

Consumer Surplus from Suppliers: How Big is it and Does it Matter for Growth?

David Baqaee
UCLA

Ariel Burstein
UCLA

Cédric Duprez
NBB and UMons*

Emmanuel Farhi[†]

May 2025

Abstract

Consumer surplus, the area between the demand curve and the price, plays a key role in many models of trade and growth. Quantifying it typically requires estimating and extrapolating demand curves. This paper provides an alternative approach to measuring consumer surplus by focusing on firms as consumers of inputs. We show that the elasticity of a downstream firm's marginal cost to supplier additions and separations measures the downstream firm's consumer surplus relative to its input costs. Using Belgian data and instrumenting for changes in supplier access, we find that for every 1% of suppliers gained or lost, the marginal cost of downstream firms falls or rises by roughly 0.3%. Our estimates are directly informative about the strength of love-of-variety effects and the gains from movements along quality-ladders. We use our microeconomic estimates of consumer surplus to assess the macroeconomic importance of supplier additions and separations in a growth-accounting framework. We find that supplier churn plausibly accounts for about half of aggregate productivity growth.

**First version: June, 2022, previously circulated as "Supplier Churn and Growth: A Micro-to-Macro Analysis."*
We thank Costas Arkolakis, Andy Atkeson, Pablo Fajgelbaum, Kyle Herkenhoff, Jim Hamilton, Federico Huneus, Jozef Konings, Marc Melitz, Guido Menzio, Ezra Oberfield, John Romalis, Todd Schoellman, Felix Tintelnot, Gianluca Violante, Venky Venkateswaran, Jonathan Vogel, and Alwyn Young for comments. We acknowledge financial support from NSF grant No. 1947611 and the Alfred P. Sloan Foundation. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the National Bank of Belgium.

[†]Emmanuel Farhi tragically passed away in July, 2020.

1 Introduction

Consumer surplus is the benefit customers get from having access to suppliers. The presence of consumer surplus motivates buyers to trade with new suppliers, and the continual appearance of new suppliers, each generating additional surplus, can then fuel growth. Despite the critical role consumer surplus plays in theories of trade and endogenous growth, there is little direct evidence on its magnitude. This study measures buyers' surplus in firm-to-firm trade by estimating how supplier access affects buyers' marginal costs. We then use these estimates to quantify the importance of changes in supplier availability for aggregate productivity growth.

The object of interest in this paper is the “consumer surplus ratio,” denoted by δ , and defined to be the area under the demand curve above the price relative to expenditures. This measures the gap between what consumers are willing to pay and what they actually pay for goods. Welfare and output counterfactuals depend on δ in most models of growth and trade, including expanding-variety and quality-ladder models.¹ The consumer surplus ratio is also an important statistic in the industrial organization literature, as it measures the fraction of gross surplus that sellers cannot capture (Spence, 1976; Mankiw and Whinston, 1986). In fact, as pointed out by Makowski and Ostroy (2001), $\delta = 0$ is a defining characteristic of perfect competition.

Our study consists of a microeconomic and a macroeconomic part. We discuss the two parts of the paper in turn. To estimate the surplus ratio, δ , at the micro-level, we employ a unique approach that enables us to estimate the area under the input demand curve without specifying the demand system itself. Traditionally, inferring consumer surplus requires estimating and integrating demand curves. Demand estimation considers how quantities respond to prices. Using this variation, one can estimate the price elasticity of demand over the region where prices and quantities vary. Given these estimates, and a functional form, one can then integrate the demand curve up to the choke price to arrive at an estimate for consumer surplus.

A shortcoming of the standard approach is that two demand curves can look similar locally, over the region where price and quantity variation is observed, but yield very

¹In expanding-variety models (e.g. Dixit and Stiglitz, 1977; Krugman, 1979; Romer, 1987; Melitz, 2003), we show that δ captures what is commonly called the degree of preference for variety (Vives, 1999) or love of variety. In these models, efficiency of the equilibrium depends on a comparison of δ with the markup (see, e.g., Matsuyama and Ushchev, 2020b). Furthermore, optimal industrial policy and the response of aggregate output to shocks also depend on the value of δ (see Dhingra and Morrow, 2019, Baqaee et al., 2024, or Baqaee and Farhi, 2020 for some examples). In quality-ladder models (e.g. Aghion and Howitt, 1992; Grossman and Helpman, 1993), the consumer surplus associated with a movement along the quality-ladder is determined by δ and expenditures at each rung of the ladder.

different amounts of consumer surplus due to extrapolation error. For example, a translog and CES demand curve can have the same value and shape locally, but imply different amounts of consumer surplus.

Our empirical strategy is instead based on the following theoretical result: we show that for any input demand system, the elasticity of a downstream firms' marginal cost with respect to upstream entry and exit is equal to δ . If $\delta > 0$, then firms' marginal costs fall (or rise) as they gain (or lose) access to suppliers. Because marginal cost is, at least in principle, observable we can use this result to estimate δ without specifying the demand system.²

We implement our approach using a detailed survey of manufacturing firms in Belgium, called Prodcom, that tracks sales and output quantities. We merge Prodcom with firm-to-firm input-output linkage information from value-added tax (VAT) returns. Using this tax information, we observe at annual frequency almost all domestic suppliers of the firms in Prodcom. We regress changes in marginal costs on supplier additions and separations. We show that, when this regression is consistently estimated, the coefficients identify the average consumer surplus ratio for additions and for separations.

To achieve consistent estimation, we instrument the addition and subtraction of suppliers using firm births and deaths. To ensure that births and deaths of upstream suppliers are not driven by idiosyncratic shocks to their downstream customers, we restrict attention to entry and exits of suppliers for whom the downstream firm is small as a share of their customer base (e.g., less than 5%). Identification requires that additions and separations of suppliers caused by our instruments are not related to idiosyncratic shocks to the downstream firms' marginal costs, like the downstream firm's productivity shocks. We also control for other input prices and include 6-digit product-by-year fixed effects to absorb market-level shocks.

We find sizeable microeconomic effects of supply linkage destruction and creation on downstream marginal costs. That is, we reject the perfectly competitive benchmark where each supplier faces perfectly elastic demand and $\delta = 0$. Our baseline estimate is $\delta \approx 0.3$, so if 1 percentage point of a firm's suppliers, in terms of its variable costs share, are lost or added, then this raises or lowers its marginal cost by around 0.3 percentage points. If demand for inputs is CES, then the consumer surplus ratio is pinned down by the elasticity of substitution. Under this assumption, our estimates of the consumer surplus ratio

²To apply this approach to a household, one would have to use the ideal price index for consumption in place of marginal cost. Unlike marginal costs of production, which we measure by dividing variable costs by quantity of output, the ideal price index of consumption is much harder to measure without knowing the demand system. This ultimately stems from the fact that, unlike quantity of output, utility is only defined up to monotone transformations.

are consistent with an elasticity of substitution of roughly 4.5. However, our estimates can also be used to calibrate parameters of more flexible demand systems where the consumer surplus ratio is not so tightly connected to the (local) elasticity of substitution.³ We also find some evidence of heterogeneity: δ is smaller if the supplier is larger relative to the average firm in its industry.

We now turn to the macroeconomic implications. Unlike consumer surplus for households, which matters for welfare but may not be fully reflected in national income statistics, the area under firms' input demand curves does show up in measured aggregate output growth. In the macroeconomic part of the paper, we develop a growth-accounting framework to quantify the importance of supplier churn for measured aggregate growth, adding an extensive margin for supplier additions and separations to otherwise standard growth accounting formulas (i.e. Solow, 1957; Hulten, 1978; Basu and Fernald, 2002; Baqaee and Farhi, 2019).⁴ We take into account how the formation and separation of supplier links affects the prices of downstream firms, and how these price changes are transmitted along existing supply chains from supplying firms to purchasing firms, all the way down to final consumers. This accounting exercise does not require a fully spelled-out model of market structure, factor markets, or link formation but is consistent with many different structural models.

To implement our growth accounting formula, we require the firm-to-firm input-output matrix over time and estimates of the consumer surplus ratio δ . We discipline the former using Belgian VAT data and the latter by extrapolating our microeconomic estimates of δ from Prodcom firms to the whole Belgian economy. Although we accommodate both expanding-variety and quality-ladder models, our macro growth-accounting exercise requires an additional assumption, not needed in the micro estimation exercise: we rule out additions and separations of suppliers caused by discontinuous movements of input demand curves, e.g. due to biased downstream technology shocks or changes to continuing suppliers' prices. We find that around half of aggregate productivity growth in Belgium between 2002 and 2018 can plausibly be accounted for by supplier churn.

The structure of the paper is as follows. Section 2 shows how the elasticity of marginal cost to supplier access can be used to measure δ . This motivates our empirical strategy, which we describe and report in Section 3. Section 4 introduces the aggregation framework and presents our growth accounting formula. We use these results, and our earlier

³We also find a reduced-form pass-through from marginal costs into prices of around 60%. That is, a little over half the changes in marginal costs are passed onto downstream customers while the remaining 40% are absorbed by markups.

⁴By extensive margin of additions and separations, we specifically mean a case where expenditure shares change discontinuously when suppliers are added or dropped. If expenditure shares change smoothly to or from zero, then standard growth accounting formulas apply without change.

microeconomic estimates, to decompose aggregate growth in our data in Section 5. We conclude in Section 6. There is an accompanying online appendix containing additional proofs, details about the data, and robustness checks.

Related literature. Our paper is related to several different literatures. First, as discussed above, our analysis contributes to the literature on growth and trade with an extensive margin of inputs. A key object of interest and source of welfare gains in this literature comes from the love for product variety.⁵ In models with monopolistic competition, the love-of-variety effect is usually defined using the elasticity of the utility function with respect to quantity, e.g. Vives (1999), Benassy (1996), Zhelobodko et al. (2012), and Dhingra and Morrow (2019). This elasticity is inherently unobservable since utility is only defined up to monotone transformations. We characterize the love-of-variety effect in terms of the area under the demand curve instead, which depends only on observables, and our characterization does not require a separable demand system. Although our empirical strategy is valid only for production, the theoretical relationship we establish between love-of-variety and the area under the demand curve also applies to consumption.

We further contribute to this literature by estimating changes in marginal cost as firms lose or gain access to suppliers, and showing that this identifies δ . We can do this because our data allows us to track costs, output quantities, and firms' suppliers. In lieu of this data, researchers have typically relied on very indirect evidence to discipline the consumer surplus from new suppliers in their models. For example, expanding-varieties models typically use a CES demand system, where the price elasticity of residual demand at any point on the demand curve also controls the love-of-variety effect. Similarly, quality-ladder models are often disciplined by indirect inference via matching moments on firm employment dynamics, patents, and growth (see Garcia-Macia et al., 2019 and Akcigit and Kerr, 2018 for example).⁶

Our paper is also related to the literature on production networks. Empirical studies by Jacobson and Von Schedvin (2015), Carvalho et al. (2014), and Miyauchi (2018) have shown firm failures are transmitted across supply chains and affect the sales of other firms in neighboring parts of the production network using reduced-form methods. Compared to these papers, we use an instrumental variable strategy to study the causal effect of firm

⁵The love-of-variety effect has been theoretically studied by Zhelobodko et al. (2012), Dhingra and Morrow (2019), Baqaee et al. (2024), and Matsuyama and Ushchev (2020b, 2023) amongst many others.

⁶There is a large literature that provides reduced-form evidence of how changes in policies (e.g. import tariffs or market access) impact firm outcomes such as size, productivity, markups, and firm product-scope. See, for example, Amiti and Konings (2007), Brandt et al. (2017), Goldberg et al. (2010), Bernard et al. (2019), and De Loecker et al. (2016). Although this literature provides evidence that input variety matters for firm-level outcomes, it does not provide an estimate of δ for the downstream firm.

failures on marginal cost (rather than sales), and we link our estimates to δ , which is an important statistic in many models with an extensive margin. For example, Baqaee (2018) and Baqaee and Farhi (2020) show that the response of aggregate output to microeconomic shocks depends on the consumer surplus ratio.

Other papers that emphasize the extensive margin of firm-to-firm linkages include Oberfield (2018), Lim (2018), Tintelnot et al. (2018), Elliott et al. (2022), Taschereau-Dumouchel (2020), Acemoglu and Tahbaz-Salehi (2020), Bernard et al. (2022), Huneus (2018), Miyauchi (2018), Boehm and Oberfield (2020), and Arkolakis et al. (2021).⁷ Our paper complements this literature in two ways: our micro estimates of the value of link formation can be used to discipline these models, and our growth accounting exercise provides moments about the aggregate importance of supplier churn that can be used as calibration targets. Unlike the structural literature, we take changes in firms' sizes and the formation and separation of links between firms as given (i.e. we take them from the data). Hence, we do not provide a fully specified model for counterfactuals. Since we do not model why firms form and break links, our exercise does not take a stance on the ultimate causes of firm growth (e.g. higher productivity or better ability to find matches).⁸

Finally, our paper is also related to a deep literature on correcting price indices to account for the entry and exit of goods. Our macroeconomic exercise quantifies the importance of supplier entry and exit for *measured* aggregate growth.⁹ The macroeconomic and trade literatures on the importance of entry and exit, which trace their origins to Hicks (1940), have been greatly influenced by Feenstra (1994) who introduced a methodology for accounting for product entry and exit under a CES demand system.

This CES methodology owes its popularity to its simplicity and non-demanding information requirements. For example, Broda and Weinstein (2006) apply it to calculate welfare gains from trade due to newly imported varieties, and Broda and Weinstein (2010) compute the unmeasured welfare gains from changes in varieties in consumer non-durables. Using a similar methodology, Jaravel (2019) calculates the gains from consumer product variety across the income distribution, while Gopinath and Neiman (2014), Halpern et al.

⁷Some papers in the literature model firm-to-firm link formation as the outcome of firms choosing amongst alternative production recipes, for example Boehm and Oberfield (2020), Acemoglu and Azar (2020), and Kopytov et al. (2024). In these models, once we minimize costs over all possible recipes, there is an induced cost function that maps input prices and output quantity to total cost. Our notion of surplus and our empirical strategy are applicable to the induced cost-function in such models.

⁸In this sense, our results are not inconsistent with the findings of Bernard et al. (2022) who show that firms tend to grow primarily by adding new customers.

⁹There is a large body of work that decomposes changes in a weighted-average of firm-level productivities into reallocation, entry, and exit terms (see e.g. Baily et al., 1992; Foster et al. 2001). However, the object these studies decompose is not aggregate productivity in a growth accounting sense — that is, it does not measure the gap between real output and real input growth. See Petrin and Levinsohn (2012), Hsieh et al. (2018), Baqaee and Farhi (2019), and Baqaee et al. (2024) for more details.

(2015), and Blaum et al. (2018) study the welfare gains from trade in intermediate inputs.¹⁰ Hottman et al. (2016) build on the Feenstra (1994) methodology to adjust firm-level productivity estimates to account for quality and product scope, whereas Aghion et al. (2019) use it to correct aggregate growth rates for expanding varieties and unmeasured quality growth. Hausman (1996), Feenstra and Weinstein (2017), and Foley (2022) provide alternative price index corrections that use non-CES demand systems.

A universal theme in this literature is to estimate or calibrate price elasticities of demand and infer the value of entering and exiting products by inverting or integrating demand curves under parametric restrictions (e.g. isoelastic, linear, or translog demand). Our approach differs from this literature in that we attempt to identify the area under the input demand curve directly through its effect on downstream marginal costs rather than via implicit or explicit integration of demand curves.

Although we do not estimate price elasticities, our paper is related to the broader objective of estimating demand curves. Whereas we focus on the integral of demand, the literature on demand estimation tends to focus on its derivatives. For example, the first derivative of demand is estimated by the literature on price elasticities (see, e.g., Berry and Haile, 2021); the second derivative is estimated by the literature on pass-through of marginal cost into the price (see, e.g., Burstein and Gopinath, 2014); and the third derivative is estimated by the literature on variable pass-through (see, e.g., Amiti et al., 2019). We directly estimate the integral, or antiderivative, of demand relative to expenditures, which determines the value of new goods. Much like the price elasticity and the degree of pass-through, the consumer surplus ratio is generically a complicated object that depends on where the perturbation occurs. However, as with estimates of price elasticities and pass-throughs, our estimates can help to pin down deeper parameters of the cost function given different parametric assumptions.

2 Microeconomic Results: Theory

In this section, we derive expressions for how supplier addition and separation affect a downstream firm's marginal cost. The partial equilibrium results in this section serve as the basis for our firm-level regressions in Section 3. We delay general equilibrium and aggregation to Sections 4 and 5.

¹⁰The methodology of Feenstra (1994) requires knowledge of the elasticity of substitution, which is typically estimated using data on expenditure switching. Blaum et al. (2018) instead uses changes in the buying firm's revenues (and parametric assumptions on the production function and demand for the buying firms' output) to estimate the elasticity of substitution between imports and domestic inputs.

Consider a *downstream* firm, indexed by i , whose variable cost function is

$$C_i(\mathbf{p}, A_i, q_i) = mc_i(\mathbf{p}, A_i) q_i,$$

where \mathbf{p} is the vector of quality-adjusted input prices (including primary factor prices), A_i indexes technology, and q_i is the total quantity of output.¹¹ We allow the firm to have fixed costs of operation, but assume that variable production has constant returns to scale.

If an input is unavailable to the downstream firm, it is as-if its price is infinite. When the price of an input jumps to infinity, we say that it became unavailable; when the price becomes finite, we say that it became available. The prices that jump to or from infinity could be market prices or shadow prices. For example, entry or exit of suppliers cause market prices to jump. On the other hand, the formation or destruction of bilateral matches or changes in the downstream firm’s decision to pay fixed costs associated with linking to some supplier cause shadow prices to jump.¹²

Assume that inputs are grouped into types. The cost function is symmetric in input prices that belong to the same type but not necessarily symmetric across types. Formally, two inputs are the same *type* if swapping their prices does not affect variable cost. This assumption ensures that the downstream firm’s input demand curve for all varieties of a given type J are the same function $x_{ij}(\mathbf{p}, A_i, q_i)$. We do not restrict own-type or cross-type price elasticities. Without loss of generality, we assume that inputs of the same type also have the same initial price when they are available (if they have different prices when available, treat them as different types). For notational convenience, suppose there is a countable number of types. Almost all popular production technologies used in the macroeconomics and trade literatures feature a notion of “types.”¹³

¹¹In the body of the paper, we assume that firms take input prices as given. In Appendix B, we show that, under some additional assumptions, our empirical strategy to estimate consumer surplus ratio remains valid if firms face a schedule of input prices as a function of input quantities instead. This input price schedule, which we take as given, could, for example, be the outcome of price discrimination or bargaining.

¹²The availability of varieties to downstream firm i may be exogenous or endogenous to i . For example, it could be that the mass of varieties i has access to responds to i ’s productivity. This could be because of decisions made by i ’s suppliers if more suppliers choose to make their variety available to i when i is more productive. Or it could be because of decisions made by i , who may be willing to pay the fixed costs necessary for gaining access to more suppliers when it is more productive. We do not endogenize the availability of inputs and only consider i ’s variable cost minimization taking the availability of varieties as given. A fully specified model would be required for counterfactuals.

¹³For example, for CES, we say two inputs have the same type if they have the same share parameter and price. In Melitz (2003), two varieties are the same type if they have the same productivity draw. For the homothetic demand systems in Matsuyama and Ushchev (2017), and the separable demand system in Fally (2022), we say that two inputs have the same type if they share the same residual demand function and the same price. Our results apply to these demand systems even though they are usually parameterized with an uncountable number of types.

Define the *consumer surplus ratio* associated with each input of type J to be:¹⁴

$$\delta_{iJ}(p) = \frac{\int_p^\infty x_{iJ}(\xi) d\xi}{px_{iJ}(p)} \geq 0. \quad (1)$$

Equation (1) is the area under the demand curve for a supplier of type J above the price p per unit of expenditures. This is depicted graphically in Figure 1 as the ratio of A to B . As long as the demand curve is strictly downward sloping, δ_{iJ} is strictly positive. If the demand curve for an individual supplier is perfectly horizontal, which is consistent with perfect competition, then $\delta_{iJ} = 0$.

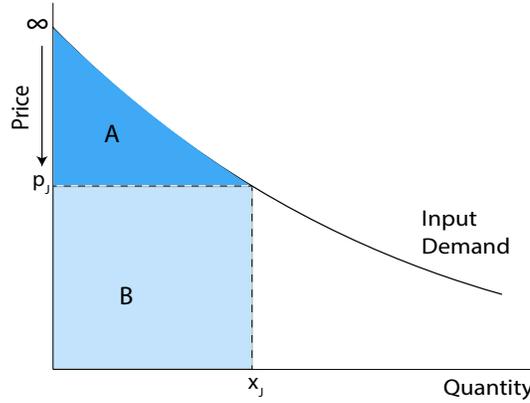


Figure 1: Consumer surplus ratio, δ , is the ratio of A to B .

Let the share in variable costs of each type- J input purchased by firm i be

$$\Omega_{iJ} = \frac{p_J x_{iJ}(\mathbf{p}, A, q_i)}{C_i(\mathbf{p}, A_i, q_i)}.$$

Suppose that there is a continuum of inputs of each type, so that we can continuously perturb their mass. Let ΔM_{iJ}^{add} be the mass of inputs of type J that i gains access to (i.e. the price jumps down from infinity to $p_J < \infty$), and ΔM_{iJ}^{sep} be the mass of inputs of type J that i loses access to (i.e. the price jumps to infinity from p_J).

The next proposition loglinearizes the downstream firm's marginal cost and decomposes it into marginal price changes, inframarginal price jumps, and changes in the firm's own technology.

¹⁴In equation (1), we suppress dependence of the conditional input demand x_{iJ} on arguments other than the price of that supplier since those other arguments are being held constant. We include the additional arguments when it helps the exposition.

Proposition 1 (Downstream Marginal Cost). *Consider a downstream firm i facing a change in the vector of input prices by type $\Delta \mathbf{p}$, the measure of available inputs by type ΔM_i^{add} and ΔM_i^{sep} , and the technology parameter ΔA_i . To a first-order approximation in these primitives, the change in the downstream firm's marginal cost is*

$$\Delta \log mc_i \approx \underbrace{\sum_J \Omega_{iJ} M_{iJ} \Delta \log p_J}_{\text{marginal changes}} - \underbrace{\sum_J \delta_{iJ} \Omega_{iJ} \Delta M_{iJ}}_{\text{inframarginal changes}} + \underbrace{\frac{\partial \log C_i}{\partial \log A_i} \Delta \log A_i}_{\text{own technology}} \quad (2)$$

where M_{iJ} is the initial mass of inputs of type J and $\Delta M_{iJ} = \Delta M_{iJ}^{add} - \Delta M_{iJ}^{sep}$ is the net change in the mass of available inputs of type J .

In words, the log change in the marginal cost of the downstream firm depends on the costs of its inputs, captured by the first two summands, as well as its own technology, captured by the last summand. The price of inputs can change on the margin or they can jump. If the change in input prices is small, then their effect on the downstream firm's marginal cost depends on the expenditures on the input.¹⁵ On the other hand, if input prices jump discretely due to changes in product availability, then their effect on the downstream firm's marginal cost depends on the area under the input demand, which is captured by the product of δ_{iJ} , and expenditures on the inputs whose price jumps $\Omega_{iJ} \Delta M_{iJ}$. That is, changes in the availability of inputs generate surplus for the downstream producer according to the total area under the input demand curve above the price. Note that consumer surplus is not δ_{iJ} (the object we estimate) but the product of δ_{iJ} and the expenditure share Ω_{iJ} (the latter is directly observable).

