the average size among continuing firms ($\overline{x}^C$). This term captures the shift in production activities between continuers with differing log ENI. Its magnitude points to the importance of firm heterogeneity in explaining the aggregate change in diversification.

Using the panel of manufacturing firms in the LBD, we calculate the annual change for each term on both sides of equation (3) from 1977 to 2017. We then calculate the cumulative sum of the series, which converts annual changes into differences from their 1977 level. The resulting series are plotted in Figure 4, where the light solid line shows the total change in aggregate log ENI, and the dark lines depict the three components on the right-hand side of equation (3). Conceptually, each dark line represents a hypothetical trend of aggregate diversification shutting off the contributions of other components. To construct these contributions, we project the series of each component on that of the total change in aggregate log ENI and take the resulting slopes as their contributions, which are reported in parentheses following the line labels in the figure.

Overall, the average log ENI declines by about 0.4 log points, indicating that, on average, firms cut about 33% of their effective number of industries. Without the changes from continuing firms, the net entry margin slightly lowers aggregate diversification in earlier years but later increases it by about 0.07 log points relative to the 1977 level. This is because the (employment-weighted) entry rate is declining while the exit rate is stable (see Figure 2 in Decker, Haltiwanger, Jarmin, and Miranda, 2016). On average, entering and exiting firms are less diversified than continuing firms. In earlier years, the entry rate was high enough to dominate exiting firms' impact and bring down the average diversification. The entry rate declines over time, and the exit component dominates in latter years, which drives up the aggregate diversification.

Figure 4: Decomposing changes in average log ENI from 1977

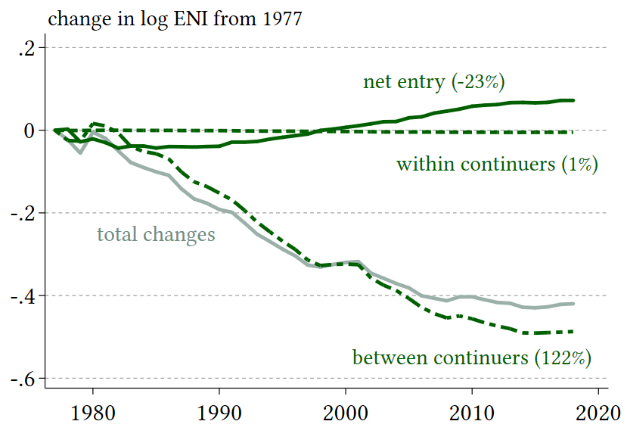

Notes: Figure shows the accounting results decomposing the change in employment-weighted average log ENI according to equation (3). Data source: LBD.

On the other hand, shutting off entry and exit, the average log ENI among continuers declines by about 0.5 log points, which means that continuing firms, on average, cut about 40% of their effective number of industries. The within-firm component only accounts for 1% of the aggregate trend. The primary driver of the decline in aggregate log ENI is instead the between-firm reallocation among continuing firms, accounting for 122% of the total change, where the extra change beyond the total is offset by the entry and exit margin. The accounting results are consistent with our earlier observation from Figures 2 and 3 that the unweighted distribution of log ENI is unchanging over the last 40 years, whereas the employment-weighted distribution shifts dramatically to the left. More importantly, the accounting results confirm that the reallocation toward specialized firms occurs constantly throughout our sample period and is not driven by net entry.

3.3 Reallocation reflects faster growth of specialized firms

We next exploit the panel dimension of our data set to show that the reallocation toward specialized firms largely reflects the faster growth of specialized firms rather than changes in how firms diversify over time. We view this result as key because it allows us to rule out certain conjectures as important drivers of the trend.

To implement this comparison, we ask whether the changing covariance between firm size and diversification comes from (i) large firms decreasing diversification relative to the average

Figure 5: Decomposing changes in average log ENI from 1977

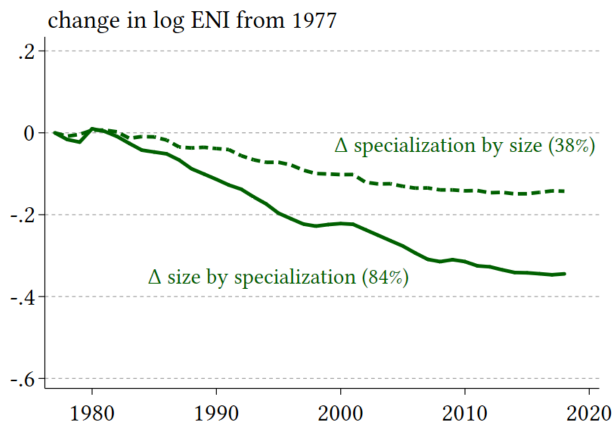

Notes: Figure shows the accounting results decomposing the change in employment-weighted average log ENI according to equation (4). Data source: LBD.

firm or (ii) more diversified firms growing more slowly than more specialized firms. Letting $\tilde{x}_{i,t} = x_{i,t} \,/\, \overline{x}_t^C$, we decompose the between-firm component into the following two terms:

$$\Delta \mathrm{cov}_C\left(\tilde{x}_i,\ n_i\right) \;=\; \underbrace{\mathrm{cov}_C\left[\Delta n_i, \left(\frac{\tilde{x}_{i,1}+\tilde{x}_{i,2}}{2}\right)\right]}_{\Delta \text{ specialization by size}} + \underbrace{\mathrm{cov}_C\left[\Delta \tilde{x}_i, \left(\frac{n_{i,1}+n_{i,2}}{2}\right)\right]}_{\Delta \text{ size by specialization}}. \tag{4}$$

For the first term, we fix firms' size at their average between the two periods and let log ENI change. Given that the unweighted mean does not change, a negative covariance indicates that (initially) larger firms tend to reduce their ENI, and (initially) smaller firms tend to increase their ENI. The magnitude of this term emphasizes the importance of firms' changes in diversification given their size. For the second term, on the other hand, we fix the firms' log ENI at their average between the two periods and let firm size change. In this case, a negative covariance indicates that (initially) more diversified firms become relatively smaller, and (initially) more specialized firms grow relatively larger. The magnitude of this term emphasizes the importance of firms' growth rates given their diversification in driving the observed reallocation in production activity.

We plot the accounting results in Figure 5. Each line depicts a component on the right-hand side of equation (4), which is converted into its difference from the 1977 level. In the parentheses following the line labels, we report their contributions to the aggregate log ENI change, so the two numbers add up to 122%, the total contribution of continuing firms' between-firm reallocation. We can see that both components play non-trivial roles in driving the reallocation, while the second term, which captures the faster growth of specialized firms, contributes about 70% of the shift among continuing firms and 84% of the total change.

How can this pattern sharpen our understanding of the reallocation toward specialized firms? First, it suggests that rising entry barriers are likely not a significant driver of the trend. The logic is as follows: if entry barriers are increasing, we should expect that firms diversify less over their life-cycle. However, this would show up as a fall in the first term of equation (4), which we show is not the main driver of the trend. Second, we do not see large firms' increasing focus on their core competencies as being the main driver of the trend. The argument is similar: initially large firms cutting their diversification due to, e.g., higher competition forcing them out of their marginal markets (as argued in Hsieh and Rossi-Hansberg, 2019) would result in a decrease in the first term of equation (4).

3.4 Reallocation occurs within almost all industries

We next examine industry-level changes in average diversification. We first show that the reallocation toward specialized firms occurs within almost all industries. Second, we show that a negligible amount of the overall reallocation toward specialized firms is due to the expansion of industries whose firms tend to be more specialized.

To proceed we first define $n_j = \sum_{i \in I_j} (x_{ij}/x_j) n_i$, the size-weighted average log ENI within industry j. The aggregate log ENI, namely, the size-weighted average across firms, can also be written as industry averages weighted by industry size, $n = \sum_{j \in J} (x_j/x) n_j$. Panel (a) of Figure 6 contrasts the industry-level log ENI (n_j) in 2017 against its value in 1977. The black dashed line is the 45-degree line. We see that almost all dots are located below the 45-degree line, indicating that average diversification decreases in almost all industries. As mentioned, the aggregate log ENI declines by about 0.4 log points from 1977 to 2017. Here we find that about 90% of industries (75 out of 85) experience a decline in log ENI greater than 0.2 log points (cutting about 20% of ENI). Note that we construct the average log ENI at the 4-digit NAICS level; ideally, we would show this scatter plot at the 6-digit NAICS level, but we have to aggregate to the 4-digit level to avoid disclosure risk for the Census data, and we so present the 4-digit level plot as an approximation to the 6-digit level plot.

The main takeaway so far is that most industries experience a substantial decline in average firm diversification. From Panel (a), we also observe significant cross-industry variation in diversification. The log ENI of most industries lies between 0.5 and 1.5 in 1977, versus between

Figure 6: Change in specialization: within vs. between industries

(a) 4-digit NAICS, 1977 vs. 2017

(b) Average across 6-digit NAICS, 1977-2017

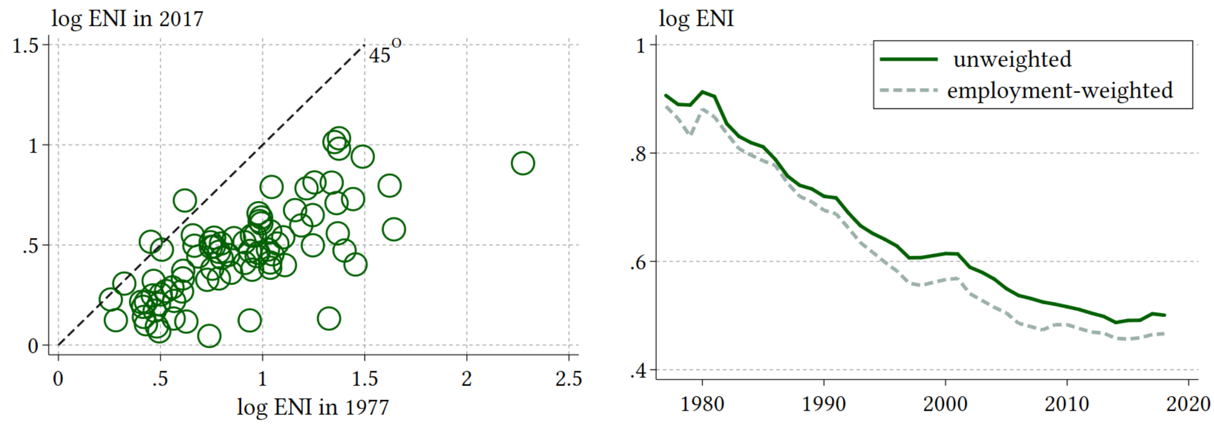

Notes: Panel (a) plots each 4-digit NAICS industry's average log ENI in 2017 against their value in 1977. Panel (b) reports the accounting results decomposing the change in employment weighted average log ENI according to equation (5). Data source: LBD.

0 and 1 in 2017. The cross-industry variation is thus at least as large as the cross-time change within industries. This points to another potential driver behind the reallocation toward specialized firms—an expansion of more specialized industries. To calculate the contribution of this channel, we again make use of the accounting framework developed by Olley and Pakes (1996):

$$n = \sum_{j \in J} \left(\frac{x_j}{x} \right) n_j = \underbrace{\frac{1}{|J|} \sum_{j \in J} n_j}_{\text{within-ind.}} + \underbrace{\operatorname{cov} \left[\left(\frac{x_j}{\bar{x}} \right), n_j \right]}_{\text{between-ind.}}. \tag{5}$$

The first term on the right-hand side is the unweighted average across industries. The change in this term captures the within-industry change in diversification. The second term is the covariance between industry size and diversification, where industry size is normalized by the average $\bar{x} = x/|J|$. This term captures the reallocation between industries with differing levels of diversification, i.e., the expansion of more specialized industries.

Panel (b) of Figure 6 plots the time series of average log ENI (measured by employment) from 1977 to 2017 using our LBD sample. The light dashed line depicts the average weighted by industry size, and the dark solid line depicts the unweighted average, i.e., the within-industry component. The change in the unweighted average accounts for almost all of the change in aggregate diversification. Projecting the unweighted average on the weighted average, we find that about 96% of the change in aggregate diversification occurs within industries. The between-industry component is negative but small in magnitude. Industries with more specialized firms are slightly larger, and this pattern strengthens over time, but the magnitude of this change only accounts for 4% of the aggregate change.

3.5 Reallocation toward specialized firms increases industry concentration

So far, we have shown that specialized firms now account for a greater share of production than 40 years ago, and that this trend occurs in most industries. How large have specialized firms become? Have they grown large enough to alter the size ranking in their industries and become industry leaders? How are these patterns related to the widely-discussed rise in industry concentration?

Following the literature, we define industry leaders as the 4 largest firms in each industry, denoted as T_j. The average industry concentration ratio (CR) is then defined as the share of sales accounted for by industry leaders:

$$\mathrm{CR} = \frac{1}{x} \sum_{j \in J} \sum_{i \in T_j} x_{ij}, \tag{6}$$

where $x_i j$ denotes sales of firm i in industry j. As documented by Autor et al. (2017, 2020), CR has increased in many sectors over the past decades. To connect our findings with this trend, we study how much the distribution of diversification *among industry leaders* changes from 1977 to 2017. To do so, we again categorize firms into 16 groups by their log ENI. Specifically, indexing the groupings by g, we denote the set of firms in each group as $I_g = \{i \in I \,|\, 0.25(g-1) < \ln \mathrm{ENI}_i \leq 0.25g\}$ for $g = 0, 1, 2, \ldots, 14$, and $I_{15} = \{i \in I \mid \ln \mathrm{ENI}_i > 2.75\}$. For each group g, we calculate the share of industry leaders whose log ENI is in the range defining g:

$$\text{unweighted share among leaders}_g = \frac{\sum_{j \in J} |T_j \cap I_g|}{\sum_{j \in J} |T_j|}. \tag{7}$$

The results are depicted in Panel (a) of Figure 7. Each bar represents a group g, and the height shows the group's share among industry leaders. The filled bars depict the shares in 1977, and the hollow bars depict the analogous distribution in 2017. We find that specialized firms' presence as industry leaders has increased dramatically from 1977 to 2017. In 1977, industry leaders tended to be more diversified firms, but the distribution has shifted substantially to the left. In 2017, about one-fourth of industry leaders are single-industry firms.

Figure 7: Share of top 4 industry leaders by log ENI

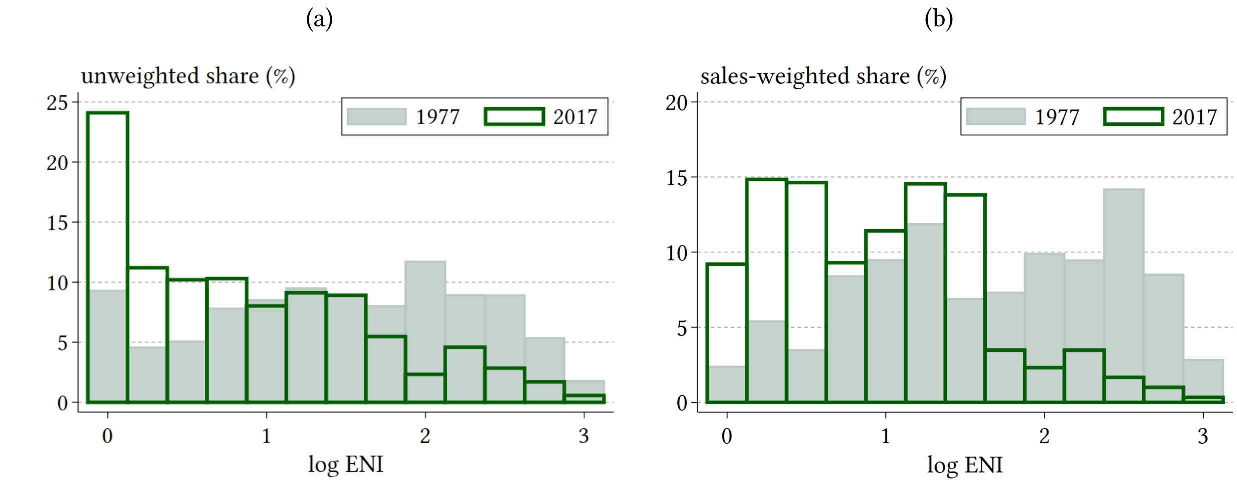

Notes: Firms are grouped by their log ENI. The grouping cutoffs are 0.25, 0.5, ..., and 3, where each group collects firms whose log ENI is weakly smaller than the cutoff and is strictly greater than the previous cutoff. To avoid disclosure risks for the Census data, we distort the right tail by (i) aggregating firms whose log ENI is greater than 2.5 and (ii) reassigning the shares to group 2.5, 2.75, and 3 by 50%, 30%, and 20% of the aggregated value. Data source: CMF.

In Panel (b) of Figure 7, we report each group's share of sales among industry leaders:

$$\text{sales share among leaders}_g = \frac{\sum_{j \in J} \sum_{i \in T_j \cap I_g} x_{ij}}{\sum_{j \in J} \sum_{i \in T_j} x_{ij}}. \tag{8}$$

This share shows each group's contribution to industry concentration, as measured by the concentration ratio. The figure shows that specialized firms now contribute much more to industry concentration than 40 years ago. To demonstrate the magnitude of this change, take $\ln \text{ENI}_i = 1.5$ as a reference point separating specialized and diversified firms, then about 40% of sales among industry leaders has been reallocated from diversified firms to specialized firms. Compared with Figure 2 and 3, the shift among industry leaders is larger than the overall reallocation toward specialized firms.

So far we have shown that specialized firms have grown sufficiently fast to largely displace diversified firms as industry leaders, but this is insufficient to conclude that the rise of specialized firms has resulted in the documented increase in industry concentration. If, for instance, diversified and specialized industry leaders have on average similar sizes, then simply replacing diversified leaders with specialized leaders would not tend to increase industry concentration.

