Micro PPI-Based Real Output Forensics*

Timo Boppart[†] Mikael Carlsson[‡] Markus Kondziella[§] Markus Peters[¶]

May 11, 2023

Abstract

We analyze producer price index micro data on total private goods and services production in Sweden to quantify the implications of methods of price index construction on the measured aggregate inflation rate. We document large quantitative effects of different methods of lower-level aggregation, i.e., the aggregation of price changes of different products into an index at the 5-digit product group level. Moving from an arithmetic index to geometric averaging across items decreases annual goods and services inflation by 0.5 and 0.4 percentage points, respectively. We contrast the results of these statistical indices with an economic theory-based index relying on a nested-CES structure. Estimating elasticities of substitution across goods within industries implies that such a theory-based index results in an annual inflation rate that is 3.9 and 3.1 percentage points lower for goods and services, respectively. Our results pose a challenge for the comparability of inflation rates and real output growth rates across countries as well as a tension between (economic) theory and (statistical) measurement. Therefore, we recommend that statistical offices report three moments for each product group instead of the published single index number. This would allow users to approximate any nested-CES index under the assumption of a joint log-normal distribution of price growth factors and weights.

Keywords: Output growth, inflation measurement, price index, PPI micro data, Baumol's cost disease.
JEL classification: E23, E31, N14, O47.

* We are grateful to Marcus Fridén at Statistics Sweden for sharing his insights on the construction of the goods and services PPI index. We thank Filip Mellgren for excellent research assistance. Boppart thanks Vetenskaprådet (grant 2016-02194) and the Global Center at Unisg, Carlsson thanks Ragnar Söderbergs Stiftelse (grant E52-13), and Handelsbankens Forskningsstiftelser (grant P18-0082), and Kondziella and Peters thank Jan Wallander's and Tom Hedelius' foundation and Tore Browaldh's foundation for financial support. We thank seminar participants at the EEA 2020, SED 2021, Boston University, Helsinki Graduate School of Economics, Hitotsubashi University, Konjunkturinstitutet and Uppsala University for comments. The views expressed in this article is solely the responsibility of the authors and should not be interpreted as reflecting the views of Sveriges Riksbank.

[†]Institute for International Economic Studies (IIES), University of St. Gallen, and CEPR. E-mail: timo.boppart@iies.su.se

[‡]Uppsala University, UCLS and Sveriges Riksbank. E-mail: mikael.carlsson@nek.uu.se

[§]IIES. E-mail: markus.kondziella@iies.su.se

[¶]IIES. E-mail: markus.peters@iies.su.se

1 Introduction

Comparing economic growth over time and across space requires accurately measuring real output. The standard approach to constructing such a measure is first to collect information on output at current prices and then to deflate this nominal output by an appropriately defined price index. Although the collection of data on nominal output requires significant resources from statistical offices, its measurement is—at least in advanced economies—relatively straightforward. It is indeed the construction of a price deflator that poses the main challenge of real output computation. In addition to being a key component for measuring real output, price indices also play a direct role in economic policy, e.g., as an (inflation) target to guide monetary policy or as the basis to index social security systems and tax brackets. Furthermore, headline price index numbers have direct consequences as they serve as a reference point for collective wage bargaining.

In this paper, we shed light on the details that go into price index construction and their quantitative consequences for the measurement of inflation and output. The price index construction is the mapping of a surveyed panel of prices into a number of the aggregate price level. If prices and sales weights of the different items were observed continuously, the choice of the index number would be straightforward: a Divisia (1926) index would be the obvious answer. In reality, however, prices and weights are collected by the statistical office at fixed intervals.[1] This makes the choice of the precise index non-trivial. In this paper, we show that the choice matters quantitatively.

In practice, the index number aggregation is done in two steps: First, prices of narrowly defined product groups (the elementary group level) are aggregated into an index number, called *lower level aggregation*. These product group-level indices are the most disaggregate numbers published by the statistical office. Then, the elementary level indices are further aggregated in the *higher level aggregation* to arrive at the aggregate producer price index (PPI).

This paper analyzes the micro data of the Swedish goods producer price index and the more recently established services producer price index. Our main focus is on the period 2004–2019 (goods sector) and 2013–2019 (service sector), for which we observe the complete picture of the underlying micro data on prices and weights. We replicate the official index numbers published by Statistics Sweden and trace all the assumptions that went into their construction. Importantly, the completeness of the data allows us to change assumptions underlying the index construction and run counterfactual experiments. Finally, by aggregating up, we can quantify the effects of these counterfactuals for

[1] In Sweden, for the producer price index (PPI), prices are reported monthly for goods and quarterly for services, whereas the (sales) weights are updated once a year. In the U.S. the PPI weights are only updated every five years.

measured inflation *at the aggregate level*.

We first focus on the methods underlying the aggregation of prices of different product items into an index at the elementary group level (lower-level aggregation). The method used for lower-level aggregation varies across countries. The Swedish goods PPI is constructed using a weighted arithmetic mean of item price changes with annually updated weights. In other countries, the same lower level aggregation in goods PPI is done instead with an unweighted average or by using a geometric average, or both. A geometric mean will result in a lower inflation rate than an arithmetic mean, but we can quantify this effect in the Swedish data. We change the method of lower-level aggregation accordingly in the micro data, leave the higher-level aggregation unchanged, and compute a counterfactual index. Applying a weighted geometric mean at the lower level (as done in Denmark) results in an, on average, 0.48 pp. lower aggregate inflation rate per year compared to a weighted arithmetic mean. When moving to an unweighted arithmetic mean (a so called Carli index as used in Spain or Slovenia), the aggregate goods PPI inflation rate is, instead, on average 0.49 pp. *higher* per year. Indeed, in the Swedish data, product items with a higher (sales) weight have a systematically lower inflation rate. We document similar effects for services. Interestingly, Statistics Sweden uses a weighted *geometric* mean for the lower-level aggregation in services. Moving to an arithmetic mean would result in an, on average, 0.40 pp. higher aggregate inflation rate per year. As a next step, we run counterfactuals in which we change the updating of the weights at the lower level. If the weights in the Swedish goods PPI are updated instead of annually only every five years as, e.g., done in the U.S. the average aggregate inflation rate increases by 0.22 pp. per year. Overall, the results on the counterfactuals we run at lower level aggregation put a question mark on the international comparability of inflation and real growth rates.

We next turn to the higher level aggregation, i.e., how the different elementary-level indices are aggregated up to an aggregate number. Compared to lower-level aggregation, the method for higher-level aggregation does not differ as widely across countries. Weights are applied everywhere and arithmetic averaging seems to be the norm. Changing the weighted arithmetic mean used in Sweden to a weighted geometric mean lowers the calculated inflation rate at the aggregate level by an average of 0.54 pp. per year.

We contrast the results from the statistical indices with an economic theory-based approach in which a nested-CES structure is imposed, the elasticities are calibrated or estimated, and the ideal price index and the inflation therein are computed. In line with the two-tier structure of the statistical indices, we consider a nested-CES specification with one elasticity of substitution across different product groups and different elasticities within each of the 318 elementary groups in the case of the goods PPI and the 298 groups in the services PPI. We show that the statistical indices are obtained as special cases within this nested-CES structure. The ideal theory-based index corresponds to a weighted arithmetic mean if the elasticities of substitution are set equal to zero (perfect complements), whereas the weighted geometric mean obtains with unitary elasticities (Cobb-Douglas). By Jensen's inequality, the computed inflation rate of an ideal index is monotonically decreasing in the elasticities of substitution. This effect is bigger the larger the dispersion of price growth factors. We quantify this effect again expressed at the aggregate level in the Swedish micro data. As there is much dispersion in price growth across items in the data within an elementary group in both goods and services, we find that the results are sensitive to the choice of the elasticity of substitution. Leaving the elasticity of substitution across groups at zero (implying arithmetic averaging as currently done by Statistics Sweden) but increasing the elasticity of substitution across items within each group to 2, 4, or 8 decreases the annual inflation rate compared to the baseline with an elasticity of zero by 0.97, 1.94 and 3.77 pp. respectively for the aggregate goods PPI. For services, moving the elasticity of substitution across items within an elementary group to a value of 2, 4, or 8 decreases the average aggregate inflation rate per year by 0.81, 1.87 and 4.67 pp. respectively.

As a next step, we estimate elasticities of substitution across items for each elementary level with the estimation procedure as in Hottman et al. (2016). We find a median elasticity of substitution of 4.37 for goods and a first and ninth decile of 2.25 and 12.02, respectively, which is well within the ballpark of other estimates in the literature. For services, we estimate comparable elasticities with a median elasticity of substitution of 5.51 and a first and ninth decile of 1.78 and 13.60. In the case of services, systematic estimates of elasticities of substitution are largely missing in the literature. We see it as one of our contributions to provide such estimates for all service product categories in Sweden. For the elasticity of substitution across product groups, we obtain estimates close to one; 1.24 in the case of goods and 1.07 for services. Hence across groups of goods and across groups of services a Cobb-Douglas assumption seems to be a good approximation. When using the estimated elasticities at the lower and higher level in the nested-CES index computation, the implied inflation rate at the aggregate level decreases compared to the baseline by, on average, 3.94 pp. per year for goods and by 3.11 pp. per year for services. Hence these counterfactual indices would imply a significant overall *deflation* in both goods and services of the considered time periods. The large difference between the economic theory-based price index and the statistical indices used by statistical offices raises issues of mapping economic theory to measurement.

Finally, we investigate the extent of a cost disease à la Baumol (1967) present in Sweden within the goods and the service sector. This is done by fixing product group weights at their initial level in the first sample year instead of annual updating as done in the official PPI. This counterfactual results in a higher average inflation rate of 0.20 pp. and 0.15 pp. per year for goods and services, respectively. Hence, we find no Baumol effect across the Swedish goods-producing industries or across the Swedish service-producing industries.[2]

In sum, our results show that because price growth rates are in the data very dispersed across items within a product group, relatively minor differences in elasticities of substitution have a large impact on the implied price index. This implies a tension for research that aims to map the growth rate in a theory to the theoretically consistent counterpart in terms of measurement. It also poses a dilemma for statistical agencies to decide which index number(s) to publish. As a concrete suggestion, we recommend statistical offices publish three moments of the distribution of items' price growth and weights within each product group (the unweighted mean across items of the log of the gross price growth factor, the unweighted variance across items of the log of the gross price growth factor, and the covariance across items between the log of the gross price growth factor and the log of the sales weight). This solution is practical because it only requires publishing these three moments for each elementary group. Under the assumption that price growth factors and weights follow a joint log-normal distribution the moments are sufficient for computing exact indices for any given distribution of elasticities of substitution. We show that the error from making this assumption is relatively small, ranging between 0.15–0.21 pp. per year in absolute terms.

The rest of the paper is structured as follows. Section 2 presents how price indices are computed in practice and formalizes an economic theory approach to index computation. Section 3 provides a description of the PPI micro data and summary statistics in terms of price growth factors and weights. Section 4 reports results from changing the applied method of index construction. Finally, Section 5 concludes the paper.

[2]This result is reminiscent of the close-to-unitary estimates of the elasticity of substitution across product groups in both goods and services. We do not estimate the elasticity of substitution between goods and services. This elasticity might still be significantly below one giving the Baumol's cost disease potentially some bite between goods and services.

1.1 Related literature

Price indices originate from the attempt to measure changes in the cost of fixed consumption baskets relative to a base period, where the cost of the consumption basket is computed as the product of a price- and a (fixed) quantity vector. This approach dates back to at least Lowe (1824).[3] Whereas Laspeyres (1871) suggested using fixed quantities from the base period, Paasche (1874) proposed to use current period quantities (Laspeyres or Paasche price indices, respectively). Fisher (1922) instead suggested using the geometric mean of the Laspeyres and Paasche price index. Based on the view that (except for random fluctuations) prices increase proportionately with the money supply, Jevons (1884) and Gian Rinaldo Carli suggested using geometric and arithmetic averages of item price changes as the respective price index.

Yet another approach to index construction relies on economic theory, particularly on consumers' or producers' expenditure or cost-minimizing behavior. For a given set of preferences, the price index is implicitly defined as the expenditure change for the basket of goods that obtains a fixed utility level under minimal costs (Bennet, 1920; Konüs, 1939; Lloyd, 1975; Moulton, 1996). Under this approach, the quantities chosen for each product result from an economic optimization problem, whereas product prices are considered as given. For the case of constant elasticity of substitution (CES) preferences, Sato (1976)-Vartia (1976) derive the implicit price index that has spurred a literature estimating the preference parameters that underlie the CES price index (Feenstra, 1994; Broda and Weinstein, 2006, 2010; Aghion et al., 2019).[4] Building on Hottman et al. (2016), we estimate the preference parameters of a CES price index in Section 4.4.

The different approaches to index construction are directly related. For example, Konüs and Byushgens (1926) show that the implicit price index under Leontief preferences is a (weighted) arithmetic mean of price changes.[5] Similarly, they establish equivalence between the implicit price index under Cobb-Douglas preferences and a (weighted) geometric Jevons index.[6] In Section 2.2, we show how a two-tier structure of CES preferences nests different price indices used by statistical agencies in practice.

The different price index counterfactuals we run speak to the literature that has quantified the bias

[3]See Diewert (1988, 1995, 2021) for an excellent overview of the history of price index research.

[4]For recent examples of the economic theory approach to price index construction, see Martin (2020) or Redding and Weinstein (2020) who develop the implicit price index under CES preferences with taste shocks.

[5]The weights correspond to expenditure shares. In return, a Laspeyres price index can be reformulated as a weighted average of price changes.

[6]For the price indices derived from economic theory, Diewert (1976) further shows that a quadratic mean of order r approximates arbitrary twice continuously differentiable linearly homogeneous preferences to the second order.

in aggregate inflation arising from substitution between products. Moulton and Smedley (1995) use a sample that covers 96% of the U.S. CPI between 1992–1994 and find a difference of 0.49 pp. per year between an arithmetic and a geometric average of item price changes at the lower level. We find a similar difference for Sweden's good PPI between 2004 and 2019 and the service PPI between 2013–2019. In line with Moulton and Smedley (1995), Reinsdorf and Moulton (1996) use a sample covering 70% of the U.S. CPI 1992–1993 and find a difference of 0.47 pp for the same counterfactual. Boskin et al. (1996) revise down the number in Moulton and Smedley (1995) and report a lower-level substitution bias of 0.25 pp. For the upper level, Boskin et al. (1996) quantify the bias at 0.15 pp. (difference between a Laspeyres and Törnqvist index). Interestingly, already Braithwait (1980) elaborates that the substitution bias of a statistical index increases with the substitutability between products and compares measures of price dispersion with the elasticity of substitution for different product categories for the U.S. CPI.

2 Index computation

Producer price indices are commonly constructed in two steps. First, prices of individual items within the same product group are aggregated to form an elementary price index. In a second step, elementary price indices are aggregated higher up across product groups to build an aggregate price index. For the first step of aggregating item prices at the elementary level, generally one of three aggregation methods is used in practice: Arithmetic mean, geometric mean or the ratio of arithmetic means. If available, item specific weights can be applied using either method. For an item i, let the item price in period t be denoted by $p_{t,i}$. Moreover, let the item price and weight in a given base period be denoted by $p_{0,i}$ and $w_{0,i}$, respectively. As weights are typically updated less frequently than prices, we focus on the case of practical relevance in which weights are only updated in the base period. The three aggregation methods, including a specification of countries using each method, are then formally given by the list below.[7]

1. Arithmetic mean
   - Unweighted (Carli): $\frac{1}{n}\sum_i \frac{p_{t,i}}{p_{0,i}}$
     Greece, Hungary, Portugal, Slovenia, Spain
   - Weighted: $\sum_i w_i \frac{p_{t,i}}{p_{0,i}}$
     Australia, Czech Republic, France, Germany, Ireland, Israel, Korea, Luxembourg, New Zealand, Slovak Republic, **Sweden**, Turkey, U.S.

[7]Source: http://www.oecd.org/sdd/prices-ppp/48370389.pdf (OECD, 2011)

2. Geometric mean

   - Unweighted (Jevons): $\prod_i \left(\frac{p_{t,i}}{p_{0,i}}\right)^{\frac{1}{n}}$
     Austria, Chile, Finland, Italy, Netherlands, Norway, Switzerland
   - Weighted: $\prod_i \left(\frac{p_{t,i}}{p_{0,i}}\right)^{w_i}$
     Denmark, Japan

3. Ratio of arithmetic means:

   - Unweighted (Dutot): $\frac{\frac{1}{n}\sum_i p_{t,i}}{\frac{1}{n}\sum_i p_{0,i}}$
     Estonia, Poland, U.K.
   - Weighted: $\frac{\sum_i w_i p_{t,i}}{\sum_i w_i p_{0,i}}$
     Canada

Even though most countries favor the weighted arithmetic mean, all three aggregation methods are used in practice. At the second step when elementary price indices are aggregated across product groups, all three methods could potentially be applied as well. However the arithmetic mean is generally used to aggregate elementary level price indices up to the highest level.

In Sweden, the official PPI for goods is given by first averaging item price changes within product groups arithmetically to obtain the growth factor of each elementary level price index which in a second step are averaged arithmetically across all product groups. For services, elementary indices are instead computed using a geometric average, while arithmetic averaging is also used at the second step. This yields an aggregate growth factor. Item sales and product group sales are used as weights. The aggregate index is built from its growth factors by chaining annual aggregate price growth factors from December to December in two consecutive years and normalizing by the average growth factor in a base year, which currently is set to 2015. Multiplying by 100 yields the actual index.

There are two alternative approaches to evaluate which aggregation method to use in practice. One is based on axiomatic properties of price indices and the other is based on economic theory. With the first approach one selects the aggregation method that fulfills desirable properties of a price index, whereas with the second approach the exact price index is derived from an economic model. The next two sections relates these two approaches when the economic model takes the form of a nested CES structure.

2.1 Axiomatic approach

The International Monetary Fund states four desirable properties of price indices in the producer price manual, IMF (2004):

1. Proportionality: multiplying all prices by a factor of x in period t changes the index by a factor of x in period t.
2. Changes in the units of measurement: changing the units of measurements of a single item leaves the index unaffected.
3. Time reversal: computing the index for period t with base period 0 results in the inverse index of computing the index in period 0 with base period t.[8]
4. Transitivity in chaining: the direct index growth between period 0 and t is identical to the chained index growth between the periods.[9]

Which of the four axioms is satisfied with each aggregation method is listed in Table 1. The geometric mean is the only aggregation method that fulfills all four axioms. IMF (2004) (p. 219) summarizes:

"*It may be concluded that from an axiomatic viewpoint, both the Carli [arithmetic mean] and the Dutot [ratio of arithmetic means] indices, although they have been and still are widely used by statistical offices, have serious disadvantages. [...] From an axiomatic point of view, the Jevons index [geometric mean] is clearly the index with the best properties, even though it may not have been used much until recently.*"

Axiomatic properties of price indices

	$\frac{\sum_i w_i p_{t,i}}{\sum_i w_i p_{0,i}}$	$\sum_i w_i \frac{p_{t,i}}{p_{0,i}}$	$\prod_i \left(\frac{p_{t,i}}{p_{0,i}}\right)^{w_i}$
Proportionality	✓	✓	✓
Change in units	✗	✓	✓
Time reversal	✓	✗	✓
Transitivity	✓	✗	✓

Table 1: Source: IMF (2004), Producer Price Index Manual

[8] To illustrate the time reversal property assume that March prices are up 20% relative to January prices. With base period January (= 100) the price index in March stands at 120. Computing the index with base period March (= 100), a price index that fulfills the time reversal property stands in January at $\frac{100}{1.2}$.

[9] The transitivity in chaining axiom implies that linking monthly price growth factors within a year yields the same inflation rate as linking Dec,t prices to Jan,t prices directly.