Additions and subtractions of suppliers that happen smoothly, without a discontinuous change in the price, do not affect the marginal cost of the downstream firm to a first-order. The expenditure share on varieties that are added or dropped in this way is zero at the choke price where they are added or dropped. Hence, their impact on the downstream firm's marginal cost is also zero to a first order by Shephard's lemma. This comment also applies to additions and separations that are caused by shifts of the input demand curve (as opposed to movements along the input demand curve). That is, if a shock to other suppliers or technology causes a given supplier to be added or dropped by moving its input demand curve in a continuous fashion, then this has no first-order effect on the overall addition and separation share and does not affect (2).

To better understand Proposition 1, we work through some simple examples.

¹⁵The non-linear impact of changes in prices of continuing inputs on marginal cost can be approximated using Tornqvist or other chaining procedures. We do this in our growth accounting analysis, as described in Footnote 39.

Example 1 (Expanding Varieties with CES). Consider the CES special case where demand for inputs of type J is

$$x_{ij}(p_J) = \frac{\omega_{ij} p_J^{-\sigma}}{\left(\sum_{J'} \omega_{iK} p_{J'}^{1-\sigma} M_{iJ'}\right)^{\frac{-\sigma}{1-\sigma}}} q_i, \quad (3)$$

where ω_{ij} and ω_{iK} are exogenous parameters, $\sigma > 1$ and q_i is total quantity. In this case, the integral in the definition of the consumer surplus ratio can be evaluated in closed form:

$$\delta_{ij} = \frac{\int_{p_J}^{\infty} x_{ij}(\xi) d\xi}{p_J x_{ij}} = \frac{1}{\sigma - 1} \geq 0.$$

Hence, in response to a change in the availability of some varieties of type J , the change in the downstream marginal cost is

$$\Delta \log mc_i \approx -\Omega_{ij} \Delta M_{ij} \delta_{ij} = -\Omega_{ij} \Delta M_{ij} \frac{1}{\sigma - 1}. \quad (4)$$

That is, the standard “love-of-variety” effect is measured by δ_{ij} .¹⁶

Once we depart from CES, the consumer surplus ratio need not be the same for all input types, as the example below demonstrates.

Example 2 (Expanding Varieties with Heterogenous Surplus Ratios). Consider a unit cost function defined implicitly by¹⁷

$$1 = \sum_J M_{ij} \frac{\omega_{ij}}{\sigma_J - 1} \left(\frac{p_J}{mc_i}\right)^{1-\sigma_J}.$$

If $\sigma_J = \sigma$ for every J , then this is a CES technology. The input demand curve for a type J variety is

$$x_{ij}(p_J) = \frac{\omega_{ij} \left(\frac{p_J}{mc_i}\right)^{-\sigma_J} q_i}{\sum_{J'} M_{iJ'} \omega_{iJ'} \left(\frac{p_{J'}}{mc_i}\right)^{1-\sigma_{J'}}}.$$

¹⁶Instead of using Proposition 1, one can also derive (4) by differentiating the closed-form expression for CES marginal cost,

$$\log mc_i = \frac{1}{A_i} \left[\sum_J \int_0^{M_{ij}} \omega_{ij} p_J^{1-\sigma} dj \right]^{\frac{1}{1-\sigma}} = \frac{1}{A_i} \left[\sum_J \omega_{ij} M_{ij} p_J^{1-\sigma} \right]^{\frac{1}{1-\sigma}},$$

with respect to the mass of varieties, M_{ij} , and using the fact that $\Omega_{ij} = \omega_{ij} p_J^{1-\sigma} / \left(\sum_{J'} \omega_{iK} p_{J'}^{1-\sigma} M_{iJ'}\right)$.

¹⁷See Matsuyama and Ushchev (2020a) for more information about this demand system.

The consumer surplus ratio associated with input J is

$$\delta_{iJ} = \frac{\int_{p_J}^{\infty} x_{iJ}(\xi) d\xi}{p_J x_{iJ}} = \frac{1}{\sigma_J - 1}.$$

In this case, δ_{iJ} varies by J but not by spending shares. The total consumer surplus associated with J is given by $\delta_{iJ} \Omega_{iJ}$.

Due to the near-ubiquitous use of the CES demand system, “love-of-variety” is sometimes conflated with the price elasticity of demand. However, as pointed out by Dixit and Stiglitz (1977), outside of the expanding-variety CES model, these two statistics are not the same. In fact, under a plausible condition (Marshall’s second law of demand), the surplus produced by new varieties is strictly lower than that implied by the CES demand system.

Example 3 (Expanding Varieties with Marshall’s Second Law). Denote the own-price elasticity of i ’s demand for each input of type J by $\sigma_{iJ}(\mathbf{p}) = -\frac{\partial \log x_{iJ}}{\partial \log p_J} > 1$. Marshall’s second law of demand holds if $\partial \sigma_{iJ} / \partial p_J \geq 0$. Under this condition,

$$\delta_{iJ}(\mathbf{p}) \leq \frac{1}{\sigma_{iJ}(\mathbf{p}) - 1} \tag{5}$$

as long as $\sigma_{iJ}(\mathbf{p}) \geq 1$. The right-hand side of (5) is the consumer surplus ratio implied by a CES demand system calibrated to match the same price elasticity of demand.¹⁸ Hence, under Marshall’s second law of demand and matching a given elasticity of demand, an isoelastic demand curve maximizes the consumer surplus ratio.

In this example, consumer surplus ratios can vary even if all types face the same demand curve (since δ_{iJ} is a function of the price). In the appendix we also show that when Marshall’s second law holds strictly, the consumer surplus ratio is strictly declining in the price of the input: $\frac{\partial \delta_{iJ}}{\partial p_J} < 0$.

Proposition 1 can also be applied to quality-ladder models as the following simplified example shows.

Example 4 (Quality Ladders with CES). Suppose the downstream firm’s marginal cost (suppressing the index i for the downstream firm) is

$$mc = \left(\int_0^1 p(s)^{1-\sigma} ds \right)^{\frac{1}{1-\sigma}},$$

¹⁸The proof of (5) in the appendix builds on ideas from Matsuyama and Ushchev (2020b) and Grossman et al. (2023).

where $p(s)$ is the price of each input variety s . Each input $s \in [0, 1]$ is either supplied at price p_J , or, there are two perfectly substitutable suppliers with prices p_J and p_{J+1} with $p_{J+1} < p_J$. We refer to the difference between p_J and p_{J+1} as the step-size.

Given symmetry of the cost function, we group suppliers into two types indexed by J and $J + 1$. The mass of type J suppliers is 1 and the mass of type $J + 1$ suppliers is $M < 1$. Since J suppliers sell nothing if a $J + 1$ supplier is available, marginal cost can be rewritten as

$$mc = \left((1 - M)p_J^{1-\sigma} + Mp_{J+1}^{1-\sigma} \right)^{\frac{1}{1-\sigma}}.$$

Demand for any input $s \in [0, 1]$ with price p is $x(p) = \left(\frac{p}{mc}\right)^{-\sigma} q$. The demand curve for an individual supplier of type $J + 1$ is $x_{J+1}(p) = x(p)\mathbf{1}\{p \leq p_J\}$. That is, sales of type $J + 1$ suppliers would be zero if their price were to exceed p_J .

Consider a movement along the quality ladder, where some J types are creatively destroyed and replaced by $J + 1$ types. We represent this using an increase in the mass of type $J + 1$ suppliers: $\Delta M > 0$. Applying Proposition 1 yields

$$\Delta \log mc \approx -\Omega_{J+1} \underbrace{\frac{\int_{p_{J+1}}^{\infty} x_{J+1}(\xi) d\xi}{p_{J+1} x_{J+1}}}_{\delta_{J+1}} \Delta M = -\Omega_{J+1} \left(1 - \left(\frac{p_J}{p_{J+1}} \right)^{1-\sigma} \right) \frac{1}{\sigma - 1} \Delta M. \quad (6)$$

In this case, δ_{J+1} reflects the curvature of the demand curve and the step-size, and it is illustrated graphically in Figure 2, as $(B + E)/(C + D)$. In the $\sigma \rightarrow 1$ limit, δ_{J+1} converges to the step size: $\log(p_{J+1}/p_J)$. In the limit where the step size goes to infinity, δ_{J+1} converges to $1/(\sigma - 1)$, as in the expanding variety model in equation (4).¹⁹

Our final example illustrates how movements along a quality-ladder can be represented as-if they are due to simultaneous additions and separations in an expanding variety model. This is useful for interpreting our regressions under both class of models.

Example 5 (Quality ladders as Expanding Varieties). The integral in (6) can be split in two:

$$\Delta \log mc \approx -\Omega_{J+1} \frac{\int_{p_{J+1}}^{p_J} x(\xi) d\xi}{p_{J+1} x_{J+1}} \Delta M = -\Omega_{J+1} \underbrace{\frac{\int_{p_{J+1}}^{\infty} x(\xi) d\xi}{p_{J+1} x_{J+1}}}_{\delta_{J+1}^{EV}} \Delta M + \Omega_J \underbrace{\frac{\int_{p_J}^{\infty} x(\xi) d\xi}{p_J x_J}}_{\delta_J^{EV}} \Delta M, \quad (7)$$

¹⁹A similar logic applies in models with matching frictions as in Miyauchi (2018) and Fontaine et al. (2023), in which downstream firms buy inputs from their lowest-cost matched supplier. The exit of low cost suppliers ($\Delta M < 0$) forces firms to switch to their second lowest cost suppliers, raising marginal cost according to $-\Omega_{J+1} \delta_{J+1} \Delta M$ where $\log(p_{J+1}/p_J) < 0$ is the price gap between the first and second lowest cost matched suppliers. The logic in Example 5 also applies.

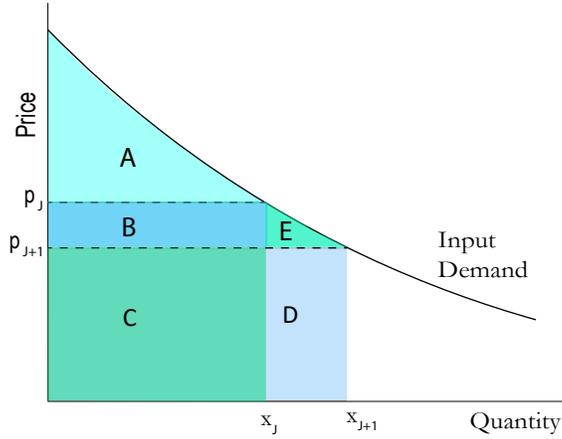


Figure 2: The consumer surplus ratio for type $J + 1$ suppliers is $\delta_{J+1} = \frac{B+E}{C+D}$. The consumer surplus ratio for the input is $\delta_J^{EV} = A/(B + C)$ when the price of the input is p_J and $\delta_{J+1}^{EV} = (A + B + E)/(C + D)$ when the price of the input is p_{J+1} . The decline in marginal cost caused by the addition of an individual supplier of type $J + 1$ is the same as the simultaneous addition of a variety with price p_{J+1} and the separation of a variety with price p_J . In both cases, the change in marginal cost is $(B + E)\Delta M$.

where x is the demand for an input variety rather than an individual supplier, and we have assumed that δ_J^{EV} and δ_{J+1}^{EV} are finite. For this CES example, this requires that $\sigma > 1$, in which case, the consumer surplus ratio associated with varieties, $\delta_J^{EV} = \delta_{J+1}^{EV} = 1/(\sigma - 1)$, is greater than the consumer surplus ratio associated with individual suppliers, δ_{J+1} in (6). The consumer surplus ratio of the variety, δ^{EV} , is higher than the one for the individual supplier, δ_{J+1} , since if $J + 1$ is unavailable, the input can still be purchased at price p_J . The surplus ratio is graphically depicted in Figure 2 as $\delta^{EV} = A/(B + C) = (A + B + E)/(C + D)$ and $\delta_{J+1} = (B + E)/(C + D)$.

In words, the change in marginal cost due to a movement along the quality ladder is equivalent to the addition of a variety with price p_{J+1} and the separation of a variety with price p_J . The difference in consumer surplus from the gain and loss of the two varieties equals the consumer surplus generated by the movement along a quality ladder.

For the remainder of the paper, we represent movements along quality ladders using simultaneous additions and separations in an expanding variety model.²⁰

²⁰This means that the specification in Section 3 estimates consumer surplus ratios associated with varieties rather than individual suppliers on a quality ladder. In an expanding variety model, these are the same, but in a quality-ladder model they differ (δ^{EV} versus δ_{J+1} in the example). Consumer surplus from movements along a quality ladder are then given by the difference in consumer surplus for added relative to separating varieties (i.e. expenditures on added suppliers times δ^{EV} for added suppliers minus expenditures on separating suppliers times δ^{EV} for separating suppliers). This observation implies that our growth accounting results in Section 5 apply to both classes of models. We also consider alternative specifications that estimate

3 Microeconomic Results: Empirics

In this section, we consider regressions aimed at identifying the consumer surplus ratio associated with gaining and losing access to suppliers. We begin by describing our specification, which is motivated by Proposition 1. We then describe the instruments, discuss the identification assumptions, and describe our data. We end the section with our regression results.

3.1 From Theory to Baseline Regression

To begin with, we index variables by time, t . Define

$$\Delta M_{ij,t}^{add} = \sum_{j \in J} \mathbf{1}(\Omega_{iJ,t+1} > 0) \mathbf{1}(p_{ij,t} = \infty) \mathbf{1}(p_{ij,t+1} = p_{iJ,t+1}) \quad (8)$$

to be the mass of inputs of type J that i did not have access to in period t but does have access to in period $t + 1$. The notation $j \in J$ above means that j is an individual supplier of type J . Similarly, define the mass of suppliers that i loses access to be

$$\Delta M_{ij,t}^{sep} = \sum_{j \in J} \mathbf{1}(\Omega_{iJ,t} > 0) \mathbf{1}(p_{ij,t} = p_{iJ,t}) \mathbf{1}(p_{ij,t+1} = \infty). \quad (9)$$

This is the mass of suppliers with positive demand in t whose price goes to infinity at $t + 1$ and are no longer available to i .

In our empirical work, we do not explicitly map suppliers to types but instead estimate average effects over all types. To that end, define the (weighted) average consumer surplus ratio associated with additions and separations as

$$\bar{\delta}_{i,t}^{add} = \sum_J \left(\frac{\Omega_{iJ,t} \Delta M_{ij,t}^{add}}{\sum_{J'} \Omega_{iJ',t} \Delta M_{ij',t}^{add}} \delta_{iJ,t+1} \right), \quad \bar{\delta}_{i,t}^{sep} = \sum_J \left(\frac{\Omega_{iJ,t} \Delta M_{ij,t}^{sep}}{\sum_{J'} \Omega_{iJ',t} \Delta M_{ij',t}^{sep}} \delta_{iJ,t} \right).$$

As an example, if the cost function is CES with elasticity σ , then $\bar{\delta}_{i,t}^{add} = \bar{\delta}_{i,t}^{sep} = 1/(\sigma - 1)$.

Given these definitions, we can rewrite Proposition 1 as

$$\Delta \log mc_{i,t} \approx -\bar{\delta}_{i,t}^{add} \sum_J \Omega_{iJ,t+1} \Delta M_{ij,t}^{add} + \bar{\delta}_{i,t}^{sep} \sum_J \Omega_{iJ,t} \Delta M_{ij,t}^{sep} + \sum_J \Omega_{iJ,t} M_{ij,t} \Delta \log p_{J,t} + \Delta \log A_{i,t}, \quad (10)$$

δ for individual suppliers, rather than for varieties, if we commit to a quality-ladder model (see Footnote 32).

where, without loss, we normalize the elasticity of marginal cost with respect to the productivity parameter to be one and we ignore higher order terms.

The first two terms on the right-hand side of (10) capture the effect of gaining and losing access to varieties. In the expression above, the per-variety expenditure share of added suppliers is measured at $t + 1$, whereas the per-variety expenditure share of separating suppliers is measured at t . Since we work with a first-order approximation, using elasticities before the shock, at t , or after the shock, at $t + 1$, are both valid. We use the expenditure share of added suppliers in $t + 1$ because the type-specific expenditure share, $\Omega_{ij,t}$, for a variety that is added in $t + 1$ is not known in t (unless one maps a supplier added at $t + 1$ to suppliers who are of the same type at t). Similarly, we use the expenditure share of separating suppliers in t because the type-specific expenditure share, $\Omega_{ij,t+1}$, of a variety that separates in t is not known in $t + 1$.

We wish to use a regression to identify the average consumer surplus ratios $\bar{\delta}^{add}$ and $\bar{\delta}^{sep}$ in (10). Unfortunately, we cannot perfectly observe any of the right-hand variables. The potential confounders in (10) are the changes in the prices of continuing suppliers, $\sum_K \Omega_{iK,t} M_{iK,t} \Delta \log p_{K,t}$, and own technology shocks, $\Delta \log A_{i,t}$. Since we do not observe all continuing input price changes and technology shocks, a simple regression can suffer from omitted variable bias.

More subtly, we also may not be directly observing the addition and separation regressors in (10). In the data, we observe overall additions and separations of suppliers. In principle, we do not know if these additions and separations are due to movements along the input demand curve, as in (8) and (9), or due to shifts of the input demand curve.

As explained in Section 2, additions and separations that happen smoothly due to shifts of the input demand curve, without a jump in expenditure shares, do not affect either the marginal cost of the downstream firm or the addition and separation shares (since expenditure shares on these suppliers are zero to a first-order).

However, we might worry that some additions and separations are caused by discontinuous shifts of the input demand curve rather than by price jumps. To allow for this possibility, in this section we enrich the model to allow for discontinuous jumps in input demand curves due to downstream technology shocks that are biased towards specific inputs. Additions and separations caused by these input demand shocks affect the addition and separation shares but have no independent first order effect on the downstream firm's marginal cost (beyond the direct effect of the shock that caused the demand curve to shift in the first place). That is, when biased discrete downstream technology shocks are present, equation (10) is unaltered except that the direct effect of these shocks is included in the technology term $\Delta \log A_{i,t}$.

Define the overall addition share

$$\Delta \tilde{M}_{ij,t}^{add} = \sum_{j \in J} \mathbf{1}(\Omega_{ij,t+1} > 0) \mathbf{1}(\Omega_{ij,t} = 0)$$

and the overall separation share

$$\Delta \tilde{M}_{ij,t}^{sep} = \sum_{j \in J} \mathbf{1}(\Omega_{ij,t} > 0) \mathbf{1}(\Omega_{ij,t+1} = 0)$$

to be the measure of suppliers that i adds and separates from between t and $t + 1$. Unlike (8) and (9), the overall addition and separation shares are directly observable. However, due to the possibility that some separations and additions may be caused by biased downstream shocks, $\Delta \tilde{M}_{ij,t}^{add}$ and $\Delta \tilde{M}_{ij,t}^{sep}$ are not necessarily equal to $\Delta M_{ij,t}^{add}$ and $\Delta M_{ij,t}^{sep}$. The difference is additions and separations caused by discontinuous shifts of the input demand curve.

We consider a regression of the form

$$\Delta \log mc_{i,t} = -\hat{\delta}^{add} \underbrace{\sum_J \Omega_{ij,t+1} \Delta \tilde{M}_{ij,t}^{add}}_{\text{addition share}_{i,t}} + \hat{\delta}^{sep} \underbrace{\sum_J \Omega_{ij,t} \Delta \tilde{M}_{ij,t}^{sep}}_{\text{separation share}_{i,t}} + \hat{\gamma}' W_{i,t} + \varepsilon_{i,t}, \quad (11)$$

where $W_{i,t}$ are controls. The error term contains the same potential confounds as (10), the additional terms associated with $\Delta \tilde{M}_{ij,t}^{add} - \Delta M_{ij,t}^{add}$ and $\Delta \tilde{M}_{ij,t}^{sep} - \Delta M_{ij,t}^{sep}$, errors from the first order approximation, and measurement error.

To overcome the identification challenges, we use an instrumental variables strategy that we describe in the next section.

3.2 Identification Strategy

In this section we describe our identification strategy and our instruments. Since we have two regressors, we need two instruments. We instrument for separations and additions using a subset of firm deaths and births. Let $S_{j,t}$ be the sales of supplier firm j in period t . For each Prodcum firm i in year t , our first instrument is

$$Z_{i,t}^{death} = \sum_j \Omega_{ij,t} \mathbf{1}(S_{j,t+1} = 0) \mathbf{1}(p_{j,t} x_{ij,t} / S_{j,t} < \text{cutoff}). \quad (12)$$

In words, we add up the expenditure share relative to variable costs, $\Omega_{ij,t}$, of i 's suppliers that exit the market ("die") between t and $t + 1$ and for whom i is a small customer in the

sense that i 's purchases from j as a fraction of j 's total sales are lower than some cutoff (in our benchmark results, 5%). We call this the *restricted death* instrument.

Our second instrument is

$$Z_{i,t}^{birth} = \sum_j \Omega_{ij,t+1} \mathbf{1}(S_{j,t} = 0) \mathbf{1}(p_{j,t+1} x_{ij,t+1} / S_{j,t+1} < \text{cutoff}). \quad (13)$$

In words, we add up the expenditure share relative to variable costs, $\Omega_{ij,t+1}$, of i 's suppliers that enter the market (are "born") between t and $t + 1$ and for whom i is a small customer (in our benchmark results, less than 5% of j 's sales). We call this the *restricted birth* instrument.

The following proposition formalizes our identification strategy.

Proposition 2 (Identification). *Consider the regression in (11). Suppose that, conditional on the controls $W_{i,t}$, the instruments are mutually independent of the error term in the first and second stage as well as $\bar{\delta}_{i,t}^{add}$ and $\bar{\delta}_{i,t}^{sep}$. Then the estimates $\hat{\delta}^{add}$ and $\hat{\delta}^{sep}$ consistently estimate $\mathbb{E}[\bar{\delta}_{i,t}^{add}]$ and $\mathbb{E}[\bar{\delta}_{i,t}^{sep}]$, where the expectation is over downstream firms and periods.*

Our instruments isolate churn due to births and deaths of suppliers. This is to ensure that those additions and separations reflect a movement along the input demand curve rather than a shift of the input demand curve. That is, if a supplier separates because it ceased operations or a supplier is added because it began operations, the price of the inputs the supplier provides must be jumping from infinity to finite values (for additions) or vice versa (for separations).

Although the birth or death of a supplier causes its price to jump, there is no guarantee that this price jump is uncorrelated with idiosyncratic shocks to the downstream firm. For example, a supplier may cease or begin operations because its main client received a technology shock. The requirement that the downstream firm be a small customer for the supplier is to ensure that idiosyncratic shocks to the downstream firm do not cause the upstream firm to enter or exit the marketplace.²¹

We control for prices of continuing suppliers (if we observe them) and industry by time fixed effects (defined as the product code of each firm's best selling product). Prices of continuing suppliers and industry by year fixed effects control for the possibility that suppliers' decisions to exit or enter the market may be caused by shocks to competitors.

²¹Even if downstream productivity shocks are independent of restricted births, the firm's adoption or link formation decision may be correlated with own productivity shocks. In this case, firm births would predict adoption not only of newly-born suppliers but also of pre-existing suppliers. However, restricted births do not predict additions of non-newly-entering suppliers (see the second row of columns (iii) and (iv) of Table A9).

Since the average consumer surplus ratio for additions and separations for each downstream firm is allowed to be itself a random variable, we require somewhat stronger assumptions than the typical linear IV regression (i.e. uncorrelatedness of the instrument with the error in the second stage). Specifically, we require that, conditional on the controls, our two instruments be independent of the errors in the second stage (e.g. own-technology shocks, changes in unobserved prices of competing inputs, and additions and separations that are due to shifts of the demand curve) and the first stage, as well as the random variables $\bar{\delta}^{add}$ and $\bar{\delta}^{sep}$ (this is automatic if δ is a constant, as in CES). We do not assume that δ_{ij} is unrelated to expenditure shares within types, Ω_{ij} . Instead, we require that the average δ for the downstream firm be independent of total (rather than per variety) expenditures on restricted births and deaths. Appendix F provides Monte Carlo simulations showing that even when $\bar{\delta}_{i,t}^{add}$ and $\bar{\delta}_{i,t}^{sep}$ are correlated with the instruments, the bias in our estimates is quite small.