To approach this issue, we next consider the following counterfactual scenario: if specialized firms' market share had not increased, would industry concentration still have increased? To construct this counterfactual, we again take $\ln \mathrm{ENI}_i = 1.5$ as a reference point separating specialized and diversified firms. The concentration ratio can be written as:

$$\mathrm{CR}_t = \frac{x_{l,t}}{x_t} = \frac{x_{l,t}^{\mathrm{spec}} + x_{l,t}^{\mathrm{divrs}}}{x_t^{\mathrm{spec}} + x_t^{\mathrm{divrs}}},$$

where x_l denotes the total sales of industry leaders, $x_l^{\mathrm{spec}}, x_l^{\mathrm{divrs}}$ denote the total sales of specialized and diversified leaders, respectively, and x^{spec}, x^{divrs} denote the total sales of specialized and diversified firms (i.e. not restricted to leaders). To construct the concentration ratio in the counterfactual scenario, we can fix the average relative size of specialized firms and specialized leaders to their 1977 level, while allowing diversified firms and leaders to reach their size in each year t. The counterfactual concentration ratio is thus:

$$\mathrm{CR}_t^{\mathrm{cf}} = \frac{x_{l,t}\left(x_{l,1977}^{\mathrm{spec}}/x_{l,1977}\right) + x_{l,t}^{\mathrm{divrs}}}{x_t\left(x_{1977}^{\mathrm{spec}}/x_{1977}\right) + x_t^{\mathrm{divrs}}}. \tag{9}$$

Figure 8 compares the counterfactual industry concentration ratio against its original value for U.S. manufacturing, 1977-2017. Panel (a) defines leaders as the top 4 largest firms. The dashed lines depict the original concentration ratio (CR_t), and the solid lines depict their counterfactual values ($\mathrm{CR}_t^{\mathrm{cf}}$). Over the past 40 years, the top 4 concentration ratio increased from around 40% to 45%. On the other hand, in the counterfactual scenario where the rise of specialized firms does not take place, the concentration ratio would have actually decreased: $\mathrm{CR}_{2017}^{\mathrm{cf}} = 35.5\%$. The gap between the two lines shows the contribution: the faster growth of specialized firms, and particularly the faster growth of specialized leaders, accounts for about a 10 percentage point increase in industry concentration. As shown in Panel (b), this result is robust to defining leaders as the top 20 largest firms.

3.6 Summary

We have shown above that there is a secular trend toward firm-level specialization in U.S. manufacturing: the average firm-level effective number of industries falls by close to 50%. This trend in turn is due to a reallocation of activity among continuing firms rather than a fall in most firms' level of diversification. The reallocation is largely due to the faster growth of specialized firms rather than large firms specializing. It is broad-based, occurring in nearly every manufacturing

Figure 8: Counterfactual industry concentration ratio, 1977-2017

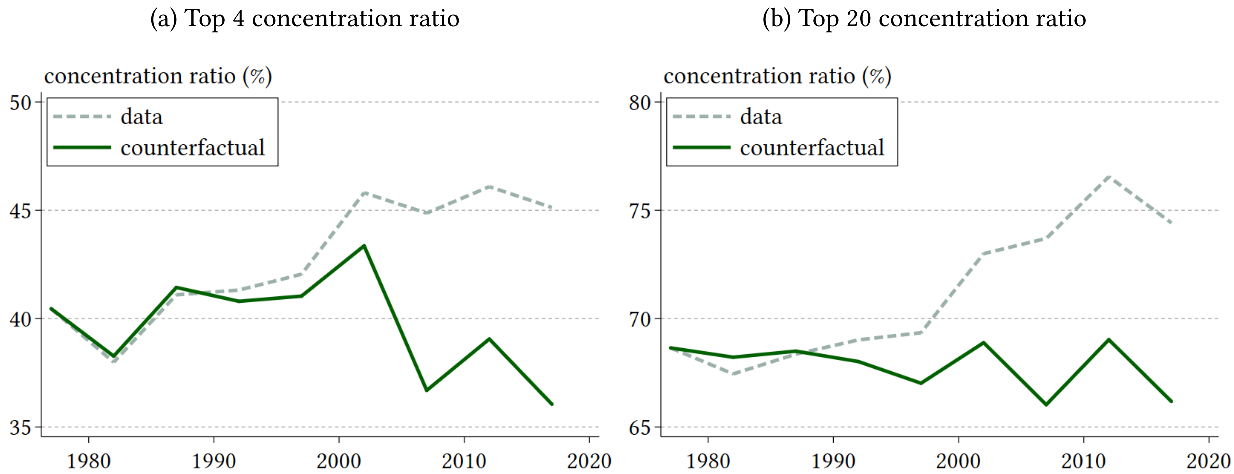

Notes: The counterfactual concentration ratios are defined in equation (9). Data source: CMF.

industry. Finally, it is particularly pronounced among industry leaders and can account for more than the total observed change in industry concentration.

We derive two additional lessons from these empirical findings. First, that the reallocation occurs over nearly the entire period over which it is possible to observe firms' diversification (the LBD begins in 1976, the CMF in 1977) motivates our view of this trend as secular rather than reflecting some movement along a transition path. Adding plausibility to this view is that the reallocation is observed in nearly all industries.

Secondly, we do not interpret the trend as reflecting large diversified firms' increasing focus on their core competencies. Our argument is two-fold: first, such a change should reflect on the distribution of diversification among multi-industry firms, the overwhelming majority of which are large firms. However, as shown in the left-hand panel of Figure 3, this is not the case. Moreover, in section 3.2 we document that the trend reflects specialized firms growing rather than large firms specializing. To be clear, we are not ruling out that large firms are increasingly focusing on their core competencies, but we can only attribute a small portion of the overall rise of specialized firms to this channel.

In addition, we want to clarify that market concentration only increases at the detailed industry level: when production over industries is aggregated to the firm level, the share of largest firms in the U.S. has not increased over time (as documented by Hsieh and Rossi-Hansberg, 2019). This is because industry leaders may not be leaders nationally. Particularly now industry leaders are more specialized, and less likely to be leaders nationally. Some studies combine the rising industry concentration with other firm-level patterns, such as those regarding R&D or patents, to support their economic models. We remark that such attempts would be misleading because their unit of observation are inconsistent, and the resulting misconception worsens over time because of the expansion of specialized firms.

3.7 Other empirical results

3.7.1 Variation in diversification reflects horizontal diversification

Some of the variation in our measure of diversification may reflect variation in vertical diversification (sometimes called vertical integration), which refers to a firm's expansion into industries in order to supply its other industries with inputs, rather than horizontal diversification, which refers to a firm's expansion into industries that do not serve as input or output for one another. Our interpretation of the reallocation toward specialized firms is based on the assumption that most of the variation in diversification comes from horizontal diversification. Below, we confirm that this is the case.

We exploit the interplant transfer variable to approximate horizontal and vertical diversification. To measure output, the CMF asks establishments to report their total value of shipments. We and most studies refer to this variable as total sales, but it does not only count the value of products sold to the market. Instead total shipments also include interplant transfers, the estimated market value of products transferred to other establishments of the same firm.[13]

Data on interplant transfers is useful to separate horizontal and vertical diversification: the former does not create interplant transfers, but the latter does. After aggregating the CMF data to the firm-level, we define for each firm:

$$\text{market sales} = \text{total shipments} - \text{interplant transfers}.$$

We then calculate the effective number of industries (ENI) for each firm using the market sales variable. The resulting measure only captures horizontal diversification. We can hence assess the relative magnitude of the two types of diversification by comparing the ENI measured by total

[13] Definition of interplant transfer in Census Glossary: "In the case of multiunit companies, the manufacturer was requested to report the value of products transferred to other establishments of the same company at full economic or commercial value, including not only the direct cost of production but also a reasonable proportion of "all other costs" (including company overhead) and profit.' See: https://www.census.gov/glossary/.

Figure 9: Aggregate ENI, total shipment versus external sales

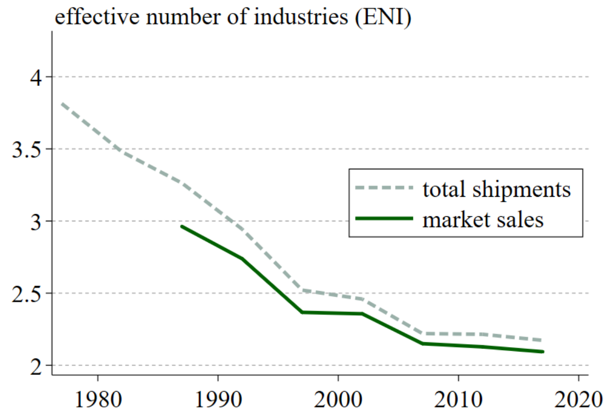

Notes: The light dash line is identical to the one in Panel (b) of Figure 1. Data source: CMF.

shipments versus that measured by market sales.

To do so, we first project the market sales-based ENI on the total shipments-based ENI using all firms in our CMF sample (weighted by total shipments). The resulting slope is 0.95 (standard error 0.02), indicating that about 95% of the cross-firm variation in firm diversification is horizontal.[14] Secondly, we compare the time series of the two average ENIs. In Figure 9, the dashed line depicts the average total shipments-based ENI where firms are weighted by total shipments, and the solid line depicts the average market sales-based ENI where firms are weighted by market sales. The two lines are close and co-move closely. Dividing the market sales-based ENI by the total shipments-based ENI for each year and taking the average over time, we find that horizontal diversification accounts for about 95% of the "level' of aggregate firm diversification. We then divide the total change of market sales-based aggregate ENI by that of the shipment-based aggregate ENI. The resulting ratio is 0.8, indicating that 80% of the aggregate change in firm diversification reflects the change in horizontal diversification. In summary, most of our documented patterns in firm diversification are horizontal. Vertical integration is present and changes over time, but its impact on the overall trend is much smaller than that of horizontal diversification.

[14]Taking log of the two ENIs before projection, the resulting slope is 0.97 (standard error 0.01).

3.7.2 Reallocation is robust to definition of manufacturing firms

The baseline sample consists of firms that produce in at least one manufacturing industry (hereafter referred to as manufacturing firms), and only counts their employment and sales in manufacturing industries when calculating their size and effective number of industries (ENI). However, some manufacturing firms also produce in other industries.[15] If diversified firms increase their non-manufacturing share while decreasing their manufacturing share, this could show up as a fall in diversification using our baseline definition of diversification. To show that this is not the case, we take the firms in the baseline sample, include all of their 6-digit NAICS industries in the LBD, and recalculate their size and ENI. We call this the extended manufacturing sample.

Panel (b) of Figure 10 compares the trend of employment-weighted average ENI in the extended manufacturing sample to the trend in the baseline sample. We then repeat the analysis of sections 3.2 and 3.3 for the extended manufacturing sample. The first two lines of the left panel of Figure 11 report the results for equation (3), and the first two lines of the right panel report that for equation (4). The length of each bar represents the term's percentage contribution to the aggregate change in log ENI.

Not surprisingly, incorporating production in other sectors substantially raises the level of ENI for manufacturing firms. A more important takeaway is that our baseline results in Section 3.2 are fully robust to extending the definition of the baseline sample. The rise of specialized firms is thus not driven by diversified firms' shifting production to other sectors.

3.7.3 Reallocation not unique to manufacturing

The above sections focus exclusively on U.S. manufacturing. Is the reallocation toward specialized firms observed in other sectors? To answer this question, we focus on four other sectors: retail, wholesale, non-tradable services (NT-services), and tradable services (T-services).[16] For each sector, we follow our baseline approach keeping only employment and sales in that sector when calculating firm size and ENI (see Appendix A.1 for details).

Panel (a) of Figure 10 shows the trend of employment-weighted average ENI for each of these sectors. We again repeat the analysis of sections 3.2 and 3.3 for the four non-manufacturing

[15]See Ding, Fort, Redding, and Schott (2022) for a comprehensive study of the cross-sector activities of manufacturing firms.

[16]The sectors are defined using the first 2 digits of the NAICS code. Retail refers to 44-45; wholesale is 42; non-tradable services are 53, 56, 71, 72, 81; and tradable services are 51, 52, 54, 55.

Figure 10: Aggregate trend of diversification beyond manufacturing, 1977-2017

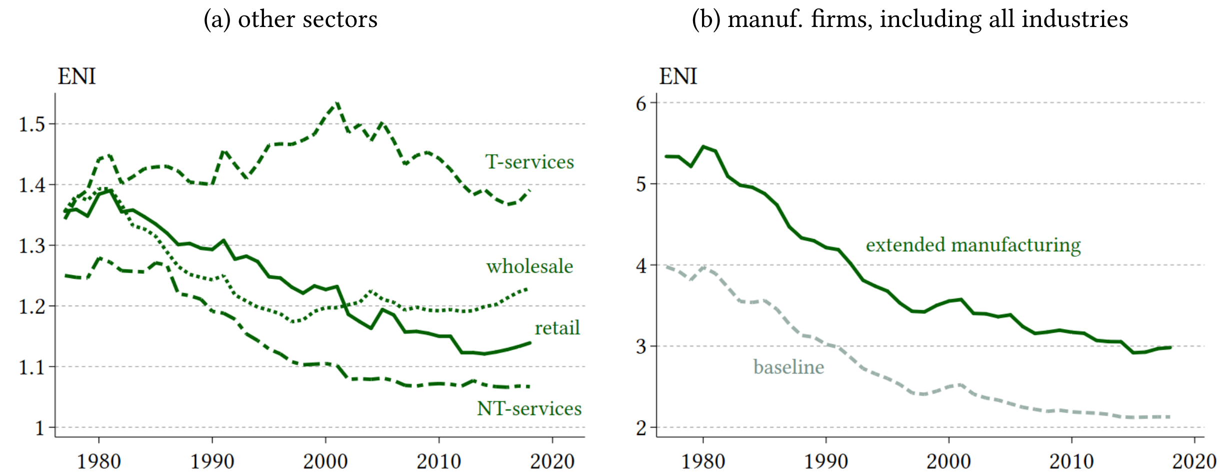

Notes: The light dash line in Panel (b) is identical to the dark solid line in Panel (b) of Figure 1. Data source: LBD.

Figure 11: Melitz-Polanec Decomposition, beyond manufacturing

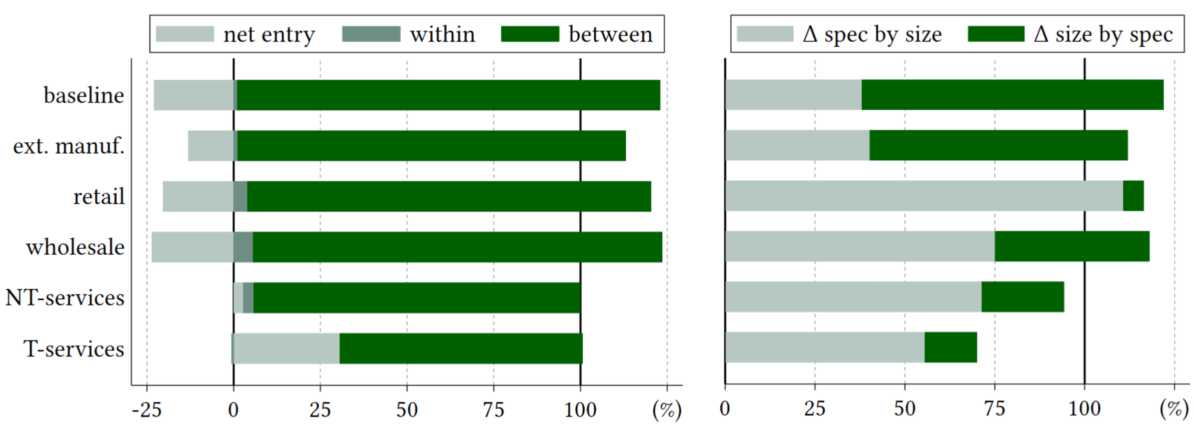

Notes: The left panel shows the Melitz-Polanec decomposition results according to equation (3), and the right panel shows the results for equation (4). The length of each bar represents the term's percentage contribution to the aggregate change in log ENI. Data source: LBD.

sectors. The bottom four lines of the left panel of Figure 11 report the results for equation (3), and the bottom four lines of the right panel report that for equation (4). The length of each bar represents the term's percentage contribution to the aggregate change in log ENI.

On average, non-manufacturing firms are much more specialized than those in the manufacturing sector, though this likely is partially a reflection of 6-digit NAICS being much coarser within non-manufacturing sectors. Three sectors (retail, wholesale, and non-tradable services) also exhibit declining trends in diversification driven by a reallocation of production toward specialized firms, though the magnitudes are smaller than in manufacturing. The main difference is that the reallocation in these sectors is mainly due to changes in specialization by firms with different sizes (the right panel of Figure 11). One sector stands out—tradable services' aggregate diversification increases before the 2000s and then falls back to its 1977 level. Our decomposition does not apply to this case because the overall change (denominator) is zero.

We next turn to the extended manufacturing sample. Perhaps not surprisingly, incorporating production in other sectors substantially raises the level of ENI for manufacturing firms. A more important takeaway is that our baseline results in Section 3.2 are fully robust in the extended manufacturing sample. The rise of specialized firms is thus not driven by diversified firms' shifting production to other sectors.

4 Establishing a link between specialization and product quality

Our proposed theory relies on specialized firms producing higher quality goods, in which case growth, by reallocating activity toward firms producing high-quality goods, leads to a rise in specialized firms. In this section we first support this relation by examining a commonly-used proxy for product quality—unit prices. We then turn toward studying firms' responses to the rapid growth in Chinese imports beginning in the 1990s.

4.1 Specialized firms charge higher unit prices

The product trailer of the CMF collects information on the physical quantity of shipments for some establishments when a meaningful metric is available. This variable is reported based on the North American Product Classification System (NAPCS), a 10-digit product code whose first

6 digits are NAICS.[17] To construct firm-product level unit prices, we divide each firm-product's market value of shipments (sales) by its physical quantity. We use non-imputed observations from the 1997-2012 CMF product trailers (as suggested by Kehrig and Vincent, 2021; White, Reiter, and Petrin, 2018). See Appendix A.2 for details on sample selection.