2.2 Economic theory-based approach

As an alternative to the axiomatic approach, we derive an exact price index from economic theory.[10] Suppose the economy can be described by a nested CES-structure where individual items are aggregated at the product group level to form elementary level output and product groups are aggregated further up to form total output of the economy. In particular, total output of the economy Y and production at the product group level, Y_g, are described by

$$Y = \left(\sum_{g \in G} \gamma_g^{\frac{1}{\sigma}} Y_g^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} \quad \text{and} \quad Y_g = \left(\sum_{i \in I_g} \gamma_i^{\frac{1}{\sigma_g}} y_i^{\frac{\sigma_g-1}{\sigma_g}} \right)^{\frac{\sigma_g}{\sigma_g-1}}, \tag{1}$$

where y_i denotes item i with price p_i and technology weight γ_i and σ_g denotes the elasticity of substitution between items within a product group. I_g is the set of items in product group g and G is the set of product groups. γ_g and σ denote the technology weight of product group g and the elasticity of substitution between product groups. Cost minimization at the product group level leads to the known expression for the elementary level exact price index

$$P_g = \left(\sum_{i \in I_g} \gamma_i p_i^{1-\sigma_g} \right)^{\frac{1}{1-\sigma_g}} \tag{2}$$

with sales shares w_i given by

$$w_i = \gamma_i \left(\frac{p_i}{P_g} \right)^{1-\sigma_g}. \tag{3}$$

We are interested in the gross growth rate of the exact price index. Introducing time subscripts to all variables except the elasticities of substitution, this is given by

$$\frac{P_{t,g}}{P_{0,g}} = \frac{\left(\sum_{i \in I_g} \gamma_{0,i} p_{t,i}^{1-\sigma_g} \right)^{\frac{1}{1-\sigma_g}}}{\left(\sum_{i \in I_g} \gamma_{0,i} p_{0,i}^{1-\sigma_g} \right)^{\frac{1}{1-\sigma_g}}}. \tag{4}$$

Since we observe sales shares, $w_{0,i}$, instead of technology weights, $\gamma_{0,i}$, we use (3) to substitute $\gamma_{0,i}$ out of (4). We obtain the following expression for the growth rate of the exact price index at the product group level

$$\frac{P_{t,g}}{P_{0,g}} = \left(\sum_{i \in I_g} w_{0,i} \left(\frac{p_{t,i}}{p_{0,i}} \right)^{1-\sigma_g} \right)^{\frac{1}{1-\sigma_g}}. \tag{5}$$

As we observe nominal prices and sales shares in the data, for a given value of σ_g, we can compute equation (5) using the PPI micro data. One counterfactual exercise that we run is to aggregate item prices using (5) for different values of the elasticity of substitution. For $\sigma_g = 0$ and $\sigma_g = 1$ equation (5) nests the arithmetic and geometric price aggregation as special cases:

[10] In particular, a so-called Loyd-Moulton index is exact for the nested CES structure of the economic model that we use here (Lloyd, 1975; Moulton, 1996)

1. $\sigma_g = 0$: Leontief production and the price index becomes the weighted arithmetic mean

$$\frac{P_{t,g}}{P_{0,g}} = \sum_{i \in I_g} w_{0,i} \left(\frac{p_{t,i}}{p_{0,i}} \right) \tag{6}$$

2. $\sigma_g = 1$: Cobb-Douglas production and the price index becomes the weighted geometric mean

$$\frac{P_{t,g}}{P_{0,g}} = \prod_{i \in I_g} \left(\frac{p_{t,i}}{p_{0,i}} \right)^{w_{0,i}} \tag{7}$$

Equation (5) is at the product group level. To aggregate the elementary price indices to an aggregate price index, note that expenditure minimization at the aggregate level leads to the following expression for the aggregate price index

$$P = \left(\sum_{g \in G} \gamma_g P_g^{1-\sigma} \right)^{\frac{1}{1-\sigma}}. \tag{8}$$

Substituting out product group weights by sales shares, W_g, and introducing time indexation as before we obtain aggregate price index growth of

$$\frac{P_t}{P_0} = \left(\sum_{g \in G} W_{0,g} \left(\frac{P_{t,g}}{P_{0,g}} \right)^{1-\sigma} \right)^{\frac{1}{1-\sigma}}. \tag{9}$$

Equation (9) is the equivalent of (5) at the aggregate level. At this level of aggregation, $\sigma = 0$ implies arithmetic and $\sigma = 1$ geometric aggregation of product group price index growth factors. To compute the theory-based PPI we take equations (5) and (9) to the data.

Note that the CES price index in equation (5) fulfills the proportionality and changes in units of measurement axiom for any value of σ. For the time reversal and transitivity in chaining axiom such general statements cannot be made as the weighted arithmetic mean aggregation (as a subcase of the CES price index) fails the time reversal and transitivity axiom.

2.3 Economic theory-based approach when price growth factors and sales shares follow a joint log-normal distribution

As shown above, the nested CES price index replicates two common indices used in pactice—the weighted arithmetic and the weighted geometric mean—for elasticities of substitution equal to zero and unity. For any other values of elasticities of substitution, the nested CES price index will give a different result than the official indices. A natural question is then whether one must access the micro data underlying the official price index in order to compute the aggregate index for other elasticity values. Proposition 1 states that, under the assumption that item-specific price growth rates and weights follow a joint log-normal distribution, only three moments are needed to compute the elementary-level price index (5).

Proposition 1. *Suppose that price growth factors* $\frac{p_{t,i}}{p_{0,i}}$ *and weights* $w_{0,i}$ *follow a joint log-normal distribution. The price index (5) for product group* g *is then given by*

$$\frac{P_{t,g}}{P_{0,g}} = \exp\left(\mu_{t,g} + \frac{1-\sigma_g}{2}\delta_{t,g} + \rho_{t,g}\right) \tag{10}$$

where $\mu_{t,g}$ *denotes the mean of* $\ln\left(\frac{p_{t,i}}{p_{0,i}}\right)$, $\delta_{t,g}$ *the variance of* $\ln\left(\frac{p_{t,i}}{p_{0,i}}\right)$, *and* $\rho_{t,g}$ *the covariance between* $\ln\left(\frac{p_{t,i}}{p_{0,i}}\right)$ *and* $\ln(w_{0,i})$, *where* i *denotes and item belonging to group* g.

See Appendix D for a proof. The first statistic $\mu_{t,g}$ is the expectation of log price growth factors. Note that the exponential of the finite sample counterpart of this statistic is equal to the geometric sample mean with equal weights $1/N_{t,g}$, where $N_{t,g}$ is the number of items i in group g at period t. If the statistical agency applies an unweighted geometric average to compute the price index, then the first statistic is already provided. In the Swedish case, Statistics Sweden needs to publish this first statistic in addition to the group indices already published, since the group indices are either weighted arithmetic averages as in the case of goods or weighted geometric averages as in the case of services.

If price growth factors and sales shares are assumed to be independent, such that they have a covariance $\rho_{t,g}$ equal to zero, only one additional statistic in the form of the variance of log price growth factors, $\delta_{t,g}$, is needed to compute the price index. Importantly, the variation in price growth rates determines how the elasticity of substitution quantitatively alters the price index, since the variance is interacted by the term $1-\sigma_g$. When price growth factors and sales shares are dependent, the statistical agency should also publish the covariance of their logarithms.

Having indices on the lower level as computed in (10), the aggregate index is given by using equation (9).

In Section 4, we provide an investigation into how sensitive the exact aggregate price index is to various levels of elasticities of substitution as well as to the assumption of a joint log-normal distribution of price growth factors and weights. Before doing so, the next section describes the micro data underlying the goods and services PPI in Sweden.

3 Data

The Swedish statistical agency Statistics Sweden collects information on prices in order to construct several price indices. On the most disaggregated level, a data observation shows monthly updated prices of an item that is produced by a particular firm, classified as a type of product as given by an 8-digit Combined Nomenclature product code (CN8).[11] We use the total scope of these micro data,

[11] The item level is indeed detailed since it represents an observation at a more disaggregated level than the CN8 product code. As mentioned in Carlsson and Nordström Skans (2012), "an example of a product code is 84181010 for a combined freezer and cooler with separate exterior doors with a volume exceeding 340 liters intended for use in civilian aircrafts".

which can be aggregated to compute the goods producer price index (PPI), the services PPI, and the import price index. We refer to these micro data as the *goods PPI micro data* and the *services PPI micro data*, covering the private markets for goods and services, respectively. These data are key inputs for measuring real GDP since they are used to deflate nominal value added as measured by the output approach in the national product accounts. This section provides a description of the micro datasets on prices and reports summary statistics in the form of frequencies and sizes of price changes as well as price growth factors and weights. As we are interested in domestic production, we focus on prices of items sold domestically or as exports. For more detailed descriptions and statistics, see Appendices A and B. In addition to the micro data on prices, we access micro data used to measure sales in the production accounts. We describe these data in more detail when applying them to estimate elasticities of substitution in Section 4.4.

3.1 Coverage

The target population of the goods and services PPI consists of all transactions of domestic production as well as foreign production imported to Sweden. The PPI micro data include price information on domestic and foreign production of, for example, agriculture, materials, energy, manufacturing, transportation, information and communication, real estate, as well as professional and scientific services such as legal and management services, and personal services provided by, e.g., restaurants and hairdressers. However, there are a couple of sectors that are not covered. These include most public services, i.e., health care, education, national defence, and social insurance. Also, price information for the construction sector and partly the financial sector is not collected. Although prices in these sectors are not directly measured, the PPI micro data are still used for measuring their real output and value added for the national accounts. For example, deflation could be done by imputation, e.g., by using the GDP-deflator, and whenever production is measured through input costs, which is the case for several public services, the PPI micro data are used to deflate such costs. Hence the PPI micro data are used to deflate nominal output in virtually all sectors of the economy.

Due to the broad scope of the population, the sample frames are defined by four other comprehensive surveys containing product-level information: The industrial production survey (Industrins varuproduktion, IVP), the firm income statements and balance sheets survey (Företagens ekonomi, FEK), as well as surveys on foreign trade of goods (Utrikeshandel med varor, UHV) and foreign trade of services (Utrikeshandel med tjänster, UHT). For example, IVP is used as sampling frame for the goods PPI and covers all firms with at least 20 employees for most sectors. For some sectors, the firm cutoff is instead at 10 employees and 75 million SEK ($\approx$ 7.9 million USD) in revenue. FEK, used as sampling frame for the services PPI, covers the universe of Swedish private firms (except the financial sector) through register data. In addition, FEK consists of a survey asking for detailed information on income statements are sent out to about 16,000 firms.

Before drawing the samples, the various data sources above are combined with additional information from the firms in question to construct a measure of sales of a given product that is produced by a given firm and sold in a given market, i.e. domestically or as exports. Arriving at the correct sales measure is an intricate process. Consider the case of measuring domestic sales of a domestically produced good by a given firm. The firm's total sales of domestic production of this good is reported in IVP. Total sales includes both domestic sales and exports, so to get domestic sales the price statisticians at Statistics Sweden must subtract exports as given in UHV. However, reported sales and exports for a given firm in IVP and UHV can be different from their true values for various reasons. For example, corporate groups can coordinate such that production takes place in one firm, which then sells the goods at cost of production to another firm in the same group which finally sell the goods as exports. The reported sales and exports in IVP and UHV do not adjust for such internal transactions within the group and must be adjusted by the price statisticians. These adjustments are made annually before drawing the samples. Having done so, adjusted sales are used to compute weights in the index construction. If a firm's sales of a given type of product is sufficiently large compared to total sales in a product group, that firm-product is added to the sample. Firm-product pairs with smaller sales shares are drawn with probability proportional to size (PPS) sampling.

Firms that are included in the sample are required to respond by law. They are contacted by mail, in which they are asked to choose a representative product within a given product code. This firm-specific representative product, which we refer to as *item*, is the unit of observation in the goods and services PPI micro data. Respondents are then in each period asked to fill in a form electronically, by e-mail, or mail. Respondents are asked whether the item is still considered as representative for the firm's sales within the given product code. If it is, the firm reports the average transaction price of the item, and it indicates if product characteristics have changed, e.g., by terms of delivery, or the customer to which the item is sold. If it is not, the firm is asked to report another product of the same type, as given by the product code. The firm is then also asked to report the last period's price of the new item and how it differs from the previously reported item. If the item was not transacted in the given period, the firm can report its listed price (Statistics Sweden, 2019*b*,*a*).

3.2 Goods PPI micro data

The unit of observation in the goods PPI micro data is an item observed at a monthly frequency. In a given month there is information on the item's reported price, the quality/quantity-adjusted price, the type of product as given by the CN8 code and a broader product group as given by a 5-digit code (Standard för svensk produktindelning efter näringsgren, SPIN5), the currency and month-specific exchange rate, and the market on which the item was sold, i.e., domestically or as exports or imports. SPIN corresponds to the Classification of Products by Activities (CPA), which in turn can be translated to the Combined Nomenclature. However, the two systems are not directly related in the sense that the first five digits of a product's CN8 code are not necessarily the same as its SPIN5 code. We refer to the CN8 and SPIN5 as the *product code* and the *product group code*, respectively.

The year-specific item and product group weights used to aggregate the micro data up to various levels of aggregation, including the final indices published by Statistics Sweden, are also reported. The weights are related to sales, but not necessarily equal to sales for all items. In particular, items with a sufficiently large sales share in relation to total sales in the same product group are drawn to the PPI sample with certainty and get a weight equal to sales. Items with smaller sales shares are drawn with probability according to PPS sampling. Items that are drawn with probability get a weight proportional to the remainder of total product group sales after subtracting the sum of sales of items drawn with certainty. The remainder is divided proportionally across items drawn with probability. Thus, weights of items with relatively small sales shares are greater than, and not equal to, their corresponding item sales. For a detailed description of the relationship between weights in the PPI micro data and actual sales, see Appendix B.

Some observations (less than 0.1% of the unweighted sample) do not represent items but are instead imputed indices representing a group of products. We remove those imputations when reporting summary statistics and estimating elasticities of substitution, but keep them in the sample when aggregating the micro data to a final aggregate index. For years 1992–2019, the data then consists of 1,543,145 item-month observations and represents the Swedish private goods market, including goods produced and sold domestically within Sweden as well as exports and imports. The weighted share of the domestic market is 35%, while exports and imports make up 33% and 32%, respectively. The sample covers in year 2019 almost 6,500 items across 2,401 firms and 2,248 products, as defined by the 8-digit product code. Although most firms report only one or two items per 8-digit product code, firms with large market shares are asked to report prices for more items. The official aggregate goods PPI, which excludes imports, can be replicated with our data for 2004–2019. We therefore restrict the focus of the main analysis to exports and domestic sales during this period, leading to a sample size of 611,049 item-month observations. The data also includes a flag for transactions, indicating if an item was not sold, or if the firm did not respond to the survey in a given month, which can for instance happen during summer vacations. We remove these observations when reporting summary statistics. The weighted share of such missing reporting accounts for 5.5% of the monthly observations in 2004–2019, and removing them reduces the sample to 562,632 observations.

In Appendix A, we provide detailed summary statistics in terms of how new items enter and exit the sample, the share of product substitution, as well as price change frequencies and sizes of price change. About 1% of the sample are new items. Out of existing items, 2% undergo a product substitution, and 42% experience a price change. The remaining part of the sample, 55% consists of existing items that do not experience substitution or a price change. For the subset of items that do not undergo a product substitution, the weighted median and mean of monthly price change frequencies are 22% and 44%, respectively. As in Nakamura and Steinsson (2008), a very small fraction of products have price change frequencies close to the median. For example, the 45th percentile of price change frequencies is 11%, while the 55th percentile amounts to 42%. This implies that the weighted medians for various subsamples are markedly different. For example, the weighted median frequency in the energy industry is 91%, while only 17% in manufacturing. The weighted mean is relatively more robust, but still varies from around 43% in manufacturing to 78% in energy. Conditional on a price change, the weighted median and mean of the size of price change are 3% and 5%, respectively, and also vary substantially across industries.

In comparison to other countries, prices change more frequently in Sweden. For example, when restricting the sample to become similar to the data for the US used in Goldberg and Hellerstein (2011), we still find similar median and mean frequencies of price changes in Sweden, given by 25% and 44%. This implies that the price change frequency is higher than in the US, where the weighted median and mean equals 17% and 37%.

3.3 Services PPI micro data

Unlike for goods, the PPI micro data for services is not collected monthly but quarterly. Existing micro data for services also covers a shorter time frame, from year 2013 to 2019, resulting in 101,269 item-quarter observations.[12] The data contains services traded in the local market between 2013–2019 as well as exports and imports between 2018–2019. As for the goods data, we drop imports in the main analysis of services. In 2019, sales in the local market cover 78% of total sales, and exports and imports cover 14% and 8%, respectively. In the same year, the sample covers 4,978 items across 2,665 firms. Unlike for goods, there is no 8-digit Combined Nomenclature code in the services PPI because of a different product classification scheme. In particular, Statistics Sweden collects information on a 7-digit product group level, SPIN7. Thus, we observe the 7-digit product group, instead of the 8-digit Combined Nomenclature code as in the goods PPI. However, the unit of observation is still on the finer item level, and a firm can report several items to the services PPI in a given quarter.

The services PPI micro data has been continuously developed during the past years, arguably to a greater extent than the goods PPI, which has been collected over a longer period. The developments in collecting services PPI result in larger shares of new product groups and items that are added over time, amounting to 12% of new items in a given quarter. 4% of items undergo a product substitution, and 35% experience a price change. The remaining part of the sample, 49%, consists of existing items that do not experience substitution or price change. For the subset of items that do not undergo a product substitution, the weighted median and mean of quarterly price change frequencies are 25% and 42%, respectively. Aggregating the time variable in the goods PPI micro data to a quarterly frequency, the corresponding rates are 33% and 53%. Conditional on a price change, the weighted median and mean of the size of price change are 3% and 7%, respectively, while 4% and 8% for goods. Thus, goods prices tend to change more often than services. This is consistent with a comparison on a monthly frequency in the US by Goldberg and Hellerstein (2011). However, the larger price change sizes in goods than in services in the Swedish data is in contrast to their finding for the US, which suggest the opposite relationship. Just as for goods,

[12] As for the goods data, we remove observations that are imputed indices representing a group of products, rather than an item. Such imputed indices represent 0.3% of the sample.

price change frequencies and sizes vary substantially across industries for services. More summary statistics on price change frequencies and sizes are provided in Appendix A.

3.4 Summary statistics of price growth factors and weights

This section provides summary statistics on the underlying micro data on price growth factors $\frac{p_{t,i}}{p_{0,i}}$ and weights $w_{0,i}$ used to construct aggregate price indices. These summary statistics also speak to the moments used to compute log-normal approximations of aggregate indices, since they are determined by the averages, variances, and covariances of the logarithms of price growth factors and weights across items within the same product group. The micro data that we access include more product groups than are used when constructing the official goods and services PPI. For example, farming, forestry and fishing are not included in the official goods PPI. In addition to the sample restrictions mentioned above, we here also remove those product groups, leading to a sample size of 546,056 for goods and 80,226 for services. We focus on the average and standard deviation of price growth factors and weights, as well as the correlation between the two, across items within a given product group and year. We only include group-years with at least four items.

Table 2 shows the distributions of averages and standard deviations of annual price growth rates and weights of items within product groups, pooled across years. It reveals substantial heterogeneity in annual price growth and weights, both within and across product groups. For example, the interquartile range of average price growth rates across groups spans almost 6 percentage points in goods, and the standard deviation of price growth within product groups is on average 8.4 percentage points for services and even higher for goods. Weights are smaller in services than in goods. Like price growth rates, weights vary across and within groups.

Goods (2004–2019)	25^{th}	50^{th}	75^{th}	Mean
Av. price growth	-0.69	1.93	5.03	2.69
Sd of price growth	4.21	7.04	11.27	8.87
Av. weight	7.14	12.50	20.00	13.41
Sd of weight	1.92	4.35	8.39	6.18

Services (2013–2019)	25^{th}	50^{th}	75^{th}	Mean
Av. price growth	0.45	1.82	3.25	2.03
Sd of price growth	4.28	6.89	10.70	8.37
Av. weight	2.04	3.70	7.14	5.46
Sd of weight	1.16	2.28	4.16	3.43

Table 2: Percentiles and means of average and standard deviation of annual price growth rates and weights, for goods (left panel) and services (right panel), respectively, pooled over years and product groups. We only include group-year observations containing at least 4 items. All numbers denote percentages.