3.3 Data

In this section we describe how we map our model to data and how we construct the terms in the baseline regression (11). Our empirical analysis makes use of a rich micro-level data structure on Belgian firms in the period 2002-2018. The data structure brings together information drawn from six comprehensive panel-level data sets: (i) the National Bank of Belgium’s (NBB) Central Balance Sheet Office (CBSO), which we refer to as the annual accounts; (ii) the Belgian Prodcom Survey, which covers firms that produce goods covered by the Prodcom classification and that have at least 20 employees or 5 million euros turnover in the previous reference year; (iii) the NBB Business-to-Business (B2B) Transactions data; (iv) International Trade data at the NBB; (v) VAT returns; and (vi) the Crossroads Bank of Enterprises (CBE) which we use to identify mergers and acquisitions.²² Additional details are provided in Appendix C.

Downstream firms. Our sample of downstream firms are firms in the Prodcom survey, where we observe data on quantities sold (which we use to construct marginal costs). We restrict the sample to non-financial corporations that file the annual accounts. To ensure that Prodcom variables are representative of a firm’s overall activities, we restrict the sample to those firms whose Prodcom sales are at least 30% of the firm’s total sales.²³ Our

²²See https://www.nbb.be/doc/dq/e_method/gni_methodological_inventory_belgium_version_2022_publication.pdf for a description of the annual accounts (page 589), VAT returns (page 589), and Prodcom (page 603) data sets.

²³Total sales may differ from Prodcom sales because, for example, firms sell products that they do not produce (Bernard et al. 2019) or they sell services along with the goods they produce (Ariu et al. 2020). The

micro sample contains roughly between 2,000 and 4,000 downstream firms per year. We now describe how we measure a number of key variables for these downstream firms.

Sales and value-added. We obtain value added from the annual accounts, which is used to construct the National Income and Product Accounts in Belgium.²⁴ We define firms' total sales as the highest value between sales reported in the annual accounts (reported mainly by large firms) and sales reported in the VAT returns. We replace this measure of sales by the sum of exports reported in the international trade data set and sales to other Belgian firms reported in the B2B VAT data set if the latter exceeds the former. We drop observations where value added exceeds sales.

Total variable costs. Firms' input costs consist of purchases of intermediates, labor costs, and the user cost of capital. Intermediate input purchases are defined to be sales minus value added, measured as defined above. Labor costs are reported in the annual accounts. The cost of capital is defined as the product of the capital stock reported by firms in the annual accounts (which includes plants, property, equipment, and intellectual property) and an industry-specific user cost of capital. The latter is the sum of a risk premium (set as 5 percent) and the risk-free real rate (defined as the corresponding governmental 10 year-bonds nominal rate minus consumer price inflation at that time period) minus the industry-level depreciation rate, $(1 - d) \times g$, where d is the industry level depreciation rate (defined as consumption of fixed capital as a ratio of net capital stock) and g is the expected growth of the relative price of capital at the industry level (defined as the growth in the relative price of capital computed from the industry-specific investment price index relative to the consumer prices index in each year).

We assume that intermediate inputs are fully variable, but allow a fraction of labor and capital costs to be overhead. Denote by ϕ the fraction of labor and capital costs that are variable with, $1 - \phi$, being overhead costs. To calibrate ϕ , we follow a similar strategy to Dhyne et al. (2022). We regress the change in labor and capital costs on the change in intermediate costs (which we assume are fully variable) instrumented using a demand shock. We set $\phi = 0.5$ because our estimates indicate that labor and capital costs rise by roughly 0.5 percent when intermediate purchases rise by 1 percent in response to a demand shock. (See Appendix C for more details). Our estimate of ϕ is similar to that

ratio of Prodcom sales to total sales is 0.89 for the median firm in our sample.

²⁴Page 81 in https://www.nbb.be/doc/dq/e_method/gni_methodological_inventory_belgium_version_2022_publication.pdf states that the annual accounts are the preferred source for estimating aggregates of the production and primary distribution of income account of non-financial corporations. The empirical results are similar if we measure sales using values reported in the annual accounts and, if the latter is missing, using values reported in the VAT returns.

found by Dhyne et al. (2022).²⁵

Given uncertainty over the extent of overhead costs, in Appendix D we report our results under alternative assumptions. First, we set $\phi = 0.4$. Second, we set $\phi = 0.6$. Third, we assume that capital costs are all overhead and keep $\phi = 0.5$ for labor costs. Fourth, we abstract from overhead costs all together, setting $\phi = 1$. The results are quite robust to the value of ϕ .

Prodcom quantities and unit values. We construct changes in output quantities and unit values for the sample of firms in the Prodcom survey. Products are identified at the 8-digit level of the Prodcom product code (PC) classification, which is common to all EU member states.²⁶ Sales values (in euros) and quantities are available at the firm-PC8-month level. Quantities are reported in one of several measurement units (over two thirds of observation are in kilograms; other units include liters, meters, square meters, kilowatt, and kg of active substance). We aggregate monthly observations to yearly values to match the other data sets, and calculate log differences in quantities and unit values by PC8 product from year t to $t + 1$. As quantities and unit values can be noisy, we trim changes in these two variables at the 5-95th percentile level.

For multi-product firms (defined as Prodcom firms that produce multiple PC8 products), we aggregate changes in quantities of individual products to the firm-level using a Divisia index, with weights given by the firm’s sales share of each product in the corresponding year. This quantity index is valid if we assume that demand for multi-product firms in Prodcom is homothetic. In this case, a Divisia index reliably aggregates multiple products into a single product bundle. For each firm, we also construct changes in unit values as log changes in Prodcom sales minus the Divisia quantity index.²⁷ We assign a product code to each firm according to its highest-selling product.

Marginal cost. For each firm in the Prodcom survey, we calculate the log change in marginal cost as

$$\Delta \log mc = \Delta \log \text{total variable costs} - \Delta \log \text{total quantity}, \quad (14)$$

which is valid as long as the scale elasticity of the variable cost function is constant. Unfortunately, we only observe changes in Prodcom quantities and not changes in total quanti-

²⁵Dhyne et al. (2022) also show that ϕ is not correlated with firm size, which is consistent with our approach.

²⁶As product codes tend to vary from year to year, we use the correspondence of 8-digit products in the Prodcom classifications that trace products over time used by Duprez and Magerman (2018).

²⁷We obtain very similar results if we calculate changes in unit values as a Divisia index (sales-weighted) of changes in unit values by product rather than deflating sales by the quantity Divisia index.

ties. To address this, we use the following identity:

$$\Delta \log \text{total quantity} = \Delta \log \text{Prodcom quantity} + \Delta \log \frac{\text{total sales}}{\text{Prodcom sales}} + \text{error},$$

where the error term is unobserved. Since sales are equal to the product of quantity and unit values, the error term is equal to the difference in log changes of average unit values between Prodcom and non-Prodcom sales of each firm. We use this equation to impute the log change in total quantity, setting the unobserved error term to zero, which we then use in (14). This imputation is innocuous as long as the unobserved error term is uncorrelated with our instruments.²⁸

Having described how we construct the left-hand side variable in (11), we now discuss how we construct the right-hand side variables. Constructing the right-hand side variables requires knowing the input shares of the downstream firms.

Intermediate input shares. We construct input shares of Prodcom firms using the confidential NBB B2B Transactions data set. At the end of every calendar year, all VAT-liable firms in Belgium are required to file a complete listing of their Belgian VAT-liable customers over that year. An observation in this data set refers to the sales value in euros of enterprise j selling to enterprise i within Belgium, excluding the VAT amount due on these sales. The reported value is the sum of invoices from j to i in a given calendar year. As every firm in Belgium is required to report VAT on all sales of at least 250 euros, the data has nearly universal coverage of all businesses active in Belgium. To control for misreporting errors, we drop a transaction if its value is higher than the seller's aggregate sales and higher than the buyer's total intermediate input purchases (which is reported separately). Since we are interested in variable inputs, we exclude suppliers that produce capital goods, identified from the Main Industrial Groupings (MIG) Classification of the EU (we report sensitivity to including these suppliers in the network). We also drop suppliers with unknown VAT numbers or that are part of the downstream firm (due to mergers and acquisitions).

Separation and addition share. For each Prodcom firm i and period t , using the B2B data, we identify the set of separating suppliers as those the firm buys from in t but does not buy from in $t + 1$. Similarly, the set of added suppliers are those that i does not buy

²⁸We also obtain similar results if we measure changes in marginal costs as log changes in Prodcom unit values minus log changes in markups, where markups are estimated following De Loecker and Warzynski (2012) with production function estimates as in Levinsohn and Petrin (2003) (see Table A3).

from in t but does buy from in $t + 1$. We calculate the separation share $_{i,t}$ as the ratio of purchases of i from separating suppliers relative to variable costs at t . We calculate the addition share $_{i,t}$ as the ratio of purchases of i from added suppliers relative to variable costs at $t + 1$. In our regressions we drop observations with separation or addition shares higher than 0.5, and perform sensitivity analysis to this cutoff.²⁹

Restricted death and birth shares. We construct the instruments defined in (12) and (13) by calculating for each downstream firm separation and addition shares over a restricted set of suppliers: those that exit or enter the market (firm deaths and births) and for whom the downstream firm accounts for less than 5% of the suppliers sales.³⁰ We perform extensive sensitivity analysis to the value of this cutoff in Table 2.

Controls. In our regressions, we control for changes in the other components of marginal cost to the extent possible. For continuing upstream suppliers that happen to belong to Prodcom, we construct and control for the change in the unit values (see Duprez and Magerman, 2018 and Cherchye et al., 2021). We also measure and control for the price of labor by dividing total labor costs by total full time employed workers. We measure and control for the price of capital services via the user cost of capital as described above. We measure and control for changes in unit values of imported inputs using a firm-level Divisia index of changes in unit values faced by firm i at the CN8 product level, trimming changes in unit values at the 5th-95th percentile. We also construct, for each Prodcom firm, a price index of general input costs using industry-level price indices from Eurostat, with weights given by the firm's industry shares in non-Prodcom input purchases.

Summary statistics. Table A5 in Appendix E reports summary statistics for our Prodcom sample on the share of factors and intermediate inputs in variable costs, the number of suppliers, separation and addition share, and restricted death and birth shares. The average firm has 227 suppliers, and the average addition and separation share are around 6%.

²⁹Our data is annual, so the separation and addition share depend on the specific month that a supplier is added or subtracted. For example, a supplier that is dropped in the middle of the year contributes less to the separation share than a supplier that is dropped at the end of the year. However, because our measure of marginal cost is also based on annual data, the increase in marginal cost is also smaller if the supplier exits in the middle of the year than if the supplier exits at the end. This means that, up to the first order approximation, our estimates are not contaminated by the fact that suppliers may enter and exit at different points in time during the year.

³⁰In practice, a few suppliers that entered in $t + 1$ were active in $t - 1$ or earlier and a few of the suppliers that exited in t re-entered in $t + 2$ or later. These suppliers represent less than 1% of all restricted births and deaths (and less than 0.5% in value terms). Our results are almost unchanged if we construct the instrument excluding such suppliers.

The restricted birth and death shares are much smaller, averaging around 0.2%.³¹

Table A6 in Appendix E reports correlations of firm size (employment and sales) with the number of suppliers, additions and separations, and our instruments. Larger firms are connected to a higher number of suppliers. We also find that additions and separations are slightly negatively correlated with the size of the downstream firms (the addition and separation shares are lower for larger firms). Importantly, our instruments are not correlated with downstream firm size. This suggests that our instruments do not differentially cause exogenous variation in additions and separations for large versus small downstream firms.

3.4 Results

Having discussed how the terms in the baseline regression (11) are constructed, we now turn our attention to the results.

The baseline results are shown in Table 1. Column (i) is an OLS regression of the overall addition and separation shares on the change in marginal costs with all controls. We find that both separations and additions have small and insignificant effects on marginal costs. Of course, there is good reason to expect that this regression does not have a causal interpretation due to omitted variable bias. Column (ii) is a reduced-form regression of changes in marginal cost directly on the instruments. As expected, an increase in the restricted death share raises marginal costs and an increase in the restricted birth share lowers marginal costs.

Columns (iii) and (iv) are the first stage regressions showing that restricted death primarily predicts separations and restricted birth primarily predicts additions. However, restricted birth also has an effect on separations (for example, due to creative destruction). Similarly, restricted death has an effect on additions (due to replacements of the exiting suppliers). Table A9 shows that restricted births positively predict separations from suppliers that continue to operate. Similarly, restricted supplier deaths positively predict additions of suppliers that previously operated. Hence, our instruments do not solely affect additions of newly-born and separations from dying suppliers. Instead, restricted births predict separations from continuing suppliers and restricted deaths predict additions from continuing suppliers.

Columns (v) and (vi) are the IV regressions without and with additional controls. The point estimates are quite insensitive to the inclusion of the controls. In the specification

³¹The average duration of continuous relationships between firms is 2.6 years (without any adjustments for censoring), and this distribution is heavily skewed to the right. The average duration of continuous relationships triggered by our instruments is slightly higher and close to 3 years.

Table 1: Baseline estimates of δ

	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)
	$\Delta \log mc$		First stage		$\Delta \log mc$		$\Delta \log p$
			Separat.	Addit.			
Separation share	-0.013 (0.013)				0.279*** (0.090)	0.268*** (0.091)	0.163** (0.076)
Addition share	0.016 (0.012)				-0.280*** (0.079)	-0.283*** (0.079)	-0.177*** (0.063)
Restricted death share		0.199** (0.083)	1.047*** (0.052)	0.290*** (0.061)			
Restricted birth share		-0.230*** (0.075)	0.377*** (0.068)	1.169*** (0.052)			
Specification	OLS	OLS	OLS	OLS	IV	IV	IV
F-stat			283	295	115	111	111
Controls	Y	Y	Y	Y	N	Y	Y
Industry \times year FE	Y	Y	Y	Y	Y	Y	Y
Observations	38,670	38,670	38,670	38,670	38,670	38,670	38,670

Notes: Columns (i), (v), and (vi) report estimates of regression (11), column (ii) is the reduced form, columns (iii) and (iv) show the first stage, and column (vii) uses changes in unit values instead of marginal cost. Restricted exit share and restricted entry share are the instruments, $Z_{i,t}^{death}$ and $Z_{i,t}^{birth}$, defined by equations (12) and (13). Controls are log changes in the price of imported inputs, log changes in the price of inputs purchased from other Prodcom firms and from other industries, changes in log wages, and changes in the log user cost of capital. Regressions are unweighted. Industry-by-time fixed effects are at the 6-digit product code level. Standard errors are clustered at the firm-level. The F-stats for the first-stage and second-stage are the Sanderson-Windmeijer (SW) and the Kleibergen-Paap (KP) statistics.

without controls, a 1% increase in the separating share raises marginal costs by around 0.28%. Similarly, a 1% increase in the addition share lowers marginal costs by around 0.28%. With controls, the point estimates for additions is slightly bigger in magnitude than the one for separations, but we cannot reject that they are equal.^{32,33}

³²Table A10 in the appendix reports results from a specification of (11) where we regress changes in marginal cost on separations and additions separately. This specification is useful to identify the consumer surplus ratio of moving along a quality ladder. If supplier births are not associated with separations, and if supplier deaths are not associated with additions, as in a pure expanding varieties model, then the joint regression and the univariate regressions give the same estimates for $\bar{\delta}^{sep}$ and $\bar{\delta}^{add}$. However, as shown in columns (iii) and (iv) of Table 1, restricted death does have a small effect on additions and restricted birth does have a small effect on separations. The point estimates in the univariate regression are somewhat smaller in magnitude (0.19 rather than 0.28). This is consistent with a quality-ladder model as in Examples 4 and 5. The univariate regressions identify $\frac{1}{\sigma-1} ((p_{J+1}/p_J)^{1-\sigma} - 1)$ (in the separation regression) and $\frac{1}{\sigma-1} (1 - (p_J/p_{J+1})^{1-\sigma})$ (in the addition regression) where p_{J+1}/p_J is the step-size in the quality-ladder (see equation (6) in Example 4). These are necessarily smaller than $1/(\sigma - 1)$, which is what the joint regression estimates (see equation (7) in Example 5).

³³Under CES demand, $\delta = 1/(\sigma - 1)$, so a point estimate of 0.28 implies that the price elasticity of demand

The final column, (vii), replaces the change in marginal cost on the left-hand side with the change in the price (unit value) charged by the downstream firm. The effect of separations and additions on the price is smaller in magnitude than on marginal cost. The reduced-form pass-through of marginal cost into prices implied by this regression is about 60%. This is very close to the pass-through estimates from Amiti et al. (2019), who also study Prodcos firms in Belgium but use a very different identification strategy.

Table 2 displays the results of the IV regression for different cut-off values of what constitutes a “small” customer in (12) and (13). The benchmark results in Table 1 use 5%. Table 2 shows that our results are reasonably robust to this choice and the point estimates remain between 0.20 and 0.28 for both additions and separations as long as the cut-off value is not too high (less than 15%).

The point estimates do start to change if the cut-off value becomes too large, however. Column (xi) shows the results if we use unconditional entry and exit of suppliers as instruments. The point estimates are very different in this case, where we include birth and death of suppliers who are heavily reliant on the downstream firm for their sales. In this case, shocks to the downstream firm can be responsible for supplier entry and exit, confounding our point estimates. The final column, column (xii), uses all separations and additions below a 5% cut-off, rather than separations and additions associated with birth and death of suppliers, as instruments. The point estimates are both zero — again, this reflects the fact that supplier addition and separation can be endogenous to other shocks that hit the downstream firm and shifts of the input demand curve, even if the downstream firm is small as a share of those suppliers’ overall sales.

Table 3 reports estimates of regression (11) where we replace $\Delta \log(mc_{t+1}/mc_t)$ by $\Delta \log(mc_{t+s}/mc_t)$ for different values of s . Column (iv) is our benchmark specification, which sets $s = 1$. Columns (i)-(iii) are placebo tests based on changes in marginal cost to future instrumented additions and separations. The estimates are small and insignificant, ruling out pre-trends in marginal cost changes that are correlated with our instruments.

Columns (v) and (vi) in Table 3 consider two and three year cumulative changes in marginal costs. We find that the effects are persistent, without evidence of mean reversion. Point estimates are slightly larger over longer horizons, but the difference is not statistically significant. The fact that our estimates are not significantly larger at longer horizons suggests that inventories do not play a large role in helping the downstream

is roughly 4.5. If we commit to CES demand, then following the logic of Feenstra (1994), we can run an alternative regression $\Delta \log mc_{it} = \hat{\beta} \times \Delta \log \text{continuing share}_{it} + \text{controls}_{it} + \varepsilon_{it}$, where $\hat{\beta}$ estimates $\delta = 1/(\sigma - 1)$. The results of this regression, using the same IV strategy, are reported in Table A11. We find $\hat{\beta} \approx 0.265$, which implies $\sigma \approx 4.8$. Note that this is not the way Feenstra (1994) identifies δ since Feenstra (1994) estimates σ from expenditure-switching and then imposes that $\delta = 1/(\sigma - 1)$.

Table 2: Sensitivity of point estimate of δ to cut-off for small customer

	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)	(viii)	(ix)	(x)	(xi)	(xii)
	$\Delta \log mc$											
Separation share	0.280*** (0.108)	0.282*** (0.093)	0.268*** (0.091)	0.275*** (0.086)	0.277*** (0.083)	0.265*** (0.080)	0.241*** (0.076)	0.205*** (0.073)	0.165** (0.069)	0.070 (0.051)	0.077* (0.040)	0.005 (0.017)
Addition share	-0.221** (0.093)	-0.230*** (0.080)	-0.283*** (0.079)	-0.269*** (0.076)	-0.258*** (0.072)	-0.241*** (0.067)	-0.229*** (0.065)	-0.209*** (0.064)	-0.190*** (0.056)	-0.037 (0.046)	0.039 (0.039)	-0.007 (0.013)
Specification	IV	IV	IV	IV	IV	IV	IV	IV	IV	IV	IV	IV
F-stat	67	98	111	125	136	149	156	175	180	371	916	21,772
Controls	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Industry \times year FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Cutoff	3	4	5	6	7	8	9	10	15	50	100	5
Suppliers	D&B	D&B	D&B	D&B	D&B	D&B	D&B	D&B	D&B	D&B	D&B	All

Notes: Columns (i)-(xi) re-run the benchmark regression, column (vi) in Table 1, but vary the cut-off value from 3% to 100% for what constitutes a small customer for exiting and entering suppliers when defining restricted deaths and births (labeled D&B in the table). The benchmark regressions use a value of 5%. Column (xii) uses all separations and additions with a 5% cutoff, rather than only D&B. The number of observations is 38,670 in all regressions.

Table 3: Pre-trends and persistence of marginal cost changes

	(i)	(ii)	(iii)	(iv)	(v)	(vi)
	$t - 3$	$t - 2$	$t - 1$	$t + 1$	$t + 2$	$t + 3$
Separation share	-0.049 (0.323)	-0.103 (0.290)	-0.007 (0.139)	0.268*** (0.091)	0.335*** (0.116)	0.375** (0.162)
Addition share	0.100 (0.196)	-0.104 (0.150)	-0.029 (0.090)	-0.283*** (0.079)	-0.313*** (0.115)	-0.447*** (0.132)
Specification	IV	IV	IV	IV	IV	IV
F-stat	32	51	77	111	92	77
Controls	Y	Y	Y	Y	Y	Y
Industry \times year FE	Y	Y	Y	Y	Y	Y
Observ.	21,931	26,410	31,999	38,670	32,052	26,502

Notes: Columns (i)-(vi) report estimates of regression (11), with $\Delta \log(mc_{t+1}/mc_t)$ replaced by $\Delta \log(mc_{t+s}/mc_t)$ for $s = \{-3, -2, -1, 1, 2, 3\}$. Column (iv), with $s = 1$, is our benchmark. Other controls are as in Table 1. Results are similar if we do not include the other controls. All regressions are unweighted. Standard errors are clustered at the firm-level, and F-stat is the Kleibergen-Paap (KP) statistic.

firm smooth shocks in the short run. The fact that our estimates do not shrink at longer horizons suggests that downstream firms are unable to find better replacements over time when they lose suppliers. Similarly, the benefits of new suppliers to the downstream firm do not disappear over time (e.g. due to expiring discounts for new customers).³⁴

Other sensitivity analyses. In Appendix D, we consider an extensive set of robustness checks. Table A2 presents results with more and less disaggregated industry-by-year fixed effects, dropping industry-by-year fixed effects, and including a firm fixed effect. Table A3 considers alternative measures of marginal cost. Table A4 varies the sample of firms and other choices, like weighting schemes and treatment of outliers. Our results are fairly robust across the board.