4.1.1 Empirical design and baseline results

Quantities are only comparable within product categories, and thus so are unit prices. Accordingly, we control for year-by-product fixed effects to capture the deviation of log price from the year-product mean. Indexing product categories by c, we run the following regression:

$$\ln \text{unit price}_{ict} = \text{FE}_{ct} + \beta_1 \ln \text{ENI}_{it} + \beta_2 \ln \text{sales}_{ict} + r_{ict}, \tag{10}$$

where FE_{ct} denotes the fixed effects and r_{ict} the regression residual. Our coefficient of interest is β_1, which captures how differently diversified firms charge different unit prices. We control for firm-by-product level sales for two reasons. First, firms' pricing behavior may change along their life cycle, and, as documented in the previous section, diversification is correlated with firm size in a non-stationary manner. Second, a firm's higher price might reflect its higher markup, and markup depends on market share in many models in the literature. Note that our definition for diversification is unchanged relative to previous sections: each firm's ENI is still calculated at the level of 6-digit NAICS industries.

The results are reported in Table 1. The first column reports the estimates and robust standard errors for simply projecting unit prices on log ENI without controlling for sales. The second column then reports results for the full regression described above. We multiply the dependent variable by 100, so the estimated β_1 is the percentage change in unit price predicted by a log point (about 2.7 times) increase in ENI. For both specifications, we find that more specialized firms tend to charge higher unit prices within narrowly-defined product categories. According to our preferred specification (Column 2), firms with a log point more ENI are predicted to charge a 3.65% lower unit price. To put this number in perspective, the interquartile range of the employment-weighted log ENI distribution among multi-industry firms in 2017 is about 1.5 log points, as indicated by Panel (b) of Figure 3, which implies a 5.5% price differential between the first and third quartiles.

[17]In the CMF and LBD, each establishment has only one NAICS code, but it could have multiple NAPCS products in the product trailer.

Table 1: Regression results for unit price analysis

specification	(1)	(2)
dependent variable	ln price × 100	
ln ENI	-1.61	-3.65
	(0.69)	(0.70)
ln sales		0.58
		(0.18)
observations	133,000	

Notes: The sample size is truncated to the nearest thousand in accordance with the Census disclosure requirement. Robust standard errors are reported in parentheses right below their estimates. Regressions control for year-by-product level fixed effects. Data source: CMF.

4.1.2 Ruling out alternative interpretations

We interpret specialized firms' higher unit prices as indicating their higher product quality. However, firms' higher prices might also reflect higher distortions in the input market (e.g., financial frictions) or product market (e.g., markup). This subsection attempts to rule out these alternative interpretations by looking at specialized firms' payments to factors of production, e.g., labor, capital, and materials.

The insight is that firms facing higher distortions should have a larger gap between their unit prices and their factor costs per unit of production (as argued in, e.g., Hsieh and Klenow, 2009). In other words, these firms charge higher unit prices without paying as much to their factors of production. Following this logic, we would ideally project unit costs (i.e., total factor costs divided by physical quantity) on log ENI and compare the projection with that of unit prices in Table 1.[18] However, this regression is not feasible because we do not observe factor costs at the firm-by-product level in the CMF product trailer.

Instead, we look at the costs-to-sales ratio (total factor costs divided by total sales) at the firm-by-industry level using the full CMF sample.[19] If specialized firms face higher market dis-

[18] If specialized firms' higher unit prices mostly reflect higher quality, we would expect the two projections to be similar; on the other hand, if they mostly reflect higher distortions, we would expect the slope resulting from the projection of unit costs on log ENI to be closer to zero or even positive.

[19] Industries are again defined by 6-digit NAICS.

tortions, we would expect them to have lower factor costs as a share of sales.

To implement this comparison, we consider three factors of production: labor, capital, and material. We use total payroll to approach labor costs, and take the average of the value of total assets at the beginning and end of the Census year to approach capital income. Material costs are directly taken from the CMF. We project the costs-to-sales ratio on firm diversification for each Census year and plot the change of this relation over time.

$$\ln\left(\frac{\text{factor costs}_{ij}}{\text{sales}_{ij}}\right) = \text{FE}_j + \alpha \ln \text{ENI}_i + r_{ij}, \tag{11}$$

where FE_j denotes industry fixed effects, and r_{ij} denotes the regression residual. Panel (a) of Figure 12 depicts the estimated α for each year, where the shaded area depicts the robust 95% confidence interval. We find no evidence suggesting a lower costs-to-sales ratio for specialized firms: the slope estimates are not significantly different from zero throughout the whole sample period. This suggests that specialized firms' higher unit prices do not reflect higher distortions, and hence supports our interpretation.

Assuming no distortions in variable input markets, some studies in the literature suggest using the sales share of variable inputs to capture variation in markups within industries (e.g., the share of labor and material as in De Loecker and Warzynski, 2012).[20] Toward this end, we separately project the revenue share of each factor of production on diversification as in equation (11). Panel (b) of Figure 12 reports the estimated slopes. We see that more diversified firms tend to be more capital-intensive, slightly more material-intensive, and less labor-intensive. Over time, specialized firms have become increasingly more labor intensive relative to diversified firms, but the relation between capital share and diversification has remained roughly stable. Summing up labor and material, specialized firms have a lower share of variable inputs, and this relation strengthens over time.[21]

Bypassing valid concerns for the relation between factor shares and markups, our finding does not favor a higher markup for specialized firms.[22] If specialized firms charge higher

[20]See equation (7) in De Loecker and Warzynski (2012). In their framework, the log difference in markups for any two firms with the same output elasticities is given by the negative of the log difference in their revenue share of variable inputs.

[21]Specialized firms' higher labor shares could also reflect a more labor-intensive technology. This interpretation could support our view provided higher-quality production is more labor intensive, as argued in e.g. Verhoogen (2008). Moreover, the stable relation between capital intensity and diversification could also be seen as going against conjectures emphasizing changing financial conditions.

[22]See Bond, Hashemi, Kaplan, and Zoch (2021) for a discussion of the assumptions underlying this measure of markups.

Figure 12: Projection of log factor shares on log ENI × 100

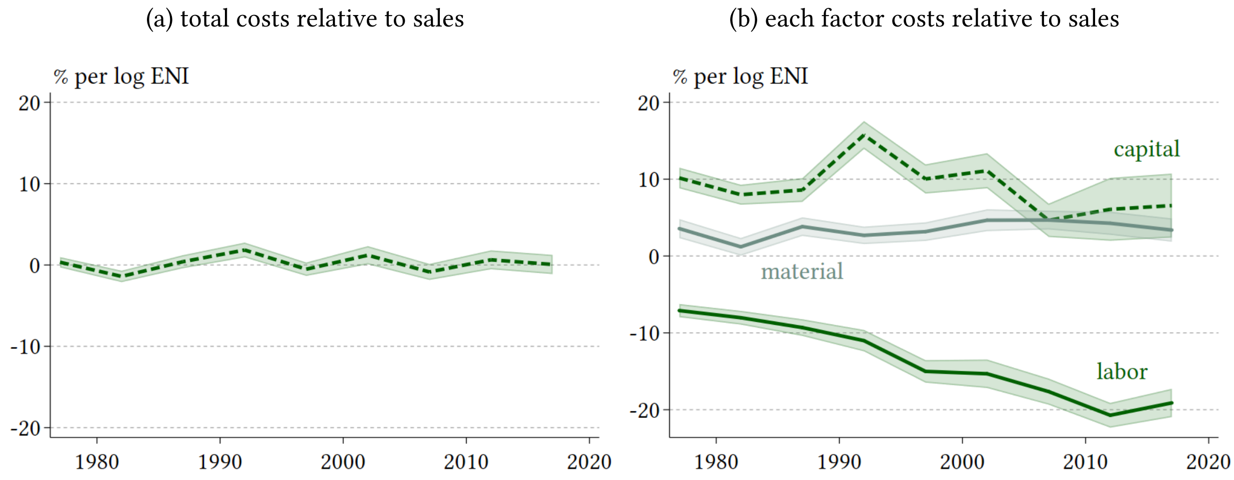

Notes: Figure reports the projection of log factor share (cost-to-sales ratio) on log ENI for each 5-year from 1977 to 2017 according to equation (11), i.e., $\alpha \times 100$. Data source: CMF.

markups, the framework used in De Loecker and Warzynski (2012) would expect specialized firms to reveal a higher share of variable inputs. This result reinforces our interpretation that specialized firms' higher unit prices reflect their higher product quality.

4.2 Specialized firms are more insulated against Chinese import competition

We study the differential responses of specialized and diversified firms' employment to changes in Chinese import competition between 1997 and 2012. This setting has three key advantages. First, Chinese exports grew rapidly over this time period, with significant variation across 6-digit NAICS industries (Autor, Dorn, and Hanson, 2016).[23] Secondly, our sample period coincides with an intense period of Chinese economic reform initiated by General Secretary Deng Xiaoping; it is thus likely that the growth of imports from China over this time period reflects increases in Chinese supply rather than increases in U.S. demand (as argued in Autor, Dorn, and Hanson, 2013).[24] Finally, and relatedly, given that the growth of Chinese exports reflects domestic reforms, it is plausible that U.S. firms did not fully anticipate which industries would be particularly affected by Chinese import competition over this time period, and thus did not position their production across industries so as to minimize their exposure. This will be important in identifying the dif-

[23] Chinese manufacturing exports as a share of global manufacturing exports grew from under 5% in 1997 to over 15% (Autor, Dorn, and Hanson, 2016).

[24] The initiation of these reforms is generally traced to Deng's "southern tour" in 1992.

ferential effect of Chinese import competition along the dimension of specialization, as detailed below.

4.2.1 Empirical design

Our empirical analyses are based on regressions of the form:

$$\Delta \ln \text{emp}_{it} = \beta_1 \Delta \text{CIE}_{it} + \beta_2 \ln \text{ENI}_{it} + \beta_3 \Delta \text{CIE}_{it} \times \ln \text{ENI}_{it} + \boldsymbol{\gamma} \boldsymbol{X}_{it} + r_{it}, \tag{12}$$

where emp_{it} is a firm's employment, ΔCIE_{it} is a measure of a firm's exposure to changes in Chinese import competition, ENI_{it} is its effective number of industries as defined in Section 2.2, $\boldsymbol{X}_{it}$ denotes a vector of controls including a constant, and r_{it} is the residual. The symbol Δ denotes forward changes over time. We measure a firm's exposure to changes in Chinese import competition as the employment-weighted average of the change in the Chinese import penetration ratio across the industries the firm operates within. Formally, for a firm i operating in J_{it} industries in period t, our exposure measure is:

$$\Delta \text{CIE}_{it} = \sum_{j \in J_{it}} \left(\frac{\text{emp}_{ijt}}{\text{emp}_{it}} \right) \Delta \text{CIP}_{jt}, \tag{13}$$

$$\text{CIP}_{jt} = \frac{\text{imports from China}_{jt}}{\text{sales}_{jt} + \text{imports}_{jt} - \text{exports}_{jt}}. \tag{14}$$

A firm thus has high exposure to changes in Chinese import competition if its employment is concentrated in industries with large changes in the Chinese import penetration ratio.[25] Our coefficient of interest is β_3: the differential response of employment growth for firms with different levels of specialization. Identification concerns in estimating this equation can be grouped into two areas: i) concerns stemming from using changes in the Chinese import penetration ratio to proxy for changes in Chinese import competition, and ii) confounders that would bias our estimate of the differential response (β_3) and/or the level response (β_2).

First, using (13) to measure exposure to import competition relies on the assumption that changes in the Chinese import penetration ratio (14) reflect changes in Chinese supply factors rather than changes in U.S. demand. As mentioned, the period of reform experienced by the Chinese economy over this time period lends plausibility to this assumption.

[25] For a perfectly specialized firm, i.e. one operating in a single industry, ΔCIE is just the change in the Chinese import penetration ratio of its industry.

As an added check, we adapt the instrument used in Autor, Dorn, and Hanson (2013) to our setting. As the instrument, we construct an alternative measure of import exposure by replacing changes in the U.S. Chinese import penetration ratio with changes in the Chinese import penetration ratio to other developed countries.[26] We thus instrument for firm exposure to changes in the U.S. Chinese import penetration using firm exposure to changes in other developed countries' Chinese import penetration ratio. Formally, our instrument for (13) is:

$$\Delta \mathrm{CIE}^{o}_{it} = \sum_{j \in J_{it}} \left(\frac{\mathrm{emp}_{ijt}}{\mathrm{emp}_{it}} \right) \Delta \mathrm{CIP}^{o}_{jt}, \tag{15}$$

$$\mathrm{CIP}^{o}_{jt} = \frac{\text{other developed countries' imports from China}_{jt}}{\text{sales}_{jt} + \text{imports}_{jt} - \text{exports}_{jt}}, \tag{16}$$

where the summation again runs over NAICS 6-digit industries.[27] The identifying assumption for this instrument is that import demand shocks at the 6-digit level are not correlated across developed countries and that there are no strong increasing, or decreasing, returns to scale in Chinese manufacturing, which could cause changes in U.S. import demand to lead to changes in Chinese exports to other countries.[28]

Secondly, given the shift-share construction of the firm-level import competition measure, the level effect (β_2) is biased if initial-period firm employment shares are spuriously correlated with the unexplained portion of firm growth.[29] For instance, if industries experiencing high growth of Chinese imports over our sample period are also those implementing labor-saving innovations, then firms more exposed to Chinese import competition would shrink employment more dramatically even in the absence of a direct effect of Chinese imports on U.S. firm growth. We will address such concerns in more detail in our robustness Section A.5.1, but we note here that for this to substantially bias our estimate of the coefficient of interest, it must be that the bias projects on firm specialization itself. Returning to the above example, firms with different levels of specialization would have to adopt the correlated labor-saving technology with systematically different intensities in order for our estimate of β_3 to be biased. This argument also holds for bias stemming from anticipation effects: though it is likely true that firms anticipated the growth of Chinese imports to some extent, there is no reason to think that such expectation bias would

[26] We use Australia, Denmark, Finland, Germany, Japan, New Zealand, Spain, and Switzerland.

[27] Note that the denominator in (16) is unchanged relative to the baseline measure; this is a necessary restriction due to data availability.

[28] This instrument has seen wide-spread usage since its introduction. See e.g.: Ding et al. (2022), Acemoglu, Autor, Dorn, Hanson, and Price (2016), Argente et al. (2021), Amiti and Heise (2021).

[29] See Goldsmith-Pinkham, Sorkin, and Swift (2020), Kirill Borusyak and Jaravel (2022) for identification in regressions with shift-share designs, Adao, Michal Kolesár, and Morales (2019) for inference.

necessarily project on specialization.[30]

4.2.2 Baseline results

We implement our empirical specifications using 5-year changes between 1997 and 2012 at an aggregated version of the 6-digit NAICS industry level. We use the 2007 vintage of NAICS, and thus crosswalk all relevant data to the 6-digit 2007 NAICS industry level. We aggregate the 2007 NAICS code slightly to ensure that no industry has zero Chinese imports by construction. Further details on data sources and sample selection can be found in Appendix A.3. Summary statistics for the regression variables are in Appendix A.3.2. Our estimation relies on variation in Chinese import exposure conditional on differing levels of diversification; we thus also examine the span of ΔCIE_{it} conditional on log ENI in Appendix A.3.2.

Table 2 shows the results of estimating equation (12) using either ordinary least squares (OLS) or two-stage least squares (2SLS) using other developed countries' changes in Chinese import penetration to instrument for Chinese supply as described above. All regressions include a year fixed effect and firm life-cycle and exporting controls for the initial period, namely a quadratic in employment, a quadratic in age, a dummy for exporter status, and exports. Additional results beyond the baseline specification are collected in section A.5.1.

We consider three specifications. The first column excludes all variables other than controls and the log of the effective number of industries. This regression thus simply restates the result shown in Section 3: specialized firms grow relative to diversified firms over our sample period. Note that we scale the dependent variable—change in log employment—by 100, so that the coefficient is approximately the percentage point difference in firm growth rates between firms with a log point difference in ENI (i.e. 2.7 times higher ENI). The coefficient thus shows that the predicted difference in employment growth rates (over a 5-year period) between two firms whose ENI differs by a log point is 15.39 percentage points (pp) in favor of the relatively specialized firm. To put this number in perspective, the interquartile range of log ENI among multi-industry firms is about 1.5 in 2017, as can be seen from Panel (b) of Figure 3.

The second column excludes all right-hand side variables other than controls and our measure of exposure to changes in Chinese import competition (13). Note that we also scale our measure of exposure by 100, so that the resulting coefficient measures the percentage point difference in firm growth rates resulting from a percentage point increase in our measure of import

[30] See Alessandria, Khan, Khederlarian, Ruhl, and Steinberg (2021) for an in-depth look at such issues.

Table 2: Responses to import competition: baseline estimation results

specification	(1)	(2)		(3)	
method	OLS	OLS	2SLS	OLS	2SLS
dependent variable	Δ ln employment $\times$ 100				
ln ENI	−15.39			−10.68	−9.31
	(0.98)			(1.23)	(1.37)
ΔCIE ×100		−0.84	−1.39	−0.79	−1.31
		(0.04)	(0.06)	(0.04)	(0.06)
ln ENI × ΔCIE × 100				-2.97	-3.85
				(0.54)	(0.66)
F statistic			22, 640		10, 920
observations			235, 000		

Notes: The sample size is truncated to the nearest thousand in accordance with Census disclosure requirement. Robust standard errors are presented in parentheses. The F statistic is the Kleibergen-Paap rk Wald F statistic for weak instruments. All regressions control for a year-fixed effect, a quadratic in age, a quadratic in employment, exports, and a dummy for whether a firm is an exporter or not. Data sources: CMF, LBD, LFTTD, and BACI.

competition, conditional on controls. The resulting estimate is negative and significant: firms more exposed to changes in Chinese import competition shrink relative to those less exposed. Column three repeats the same specification but estimated via 2SLS rather than OLS. The coefficient on our exposure measure increases by about 50%, suggesting that some of the variation in changes in Chinese imports to the U.S. reflects U.S. demand factors rather than Chinese supply factors. The magnitude of the coefficient is significant across both specifications. Focusing on the 2SLS regression as a benchmark, a 1pp increase in Chinese import exposure predicts 1.39pp less employment growth over a 5-year period. The average change in Chinese import exposure is 1.58pp, so the average effect is about 2.22pp.