Table 3 shows the distribution of correlations between price growth rates and weights. On average, the correlation is slightly negative but close to zero in both goods and services. However, the average conceils great heterogeneity since the interquartile range of correlation coefficients goes from -0.35 to 0.24 for goods, and from -0.16 to 0.15 for services.

Goods (2004–2019)	25th	50th	75th	Mean
Correlation	−0.35	−0.06	0.24	−0.05

Services (2013–2019)	25th	50th	75th	Mean
Correlation	−0.16	−0.01	0.15	−0.00

Table 3: Percentiles and means of correlation between annual price growth factors and weights, for goods (left panel) and services (right panel), respectively, pooled over years and product groups. We only include group-year observations containing at least 4 items, and group-year observations with non-zero standard deviations of annual price growth and weights.

4 Index counterfactuals

In this section we compute counterfactuals for the two aggregate indices of goods and services using alternative aggregation methods, estimated elasticities of substitution, alternative weighting schemes and distributional assumptions. First, we contrast the previously outlined weighted and unweighted arithmetic mean, geometric mean and ratio of arithmetic means. Second, we return to the CES production framework and compute the indices for different parameterizations of the elasticity of substitution at the lower and upper level. Third, we let the data speak for itself and compute the indices using estimated elasticities of substitution. As a fourth exercise we compute the indices using alternative weighting schemes, either holding product group weights fixed or updating them in five year intervals. Lastly, we compute an approximation to the true price index using distributional assumptions about price growth rates and item weights.

We use Statistics Sweden's convention to refer to the aggregate index for goods as *the PPI* and the aggregate index for services as *the TPI*.[13] Each counterfactual is computed for both the PPI and the TPI. The TPI differs from the PPI in terms of how item price changes are aggregated. Before presenting results on the index counterfactuals starting in section 4.2 we outline how the PPI and TPI computation differs in practice.

4.1 Differences between PPI and TPI computation

In the PPI item price changes are first aggregated arithmetically to the 5-digit item classification level (SPIN5). In a second step, price growth rates at the SPIN5 level are aggregated up arithmetically to the highest level. In both steps item sales or SPIN5 sales are used as weights for the aggregation. The aggregation method for the TPI differs from that of the PPI and even changes over the years. After 2014, service prices are first aggregated up geometrically to the 7-digit item classification level (SPIN7). Sales are used as weights in this step. Second, growth rates at the SPIN7 level are aggregated up to the highest level arithmetically using SPIN7 sales as weights. For the years before 2015 the TPI was computed using an alternative aggregation method: first, within each SPIN5 firm price growth rates are computed as unweighted geometric averages of item price changes of that firm. Second, firm price growth rates are averaged geometrically within a

[13] TPI is an abbreviation of *Tjänsteprisindex*, meaning "services price index".

SPIN5 using firm sales as weights. Lastly, SPIN5 growth rates are aggregated up to the highest level arithmetically using SPIN5 sales as weights. Using the outlined procedure we are able to replicate the PPI and TPI published by Statistics Sweden with our data, see Appendix (Figure 14 and Figure 15).

The change in the level of aggregation in 2014 introduces an inconsistency to the index computation. For the counterfactuals we are interested in changing the method of aggregation. Since the level at which item price growth rates are first aggregated changes in 2014, counterfactual methods of aggregation would be applied at different levels before and after the change. We therefore construct an alternative TPI that first uses a weighted geometric aggregation of item price growth rates up to the SPIN5 level for all years, and then a weighted arithmetic aggregation of SPIN5 growth rates up to the aggregate level in the second step. This structure of aggregation is in line with how aggregation is done for the PPI. Note however that in contrast to the PPI, in the TPI item price growth rates are geometrically aggregated in the first step. As shown in Figure 16 in the Appendix, the TPI with geometric aggregation up to SPIN5 for all years is indistinguishable from the original TPI. A second concern for the counterfactuals is that Statistics Sweden includes sub-indices of other price indices in the TPI. For example several product group indices from the CPI are included in the TPI. Those indices enter our data as pre-compiled indices without any underlying micro data. The imported product group index appears in our data as a product group consisting of a single "item" observation with price growth equal to the index price growth. For product groups with only one observation product group price growth automatically equals price growth of the single observation. Changing the lower-level aggregation method does not change price growth at the product group level for those groups. This biases differences in inflation rates at the aggregate level for different lower-level aggregation methods to zero. For the counterfactual exercises we hence use a TPI that is constructed from micro data only and excludes all imported indices. Figure 17 in the Appendix shows the comparison between this alternative TPI and the original. The match is fairly close. We therefore use the TPI consisting of micro data only with item price growth aggregated at the SPIN5 level as the baseline reference for the counterfactual exercises.

4.2 Axiomatic approach

Figure (1) contains PPI counterfactuals using the arithmetic mean, geometric mean and the ratio of arithmetic means for the price aggregation. All three methods are compared using item specific weights or applying equal weights to all items in the same elementary group giving a total of six different price indices which all are used in practice. Statistics Sweden computes the PPI using weighted arithmetic price aggregation which we refer to as "Baseline". The remaining five indices form the counterfactuals. Given the desirable axiomatic properties of geometric price aggregation one interesting comparison is between arithmetic and geometric aggregation. Following Jensen's inequality, arithmetic price aggregation leads to a systematic weakly higher inflation rate than geometric aggregation. Inflation rates at the elementary level only coincide for both aggregation methods if all item price growth factors within the product group are identical. Any positive variance of item price growth factors across items will lead to a lower product group average inflation rate using the geometric mean. This is displayed in Figure 1 and quantified in Table 4. Under weighted arithmetic averaging the annual inflation rate during 2004–2019 is 2.01%. Changing the lower level aggregation method from weighted arithmetic to weighted geometric mean lowers the annual inflation rate by 0.48 percentage points (pp.). Hence, under weighted geometric averaging the inflation rate stands at 1.53%. This comparison shows that if Sweden's and Denmark's statistical offices compute the PPI on the very same data, Denmark's inflation rate would be 0.48 pp. lower than Sweden's by construction. Since inflation rates carry through one-to-one to real growth, if nominal output growth of the goods sector in Sweden and Denmark were identical, it necessarily follows that Denmark's real output growth were 0.48 pp. higher. Surprisingly, weighted arithmetic aggregation and the ratio of weighted arithmetic means almost give the same inflation rate. The two indices almost perfectly overlap in Figure 1 and the difference is quantified at -0.04 pp.

What happens instead if we drop the item specific weights in the index computation? Comparing the baseline with unweighted arithmetic aggregation shows that applying equal weights increases the inflation rate by 0.49 pp. relative to Baseline. Hence, if Sweden adapts the PPI computation method of Spain, Swedish PPI inflation would be 0.49 pp. higher and real output growth lower by the same amount. Dropping the weights also leads to higher inflation using the geometric mean and the ratio of arithmetic means as aggregation methods. This increase in inflation is, however, not by construction and is solely driven by items with smaller weights displaying a systematically higher inflation rate in the data. This comparison shows that differences in inflation rates arising from the lower level aggregation method are sizable. Comparing PPI inflation rates or real output growth rates internationally without taking the differences arising from the aggregation method into account is questionable.

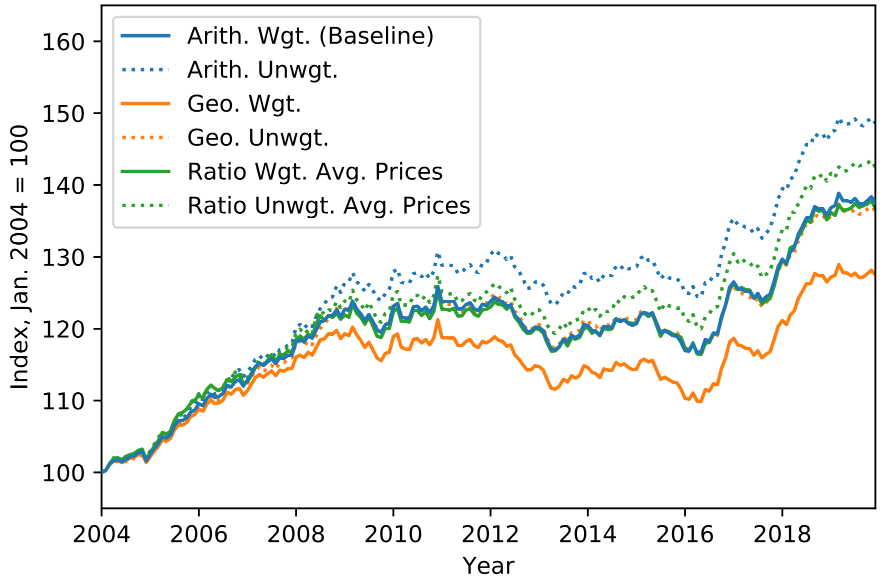

Figure 1: PPI for different aggregation methods at the elementary level

Counterfactual inflation: elementary level aggregation method

	PPI (2004–2019)	TPI (2013–2019)
Arith. Wgt.	2.01%	1.62%
Difference to Arith. Wgt. in pp.		
Arith. Unwgt.	0.49	0.16
Geo. Wgt.	-0.48	-0.40
Geo. Unwgt.	-0.05	-0.31
Ratio Wgt. Avg. Prices	-0.04	-0.18
Ratio Unwgt. Avg. Prices	0.22	-0.27

Table 4: Annualized inflation rate for alternative aggregation methods at the elementary level. Statistics Sweden uses Arith. Wgt. as the aggregation method for the PPI and Geo. Wgt. for the TPI.

How does adjusting the upper level aggregation method alter the PPI? Figure 2 contains the PPI with both arithmetic averaging at the lower and higher level, and the counterfactual with arithmetic averaging at the lower level combined with geometric averaging at the higher level. We also show the index using geometric averaging all the way both at the lower and higher level. By Jensen's inequality going from arithmetic averaging at both lower and higher level to geometric averaging at any level necessarily leads to a weakly lower inflation rate. Geometric averaging at both lower and higher level leads to the lowest inflation rate in Figure 2.

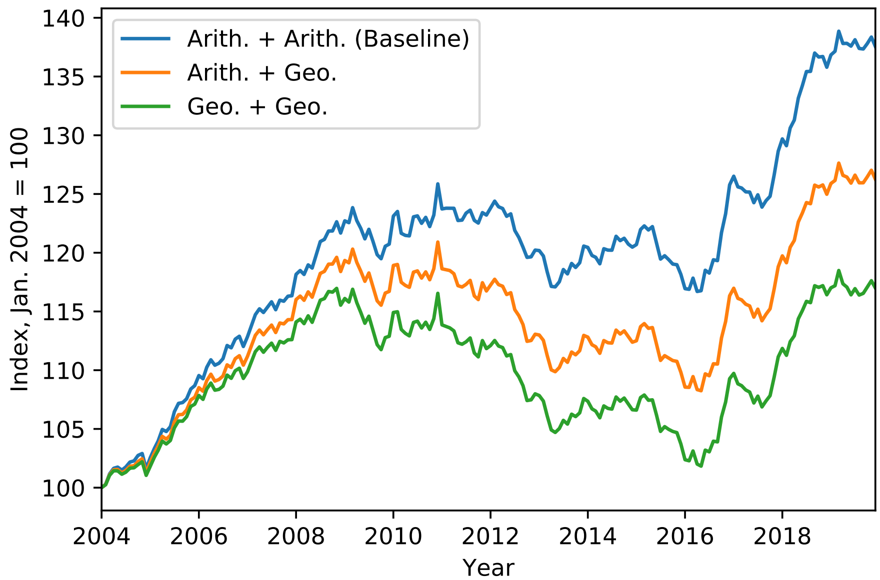

Figure 2: PPI for different aggregation methods at the higher level. "Arith. + Geo." refers to arithmetic averaging at the elementary level and geometric averaging at the higher level.

Counterfactual inflation: higher level aggregation method

	PPI (2004–2019)
Arith. + Arith. (Baseline)	2.01%
Difference to Baseline in pp.	
Arith. + Geo.	-0.54
Geo. + Geo.	-1.03

Table 5: Annualized inflation rate for the baseline method and difference in inflation rate when using geometric averaging at the higher level. "Arith. + Geo." refers to arithmetic averaging at the elementary level and geometric averaging at the higher level.

In sum, alternative aggregation methods at the lower level lead to sizable differences in PPI inflation with arithmetic aggregation resulting in a higher inflation rate than with geometric aggregation. The method of aggregation at the higher level also has a sizable effect on the inflation rate. The counterfactual without item specific weights shows that weighting tends to reduce inflation: larger weight items have a systematically lower inflation rate.

Similar to the PPI we repeat the same counterfactuals for the TPI. The counterfactuals are shown in Figure 3. By construction arithmetic aggregation results in at least as much inflation as geometric aggregation: the blue line lies above the orange one for all years. As for the PPI we observe that ignoring item weights leads to higher inflation when comparing 2019 with 2013 prices. Items with higher weights display lower inflation rates such that unweighted aggregation results in higher inflation. That this is nothing systematical, but rather just a feature of the data can be seen from the fact that for selected sub-periods unweighted aggregation leads to less inflation. Table 4 quantifies the differences. Weighted arithmetic aggregation gives an annual inflation rate of 1.62%. Ignoring item specific weights leads to 0.16 pp. higher inflation per year. Switching from arithmetic to geometric aggregation which is how the TPI is computed in practice reduces inflation by 0.4 pp. annually. This is similar in size to the 0.48 pp. reduction found for the PPI. Using the ratio of weighted average prices instead of arithmetic aggregation decreases inflation by 0.18 pp. per year.

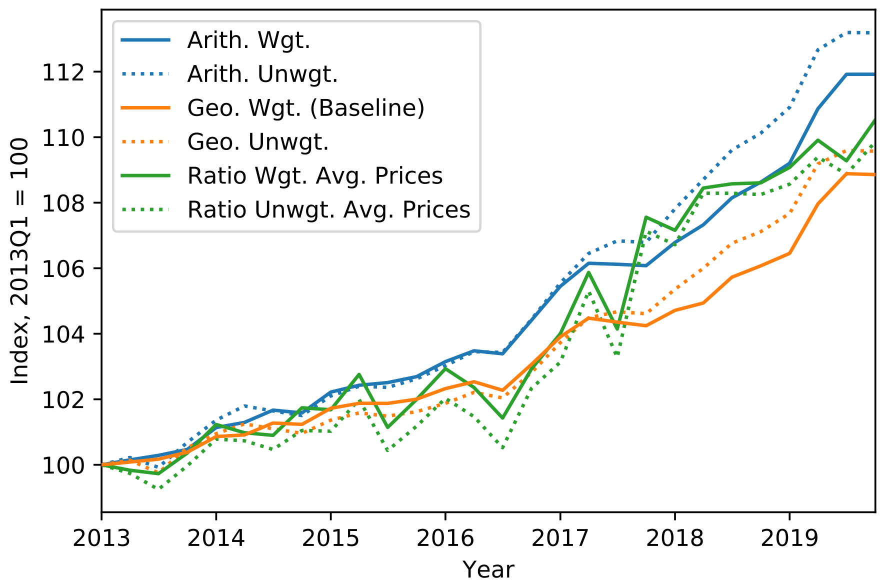

Figure 3: TPI for different aggregation methods at the elementary level

In a second counterfactual exercise we change the method of aggregation which is used to aggregate product group price changes to the aggregate level. The counterfactual for the higher level-aggregation is shown in Figure 4. As expected geometric aggregation at the higher level leads to less inflation than arithmetic aggregation. Somewhat surprising is how little the aggregate index changes compared to the change observed for the PPI. For the PPI going from arithmetic to geometric aggregation at the higher level decreases annual inflation by 0.54 pp. Table 6 shows that for the TPI inflation only decreases by 0.04 pp. when switching the higher level aggregation method. Note though that for the PPI item price changes within product groups are aggregated arithmetically whereas this is done geometrically for the TPI. If we were to aggregate item price changes in the PPI as is done in the TPI we observe a smaller change in annual inflation when varying the higher level aggregation method (not shown).

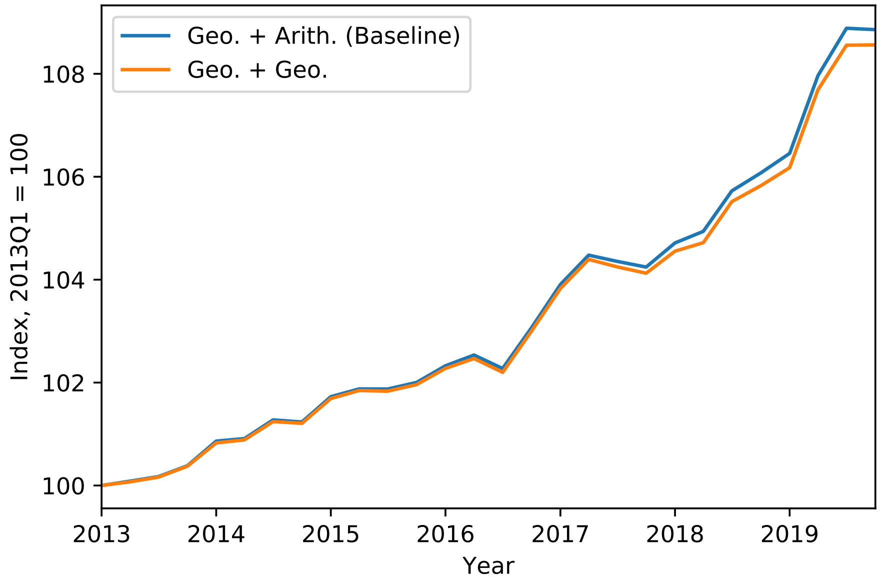

Figure 4: TPI for different aggregation methods at the higher level. "Geo. + Arith." refers to geometric averaging at the elementary level and arithmetic averaging at the higher level.

Counterfactual inflation: higher level aggregation method

	TPI (2013–2019)
Geo. + Arith. (Baseline)	1.22%
Difference to Baseline in pp.	
Geo. + Geo.	-0.04

Table 6: Annualized TPI inflation rate for the baseline method and difference in inflation rate when using geometric averaging at the higher level. "Geo. + Arith." refers to geometric averaging at the elementary level and arithmetic averaging at the higher level.

4.3 Economic theory-based approach

For the counterfactuals in this section we return to the CES-framework outlined in section 2.2. Equations (5) and (9) are used to compute the theory-based index. If not mentioned otherwise, we keep the elasticity of substitution between product groups, σ, fixed at zero and vary σ_g the elasticity of substitution between items within a product group. $\sigma = 0$ coincides with arithmetic averaging at the higher level. The exact price index for different values of σ_g is shown in Figure 5. The corresponding differences in inflation rates for alternative values of σ_g are quantified in Table 7. $\sigma_g = 1$ results in 0.48 pp. lower annual inflation relative to $\sigma_g = 0$. This difference is necessarily the same as the difference in inflation between arithmetic and geometric aggregation from the previous counterfactual since eq. (5) boils down to arithmetic and geometric aggregation of item price changes for the special cases of $\sigma_g = 0$ and $\sigma_g = 1$. Higher values of σ_g lead to lower inflation as shown in Figure 5 which is due to Jensen's inequality. The proof is contained in the Appendix D.1. Intuitively, the higher the value of σ the higher the substitutability of items in production and items with relatively higher inflation rates are replaced by items with lower inflation rates. Already an elasticity of substitution equal to four cancels out all inflation. Over the period 2004-2019 the annual inflation rate with $\sigma = 4$ is almost 2 pp. lower than under Baseline. With $\sigma = 8$ annual inflation is 3.77 pp. lower than with Baseline. To conclude, we reiterate that the PPI computation is sensitive with respect to the lower level aggregation method.