Table 4: Heterogeneity of δ by supplier characteristics

	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)	(viii)	(ix)
	Zero slope	Relative sales of supplier	Relative value-added of supplier	Cost share	Z_{ijt} Material share	Distance btw. firms	Age of relationship	Relative empl. of customer	Relative sales of customer
Intercept $\bar{\delta}_0$	0.278*** (0.074)	0.505*** (0.114)	0.324*** (0.076)	0.372*** (0.110)	0.363*** (0.107)	0.423* (0.219)	0.289*** (0.104)	0.246*** (0.079)	0.192** (0.084)
Slope $\bar{\delta}_1$	0	-0.121*** (0.042)	-0.084* (0.047)	-0.921 (0.620)	-0.700 (0.474)	-0.032 (0.056)	-0.002 (0.015)	-0.035 (0.061)	-0.065 (0.049)
Specification	IV	IV	IV	IV	IV	IV	IV	IV	IV
F-stat	234	86	52	115	141	90	105	66	57
Controls	Y	Y	Y	Y	Y	Y	Y	Y	Y
Industry \times year FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Observ.	38,670	38,670	38,670	38,670	38,670	33,233	38,670	38,670	38,670

Notes: This table displays estimates of regression (11) replacing the first two terms on the right-hand side by $-\sum_J [\delta_0 + \delta_1 Z_{ij,t+1}] \Omega_{ij,t+1} \Delta \tilde{M}_{ij,t}^{add} + \sum_J [\delta_0 + \delta_1 Z_{ij,t}] \Omega_{ij,t} \Delta \tilde{M}_{ij,t}^{sep}$ where we have assumed that $\delta_{ijt}^{sep} = \delta_{ijt}^{add} = \bar{\delta}_0 + \bar{\delta}_1 Z_{ijt}$ for different choices of Z . Relative sales and value-added of the supplier in columns (ii) and (iii) are defined as the log ratio of the sales or value-added of the supplier relative to the average firm in the same 2 digit industry. Cost share and material share, in columns (iii) and (iv), are the supplier's share of total variable cost and materials. Distance between firms in column (vi) is the log geographic distance between the headquarters of the downstream firm and each supplier. The age of the relationship in column (vii) is the number of consecutive years the downstream firm and the supplier have transacted before separation or after addition. Relative employment and relative sales of customer, in columns (viii) and (ix), are the log ratio of the employment or sales of the customer relative to the average firm in the same 2 digit industry. Controls are as in Table 1. All regressions are unweighted. Industry fixed effects are at the 6-digit product code level. Standard errors are clustered at the firm-level, and F-stat is the Kleibergen-Paap (KP) statistic.

Heterogeneity analysis. We end this section by probing for heterogeneity in δ in a reduced-form way. Table 4 reports estimates where we allow δ to vary as a linear function $\delta_{ijt} =$

³⁴At longer time horizons, we continue to find incomplete pass-through of changes in the downstream firm's marginal cost into its price.

$\bar{\delta}_0 + \bar{\delta}_1 Z_{ijt}$ of relationship characteristics Z_{ijt} . In this table, we also impose that the functional form for δ is the same for both additions and separations. Column (i) is our benchmark where δ is assumed to be constant and equal for additions and separations — the estimate is around 0.28, similar to columns (v) and (vi) in Table 1. Columns (ii) and (iii) allow δ to vary as a function of the relative size of the supplying firms in their industries. Since the slope coefficient is negative, this suggests that the consumer surplus ratio is smaller for downstream firms when the upstream supplier is relatively large in their industry. This means that larger suppliers are able to appropriate a larger share of the surplus they generate per Euro of spending.³⁵ Columns (iv) and (v) allow δ to vary as a function of the cost share of the supplier for the downstream firm. The point estimates for the slope coefficients are negative but imprecise and statistically insignificant. Columns (vi) and (vii) allow δ to vary as a function of the log geographic distance between the headquarters and the age (in years) of the relationship between two firms. Columns (viii) and (ix) condition on the size of the downstream firm. These slope coefficients are all statistically insignificant.

Table A12 in the appendix considers how our results change if we exclude subsets of suppliers. This table shows that the effects are strongest when focusing on service-providing suppliers (including wholesale and retail traders who are in service sectors, even though they sell goods). This is expected given that suppliers in the service sector account for the majority of intermediate inputs in our data (see the summary statistics in Table A5). Additionally, our separation and addition instruments do not include imports, which plausibly are a very important source of goods trade for Belgian manufacturers. Table A12 also provides estimates if we include suppliers of capital goods as part of materials. Our benchmark excludes these suppliers because variable cost only includes the user cost of capital not investment. Nevertheless, including capital goods as part of variable costs barely affects our benchmark estimates. Table A12 also provides estimates where we exclude suppliers that are self-employed, government, and financial entities. This slightly lowers the magnitude of our point estimates.

4 Macroeconomic Results: Theory

In the previous section, we show that input suppliers generate a considerable amount of inframarginal surplus per unit of spending for their downstream customers. In this section we develop a growth accounting framework to decompose the fraction of aggregate

³⁵This does not imply that total consumer surplus from larger suppliers is smaller, since total consumer surplus from a supplier is the product of δ and spending on that supplier.

productivity growth that can be accounted for by observed churn in supply chains. The model explicitly accounts for how changes in one firm’s marginal cost, due to additions and separations of suppliers, spill over to that firms’ customers, customers’ customers, and so on.

We discipline our macro growth accounting results using estimates from the micro regressions which, recall, are estimated using only the Prodcom sample of manufacturing firms. However, we apply our growth accounting formulas to a much larger sample of Belgian firms.

We specify minimal structure on the aggregative model and do not fully specify the environment. This is because we take advantage of the fact that endogenous variables, like changes in factor prices, are directly observable and capture whatever resource constraints the economy is subject to.

4.1 Environment

Consider a set of producers, denoted by N , called the *network*. There is a set of *external inputs* denoted by F . An external input is an input used by producers in the network, N , that those producers do not themselves produce. In practice, the set F includes labor, capital, and intermediate inputs purchased from firms not in the network N . The firms in N collectively produce *final outputs*. Final output is the production by firms in N that firms in N do not themselves use. A stylized representation is given in Figure 3 showing the flow of goods and services.

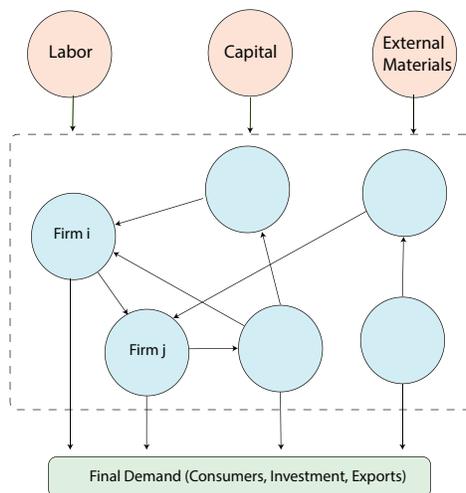


Figure 3: Graphical illustration of the economy. The set N is depicted by the dotted line.

Production. Each producer $i \in N$ has a constant-returns-to-scale production technology in period t given by

$$q_{i,t} = A_{i,t} F_{i,t} \left(\{x_{ij,t}\}_{j \in N}, \{l_{if,t}\}_{f \in F} \right).$$

In the expression above, $l_{if,t}$ is the quantity of external input f and $x_{ij,t}$ is the quantity of intermediate input j used by i at time t . The exogenous parameter $A_{i,t}$ is a technological shifter. There may be fixed overhead costs that must be paid in addition to the variable production technology defined above. We abstract from multi-product firms and associate each firm with a single output.³⁶

After having paid fixed costs, which could include the costs required to access specific inputs, the total variable costs of production paid by firm i are

$$\sum_{j \in N} p_{j,t} x_{ij,t} + \sum_{f \in F} w_{f,t} l_{if,t},$$

where $p_{j,t}$ and $w_{f,t}$ are the prices of internal and external inputs. The markup charged by each producer i , $\mu_{i,t}$, is defined to be the ratio of its price $p_{i,t}$ and its marginal cost of production.

We say that i is *continuing* between t and $t + 1$ if i has positive sales in both t and $t + 1$. Denote by C_t the set of all goods who are continuing at time t .

Resource constraints. We construct a measure of net or final production by the set of continuing, C_t , firms. Let the total quantity of external inputs used by continuing firms be

$$L_{f,t} = \sum_{i \in C_t} l_{if,t} + \sum_{i \in C_t} l_{if,t}^{\text{fixed}},$$

where $l_{if,t}$ is used in variable production and $l_{if,t}^{\text{fixed}}$ are fixed costs. Firm i 's final output is defined to be the quantity of its production that is not sold to other firms in C_t :

$$y_{i,t} = q_{i,t} - \sum_{j \in C_t} x_{ji,t}.$$

That is, final output of good $i \in C_t$, denoted by $y_{i,t}$, is the quantity produced of i that is not used by any $j \in C_t$ and is either consumed by households, used for investment, sold as exports, or sold to other suppliers that are not in the network of continuing producers.

³⁶More precisely, we assume that demand for the products of multi-product firms can be aggregated into a single representative bundle.

Aggregate growth. We measure aggregate growth by deflating nominal final output by a price index. Growth in *real final output* of the set of continuing goods, denoted by $\Delta \log Y_t$, is the change in nominal final output minus the final output price deflator:

$$\Delta \log Y_t = \Delta \log \left(\sum_{i \in C_t} p_{i,t} y_{i,t} \right) - \Delta \log P_t^Y. \quad (15)$$

The change in the *final output price deflator* between t and $t + 1$ is defined to be the share-weighted change in the price of continuing goods

$$\Delta \log P_t^Y = \sum_{i \in C_t} b_{i,t} \Delta \log p_{i,t},$$

where the weights are final output shares

$$b_{i,t} = \frac{p_{i,t} y_{i,t}}{\sum_{j \in C_t} p_{j,t} y_{j,t}}.$$

To calculate growth in real final output between t and $t + T$, we cumulate $\Delta \log Y$:

$$\log Y_{t+T} - \log Y_t = \sum_{s=t}^{t+T} \Delta \log Y_s.$$

Theoretically, this measure of aggregate growth accurately reflects social welfare over continuing goods if final demand for continuing goods is derived from a homothetic aggregator (see Baqaee and Burstein, 2023 for a discussion of the necessary assumptions). Of course, this measure does not capture welfare from entry and exit of goods in final demand.³⁷ Empirically, this measure of aggregate growth is constructed in a way that is similar to how real GDP is constructed. The primary difference is in how we treat external intermediate inputs (e.g. imported intermediate inputs). GDP-style measures subtract the value of imported intermediate inputs from final output. By not subtracting the value of external materials from final output, we treat external materials like factors of production (labor and capital).³⁸ The objective of this section is to decompose the contribution of supplier churn to growth in real final output.

³⁷By focusing on continuing products in final demand, we abstract from the direct welfare effect of entry and exit of goods on final consumers. To quantify these effects, we would need to estimate consumer surplus of new goods for final consumers, which is beyond the scope of this paper.

³⁸If we subtract the value of external materials from final output, then our growth accounting expressions have an additional term involving the difference between expenditures on external materials and the elasticity of aggregate output with respect to external materials. This difference is nonzero in the presence of markups. See Baqaee and Farhi (2024) for more details.

4.2 Theoretical Results

To state our decomposition result, we need to set up some input-output notation. Define the $C_t \times C_t$ cost-based input-output network of continuing firms with ij th element

$$\Omega_{ij,t} = \frac{p_{j,t}x_{ij,t}}{\sum_{k \in C_t} p_{k,t}x_{ik,t} + \sum_{f \in F} w_{f,t}l_{if,t}}.$$

Let Ω^F be the $C_t \times F$ matrix of external input usages, where the if th element is

$$\Omega_{if,t}^F = \frac{w_{f,t}l_{if,t}}{\sum_{k \in C_t} p_{k,t}x_{ik,t} + \sum_{f \in F} w_{f,t}l_{if,t}}.$$

We build on Proposition 1, which is about a single firm, to decompose aggregate growth $d \log Y_t$. To do this, rewrite Proposition 1 for all firms in C_t in matrix notation as

$$\Delta \log \mathbf{p}_t \approx \Delta \log \boldsymbol{\mu}_t - \Delta \log \mathbf{A}_t + \Omega_t \Delta \log \mathbf{p}_t + \Omega_t^F \Delta \log \mathbf{w}_t + \bar{\delta}_t^{\text{sep}} \Delta \mathcal{X}_t - \bar{\delta}_t^{\text{add}} \Delta \mathcal{E}_t,$$

where $\mu_{i,t}$ is the markup of firm i , the ratio of price to marginal cost, $\Delta \mathcal{X}_{i,t} = \sum_J \Omega_{iJ,t} \Delta M_{iJ,t}^{\text{sep}}$ is the cost share of suppliers who separate due to price jumps, and $\Delta \mathcal{E}_{i,t} = \sum_J \Omega_{iJ,t+1} \Delta M_{iJ,t}^{\text{add}}$ is the cost share of suppliers who are added due to price jumps. In the expression above, we normalize the elasticity of the cost function with respect to the productivity shock to be one. Solve out for changes in the prices of continuing firms:

$$\Delta \log \mathbf{p}_t \approx \Psi_t \left[\Delta \log \boldsymbol{\mu}_t - \Delta \log \mathbf{A}_t + \Omega_t^F \Delta \log \mathbf{w}_t + \bar{\delta}_t^{\text{sep}} \Delta \mathcal{X}_t - \bar{\delta}_t^{\text{add}} \Delta \mathcal{E}_t \right], \quad (16)$$

where Ψ_t is the *cost-based continuing* Leontief inverse

$$\Psi_t = (I - \Omega_t)^{-1} = \sum_{s=0}^{\infty} \Omega_t^s.$$

Equation (16) shows that changes in the price of continuing goods depend on changes in markups, $\Delta \log \boldsymbol{\mu}_t$, productivity shifters, $\Delta \log \mathbf{A}_t$, prices of external inputs, $\Delta \log \mathbf{w}_t$, as well as the extensive margin terms, $\Delta \mathcal{X}_t$ and $\Delta \mathcal{E}_t$. All of these effects are mediated by the forward linkages in the Leontief inverse Ψ_t .

Define the *cost-based continuing Domar weights* of $i \in C_t$ and $f \in F$ to be

$$\lambda_{i,t} = \sum_{j \in C_t} b_{j,t} \Psi_{ji,t}, \quad \text{and} \quad \Lambda_{f,t} = \sum_{j \in C_t} \lambda_{j,t} \Omega_{j,t}^F.$$

The cost-based continuing Domar weight $\lambda_{i,t}$ measures the exposure of each continuing firm j to each continuing supplier i , captured by $\Psi_{ji,t}$, and averages this exposure by j 's share in final output $b_{j,t}$. Substituting (16) into the definition of the final output price deflator yields the following first order approximation for the change in the output price deflator

$$\Delta \log P_t^Y \approx \sum_{i \in C_t} \lambda_{i,t} \left[\Delta \log \frac{\mu_{i,t}}{A_{i,t}} + \bar{\delta}_{i,t}^{\text{sep}} \Delta \mathcal{X}_{i,t} - \bar{\delta}_{i,t}^{\text{add}} \Delta \mathcal{E}_{i,t} \right] + \sum_{f \in F} \Lambda_{f,t} \Delta \log w_{f,t}.$$

That is, shocks to i are transmitted into the final output price according to the cost-based Domar weight $\lambda_{i,t}$. Similarly, changes in the price of external input f affects the final output price deflator according to its cost-based Domar weight $\Lambda_{f,t}$.

Plugging this into the definition of real final output in equation (15) yields the following decomposition.

Proposition 3 (Growth-accounting with supplier-churn). *The change in real final output is given, to a first-order, by*

$$\begin{aligned} \Delta \log Y_t \approx & \underbrace{\sum_{i \in C_t} \lambda_{i,t} \Delta \log A_{i,t}}_{\text{technology}} + \underbrace{\sum_{f \in F} \Lambda_{f,t} \Delta \log L_{f,t}}_{\text{factor quantities}} \\ & - \underbrace{\sum_{i \in C_t} \lambda_{i,t} \Delta \log \mu_{i,t}}_{\text{markups}} - \underbrace{\sum_{f \in F} \Lambda_{f,t} \Delta \log \check{\Lambda}_{f,t}}_{\text{factor shares}} \\ & + \underbrace{\sum_{i \in C_t} \lambda_{i,t} \left(\bar{\delta}_{i,t}^{\text{add}} \Delta \mathcal{E}_{i,t} - \bar{\delta}_{i,t}^{\text{sep}} \Delta \mathcal{X}_{i,t} \right)}_{\text{supplier churn due to price jumps}}, \end{aligned}$$

where $\check{\Lambda}_{f,t} = w_{f,t} L_{f,t} / \sum_{j \in C_t} p_{j,t} y_{j,t}$ is the income share of factor f at t .

As discussed in Section 2, since Proposition 1 nests both expanding-variety and quality-ladder models, Proposition 3 also applies to both classes of models.

To better understand this proposition, start by considering the neoclassical benchmark where consumer surplus ratio is zero, $\bar{\delta}^{\text{add}} = \bar{\delta}^{\text{sep}} = 0$, and prices are equal to marginal cost, $\mu = 1$. In this case, Proposition 3 collapses to the standard Solow-Hulten formula, where only the first line is non-zero and cost-based Domar weights are equal to sales shares.

In the more general case, aggregate output growth can be broken down into the following terms. The first term is exogenous productivity growth weighted by cost-based Domar

weights. This accounts for how exogenous improvements in technology affect output, taking into account the fact that improvements in each firm’s technology will mechanically raise production by its consumers, and its consumers’ consumers, and so on. The second term captures a similar effect but for changes in factor quantities — if the quantity of factor f rises, then that raises the production of all firms that use factor f , which raises the production of all firms that use the products of factor f , and so on.³⁹

The second line captures the way changes in markups and factor prices affect output. An increase in i ’s markup will raise i ’s price, which raises the costs of production for i ’s consumers, and i ’s consumers’ consumers, and so on. Similarly, if the factor share $\check{\Lambda}_f$ of factor f rises more quickly than the quantity L_f of factor f , then this means that the relative price of factor f has increased. An increase in f ’s price will raise the costs of production for all firms. An increase in markups or factor shares, therefore, raises prices and lowers output.

The last line is what this paper is focused on and captures the effects of supplier churn on output. It measures the reduction in the final-goods price deflator caused by jumps in input prices due to supplier churn, holding fixed technologies of continuing firms, markups, and factor prices. Churn at the level of each individual firm percolates to the rest of the economy through the input-output network and this effect is captured by weighing the extensive margin terms by the cost-based Domar weight of each firm and summing across all firms. This captures the idea that if one firm’s marginal costs change from separations and additions of suppliers, then those marginal cost changes will propagate to that firm’s consumers, its consumers’ consumers, and so on. The elasticity of aggregate output with respect to additions and separations for firm i is $\lambda_i \bar{\delta}^{add}$ and $\lambda_i \bar{\delta}^{sep}$.

In general equilibrium, the different terms of Proposition 3 are interdependent (e.g. technology shocks may result in changes markups, factor shares, and supplier churn). Answering counterfactual questions requires solving for these terms in equilibrium, which in turn may require modeling the specifics of fixed costs, extensive margin decisions, market structure, etc. However, conditional on the terms in Proposition 3, we do not need to specify such details.

In the next section, we show results for different values of $\bar{\delta}$ given observed additions and separations. This exercise is analogous to computing the contribution of, say, capital to growth given observed investment under different assumptions about the elasticity of output with respect to capital. Although useful for inspecting the mechanisms driving

³⁹When we apply Proposition 3, we use a Tornqvist second-order adjustment. That is, although Proposition 3 is a first order approximation, we average the t and $t + 1$ coefficients on each shock to provide a second order approximation. For example, we weigh $\Delta \log L_{f,t}$, the change in factor quantity f between t and $t + 1$, using the average of $\Lambda_{f,t}$ and $\Lambda_{f,t+1}$.

aggregate growth, as in standard growth accounting, the results cannot be used to make counterfactual statements since, just like capital, supplier churn is endogenous.

5 Macroeconomic Results: Empirics

In this section, we apply Proposition 3 to decompose aggregate growth for a large subset of the Belgian economy. In the first part of this section, we describe how we map the data to the terms in Proposition 3. In the second part of this section, we show the results.

5.1 Mapping to Data

Proposition 3 is exact in continuous time if the primitive shocks are smooth functions of time. Following standard practice in the growth accounting literature (Solow, 1957), we map our model to data using a discrete-time approximation of the continuous time limit. To apply Proposition 3, we need to define the set of continuing firms C_t , the matrices Ω_t and Ω_t^F , the average consumer surplus parameters $\bar{\delta}_{i,t}^{\text{add}}$ and $\bar{\delta}_{i,t}^{\text{sep}}$, the share of additions and separations due to price jumps, $\Delta \mathcal{E}_t$ and $\Delta \mathcal{X}_t$, changes in markups $\Delta \log \mu_{i,t}$, the growth in external input quantities (labor, capital, and external materials), and the growth in final real output. The exogenous technology term in Proposition 3 is a residual. We discuss how we construct these terms in turn.

Assigning the continuing network set. We construct the network of domestic firms using the NBB B2B Transactions data set, which has near-universal coverage of domestic firms. This data set contains the values of yearly sales relationships among all VAT-liable companies for the years 2002 to 2018, and is based on the VAT listings collected by the tax authorities. We calculate an output measure for a subset of continuing, non-financial domestic Belgian corporations. We exclude self-employed, government, financial entities, and non-financial corporations in non-market services (NACE codes 84 and higher) because these sectors are not well-covered by VAT data (for example, hospitals and health centers are not required to submit VAT returns) and markups are hard to measure.⁴⁰ Even though we exclude from N self-employed, government, financial entities and non-market services, we include purchases from these suppliers in variable costs and treat them as a separate external factor.

⁴⁰We exclude self-employed because of data-privacy considerations. Government (including education) and non-market services, such as health, art, and entertainment are not well-covered by VAT data. We exclude financial entities because (i) banks fill special annual accounts that we do not have access to, and (ii) interest receipts by banks and insurance premia receipts by insurance companies are not included in the VAT data. Our micro estimates are slightly smaller than our baseline if we exclude input purchases from self-employed, government, and finance suppliers (see Table A12).

We define a firm in N to be continuing in t if the following conditions are met: its sales, employment, capital stock, and intermediate inputs are positive in t and $t + 1$. This gives us the set C_t , which covers around 70% of both value-added and total employment of the non-financial corporate sectors in Belgium as measured by the National Accounts Institute (see Table A7). Crucially, our output measure is much broader than the Prodcom sample that we used in Section 3. Whereas our Prodcom sample contains roughly 3,000 downstream firms per year, the growth accounting sample contains roughly 100,000 firms per year.

Calibrating input-output shares and markups. We construct the $C_t \times C_t$ network of domestic suppliers of Belgian firms using the NBB B2B Transactions data set. As mentioned before, almost all firms in Belgium are required to report sales of at least 250 euros, and the data has universal coverage of all businesses in C_t . We drop from the network purchases of capital inputs and outlier transactions as described in Section 3. There are four external inputs: labor, capital, imported materials, and materials from outside the set N (i.e. purchased from self-employed firms, finance, and government entities).⁴¹ We construct the $C_t \times F$ matrix of external input requirements using data from the annual accounts, B2B transactions, and customs declarations. For capital, as in Section 3, we multiply the industry-specific user cost of capital by firms' reported capital stocks. We measure firm-level markups by dividing sales by total variable costs. Total variable costs is the sum of intermediate inputs and the non-overhead component of the wage bill and the cost of capital (which we assume is a fraction $\phi = 0.50$ of labor and capital costs). Any other expenditures the firm incurs are treated as overhead costs.⁴²

Calibrating final output. Final output is defined to be the sales of continuing firms in the network, C_t , minus sales of materials to other firms in the production network. That is, final output are sales to households, exports, investment, and any other sales that are not considered to be intermediate purchases by firms in N .⁴³ We convert nominal final output into a real measure by deflating nominal growth in final output using the Belgian

⁴¹We also include in this external factor purchases from suppliers that do not report VAT, intra-firm purchases (due to mergers and acquisitions), and purchases from zero-employment continuing suppliers.