Finally, columns four and five estimate our baseline specification (12), adding log effective number of industries and an interaction between log effective number of industries and the import exposure measure as independent variables. Column 5 estimates (12) using OLS and column 6 repeats the same specification using 2SLS. The resulting coefficient on the interaction is significantly negative in both regressions. Faced with a 1pp increase in Chinese import exposure, the predicted difference in employment growth rates (over a 5-year period) between two firms whose

ENI differs by one log point is 3.85pp in favor of the relatively specialized firm in the 2SLS regression. Because the average ΔCIE is about 1.58pp (see Appendix A.3.2), this implies a 6.09pp 5-year growth rate differential between two firms whose ENI differs by one log point resulting from changes in the exposure to Chinese import competition. Compared with Column 2, exposure to changes in Chinese import competition can thus account for around 40% of the average 5-year difference in employment growth between differently specialized firms.[31] Note that the 2SLS regression increases both the coefficient on changes in Chinese import exposure (by around 66%) and the coefficient on the interaction term (by around 30%), again suggesting that variation in changes in Chinese imports to the U.S. is partially driven by U.S. demand factors that are purged when instrumenting.

4.2.3 Ruling out alternative interpretations

We interpret the differential response of more diversified firms documented above to selection along the dimension of product quality: diversified firms produce lower quality goods and their output is thus more substitutable with lower quality Chinese imports than that of diversified firms (Hummels and Klenow, 2005; Hallak and Schott, 2011). We examine two alternative interpretations. First, Holmes and Stevens (2014) show that initially large plants in manufacturing shrank much more in response to Chinese import competition. They argue that this is because large plants produce standardized goods for the national market, whereas small plants tend to produce specialty goods for local markets and are thus relatively insulated from import competition. To the extent that the log effective number of industries is correlated with average plant size, our estimated interaction coefficient may thus partially reflect such variation. To address this we explicitly control for an interaction term between employment and exposure to change in Chinese competition: $\text{emp}_{it} \times \Delta\text{CIE}_{it}$. Figure A2 in Appendix A.5.1 shows that the resulting point estimate on the coefficient of interest is virtually unchanged.

Secondly, developing countries tend to export more labor-intensive goods when trading with developed countries (see e.g., Pierce and Schott, 2016); it is thus plausible that our estimated differential response is driven by systematic differences in factor shares along the dimension of specialization. To examine this, we add controls for the firm's labor income share LS_{it} and its interaction with changes in Chinese import competition, $\text{LS}_{it} \times \Delta\text{CIE}_{it}$, where the labor income shares used are de-meaned by industry prior to aggregating to the firm level. Figure A2 again shows that the resulting point estimate on the interaction term of interest is unchanged.

[31] $(6.09/15.39) \times 100 = 39.6\%$.

5 Theory

In this section, we propose a general equilibrium model motivated by our empirical findings to show that one should expect industry concentration to eventually increase in a growing economy. We define quality as a product characteristic under the control of the firm; improving quality increases the demand for a firm's product.

The key innovation in our model is that quality is a luxury: demand over different qualities is non-homothetic, so the income elasticity of demand increases with product quality. This simply captures that richer consumers demand higher quality goods (Bils and Klenow, 2001; Fajgelbaum, Grossman, and Helpman, 2011). Under this assumption, relative demand shifts toward higher-quality products when aggregate income grows. To do this, we partition the set of firms within an industry into segments that group firms with the same product quality and impose a non-homothetic demand structure over these quality segments.[32] We discipline the strength of said income effect by the estimated slope of the quality Engel curve, i.e., the relationship between consumers' total expenditure and the unit prices of products consumed, using the Consumer Expenditures Survey following Bils and Klenow (2001).

Given the demand structure, firms choose to invest in two types of technology: product quality and process efficiency. The key distinction between these two types of technology is that they differ in terms of their economies of scope, i.e., in how the average cost of technology increases with a firm's number of industries. We discipline the economies of scope for the two types of technology by the empirical relation between unit prices and firm scope from Section 4, together with the size-weighted distribution of firm scopes. Our estimated model implies economies of scope in improving process efficiency and diseconomies of scope in improving product quality. As a result, more specialized firms have a greater incentive to improve product quality, while more diversified firms have a greater incentive to improve process efficiency.

These two properties together imply that income growth reallocates production toward more specialized firms. We, in turn, investigate the quantitative relevance of our proposed mechanism.[33] We confirm that our calibrated model generates a reallocation toward specialized firms that is about one-third of our empirical findings along the U.S. growth path from 1977 to 2017.

[32]This is similar in spirit to Holmes and Stevens (2014) in that we segment firms by a characteristic of their products; however, there are two key differences. First, demand over segments in our model is non-homothetic, whereas demand over segments is homothetic in their model. Second, we allow firms to endogenously sort across said segments, whereas the distribution of firms over segments is exogenous in their paper.

[33]We focus on a numerical exercise because we allow firms to endogenously sort into quality segments, which makes our model difficult to characterize analytically.

Moreover, it predicts an upward trend in industry concentration as large as observed in U.S. manufacturing over the period 1977-2017. We conclude that one should expect to eventually observe a rising trend of industry concentration in a growing economy; other implications of concentration should be viewed relative to this secular trend.

5.1 Economic environment and notion of quality

The economy consists of a representative household and a continuum of firms that are indexed by $i \in I$. The only factor of production is labor. The household owns all firms, so that its income is equal to the sum of firm profits and labor income. There is no savings technology, and bonds are in net zero supply so that there is no borrowing and lending in equilibrium. The model can thus be analyzed period by period. Accordingly, we will drop the time subscripts for what follows.

There are two stages in the firm problem. In the initial stage, firms make technology choices: process efficiency, A_i, and product quality, q_i. In the second stage, firms produce and sell their products to the household in monopolistically competitive markets. Industries are symmetric, and thus the number of industries a firm operates, N_i, corresponds to the effective number of industries used in the empirical section of this paper.[34] For simplicity's sake, we will often refer to N_i as a firm's scope of production. We take firms' scope as given. We thus focus on modeling differing growth among differently-specialized firms; as shown in section 3.1, this channel accounts for the majority of the overall change.

5.1.1 Product demand: quality is a luxury

Within each industry, product demand is rationalized by a representative household's utility over the products of firms $\{c_i\}_{i \in I}$. Firms are partitioned into quality segments that group firms with the same product quality: $I(q) = \{i \in I | q_i = q\}$ for $q \in Q = [\underline{q}, \overline{q}]$. The distribution of quality, including the support, Q, is endogenously determined in general equilibrium by firms' technology choices. The utility function consists of two layers. The first layer is a non-homothetic constant elasticity of substitution (NH-CES) function over quality segments, implicitly defined by the equation:

$$1 = \int_Q q^{\frac{1}{\sigma}} C(q)^{\frac{\sigma-1}{\sigma}} U^{\frac{\phi(q)}{\sigma}} \, dq, \tag{17}$$

[34]This assumption accords with our findings, reported in Section 3.4, that cross-industry differences explain a negligible share of overall change in diversification.

where U denotes the utility.[35] The function $C(q)$ denotes the aggregate consumption of segment q. The second layer is a homothetic CES (H-CES) aggregator over firms within each segment:

$$C(q) = \left(\int_{I(q)} c_i^{\frac{\epsilon-1}{\epsilon}} \, di \right)^{\frac{\epsilon}{\epsilon-1}} . \tag{18}$$

In equation (17), $\sigma > 1$ is the elasticity of substitution across quality segments. For each quality segment, $\phi(q) < 0$ governs (but is not equal to) the income elasticity of demand for the segment, which can be seen more clearly from the demand function below. For the specific case where $\phi(q) = 1 - \sigma$, the function degenerates to the H-CES form: $U = [\int_0^{\overline{q}} q \, C(q)^{(\sigma-1)/\sigma} \, dq]^{\sigma/(\sigma-1)}$. In equation (18), the parameter $\epsilon \geq 1$ denotes the elasticity of substitution across firms within each quality segment. We assume $\epsilon > \sigma$ so that products in the same quality segment are more substitutable to one another than to products in other quality segments. We thus refer to firms in the same segment as direct competitors.

The household is a price taker. Given the joint distribution of quality and prices $\mathcal{P} = \{q_i, p_i\}_{i \in I}$, the household chooses a consumption bundle to maximize utility subject to the budget constraint $\int_I p_i c_i \leq E$, where E is the household's total income, equal to the sum of total profits and labor income.[36] Note that, in our model, E and $\mathcal{P}$ are endogenously determined in a general equilibrium. Let $U(E, \mathcal{P})$ be the indirect utility from the household's optimization. We characterize demand using the dual cost-minimization problem. The resulting (Hicksian) demand function for a firm with quality q_i is:

$$c_i = \underbrace{q_i \cdot P(q_i)^{\epsilon-\sigma} \cdot E^{\sigma} U(E, \mathcal{P})^{\phi(q_i)}}_{\lambda(q_i; E, \mathcal{P})} \cdot p_i^{-\epsilon}. \tag{19}$$

We refer to the function $\lambda(q; E, \mathcal{P})$ as the quality demand shifter because it controls how quality shifts up individual firms' residual demand. An important distinction of the way we introduce quality in our model relative to what is usually done in the literature is that the quality demand shifter depends on the household's income and the joint distribution of quality and prices.

The form of $\lambda(q; E, \mathcal{P})$ indicates three channels through which a firm's product quality shifts its residual demand, holding constant the joint distribution of quality and prices. The first is the direct utility gain from higher quality. This is reflected in the first term on the right-hand side of equation (19). Second, quality improvement changes the composition of a firm's direct

[35]This functional form is used by Comin, Lashkari, and Mestieri (2021) to model the reallocation of production among agriculture, manufacturing, and services over the growth path.

[36]Below we will normalize wages to 1 so that the numeraire is the real wage.

competitors, and hence determines the tightness of competition the firm faces. This channel is captured by the q-specific H-CES price index $P(q) = \int_{I(q)} ({p_i}^{1-\epsilon} di)^{1/(1-\epsilon)}$. A lower price index indicates tighter competition in the segment and, hence, a lower residual demand. For instance, if $P'(q) > 0$, firms in higher quality segments face less competition. In our model the shape of $P(q)$ is endogenously determined in a general equilibrium.

Finally, and most importantly, quality is a luxury, i.e., relative demand for higher-quality products increases with the household's income. We refer to this channel as the income effect. The income effect is captured by the term $E^{\sigma} U(E, \mathcal{P})^{\phi(q)}$ under the assumption $\phi'(q) > 0$. Formally, if $\phi'(q) > 0$, the income elasticity of demand increases with quality.

$$\frac{\partial}{\partial q} \left(\frac{\partial \ln \lambda}{\partial \ln E} \right) = \sigma \left(\frac{\partial \ln U}{\partial \ln E} \right) \phi'(q) > 0. \tag{20}$$

The proof is straightforward given that the indirect utility function, $U(E, \mathcal{P})$ is strictly increasing in income, E. In Section 5.2, we parameterize the $\phi(q)$ function by

$$\phi(q) = (1 - \sigma) q^{-\phi},$$

where $\phi > 0$ denotes the degree of nonhomotheticity over quality segments. The parameter thus governs the magnitude of the income effect on the relative demand of higher-quality products; a larger ϕ implies a larger income effect, whereas the household's demand degenerates to the homothetic case as ϕ goes to zero.

Figure 13 provides an illustrative example to summarize said three roles of product quality. The figure depicts two quality demand shifters under different levels of household income, $\lambda(q; E_h, \mathcal{P})$ and $\lambda(q; E_\ell, \mathcal{P})$, where $E_h > E_\ell$. Each demand shifter is increasing with respect to quality. This is determined by the first two channels: the direct utility effect and the competition effect. The gap between the two curves shows the income effect. Higher-quality firms gain more demand than lower-quality firms when income increases. Income growth thus shifts demand towards higher-quality products.

5.1.2 Firms' problem in the product market

A firm with process efficiency $A_i > 0$ and product quality $q_i > 0$ chooses output y_i in each of its industries subject to the demand function defined in equation (19). Firms' production functions are linear with labor as the only factor of production and do not differ across different levels of

Figure 13: Example for how quality demand shifter changes with total expenditure

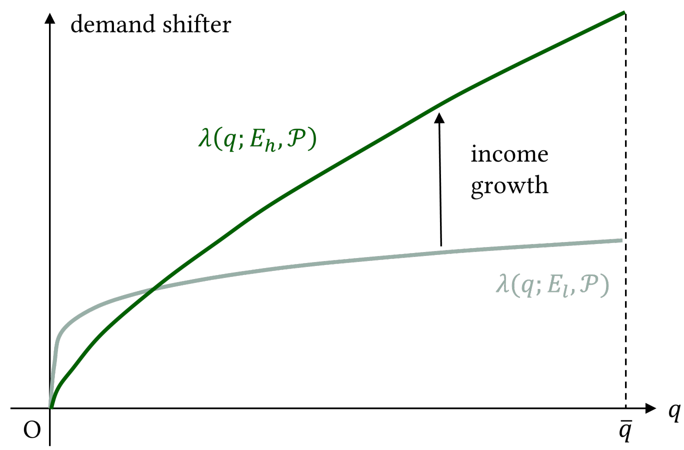

Notes: This figure shows an example of how the shape of the quality-specific demand shifter changes with aggregate expenditure. The graph shows the quality-specific demand shifter for the two economies. Subscript ℓ (h) refers to the low-expenditure (high-expenditure) economy.

quality. We normalize the wage rate to one. Firm production is independent across industries, and thus the profit maximization problem within an industry is standard:

$$\max_{\{p_i, y_i\}} p_i y_i - \frac{y_i}{A_i} \quad \text{s.t.} \quad y_i = \lambda(q_i; E, \mathcal{P})\, p_i^{-\epsilon}. \tag{21}$$

Higher process efficiency enables the firm to produce with lower marginal cost. At the optimum, firms charge $p_i = \left(\frac{\epsilon}{\epsilon-1}\right)\frac{1}{A_i}$, a constant markup over their marginal cost. Optimal firm revenue is $p_i y_i = \left(\frac{\epsilon}{\epsilon-1}\right)^{1-\epsilon} A_i^{\epsilon-1}\lambda(q_i; E, \mathcal{P})$, and the resulting indirect profit function is:

$$\pi(A_i, q_i) = \left(\frac{(\epsilon-1)^{\epsilon-1}}{\epsilon^{\epsilon}}\right) \cdot A_i^{\epsilon-1}\, \lambda(q_i; E, \mathcal{P}). \tag{22}$$

Because we assume that industries are symmetric, the total profits for a multi-industry firm with N_i industries are simply $N_i\, \pi(A_i, q_i)$, where q_i now denotes firm i's average quality per industry, and A_i the average process efficiency. Given this symmetry, N_i is the theoretical analog to the effective number of industries (ENI) used in our empirical sections.

5.1.3 (Dis)economies of scope in technology choices

Prior to production, firms face a tradeoff between process efficiency (A_i) and product quality (q_i). Quality improvement makes use of the same scarce resources used in improving process efficiency. Formally, we assume that firms maximize profits, $N_i\, \pi(A_i, q_i)$, subject to a feasible set of (A_i, q_i) described by the following technology frontier:

$$\left[\left(\frac{N_i^{\alpha_A}}{\gamma_{A,i}}\right) A_i^{\omega} + \left(\frac{N_i^{\alpha_q}}{\gamma_{q,i}}\right) q_i^{\omega}\right]^{\frac{1}{\omega}} \leq \Gamma. \tag{23}$$

This form was initially proposed by Caselli and Coleman (2006). The shape parameter $\omega > 1$ so that each firm's technology frontier is concave. We impose that ω is large enough so that the firm's maximization problem is well-defined.

The parameter Γ denotes the bound on the aggregate technology frontier. We think of it as capturing the fundamental state of knowledge in the economy. At a given point in time, higher quality products can only be produced with lower process efficiency; though there is heterogeneity in this tradeoff across firms, Γ summarizes those tradeoffs that are common to all firms. More precisely, changes in Γ shift all firms' technology frontiers uniformly; accordingly, we will treat it as the source of growth in this economy. In the next subsection, we will let Γ grow at a rate such that the equilibrium consumption in the model grows at the same rate as U.S. real consumption per capita over the period 1977-2017.

We view $N_i^{\alpha_A}/\gamma_{A,i}$ and $N_i^{\alpha_q}/\gamma_{q,i}$ as capturing the average productivity of technology improvement across industries, where the parameters γ_{Ai} and γ_{qi} refer to firm-specific components of the technology productivity. The parameters α_A and α_q capture economies of scope, which determine how a firm's average productivity of technology improvement changes with the number of industries it operates in.

If a technology's scope parameter is less than 0, then there are economies of scope in improving that technology: firms operating in multiple industries face lower rates of substitution in improving that technology than single-industry firms. This case can stem from the existence of some "quasi-public" input at the level of the firm, i.e. an input that can be shared across industries without complete congestion (as introduced in Panzar and Willig, 1981). For instance, if there is considerable overlap between the production processes used in different industries, then the non-rivalry of ideas creates an economy of scope in technology improvement: an idea that leads to better process efficiency in one industry can be applied to another without further modification, and thus increases the average productivity across industries.

If instead a technology's scope parameter is greater than 0, then there are diseconomies of scope in improving that technology: firms operating in multiple industries face higher rates of substitution in improving that technology than single-industry firms. This case can occur if there are returns to specialization at the firm level. For instance, if firms have core competencies in technology improvement, i.e. they are not equally productive at improving technology in all industries, then the average productivity of technology improvement is naturally decreasing in a firm's scope. Another possible source of returns to specialization is a per-industry fixed cost incurred by technology improvement which is increasing in a firm's scope (e.g., costs of coordination as in Becker and Murphy, 1992). Both sources of diseconomies of scope are analogous to traditional arguments for returns to the division of labor advanced in the labor literature, e.g., Rosen (1974) and Borland and Yang (1992).