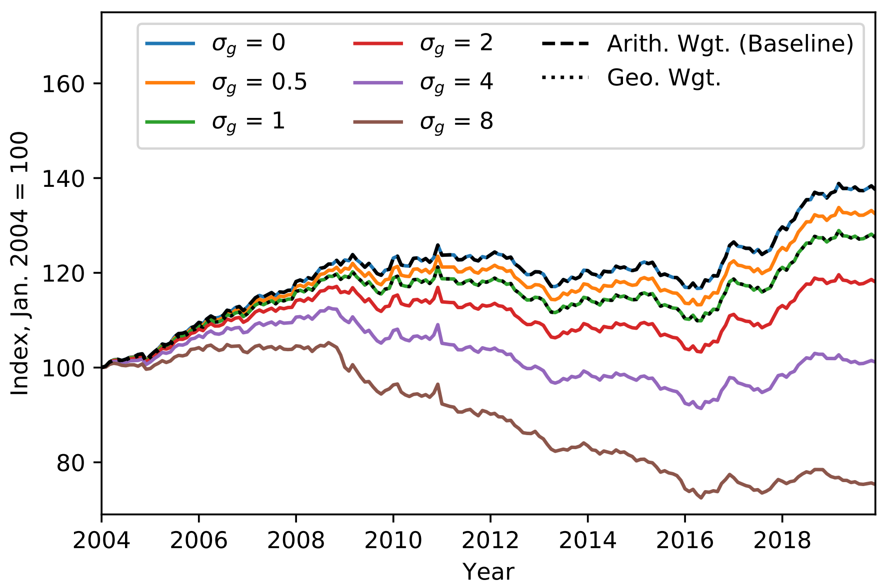

Figure 5: PPI for different elasticities of substitution at the elementary level

Counterfactual inflation: elementary level elasticity of substitution

	PPI (2004–2019)	TPI (2013-2019)
$\sigma_g = 0$	2.01%	1.62%
Difference to $\sigma = 0$ in pp.		
$\sigma_g = 0.5$	-0.24	-0.21
$\sigma_g = 1$	-0.48	-0.40
$\sigma_g = 2$	-0.97	-0.81
$\sigma_g = 4$	-1.94	-1.87
$\sigma_g = 8$	-3.77	-4.67

Table 7: Annualized inflation rate for the baseline method and difference in inflation rates when using alternative elasticities of substitutions at the elementary level.

Similarly we compute the TPI for different values of σ_g. Figure 6 displays the indices and Table 7 quantifies the differences. Moving from $\sigma_g = 0$ to $\sigma_g = 1$ reduces inflation by 0.4 pp. per year which is (by construction) equal to the difference between weighted arithmetic and weighted geometric aggregation shown in Table 4. Already $\sigma_g = 4$ leads to a negative inflation rate and $\sigma_g = 8$ results in 4.67 pp. lower inflation per year than with arithmetic aggregation. The differences in inflation rates across σ_g line up well with the differences observed for the PPI.

One interesting observation stands out from Figure 6. We argued before that higher values of σ_g lead to systematically weakly lower inflation rates. This holds true for inflation rates with respect to the baseline period the index is chained at which for the PPI is December and for the TPI Q4 of a given year. Hence, for the TPI the Q4_{t-1}-Q1_t, Q4_{t-1}-Q2_t, Q4_{t-1}-Q3_t and the Q4_{t-1}-Q4_t inflation rates are decreasing in σ_g. Inflation rates with respect to a quarter that is different from Q4 can indeed be increasing in σ_g as is shown in Figure 6 for the period Q2_{2014}-Q3_{2014}. The counterfactual with $\sigma_g = 8$ displays the highest inflation rate for this quarter. This observation can arise as a result of mean reversion of price changes within product groups. See Appendix D.1 for an example. If Statistics Sweden were to chain the TPI at quarterly frequency (or the PPI and monthly frequency) instead of annual frequency, all inflation rates for any given time interval were weakly decreasing in σ_g.

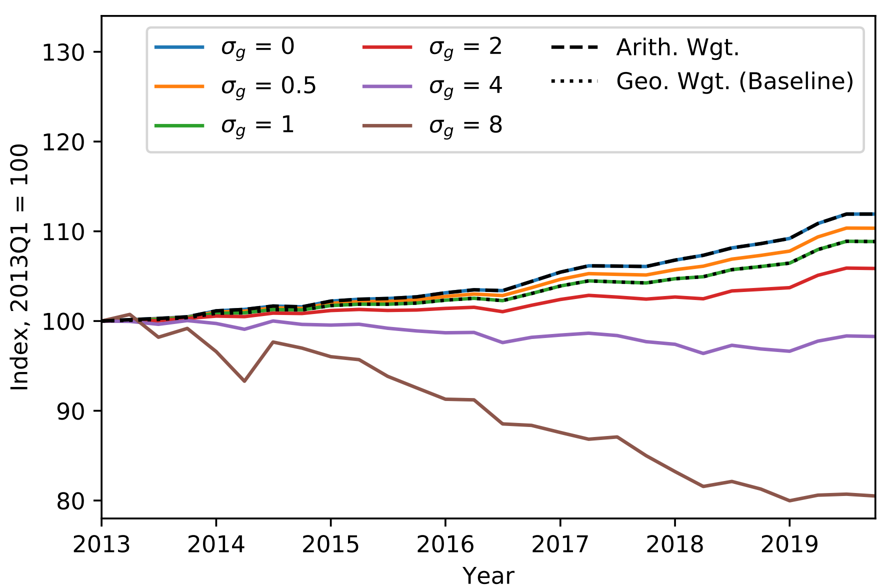

Figure 6: TPI for different elasticities of substitution at the elementary level

4.4 Estimated elasticities of substitution

The counterfactuals in the previous section show how the PPI varies for different values of the elasticity of substitution at the lower level. In this section we estimate the elasticity of substitution from micro data and contrast the official PPI to the one obtained using the estimated elasticities of substitution. When estimating the elasticities of substitution, we focus on domestic sales and impute elasticities for exports. A detailed outline of the data and the estimation strategy can be found in Appendix C.

Data, strategy, and results from estimation

The data used for estimating the elasticities of subsitution across goods is the Swedish industrial goods production survey, IVP—an encompassing database containing information on yearly production values and quantities of goods reported by a sample of Swedish production-units, active last year in the tax records—and trade data from Swedish customs, UHV. We use this data because it constitutes the underlying source for determining weights in the goods PPI micro data. As described in Section 3, sales and exports reported to IVP and UHV does not necessarily represent their true values, e.g. due to transactions between firms within the same corporate group. There are also complications with merging IVP and UHV data because UHV reports exports in a calendar year while IVP reports production values of a firm's fiscal year. Since we do not have additional information from the firms in question, we can not fully account for these errors in measurement. However, we constrain the sample to firms with a non-negative difference between the production value in IVP and exports in UHV, and assume that goods for which there is no match in the trade data are only sold in the domestic market.

In contrast to goods, we do not access the underlying data sources used by Statistics Sweden to compute the weights observed in the services PPI micro data. We therefore use the weights in the services PPI micro data for estimation. These weights do not necessarily represent true sales for a given service, implying potential mismeasurement and bias when estimating elasticities of substitution. First, while products with sufficiently high sales in relation to total sales in the product group are assigned their true sales as weights, products with low sales are given a weight that equals the sample average of sales of all products with low sales. Thus, the weight of small producers are not representative for the true sales shares of those producers. Since we can not identify which items in the PPI micro data that do not have true sales as weights, these are affecting our estimates. Second, all weights in the PPI micro data in a given year a are collected from sales data in year $a-2$. These sales are inflated by multiplying by a price index that captures price growth during year $a-1$. Depending on data availability, the index can be given by product-specific price growth, or by an index capturing price growth for a collection of products, e.g. within the same product group.

The estimation strategy builds on Hottman et al. (2016) who introduce a nested CES-structure aggregating items within a firm to firm output and aggregating up firm output to the aggregate level. Key parameters for aggregation are the elasticity of substitution between items within a firm and the elasticity of substitution between firms. Identification for estimating the within firm elasticity comes from the assumption that supply and demand shocks to double-differenced item price changes within the same firm are orthogonal. Double-differencing here refers to subtracting price changes across time for a given item to the price change of a reference item within the same firm, with the aim of removing time-specific and firm-specific shocks.

We adjust the estimation framework above to our setting by changing the level of CES aggregation such that instead of aggregating items within firms we aggregate products within product groups. Hence, the identifying assumption for estimating the within product group elasticity of substitution between products is that supply and demand shocks to double-differenced item price changes within the same elementary-level product group code, e.g. SPIN5 for goods, are orthogonal. Using this approach we obtain an estimate of the elasticity of substitution between products for each product group. In some cases, for example if too few observations in a SPIN5 product group are available, we group the double-differenced item price changes for different SPIN5s within the same SPIN4 to obtain a σ_g estimate for that SPIN5. If estimation again turns out unsuccessful at this level we group SPIN5s within the same SPIN3 and so on. For goods, the product group code on the elementary level is given by SPIN5, while for services it is given by SPIN7.

With estimated elasticities at the product group level we then apply the Hottman et al. (2016) approach for the higher level of aggregation to arrive at an estimated elasticity across product groups, σ, for goods and services, respectively.

Table 8 shows the number of estimates pooled across different levels of aggregation of product codes. For goods, 122 of the 318 estimated σ_g are obtained without pooling across product groups. For services, the corresponding number is 44 out of 298 estimated elasticities.

Pooling level of estimated elasticities across products

Level	N. estimated at level for goods	N. estimated at level for services
SPIN0	2	44
SPIN1	41	116
SPIN2	106	41
SPIN3	39	20
SPIN4	8	17
SPIN5	122	2
SPIN6	–	34
SPIN7	–	24
Tot. estimates:	318	298

Table 8: Number of elasticities across products within 7-digit product groups that are estimated using products within a 7-digit group ($SPIN7$), or using products within a 6-digit group ($SPIN6$), within a 5-digit group ($SPIN5$), and so on. *Tot. estimates* equals the number of unique 5-digit product groups in the PPI micro data for goods, and the number of unique 7-digit product groups in the PPI micro data for services.

Table 9 shows the distribution of estimated σ_g. Recall that an elasticity of substitution of zero coincides with arithmetic averaging used in the official index for the goods sector, and an elasticity of substitution of one coincides with geometric averaging used in the services sector. Our estimates show substantially higher elasticities in both sectors. While the median elasticity is greater for services than for goods, 5.5 compared to 4.4, the range of values is wider in the services sector. For example, the top and bottom 5th percentiles are 27.0 and 1.4 for services, and 18.8 and 2.8 for goods. This pattern remains when we only include elasticities estimated at the elementary levels of product groups.

Distribution of estimated elasticities across products

	Goods		*Services*	
Ranked Percentile	Estimate	Estimate (SPIN5 only)	Estimate	Estimate (SPIN7 only)
99	1.38	1.23	1.29	0.25
95	1.77	1.56	1.37	0.49
90	2.25	1.75	1.78	1.20
75	3.12	2.72	2.22	2.61
50	4.37	3.85	5.51	4.85
25	6.97	6.42	6.60	9.60
10	12.02	12.60	13.60	26.78
5	18.78	16.95	27.04	32.50
1	43.11	28.90	45.03	43.23
N. estimates:	318	122	298	24

Table 9: Percentiles of elasticities of substitution across products within product groups. *Estimate* shows the percentiles using all elasticities, including those estimated using higher level groupings than SPIN5 for goods and SPIN7 for services. *Estimate (SPIN5 only)* and *Estimate (SPIN7 only)* only includes elasticities estimated for groupings at the elementary level.

Finally, using the estimated elasticities of substitution on the lower level, we estimate also the elasticity on the higher level. We estimate σ to be 1.237 for goods and 1.071 for services. The estimate for goods is statistically different from 1, while the estimate for services is not. Just as for the elasticities on the lower level of aggregation, the elasticity on the higher level of aggregation is higher than the one used for the official indices, which equals 0 for both goods and services.

Counterfactuals with estimated elasticities

After matching the estimated elasticities of substitutions with the corresponding product groups in the PPI micro data we compute the PPI as outlined in Section 2.2. Figure 7 contains the results. Baseline refers to the Statistics Sweden index with arithmetic averaging at the lower and higher level ($\sigma_g = \sigma = 0$). The counterfactuals where σ_g is estimated are added for two different higher aggregation methods. The first shows the PPI with estimated elasticities of substution within product groups and arithmetic averaging at the higher level ($\sigma = 0$) to illustrate how the baseline PPI changes if instead of assuming $\sigma_g = 0$, estimates for the lower level elasticity of substitution are obtained from the data while keeping higher aggregation unchanged. The second counterfactual shows the PPI with estimated elasticities of substitution both at the lower and higher level. This counterfactual shows how much inflation changes if on top of introducing estimated σ_g we let the data determine the substitutability between product groups. The counterfactual for $\sigma_g = 4, \sigma = 0$ is added for comparison.

The counterfactuals with estimated elasticities show less inflation than the baseline index. The differences in inflation rates are quantified in Table 10. Replacing σ_g by its estimated counterpart decreases inflation by 3.06 pp. per year relative to the baseline PPI. If on top of that the elasticity of substitution between product groups, σ, is estimated, inflation relative to baseline decreases by a total of 3.94 pp. per year. Since the baseline index has an elasticity of substitution equal to zero at both aggregation levels, allowing for more substitutability either between items within a product group or between product groups leads to lower price growth. Holding σ fixed at zero the counterfactuals show that the index with estimated σ_g displays even less inflation than the counterfactual index where σ_g is set to 4 for every product group.

PPI inflation with estimated elasticities is 3.96 pp. lower per year than inflation under the baseline index that uses elasticities of substitution equal to zero at the lower and higher level. To judge whether the difference between the baseline index and the index with estimated elasticities is statistically significant, we bootstrap the index with estimated elasticities as follows. We re-compute the index 1000 times, each time drawing product group specific elasticities and an elasticity for the higher level aggregation from the distribution of estimated σ_g and σ. Confidence bands for the distribution of indices are shown in Appendix C.4. The range covered by the confidence bands compares small to the differences in price levels that arise from different aggregation methods.

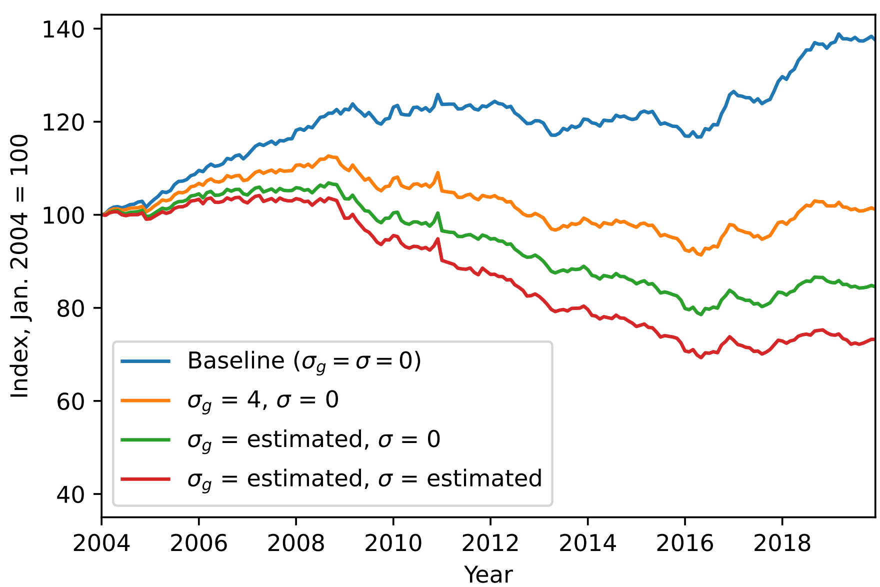

Figure 7: PPI with estimated elasticities of substitution at the higher and lower level. For an overview of the distribution of estimated σ_g see Table 9. σ is estimated at 1.237

Counterfactual inflation: estimated elasticities

	PPI (2004–2019)	TPI (2013–2019)
$\sigma_g = 0$, $\sigma = 0$	2.01%	1.62%
Difference to $\sigma_g = 0$, $\sigma = 0$ in pp.		
$\sigma_g = 4$, $\sigma = 0$	-1.94	-1.87
$\sigma_g =$ estimated, $\sigma = 0$	-3.06	-2.52
$\sigma_g =$ estimated, $\sigma =$ estimated	-3.94	-3.11

Table 10: Annualized inflation rate for the baseline method and difference in inflation rate with estimated elasticities at the higher and lower level. For an overview of the distribution of estimated σ_g see Table 9. σ is estimated at 1.237 for the PPI and at 1.071 for the TPI.

Having estimated the elasticities of substitution at the lower and higher level of aggregation for services we compute the TPI with the obtained elasticities. We run two counterfactuals. The first uses the estimated elasticities at the lower level of aggregation while keeping the arithmetic aggregation at the higher level ($\sigma = 0$) as in the baseline TPI. In a second counterfactual we introduce the estimated elasticity of substitution at the higher level on top of the estimated elasticities of substitution at the lower level. The counterfactual with $\sigma_g = 4, \sigma = 0$ is added for comparison. The counterfactuals are shown in Figure 8 and differences in annual inflation rates are quantified in Table 10. Going from an elasticity of substitution of zero at both lower and higher level aggregation to estimated elasticities at the lower level decreases annual inflation by 2.52 pp. The reduction in inflation is greater than if σ_g had been set to four for all product groups instead of estimating the elasticity for each product group separately. So far, σ has been held constant at zero as in the baseline TPI. Using the estimated elasticities at both lower and higher level in the index computation results in 3.11 pp. less inflation per year than under the baseline TPI. Since σ is estimated at 1.071, aggregating with an estimated σ is roughly equivalent to geometric aggregation at the higher level. The change in annual inflation when changing the higher level aggregation from zero to (roughly) one in Table 10 is larger than in Table 6. Note, however, that in Table 6 the lower level elasticity is set to one whereas in Table 10 it is estimated. As for the PPI, a price index with reasonable elasticities of substitution that is derived from economic theory leads to substantially less inflation than price indices with arithmetic or geometric averaging that are generally used by statistical agencies. In Appendix C.4, we bootstrap the TPI with estimated elasticities of substitution similarly as we did for the PPI in the previous paragraph and report results on the distribution of indices. As before, differences in price levels that arise from different aggregation methods compare large to differences in price levels attributed to uncertainty in the estimated elasticities of substitution.

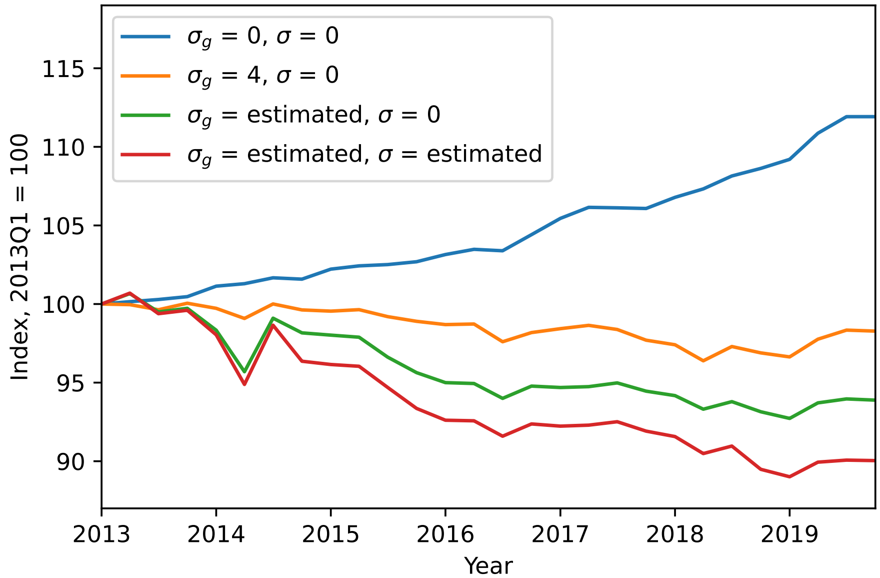

Figure 8: TPI with estimated elasticities of substitution at the higher and lower level. For an overview of the distribution of estimated σ_g see Table 9. σ is estimated at 1.071

4.5 Weighting schemes

The final counterfactuals speak towards the question: how do different weighting schemes affect the PPI? The counterfactuals on the aggregation methods used in practice already showed that an equal weighting where all item price growth factors within a product group receive the same weight leads to more inflation. In this section we compute the PPI for different weighting schemes at the product group level. One approach is to fix group weights at a constant level, e.g. at their 2004 levels. An alternative is to update the weights in fixed intervals. To mimic the US-PPI we choose to update weights every five years. When fixing the weights at their 2004 levels the problem arises that not all product groups are present in 2004. For the 121 (out of 502) product groups that enter the PPI after 2004 we take the earliest available weight, i.e. the weight from the year the product group enters the data. We include specifications where only product groups that exist in every year enter the counterfactual since in this specification all product groups receive the 2004 weight. Explicitly, we hold product group weights fixed either permanently or for five years, but let item weights within a product group vary over time as we observe them in the data. Running counterfactuals for the item specific weights is another interesting counterfactual, however given that items leave and enter the sample at a higher frequency than complete product groups, the problem that the unit of observation is not observed in the year when we hold weights constant is magnified when computing counterfactuals for item instead of product group weights. Therefore, the counterfactuals that we compute in this section concern only the product group weights.