⁴²To ensure consistency across datasets, when we construct the input-output table $\Omega_{ij,t}$, we rescale each firm's intermediate purchases from the B2B network and intermediate imports to ensure that their sum equals our measure of total intermediate input purchases (sales minus value added from the annual accounts). We check that using these rescaled values of intermediate purchases to calculate addition and separation shares produces very similar micro estimates of δ as in our baseline regressions.

⁴³Given data on sales ($p_i q_i$) for each firm $i \in C_t$, and the input-output matrix relative to sales, $\Omega_{ij}^s = \frac{p_j x_{ij}}{p_i q_i}$, we calculate total final output as $E = \sum_{i \in C_t} p_i q_i - \sum_{i \in C_t} p_i q_i \sum_{j \in C_t} \Omega_{ij}^s$. Final output shares are given by $b_i = (p_i q_i - \sum_{j \in C_t} \Omega_{ji}^s p_j q_j) / E$.

GDP deflator from the national accounts. That is, we assume that the price deflator of our measure of final output grows at the same rate as the Belgian GDP deflator.

Calibrating external input quantities. We measure growth in labor quantity using total equivalent full time employees for firms in our sample. We measure growth in the capital stock of each firm by deflating the nominal value of its capital stock (which includes plants, property, equipment, and intellectual property) using the aggregate investment price deflator from the national accounts of Belgium. We measure the growth in imported materials by deflating the nominal imported material input growth with the import price deflator used for constructing the national accounts in Belgium. We cannot measure growth in the quantity of materials purchased from excluded domestic firms (self-employed, finance, and government entities, as well as continuing zero employment suppliers), so growth in the quantity of these materials is part of the residual.

Calibrating addition and separation shares. To apply Proposition 3, we need the variable cost share of additions and separations due to price jumps, $\Delta\mathcal{X}$ and $\Delta\mathcal{E}$, at the firm level. For the growth accounting exercise in this section, we assume that all discontinuous additions and separations, where the expenditure shares jump, are due to price jumps.

This assumption accommodates both expanding-variety and quality-ladder models. However, it does not accommodate additions and separations caused by discontinuous shifts in the input demand curve due to other factors, like biased shocks to downstream technology or continuing suppliers' prices. For example, this rules out an infinite cross-price elasticity of demand between separating and continuing suppliers. In other words, additions and separations caused by factors other than price jumps must happen smoothly (the input demand curve continuously shifts until the choke price is below the input price). In the continuous-time limit we consider, such additions and separations have no effect on the addition and separation shares since the expenditure share on inputs added or dropped in this way is zero at the moment when they are added or dropped.⁴⁴

In this limit, we can set

$$\Delta\mathcal{X}_{i,t} = \underbrace{\left(\sum_{J \in \mathcal{J}_i} M_{iJ,t} \Omega_{iJ,t} \right)}_{\text{interm. input share of total variable cost}} \underbrace{\left(1 - \frac{\sum_{j \in C_{i,t}} p_{j,t} x_{ij,t}}{\sum_k p_{k,t} x_{ik,t}} \right)}_{\text{interm. input share of separating suppliers}} \geq 0,$$

⁴⁴If this assumption is violated, then the effect on our growth accounting results is ambiguous. In the case where $\delta^{add} = \delta^{sep}$, only the net addition share matters. Hence, the bias of the supplier churn term depends on whether there are more additions or separations due to factors other than price jumps. In the former case, there is an upward bias, and in the latter case, there is a downward bias. If, on average, additions and separations due to factors other than price jumps are the same, then there is no bias.

where $C_{i,t}$ is the set of continuing suppliers for firm i : $C_{i,t} = \{j \in C_t : x_{ij,t} \times x_{ij,t+1} > 0\}$. That is $\Delta\mathcal{X}_{i,t}$ is the share of firm i 's variable cost spent on suppliers that are lost between t and $t + 1$. Similarly, we can set

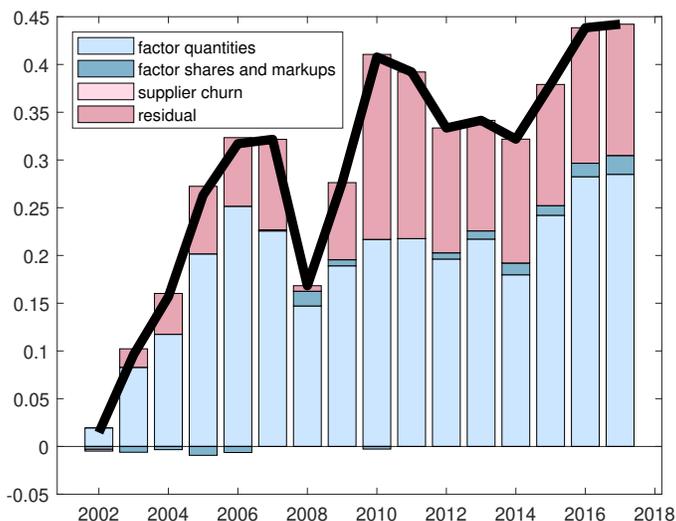
$$\Delta\mathcal{E}_{i,t} = \underbrace{\left(\sum_{J \in \mathcal{J}_i} M_{iJ,t} \Omega_{iJ,t} \right)}_{\text{interm. input share of total variable cost}} \underbrace{\left(1 - \frac{\sum_{j \in C_{i,t}} p_{j,t+1} x_{ij,t+1}}{\sum_k p_{k,t+1} x_{ik,t+1}} \right)}_{\text{interm. input share of added suppliers}} \geq 0.$$

This is the share of firm i 's variable cost spent on suppliers added between t and $t + 1$.⁴⁵

Calibrating $\bar{\delta}_{i,t}^{\text{add}}$ and $\bar{\delta}_{i,t}^{\text{sep}}$. We calibrate the average consumer surplus over additions and separations of suppliers per unit of expenditures using our microeconomic estimates from Section 3. We consider a few different cases: first, we set $\bar{\delta}_{i,t}^{\text{add}} = \bar{\delta}_{i,t}^{\text{sep}} = 0$, which turns off consumer surplus from supplier churn. Second, we set $\bar{\delta}_{i,t}^{\text{add}} = \bar{\delta}_{i,t}^{\text{sep}} = 0.28$, which is column (v) in Table 1, the IV estimates without additional controls. Finally, we set $\bar{\delta}_{i,t}^{\text{add}} = 0.283$ and $\bar{\delta}_{i,t}^{\text{sep}} = 0.268$, which is column (vi) in Table 1, the IV estimates with additional controls.⁴⁶ Table A13 in Appendix E shows how the results vary away from these cases. Figure A1, in the same appendix, shows how our results change if we let δ vary as a function of supplier size, as in column (ii) of Table 4.

Additional summary statistics. Table A7 in Appendix E reports information on the fraction of Belgian value-added in our sample and compares how aggregate growth rates in our sample compare to Belgian national accounts data. Table A8 in Appendix E reports basic statistics for the growth accounting sample of firms on the cost share of factors and intermediate inputs, the number of suppliers each firm has, and the separation and addition shares (relative to domestic material spending). Each firm has, on average, 68 suppliers (not reported in the table) while the sales-weighted average number of suppliers is 675. Table A8 also shows that addition shares are higher than separation shares.

Figure 4: Growth accounting with $\bar{\delta}^{\text{add}} = \bar{\delta}^{\text{sep}} = 0$.



5.2 Results

We start with a special case of Proposition 3 where the extensive margin is irrelevant, $\bar{\delta}^{\text{add}} = \bar{\delta}^{\text{sep}} = 0$. That is, Figure 4 implements a Baqaee and Farhi (2019) style decomposition. This is a generalization of Solow-Hulten growth decompositions to an environment with markups. The markup and factor share terms, which capture reallocations (see Baqaee and Farhi, 2019), do not play a large role in cumulative growth rates in this data set. The “unexplained” technology residual is large and accounts for about 14 log points of cumulative growth — roughly 1% per year.

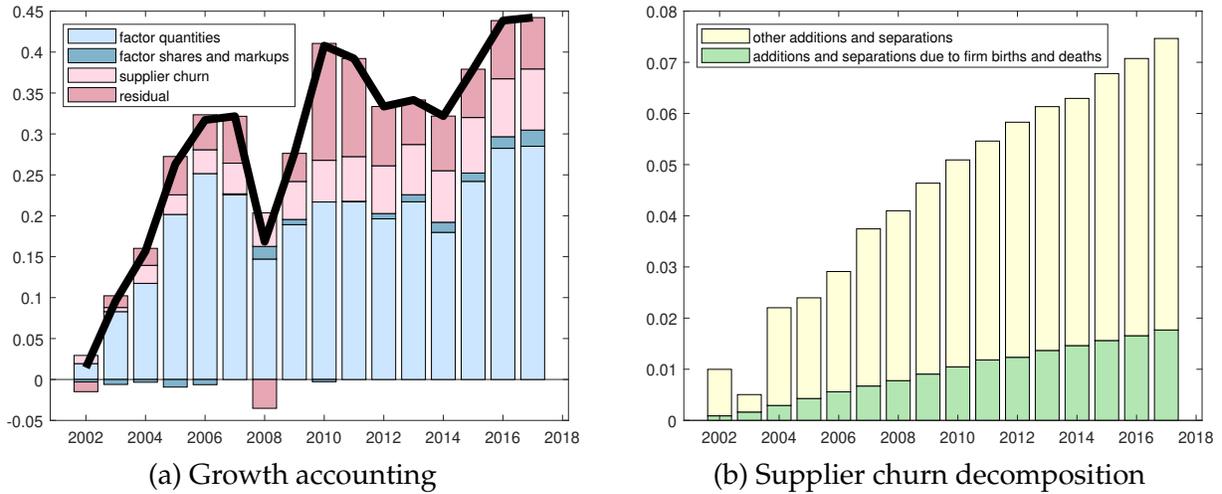
Figure 5 sets $\bar{\delta}^{\text{add}} = \bar{\delta}^{\text{sep}} = 0.28$. The left panel shows that the extensive margin of supplier addition and separation accounts for 7.5 log points out of a total of 14 log points of unexplained cumulative growth in the technology residual over the sample period. The extensive margin effect more than halves the size of the technology residual.⁴⁷ The

⁴⁵We construct $\sum_{J \in \mathcal{J}_i} M_{iJ,t} \Omega_{iJ,t}$ by adding up the cost share of i across all its suppliers in the included set N . That is, to construct this object, we do not need to explicitly group suppliers into different types, because only the sum matters.

⁴⁶Although we use all observed additions and separations to measure the addition and separation share in this section, we do not use the $\bar{\delta}$ estimated from an OLS regression of all additions and separations on marginal cost due to the endogeneity concerns described in Section 3.2. That is, if downstream technology shocks or changes in continuing suppliers’ prices are correlated with additions and separations, then Proposition 3 applies but the coefficients in an OLS regression are biased and cannot be used.

⁴⁷This does not mean that in a counterfactual where firms cannot add or drop suppliers aggregate productivity growth is 7.5 log points lower. In such a counterfactual, the remaining terms in Proposition 3 (the markup term, factor price changes, factor quantities, and the technology shocks) may all be different. The logic is similar to how in traditional growth accounting, shutting down productivity growth can affect, say, employment or capital accumulation. Instead, our growth accounting expression measures the technology residual given the observed patterns in the data and the calibrated values of δ .

Figure 5: $\bar{\delta}^{\text{entry}} = \bar{\delta}^{\text{exit}} = 0.28$



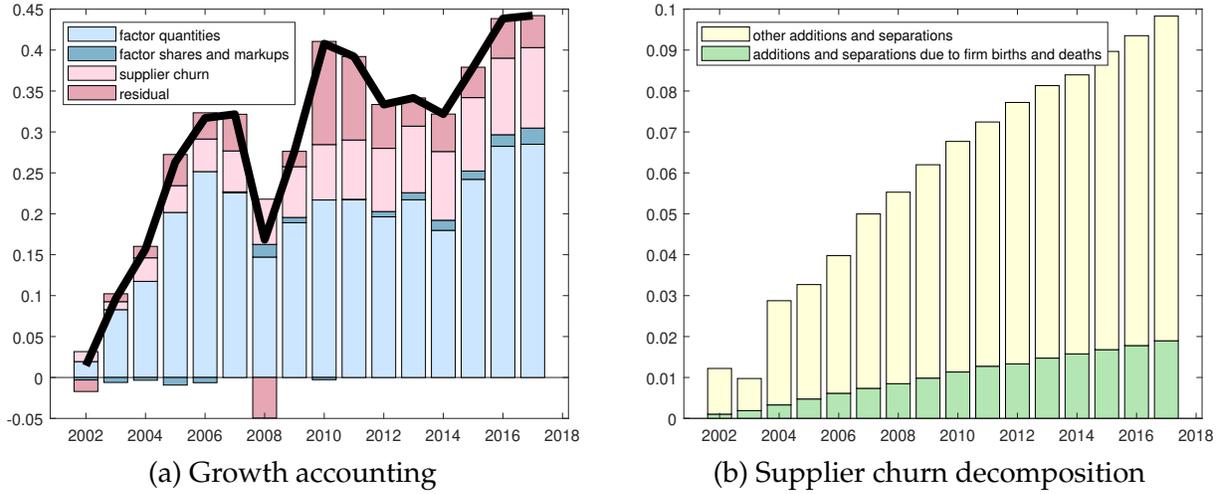
extensive margin effect is positive, even though $\bar{\delta}^{\text{add}}$ and $\bar{\delta}^{\text{sep}}$ are equal, because on balance additions are larger than separations (see Table A8). That is, the expenditure share on suppliers that continue from one year to the next declines on average. These results are consistent with both expanding variety models (on average, added varieties generate more consumer surplus than separating varieties) and quality-ladder models (on average, added suppliers are better than separating suppliers).

The right panel breaks down the extensive margin term into additions and separations associated with firm entry and exit (births and deaths) and the rest. Roughly one quarter is attributable to birth and death of firms, and the remaining three quarters is from additions and separations of firms that are continuing. Moreover, out of the 7.5 log points of the supplier churn term, three quarters is accounted for by services-producing downstream firms and a quarter by goods-producing downstream firms.

Whereas supplier churn is important for long-run growth in the period 2002-2018, it is not as important for explaining cyclical fluctuations. For example, the supplier churn term plays a small role for explaining the decline in aggregate output following the 2008 financial crisis. More formally, at annual frequency, the standard deviation of fluctuations in the residual is almost 10 times larger than that of the supplier churn term.

Figure 6 shows results using the point estimates from column (vi) of Table 1: $\bar{\delta}^{\text{add}} = 0.283$ and $\bar{\delta}^{\text{sep}} = 0.268$. Since additions are more valuable than separations, this enlarges the extensive margin term so that it accounts for almost 10 log points of growth. This reduces the technology residual's cumulative role in growth over our 16 year sample to roughly 4 log points. The right panel breaks down the extensive margin effect into additions and separations due to firm birth and death, and the rest.

Figure 6: Growth accounting with $\bar{\delta}^{\text{entry}} = 0.283$ and $\bar{\delta}^{\text{exit}} = 0.268$



Sensitivity analysis. Since the supplier churn term is not a residual, it is not affected by measurement error in the other terms in the growth accounting expression. It does depend on $\bar{\delta}^{\text{add}}$ and $\bar{\delta}^{\text{sep}}$. Table A13 in Appendix E provides robustness. Our results are fairly insensitive to different values of $\bar{\delta}^{\text{add}} = \bar{\delta}^{\text{sep}}$ in the range we estimate, but are sensitive to the gap between $\bar{\delta}^{\text{add}}$ and $\bar{\delta}^{\text{sep}}$. The magnitude of the supplier churn term is increasing in $\bar{\delta}^{\text{add}} - \bar{\delta}^{\text{sep}}$. This is because $\bar{\delta}^{\text{add}} - \bar{\delta}^{\text{sep}}$ dictates whether or not suppliers who are added are, on balance, more valuable per unit of expenditures than suppliers who separate. As we discuss in the appendix, the difference between $\bar{\delta}^{\text{add}}$ and $\bar{\delta}^{\text{sep}}$ matters a lot for the size of the supplier churn term because the gross level of additions and separations is high. Table A13 also shows the portion of supplier churn attributable to supplier births and deaths. These numbers are less sensitive to the precise values of $\bar{\delta}^{\text{add}}$ and $\bar{\delta}^{\text{sep}}$.

Figure A1 provides a different sensitivity where we allow δ to vary as a function of supplier size, as in column (ii) of Table 4, which is the only dimension of heterogeneity along which we find significant effects. Although the resulting consumer surplus ratios are quite different across individual suppliers once we allow them to vary by the size of the upstream firm, the aggregate consequences for productivity growth are quite similar.

Of course, our growth accounting results are suggestive since they involve extrapolating estimates from the Prodcom manufacturing sample of firms to a much broader subset of Belgian firms (including ones outside the manufacturing sector). However, with these caveats in mind, our aggregation exercise suggests that the extensive margin of supplier entry and exit is plausibly an important driver of aggregate productivity growth.

6 Conclusion

This paper analyzes and quantifies the microeconomic and macroeconomic importance of creation and destruction of supply linkages. Our analysis shows that downstream firms' marginal costs are significantly affected by supplier entry and exits, and this enables us to directly calculate the area under the input demand curve. The reduced form statistic we estimate shapes counterfactuals in many theories with an extensive margin. For example, it disciplines the welfare effect of changes in market size (e.g. Krugman, 1979), gains from trade (e.g. Melitz, 2003), efficiency of the decentralized equilibrium with entry (e.g. Matsuyama and Ushchev, 2020b), and optimal innovation subsidies (e.g. Baqaee and Farhi, 2020).

Our growth accounting results demonstrate that supplier additions and separations plausibly account for a large portion of the long-run aggregate productivity growth in a Solow (1957)-style growth accounting exercise. That is, inframarginal surplus associated with supplier churn can be an important channel through which aggregate productivity grows. These macroeconomic moments can be used as targeted moments for disciplining structural models of endogenous network formation and growth.

References

- Acemoglu, D. and P. D. Azar (2020). Endogenous production networks. *Econometrica* 88(1), 33–82.
- Acemoglu, D. and A. Tahbaz-Salehi (2020). Firms, failures, and fluctuations: the macroeconomics of supply chain disruptions. Technical report, National Bureau of Economic Research.
- Aghion, P., A. Bergeaud, T. Boppart, P. J. Klenow, and H. Li (2019). Missing growth from creative destruction. *American Economic Review* 109(8), 2795–2822.
- Aghion, P. and P. Howitt (1992). A model of growth through creative destruction. *Econometrica: Journal of the Econometric Society*, 323–351.
- Akcigit, U. and W. R. Kerr (2018). Growth through heterogeneous innovations. *Journal of Political Economy* 126(4), 1374–1443.
- Amiti, M., O. Itskhoki, and J. Konings (2019). International shocks, variable markups, and domestic prices. *The Review of Economic Studies* 86(6), 2356–2402.
- Amiti, M. and J. Konings (2007, December). Trade Liberalization, Intermediate Inputs, and Productivity: Evidence from Indonesia. *American Economic Review* 97(5), 1611–1638.
- Amiti, M., S. J. Redding, and D. E. Weinstein (2019). The impact of the 2018 trade war on us prices and welfare.
- Ariu, A., F. Mayneris, and M. Parenti (2020). One way to the top: How services boost the demand for goods. *Journal of International Economics* 123, 103278.
- Arkolakis, C., F. Huneus, and Y. Miyachi (2021). Spatial production networks. *Unpublished, Yale University*.
- Baily, M. N., C. Hulten, D. Campbell, T. Bresnahan, and R. E. Caves (1992). Productivity dynamics in manufacturing plants. *Brookings papers on economic activity. Microeconomics* 1992, 187–267.

- Baqae, D. and E. Farhi (2020). Entry versus rents: Aggregation with scale economies. Technical report.
- Baqae, D. R. (2018). Cascading failures in production networks. *Econometrica* 86(5), 1819–1838.
- Baqae, D. R. and A. Burstein (2023). Welfare and output with income effects and taste shocks. *The Quarterly Journal of Economics* 138(2), 769–834.
- Baqae, D. R. and E. Farhi (2019). Productivity and misallocation in general equilibrium. *The Quarterly Journal of Economics* 135(1), 105–163.
- Baqae, D. R. and E. Farhi (2024). Networks, barriers, and trade. *Econometrica* 92(2), 505–541.
- Baqae, D. R., E. Farhi, and K. Sangani (2024). The darwinian returns to scale. *Review of Economic Studies* 91(3), 1373–1405.
- Basu, S. and J. G. Fernald (2002). Aggregate productivity and aggregate technology. *European Economic Review* 46(6), 963–991.
- Benassy, J.-P. (1996). Taste for variety and optimum production patterns in monopolistic competition. *Economics Letters* 52.
- Bernard, A. B., E. J. Blanchard, I. Van Beveren, and H. Vandenbussche (2019). Carry-along trade. *The Review of Economic Studies* 86(2), 526–563.
- Bernard, A. B., E. Dhyne, G. Magerman, K. Manova, and A. Moxnes (2022). The origins of firm heterogeneity: A production network approach. *Journal of Political Economy* 130(7), 1765–1804.
- Bernard, A. B., A. Moxnes, and Y. U. Saito (2019). Production networks, geography, and firm performance. *Journal of Political Economy* 127(2), 639–688.
- Berry, S. T. and P. A. Haile (2021). Foundations of demand estimation. In *Handbook of Industrial Organization*, Volume 4, pp. 1–62. Elsevier.
- Blaum, J., C. Lelarge, and M. Peters (2018). The gains from input trade with heterogeneous importers. *American Economic Journal: Macroeconomics* 10(4), 77–127.
- Boehm, J. and E. Oberfield (2020). Misallocation in the market for inputs: Enforcement and the organization of production. *The Quarterly Journal of Economics* 135(4), 2007–2058.
- Brandt, L., J. Van Biesebroeck, L. Wang, and Y. Zhang (2017, September). Wto accession and performance of chinese manufacturing firms. *American Economic Review* 107(9), 2784–2820.
- Broda, C. and D. E. Weinstein (2006). Globalization and the gains from variety. *The Quarterly journal of economics* 121(2), 541–585.
- Broda, C. and D. E. Weinstein (2010). Product creation and destruction: Evidence and price implications. *The American economic review* 100(3), 691–723.
- Burstein, A. and G. Gopinath (2014). International prices and exchange rates. In *Handbook of international economics*, Volume 4, pp. 391–451. Elsevier.
- Carvalho, V. M., M. Nirei, and Y. Saito (2014, June). Supply chain disruptions: Evidence from the Great East Japan Earthquake. Technical Report 35, RIETI.
- Cherchye, L., T. Demuyne, B. De Rock, C. Duprez, G. Magerman, and M. Vershelde (2021). Structural identification of productivity under biased technological change. Technical report, ECARES.
- De Loecker, J., P. K. Goldberg, A. K. Khandelwal, and N. Pavcnik (2016). Prices, markups, and trade reform. *Econometrica* 84(2), 445–510.
- De Loecker, J. and F. Warzynski (2012). Markups and firm-level export status. *The American Economic Review* 102(6), 2437–2471.
- Dhingra, S. and J. Morrow (2019). Monopolistic competition and optimum product diversity under firm heterogeneity. *Journal of Political Economy* 127(1), 196–232.
- Dhyne, E., A. K. Kikkawa, T. Komatsu, M. Mogstad, and F. Tintelnot (2022). Foreign demand shocks to pro-

- duction networks: Firm responses and worker impacts. Technical report, National Bureau of Economic Research.
- Dixit, A. K. and J. E. Stiglitz (1977). Monopolistic competition and optimum product diversity. *The American Economic Review*, 297–308.
- Duprez, C. and G. Magerman (2018, October). Price updating in production networks. national bank of belgium working paper no. 352.
- Elliott, M., B. Golub, and M. V. Leduc (2022). Supply network formation and fragility. *American Economic Review* 112(8), 2701–2747.
- Fally, T. (2022). Generalized separability and integrability: Consumer demand with a price aggregator. *Journal of Economic Theory* 203, 105471.
- Feenstra, R. C. (1994). New product varieties and the measurement of international prices. *The American Economic Review*, 157–177.
- Feenstra, R. C. and D. E. Weinstein (2017). Globalization, markups, and us welfare. *Journal of Political Economy* 125(4), 1040–1074.
- Foley, C. (2022). Flexible entry/exit adjustment for price indices. Technical report, UCLA.
- Fontaine, F., J. Martin, and I. Mejean (2023). Frictions and adjustments in firm-to-firm trade.
- Foster, L., J. C. Haltiwanger, and C. J. Krizan (2001). Aggregate productivity growth: Lessons from microeconomic evidence. In *New developments in productivity analysis*, pp. 303–372. University of Chicago Press.
- Garcia-Macia, D., C.-T. Hsieh, and P. J. Klenow (2019). How destructive is innovation? *Econometrica* 87(5), 1507–1541.
- Goldberg, P. K., A. K. Khandelwal, N. Pavcnik, and P. Topalova (2010). Imported intermediate inputs and domestic product growth: Evidence from india. *The Quarterly journal of economics* 125(4), 1727–1767.
- Gopinath, G. and B. Neiman (2014). Trade adjustment and productivity in large crises. *American Economic Review* 104(3), 793–831.
- Grossman, G. M. and E. Helpman (1993). *Innovation and growth in the global economy*. MIT press.
- Grossman, G. M., E. Helpman, and H. Lhuillier (2023). Supply chain resilience: Should policy promote international diversification or reshoring? *Journal of Political Economy* 131(12), 3462–3496.
- Halpern, L., M. Koren, and A. Szeidl (2015). Imported inputs and productivity. *American Economic Review* 105(12), 3660–3703.
- Hausman, J. A. (1996). Valuation of new goods under perfect and imperfect competition. In *The economics of new goods*, pp. 207–248. University of Chicago Press.
- Hicks, J. R. (1940). The valuation of the social income. *Economica* 7(26), 105–124.
- Hottman, C. J., S. J. Redding, and D. E. Weinstein (2016). Quantifying the sources of firm heterogeneity. *The Quarterly Journal of Economics* 131(3), 1291–1364.
- Hsieh, C.-T., P. J. Klenow, et al. (2018). The reallocation myth. *Center for Economic Studies Working Paper* 18, 1–25.
- Hulten, C. R. (1978). Growth accounting with intermediate inputs. *The Review of Economic Studies*, 511–518.
- Huneus, F. (2018). Production network dynamics and the propagation of shocks. *Graduate thesis, Princeton University, Princeton, NJ*.
- Jacobson, T. and E. Von Schedvin (2015). Trade credit and the propagation of corporate failure: An empirical analysis. *Econometrica* 83(4), 1315–1371.
- Jaravel, X. (2019). The unequal gains from product innovations: Evidence from the us retail sector. *The Quarterly Journal of Economics* 134(2), 715–783.