(Dis)economies of scope in technology choices create a connection between a firm's production scope and its average product quality. If there are diseconomies of scope in quality improvement ($\alpha_q > 0$), firms with narrower production scope—more specialized firms—will endogenously choose to produce higher quality goods. The intuition is straightforward: under diseconomies of scope in quality improvement, more specialized firms face lower rates of substitution in improving product quality relative to more diversified firms.

In what follows, we will not impose any restrictions on α_A, α_q. Instead, we will use the empirical relationship between unit price, firm size, and scope of production to jointly discipline these parameters. As we will show in Section 5.2.1, the resulting scope parameters imply economies of scope in process efficiency improvement ($\alpha_A < 0$) and diseconomies of scope in product quality improvement ($\alpha_q > 0$).

5.1.4 Definition of general equilibrium

Overall, there are three fundamental sources of firm heterogeneity in this model: γ_{Ai}, γ_{qi}, and N_i. Note that a key component of the firm problem is the quality demand shifter, $\lambda(q; E, \mathcal{P})$. Therefore, each firm's optimal choices are determined not only by its individual characteristics, but also by the joint distribution of quality and prices, $\mathcal{P}$, which is endogenously determined in general equilibrium; we close this subsection by defining a general equilibrium.

Given a technology frontier, Γ, and a distribution of firm characteristics, $\{\gamma_{Ai}, \gamma_{qi}, N_i\}_{i\in I}$, a general equilibrium is an allocation $\{c_i^*, y_i^*\}_{i\in I}$, technology $\{A_i^*, q_i^*\}_{i\in I}$, prices $\{p_i^*\}_{i\in I}$, and income

E^* such that:

(a) given E^* and $\mathcal{P}^* = \{p_i^*, q_i^*\}_{i\in I}$, the bundle $\{c_i^*\}_{i\in I}$ solves the household's optimization problem;

(b) given E^*, $\mathcal{P}^*$, A_i^*, and q_i^*, each firm's p_i^* and y_i^* solves its product market problem with;

(c) given E^* and $\mathcal{P}^*$, each firm's A_i^* and q_i^* solves its technology problem;

(d) markets clear: $c_i^* = y_i^*$ for each $i \in I$, and $E^* = \int_I (\pi_i^* + y_i^*/A_i^*) di$.

5.2 Quantifying the impacts of economic growth

A key implication of our model is that demand shifts toward more specialized firms along a growth path because they tend to produce higher-quality products. Is the proposed mechanism economically significant? Specifically, can growth generate a reallocation toward specialized firms that is comparable to what we observe in the data? Does growth result in rising industry concentration in our model?

This subsection answers these questions in the following three steps. First, we calibrate the model so that it replicates relevant empirical patterns. We target moments that discipline (i) the income effect on the relative demand for higher-quality products and (ii) the economies of scope in technology.

Second, we show that, in the calibrated model, growth can generate a rise of specialized firms that is one-third of what we observe in the data. To do so, we let the frontier of technology (Γ) grow so that the model-generated path of real consumption per capita replicates that of the U.S. economy between 1977 and 2017, while we fix all other parameters at their calibrated values. Finally, we show that our model predicts an upward trend in industry concentration over this time period. Note that we do not use these trends as targeted moments in calibration.

5.2.1 Calibration

We discipline the income effect by the slope of the quality Engel curve, which traces out the price of consumed products against consumers' total expenditure. In the model, we define the slope of the quality Engel curve as the income elasticity of industry price index in a partial equilibrium, $\partial \ln \overline{P} / \partial \ln E$, where

$$\overline{P} \equiv \frac{E}{U} = \int_0^{\overline{q}} U(E, \mathcal{P})^{(\sigma-1)q^{-\phi}-1} E^{\sigma} q P(q)^{1-\sigma} dq. \tag{24}$$

See Appendix B.3 for derivation. Holding constant $P(q)$, this equation shows that higher income shifts the relative demand towards higher-quality products, as is reflected by the $U^{(\sigma-1)q^{-\phi}-1}$ term in the integral. Recall that $P(q)$ is the price index for quality segment q, and is increasing in quality, $P'(q) > 0$. Income growth hence drives up the price index, $\overline{P}$, by reallocating demand towards segments with higher price indices, $P(q)$. This equation also tells us that the quality Engel curve is specifically informative in pinning down two parameters: the income elasticity, ϕ, and the elasticity of substitution between quality segments, σ. Appendix B.3 describes how we compute the slope of the quality Engel curve in our quantitative model.

To pin down this moment, we use the 1990-2021 U.S. Consumer Expenditure Survey (CEX). We estimate the product-specific elasticity of unit prices paid by consumers with respect to their total expenditure following Bils and Klenow (2001)'s empirical approach. To map to the model quality Engel curve, we then crosswalk the CEX products to 4-digit NAICS and take the average slope over industries weighted by employment (see Appendix A.6 for details).

The other key component of our model is the economies of scope in technology choices, governed by the parameters of the technology frontier, α_A, α_q, and ω, as specified in equation (23). Ideally, we would discipline them by how firms' process efficiency (A) and product quality (q) project on their scope of production. But we cannot observe technology. Instead, we discipline the economies of scope by how firms' size and unit price relate to their scope, as measured by the effective number of industries (ENI). Better process efficiency increases firm size via lower prices with magnitude controlled by ϵ, the price elasticity of demand. Product quality improvement, on the other hand, raises firm size by shifting up firm-level residual demand with magnitude governed by the consumer's preferences as summarized by $\lambda(q)$, the quality demand shifter. Therefore, the empirical relation of firms' scope with their size and unit price, in conjunction with the quality Engel curve, jointly discipline the technology frontier and the demand elasticity. Specifically, we target the projection of log unit prices on log ENI reported in Column 2 of Table 1 and the 1977 employment-weighted distribution of log ENI reported in Figure 2 and Panel (c) of Figure 3.

To sum up, we pick parameters, $(\phi, \sigma, \alpha_q, \alpha_A, \omega, \epsilon)$, jointly that replicates (i) the slope of the quality Engel curve, (ii) the empirical relation between unit price and ENI, and (iii) the 1977 employment-weighted distribution of ENI in U.S. manufacturing, while normalizing the 1977 technology frontier, $\Gamma = 1$. As reported in Table 3 and Figure 14 our model closely matches the targeted moments. Note that we not only replicate the average ENI in 1977 but the entire distribution as depicted in Figure 14. We report other computational details in the appendices:

Table 3: Targeted moments: data vs. model (part 1)

moment	data	model
slope of the quality Engel curve	0.66	0.67
projection of log price on log ENI ×100	-3.65	-3.68
average ENI (employment weighted) in 1977	4.0	3.98
employment share of single-industry firms in 1977	34.0	34.36

Notes: The quality Engel curve is estimated using the U.S. Consumer Expenditure Survey (CEX). Details are reported in Appendix A.6. Other moments are calculated using the CMF. For details, see Section 4.1 for the projection of log price on log ENI and Section 3.1 for the employment-weighted distribution of log ENI.

Figure 14: Targeted moments: data vs. model (part 2)

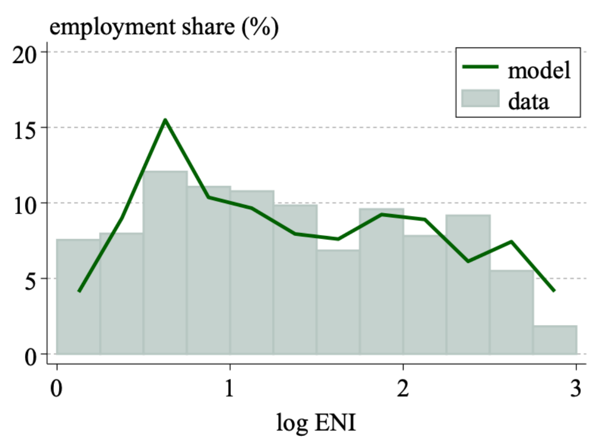

Notes: Figure shows employment share among multi-industry firms in 1977 for each log ENI group. Firms are grouped by their log ENI. The grouping cutoffs are 0.25, 0.5, ..., and 3, where each group collects firms whose log ENI is weakly smaller than the cutoff and is strictly greater than the previous cutoff. The data distribution is identical to that reported in Panel (c) of Figure 3. Data source: CMF.

Appendix B.1 describes how we parameterize the distributions of γ_A and γ_q; Appendix B.2 describes the computational algorithm that solves a general equilibrium.

The calibrated parameters are reported in Table 4. We highlight two sets of parameters. First, the firm-level price elasticity of demand, $\epsilon = 4.6$, is consistent with those estimated by Hottman, Redding, and Weinstein (2016) using barcode data (see their Table V). Second, our calibration implies diseconomies of scope in product quality improvement, $\alpha_q > 0$, and economies of scope in process efficiency improvement, $\alpha_A < 0$. As discussed above, these parameters are mainly disciplined by (i) specialized firms' higher unit prices and (ii) the fact that diversified firms are relatively large in 1977. The implications of the calibrated technology parameters are shown in Figure 15. The left panel shows that, in the calibrated model, more specialized firms operate in higher-quality segments, whereas the right panel shows that diversified firms operate with higher process efficiency. As relative demand shifts toward these segments, more specialized firms will thus grow disproportionately.

5.2.2 Reallocation toward specialized firms

We now examine the general equilibrium effect of income growth. To do so, we take annual U.S. real consumption per capita from 1977 to 2017 and normalize the data by the 1977 value. We set a growth path of Γ such that the growth path of real consumption in the model replicates that in the data. Note that we hold all other parameters constant so that technological progress through Γ is the only exogenous factor changing over time. We want to highlight that no other trends are used as targeted moments.

Panel (a) of Figure 16 compares the model consumption path with its data counterpart, where the vertical axis indicates the factor change from 1977. Real consumption has doubled over our sample period. We then compare the model-generated reallocation towards specialized firms with its data counterpart. We first look at the changes in aggregate firm scope over the period 1977-2017 that are accounted for by the rise of specialized firms. In the data, this is captured by the covariance between log ENI and changes in employment share among continuing firms, as specified in equation (4). In our model, this is simply captured by the employment-weighted average log ENI because reallocation among surviving firms is the only source of changes in aggregate specialization.[37]

[37] This is because (i) there is no entry and exit, and (ii) we fixed the unweighted distribution of log ENI over time. In this case, combining equation (3) and (4) results in $\Delta n = \text{cov}_C[\Delta \tilde{x}_i, (n_{i,1} + n_{i,2})/2]$.

Table 4: Calibrated model parameters

notation	description	value
ϕ	degree of nonhomotheticity over quality segments	1.10
σ	elasticity of substitution among quality segments	1.20
α_q	scope parameter for improving product quality	1.27
α_A	scope parameter for improving process efficiency	-1.67
ω	shape parameter of the technology frontier	6.00
ϵ	elasticity of substitution among firms	4.60

Notes: Parameters are calibrated to match the moments reported in Table 3 and Figure 14

Figure 15: Firm technology choices by scope in 1977 calibrated model

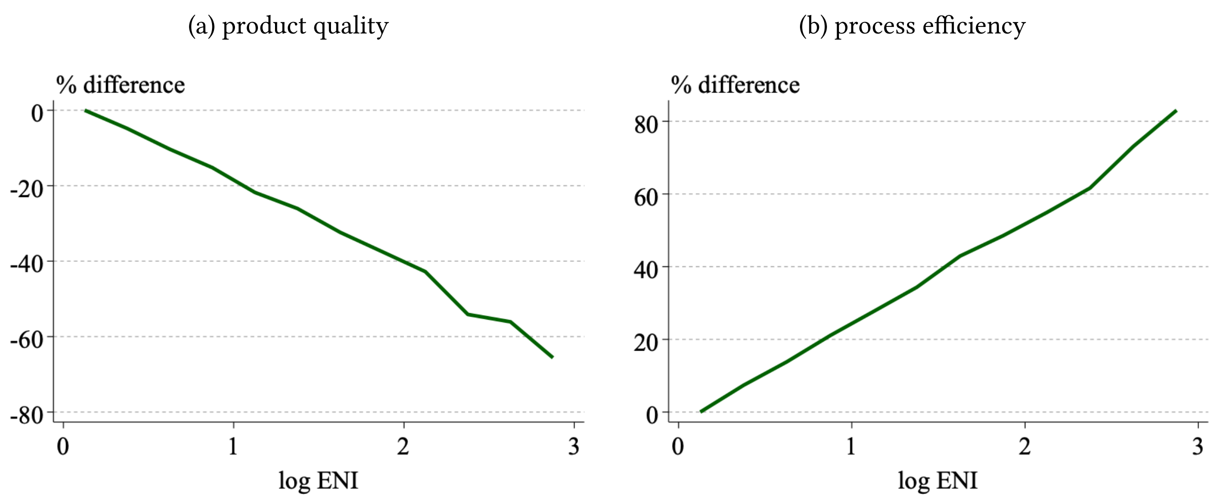

Notes: Panel (a) reports the log employment-weighted average product quality (q_i) by log ENI groups in the 1977 calibrated model. Panel (b) reports the log employment-weighted average process efficiency (A_i) by log ENI groups. Firms are grouped by their log ENI. The grouping cutoffs are 0.25, 0.5, ..., and 3, where each group collects firms whose log ENI is weakly smaller than the cutoff and is strictly greater than the previous cutoff.

Figure 16: Income growth and reallocation toward specialized firms, model vs. data

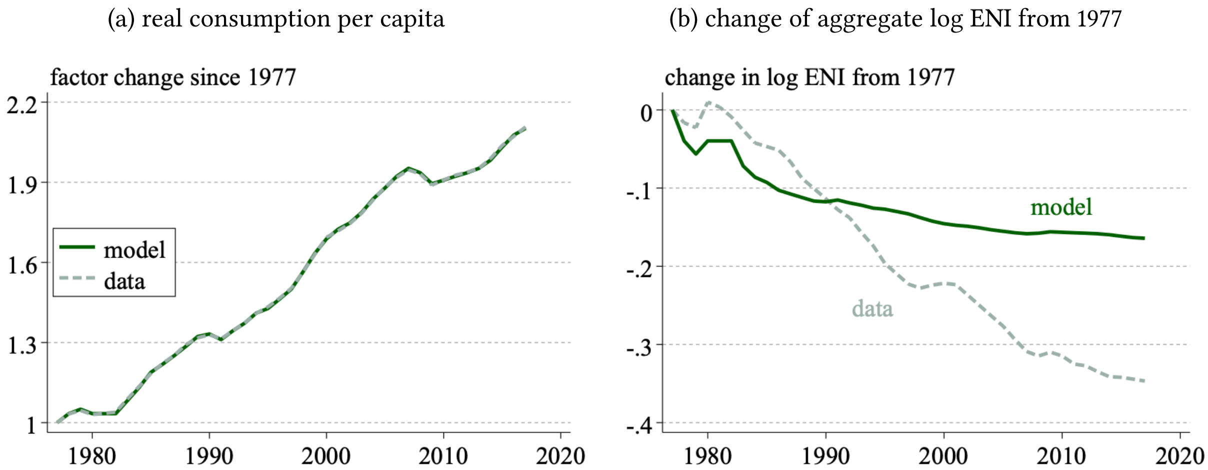

Notes: Data source for Panel (a) is U.S. Bureau of Economic Analysis, real personal consumption expenditures per capita, retrieved from FRED: https://fred.stlouisfed.org/series/A794RX0Q048SBEA. We choose the sum of quarterly value as the annual value. Data source for Panel (b) is the LBD. We plot the changes in employment-weighted average log ENI from 1977 that can be accounted for by the faster growth of more specialized firms, i.e., $\text{cov}_C[\Delta\tilde{x}_i, (n_{i,1}+n_{i,2})/2]$, as defined in equation (4).

Panel (b) of Figure 16 plots the results, where the dark solid line represents the model, and the light dashed line represents the data. Our model predicts that, over the past 40 years, the cumulative growth of income leads to a 0.16 log point decline in employment-weighted average log ENI, which is about half of the decline in the data.

Our model also generates an overall shift in the employment-weighted distribution of log ENI between 1977 and 2017. In accordance with our empirical section, we split this comparison into two, comparing first the employment share of single-industry firms and secondly the employment-weighted distribution of log ENI among multi-industry firms. Figure 17 shows the employment share of single-industry firms in 1977 and 2017; the light bars depict the data values and the dark bars depict the model values. Figure 18 shows the employment-weighted distribution of log ENI among multi-industry firms for 1977 and 2017; Panel (a) depicts the data values, and Panel (b) depicts the model values. The two figures confirm that the decline in aggregate scope in our model is due to a reallocation towards more specialized firms.

Figure 17: Employment share of single industry firms, model vs. data

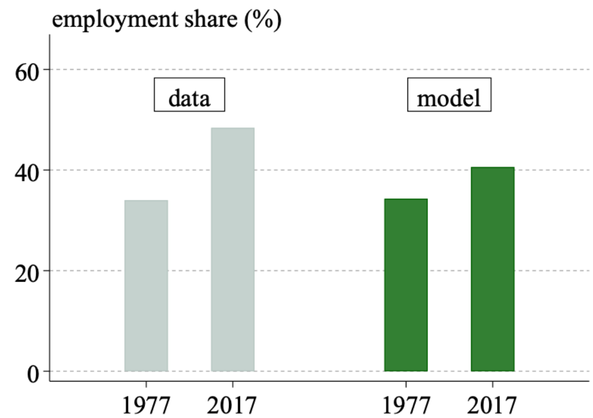

Notes: The data values are identical to those reported in Figure 2. Data source: CMF.