The counterfactuals are shown in Figure 9. "Baseline" refers to the PPI computed using Statistics

Sweden's method with annually updated weights, "2004 weights" is the counterfactual with product group weights fixed at their 2004 levels and "5 year weights" refers to the counterfactual with updated weights in 2004, 2009, 2014 and 2019. The specification where only product groups that are present in the data in all years are included in the PPI computation is labelled "const. groups" (constant groups). Figure 9 shows that holding weights constant, either at the 2004 levels or for five year periods, leads to higher inflation relative to the baseline method. The extend to which the weighting schemes result in inflation that is different from the baseline method is quantified in Table 11. Fixing the weights at the 2004 levels results in 0.2 pp. higher inflation every year compared to the baseline specification. Five year updating of the weights yields 0.22 pp. higher relative inflation. It turns out that conditioning on the set of product groups that are present in all years doesn't affect the inflation numbers too much.

From a theoretical point of view, in a nested CES framework as in equation (1), if the elasticity of substitution between product groups, σ, is larger than one and some product groups display persistently higher inflation rates over time relative to other product groups, relative expenditure shares of the product groups with higher inflation will decrease. In this setting, fixing expenditure shares in the index computation at their initial levels or holding them constant for specified periods will hence introduce inflation. This is exactly what the counterfactuals in Figure 9 show. That expenditure shares of product groups with persistently higher price growth rates decline provides evidence against a Baumol effect within the Swedish goods producing industries. Baumol's cost disease suggests that prices in sectors with low productivity growth grow faster than in sectors with high productivity growth. If expenditure shares of sectors with high price and low productivity growth were to increase over time, this would result in a decrease in aggregate productivity growth. The counterfactuals show that fixing expenditure shares at their initial level introduces additional inflation to the aggregate index which implies that in the baseline index where weights are updated every year there is substitution away from high price (and low productivity) growth sectors. We therefore find no evidence that sectors with low productivity growth receive larger expenditure shares over time. Note that this is consistent with an elasticity of substitution between product groups that is larger than one and that we quantified at 1.237 in section 4.4 for the PPI.

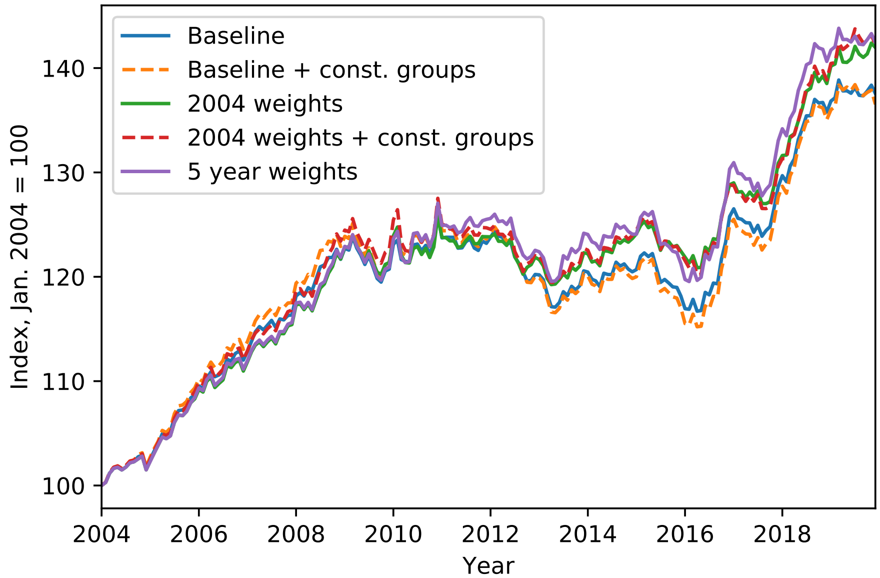

Figure 9: PPI for different weighting schemes at the product group level

Counterfactual inflation: weighting schemes

	PPI (2004–2019)	TPI (2013–2019)
Baseline	2.01%	1.22%
Difference to Baseline in pp.		
Baseline + const. groups	-0.05	-0.02
2004 weights	0.20	
2004 (PPI)/2013 (TPI) weights + const. groups	0.21	0.15
5 year weights	0.22	0.14

Table 11: Annualized inflation rate for the baseline method and difference in inflation rates when using alternative weighting schemes at the product group level. There is no TPI counterfactual with fixed 2013 weights due to the inclusion of further product groups in the index after 2013 and the trend break in weight levels in 2018 as explained in the main text.

We next repeat the weight counterfactuals for the TPI. As for the PPI, Statistics Sweden updates weights annually for the TPI computation. In the counterfactuals we run we either hold weights constant at the 2013 level or update them every five years, akin to the US-PPI, which for the sample period 2013–2019 implies a single updating of weights in the year 2018. As before we adjust weights at the product group level, while relative item weights within a product group remain the same. Only the relative weights of product groups are affected by our adjustments.

Not all product groups are included in the TPI for all years, i.e., some product groups enter after

2013. In the PPI counterfactual where weights were held constant at the 2004 level, for product groups that entered after 2004, we held the weights of their entering year constant. In the TPI there is the further complication that there is a trend break in the level of weights in 2018. This does not affect relative weights for a given year, however creates problems when holding weights fixed from different years for different product groups. When we hold weights constant at the 2013 level we therefore only include product groups that are present from 2013 on. For the counterfactual where we update after five years all product groups that are present in 2013 are included from 2013 onwards and the product groups that enter afterwards are included from 2018 on with the updating of the weights.

Figure 10 shows the counterfactuals. Similarly to the PPI counterfactuals, holding weights constant for all years or updating them every five years leads to higher inflation. Holding weights constant for all years adds an additional 0.15 pp. inflation per year whereas the five-year updating adds 0.14 pp., see Table 11. That weights fixed at their initial levels leads to more inflation implies that in the baseline index where expenditure shares are updated every year there is substitution away from product groups with permanently higher price growth rates. Hence we conclude that similarly as for the PPI we find no evidence for Baumol's cost disease within Swedish service industries.

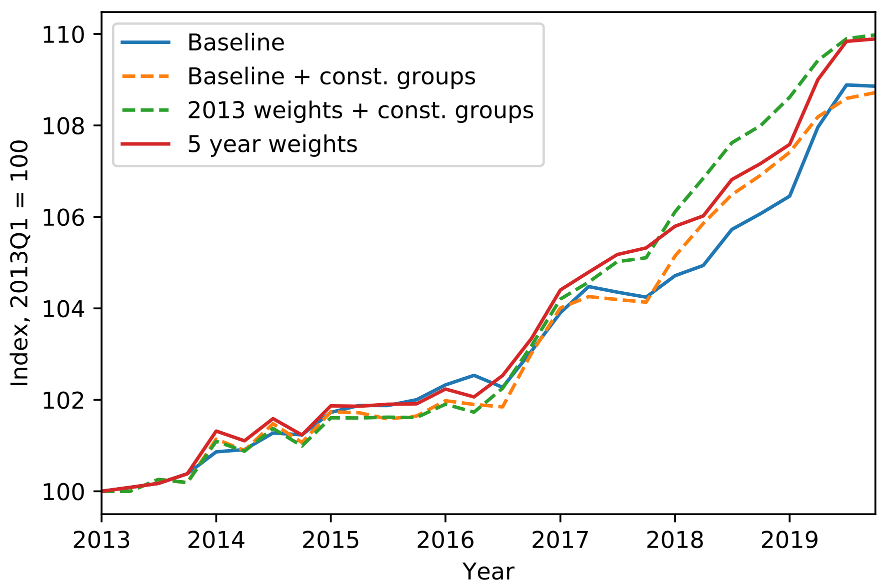

Figure 10: TPI for different weighting schemes at the product group level

4.6 Log-normal approximation

Our main result from the analysis so far is that the aggregate price index is sensitive to the values of elasticities of substitution. In general, published price indices assume low values of elastities compared to those estimated from data, and typically used in the literature on e.g. economic growth. Can this tension be overcome by assuming a joint log-normal distribution of price growth factors and weights? This section provide counterfactual indices that make this assumption, using Proposition 1. The price index for a product group g given by a log-normal approximation as in equation (10) is denoted $\widehat{\frac{P_{t,g}}{P_{0,g}}}$ and the aggregate index is denoted by $\widehat{\frac{P_t}{P_0}}$. When reporting the distributions of approximation errors we use the same sample restrictions as in Section 3.4. This is to ensure that approximation errors do not have a mass point at zero due to e.g. imputation of item prices. However, we do not apply these restrictions when comparing the aggregate index under a bivariate log-normal distribution assumption to the official index.

Table 12 and Figure 11 show the distribution of approximation errors of annual price growth factors on the product group level, as given by $100 \times \left(\widehat{\frac{P_{t,g}}{P_{0,g}}} - \frac{P_{t,g}}{P_{0,g}}\right) / \frac{P_{t,g}}{P_{0,g}}$, for the baseline index with all elasticities of substitution across items within the same group set to $\sigma_g = 0$ for goods and to $\sigma_g = 1$ for services. Errors are small throughout the distribution, with medians of about zero percent and means of 0.13 and 0.25 percent for goods and services, respectively.

Group-level approximation error distribution

	25^{th}	50^{th}	75^{th}	Mean
Goods (%)	-0.23	0.02	0.35	0.13
Services (%)	-0.20	-0.00	0.26	0.25

Table 12: Percentiles and mean of group-level annual price growth approximation errors. Each observation is given by a group-year pair, and all elasticities of substitution between products within product groups are set to $\sigma_g = 0$ for goods and $\sigma_g = 1$ for services. We only include group-year observations containing at least 4 items.

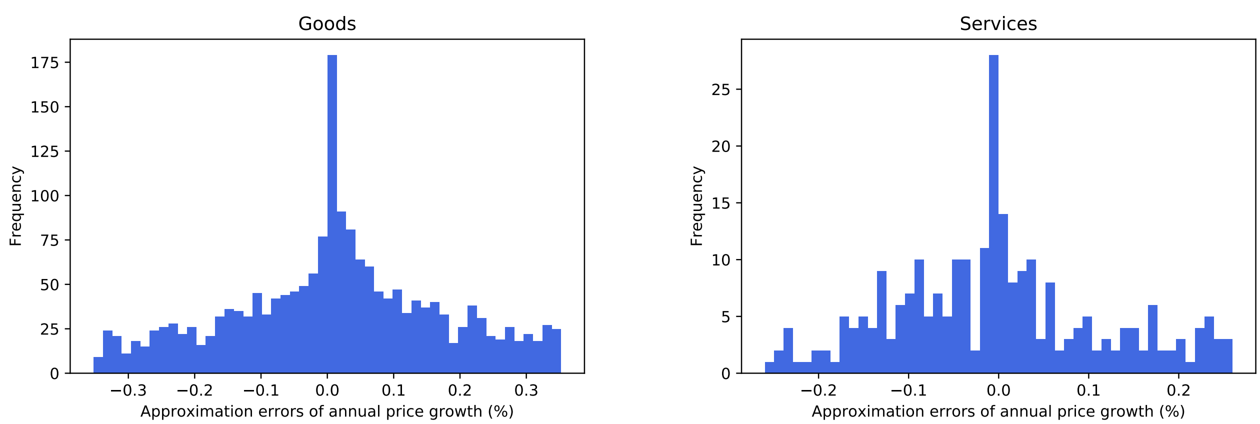

Figure 11: Distribution of group-level annual price growth approximation errors for goods (left figure) and services (right figure). Each observation is given by a group-year pair, and all elasticities of substitution between products within product groups are set to $\sigma_g = 0$ for goods and $\sigma_g = 1$ for services. The horizontal axis is cut off symmetrically around zero by the limits $\pm \max(|25^{\text{th}}$ percentile$|, |75^{\text{th}}$ percentile$|)$. We only include group-year observations containing at least 4 items.

Figures 12 and 13 show time series for the aggregate indices under the log-normal assumption, compared to the true indices with $\sigma_g = 0$ and $\sigma_g = 4$ for goods, and with $\sigma_g = 1$ and $\sigma_g = 4$ for services. These figures show graphically that the errors from approximation are small, and substantially smaller than the difference between the true indices with varying elasticities of substitution. Table 13 shows the differences in average annual inflation rates. The error arising from the log-normal approximation is in absolute terms around an order of magnitude lower than the difference between the true index with varying σ_g.

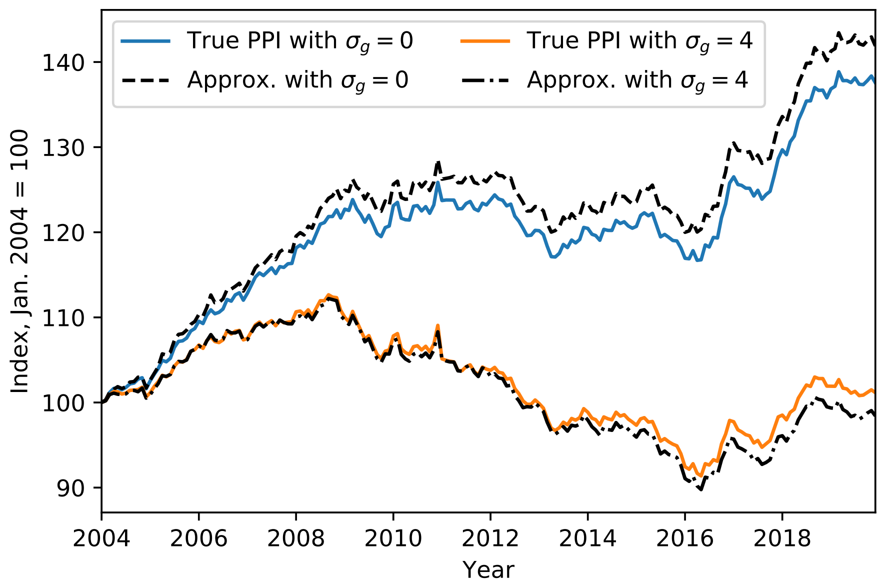

Figure 12: Comparison of true goods PPI and log-normal approximation with lower level elasticities of substitution set to $\sigma_g = 0$ and $\sigma_g = 4$, respectively. The elasticity of substitution at the higher level is set to $\sigma = 0$ for all specifications.

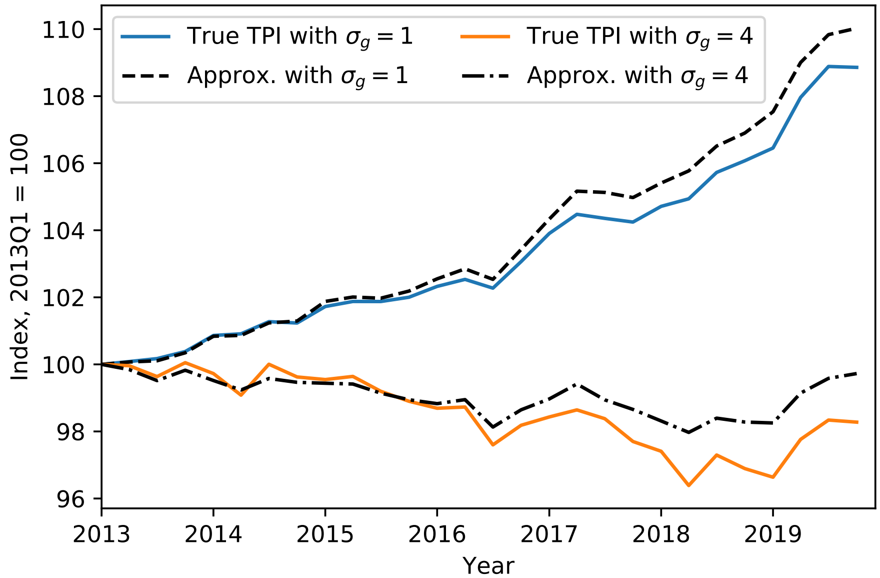

Figure 13: Comparison of true services PPI and log-normal approximation with lower level elasticities of substitution set to $\sigma_g = 1$ and $\sigma_g = 4$, respectively. The elasticity of substitution at the higher level is set to $\sigma = 0$ for all specifications.

Approximation error

	Goods		*Services*	
	$\sigma_g = 0$	$\sigma_g = 4$	$\sigma_g = 1$	$\sigma_g = 4$
True index	2.01%	0.07%	1.22%	−0.25%
Difference to true index in pp.				
Log-normal approximation	0.20	−0.17	0.15	0.21

Table 13: Annualized inflation rate for the true index and difference in inflation rate of log-normal approximation between 2004–2019 for goods and 2013–2019 for services. Elasticities of substitution between products within product groups are set to $\sigma_g = 0$ and $\sigma_g = 4$ for goods, and $\sigma_g = 1$ and $\sigma_g = 4$ for services. The elasticity of substitution across product groups is set to $\sigma = 0$ in all columns.

From the analysis above, we conclude that the error from a log-normal approximation of the aggregate price index shows non-trivial behavior, in particular because it is hard to know ex ante how the sign of the error changes with the levels of elasticities of substitution. However, for both goods and services, and for different assumptions of the values of elasticities, the error ranges between 0.15–0.21 percentage points of annual inflation in absolute terms. This is a relatively modest error when compared to not using a log-normal approximation. For example, we found in Section 2 that increasing σ_g to 4 lead to a bias of the baseline index of about 1.5–1.9 percentage points.

5 Conclusion

This study quantitatively assesses the implications of different methods of price index construction for aggregate inflation. We show that various statistical methods used worldwide result in significant differences in measured inflation when applied to the same data. Using the example of the producer price index (PPI), changing the index construction method at the lower level from the one used in Sweden to the one used in Spain or Denmark increases or decreases the aggregate inflation rate, on the same underlying micro data, by almost half a percentage point per year, respectively.

We build an economic framework, a two-tier nested-CES structure, that nests as the ideal price index the different statistical indices used in practice. Then, using Swedish micro data, we estimate the elasticities of substitution between products within narrowly-defined product groups and across product groups for both goods and services. Measuring inflation through the implied ideal price index with estimated elasticities lowers the aggregate inflation rate relative to the official price index by 3.9 percentage points per year in the case of goods, resulting in an annual inflation rate of -1.9% in 2004–2019. This result is driven by the fact that the official price index assumes no substitutability between products. In contrast, the median elasticity of substitution across items within product groups that we estimate for goods is 4.37. Finally, we repeat the same exercises for the Swedish Service Price Index and find quantitatively similar results.

The fact that different methods to construct the price index have a sizeable effect on the aggregate inflation rate hampers the comparability of inflation rates across countries. We show that the approximation error from knowing three moments and making a joint log-normality assumption of product price changes and their weights is small compared to the true index. With information on the mean and variance of log price changes as well as their covariance with the product weights, the implied inflation rate of any CES structure can be approximated accurately without access to the underlying micro data. Those three moments could be provided by statistical agencies along with the usually reported index value allowing for a method-consistent international comparison of inflation rates.

In this paper, we use the PPI to study the implications of the index construction method for aggregate inflation and, through its role as a nominal output deflator, on real output growth. It would be interesting to extend our exercises to the Consumer Price Index, commonly used as a monetary policy target or in the wage bargaining process by labor unions. We leave this exercise for future research.

References

Aghion, P., Bergeaud, A., Boppart, T., Klenow, P. J. and Li, H. (2019), 'Missing growth from creative destruction', *American Economic Review* **109**(8), 2795–2822.

Baumol, W. J. (1967), 'Macroeconomics of unbalanced growth: The anatomy of urban crisis', *The American Economic Review* **57**(3), 415–426.