- Kopytov, A., B. Mishra, K. Nimark, and M. Taschereau-Dumouchel (2024). Endogenous production networks under supply chain uncertainty. *Econometrica* 92(5), 1621–1659.
- Krugman, P. R. (1979). Increasing returns, monopolistic competition, and international trade. *Journal of international Economics* 9(4), 469–479.
- Levinsohn, J. and A. Petrin (2003). Estimating production functions using inputs to control for unobservables. *The review of economic studies* 70(2), 317–341.
- Lim, K. (2018). Endogenous production networks and the business cycle. *Work. Pap.*
- Makowski, L. and J. M. Ostroy (2001). Perfect competition and the creativity of the market. *Journal of economic literature* 39(2), 479–535.
- Mankiw, N. G. and M. D. Whinston (1986). Free entry and social inefficiency. *The RAND Journal of Economics*, 48–58.
- Matsuyama, K. and P. Ushchev (2017). Beyond CES: three alternative classes of flexible homothetic demand systems. *Global Poverty Research Lab Working Paper* (17-109).
- Matsuyama, K. and P. Ushchev (2020a). Constant pass-through.
- Matsuyama, K. and P. Ushchev (2020b). When does procompetitive entry imply excessive entry?
- Matsuyama, K. and P. Ushchev (2023). Love-for-variety. *CEPR DP 18184*.
- Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica* 71(6), 1695–1725.
- Miyauchi, Y. (2018). Matching and agglomeration: Theory and evidence from Japanese firm-to-firm trade. Technical report, working paper.
- Oberfield, E. (2018). A theory of input–output architecture. *Econometrica* 86(2), 559–589.
- Petrin, A. and J. Levinsohn (2012). Measuring aggregate productivity growth using plant-level data. *The RAND Journal of Economics* 43(4), 705–725.
- Romer, P. M. (1987). Growth based on increasing returns due to specialization. *The American Economic Review* 77(2), 56–62.
- Solow, R. M. (1957). Technical change and the aggregate production function. *The review of Economics and Statistics*, 312–320.
- Spence, M. (1976). Product selection, fixed costs, and monopolistic competition. *The Review of economic studies* 43(2), 217–235.
- Taschereau-Dumouchel, M. (2020). Cascades and fluctuations in an economy with an endogenous production network. *Available at SSRN 3115854*.
- Tintelnot, F., A. K. Kikkawa, M. Mogstad, and E. Dhyne (2018). Trade and domestic production networks. Technical report, National Bureau of Economic Research.
- Vives, X. (1999). *Oligopoly pricing: old ideas and new tools*. MIT press.
- Zhelobodko, E., S. Kokovin, M. Parenti, and J.-F. Thisse (2012). Monopolistic competition: Beyond the constant elasticity of substitution. *Econometrica* 80(6), 2765–2784.

Online Appendix

Consumer Surplus from Suppliers: How Big is it and Does it Matter for Growth?

David Baqaee, Ariel Burstein, Cédric Duprez, and Emmanuel Farhi

May 2025

A Proofs	A1
B Monopsonistic Downstream Firms	A8
C Additional Data Details	A11
D Sensitivity Analysis for Section 3	A14
E Additional Tables and Figures	A17
F Monte Carlo Simulations	A22

Appendix A Proofs

Proof of Proposition 1. Consider some downstream firm, and omit subscripts for that downstream firm throughout the proof. For notational simplicity, focus on changes in the input prices of a single input type J . We will take advantage of the fact that the total derivative of the marginal cost function is the sum of the derivatives with respect to changes in the prices of each input type.

Index inputs of type J by real numbers j . Let $p_J(j)$ be the price of input variety j of type J . Define the input price function by

$$p_J(j) = \begin{cases} p_J^0 & j < M_J - M_J^{sep} \\ p_J^1 & j \in [M_J - M_J^{sep}, M_J] \\ p_J^2 & j \in [M_J, M_J + M_J^{add}] \\ \infty & j > M_J + M_J^{add} \end{cases}.$$

Hence, the price function for inputs of type J is parameterized by the scalars: $p_J^0, M_J, p_J^1, M_J^{sep}, p_J^2,$ and M_J^{add}). We suppress the dependence of the marginal cost function on variables except write $mc(p_J^0, M_J, p_J^1, M_J^{sep}, p_J^2, M_J^{add})$ since all other variables are being held constant for the perturbation.

To capture the separations or additions of suppliers, we consider the change in marginal cost as p_j^1 moves from p_j^0 to ∞ or p_j^2 moves from ∞ to p_j^0 . To capture changes in continuing suppliers' prices, we consider the change in marginal cost from changes in p_j^0 . We take derivatives of changes in marginal cost with respect to M_J^{sep} , M_J^{add} , and p_j^0 , and evaluate these derivatives at the initial point $(p_j^0, M_J, p_j^1, M_J^{sep}, p_j^2, M_J^{add}) = (p_j^0, M_J, \infty, 0, p_j^0, 0)$.

To start, consider the change in marginal cost caused by some separations starting at the initial point. The change in the marginal cost is

$$\begin{aligned} \log \frac{mc(p_j^0, M_J, \infty, M_J^{sep}, p_j^0, 0)}{mc(p_j^0, M_J, p_j^0, M_J^{sep}, p_j^0, 0)} &= \int_{p_j^0}^{\infty} \int_{j \in [M_J - M_J^{sep}, M_J]} \Omega_J(p_j^0, M_J, \zeta, M_J^{sep}, p_j^0, 0) dj d \log \zeta \\ &= M_J^{sep} \int_{p_j^0}^{\infty} \Omega_J(p_j^0, M_J, \zeta, M_J^{sep}, p_j^0, 0) d \log \zeta, \end{aligned}$$

where Ω_j is the share of any input $j \in [M_J - M_J^{sep}, M_J]$ in total variable cost (which is the same for all $j \in [M_J - M_J^{sep}, M_J]$ by symmetry of the cost function). The first equality follows from the fundamental theorem of calculus for line integrals and Shephard's lemma. Rewrite the previous equation as

$$\begin{aligned} \log mc(p_j^0, M_J, \infty, M_J^{sep}, p_j^0, 0) &= M_J^{sep} \int_{p_j^0}^{\infty} \Omega_J(p_j^0, M_J, \zeta, M_J^{sep}, p_j^0, 0) d \log \zeta \\ &\quad + \log mc(p_j^0, M_J, p_j^0, M_J^{sep}, p_j^0, 0). \end{aligned}$$

We now approximate this exact expression as M_J^{sep} rises, capturing the separation of more varieties. The derivative of the marginal cost function with respect to M_J^{sep} is

$$\begin{aligned} d \log mc(p_j^0, M_J, \infty, M_J^{sep}, p_j^0, 0) &= dM_J^{sep} \int_{p_j^0}^{\infty} \Omega_J(p_j^0, M_J, \zeta, M_J^{sep}, p_j^0, 0) d \log \zeta \\ &\quad + M_J^{sep} \int_{p_j^0}^{\infty} \left[\frac{\partial \Omega_J(p_j^0, M_J, \zeta, M_J^{sep}, p_j^0, 0)}{\partial M_J^{sep}} dM_J^{sep} \right] d \log \zeta, \end{aligned}$$

where we use the fact that $\partial \log mc(p_j^0, M_J, p_j^0, M_J^{sep}, p_j^0, 0) / \partial M_J^{sep} = 0$. Evaluating the derivative above at $M_J^{sep} = 0$ and suppressing arguments gives

$$d \log mc = dM_J^{sep} \int_{p_j^0}^{\infty} \Omega_J(p_j^0, M_J, \zeta, 0, p_j^0, 0) d \log \zeta.$$

Using the definition of Ω_J , rewrite the previous equation as

$$\begin{aligned} d \log mc &= dM_J^{sep} \int_{p_J^0}^{\infty} \frac{\xi x_J(p_J^0, M_J, \xi, 0, p_J^0, 0)}{\mathcal{C}(p_J^0, M_J, \xi, 0, p_J^0, 0)} d \log \xi, \\ &= dM_J^{sep} \int_{p_J^0}^{\infty} \frac{x_J(p_J^0, M_J, \xi, 0, p_J^0, 0)}{\mathcal{C}(p_J^0, M_J, \xi, 0, p_J^0, 0)} d\xi, \end{aligned}$$

where $x_J(p_J^0, M_J, \xi, 0, p_J^0, 0)$ is the quantity demanded of an input of type J with price ξ (omitting arguments that are constant in the perturbation, e.g. other input prices and output quantity). Using the fact that $\mathcal{C}(p_J^0, M_J, \xi, 0, p_J^0, 0) = \mathcal{C}(p_J^0, M_J, \infty, 0, p_J^0, 0)$ for any value of ξ , rewrite the right-hand side as

$$\begin{aligned} d \log mc &= dM_J^{sep} \int_{p_J^0}^{\infty} \frac{x_J(p_J^0, M_J, \xi, 0, p_J^0, 0)}{\mathcal{C}(p_J^0, M_J, \infty, 0, p_J^0, 0)} d\xi \\ &= dM_J^{sep} \frac{1}{\mathcal{C}(p_J^0, M_J, \infty, 0, p_J^0, 0)} \int_{p_J^0}^{\infty} x_J(p_J^0, M_J, \xi, 0, p_J^0, 0) d\xi \\ &= dM_J^{sep} \Omega_J \frac{\int_{p_J^0}^{\infty} x_J(p_J^0, M_J, \xi, 0, p_J^0, 0) d\xi}{p_J^0 x_J(p_J^0, M_J, p_J^0, 0, p_J^0, 0)} \\ &= dM_J^{sep} \Omega_J \delta_J. \end{aligned}$$

Hence, separations increase marginal cost in accordance to $\Omega_J \delta_J$.

Similarly, we can capture the addition of suppliers by repeating the same argument but instead considering the change in marginal cost caused by some additions starting at the initial point:

$$\begin{aligned} \log \frac{mc(p_J^0, M_J, \infty, 0, p_J^0, M_J^{add})}{mc(p_J^0, M_J, p_J^0, 0, p_J^0, 0)} &= - \int_{p_J^0}^{\infty} \int_{j \in [M_J, M_J + M_J^{add}]} \Omega_J(p_J^0, M_J, \infty, 0, \xi, M_J^{add}) dj d \log \xi \\ &= M_J^{add} \int_{p_J^0}^{\infty} \Omega_J(p_J^0, M_J, \infty, 0, \xi, M_J^{add}) d \log \xi, \end{aligned}$$

We then approximate this as M_J^{add} rises and evaluate at $M_J^{add} = 0$ to get

$$d \log mc = -dM_J^{add} \Omega_J \delta_J,$$

where $dM_J^{add} > 0$ corresponds to additions of varieties of type J .

To consider smooth (marginal) changes in the price of inputs of type J , we differentiate

marginal costs with respect to p_J^0 at $M_J^{sep} = M_J^{add} = 0$ and use Shephard's lemma:

$$d \log mc = M_J \Omega_J d \log p_J^0.$$

The final perturbation is to the technology parameter of the downstream firm, which trivially gives

$$d \log mc = \frac{\partial \log mc}{\partial \log A} d \log A.$$

Note that all of these perturbations are taken at the same initial point where $M_J^{add} = M_J^{sep} = 0$. Hence, summing these first-order perturbations in dM_J^{add} , dM_J^{sep} , and $d \log p_J^0$ across all types to the perturbation in $d \log A$ yields the proposition.^{A1} In writing the statement of the proposition, we write Δx in place of infinitesimal changes dx and replace equality signs with approximately equal signs. \square

Proof of Example 3. To prove (5), once again, we suppress the index i for the downstream firm and other arguments in conditional input demand. Observe that

$$x_J(p_J) = \frac{\frac{\partial(p_J x_J(p_J))}{\partial p_J}}{1 - \sigma_J(p_J)}.$$

Substitute this into the definition of δ_J to get

$$\delta_J = \frac{\int_{p_J}^{\infty} x_J(\xi) d\xi}{p_J x_J(p_J)} = \frac{\int_{p_J}^{\infty} \frac{\frac{\partial(\xi x_J(\xi))}{\partial \xi}}{1 - \sigma_J(\xi)} d\xi}{p_J x_J(p_J)}.$$

Marshall's second law implies that $\sigma_J(\xi) > \sigma_J(p_J)$ if $\xi > p_J$, and the fundamental theorem of calculus implies $\int_{p_J}^{\infty} \frac{\partial(\xi x_J(\xi))}{\partial \xi} d\xi = -p_J x_J(p_J)$. We thus have

$$\delta_J < \frac{\int_{p_J}^{\infty} \frac{\partial(\xi x_J(\xi))}{\partial \xi} d\xi}{p_J x_J(p_J)(1 - \sigma_J(p_J))} = \frac{-p_J x_J(p_J)}{p_J x_J(p_J)(1 - \sigma_J(p_J))} = \frac{1}{\sigma_J(p_J) - 1}.$$

To prove that $\delta'_J(p) < 0$, re-express the consumer surplus ratio as

$$\delta_J(p_J) = \frac{\int_{p_J}^{\infty} \frac{\partial(\xi x_J(\xi))}{\partial \xi} d\xi}{p_J x_J(p_J)}.$$

^{A1}If the variable cost function is isoelastic in output quantity $\mathcal{C}(p, A, q) = c(p, A)q^\alpha$, then the total derivative of marginal cost includes another term: $(\alpha - 1)d \log q$. In our benchmark with constant returns, $\alpha = 1$, this term does not appear. We do allow for $\alpha = 1.15$ in column (vi) of Table A3.

Note that

$$\begin{aligned}\delta'_J(p_J) &= -\frac{\frac{\partial(p_J x_J(p_J))}{\partial p_J}}{p_J x_J(p_J)} - \frac{\int_{p_J}^{\infty} \frac{\frac{\partial(\xi x_J(\xi))}{\partial \xi}}{1-\sigma_J(\xi)} d\xi}{p_J x_J(p_J)} \frac{\frac{\partial(p_J x_J(p_J))}{\partial p_J}}{p_J x_J(p_J)} \\ &= -\frac{\frac{\partial(p_J x_J(p_J))}{\partial p_J}}{p_J x_J(p_J)} \left[\delta_J(p_J) - \frac{1}{\sigma_J(p_J) - 1} \right].\end{aligned}$$

Hence,

$$\delta'_J(p_J) < 0$$

if

$$\frac{1}{\sigma_J(p_J) - 1} > \delta_J(p_J).$$

Note, from their definitions, that at the choke price, p_J^* , we must have $\delta_J(p_J^*) = 1/(\sigma_J(p_J^*) - 1) = 0$. For any $p_J < p_J^*$, Example 3 then guarantees that $\delta_J(p_J) < \frac{1}{\sigma_J(p_J) - 1}$. \square

Proof of Proposition 2. According to Proposition 1, and re-introducing the downstream firm i index, we can write

$$\Delta \log mc_{i,t} = -\bar{\delta}_{i,t}^{add} X_{1i,t} + \bar{\delta}_{i,t}^{sep} X_{2i,t} + W'_{i,t} \gamma + \epsilon_{i,t} \quad (\text{A1})$$

where $X_{1i,t}$ and $X_{2i,t}$ are the addition and separation share due to price jumps for firm i at time t and $W_{i,t}$ are other variables we control for, including fixed effects. The parameter γ is not necessarily a structural parameter and the error term $\epsilon_{i,t}$ is uncorrelated with $W_{i,t}$ by construction. Our first stage regression relates the addition and separation share to our instruments:

$$X_{1i,t} = \alpha_{11} Z_{1i,t} + \alpha_{12} Z_{2i,t} + W'_{i,t} \pi_1 + v_{1i,t},$$

$$X_{2i,t} = \alpha_{21} Z_{1i,t} + \alpha_{22} Z_{2i,t} + W'_{i,t} \pi_2 + v_{2i,t},$$

where $Z_{1i,t}$ and $Z_{2i,t}$ are the restricted birth and death instruments and $v_{1i,t}$ and $v_{2i,t}$ are residuals including other additions and separations due to price jumps and due to shifts of input demand. These first-stage residuals are orthogonal to the instruments by construction.

Without loss of generality, we also assume that $Z_{1i,t}$ and $Z_{2i,t}$ have been orthogonalized. That is, let $Z_{2i,t}$ be the residuals from a regression of the restricted death instrument on the restricted birth instrument so that they are uncorrelated by construction. Similarly, for each variable, say $Q_{i,t}$, let $\tilde{Q}_{i,t}$ be residuals from a regression of $Q_{i,t}$ on covariates $W_{i,t}$.

We first present some preliminary steps we use in the proof. Our assumption that the instruments are mutually independent of the error term in the second stage implies

$$\mathbb{E} [\tilde{Z}_{1,t} \epsilon_{i,t}] = \mathbb{E} [\tilde{Z}_{1,t}] \mathbb{E} [\epsilon_{i,t}] = 0, \quad (\text{A2})$$

where the second equality holds because $\mathbb{E} [\tilde{Z}_{1,t}] = 0$. A similar equation holds for $\tilde{Z}_{2,t}$. Our assumption that the instruments are mutually independent of the error terms in the first stage and also mutually independent of $\bar{\delta}_i^{add}$ and $\bar{\delta}_i^{sep}$ implies

$$\mathbb{E} [\bar{\delta}_{i,t}^{add} v_{1,t} \tilde{Z}_{1,t}] = \mathbb{E} [\bar{\delta}_{i,t}^{sep} v_{2,t} \tilde{Z}_{1,t}] = 0, \quad (\text{A3})$$

and similar for $\tilde{Z}_{2,t}$.^{A2,A3} Finally, our assumption that the instruments are mutually independent of $\bar{\delta}_{i,t}^{sep}$ and $\bar{\delta}_{i,t}^{add}$ implies that

$$\begin{aligned} \mathbb{E} [\tilde{Z}_{1,t}^2 \bar{\delta}_{i,t}^{sep}] &= \mathbb{E} [\tilde{Z}_{1,t}^2] \mathbb{E} [\bar{\delta}_{i,t}^{sep}] \\ \mathbb{E} [\tilde{Z}_{1,t}^2 \bar{\delta}_{i,t}^{add}] &= \mathbb{E} [\tilde{Z}_{1,t}^2] \mathbb{E} [\bar{\delta}_{i,t}^{add}], \end{aligned} \quad (\text{A4})$$

and similar for $\tilde{Z}_{2,t}^2$.

The estimates $\hat{\delta}^{add}$ and $\hat{\delta}^{sep}$ satisfy the moment conditions

$$\mathbb{E} \left[\left(\Delta \log \tilde{m} c_i + \hat{\delta}^{add} \tilde{X}_{1i} - \hat{\delta}^{sep} \tilde{X}_{2i} \right) \tilde{Z}_{1i} \right] = 0,$$

$$\mathbb{E} \left[\left(\Delta \log \tilde{m} c_i + \hat{\delta}^{add} \tilde{X}_{1i} - \hat{\delta}^{sep} \tilde{X}_{2i} \right) \tilde{Z}_{2i} \right] = 0,$$

where we have suppressed the time subscript for simplicity. Substituting the first stage into the second stage yields

$$\mathbb{E} \left[\left(\Delta \log \tilde{m} c_i + \hat{\delta}^{add} (\alpha_{11} \tilde{Z}_{1i} + \alpha_{12} \tilde{Z}_{2i} + v_{1i}) - \hat{\delta}^{sep} (\alpha_{21} \tilde{Z}_{1i} + \alpha_{22} \tilde{Z}_{2i} + v_{2i}) \right) \tilde{Z}_{1i} \right] = 0.$$

Simplify this equation using $\mathbb{E} [\tilde{Z}_{1i} \tilde{Z}_{2i}] = \mathbb{E} [\tilde{Z}_{1i} v_{1i}] = \mathbb{E} [\tilde{Z}_{1i} v_{2i}] = 0$ (where the two

^{A2}If $\bar{\delta}_{i,t}^{add}$ and $\bar{\delta}_{i,t}^{sep}$ are constant, then the first-stage regression implies $\mathbb{E} [v_{1,t} \tilde{Z}_{1,t}] = \mathbb{E} [v_{2,t} \tilde{Z}_{1,t}] = 0$, so (A3) does not require the assumption that the instruments are mutually independent of the error terms in the first stage.