Figure 18: Employment-weighted distribution of log ENI among multi-industry, model vs. data

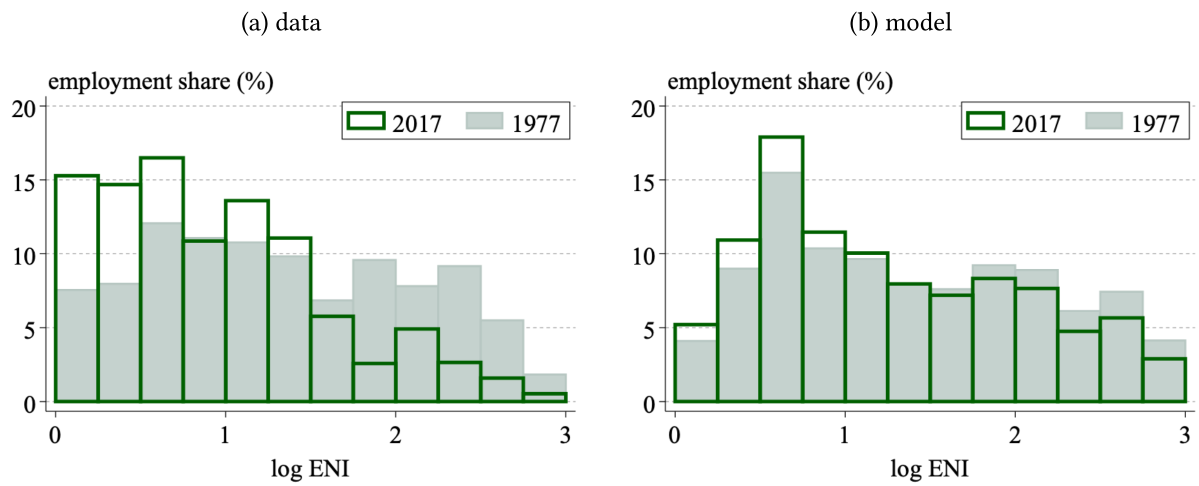

Notes: Firms are grouped by their log ENI. The grouping cutoffs are 0.25, 0.5, ..., and 3, where each group collects firms whose log ENI is weakly smaller than the cutoff and is strictly greater than the previous cutoff. The data distributions are identical to those reported in Panel (c) of Figure 3. Data source: CMF.

5.2.3 Industry concentration increases in a growing economy

We now compare the model-generated industry concentration with its data counterpart. To do so, we look at the changes in industry concentration that are accounted for by the rise of specialized firms over the period 1977-2017. In the data, this is captured by the differences between the observed concentration ratio and its counterfactual value absent the rise of specialized firms, $\mathrm{CR}_t - \mathrm{CR}_t^{\mathrm{cf}}$, as defined in Section 3.5. In the model, this is simply captured by the change in concentration ratio because the only source of changes is the growth-generated reallocation toward specialized firms.

Following our empirical analysis, we look at the top-4 and top-20 concentration ratios (CR4 and CR20) in U.S. manufacturing. These measures do not have direct counterparts in our model because there is a continuum of firms. Instead, we pick the percentile of the firm size distribution such that the share of total sales accounted for by these firms matches the data concentration ratio in 1977. Specifically, we use the top 2.85% largest firms' sales share as the model counterpart for CR4 and the top 3.95% largest firms' sales share as the model counterpart for CR20. Figure 19 plots the results. In the data, the rise of specialized firms can account for around a 9pp increase in CR4 and an 8pp increase in CR20. In our model, income growth can generate a similar amount of changes in CR4, while it only generates half of the changes in CR20.

The takeaway from this subsection is that our proposed mechanism is not only qualitatively but also quantitatively relevant to explaining the rise of specialized firms and the rising industry concentration. Restating, because quality is a luxury and there are diseconomies of scope in quality improvement, an income growth path replicating the observed growth in the U.S. can lead to a reallocation toward more specialized firms, which in turn drives up industry concentration. Most importantly, their magnitude is economically significant compared with what we observed in the data.

6 Conclusion

This article advances an alternative view of the much-discussed shift in activity toward large firms observed in the average U.S. industry over the past 50 years. Rather than reflecting structural changes in the nature of competition, we argue that this development reflects shifts in demand toward higher-quality products that inevitably occur as aggregate income grows. Because industry concentration naturally changes on the growth path, other inferences on this measure

Figure 19: Changes in industry concentration from 1977: model vs. data

(a) Top 4 concentration ratio

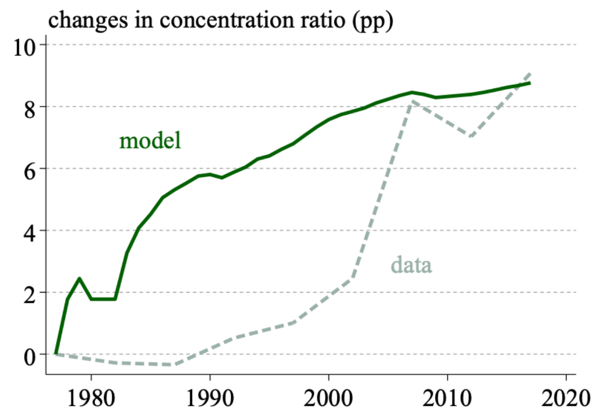

(b) Top 20 concentration ratio

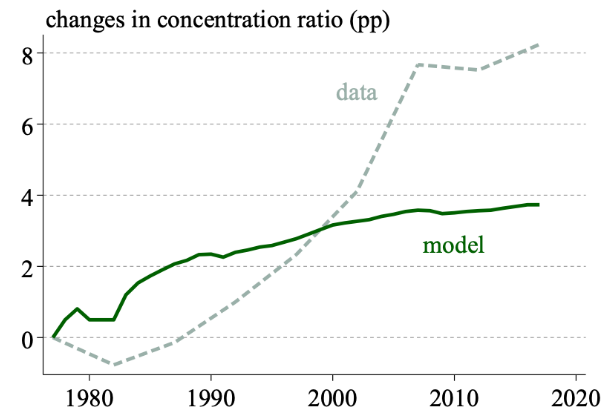

Notes: The data lines report the gap between the observed concentration ratio and its counterfactual value absent the rise of specialized firms, $\mathrm{CR}_t - \mathrm{CR}_t^{\mathrm{cf}}$, as defined in Section 3.5. Data source: CMF.

should be made relative to this secular trend.

Our interpretation is based on the observation that there has been a marked shift in activity toward more specialized firms that coincides with the rise in the revenue share of large firms within U.S. industries; this shift can be reconciled with growing relative demand for quality as long as more specialized firms tend to produce higher-quality products.

To support this view, we first study firms' unit prices within narrowly-defined product categories. We find that conditional on firm size, more specialized firms tend to charge higher unit prices, while they do not seem to face higher distortions in the input or product markets. Secondly, we study firms' responses to the rapid rise in Chinese import competition beginning in the early 1990s. The results of this analysis are stark: more diversified firms shrink dramatically relative to specialized firms when faced with the same exposure to Chinese import competition. Noting that developing countries export low-quality products to developed countries, we interpret this finding as favoring our posited link between specialization and product quality.

We construct a model to show that the average industry revenue share of large firms is naturally changing over a growth path. Because quality is a luxury, relative demand increasingly favors more specialized firms as income grows. We can thus generate an increasing path of average industry concentration which is not driven by structural changes in the nature of competition, but instead simply reflects income growth.

There is much fruitful future work that could build on ours. We have focused on one way in which firms producing high-quality products may differ from those producing low-quality products, namely industry scope, but there are others. Factor intensities, input complexity, usage of skilled labor, etc., may naturally change along a growth path. Secondly, our model raises questions about how the scope of production shapes the incentives for innovation. Though we have abstracted from dynamic innovation choices in order to focus on the cross-sectional implications, an extension incorporating such decisions could be used to better explore the sources of growth. Particularly interesting may be revisiting the externalities inherent to different types of innovation in our setting. Finally, our broader point, that aggregate income growth itself may reallocate activity between different types of firms, could surely be productively applied elsewhere.

References

Acemoglu, Daron, David Autor, David Dorn, Gordon H. Hanson, and Brendan Price (2016). "Import Competition and the Great US Employment Sag of the 2000s." *Journal of Labor Economics* 34 (S1):S141–S197.

Adao, Rodrigo, Michal Michal Kolesár, and Eduardo Morales (2019). "Shift-Share Designs: Theory and Inference." *Quarterly Journal of Economics* 134 (4):1949–2010.

Aghion, Philippe, Antonin Bergeaud, Timo Boppart, Peter J Klenow, and Huiyu Li (2023). "A Theory of Falling Growth and Rising Rents." *Review of Economic Studies* rdad016.

Aguiar, Mark A, Mark Bils, and Corina Boar (2020). "Who are the Hand-to-Mouth?" NBER Working Paper 26643.

Akcigit, Ufuk and Sina T. Ates (2021). "Ten Facts on Declining Business Dynamism and Lessons from Endogenous Growth Theory." *American Economic Journal: Macroeconomics* 13 (1):257–298.

Alessandria, George A., Shafaat Y. Khan, Armen Khederlarian, Kim J. Ruhl, and Joseph B. Steinberg (2021). "Trade-Policy Dynamics: Evidence from 60 Years of US-China Trade." NBER Working Paper 29122.

Amiti, Mary and Sebastian Heise (2021). "U.S. Market Concentration and Import Competition." CEPR Discussion Paper 16126.

Argente, David A., Sara Moreira, Ezra Oberfield, and Venky Venkateswaran (2021). "Scalable Expertise." Working Paper.

Autor, David, David Dorn, and Gordon H. Hanson (2016). "The China Shock: Learning from Labor-Market Adjustment to Large Changes in Trade." *Annual Review of Economics* 8:205–240.

Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen (2017). "Concentrating on the Fall of the Labor Share." *American Economic Review P&P* 107 (5):180–185.

Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen (2020). "The Fall of the Labor Share and the Rise of Superstar Firms." *Quarterly Journal of Economics* 135 (2):645–709.

Autor, David H., David Dorn, and Gordon H. Hanson (2013). "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review* 103 (6):2121–2168.

Barkai, Simcha (2020). "Declining Labor and Capital Shares." *Journal of Finance* 75 (5):2421–2463.

Becker, Gary S and Kevin M Murphy (1992). "The Division of Labor, Coordination Costs, and Knowledge." *Quarterly Journal of Economics* 107 (4):1137–1160.

Bernard, Andrew B., Stephen J. Redding, and Peter K. Schott (2011). "Multiproduct Firms and Trade Liberalization." *Quarterly Journal of Economics* 126 (3):1271–1318.

Bils, Mark and Peter J Klenow (2001). "Quantifying Quality Growth." *American Economic Review* 91 (4):1006–1030.

Bond, Steve, Arshia Hashemi, Greg Kaplan, and Piotr Zoch (2021). "Some Unpleasant Markup Arithmetic: Production Function Elasticities and Their Estimation from Production Data." *Journal of Monetary Economics* 121:1–14.

Borland, Jeff and Xiaokai Yang (1992). "Specialization and a New Approach to Economic Organization and Growth." *American Economic Review* 82 (2):386–391.

Caselli, Francesco and Wilbur John Coleman II (2006). "The World Technology Frontier." *American Economic Review* 96 (3):499–522.

Chow, Melissa, Teresa C. Fort, Christopher Goetz, Nathan Goldschlag, James Lawrence, Elisabeth Ruth Perlman, Martha Stinson, and T. Kirk White (2021). "Redesigning the Longitudinal Business Database." CES Working Paper 21-08.

Comin, Diego, Danial Lashkari, and Martí Mestieri (2021). "Structural Change with Long-Run Income and Price Effects." *Econometrica* 89 (1):311–374.

De Loecker, Jan and Frederic Warzynski (2012). "Markups and Firm-Level Export Status." *American Economic Review* 102 (6):2437–71.

Decker, Ryan A., John Haltiwanger, Ron S. Jarmin, and Javier Miranda (2016). "Where has All the Skewness Gone? The Decline in High-Growth (Young) Firms in the U.S." *European Economic Review* 86:4–23.

Ding, Xian, Teresa C. Fort, Stephen J. Redding, and Peter K. Schott (2022). "Structural Change Within versus Across Firms." NBER Working Paper 30127.

Fajgelbaum, Pablo, Gene M Grossman, and Elhanan Helpman (2011). "Income Distribution, Product Quality, and International Trade." *Journal of Political Economy* 119 (4):721–765.

Fort, Teresa C. and Shawn D. Klimek (2018). "The Effects of Industry Classification Changes on US Employment Composition." CES Working Paper 18-28.

Furman, Jason and Peter Orszag (2018). "Slower Productivity and Higher Inequality: Are They Related?" PIIE Working Paper 18-4.

Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift (2020). "Bartik Instruments: What, When, Why, and How." *American Economic Review* 110 (8):2586–2624.

Gutierrez, German and Thomas Philippon (2017). "Declining Competition and Investment in the U.S." NBER Working Paper 23583.

Hallak, Juan Carlos and Peter K. Schott (2011). "Estimating Cross-Country Differences in Product Quality." *Quarterly Journal of Economics* 126 (1):417–474.

Holmes, Thomas J. and John J. Stevens (2014). "An Alternative Theory of the Plant Size Distribution, with Geography and Intra- and International Trade." *Journal of Political Economy* 122 (2):369–421.

Hottman, Colin J., Stephen J. Redding, and David E. Weinstein (2016). "Quantifying the Sources of Firm Heterogeneity." *Quarterly Journal of Economics* 131 (3):1291–1364.

Hsieh, Chang-Tai and Peter J. Klenow (2009). "Misallocation and Manufacturing TFP in China and India." *Quarterly Journal of Economics* 124 (4):1403–1448.

Hsieh, Chang-Tai and Esteban Rossi-Hansberg (2019). "The Industrial Revolution in Services." NBER Working Paper 25968.

Hummels, David and Peter J. Klenow (2005). "The Variety and Quality of a Nation's Exports." *American Economic Review* 95 (3):704–723.

Kehrig, Matthias and Nicolas Vincent (2021). "The Micro-Level Anatomy of the Labor Share Decline." *Quarterly Journal of Economics* 136 (2):1031–1087.

Kirill Borusyak, and Peter Hull and Xavier Jaravel (2022). "Quasi-Experimental Shift-Share Research Designs." *Review of Economic Studies* 89 (1):181–213.

Ma, Yueyuan (2022). "Specialization in a Knowledge Economy." Working Paper.

Melitz, Marc J. and Sašo Polanec (2015). "Dynamic Olley-Pakes Productivity Decomposition with Entry and Exit." *RAND Journal of Economics* 46 (2):362–375.

Olley, G. Steven and Ariel Pakes (1996). "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica* 64 (6).

Olmstead-Rumsey, Jane (2022). "Market Concentration and the Productivity Slowdown." Working Paper.

Panzar, John C and Robert D Willig (1981). "Economies of Scope." *American Economic Review* 71 (2):268–272.

Pierce, Justin R. and Peter K. Schott (2012). "A Concordance Between U.S. Harmonized System Codes and SIC/NAICS Product Classes and Industries." *Journal of Economic and Social Measurement* 37 (1-2):61–96.

Pierce, Justin R. and Peter K. Schott (2016). "The Surprisingly Swift Decline of US Manufacturing Employment." *American Economic Review* 106 (7):1632–1662.

Rosen, Sherwin (1974). "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition." *Journal of Political Economy* 82 (1):34–55.

Rossi-Hansberg, Esteban, Pierre-Daniel Sarte, and Nicholas Trachter (2020). "Diverging Trends in National and Local Concentration." *NBER Macroeconomics Annual* 35.

Schott, Peter K. (2008). "The Relative Sophistication of Chinese Exports." *Economic policy* 23 (53):6–49.

Sui, Xiamei (2022). "Uneven Firm Growth in a Globalized World." Working Paper.

Verhoogen, Eric A (2008). "Trade, Quality Upgrading, and Wage Inequality in the Mexican Manufacturing Sector." *Quarterly Journal of Economics* 123 (2):489–530.

White, T. Kirk, Jerome P. Reiter, and Amil Petrin (2018). "Imputation in U.S. Manufacturing Data and Its Implications for Productivity Dispersion." *Review of Economics and Statistics* 100 (3):502–509.

Appendix

A Data appendix

A.1 The LBD and CMF

Here we explain the construction of the benchmark LBD and CMF samples used in Section 2 and 3. The process aims to keep all active manufacturing establishments with well-defined firm identifiers and harmonized industry codes. Operationally, we consider an establishment active only if it reports positive production, i.e., employment, payroll, and sales.

To construct the benchmark LBD sample, we use the following variables (with variable names in parentheses): employment (emp), payroll (pay), firm identifier (lbdfid), vintage-consistent NAICS codes (bds_vcnaics), and active status flag (act).[A1] We then select samples in two steps. First, we drop establishments whose payroll or employment are missing or non-positive, whose firm identifier or NAICS codes are missing, or whose active status flag is "N." This step is to keep active establishments whose industry and firm are well-defined. Second, we then drop establishments whose payroll per employment is less than 1 USD, and we drop firms whose employment is more than 1 million. This step is to drop outliers with unreasonable values. We call this establishment-level sample the *full LBD* sample.

Our benchmark sample only focuses on manufacturing. We thus take the full LBD sample and keep the establishments whose first two digits of NAICS are 31, 32, or 33. We then aggregate data to the firm-by-industry level by summing up the employment of establishments with the same firm identifier and NAICS.

In Section 3.7.3, we consider another five LBD samples beyond the benchmark manufacturing sample. All of them are based on the full LBD sample. The only difference is the selection before aggregating data to the firm-by-industry level. For the extended manufacturing sample, we keep all establishments whose firm has at least one active manufacturing establishment, regardless of their own NAICS. We keep establishments whose first two digits of NAICS are 44 or 45 for the retail sector; 42 for wholesale; 53, 56, 71, 72, or 81 for non-tradable services; and 51, 52, 54, or 55, for tradable services.

[A1] Because we use the 2018 version of the LBD, the bds_vcnaics variable assigns a 2012 NAICS code to each establishment for all LBD years. See Chow et al. (2021) for more details.

We apply a similar process to construct the benchmark CMF sample. The main variables used for sample selection are: sales (total value of shipments, tvs), employment (te), payroll (pay), firm identifier (firmid), Census tabulation flag (tabbed), administrative recording flag (ar), and vintage-consistent 2012 NAICS codes assigned by Fort and Klimek (2018).[A2] To keep active establishments, we drop those whose shipments, employment, or payroll are missing or non-positive. Following Autor et al. (2020) and Kehrig and Vincent (2021), we also drop establishments that are not used in official Census statistics (tabbed = "N') or are administrative records (ar = 1). To keep establishments with well-defined firms and industries, we drop establishments with missing firm identifiers or with missing Fort-Klimek NAICS codes. To drop outliers with unreasonable values, we drop establishments whose payroll per employment is less than 1 USD or whose employment is more than 1 million. We then aggregate the sample to the industry-by-firm level by aggregating establishments with the same firm identifier and NAICS.