Bennet, T. (1920), 'The theory of measurement of changes in cost of living', *Journal of the Royal Statistical Society* **83**(3), 455–462.

Bils, M. and Klenow, P. (2004), 'Some evidence on the importance of sticky prices', *Journal of Political Economy* **112**(5), 947–985.

Boskin, M. J., Dulberger, E. R., Gordon, R. J., Griliches, Z. and Jorgenson, D. (1996), *Toward a More Accurate Measure of the Cost of Living: Final Report to the Senate Finance Committee from the Advisory Commission to Study the Consumer Price Index december 4, 1996*.

Braithwait, S. D. (1980), 'The substitution bias of the laspeyres price index: An analysis using estimated cost-of-living indexes', *The American Economic Review* **70**(1), 64–77.

Broda, C. and Weinstein, D. E. (2006), 'Globalization and the Gains From Variety', *The Quarterly Journal of Economics* **121**(2), 541–585.

Broda, C. and Weinstein, D. E. (2010), 'Product creation and destruction: Evidence and price implications', *American Economic Review* **100**(3), 691–723.

Carlsson, M. and Nordström Skans, O. (2012), 'Evaluating microfoundations for aggregate price rigidities: Evidence from matched firm-level data on product prices and unit labor cost', *American Economic Review* **102**(4), 1571–95.

Diewert, W. E. (1976), 'Exact and superlative index numbers', *Journal of Econometrics* **4**(2), 115–145.

Diewert, W. E. (1988), 'The early history of price index research'.

Diewert, W. E. (1995), Axiomatic and economic approaches to elementary price indexes, Working Paper 5104, National Bureau of Economic Research.

Diewert, W. E. (2021), 'The economic approach to index number theory', *Consumer Price Index Theory* .

Divisia, F. (1926), 'L'indice monétaire et la théorie de la monnaie (suite et fin)', *Revue d'économie politique* **40**(1), 49–81.

Feenstra, R. C. (1994), 'New product varieties and the measurement of international prices', *The American Economic Review* **84**(1), 157–177.

Fisher, I. (1922), *The making of index numbers: a study of their varieties, tests, and reliability*, number 1, Boston: Houghton Mifflin Company, 1923 [c1922].

Goldberg, P. K. and Hellerstein, R. (2011), 'How Rigid are Producer Prices?', *Federal Reserve Bank of New York Staff Reports* **407**.

Hottman, C. J., Redding, S. J. and Weinstein, D. E. (2016), 'Quantifying the sources of firm heterogeneity', *The Quarterly Journal of Economics* **131**(3), 1291–1364.

IMF (2004), 'Producer price index manual: Theory and practice', *Published for ILO, IMF, OECD, UN, Eurostat, The World Bank by IMF, Washington, DC* .

Jevons, W. S. (1884), *Investigations in currency and finance*, Macmillan and Company.

Klenow, P. J. and Kryvtsov, O. (2008), 'State-Dependent or Time-Dependent Pricing: Does it Matter for Recent U.S. Inflation?', *The Quarterly Journal of Economics* **123**(3), 863–904.

Konüs, A. A. (1939), 'The problem of the true index of the cost of living', *Econometrica: Journal of the Econometric Society* pp. 10–29.

Konüs, A. A. and Byushgens, S. S. (1926), 'K probleme pokupatelnoi cili deneg', *Voprosi konyunkturi* **2**(1), 151–172.

Laspeyres, E. (1871), 'Die berechnung einer mittleren waarenpreissteigerung', *Jahrbücher für Nationalökonomie und Statistik* **16**, 296–314.

Lloyd, P. J. (1975), 'Substitution Effects and Biases in Nontrue Price Indices', *American Economic Review* **65**(3), 301–313.

Lowe, E. J. (1824), *The present state of England in regard to agriculture, trade and finance: with a comparison of the prospects of England and France*, E. Bliss and E. White.

Martin, R. S. (2020), *Changing Tastes Versus Specification Error in Cost-of-living Measurement*, US Department of Labor, US Bureau of Labor Statistics.

Moulton, B. R. (1996), 'Constant Elasticity Cost-of-Living Index in Share Relative Form', *Unpublished U.S. Bureau of Labor Statistics manuscript* .

Moulton, B. R. and Smedley, K. E. (1995), A comparison of estimators for elementary aggregates of the cpi, *in* 'Washington: Bureau of Labor Statistics Mimeo'.

Nakamura, E. and Steinsson, J. (2008), 'Five Facts about Prices: A Reevaluation of Menu Cost Models', *The Quarterly Journal of Economics* **123**(4), 1415–1464.

OECD (2011), 'Producer price indices – comparative methodological analysis'.

Paasche, H. (1874), 'Über die preisentwicklung der letzten jahre nach den hamburger börsennotirungen', *Jahrbücher für Nationalökonomie und Statistik* pp. 168–178.

Redding, S. J. and Weinstein, D. E. (2020), 'Measuring Aggregate Price Indices with Taste Shocks: Theory and Evidence for CES Preferences', *The Quarterly Journal of Economics* **135**(1), 503–560.

Reinsdorf, M. B. and Moulton, B. R. (1996), *The Construction of Basic Components of Cost-of-Living Indexes*, University of Chicago Press, pp. 397–436.

Sato, K. (1976), 'The ideal log-change index number', *The Review of Economics and Statistics* **58**(2), 223–28.

SOU (2002), 'Utveckling och förbättring av den ekonomiska statistiken, bilaga 3: Beräkningsrutiner för nationalräkenskaperna', *Statens Offentliga Utredningar* **2002**(118), 131–157.

Statistics Sweden (2019*a*), 'Kvalitetsdeklaration', *Prisindex i producent- och importled* .

Statistics Sweden (2019*b*), 'Statistikens framställning', *Prisindex i producent- och importled* .

Vartia, Y. O. (1976), 'Ideal log-change index numbers', *Scandinavian Journal of Statistics* **3**(3), 121–126.

Vermeulen, P., Dias, D. A., Dossche, M., Gautier, E., Hernando, I., Sabbatini, R. and Stahl, H. (2012), 'Price setting in the euro area: Some stylized facts from individual producer price data', *Journal of Money, Credit and Banking* **44**(8), 1631–1650.

Appendices

A Data

This section provides additional summary statistics. First, price change frequencies and price change sizes are reported and related to previous studies on price change frequencies and sizes in the USA and EU. Finally, summary statistics on moments used for the log-normal approximation of price indices are reported.

A.1 Replication of official aggregate index

A.1.1 Goods

The official aggregate PPI can be replicated with our micro data from 2004 due to missing observations between 1992–2003. For this reason, we focus on the period 2004–2019. Figure 14 shows the aggregation of our micro data compared to the aggregate goods PPI.

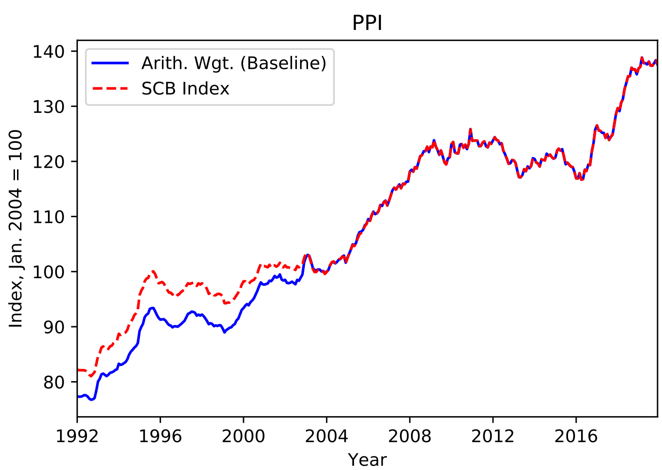

Figure 14: Statistics Sweden's published aggregate price index and the corresponding aggregate index using our micro data.

A.1.2 Services

Figure 15 shows the official TPI provided by Statistics Sweden and the replication using our data.

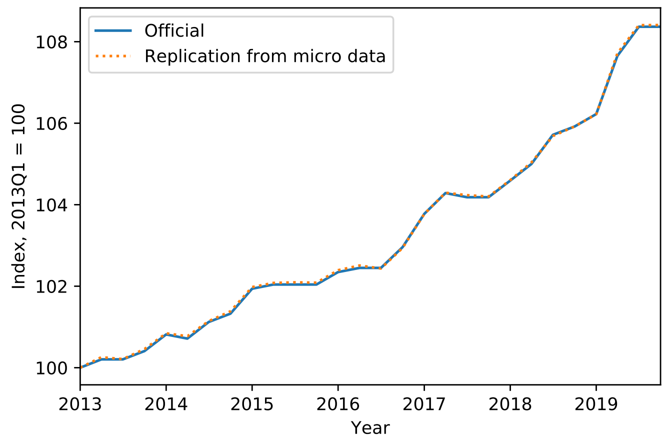

Figure 15: Statistics Sweden's published aggregate price index and the corresponding aggregate index using our micro data.

The level of aggregation for the TPI has been changing in 2014. To keep the level of aggregation constant throughout the years when changing the method of aggregation for the counterfactual exercise we compute an alternative TPI index where the level of aggregation is the same for all years. Figure 16 shows the official TPI provided by Statistics Sweden and an alternative TPI index where item prices are aggregated at the SPIN5 level for all years.

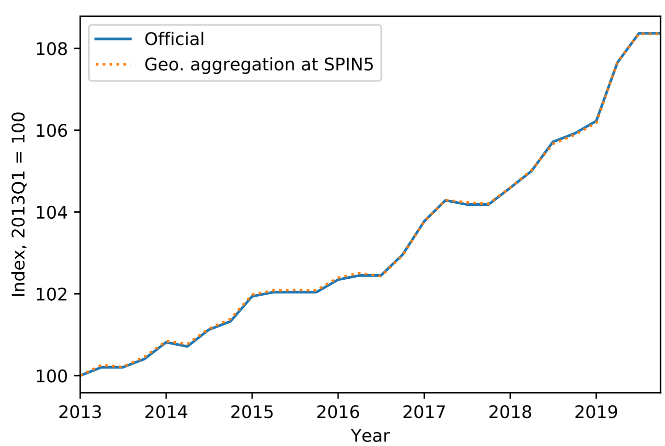

Figure 16: Comparing the official TPI with an index that uses weighted geometric aggregation of item price changes up to the SPIN5 level for all years.

Several sub indices of other aggregate price indices enter the TPI. For example Statistics Sweden imports product group indices from the CPI for the TPI computation as pre-compiled indices. Those indices would be unaffected by changing the lower-level of item price growth aggregation and hence bias the difference between different aggregation methods at the aggregate level to zero.

We hence compute an alternative TPI where we only use micro data and exclude all imported indices. Figure 17 shows the comparison.

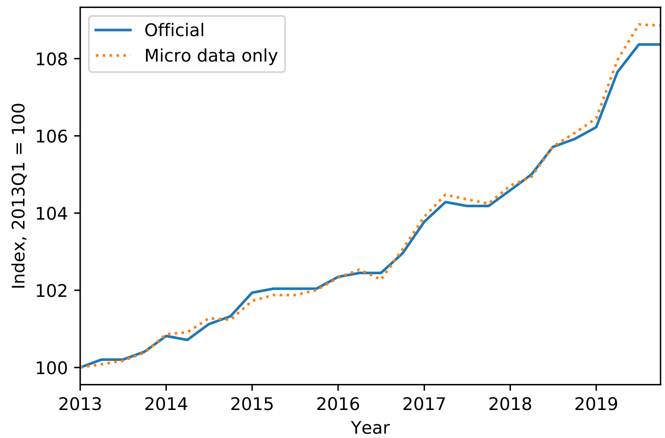

Figure 17: Comparing the official TPI with an index that excludes imported sub-indices in the TPI computation (and aggregates item price changes at the SPIN5 level for all years).

A.2 Summary statistics of shares of new items, substitutions, price change frequencies, and price change sizes

A.2.1 Goods

We break down the sample into subsets consisting of new items and existing items, and then partition the sample of existing items into product substitutions, price changes, and items experiencing no change at all, respectively. The result is provided further down in Figure 18. In the subsequent paragraphs leading up to that figure, we define the various subsets and the general structure of the sample. While the official goods PPI excludes the agriculture, forestry, and fishing industries, we include them in the sample used in this section.

The sample is constructed on a yearly basis, at the same time as new sales, and hence weights, are computed. In each year, new item identifiers (IDs) that were previously not visible enter the sample, which represent the newly introduced items in that year. The other item IDs, which were observed also in the previous year, represent existing, or continuing, items. The fact that re-weighting of the whole sample happens once per year has the implication that about 90% of new items enter the sample in January. The sample of new and continuing items is constructed as follows. Items are drawn with certainty if their sales share is sufficiently large in relation to the total sales of similar products, defined by one or a collection of 5-digit product groups. Items with smaller sales shares are instead drawn with probability, according to PPS, and are then assigned imputed sales so to represent a set of smaller items.[14] To minimize the small producers' burden of reporting, part of the sample of items drawn with PPS is rotated every year. Since items that are rotated

[14] In particular, items drawn with PPS get a value of sales such that the remainder of total sales, i.e., total sales

into the sample enter mechanically, new item IDs introduced to the sample due to rotation should not be interpreted as endogenous entry of new products. Thus, to get a measure of new items that enter endogenously, rotations should be excluded from the set of items with new item IDs. This is straightforward for the period 2015–2019, when the data contains an indicator variable flagging items drawn with probability, but not possible in earlier years. We therefore only report the share of new items for 2004–2019 in Figure 18 below, but emphasize that the share of new items of 1.1% in that figure does not distinguish between rotated new items and new items that enter endogenously. For the period 2015–2019, rotations make up about half of newly introduced items.

While new items show up in the PPI micro data as observations with new item IDs, products represented by existing items that are substituted to similar but not directly comparable products instead keep the same item ID. Since the ID stays the same, it is not possible to distinguish non-comparable product substitutions by merely observing new item IDs. Instead, product substitutions are identified by observing changes in an item's *base price*. When a product substitution occurs, Statistics Sweden must take a stand on the comparability between the new product and the old that is being replaced, and how to adjust prices of non-comparable substitutions. This decision shows up in the right panel of Table 14. For a given product with no change in characteristics, the base price refers to the last December price. Thus whenever the base price changes within a year, or when the base price in January does not equal the observed price in December in the year before, the substituted product is regarded as not directly comparable to the old product and its price must undergo either a quality or quantity adjustment before being compared to the price prior to substitution. Whenever there is such an unexpected change in an item's base price, we classify it as a product substitution.

To identify price changes, in the sense that an item changes its price without simultaneous changes in characteristics such as quantity or quality, we observe items that experience a change in their nominal price but not a change in their base price. This is illustrated in the left panel of Table 14.

less the sales of items drawn with certainty, are divided proportionally across items drawn with probability. For a more detailed description of the relationship between sales in the PPI micro data and actual sales, see Appendix B.

Price change					
Year	Month	Item	Sales	p	p^b
2014	Jan	1	3,000	**50**	50
2014	Feb	1	3,000	**55**	50
2014	Mar	1	3,000	57	58
⋮	⋮	⋮	⋮	⋮	
2014	Dec	1	3,000	57	58
2015	Jan	1	3,200	57	58

Product substitution					
Year	Month	Item	Sales	p	p^b
2014	Jan	1	3,000	50	50
2014	Feb	1	3,000	55	**50**
2014	Mar	1	3,000	57	**58**
⋮	⋮	⋮	⋮	⋮	⋮
2014	Dec	1	3,000	**57**	58
2015	Jan	1	3,200	57	**58**

Table 14: Example of a made-up item, with item ID given by 1. *Sales* show the value of item sales, which is updated each year. p and p^b denotes the item's observed price and its base price, respectively.

The data also includes a flag for transactions, indicating if an item was not sold, or if the firm did not respond to the survey in a given month, which can for instance happen during summer vacations. We remove these observations when reporting summary statistics. The weighted share of such missing reporting accounts for 5.5% of the monthly observations in 2004–2019, and removing them reduces the sample to 562,632 observations.

The weighted share of monthly observations of item type τ (e.g. new items, product substitutions, etc) is computed as follows

$$\text{Monthly share of observations of type } \tau \equiv \frac{1}{|A|} \sum_{a \in A} \frac{\sum_{j \in J_{\tau,a}} s_j}{\sum_{j \in J_a} s_j}. \tag{11}$$

$A = [\underline{a}, \bar{a}]$ is the set of years, and $|A|$ is its cardinality. J_a is the set of item-periods that are in the PPI sample in a given year $a \in A$. An element in J_a is an item-period pair, $j = \{i, t\}$, where a period t is given by a particular month-year combination. Hence if an item is observed in $M \in [1, 12]$ months in year a, that item also occurs M times in the sum in the denominator. In the PPI micro data, sales represent annual nominal sales, hence (11) yields a monthly rate by summing the same annual sales of an item M times.[15] $J_{\tau,a} \subset J_a$ is a subset of item-periods restricted to only include items of type τ, for example new items, existing items, substitutions, or price changes. The PPI micro data includes item sales $s_{i,t}$ that correspond to yearly nominal sales of a given item, so $s_{i,t}$ is constant for a given item for all months in the same year. Since nominal sales on average increases across years, we first compute (11) for each year separately in 2015–2019 and then take the unweighted average across years.

[15]Suppose that an item with annual sales 1,000,000 SEK is substituted once in a year. The monthly rate of substitution is then given by $\frac{1{,}000{,}000}{M \times 1{,}000{,}000}$. $M < 12$ if the item is not observed in all months, e.g. due to not being transacted in all months.

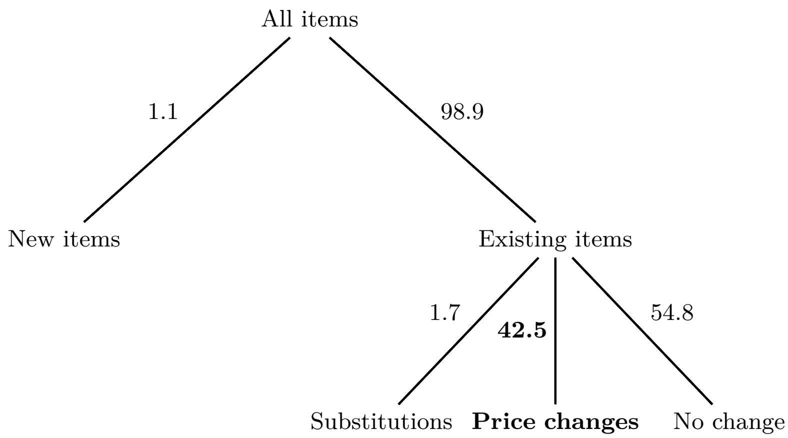

Figure 18: Shares of item adjustments on a monthly frequency in the goods PPI micro data in 2004–2019. *Substitutions* are non-comparable product substitutions that experience a quantity or quality adjustment in prices, *Price changes* are price changes that occur for the same item without a simultaneous quantity or quality change.

Figure 18 reports the monthly shares of different types of adjustments. The weighted share of items that undergo a non-comparable product substitution between two adjacent months is 1.7%, and 1.1% are newly introduced. We focus on the other type of adjustment, price changes of existing products, which amounts to 42.5% of all goods in a given month. We do a thorough analysis of the effects of quality adjustments from product substitution in a companion paper.

A.2.2 Price change frequencies and sizes

Before estimating monthly price change frequencies and sizes, we exclude items that are either new or non-comparable substitutions. Items that are newly introduced into the sample have no information on their price in the previous month before being observed in the data, so it is impossible to know if their price has changed. Non-comparable product substitutions experience price changes partly due to changes in quantity or quality, hence including them would lead to an upward bias of the price change frequency. Thus, focusing on comparable products captures pure price changes for a constant set of products over time. Table 15 illustrates how the sample for estimating price change frequencies and sizes is constructed. The example shows the first 6 observations in the raw data for the same made-up item as in Table 14. In addition to Table 14, the example includes a variable indicating whether the observation is included in the final sample. At the first period, January 2014, the item is newly introduced and thus excluded from the sample. In February, the price p is different from what it was in the initial period, indicating that the item experienced a price change. The third observation is dropped due to a comparable product substitution, as indicated by the base price p^b, in March. The observation in May is also dropped due to there being no transaction in that month. Finally, the observations in April and June show how price changes are constructed in the case when the previous adjacent period is excluded from the sample. After dropping observations due to a substitution, Δp indicates a price change only if the price changed after the substitution. Since p is the same in March and April, Δp indicates no change.