^{A3}Instead of assuming that the instruments Z are mutually independent of $\bar{\delta}$ and the error in the first stage (conditional on the controls), we could alternatively assume that the instruments Z is independent of $\bar{\delta}$ and uncorrelated with the product of $\bar{\delta}$ and the error in the first stage (conditional on the controls). This is a weaker assumption.

latter equalities are implied by the first-stage regression) to obtain

$$\mathbb{E} [\Delta \log \tilde{m} c_i \tilde{Z}_{1i}] + \hat{\delta}^{add} \alpha_{11} \mathbb{E} [\tilde{Z}_{1i}^2] - \hat{\delta}^{sep} \alpha_{21} \mathbb{E} [\tilde{Z}_{i1}^2] = 0.$$

Substitute the residualized version of (A1) for $\Delta \log \tilde{m} c$ to get

$$\mathbb{E} \left[\left[-\bar{\delta}_i^{add} \tilde{X}_{1i} + \bar{\delta}_i^{sep} \tilde{X}_{2i} + \epsilon_i \right] \tilde{Z}_{1i} \right] + \hat{\delta}^{add} \alpha_{11} \mathbb{E} [\tilde{Z}_{1i}^2] - \hat{\delta}^{sep} \alpha_{21} \mathbb{E} [\tilde{Z}_{i1}^2] = 0.$$

Substitute the first stage and use (A2) to obtain

$$\begin{aligned} \mathbb{E} \left[\left[-\bar{\delta}_i^{add} (\alpha_{11} \tilde{Z}_{i1} + \alpha_{12} \tilde{Z}_{2i} + v_{1i}) + \bar{\delta}_i^{sep} (\alpha_{21} \tilde{Z}_{i1} + \alpha_{22} \tilde{Z}_{2i} + v_{2i}) \right] \tilde{Z}_{1i} \right] \\ + \hat{\delta}^{add} \alpha_{11} \mathbb{E} [\tilde{Z}_{1i}^2] - \hat{\delta}^{sep} \alpha_{21} \mathbb{E} [\tilde{Z}_{i1}^2] = 0. \end{aligned}$$

Using $\mathbb{E} [\tilde{Z}_{1i} \tilde{Z}_{2i}] = 0$ and (A3) simplifies this expression to

$$-\alpha_{11} \mathbb{E} [\bar{\delta}_i^{add} \tilde{Z}_{i1}^2] + \alpha_{21} \mathbb{E} [\bar{\delta}_i^{sep} \tilde{Z}_{i1}^2] + \hat{\delta}^{add} \alpha_{11} \mathbb{E} [\tilde{Z}_{1i}^2] - \hat{\delta}^{sep} \alpha_{21} \mathbb{E} [\tilde{Z}_{i1}^2] = 0.$$

Using (A4) further simplifies this expression to

$$-\alpha_{11} \mathbb{E} [\bar{\delta}_i^{add}] + \alpha_{21} \mathbb{E} [\bar{\delta}_i^{sep}] = -\hat{\delta}^{add} \alpha_{11} + \hat{\delta}^{sep} \alpha_{21}.$$

Following similar steps, the second moment condition implies

$$\mathbb{E} [\Delta \log \tilde{m} c_i \tilde{Z}_{2i}] + \hat{\delta}^{add} \alpha_{21} \mathbb{E} [\tilde{Z}_{2i}^2] - \hat{\delta}^{sep} \alpha_{22} \mathbb{E} [\tilde{Z}_{i1}^2] = 0$$

which can be simplified to

$$-\alpha_{21} \mathbb{E} [\bar{\delta}_i^{add}] + \alpha_{22} \mathbb{E} [\bar{\delta}_i^{sep}] = -\hat{\delta}^{add} \alpha_{21} + \hat{\delta}^{sep} \alpha_{22}$$

So the two estimates $\hat{\delta}^{add}$ and $\hat{\delta}^{sep}$ satisfy the following two equations:

$$-\alpha_{11} \mathbb{E} [\bar{\delta}_i^{add}] + \alpha_{21} \mathbb{E} [\bar{\delta}_i^{sep}] = -\hat{\delta}^{add} \alpha_{11} + \hat{\delta}^{sep} \alpha_{21}$$

and

$$-\alpha_{21} \mathbb{E} [\bar{\delta}_i^{add}] + \alpha_{22} \mathbb{E} [\bar{\delta}_i^{sep}] = -\hat{\delta}^{add} \alpha_{21} + \hat{\delta}^{sep} \alpha_{22}.$$

This gives the desired result that $\hat{\delta}^{add} = \mathbb{E} [\bar{\delta}_i^{add}]$ and $\hat{\delta}^{sep} = \mathbb{E} [\bar{\delta}_i^{sep}]$ as long as the matrix of α 's has full rank. \square

Proof of Proposition 3. In the text we showed that, to a first-order approximation, the final output price deflator is given by

$$\Delta \log P_t^Y = \sum_{i \in C_t} \lambda_{i,t} \left[\Delta \log \frac{\mu_{i,t}}{A_{i,t}} + \bar{\delta}_{i,t}^{\text{sep}} \Delta \mathcal{X}_{i,t} - \bar{\delta}_{i,t}^{\text{add}} \Delta \mathcal{E}_{i,t} \right] + \sum_{f \in F} \Lambda_{f,t} \Delta \log w_{f,t}.$$

Substitute this into

$$\Delta \log Y = \Delta \log \left(\sum_{i \in C_t} p_{i,t} y_{i,t} \right) - \Delta \log P_t^Y$$

and use the fact that $\sum_{f \in F} \Lambda_{f,t} = 1$ and the fact that $\Delta \log w_{f,t} = \Delta \log \check{\Lambda}_{f,t} - \Delta \log L_{f,t} + \Delta \log (\sum_{i \in C_t} p_{i,t} y_{i,t})$, and we obtain the expression in the proposition. \square

Appendix B Monopsonistic Downstream Firms

In Section 2 and Proposition 1 we assumed that firms buy inputs at given prices. Here we generalize this proposition to the case where firms face a price schedule for each input. We show that our regression identifies the consumer surplus ratio, given some assumptions.

Assume that if the firm buys \mathbf{x} units of each input type, the per unit cost is given by $\mathbf{p}(\mathbf{x})$. The cost minimization problem is

$$\mathcal{C}(\mathbf{p}(\cdot), A, q) = \min_{\mathbf{x}} \int_j p_j(\mathbf{x}) x_j d_j, \text{ subject to } q = AF(\mathbf{x}),$$

where for simplicity we suppress the index for the downstream firm. Given A and q , this cost minimization problem implies a vector of input quantity choices with its implied input prices. When an input j is unavailable, $p_j(\mathbf{x}) = \infty$ for any choice of \mathbf{x} . We assume that $\mathbf{p}(\mathbf{x})$ is homogeneous of degree zero in \mathbf{x} .

Proposition 4. Consider a downstream firm facing a finite change in the price schedule for type J inputs, $\Delta p_J(\cdot)$, in the measure of available inputs of type J , $\Delta M_J = \Delta M_J^{\text{add}} - \Delta M_J^{\text{sep}}$, and in technology ΔA . To a first-order approximation in these primitives, the change in marginal cost for this firm is,

$$\Delta \log mc = \sum_J \Omega_J M_J \Delta \log p_J - \sum_J \Omega_J \delta_J \Delta M_J + \frac{\partial \log \mathcal{C}}{\partial \log A} \Delta \log A, \quad (\text{A5})$$

where δ_J is the consumer surplus ratio for input of type J ,

$$\delta_J = \frac{\int_1^\infty x_J(\xi p_J(\cdot)) d\xi}{p_J x_J}.$$

In words, δ_J is the integral of demand of input J as its price schedule smoothly rises from $p_J(\cdot)$ to ∞ and the price schedule for other inputs remains unchanged, relative to initial expenditures on this input.

Proposition 1 is a special case of (A5) when input prices do not depend on input quantities. This justifies the regression in (11) under more general assumptions.

Homogeneity of degree zero in the price schedule combined with constant returns to scale in variable inputs implies that variable costs are homogeneous of degree one in quantity:

Lemma 1. *Suppose that $F(\mathbf{x})$ has constant returns to scale in \mathbf{x} , and $\mathbf{p}(\mathbf{x})$ is homogeneous of degree zero in \mathbf{x} . Then, $\partial \log C(\mathbf{p}(\cdot), A, q) / \partial \log q = 1$, and marginal cost equals average cost equals $C(\mathbf{p}(\cdot), A, 1)$.*

Proof. Under the assumption above, we have that:

$$\begin{aligned}
C(\mathbf{p}(\cdot), A, q) &= \min_{\mathbf{x}} \{p(\mathbf{x}) \cdot \mathbf{x} : q = AF(\mathbf{x})\} \\
&= \min_{\mathbf{x}} \{q(p(\mathbf{x}/q) \cdot \mathbf{x}/q) : q = AF(\mathbf{x}/q)q\} \\
&= \min_{\mathbf{z}} \{q(p(\mathbf{z}) \cdot \mathbf{z}) : q = AF(\mathbf{z})q\} \\
&= \min_{\mathbf{z}} \{q(p(\mathbf{z}) \cdot \mathbf{z}) : 1 = AF(\mathbf{z})\} \\
&= q \min_{\mathbf{z}} \{(p(\mathbf{z}) \cdot \mathbf{z}) : 1 = AF(\mathbf{z})\} \\
&= qC(\mathbf{p}(\cdot), A, 1).
\end{aligned}$$

□

We now sketch how to modify the proof of Proposition 1 to allow for price schedules. We focus on separation of suppliers, but the same logic can be used to consider additions and marginal changes in the price menu. Index inputs of type J by real numbers j . Let $p_{Jj}(\mathbf{x})$ be the price of input j of type J given input quantities \mathbf{x} . Consider some scalar $M_J^{sep} \geq 0$ and define the input price function:

$$p_{Jj}(\mathbf{x}) = \begin{cases} p_j^0(\mathbf{x}) & j < M_J - M_J^{sep} \\ t_J p_j^0(\mathbf{x}) & j \in [M_J - M_J^{sep}, M_J] \\ \infty & j > M_J \end{cases}$$

Hence, the price function for inputs of type J is parameterized by $(p_J^0(\cdot), M_J, t_J, M_J^{sep})$.

Denote variable costs by $\mathcal{C}(p_j^0(\cdot), M_J, t_J, M_J^{sep})$, where we omit those parameters that are kept constant in the perturbation.

To capture the separation of suppliers, consider the change in marginal cost as t_J goes from 1 to ∞ . The change in variable cost (keeping output constant) is

$$\log \frac{\mathcal{C}(p_j^0(\cdot), M_J, \infty, M_J^{sep})}{\mathcal{C}(p_j^0(\cdot), M_J, 1, M_J^{sep})} = \int_1^\infty \int_{j \in [M_J - M_J^{sep}, M_J]} \Omega_J(p_j^0(\cdot), M_J, \xi, M_J^{sep}) dj d \log \xi,$$

where Ω_J is the share of input $j \in [M_J - M_J^{sep}, M_J]$ in variable costs, and the equality follows from the fundamental theorem of calculus for line integrals and Shephard's lemma. When applying Shephard's lemma we use the envelope condition that, as the price schedule changes, changes in input quantities across inputs do not have first order effects on variable costs (including their effects on input prices). Next, use the symmetry of the cost function with respect to the prices of inputs of the same type to write

$$\log \frac{\mathcal{C}(p_j^0(\cdot), M_J, \infty, M_J^{sep})}{\mathcal{C}(p_j^0(\cdot), M_J, 1, M_J^{sep})} = M_J^{sep} \int_1^\infty \Omega_J(p_j^0(\cdot), M_J, \xi, M_J^{sep}) d \log \xi,$$

so

$$\log \mathcal{C}(p_j^0(\cdot), M_J, \infty, M_J^{sep}) = M_J^{sep} \int_1^\infty \Omega_J(p_j^0(\cdot), M_J, \xi, M_J^{sep}) d \log \xi + \log \mathcal{C}(p_j^0(\cdot), M_J, 1, M_J^{sep}).$$

We now approximate this exact expression as M_J^{sep} rises, capturing the separation of more varieties. The derivative of variable cost with respect to M_J^{sep} is

$$\begin{aligned} d \log \mathcal{C}(p_j^0(\cdot), M_J, \infty, M_J^{sep}) &= dM_J^{sep} \int_1^\infty \Omega_J(p_j^0(\cdot), M_J, \xi, M_J^{sep}) d \log \xi \\ &\quad + M_J^{sep} \int_1^\infty \left[\frac{\partial \Omega_J(p_j^0(\cdot), M_J, \xi, M_J^{sep})}{\partial M_J^{sep}} dM_J^{sep} \right] d \log \xi, \end{aligned}$$

where we use the fact that $\partial \log \mathcal{C}(p_j^0(\cdot), M_J, 1, M_J^{sep}) / \partial M_J^{sep} = 0$. Evaluating the derivative above at $M_J^{sep} = 0$ and suppressing arguments gives

$$d \log \mathcal{C} = dM_J^{sep} \int_1^\infty \Omega_J(p_j^0(\cdot), M_J, \xi, 0) d \log \xi.$$

Using the definition of Ω_J , we can rewrite the previous equation as

$$\begin{aligned} d \log \mathcal{C} &= dM_J^{sep} \int_1^\infty \frac{\xi x_J(p_J^0(\cdot), M_J, \xi, 0)}{\mathcal{C}(p_J^0(\cdot), M_J, \xi, 0)} d \log \xi, \\ &= dM_J^{sep} \int_1^\infty \frac{x_J(p_J^0(\cdot), M_J, \xi, 0)}{\mathcal{C}(p_J^0(\cdot), M_J, \xi, 0)} d\xi, \end{aligned}$$

where $x_J(p_J^0(\cdot), M_J, \xi, 0)$ is the quantity demanded of any input of type J with price schedule $\xi p_J^0(\cdot)$. Use the fact that $\mathcal{C}(p_J^0(\cdot), M_J, \xi, 0) = \mathcal{C}(p_J^0(\cdot), M_J, \infty, 0)$ for any value of ξ to rewrite the right-hand side as

$$\begin{aligned} d \log \mathcal{C} &= dM_J^{sep} \int_1^\infty \frac{x_J(p_J^0(\cdot), M_J, \xi, 0)}{\mathcal{C}(p_J^0(\cdot), M_J, \infty, 0)} d\xi \\ &= dM_J^{sep} \frac{1}{\mathcal{C}(p_J^0(\cdot), M_J, \infty, 0)} \int_1^\infty x_J(p_J^0(\cdot), M_J, \xi, 0) d\xi \\ &= dM_J^{sep} \Omega_J \frac{\int_1^\infty x_J(p_J^0(\cdot), M_J, \xi, 0) d\xi}{p_J^0 x_J^0} \\ &= dM_J^{sep} \Omega_J \delta_J, \end{aligned}$$

where $x_J^0 = x_J(p_J^0(\cdot), M_J, 1, 0)$ and $p_J^0 = p_J^0(x_J^0)$ are input quantities and prices associated to the cost minimization $\mathcal{C}(p_J^0(\cdot), M_J, 1, 0)$, and

$$\delta_J = \frac{\int_1^\infty x_J(p_J^0(\cdot), M_J, \xi, 0) d\xi}{p_J^0 x_J^0}.$$

Hence, separations increase marginal cost in accordance to $\Omega_J \delta_J$. Finally, the Lemma above implies that variable and marginal cost do not depend on output quantity, so $d \log \mathcal{C}$ equals $d \log mc$. The argument to consider supplier additions is analogous.

Appendix C Additional Data Details

Mergers and acquisitions. One challenge with using data recorded at the level of the VAT identifier is the case of mergers and acquisitions, since this might blur our entry/exit analysis of suppliers.^{A4} When a firm stops its business, it reports to the Crossroads Bank of Enterprises (CBE) the reason for ceasing activities, one of which is merger and acquisition.

^{A4}Another challenge is that VAT returns are made at the unit level, which in some instances group more than one VAT identifier. In this case, we group the two (or more) VAT identifiers into a unique firm.

In such cases, we use the financial links also reported in the Crossroads Bank of Enterprises (CBE) to identify the absorbing VAT identifier and we group the two (or more) VAT identifiers into a unique firm. We choose the VAT identifier with the largest total assets. We use this head VAT identifier as the identifier of the firm. Having determined the head VAT identifier, we aggregate all the variables up to the firm level. For variables such as total sales and inputs, we adjust the aggregated variables with the amount of B2B trade that occurred within the firm, correcting for double counting. For other non-numeric variables such as firms' primary sector, we take the value of its head VAT identifier. It is important to emphasize that we group VAT identifiers only for the year of the M&A and thereafter, and not over the whole panel period.

Estimating share of variable costs in labor and capital costs To estimate the share of labor and capital costs that are variable inputs, ϕ , we consider the following regression:

$$\Delta \log (\text{labor} + \text{capital})_{i,t} = \phi \times \Delta \log (\text{intermediate inputs})_{i,t} + \text{controls}_{i,t} + \varepsilon_{i,t}. \quad (\text{A6})$$

The variable $(\text{labor} + \text{capital})_{i,t}$ denotes the sum of labor and capital costs of firm i in period t , and $\text{intermediate purchases}_{i,t}$ denotes intermediate input purchases of firm i in period t . Assuming that the variable component of labor and capital costs move one-to-one with intermediate input purchases (which we assume are fully variable) in response to firm-level demand shocks that keep technologies and relative factor prices unchanged, ϕ captures the fraction of variable labor and capital costs.

We instrument changes in intermediate purchases using a Bartik-type demand shock. For each firm i at time t , we define the instrument:

$$\text{Firm's Demand}_{i,t} = \sum_j \sum_K \Omega_{iK,t} \times \Delta \log \text{sales}_{K,t+1}, \quad (\text{A7})$$

where $\Omega_{iK,t}$ is the share of i 's sales to other domestic firms in each industry K (leaving out the firm's own industry) and $\Delta \log \text{sales}_{K,t+1}$ is the change in total sales of industry K between t and $t + 1$.

All regressions include 4 digit NACE industry by year fixed effects, which is the most disaggregated classification we can consider for the sample of manufacturing firms. Controls include a non-manufacturing input-price deflator (calculated by weighing disaggregated industry-level deflators from Eurostat using firm-level sales shares across industries) and a variant of the instrument defined in (A7) where $\Omega_{iK,t}$ is the share of i 's variable costs spent on industry K .

Table A1 displays the results. Columns (i) and (ii) report OLS results, which shows a

positive but low estimate of ϕ . However, OLS is subject to omitted variable bias because changes in intermediate purchases can result from shocks to firms' costs, such as changes in the price of intermediates or factor-biased technical change.

Columns (iii)-(vii) show the 2SLS results for different samples of firms (manufacturing, goods producing firms, all firms, and the smaller Prodcum sample) and controls. In all cases (except for the Prodcum sample) the first-stage is strong (demand shocks help predict changes in intermediate input purchases). The point estimate of ϕ is between 0.4 and 0.6, and the controls have a small impact on the estimates. In our baseline, we set $\phi = 0.5$, which is also the fraction of variable inputs in labor costs estimated by Dhyne et al. (2022) using an export-demand instrument in the Belgian data. We consider alternative values for ϕ in sensitivity analysis.

Table A1: Elasticity of labor and capital costs with respect to intermediate purchases

	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)
	$\Delta \log(\text{labor} + \text{capital})$						
$\Delta \log(\text{interm. inputs})$	0.268*** (0.006)	0.269*** (0.006)	0.575*** (0.169)	0.575*** (0.175)	0.668 (0.459)	0.481*** (0.157)	0.399*** (0.054)
Specification	OLS	OLS	IV	IV	IV	IV	IV
F-stat			62	58	3	57	653
Sample of firms	Manufact.	Manufact.	Manufact.	Manufact.	Prodcum	Goods	All
Input prices control	N	Y	N	Y	Y	Y	Y
Bartik control	N	Y	N	Y	Y	Y	Y
Industry \times year FE	Y	Y	Y	Y	Y	Y	Y
Obs.	305,158	304,421	219,992	219,892	39,149	295,916	3,105,553

Notes: This table displays estimates of regression (A6) for different samples of firms. The instrument is the firms' demand shock defined in (A7). The first control is an input price deflator, and the second control is a variant of the instrument defined in (A7) using purchases from (rather than sales to) other industries. Industry fixed effects at the 4-digit NACE level. Regressions are unweighted, and standard errors are clustered at the firm-level.

Appendix D Sensitivity Analysis for Section 3

Table A2 provides sensitivity of our estimates for different configurations of fixed effects. Column (i) is our baseline specification with 6-digit industry by year fixed effects. Column (ii) includes a firm fixed effect to allow for the possibility that our instruments are correlated with trends in the downstream firm’s marginal cost. Column (iii) replaces industry-by-year fixed effects with a year fixed effect. Columns (iv) and (v) vary the disaggregation in the industry-by-year fixed effects, considering 4 or 8 digit product codes (rather than 6 digits). Our estimates are significant and quite robust across specifications with more or less stringent fixed effects.

Table A2: Estimates of δ under different fixed effect configurations

	(i)	(ii)	(iii)	(iv)	(v)
	$\Delta \log mc$				
Separation share	0.268*** (0.091)	0.304*** (0.106)	0.196** (0.080)	0.220** (0.092)	0.232*** (0.083)
Addition share	-0.283*** (0.079)	-0.335*** (0.090)	-0.244*** (0.071)	-0.270*** (0.082)	-0.256*** (0.068)
Specification	IV	IV	IV	IV	IV
F-stat	111	108	160	96	155
Controls	Y	Y	Y	Y	Y
6d industry \times year FE	Y	Y	N	N	N
8d industry \times year FE	N	N	N	Y	N
4d industry \times year FE	N	N	N	N	Y
Year FE	N	N	Y	N	N
Firm FE	N	Y	N	N	N
Observ.	38,670	37,898	41,980	34,696	41,643

Notes: Columns (i)-(v) report estimates of regression (11) for different fixed effect configurations. Column (i) is our baseline. Other controls are as in Table 1. All regressions are unweighted. Standard errors are clustered at the firm-level, and F-stat is the Kleibergen-Paap (KP) statistic.

Table A3 provides sensitivity to alternative measures of marginal cost. We vary the fraction of labor and capital costs that are overhead and we use a production function estimation approach to measure the change in marginal cost. We find similar results to our benchmark specification. Table A3 also considers a case where we allow for decreasing returns in the production function which slightly raises the magnitude of our point estimates.^{A5}

^{A5}We assume an isoelastic cost function, $C_i(\mathbf{p}, A_i, q_i) = c_i(\mathbf{p}, A_i) q_i^{1.15}$. Log changes in average variable

Table A3: Estimates of δ for alternative measures of marginal costs

	(i)	(ii)	(iii)	(iv)	(v)	(vi)
	Capital all overhead	60% overhead	40% overhead	0% overhead	Prod. fun. estimation	Decreasing returns
Separation share	0.271*** (0.091)	0.274*** (0.090)	0.313*** (0.103)	0.274** (0.108)	0.310*** (0.113)	0.304*** (0.105)
Addition share	-0.291*** (0.079)	-0.289*** (0.078)	-0.297*** (0.082)	-0.247*** (0.084)	-0.320*** (0.098)	-0.283*** (0.088)
Specification	IV	IV	IV	IV	IV	IV
F-stat	113	112	110	107	111	111
Controls	Y	Y	Y	Y	Y	Y
Industry \times year FE	Y	Y	Y	Y	Y	Y
Observ.	38,654	38,634	38,695	38,783	38,670	38,670

Notes: This table displays estimates of regression (11) for different measures of marginal cost, where we instrument separation and additions using restricted exit and entry shares defined by equations (12) and (13). Columns (i)-(iv) use measures of marginal costs under alternative assumptions on the share of overhead costs in capital and labor, column (v) uses marginal costs obtained from Levinsohn-Petrin production function estimates, column (vi) uses marginal costs assuming decreasing returns to scale in variable production, such that variable costs are $C_i(\mathbf{p}, A_i, q_i) = c_i(\mathbf{p}, A_i) q_i^{1.15}$. Columns (vii) and (viii) use two and three-year changes in marginal cost as outcomes. Controls are as in Table 1. All regressions are unweighted. Industry fixed effects are at the 6-digit product code level. Standard errors are clustered at the firm-level, and F-stat is the Kleibergen-Paap (KP) statistic.