When calculating horizontal diversification in Section 3.7.1 and factor shares in Section 4.1.2, we use four other variables: interplant transfers (ipt), the total value of assets at the beginning (tab), the total value of assets at the end (tae), and value-added (va). To avoid disclosure risks for Census data, we assume zero interplant transfer when the variable is missing, and we impute the industry average factor shares for a small number of establishments whose assets or value-added are missing. This treatment ensures that all statistics are calculated using the same sample as the benchmark and hence prevents creating small implicit sub-samples that could reveal confidential information.

A.2 Data for unit price analysis

To construct the data used for the unit price analysis in Section 4, we merge the CMF data set described in Section A.1 with data on product-level sales and quantity sold from the CMF's product trailer. We follow Kehrig and Vincent (2021) in keeping only non-imputed observations of product-level sales and quantity sold. Specifically, we keep only those observations for which i) the reported value was not replaced by an imputed value, or ii) the value, whether reported or not, was replaced by a Census analyst using establishment-specific information. Our use of the edit flag to implement this selection follows Kehrig and Vincent (2021) exactly; interested readers should consult their online appendix.

After dropping imputed observations, we further drop observations with a missing prod-

[A2] The Census tabulation flag indicates if the observation is used in official Census statistics. This variable is only available after 1997.

uct code, missing or non-positive product-level sales, and missing or non-positive product-level quantity sold. Unit prices as used in the main text are then just given by the ratio of product-level sales to product quantity sold.

A.3 Regression sample for Chinese import competition

Data on Chinese imports to the US and other developed countries comes from the BACI database of the Centre d'Études Prospectives et d'Informations Internationales (CEPII), which provides imports at the 6-digit HS product level.[A3] To map this data to the 2007 vintage of NAICS, we first apply the crosswalk from Pierce and Schott (2012), which assigns each 10-digit HS product to a single NAICS industry. To implement this linkage, we convert their 10-digit HS to NAICS crosswalk to a 6-digit HS to NAICS crosswalk by using 1992 US import weights taken from the Longitudinal Firm Trade Transactions Database (LFTTD), which provides imports at the 10-digit HS product level.

To construct domestic absorption by NAICS-2007 industry, we take sales and exports from a version of the CMF with longitudinally-consistent NAICS2007 codes from Fort and Klimek (2018). Total imports to the US are taken from the LFTTD, and are mapped to NAICS-2007 by first applying the Pierce and Schott (2012) crosswalk, which maps each annual vintage of the 10-digit HS to the most recent NAICS vintage in that year, and then mapping all years to NAICS-2007 using a weighted NAICS concordance in which the weights are given by sales shares from the CMF. Firm-level variables are taken from the LBD as discussed in section 2.

The regression sample is constructed from four different data sets: the CMF, LBD, LFTTD, and CEPII BACI. The sample selection applied to the CMF and LBD is identical to that described in Section A.1, with the addendum that we drop observations with NAICS codes 316212, 333293, 334210, and 339912 due to implausible values for Chinese imports. When using the LFTTD, we drop observations with missing product codes, as well as observations with non-positive or missing dollar values.

A.3.1 Aggregating the 2007 NAICS code

As described in the main text, we rely on the Pierce and Schott (2012) crosswalk to go from imports recorded at the HS-6 (BACI) or HS-10 (LFTTD) level to imports at the NAICS 6-digit

[A3] Data are downloaded from the CEPII website: http://www.cepii.fr/CEPII/en.

level. A complication in implementing this mapping is that this crosswalk assigns no HS-10 (and therefore no HS-6) codes to certain NAICS industries. These industries then have 0 import exposure by construction. To ensure that this does not happen, we aggregate the 2007 vintage of 6-digit NAICS so that each NAICS industry has at least one HS-10 product mapped to it. We combine NAICS codes with 0 HS-10 products mapped to them with codes that share a 5-digit NAICS code whenever possible.

A.3.2 Summary statistics and span of Chinese import exposure by diversification

Table A1 shows summary statistics for the dependent variable (change in log employment) and change in exposure to Chinese import competition measure (ΔCIE_{it} as defined in equation 14).

The key explanatory variable in Section 4.2 is the interaction between changes in Chinese import exposure and log effective number of industries, $\Delta\text{CIE}_{it} \times \ln\text{ENI}_{it}$. This estimation relies on the variations of Chinese import exposure conditional on differing levels of diversification. Figure A1 shows the span of ΔCIE_{it} conditional on log ENI. The solid line depicts the conditional median, and the shadowed area depicts the 25^{th} to 75^{th} percentile, where the values represent the predicted value from quantile regressions of ΔCIE_{it} on a 5^{th}-degree polynomial of log ENI. The dashed lines depict the median plus and minus 1.5 times of interquartile range (IQR). The figure confirms that variation in ΔCIE_{it} comes from a wide range of log ENI in our sample.

A.4 Up- and down-stream exposure to Chinese imports

We follow Acemoglu et al. (2016) in constructing measures of the first-order up- and downstream exposure to Chinese imports. Recall the Chinese import penetration ratio for an industry g:

$$\text{CIP}_{gt} = \frac{\text{imports from China}_{gt}}{\text{sales}_{gt} + \text{imports}_{gt} - \text{exports}_{gt}}. \tag{A1}$$

For each industry j, we first define that industry's upstream exposure to Chinese imports as a weighted average of the exposure of industry j's buyers to Chinese imports. where the weight for an industry g is given by the share of the industry j's output bought by industry g. Formally, this means that:

$$\text{CIP}^{\text{up}}_{jt} = \sum_{g} w^{\text{up}}_{gj}\text{CIP}_{gt}, \quad \text{where} \quad w^{\text{up}}_{gj} = \frac{\mu^{\text{up}}_{gj}}{\sum_{g'} \mu^{\text{up}}_{g'j}}. \tag{A2}$$

Table A1: Summary statistics

Variable	Mean	Std. Dev.
$\Delta \ln \text{emp} \times 100$	-2.56	58.59
$\Delta \text{CIE} \times 100$	1.58	3.22
observations	235,000	

Notes: Summary statistics of the sample underlying the import competition regression 12. CIE refers to the import competition exposure measure defined in equation 14. The sample size is rounded to the nearest thousand in accordance with Census disclosure requirement. Data sources: CMF, LBD, LFTTD, and BACI.

Figure A1: Span of ΔCIE conditional on log ENI

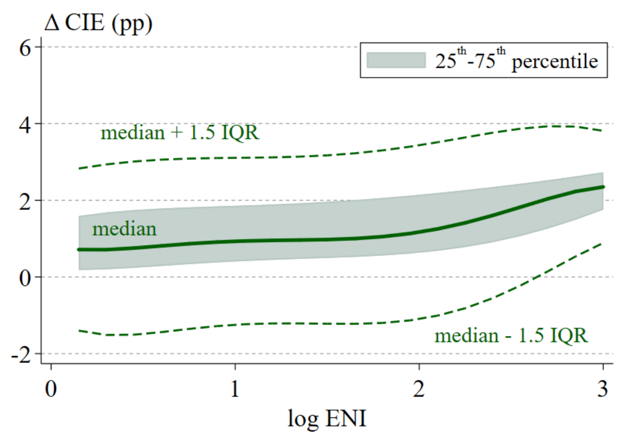

Notes: Data source: CMF, LBD, LFTTD and BACI.

Note that μ_{gj} is thus defined as the dollar value of industry j's output purchased by industry g.

Downstream exposure is similarly defined as:

$$\text{CIP}_{jt}^{\text{down}} = \sum_g w_{gj}^{\text{down}} \text{CIP}_{gt}, \quad \text{where} \quad w_{gj}^{\text{down}} = \frac{\mu_{gj}^{\text{down}}}{\sum_{j'} \mu_{gj'}^{\text{down}}}. \tag{A3}$$

In words, an industry j's downstream exposure to Chinese imports is given by a weighted average of the exposure of industry j's input-suppliers to Chinese imports, where the weight for an industry g is given by the share of industry j's intermediate inputs bought from industry g. Note that we adopt the convention used in Acemoglu et al. (2016) in naming the indirect exposure terms; accordingly, upstream exposure is exposure propagating upstream i.e. from an industry's buyers, whereas downstream exposure is exposure propagating downstream i.e. from an industry's input suppliers.

Given the industry-level up- and downstream exposure, a firm's change in upstream (or downstream) exposure is just given by the employment-weighted average of the changes in upstream (or downstream) exposures of the industries the firm operates in. Formally:

$$\Delta \text{CIE}_{it}^{\text{up}} = \sum_{j \in J_{it}} \left(\frac{\text{emp}_{ijt}}{\text{emp}_{it}} \right) \Delta \text{CIP}_{jt}^{\text{up}}, \tag{A4}$$

$$\Delta \text{CIE}_{it}^{\text{down}} = \sum_{j \in J_{it}} \left(\frac{\text{emp}_{ijt}}{\text{emp}_{it}} \right) \Delta \text{CIP}_{jt}^{\text{down}}. \tag{A5}$$

We instrument for these measures analogously to how we instrument for the direct exposure measure, i.e. we replace CIP_{gt} in the upstream and downstream exposure measures with

$$\text{CIP}_{gt}^{o} = \frac{\text{other developed countries' imports from China}_{gt}}{\text{sales}_{gt} + \text{imports}_{gt} - \text{exports}_{gt}} \tag{A6}$$

A.4.1 Implementation

To construct up- and downstream weights as defined above, we rely on the 1997 BEA supplementary use-table, which records sales flows between industries at a disaggregated level. Note that the denominators in equations A3 and A5 sum over all BEA industry codes, including non-manufacturing industries and final demand. We follow Acemoglu et al. (2016) in fixing the input-output table in an initial period (1997 in our case) and aggregate to the 4-digit NAICS level. We perform this aggregation because the industry code used by the BEA is not one-to-one with any vintage of the 6-digit NAICS code; constructing the above measures at the 6-digit NAICS level would thus require using both a weighted crosswalk between the BEA code and NAICS6 and a weighted NACIS6 crosswalk to map all observations to the 2007 vintage of NAICS6. To avoid relying too much on these crosswalks, we thus choose instead to construct the upstream and downstream exposure measures for NAICS 4-digit level, where the BEA industry code is exactly equal to NAICS after a minor modification. Finally, we collapse a handful of industries so that the resulting NAICS code is longitudinally consistent within our sample period. Specifically, we combine NAICS codes:

- 3391, 3371, 3339, 3334, and 3332
- 3315, 3314, and 3313
- 3152 and 3149
- 3342 and 3345
- 3261 and 3366
- 3113 and 3112
- 3118 and 3119
- 3312 and 3311
- 3327 and 3329

A.5 Extended regression tables

A.5.1 Import competition robustness

We examine the robustness of our results to several different specifications. The effect of all robustness checks on the coefficient of interest, the slope estimate for the interaction term ($\ln \text{ENI} \times \Delta\text{CIE} \times 100$), are reported in Figure A2.

(a) To address the potential of confounding industry trends, we replace year-fixed effects with year-industry fixed effects, where a firm's industry refers to its largest industry by employment. The point estimate of the coefficient on the interaction between the log effective number of industries and changes in exposure to Chinese import competition shrinks somewhat, but remains significant.

(b) Holmes and Stevens (2014) show that initially large plants in manufacturing shrank much more in response to Chinese import competition. They argue that this is because large plants produce standardized goods for the national market, whereas small plants tend to produce specialty goods for local markets and are thus relatively insulated from import competition. To the extent that the log effective number of industries is correlated with average plant size, our estimated interaction coefficient may thus partially reflect such variation. To address this

Figure A2: Estimation results for robustness regressions

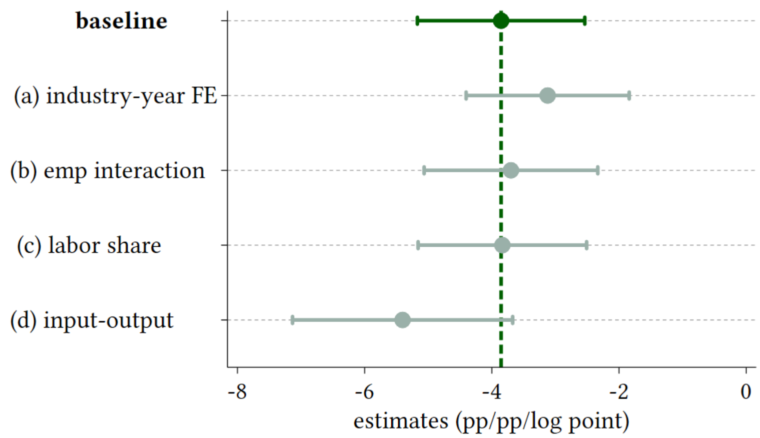

Notes: The vertical axis labels the robustness regression models as listed in Section A.5.1. The dots indicate point estimates and the capped spikes represent the 95% confidence interval. The baseline regression is identical to the last column in Table 2. The sample size for every regression is 235,000 (truncated to the nearest thousand in accordance with Census disclosure requirement). Data sources: CMF, LBD, LFTTD, and BACI.

we explicitly control for an interaction term between employment and exposure to change in Chinese competition: $\text{emp}_{it} \times \Delta\text{CIE}_{it}$. Figure A2 shows that the resulting point estimate on the interaction term of interest is virtually unchanged.

(c) Developing countries tend to export more labor-intensive goods when trading with developed countries (see e.g., Pierce and Schott, 2016). It is thus plausible that our estimated differential response is driven by systematic differences in factor shares along the dimension of specialization. To examine this, we add controls for firms' labor income share LS_{it} and its interaction with changes in Chinese import competition, $\text{LS}_{it} \times \Delta\text{CIE}_{it}$, where the labor income shares used are de-meaned by industry prior to aggregating to the firm level. Figure A2 again shows that the resulting point estimate on the interaction term of interest is unchanged.

(d) In the baseline regression we examine only the direct effect of Chinese import competition i.e. the effect stemming only from changes in import competition in the industries where firms operate. However, as shown in Acemoglu et al. (2016), Chinese import competition also had significant indirect effects stemming from changes in imports in firms' buyers and input suppliers. Excluding these indirect effects from the regression may bias estimated coefficients to the extent that indirect exposure is correlated with direct exposure. To address this, we use the BEA input-output table to construct indirect exposure measures for first-order up-and down-stream changes in Chinese import competition.[A4] Details of this construction are reported in Appendix A.3. As shown in the last line of Figure A2, adding controls for up- and down-stream exposure to Chinese imports increases the estimated coefficient of interest.

Table A2 reports extended results for 2. The first three rows of the two tables are identical. Table A3 reports the regression results for the four robustness regressions in Section A.5.1, where the point estimates and standard errors in the first rows are used in Figure A2.

A.6 Quality Engel curve estimation

Underlying our estimate of the aggregate slope of the quality Engel curve are product-specific elasticities of unit prices paid by consumers with respect to their total expenditure. To estimate these elasticities, we rely on the Consumer Expenditure Survey (CEX), which collects quarterly household expenditure data within detailed product categories. Each individual purchase is assigned to a particular month. In principle, each purchase could be reported separately; however, it is likely that some aggregation across purchases occurs. To minimize the influence of aggregation on our estimated elasticities, we follow Bils and Klenow (2001) in restricting to goods whose purchases tend to be distinct (e.g. cars, dishwashers) so that we can interpret the expenditure amount within a product category as a unit price. Our regression equations take the following form:

$$p_{ijt} = \beta_j + \gamma_j y_{it} + \mathbf{\Gamma}_j \mathbf{X}_{it} + \varepsilon_{ijt},$$

with p_{ijt} the log unit price paid by household i for good j at time t, y_{it} log total expenditure for household i at time t, and $\mathbf{X}_{it}$ denoting a matrix of household-specific controls. The coefficient of interest is γ_j, which traces out how unit prices within a product category change with the purchasing household's total expenditure.

To implement the above equation, we proceed as follows. First, we harmonize year-specific product codes across the 1990-2021 waves of the CEX. We then group product codes into nondurables and durables following Aguiar, Bils, and Boar (2020). We use total nondurable expenditure in place of total expenditure as the dependent variable to avoid introducing a mechanical correlation between the dependent and independent variables; this is particularly trenchant for product codes that plausibly represent a significant portion of total expenditure within a time

[A4] We adopt the language of Acemoglu et al. (2016); accordingly, down-stream propagation refers to effects propagating downstream from a firm's input suppliers, whereas up-stream propagation refers to effects propagating upstream from a firm's customers.

Table A2: Responses to import competition: baseline estimation results (all coefficients)

specification	(1)	(2)	(3)		(4)	
method	OLS	OLS	OLS	2SLS	OLS	2SLS
dependent variable	Δ ln employment × 100					
ln ENI × ΔCIE × 100					-2.97	-3.85
					(0.54)	(0.66)
ln ENI	−23.47	−15.39			−10.68	−9.31
	(0.84)	(0.98)			(1.23)	(1.37)
ΔCIE ×100			−0.84	−1.39	−0.79	−1.31
			(0.04)	(0.06)	(0.04)	(0.06)
firm age		−3.01	−3.00	−3.01	−3.01	−3.02
		(.068)	(.068)	(.068)	(.068)	(.068)
firm age squared		.06	.06	.06	.06	.06
		(.002)	(.002)	(.002)	(.002)	(.002)
firm emp.		−1.34	−2.73	−2.74	−1.35	−1.35
		(.220)	(.244)	(.243)	(.217)	(.216)
firm emp. squared		.005	.012	.012	.005	.005
		(.002)	(.003)	(.003)	(.002)	(.002)
exports		5.73	7.94	7.93	5.66	5.64
(coeff. & SE $\times 10^{-6}$)		(1.52)	(1.95)	(1.95)	(1.51)	(1.51)
adjusted R^2	0.01	0.03	0.03	0.03	0.04	0.03
F statistic				22, 640		10, 920
observations	235, 000					

Notes: Every regression controls for the year fixed effects, and columns 2 to 6 controls for an export dummy which equals 1 if exports are positive and 0 otherwise. We do not report fixed effects to avoid disclosure risks for Census data. The sample size is truncated to the nearest thousand in accordance with Census disclosure requirement. Robust standard errors are presented in parentheses. The F statistic is the Kleibergen-Paap rk Wald F statistic for weak instruments. Data sources: CMF, LBD, LFTTD, and BACI.