Thus, although the March observation will be excluded from the sample when estimating price change frequencies and sizes, information of the price p in March is still used to infer whether there was a price change in April. In contrast, the June observation compares the price in June to the last observed period, which in this case April. Including the April and June observations in the sample in this way is coherent with the approach in Klenow and Kryvtsov (2008) and Goldberg and Hellerstein (2011). However, in Nakamura and Steinsson (2008) a missing price observation yields two missing values in the price change variable because the price change variable is constructed by using only observations in adjacent months that are both included in the final sample.

Year	Month	Item	Sales	p	p^b	Included	Δp
2014	Jan	1	3,000	50	50	0	-
2014	Feb	1	3,000	55	50	1	1
2014	Mar	1	3,000	57	58	0	-
2014	Apr	1	3,000	57	58	1	0
2014	May	1	3,000	-	-	0	-
2014	Jun	1	3,000	57	58	1	0
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮

Table 15: Data example showing construction of sample used to estimate price change frequencies and sizes (Included), and of the price change variable Δp. *Included* shows if an observation is included (1) or not (0) in the sample. 1 in Δp indicates a price change, 0 indicates no change, and a dash (-) indicates a missing value due to not being included in the sample.

Throughout the analysis, we compare prices in terms of their invoicing currency such that changes in exchange rates do not alter the price change frequencies directly. We follow Goldberg and Hellerstein (2011) by excluding observations with log price changes exceeding 4. Such observations constitute less than 0.01% of price changes. As in the rest of the literature, we do not identify or exclude sales.

Before reporting the results, we start by a description of how weighted medians and means of frequencies and sizes are estimated. As described in the main text, items are weighted using nominal sales. We construct weights to sum to 1 within a year so that all years are weighted equally. In particular, in a given year a, the normalized weight of item i in month m is $v_{i,m} = s_{i,m}/S_a$, where $S_a = \sum_{m=1}^{12} \sum_{i \in I_a} s_{i,m}$ is the sum of sales over all months and all items observed in the given year, denoted by the set I_a. Since sales are updated every year, we estimate the price change frequency as follows. First, we compute the average frequency of price changes within an item and year (fr) by dividing the number of observations with a price change by the total number of observations for that item in that year (M).

$$fr_i = \frac{\sum_{m=1}^{12} \mathbb{1}(\Delta p_{i,m} \neq 0)}{M_i}$$

The weighted median and weighted mean are then estimated using the following two-step procedure:[16]

1. Estimate weighted median and mean across different items within 5-digit product groups, markets and years. For convenience, subscripts for group-market (g) and year (a) are suppressed.

Median: Compute weighted median using fr_i and normalized item weights $V_i = \sum_{m=1}^{12} v_{i,m}$.

Mean: Compute the weighted mean as

$$\frac{\sum_{i \in I} fr_i \times V_i}{\sum_{i \in I} V_i}$$

where I is the set of items in the given group-market and year.

2. Estimate weighted median and mean across different 5-digit product groups, markets, and years. Let the weighted frequencies across different group-markets and years from the first step be denoted by $fr_{g,a}$.

Median: Compute weighted median using $fr_{g,a}$ and normalized group-market-year weights $V_{g,a}$.

Mean: Compute the weighted mean as

$$\frac{\sum_{a \in A} \sum_{g \in G} fr_{g,a} \times V_{g,a}}{\sum_{a \in A} \sum_{g \in G} V_{g,a}}$$

where A and G are the sets of years and group-markets, respectively.

The weighted median and mean absolute size of price changes are estimated similarly as described above. First, observations with $\Delta p = 0$ are dropped such that the sample only includes price changes. Second, the absolute size of each price change is measured and averaged within items and years, yielding an item-year specific average absolute price change, pc_i. The weighted median and mean are then estimated using the two-step procedure described above, only using observations experiencing a price change and using the average absolute price change pc_i instead of the average frequency of price changes fr_i.

The results are presented in Table 16. The weighted median of 22.2% corresponds to prices being changed every fourth month. Conditional on a price change, the median absolute size of the change is 3.0%. Compared to the weighted median, the weighted mean gives substantially larger frequencies and price change sizes. On average, items change prices after about one and a half month, and

[16] This procedure is redundant when computing the weighted mean and using weights that are normalized to sum to 1 in each year, since it gives the same result as computing (11) in the main text. I.e. it is not of relevance to first compute a weighted mean within each product group, and then across groups. For the median, however, doing the two-step procedure matters.

when they do, the new price deviates by more than 5% in absolute terms. 48% of all price changes are price decreases.

Monthly price change frequency and size: Goods

	W. Median	W. Mean
Change frequency (%)	22.2	43.7
Duration (Months)	4.0	1.7
Absolute size of changes (%)	3.0	5.3
Size of positive changes (%)	3.0	6.4
Size of negative changes (%)	-2.7	-4.2

Table 16: Monthly frequency of price changes and conditional change size for existing items in the goods PPI not experiencing a product substitution between 2004–2019, including domestically produced and sold items and exports. *Duration* is the implied duration in months from the weighted median and mean, computed as in Klenow and Kryvtsov (2008).

As in Nakamura and Steinsson (2008), a very small fraction of products have price change frequencies close to the median. For example, the 45th percentile of price change frequencies is 11.1%, while the 55th percentile amounts to 41.7%. This implies that the weighted medians for various subsamples are markedly different. For example, the weighted median frequency in the energy industry is 90.9%, while only 16.7% in manufacturing. The weighted mean is relatively more robust, but still varies from around 43% in manufacturing to 78% in energy. The dispersion is yet much greater across finer levels of product groups. A more detailed comparison between different subsamples is provided in Appendix A.2, including a comparison between exports, imports, and domestic sales. We there also lay out a more detailed comparison to studies covering other countries. When restricting the sample to become similar to the data for the US used in Goldberg and Hellerstein (2011), we still find similar median and mean frequencies of price changes in Sweden, given by 25.0% and 43.4%. This implies that the price change frequency is higher than in the US, where the weighted median and mean equals 16.5% and 37.3%.

In Table 17, frequencies and sizes are reported for each market respectively, including imports. In the other tables, imports are excluded.

Price change frequencies and sizes across markets: Goods PPI

	Local	Exports	Imports	Full sample
Median frequency	16.7	33.3	16.7	16.7
Mean frequency	42.0	45.4	39.2	42.2
Median duration	5.5	2.5	5.5	5.5
Mean duration	1.8	1.7	2.0	1.8
Median absolute size	3.0	3.0	4.6	3.3
Mean absolute size	4.7	5.9	9.0	6.4
Median positive size	2.9	3.0	4.5	3.3
Mean positive size	5.1	7.6	11.5	7.9
Median negative size	-2.8	-2.6	-4.5	-3.1
Mean negative size	-4.3	-4.1	-6.2	-4.8
Market share	33.6	33.1	33.3	100.0

Table 17: Monthly frequency of price changes and change size, conditional on a price change, for existing items, excluding non-comparable product substitutions, in the goods PPI between 2004–2019. *Local* denotes the domestic Swedish market, while *Exports* and *Imports* denotes exports and imports. *Median duration* and *Mean duration* is the implied duration in months from the median and mean frequency, computed as in Klenow and Kryvtsov (2008). *Market share* is in terms of gross output in the goods PPI micro data.

Weighted medians and means mask substantial heterogeneity in frequencies and sizes across different types of goods. Table 18 shows differences across the five major groups in the goods PPI. For example, the median item in the energy and crude materials sectors change its price every month, while the median item in the manufacturing sector only changes its price two times per year.

Median price change frequencies and sizes: Goods

	Median f.	Mean f.	Median s.	Mean s.	Weight
Farming, forestry, & fishing	33.3	50.5	3.0	5.6	5.3
Crude materials & extraction	33.3	44.2	4.7	7.2	2.2
Manufacturing	16.7	42.6	2.9	5.0	87.1
Electr., gas, heating, & cooling	90.9	78.1	8.6	9.3	2.3
Water, sewerage, & waste	33.3	43.1	4.9	6.9	1.7
Information & communication	16.7	29.7	1.6	2.8	1.4
Goods PPI	22.2	43.7	3.0	6.4	100.0

Table 18: Price change frequencies (f.), absolute sizes (s.), and gross output shares of major groups, 2004–2019. Information & communication includes goods produced from that sector, e.g. from printing, while services are reported in the section below.

Price change frequencies and sizes across 2-digit product groups: Goods PPI

	Median f.	Mean f.	Median s.	Mean s.	Weight
Farming	66.7	62.3	3.7	6.2	2.7
Forestry	16.7	36.0	2.0	3.7	2.5
Fishing	100.0	87.4	13.8	14.4	0.1
Crude petr., natural gas	72.7	72.7	4.4	4.4	0.0
Metal ore	33.3	47.1	5.0	8.1	1.7
Misc. mineral extraction	16.7	35.4	3.3	3.9	0.6
Foods	66.7	53.3	2.4	4.4	7.2
Beverages	8.3	20.0	3.1	5.4	1.0
Tobacco	8.3	10.6	2.5	5.5	0.3
Textiles	8.3	24.4	3.1	6.1	0.4
Apparel	0.0	18.0	4.7	6.2	0.1
Hides, skins, leather	91.7	53.7	3.8	5.2	0.1
Lumber & wood	50.0	51.6	2.6	3.4	4.6
Paper	100.0	76.1	2.2	3.4	6.9
Printing services & recording	8.3	21.7	4.8	5.8	0.8
Coal & refined petr.	100.0	86.0	5.2	6.1	5.3
Chemicals	66.7	51.4	3.2	5.6	4.3
Pharmaceuticals	0.0	11.3	1.3	10.5	3.6
Rubber & plastic	16.7	37.9	3.2	4.8	2.4
Misc. non-metallic mineral	9.1	33.3	1.8	3.8	1.9
Metals	100.0	77.3	3.6	5.3	7.2
Misc. metal products	8.3	24.0	2.8	5.4	6.0
Computers, electronics, optics	91.7	50.4	1.4	3.8	4.7
Electric devices	8.3	31.4	3.5	7.4	3.3
Misc. machinery	8.3	18.8	3.0	8.8	9.8
Motor vehicles	9.1	27.0	2.1	4.9	11.2
Misc. transportation	9.1	29.3	0.7	3.5	1.4
Furniture	8.3	16.5	3.8	5.9	1.4
Misc. manufactured goods	0.0	21.5	3.5	5.5	1.1
Repair/install for machi. & dev.	8.3	9.5	3.0	5.1	2.0
Electr., gas, heating, cooling	90.9	78.1	8.6	9.3	2.3
Water	16.7	18.3	1.3	1.3	0.7
Waste disposal	83.3	62.3	5.3	8.2	1.0
Decontamination	0.0	16.5	17.2	13.6	0.0
Printing & publishing	16.7	30.2	1.6	2.8	1.4
Film/video/TV/audio	0.0	2.1	9.1	11.8	0.0

Table 19: Monthly price change frequencies (f.), absolute sizes (s.), and gross output shares of 2-digit product groups, 2004–2019.

A.2.3 Services

As for the goods PPI, we exclude all quarters in which items are not sold or not reported, resulting in dropping 9.7% of the sample.

The services PPI micro data has been continuously developed during the past years, possibly to a greater extent than the goods PPI, which has been collected over a longer period. The developments in collecting services PPI result in larger shares of new product groups and items that are added over time. In addition, there has been an increased coverage of different markets. In the initial years 2013–2017, market information was not collected, which lead to that observations in those years can indistinguishably refer to services sold either domestically or as exports. In 2018–2019, the data contains a market identifier and the sample covers services sold domestically, as well as exports and imports. This mechanically drives up the share of new items in the exports and imports markets to 100% in 2018. Due to these developments, we drop all imports, but keep exports in 2018–2019 since they are included but not separated from domestically sold items in 2013–2017. Statistics for various markets are presented in detail in Appendix A.2.

Figure 19 shows the adjustment shares on a quarterly frequency for services. Similar to the corresponding Figure 18 for goods, most items experience either no adjustment at all or a price change between two quarters.

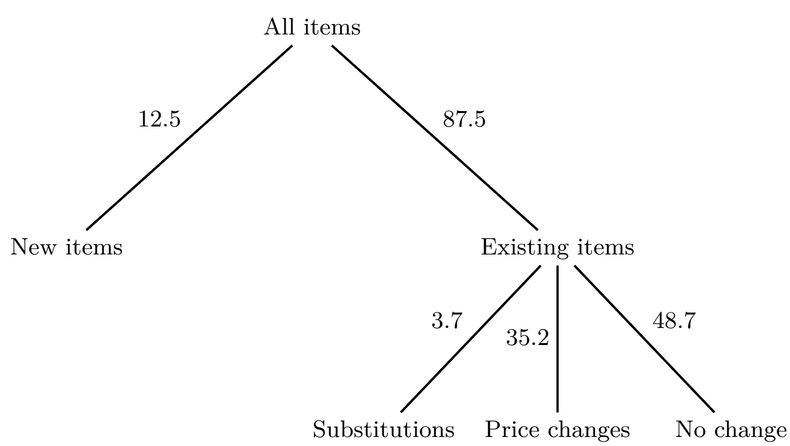

Figure 19: Shares of item adjustments on a quarterly frequency in the services PPI micro data in 2014–2019. *Substitutions* are non-comparable product substitutions that experience a quantity or quality adjustment in prices, *Price changes* are price changes that occur for the same item without a simultaneous quantity or quality change.

A.2.4 Price change frequencies and sizes

The same sample restrictions is applied for services as those for goods described above. Less than 0.1% of the sample is dropped due to excluding price changes larger than four log points. To compare goods to services prices, which are collected in the end of each quarter, we construct a dataset of quarterly frequency for the goods PPI micro data by only keeping observations reported in March, June, September, and December. Table 20 shows the quarterly price change frequencies and sizes, estimated in the same way as described above for the monthly data.

Quarterly price change frequency and size: Services vs goods

	Median		**Mean**	
	Services	Goods	Services	Goods
Change frequency (%)	25.0	33.3	41.9	52.7
Duration (Quarters)	3.5	2.5	1.8	1.3
Absolute size of changes (%)	3.5	4.1	7.3	8.4
Size of positive changes (%)	3.0	3.8	7.6	9.2
Size of negative changes (%)	-3.8	-3.9	-7.0	-7.6

Table 20: Quarterly frequency of price changes and change size, conditional on a price change, for existing items in the services and goods PPI, including domestically sold items and exports 2014–2019. *Duration* is the implied duration in quarters from the weighted median and mean, computed as in Klenow and Kryvtsov (2008).

Table 20 shows that goods price are, in general, changed more frequently than prices of services. Goods prices also experience larger changes than services. 41.6% of the price changes of services were price decreases.

Table 21 shows price change frequencies and sizes for services for domestically sold products and exports. Before 2018, there was no distinction between exports and locally sold services since every sold item was reported as *Local*, which affects the market shares in the table. The table excludes imports. In 2019, the weighted share of domestically sold products, exports, and imports are 78.3%, 13.6%, and 8.1%, respectively.

Price change frequencies and sizes across markets: Services

	Local	Exports	Full sample
Median frequency	25.0	75.0	25.0
Mean frequency	41.1	64.9	41.9
Median duration	3.5	0.7	3.5
Mean duration	1.9	1.0	1.8
Median absolute size	3.5	1.9	3.5
Mean absolute size	7.4	5.1	7.3
Median positive size	3.0	0.7	3.0
Mean positive size	7.6	5.7	7.6
Median negative size	-3.8	-1.9	-3.8
Mean negative size	-7.0	-4.5	-7.0
Market share	96.4	3.6	100.0

Table 21: Quarterly frequency of price changes and change size, conditional on a price change, for existing items, excluding non-comparable product substitutions, in the services PPI between 2014–2019 for the domestic market, and exports. *Local* denotes the domestic Swedish market, while *Exports* denotes exports. *Median duration* and *Mean duration* is the implied duration in quarters from the median and mean frequency, computed as in Klenow and Kryvtsov (2008). *Market share* is in terms of gross output in the services PPI micro data. Before 2018, there was no distinction between exports and locally sold services since every sold item was reported as "Local", which affects the market shares in the table.

Price change frequencies and sizes across markets: Services PPI

	Local	Exports	Imports	Full sample
Median frequency	25.0	75.0	33.3	33.3
Mean frequency	45.7	64.9	43.6	47.6
Median duration	3.5	0.7	2.5	2.5
Mean duration	1.6	1.0	1.7	1.5
Median absolute size	3.5	1.9	8.9	3.4
Mean absolute size	7.6	5.1	11.2	7.4
Median positive size	2.8	0.7	7.1	2.8
Mean positive size	7.9	5.7	11.2	7.8
Median negative size	-4.0	-1.9	-5.2	-3.8
Mean negative size	-7.1	-4.5	-11.2	-6.9
Market share	78.6	12.7	8.7	100.0

Table 22: Quarterly frequency of price changes and change size, conditional on a price change, for existing items, excluding non-comparable product substitutions, in the services PPI between 2018–2019 for the domestic market, exports, and imports. *Median duration* and *Mean duration* is the implied duration in quarters from the median and mean frequency, computed as in Klenow and Kryvtsov (2008). *Market share* is in terms of gross output in the services PPI micro data.

Median price change frequencies and sizes: Services

	Median f.	Mean f.	Median s.	Mean s.	Weight
Transportation & warehousing	66.7	48.4	3.7	6.2	18.8
Hotels & restaurants	100.0	77.9	10.0	13.6	3.2
Information & communication	25.0	39.0	3.2	6.2	20.6
Real estate	33.3	40.8	3.1	9.7	21.8
Legal, business, & engineering	25.0	37.1	3.3	6.1	24.7
Rental, leasing, property, & office	25.0	39.6	3.8	6.0	10.2
Misc. (sports, laundry, hairdressing, funeral)	25.0	30.6	2.6	6.5	0.7
Services PPI	25.0	41.9	3.5	7.3	100.0

Table 23: Quarterly price change frequencies (f.), absolute sizes (s.), and gross output shares of major groups, 2014–2019. Imports excluded.

Price change frequencies and sizes across 2-digit product groups: Services

	Median f.	Mean f.	Median s.	Mean s.	Weight
Land & pipeline transportation	0.0	28.8	2.6	4.8	5.5
Water transportation	100.0	69.7	4.8	7.7	1.5
Air transportation	100.0	89.1	6.6	8.0	1.9
Warehousing	75.0	53.2	3.2	6.0	7.8
Courier & postal	25.0	29.3	2.2	3.4	2.0
Hotels & accommodation	100.0	75.6	7.8	10.4	2.2
Bars & restaurants	0.0	11.7	7.0	7.0	0.0
Publishing	25.0	45.7	2.9	4.8	4.9
Audio & video	0.0	9.8	4.0	5.5	0.6
Broadcasting (& scheduling)	100.0	72.6	9.7	12.8	0.2
Telecommunication	25.0	38.6	4.6	7.0	3.1
Computer programming	25.0	38.9	2.9	6.4	11.1
Information	0.0	17.4	4.2	8.6	0.9
Banking	100.0	99.0	7.0	15.9	0.8
Real estate	33.3	40.8	3.1	9.7	21.8
Legal & accounting	25.0	39.8	4.0	5.4	3.9
Business consultancy	0.0	23.4	4.5	6.3	5.9
Architectural & technical	25.0	47.3	1.7	3.0	8.2
Advertising & marketing	0.0	24.7	6.7	10.9	3.2
Misc. legal, business	25.0	29.7	3.8	7.3	1.5
Rental & leasing	25.0	39.2	4.0	8.2	3.3
Employment & staffing	75.0	57.6	3.8	5.0	2.8
Travel	100.0	84.8	11.5	20.2	1.0
Security	25.0	26.0	3.0	4.8	1.0
Property maintenance	25.0	28.0	2.7	4.9	3.1
Office & administration	25.0	30.0	3.5	9.1	1.3
Repair (electronics, instruments, etc.)	25.0	37.3	2.6	6.8	0.3
Misc. consumer services	25.0	24.0	2.9	5.9	0.4

Table 24: Quarterly price change frequencies (f.), absolute sizes (s.), and gross output shares of 2-digit product groups, 2014–2019. Imports excluded.