Table A4 considers how results change if we vary the sample of firms. To reduce the possibility that the downstream firm changes the quality of its output in response to supplier additions and separations, column (i) restricts attention to downstream firms that do not change the mix of 8-digit products they offer and column (ii) focuses only on single product firms. In the latter case, the sample shrinks by half, and the estimated surplus ratio for separations increases but the one for additions stays similar.^{A6}

Table A4: Estimates of δ for alternative samples

	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)	(viii)	(ix)
	$\Delta \log mc$								
	Constant prod. mix	Single product	Two year cutoff	Three year cutoff	Employment weighted	Sep. & add. shares < 0.3	Sep. & add. shares < 1	Prodcom / total sales > 0.5	$ \Delta \log mc $ < 1
Separation share	0.258*** (0.093)	0.479*** (0.130)	0.250*** (0.093)	0.242** (0.107)	0.387** (0.154)	0.255** (0.102)	0.316*** (0.102)	0.284*** (0.093)	0.263*** (0.091)
Addition share	-0.297*** (0.081)	-0.355*** (0.123)	-0.285*** (0.085)	-0.286*** (0.094)	-0.239** (0.106)	-0.296*** (0.091)	-0.289*** (0.080)	-0.286*** (0.079)	-0.282*** (0.078)
Specification	IV	IV	IV	IV	IV	IV	IV	IV	IV
F-stat	105	54	88	64	86	156	75	106	111
Controls	Y	Y	Y	Y	Y	Y	Y	Y	Y
Industry \times year FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Observ.	36,163	19,097	33,306	27,990	38,670	37,792	38,961	33,978	38,656

Notes: This table displays estimates of regression (11) for different measures of marginal cost, where we instrument separation and additions using restricted exit and entry shares defined by equations (12) and (13). Column (i) drops downstream firms that switch the set of 8-digit products between years, and column (ii) drops firms that produce more than one 8-digit products. Columns (iii) and (iv) restrict the set of suppliers in the instrument to those for which the downstream firm is a small customer for two or three years (rather than one year in the baseline) before exiting or entering. Column (v) weights observations by employment of the downstream firm. Columns (vi) and (vii) drop observations in which the separation or addition share are higher than 0.3 or 1 (rather than 0.5 in the baseline). Column (viii) restricts the sample to firms whose Prodcom sales are at least 50% of total sales, and column (iv) drops observations for which the absolute size of marginal costs changes exceeds 1. Controls are as in Table 1. Industry fixed effects are at the 6-digit product code level. All regressions are unweighted except for column (v). Standard errors are clustered at the firm-level, and F-stat is the Kleibergen-Paap (KP) statistic.

Table A4 also considers a more demanding formulation of the instruments where the downstream firm has to be a small customer for exiting suppliers not just in the year the supplier exits but also in the year prior to exit (column iii) and two years prior to exit

costs are still equal to log changes in marginal costs, however, the change in marginal cost now depends on the change in output quantity, which we move to the left hand side of (11).

^{A6}If quality changes are associated with changes in product codes, then restricting attention to firms that do not change their product mix or have only a single product may help alleviate mismeasurement associated with quality change. Even though the 8 digit product codes are very detailed (e.g. “throat pastilles and cough drops consisting essentially of sugars and flavouring agents excluding pastilles or drops with flavouring agents containing medicinal properties”) a remaining concern is that, within 8 digit product codes, the downstream firm downgrades output quality in response supplier separation. In this case, we underestimate the rise in marginal cost because quality-adjusted quantity falls by more than measured quantity (and vice versa for supplier additions). In this case, our estimates of the consumer surplus ratio are biased towards zero.

(column iv). Similarly, when constructing the entry instrument, the downstream firm has to be a small customer for entering suppliers not just in the year of entry, but also the year after (column iii) or two years after (column iv) entry. The estimates are quite robust, except for the 3-year separation instrument, for which estimates lose some precision.

The remaining columns in Table A4 provide sensitivity to other choices, such as weighting observations by employment, changing the minimum threshold in the ratio of a firm's Prodcom sales to the firm's total sales from the annual accounts, and changing the treatment of outliers.

Appendix E Additional Tables and Figures

Table A5: Descriptive statistics: Prodcom sample

	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)	(viii)	(ix)	(x)	(xi)	(xii)
	Share in variable costs			Import	Service	Numb.	Share in variable costs					
	labor	capital	interm.	interm. share	interm. share	suppl.	separations	additions	deaths	births	deaths restricted	births restricted
mean	0.136	0.009	0.854	0.269	0.654	227	0.057	0.068	0.003	0.005	0.001	0.002
p25	0.071	0.003	0.805	0.000	0.522	112	0.022	0.026	0.000	0.000	0.000	0.000
p50	0.120	0.006	0.870	0.221	0.692	168	0.040	0.049	0.000	0.001	0.000	0.000
p75	0.184	0.012	0.922	0.469	0.815	257	0.073	0.087	0.002	0.003	0.001	0.001
count	41,980	41,980	41,980	41,980	41,980	41,980	41,980	41,980	41,980	41,980	41,980	41,980

Notes: The sample of firms used in this table are those used in the micro regressions in Section 3 based on the Prodcom sample. Service suppliers are those in NACE code sections F-T. Summary statistics are unweighted.

Table A6: Correlation of addition and separations with downstream firm size

	(i)	(ii)	(iii)	(iv)	(v)
	log number suppliers	separation share	addition share	restricted death share	restricted birth share
log employment	0.78	-0.23	-0.22	-0.04	-0.04
log sales	0.80	-0.30	-0.30	-0.07	-0.06

Notes: The sample of firms used in this table are those used in the micro regressions in Section 3 based on the Prodcom sample. All shares are calculated relative to variable costs of the downstream firm.

Table A7: Coverage of growth accounting sample of firms

	(i)	(ii)	(iii)	(iv)	(v)
year	count	value added	% of agg.	employment	% of agg.
2002	99,577	107,652	72%	1,574	67%
2003	102,716	114,520	74%	1,579	67%
2004	104,826	122,354	75%	1,588	67%
2005	106,476	125,755	74%	1,595	66%
2006	108,461	134,770	75%	1,636	67%
2007	109,761	142,913	75%	1,710	68%
2008	110,700	143,835	73%	1,727	67%
2009	109,413	137,080	73%	1,653	64%
2010	109,026	146,411	74%	1,640	63%
2011	110,216	150,341	73%	1,684	64%
2012	110,983	152,705	73%	1,696	64%
2013	110,168	153,660	72%	1,693	64%
2014	110,415	151,948	70%	1,633	62%
2015	106,344	155,171	69%	1,621	61%
2016	105,992	174,552	75%	1,776	65%
2017	105,948	180,709	75%	1,818	66%
avg. growth (%)		3.5	3.3	1.0	1.1

Notes: The sample of firms used in this table are those used in the growth accounting exercise (continuing corporate non-financial firms) in Section 5. Employment is in thousands of people, and value added is in millions of euros. “% agg.” is the share of value added and employment in the non-financial corporate sector reported in the national statistics calculated by the National Accounts Institute. The bottom row reports average annual growth rate for value added (in the sample and national statistics, respectively) and for employment.

Table A8: Descriptive statistics: growth-accounting sample (sales-weighted)

	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)	(viii)	(ix)	(x)
	Share in variable costs			Import	Services	Numb.	Share in domestic intermediate spending			
	labor	capital	interm.	interm. share	interm. share	suppl.	separations	additions	deaths	births
mean	0.074	0.009	0.917	0.315	0.725	675	0.096	0.110	0.005	0.009
p25	0.009	0.001	0.896	0.000	0.55	123	0.022	0.027	0.000	0.000
p50	0.037	0.002	0.958	0.148	0.846	330	0.053	0.065	0.000	0.001
p75	0.093	0.006	0.989	0.645	0.973	853	0.116	0.138	0.002	0.006
count	1,721,022	1,721,022	1,721,022	1,716,375	1,715,958	1,717,426	1,715,958	1,717,124	1,715,958	1,717,124

Notes: The sample of firms used in this table are those used in growth accounting in Section 5. Service suppliers are those in NACE code sections F-T. Summary statistics are weighted by sales.

Table A9: Separations and additions from continuing suppliers on instruments

	(i)	(ii)	(iii)	(iv)
	Separation count share from continuing suppliers		Addition count share from continuing suppliers	
Restricted death count share	-0.355*** (0.055)	-0.356*** (0.055)	0.282*** (0.062)	0.284*** (0.062)
Restricted birth count share	0.485*** (0.064)	0.484*** (0.065)	-0.005 (0.047)	-0.000 (0.048)
Specification	OLS	OLS	OLS	OLS
Controls	N	Y	N	Y
Industry \times year FE	Y	Y	Y	Y
Observations	38,670	38,670	38,670	38,670

Notes: This table shows that restricted supplier deaths predict additions from continuing suppliers — first row of columns (iii) and (iv). Restricted births predict separations from continuing suppliers — second row of columns (i) and (ii). Furthermore, restricted births do not predict additions of non-newly-entering suppliers — second row of columns (iii) and (iv). The separation (and addition) count share from continuing suppliers is the ratio of the number of suppliers who separate (or are added) by the downstream firm but continue to operate (or operated before addition) relative to the number of suppliers of the downstream firm. We use count share rather than cost share because the cost share can adjust through the intensive margin. Other controls are as in Table 1. Industry fixed effects are at the 6-digit product code level. All regressions are unweighted. Standard errors are clustered at the firm-level.

Table A10: Estimates of δ when separations and additions are regressed separately

	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)	(viii)	(ix)	(x)
	$\Delta \log mc$		First stage	$\Delta \log mc$		$\Delta \log mc$		First stage	$\Delta \log mc$	
Separation share	-0.009 (0.013)			0.190** (0.077)	0.182** (0.079)					
Additions share						0.014 (0.011)			-0.181*** (0.063)	-0.192*** (0.064)
Restricted death share		0.192** (0.083)	1.058*** (0.052)							
Restricted birth share							-0.225*** (0.075)	1.175*** (0.053)		
Specification	OLS	OLS	OLS	IV	IV	OLS	OLS	OLS	IV	IV
F-stat				463	408				496	500
Controls	Y	Y	Y	N	Y	Y	Y	Y	N	Y
Industry \times year FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Observations	38,670	38,670	38,670	38,670	38,670	38,670	38,670	38,670	38,670	38,670

Notes: Columns (i)-(v) report estimates of regression (11) where addition share and its instrument are dropped. Columns (vi)-(x) report estimates of regression (11) where separation share and its instrument are dropped. Columns (iii) and (viii) display the first-stage for each regression. Other controls are as in Table 1. Industry fixed effects are at the 6-digit product code level. All regressions are unweighted. Standard errors are clustered at the firm-level, and F-stat is the Kleibergen-Paap (KP) statistic.

Table A11: Estimates of $1/(\sigma - 1)$ assuming CES

	(i)	(ii)	(iii)
	$\Delta \log mc$		
$\Delta \log$ continuing share	0.265*** (0.078)	0.263** (0.128)	0.266*** (0.098)
Specification	IV	IV	IV
Instrument	Birth & death	Death	Birth
F-stat	48	27	82
Controls	Y	Y	Y
Industry \times year FE	Y	Y	Y
Observ.	38,670	38,670	38,670

Notes: This table reports estimates of a regression $\Delta \log mc_{it} = \hat{\beta} \times \Delta \log \text{continuing share}_{it} + \text{controls}_{it} + \varepsilon_{it}$. Column (i) instruments using both the death and birth instruments, column (ii) using the death instrument, and column (iii) using the birth instrument only. Controls are as in Table 1. Industry fixed effects are at the 6-digit product code level. All regressions are unweighted. Standard errors are clustered at the firm-level, and F-stat is the Kleibergen-Paap (KP) statistic.

Table A12: Estimates of δ for alternative set of suppliers

	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)
	$\Delta \log mc$						
	Industry	Services	Exclude utilities	Exclude retail & wholesale	Incl. capital producers	Excl. finance	Excl. self-empl., finance, govt.
Separation share	0.103 (0.119)	0.388*** (0.121)	0.261*** (0.092)	0.228* (0.123)	0.267*** (0.089)	0.265*** (0.090)	0.200** (0.090)
Addition share	-0.239 (0.153)	-0.327*** (0.095)	-0.276*** (0.081)	-0.300** (0.129)	-0.270*** (0.078)	-0.280*** (0.078)	-0.235*** (0.075)
Specification	IV	IV	IV	IV	IV	IV	IV
F-stat	100	90	107	71	109	123	120
Controls	Y	Y	Y	Y	Y	Y	Y
Industry \times year FE	Y	Y	Y	Y	Y	Y	Y
Observ.	38,968	38,819	38,675	38,872	38,623	38,679	38,702

Notes: This table displays estimates of regression (11) for different sets of suppliers. Industry suppliers are those in NACE code sections A-E, and service suppliers in sections F-T. Controls are as in Table 1. All regressions are unweighted. Industry fixed effects are at the 6-digit product code level. Standard errors are clustered at the firm-level, and F-stat is the Kleibergen-Paap (KP) statistic.

Table A13: Cumulative supplier churn term under alternative values of $\bar{\delta}^{add}$ and $\bar{\delta}^{sep}$

		$\bar{\delta}^{sep}$				
		0.26	0.27	0.28	0.29	0.30
$\bar{\delta}^{add}$	0.26	0.069	0.054	0.039	0.024	0.008
	0.27	0.087	0.072	0.057	0.042	0.026
	0.28	0.105	0.090	0.075	0.059	0.044
	0.29	0.123	0.108	0.093	0.077	0.062
	0.30	0.141	0.126	0.110	0.095	0.080

(a) All separations and additions

		$\bar{\delta}^{sep}$				
		0.26	0.27	0.28	0.29	0.30
$\bar{\delta}^{add}$	0.26	0.016	0.016	0.015	0.014	0.013
	0.27	0.018	0.017	0.016	0.016	0.015
	0.28	0.019	0.018	0.018	0.017	0.016
	0.29	0.021	0.020	0.019	0.018	0.018
	0.30	0.022	0.021	0.020	0.020	0.019

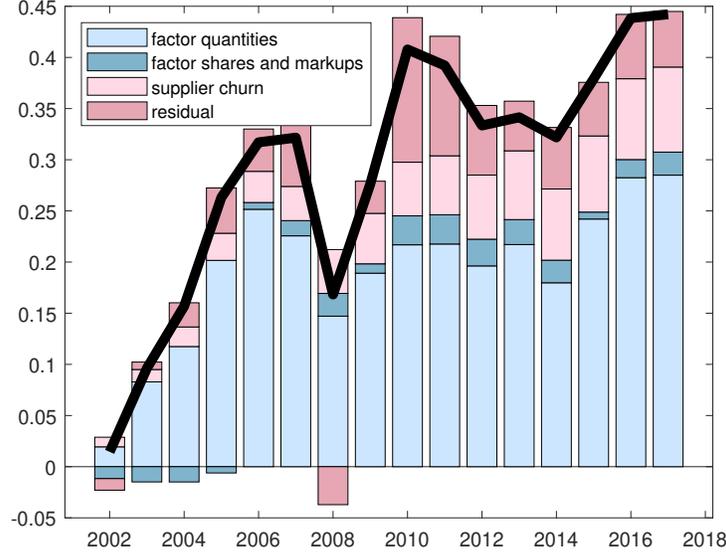
(b) Firm births and deaths

Table A13 shows how the contribution of the supplier churn term changes for different values of $\bar{\delta}^{add}$ and $\bar{\delta}^{sep}$. The top panel is quite sensitive to differences between $\bar{\delta}^{add}$ and $\bar{\delta}^{sep}$, whereas the bottom panel is not. To understand this, consider the following decomposition of the supplier churn term:

$$\bar{\delta}^{add} \Delta \mathcal{E} - \bar{\delta}^{sep} \Delta \mathcal{X} = \left[\frac{\bar{\delta}^{add} + \bar{\delta}^{sep}}{2} \right] [\Delta \mathcal{E} - \Delta \mathcal{X}] + [\bar{\delta}^{add} - \bar{\delta}^{sep}] \left[\frac{\Delta \mathcal{E} + \Delta \mathcal{X}}{2} \right],$$

where $\Delta \mathcal{E}$ and $\Delta \mathcal{X}$ are the Domar-weighted additions and separations across all downstream firms. This decomposition shows that the average level of $\bar{\delta}^{add}$ and $\bar{\delta}^{sep}$ matters to the extent that additions exceed separations. On the other hand, the difference between $\bar{\delta}^{add}$ and $\bar{\delta}^{sep}$ matters to the extent that the gross level of additions and separations is high. In the top panel of Table A13, we consider all additions and separations but in the bottom panel, we only consider additions and separations due deaths and births. Since the latter is lower than the former, this explains the differential sensitivity to the value of $\bar{\delta}^{add} - \bar{\delta}^{sep}$.

Figure A1: Growth accounting allowing for heterogenous δ



Notes: Following column (ii) of Table 4, we set $\delta_{ijt} = 0.505 - 0.121 \times \log \text{sales ratio of supplier}_{jt}$ for both additions and separations. The results of this figure are quite similar to Figure 5.

Appendix F Monte Carlo Simulations

In this appendix we report results when we run regression (11) on artificial data. We use the cost function introduced in Example 2. The marginal cost for downstream firm i is $mc_i = A_i^{-1} \tilde{m}c_i$, where A_i is a Hicks neutral productivity shifter and $\tilde{m}c_i$ solves

$$\sum_{j=1}^M \frac{\omega_{ij}}{\sigma_j - 1} \left(\frac{p_{ij}}{\tilde{m}c_i} \right)^{1-\sigma_{ij}} = \sum_{j=1}^M \frac{\omega_{ij}}{\sigma_{ij} - 1}.$$

The scalars ω_{ij} and σ_{ij} are parameters of firm i 's cost function and M is the number of potential suppliers. Inputs that are unavailable to firm i have infinite price. The spending share on supplier j by firm i is

$$\Omega_{ij} = \frac{\omega_{ij} (p_{ij} / \tilde{m}c_i)^{1-\sigma_{ij}}}{\sum_k \omega_{ik} (p_k / \tilde{m}c_i)^{1-\sigma_{ik}}}.$$

We parameterize σ_{ij} and ω_{ij} as follows so that we can control the correlation between spending shares on each input and the consumer surplus ratio of that input.

Firm i draws random variables ϵ_{kij} for $j = \{1, \dots, M\}$ and $k = 1, 2, 3$ that are uniformly distributed in the interval $[0, r_k]$. We set $\sigma_{ij} = \bar{\sigma}^{sep} + \epsilon_{1ij} + \epsilon_{2ij}$ for $j = \{1, \dots, M/2\}$, and $\sigma_{ij} = \bar{\sigma}^{add} + \epsilon_{1ij} + \epsilon_{2ij}$ for $j = \{M/2 + 1, \dots, M\}$. We set the parameters determin-

ing spending shares on each input as follows: $\tilde{\omega}_{ij} = \epsilon_{3ij} + \kappa\epsilon_{2ij}$, $\bar{\omega}_{ij} = \tilde{\omega}_{ij} / \sum_j \tilde{\omega}_{ij}$, and $\omega_{ij} = \bar{\omega}_{ij}(p_{ij}/\tilde{m}c_i)^{\sigma_{ij}-1}$. If $\kappa = 0$, spending shares are uncorrelated with σ_{ij} . If $\kappa < 0$, spending shares are negatively correlated with σ_{ij} .

Inputs $j = \{1, \dots, M/2\}$ are available in the first period, and each input has probability ρ^{sep} of becoming unavailable in the second period. All inputs $j = \{M/2 + 1, \dots, M\}$ are available in period 2, and each input has probability ρ^{add} of being unavailable in the first period. Hence, ρ^{sep} and ρ^{add} control the fraction of separating inputs and the fraction of added inputs between the first and second period. All available inputs in the first period have price equal to one. Available inputs in the second period have log-normally distributed prices with standard deviation σ^p . For each firm, changes in Hicks-neutral productivity are log-normally distributed price with standard deviation σ^A .

In our simulations, we set $M = 200$ which is close to number of suppliers for the average downstream firm. We set $\bar{\sigma}^{sep}$ and $\bar{\sigma}^{add}$ so that, conditional on the other parameters, the average δ is 0.268 for separating suppliers and 0.283 for added suppliers (consistent with our baseline estimates). We set $\rho^{sep} = 0.01$ and $\rho^{add} = 0.01$ so that the average separation and addition shares are 0.005, which is similar to the variable cost share of entering and exiting suppliers in the Prodcom sample. We set the upper bound of the uniform distribution r_1 so that the range of δ across inputs (within each of the addition and separation sets) is 0.1. We set $r_2 = 1$ without loss since we rescale the input shifters $\tilde{\omega}_{ij}$. We set $r_3 = 1$ so that the correlation between δ and cost shares Ω across inputs is 0.5 if $\kappa = -1$ and -0.5 if $\kappa = 1$. Across firms, the correlation between separation or addition share and average δ for separating or added inputs is 0.28 if $\kappa = -1$ or -0.28 if $\kappa = 1$. We report results for three sets of values of σ^p and σ^A : (i) $\sigma^p = \sigma^A = 0$, (ii) $\sigma^p = \sigma^A = 0.01$, and (iii) $\sigma^p = \sigma^A = 0.02$. We consider 100 simulations, and for each simulations draw artificial data for 35,000 firms (roughly the number of observations in our regressions). We run regression (11) without instrumenting because additions and separations are exogenous in our simulations. Table A14 reports percentile estimates across the 100 simulations.

Motivated by Proposition 2, we first consider the case where average δ firm is uncorrelated with the addition and separation shares. Columns (i)-(iii) show that the estimated coefficients are very close to the true average δ for additions and separations. They are not exactly equal because of the small errors from the first-order approximation. As expected, the sampling uncertainty of the estimates is increasing when we increase the standard deviation of productivity and continuing price shocks. The remaining columns show that, when addition and separation shares are systematically correlated with average δ , violating one of the assumptions in Proposition 2, the estimated coefficients are biased. However, for the median estimate the bias is quite small (it is of the same order as the variation

induced by sampling uncertainty).

Table A14: Monte Carlo simulations

	(i)	(ii)	(iii)	(iv)	(v)	(vi)	(vii)	(viii)	(ix)
Correlation δ, Ω		Zero			-0.5			+0.5	
Std. dev. A, p shocks	0	0.01	0.02	0	0.01	0.02	0	0.01	0.02
Addition share									
$\mathbb{E}[\bar{\delta}^{add}]$	0.283	0.283	0.283	0.283	0.283	0.283	0.283	0.283	0.283
Median estimate $\hat{\delta}^{add}$	0.285	0.286	0.287	0.274	0.273	0.273	0.297	0.297	0.297
5th percentile estimate	0.285	0.270	0.257	0.273	0.258	0.246	0.296	0.275	0.263
95th percentile estimate	0.286	0.302	0.318	0.274	0.290	0.308	0.298	0.313	0.325
Separation share									
$\mathbb{E}[\bar{\delta}^{sep}]$	0.268	0.268	0.268	0.268	0.268	0.268	0.268	0.268	0.268
Median estimate $\hat{\delta}^{sep}$	0.271	0.270	0.272	0.260	0.259	0.261	0.280	0.280	0.281
5th percentile estimate	0.270	0.256	0.243	0.259	0.244	0.226	0.280	0.265	0.245
95th percentile estimate	0.271	0.286	0.296	0.260	0.277	0.294	0.281	0.296	0.318

Notes: Table reports Monte Carlo statistics from 100 simulations with a sample of 35,000 firms in each simulation. The value of $\mathbb{E}[\bar{\delta}^{add}]$ and $\mathbb{E}[\bar{\delta}^{sep}]$ are unweighted averages of the true δ 's for additions and separations. The estimates $\hat{\delta}^{add}$ and $\hat{\delta}^{sep}$ are for regression (11), with percentiles calculated across the 100 simulations. Details of the calibration are in the text.