Table A3: Regression table for robustness regressions

specification	(a)	(b)	(c)	(d)
dependent variable	Δ ln employment × 100			
ln ENI × ΔCIE × 100	-3.12	-3.19	-3.83	-5.40
	(0.64)	(0.73)	(0.66)	(0.87)
ln ENI	-11.63	-10.24	-9.67	-12.25
	(1.36)	(1.42)	(1.37)	(1.59)
ΔCIE × 100	0.71	-1.30	-1.33	-0.94
	(0.66)	(0.06)	(0.14)	(0.07)
emp × ΔCIE × 100		-0.28		
		(0.13)		
labor share (pp)			-0.07	
			(0.02)	
labor share × ΔCIE × 100			0.003	
			(0.01)	
ΔCIE^{up} × 100				-0.28
				(0.18)
ΔCIE^{down} × 100				-4.58
				(0.44)
ln ENI × ΔCIE^{up} × 100				-2.19
				(1.33)
ln ENI × ΔCIE^{down} × 100				17.22
				(3.90)
lifecycle & exports	yes	yes	yes	yes
year-industry FEs	yes	no	no	no
adjusted R^2	0.027	0.033	0.033	0.033
Fstatistic	562.5	7,287	5.098	2,401
observations	235, 000			

Notes: We do not report fixed effects to avoid disclosure risks for Census data. The sample size is truncated to the nearest thousand in accordance with Census disclosure requirement. Robust standard errors are presented in parentheses. The F statistic is the Kleibergen-Paap rk Wald F statistic for weak instruments. Data sources: CMF, LBD, LFTTD, and BACI.

period, e.g. motor vehicles. To minimize the influence of measurement error, we follow Bils and Klenow (2001) in first aggregating to the annual level and then using the first two quarters of total nondurable expenditure to instrument for the latter two quarters of total nondurable expenditure. For each product, we control for a polynomial of the respondent's age, family size, number of children, and region-, urban-, and year-fixed effects. For a subset of goods (e.g. carpeting), we also control for the total number of rooms in the responding household's domicile. Finally, we restrict to non-imputed observations. Table A4 shows slope coefficients and standard errors by product category.

B Theory appendix

To simplify notation, we will drop the firm index i in this subsection, but we emphasize that all results are partial-equilibrium properties for an individual firm.

B.1 Parameterizing the distribution of technology productivity

The distribution of γ_q is critical: it determines the distribution of q and hence shapes the function $P(q)$. We need the distribution to be right-tailed and peak at 0 so that the firm's profit function is strictly increasing in q and the firm's technology problem is well-defined. In addition, to ensure that the household's utility is well-defined, we set the γ_q distribution to be bounded, so that an upper bound of quality $\overline{q}$ exists. We therefore parameterize the distribution using a bounded Pareto distribution with support $(0, 1)$ and shape parameter 1%.[A5]

The distribution of N is exogenously taken from the data distribution of ENI. We assume that N and γ_q are independent. The distribution of γ_A is generated in two steps. First, we independently draw each firm's γ_A from a bounded Pareto distribution with support $(0, 1)$ and shape parameter 1%. We then re-scale the γ_A for all single-industry firms by a scalar $s_A \in (0, 1)$, which gives an efficiency premium for multi-industry firms. This premium is to fit the fact that, in 1977, only 5% of firms were multi-industry, but they account for about 66% of total employment (see the last two rows of Table 3). The calibrated s_A is 4.9%.

[A5]The PDF of a bounded Pareto distribution with support $(\overline{x}, \underline{x})$ and shape $\alpha > 0$ is: $(\alpha \underline{x}^{\alpha} x^{-\alpha+1})/[1 - (\underline{x}/\overline{x})^{\alpha}]$.

B.2 Algorithm to approach a general equilibrium

This section specifies the algorithm to approximate a general equilibrium (GE) given the model parameters. The quality demand shifter is key. Recall that

$$\lambda(q; E, \mathcal{P}) = qP(q)^{\epsilon-\sigma} E^{\sigma} U(E, \mathcal{P})^{(1-\sigma)q^{-\phi}}. \tag{A7}$$

Given a quality demand shifter, $\lambda_0(q; E, \mathcal{P})$, firms' technology and product market problem gives the optimal technology $\{A_i^*, q_i^*\}_{i \in I}$ and prices $\{p_i^*\}_{i \in I}$. The optimum then forms a new quality demand shifter following the right-hand side of equation (A7). A GE quality demand shifter is then a fixed point to said operation. When the GE shifter is found, the equilibrium allocations and prices can be easily solved by the firms' and households' optimization problem.

The main task of our algorithm is thus to find the GE quality demand shifter. To make progress, we first take a log-linear approximation to the demand shifter:

$$\ln \lambda(q; E, \mathcal{P}) \approx \beta + \theta \ln q, \tag{A8}$$

where $\beta > 0$ and $\theta > 0$. The output of the algorithm is thus a pair (β^*, θ^*), and the function $\hat{\lambda}(q) = \beta^* q^{\theta^*}$ is an approximated GE quality demand shifter. Step 0 to 9 below specifies our algorithm:

Step 0. Set up:

(a) assign model parameters

(b) take the 1977 unweighted distribution of log ENI from data as the distribution for N.

(c) set the γ_q and γ_A distribution: bounded Pareto with support (0,1) and shape parameter 0.01.

(d) independently simulate K firms from said distributions: $\{\gamma_{qi}, \gamma_{Ai}, N_i\}_{i=1}^{K}$.

(e) re-scale γ_{Ai} by s_A if $N_i = 1$, where $s_A \in (0, 1)$ is defined in Appendix B.1.

(f) set a tolerance value T and a maximum number of iteration M

Step 1. Guess: $\beta = \beta_0$, and $\theta = \theta_0$.

Step 2. Given that $\lambda(q) = \beta q^{\theta}$ calculate each firm's optimal technology and price:

$$q_i^* = \left(\frac{\theta}{\theta+\epsilon-1}\right)^{\frac{1}{\omega}} \Gamma\, \gamma_{qi}{}^{\frac{1}{\omega}}\, N_i{}^{\frac{-\alpha_q}{\omega}}, \tag{A9}$$

$$A_i^* = \left(\frac{\epsilon-1}{\theta+\epsilon-1}\right)^{\frac{1}{\omega}} \Gamma\, \gamma_{Ai}{}^{\frac{1}{\omega}}\, N_i{}^{\frac{-\alpha_A}{\omega}}, \tag{A10}$$

$$p_i^* = \left(\frac{\epsilon}{\epsilon-1}\right)\frac{1}{A_i^*}.$$

Step 3. Discretize the quality space: define $\underline{q}$ and $\overline{q}$ as the 1st and 99th percentile of $\{q_i^*\}_{i=1}^K$, and partition $[\underline{q}, \overline{q}]$ into S intervals $\{Q_1, Q_2, Q_3, ..., Q_S\}$. Specifically, $Q_s = \left[(s-1)\left(\frac{\overline{q}}{S}\right), s\left(\frac{\overline{q}}{S}\right)\right]$, and the center points of the intervals are: $q_s = (s-0.5)(\overline{q}/S)$ for each s = 1, 2, ..., S.

Step 4. Approximate the CES price index function with the discretized domain:

$$P(q) = \left(\int_{I(q)} p_i{}^{1-\epsilon}\right)^{\frac{1}{1-\epsilon}} = \left[\mathbb{E}\left(p^{1-\epsilon}\,\middle|\, q\right)\cdot f(q)\right]^{\frac{1}{1-\epsilon}}.$$

Let $I_s = \{i \mid q_i^* \in Q_s\}$ for each quality bin $s = 1, 2, \ldots, S$. Calculate:

$$P_s^* = \left[\underbrace{\left(\frac{1}{|I_s|}\sum_{i\in I_s} p_i^{*\,1-\epsilon}\right)}_{\approx E(p^{1-\epsilon}|q)}\cdot\underbrace{\left(\frac{|I_s|}{K/S}\right)}_{\approx f(q)}\right]^{\frac{1}{1-\epsilon}}.$$

Step 5. Find U^* that solves the NH-CES definition equation: (see *Remark* below)

$$1 = \frac{1}{S}\sum_{s=1}^{S} E^{\sigma-1}\, q_s\, U^{*(1-\sigma)q_s{}^{-\phi}}\, P_s^{*\,1-\sigma}. \tag{A11}$$

Step 6. Calculate the discretized quality demand shifter:

$$\lambda_s^* = E^{\sigma} q_s U^{*(1-\sigma)q_s{}^{-\phi}} P_s^{*\,\epsilon-\sigma},$$

Step 7. Take the log-linear approximation by least squares: $(\ln \beta_1, \theta_1)' = (\boldsymbol{Q}'\boldsymbol{Q})^{-1}\boldsymbol{Q}'\boldsymbol{\lambda}$, where

$$\boldsymbol{Q} = \begin{bmatrix} 1 & \ln q_1 \\ \vdots & \vdots \\ 1 & \ln q_S \end{bmatrix} \quad \text{and} \quad \boldsymbol{\lambda} = \begin{bmatrix} \ln \lambda_1^* \\ \vdots \\ \ln \lambda_S^* \end{bmatrix}.$$

Step 8. compare (β_0, θ_0) and (β_1, θ_1) by calculating the distance:

$$d = \frac{1}{2}\left[\left(\frac{\beta_0 - \beta_1}{\beta_0 + \beta_1}\right)^2 + \left(\frac{\theta_0 - \theta_1}{\theta_0 + \theta_1}\right)^2\right]^{\frac{1}{2}}.$$

Step 9. Let m be the number of iterations including the current one.

- if $d \leq T$ and $m \leq M$: stop and report "a general equilibrium is found.'
- if $d > T$ and $m < M$: replace (β_0, θ_0) by $\left(\frac{\beta_0+\beta_1}{2}, \frac{\theta_0+\theta_1}{2}\right)$, and repeat step 1 to 9.
- if $d > T$ and $m = M$: stop and report "no general equilibrium is found.'

Remark: Recall that $c_i = \lambda(q){p_i}^{-\epsilon}$, where $\lambda(q) = qE^{\sigma}U^{\phi(q)}P(q)^{\epsilon-\sigma}$. We can write the segment-level demand function as:

$$C(q) = \left(\int_{I_q} {c_i}^{\frac{\epsilon-1}{\epsilon}}\, di\right)^{\frac{\epsilon}{\epsilon-1}} = qE^{\sigma}U^{\phi(q)}P(q)^{-\sigma}. \tag{A12}$$

Substitute the $C(q)$ in the definition of NH-CES utility using this equation. We then have:

$$1 = \int_0^{\overline{q}} q^{\frac{1}{\sigma}}\, C(q)^{\frac{\sigma-1}{\sigma}} U^{\frac{\phi(q)}{\sigma}}\, dq = \int_0^{\overline{q}} E^{\sigma-1} q U^{\phi(q)} P(q)^{1-\sigma}\, dq.$$

Equation (A11) is an approximation of this equation with a discretized quality space.

B.3 Computing the slope of the quality Engel curve

From equation (A12), in partial equilibrium, the household's maximization gives the segment-level demand function: $C(q) = qE^{\sigma}U^{\phi(q)}P(q)^{-\sigma}$. Plug it into the NH-CES utility function over quality segments, i.e., equation (17), we get

$$1 = E^{\sigma-1} \int_0^{\overline{q}} U^{\phi(q)} q P(q)^{1-\sigma} dq. \tag{A13}$$

Rearranging the equation, we can express the industry price index as:

$$\overline{P} \equiv \frac{E}{U} = E^{\sigma} \int_0^{\overline{q}} U(E, \mathcal{P})^{\phi(q)-1} q P(q)^{1-\sigma} dq.$$

Therefore, the slope of the quality Engel curve is:

$$\frac{\partial \ln \overline{P}}{\partial \ln E} = \sigma + \left[\frac{\int_0^{\overline{q}} \big(\phi(q) - 1\big) U^{\phi(q)-2} q P(q)^{1-\sigma} dq}{\int_0^{\overline{q}} U^{\phi(q)-1} q P(q)^{1-\sigma} dq} \right] \frac{\partial U}{\partial \ln E}. \tag{A14}$$

On the other hand, take logs on both sides of (A13) and take derivative with respect to $\ln E$.

$$\frac{\partial U}{\partial \ln E} = (1 - \sigma) \left[\frac{\int_0^{\overline{q}} U^{\phi(q)} q P(q)^{1-\sigma} dq}{\int_0^{\overline{q}} \phi(q) U^{\phi(q)-1} q P(q)^{1-\sigma} dq} \right].$$

Plugging it into equation (A14) gives:

$$\frac{\partial \ln \overline{P}}{\partial \ln E} = \sigma + (1 - \sigma) \left[\frac{\int_0^{\overline{q}} U^{\phi(q)} q P(q)^{1-\sigma} dq}{\int_0^{\overline{q}} \phi(q) U^{\phi(q)-1} q P(q)^{1-\sigma} dq} \right] \times \left[\frac{\int_0^{\overline{q}} \big(\phi(q) - 1\big) U^{\phi(q)-2} q P(q)^{1-\sigma} dq}{\int_0^{\overline{q}} U^{\phi(q)-1} q P(q)^{1-\sigma} dq} \right]. \tag{A15}$$

The integrals over quality segments are then numerically approximated in the same way as introduced in steps 4-5 in Section B.2.

Table A4: Quality Engel curve slope estimates by product

Product	Slope	Std Err
Window coverings	1.09	0.06
Jewelry	1.00	0.03
Purchase of motorized camper	1.00	0.47
Floor coverings, nonpermanent	0.92	0.04
Clocks and other household decorative items	0.91	0.02
Cars	0.91	0.04
Dinnerware, glassware, serving pieces	0.88	0.02
Trucks	0.88	0.05
Winter sports equipment	0.86	0.09
Outboard motors	0.85	0.35
Boat without motor and boat trailers	0.84	0.26
Hearing aids	0.82	0.13
Outdoor furniture	0.82	0.06
Kitchen, dining room furniture	0.82	0.06
Watches	0.81	0.03
Lamps, lighting fixtures, ceiling fans	0.77	0.03
Living room tables	0.76	0.05
Wigs and hairpieces	0.76	0.08
Photographic equipment	0.75	0.04
Women's suits	0.74	0.02
Luggage	0.72	0.03
Men's suits	0.72	0.03
Outdoor equipment	0.71	0.05

Portable dishwasher	0.71	0.25
Sofas	0.71	0.05
Men’s sportcoats, tailored jackets	0.71	0.03
Stereos, radios, speakers	0.70	0.08
Purchase of boat with motor	0.68	0.23
Office furniture for home use	0.68	0.07
Cooking stoves, ovens	0.66	0.06
Camping equipment	0.65	0.05
Women’s sportcoats, tailored jackets	0.65	0.03
Sewing machines	0.64	0.10
Women’s dresses	0.64	0.01
Infants’ furniture	0.63	0.07
Hunting and fishing equipment	0.63	0.04
Women’s footwear	0.61	0.01
Men’s shirts, sweaters, and vests	0.59	0.01
Business equipment for home use	0.59	0.08
Women’s sweaters, shirts, tops, vests	0.58	0.01
Telephones and accessories	0.58	0.02
Toys, games, arts and crafts, and tricycles	0.58	0.01
Nonelectric cookware	0.57	0.03
Athletic gear, game tables, and exercise equipment	0.57	0.02
Playground equipment	0.57	0.12
Men’s coats and jackets	0.57	0.02
Women’s coats and jackets	0.56	0.02
Women’s skirts	0.55	0.02
Living room chairs	0.54	0.05

Microwave ovens	0.54	0.04
Satellite dishes	0.53	0.10
Women’s pants and shorts	0.51	0.01
Infant dresses, outerwear	0.50	0.02
Clothes washer or dryer	0.50	0.03
Men’s footwear	0.50	0.01
Televisions	0.50	0.03
Refrigerators, freezers	0.50	0.04
Hand tools	0.49	0.04
VCR’s and video disc players	0.48	0.03
Personal digital audio players	0.47	0.06
Motorcycles	0.47	0.21
Tires - purchased, replaced, installed	0.47	0.02
Men’s pants and shorts	0.46	0.01
Electric personal care appliances	0.44	0.02
Musical instruments and accessories	0.44	0.06
Infants’ equipment	0.44	0.05
Computers and computer hardware for nonbusiness use	0.44	0.03
Girls’ active sportswear	0.43	0.03
Men’s active sportswear	0.43	0.02
Boy’s shirts and sweaters	0.43	0.02
Digital book readers	0.43	0.06
Girls’ dresses and suits	0.42	0.03
Girls’ shirts, blouses, sweaters	0.42	0.02
Video cassettes, tapes, and discs	0.41	0.01
Computer software and accessories for nonbusiness use	0.40	0.02

Boys' suits, sportcoats, vests	0.36	0.07
Portable heating and cooling equipment	0.36	0.04
Girls' coats and jackets	0.36	0.03
Video game software	0.35	0.02
Girls' skirts, pants, and shorts	0.35	0.02
Girls' footwear	0.34	0.02
Power tools	0.34	0.04
Infant coat, jacket, snowsuit	0.34	0.05
Eyeglasses and contact lenses	0.33	0.02
Electric floor cleaning equipment	0.31	0.03
Video game players and video game accessories	0.30	0.06
Trailer and other attachable campers	0.30	0.31
Lawn and garden equipment	0.30	0.04
Boys' coats and jackets	0.28	0.03
Boys' footwear	0.27	0.02
Boys' uniforms and active sportswear	0.24	0.03
Window air conditioners	0.16	0.07