A.3 Comparison to other countries

A.3.1 Goods

This section restricts the data to be more comparable to those used for other countries. We first make a comparison to studies for the USA. Since Bils and Klenow (2004) and Klenow and Kryvtsov (2008) study consumer prices, we focus on comparing results for producer prices from Nakamura and Steinsson (2008) and Goldberg and Hellerstein (2011), which we refer to as NS and GH, respectively. Finally, we make a comparison to the results for Euro countries in Vermeulen et al. (2012).

First, imports are taken out so that the data only cover domestic production as in the PPI for the USA. Second, the time frame studied is changed in order to compare a period closer to the ones in NS and GH. Third, a difference between the US and Sweden is that the US PPI is aggregated to three different indices – finished goods, intermediate goods, and crude materials – while the Swedish PPI is aggregated to one index, including different types of goods. For this reason, we exclude the crude materials and extraction industries and compare frequencies in the Swedish data to the results for finished goods in NS and GH. As in the main part of the paper, we remove non-transactions and observations with imputed prices due to non-reporting. Like Klenow and Kryvtsov (2008) and GH, but in contrast to NS, we treat a price change after a sequence of missing values due to no-transaction as a price change.

The first two columns in Table 25 states the results from GH and NS. The third column shows the results for the Swedish PPI, subject to the sample restrictions mentioned above. These columns suggest that Swedish producer prices change more frequently than in the USA. The median item changes its price almost 5 times per year in Sweden, while about 2 times per year in the USA, as estimated by GH, or even less frequently, as estimated by NS. Conversely, the implied median duration is 2.5 months in Sweden and 5.6 months in the USA.

Monthly price change frequency and size: USA vs Sweden

	USA		**Sweden**	
	GH	NS	Including exports	Excluding exports
Weighted median (%)	16.5	10.8	33.3	25.0
Weighted mean (%)	37.3	24.7	47.3	44.0
Median duration (Months)	5.6	8.8	2.5	3.5

Table 25: Percentage points (except for durations reported in months). Swedish data: Local transactions and exports between 1993–2008, including quality or quantity adjustments but not new items. Price changes are identified by a change in the quality adjusted price. US data: GH refers to Goldberg and Hellerstein (2011) with data between 1987–2008 and NS refers to Nakamura and Steinsson (2008) with data between 1998–2005. Durations are computed as in Klenow and Kryvtsov (2008). The Swedish data here differs from the data used in the rest of the paper due to (i) a different time frame, (ii) not including imports, and (iii) not including the crude materials industries. These assumptions are made to make the sample more comparable to the ones used in GH and NS.

As in the sample used in the main paper, the median frequency tends to fall when not including exports, as shown in column 4. Since prices of exports change more frequently, excluding exports reduces the median down to 25%. However, the weighted mean frequency only drops by about three percentage points.

The computation in the third and fourth columns of Table 25 differs from the first two columns in that re-weighting occurs annually, as compared to every 5–7 years in GH. NS study a 5-year period with fixed weights, and apply a different method to compute price change frequencies using unweighted medians at the lower level of aggregation and sales-adjusted weights at the higher level. To gauge the implications from these differences, we make an additional analysis where we use the same method as in NS. First, we use years 1998–2005 to be consistent with their time frame. We follow their approach by first taking the (unweighted) median across items within 4-digit product groups, and then the weighted median across 4-digit product groups. We restrict our data by holding weights fixed over the period, equal to the SPIN4-market weights in 1998. Due to products entering and exiting the sample, the distribution of SPIN4-markets is not constant across years. We therefore only include the SPIN4-markets that are observed in all months and all years in 1998–2005, which implies that we drop 5% of the weighted sample. To get SPIN4-market weights, we sum the item weights for all item-month observations. The result from this exercise suggests that the differences in results between this paper and NS are not due to the methods used, since it yields a weighted median of 28.8% for the sample consisting of both domestic sales and exports, and 23.3% for domestic sales only.

Table 26 repeats the weighted mean for the Swedish PPI (excluding exports) and the corresponding numbers for other European countries from Vermeulen et al. (2012). The sample period for Belgium only covers years after the implementation of the Euro, while the period for other countries, except Spain, partly covers years before adopting the Euro. The sample period for Spain only covers years before adopting the Euro.

Monthly price change frequency: Sweden vs other European countries

	Sweden	Euro area	Belgium	France	Germany	Italy	Portugal	Spain
Weighted mean	44.0	20.8	23.6	24.8	21.2	15.3	23.1	21.4

Table 26: Percentage points. Swedish data: Local transactions between 1993–2008, excluding crude materials. Data for other European countries are from Vermeulen et al. (2012), covering Belgium 2001–2005, France 1994–2005, Germany 1997–2003, Italy 1997–2002, Portugal 1995–2000, and Spain 1991–1999.

We conclude from the analysis above that producer prices in Sweden change more often than in the US and than in other European countries reported in Vermeulen et al. (2012). While the weighted median is sensitive to alternative specifications, indicating that the median item changes its price about 4 times per year, the weighted mean is fairly robust around 45%, indicating that the mean item changes its price 5–6 times per year. The mean is robust to various specifications, and is substantially larger than in the US and other European countries.

A.3.2 Services

GH reports frequencies for the services PPI micro data in the US, but a comparison to the results in this paper is not straightforward since those frequencies are at a monthly frequency, while the Swedish data for services only exists on a quarterly frequency. In the US, the monthly median and mean frequencies for services are 11.9% and 30.1%, respectively. The corresponding figures for goods in Table 25, 16.5% and 37.3%, show that services change prices less frequently than goods.

As shown in Table 20, this relationship also holds qualitatively on a quarterly frequency in the Swedish data.

A.4 Summary statistics of moments used in log-normal approximation

Table 27 shows quartiles and means of the moments used to compute the lower level indices under the assumption of a joint log-normal distribution. Since prices increase more over longer periods in general, we only report moments for annual price growth factors. We also restrict the sample as in the main text, e.g. by dropping items with imputed prices due to no transaction, and we only keep group-year observations that contain at least 4 items. Below, we show results for when keeping all group-years, including those with 1–3 items.

Goods (2004–2019)

	25^{th}	50^{th}	75^{th}	Mean
μ	-0.010	0.017	0.045	0.019
δ	0.002	0.005	0.012	0.011
ρ	-0.008	-0.000	0.004	-0.003

Services (2013–2019)

	25^{th}	50^{th}	75^{th}	Mean
μ	0.001	0.015	0.029	0.015
δ	0.002	0.005	0.012	0.010
ρ	-0.006	0.000	0.006	0.001

Table 27: Percentiles and means of moments for log-normal approximation, for goods (left panel) and services (right panel), respectively, pooled over years and product groups. μ denotes the product group mean of item log price growth factors, δ the product group variance of item log price growth factors, and ρ the product group covariance between item log price growth factors and item log sales shares in the base period. Only annual price growth factors are included. We only include group-year observations containing at least 4 items.

While the moments in Table 27 are pooled over groups and years, Figures 20 and 23 instead show the average of the same moments for different years. For goods, the average price growth varies substantially over time, while the price growth variance and covariance between price growth and sales shares are constant over time. For services, both average price growth and variance of price growth trends upward over time, wile the covariance fluctuates around 0.

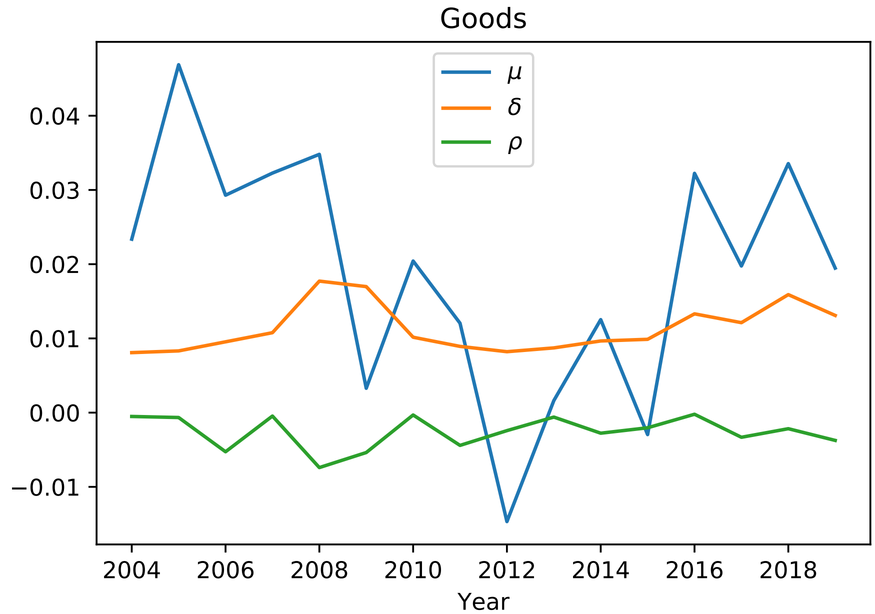

Figure 20: Average moments over time for goods. μ denotes the product group mean of item log price growth factors, δ the product group variance of item log price growth factors, and ρ the product group covariance between item log price growth factors and item log sales shares in the base period. Only annual price growth factors are included. We only include group-year observations containing at least 4 items.

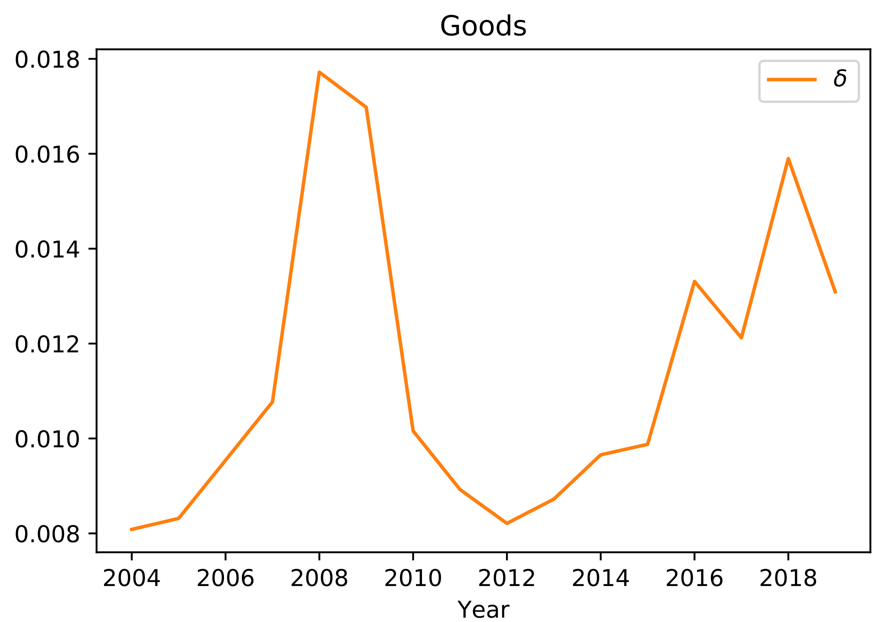

Figure 21: Average product group variance of item log price growth factors over time for goods. Only annual price growth factors are included. We only include group-year observations containing at least 4 items.

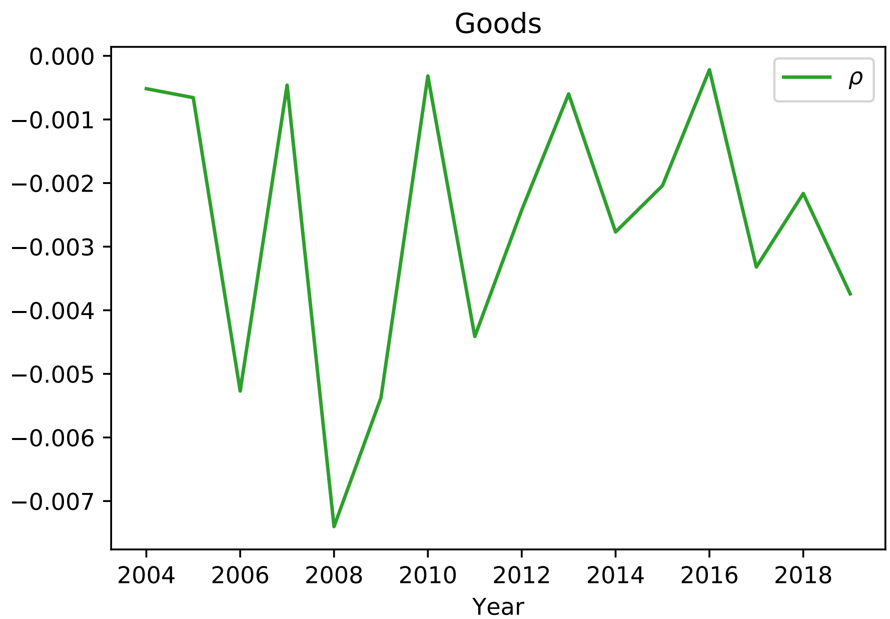

Figure 22: Average product group covariance between item log price growth factors and item log sales shares in the base period over time for goods. Only annual price growth factors are included. We only include group-year observations containing at least 4 items.

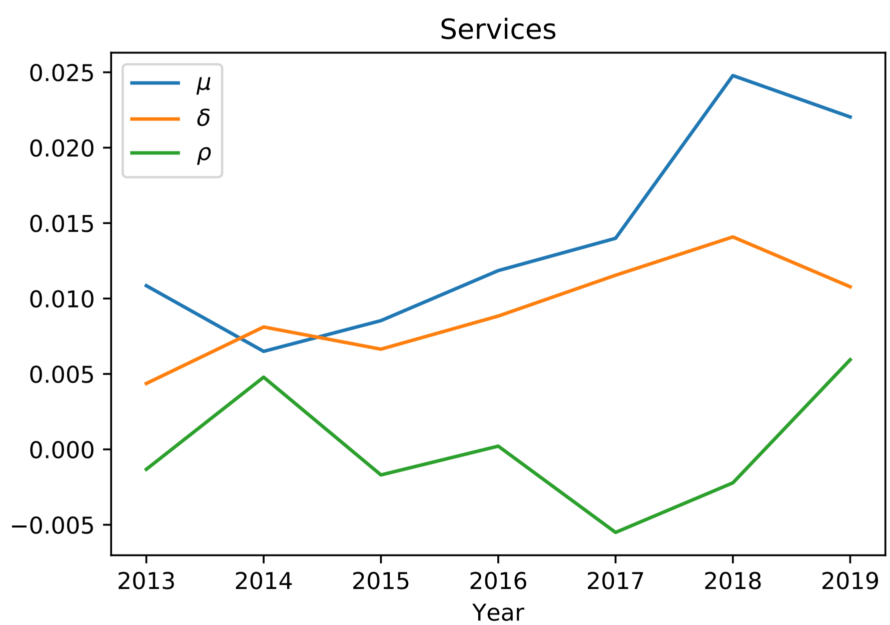

Figure 23: Average moments over time for services. μ denotes the product group mean of item log price growth factors, δ the product group variance of item log price growth factors, and ρ the product group covariance between item log price growth factors and item log sales shares in the base period. Only annual price growth factors are included. We only include group-year observations containing at least 4 items.

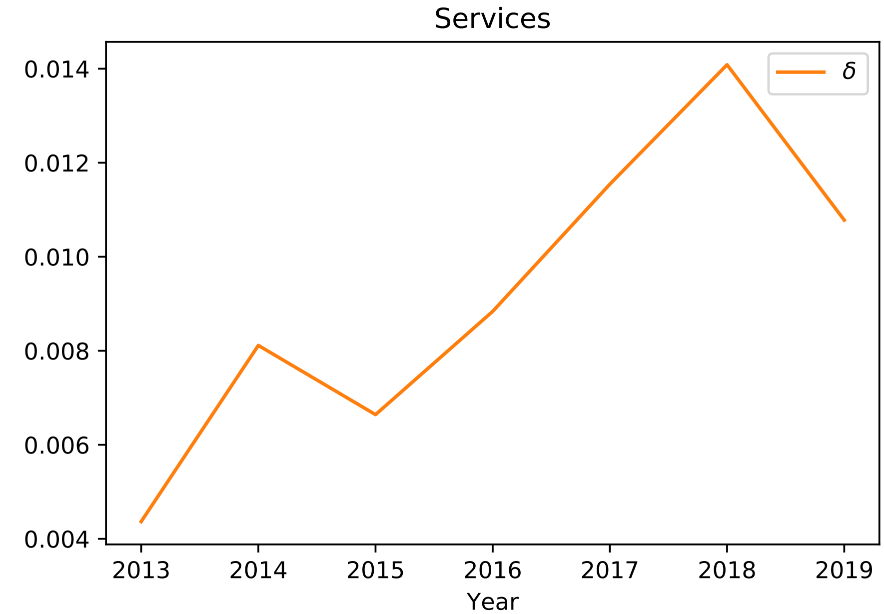

Figure 24: Average moments over time for goods. μ denotes the product group mean of item log price growth factors, δ the product group variance of item log price growth factors, and ρ the product group covariance between item log price growth factors and item log sales shares in the base period. Only annual price growth factors are included. We only include group-year observations containing at least 4 items.

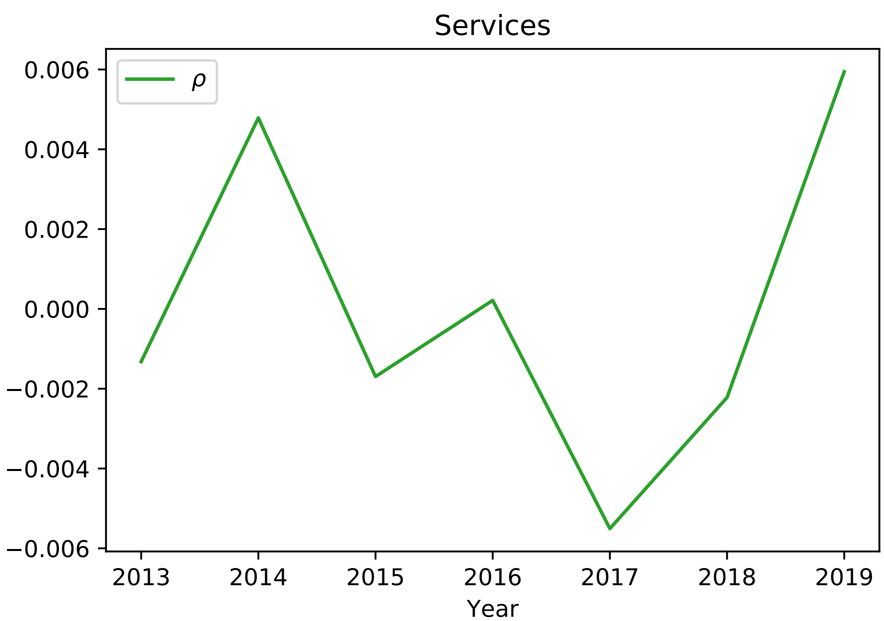

Figure 25: Average moments over time for goods. μ denotes the product group mean of item log price growth factors, δ the product group variance of item log price growth factors, and ρ the product group covariance between item log price growth factors and item log sales shares in the base period. Only annual price growth factors are included. We only include group-year observations containing at least 4 items.

A.4.1 Alternative specifications

We here show the distribution of moments when including all group-year observations, and not only when keeping the group-years containing at least 4 items.

Goods (2004–2019)

	25^{th}	50^{th}	75^{th}	Mean
μ	-0.010	0.015	0.048	0.020
δ	0.000	0.002	0.008	0.009
ρ	-0.003	0.000	0.002	-0.002

Services (2013–2019)

	25^{th}	50^{th}	75^{th}	Mean
μ	0.000	0.016	0.031	0.017
δ	0.001	0.004	0.012	0.010
ρ	-0.005	0.000	0.005	0.001

Table 28: Percentiles and means of moments for log-normal approximation, for goods (left panel) and services (right panel), respectively. μ denotes the product group mean of item log price growth factors, δ the product group variance of item log price growth factors, and ρ the product group covariance between item log price growth factors and item log sales shares in the base period. Only annual price growth factors are included.

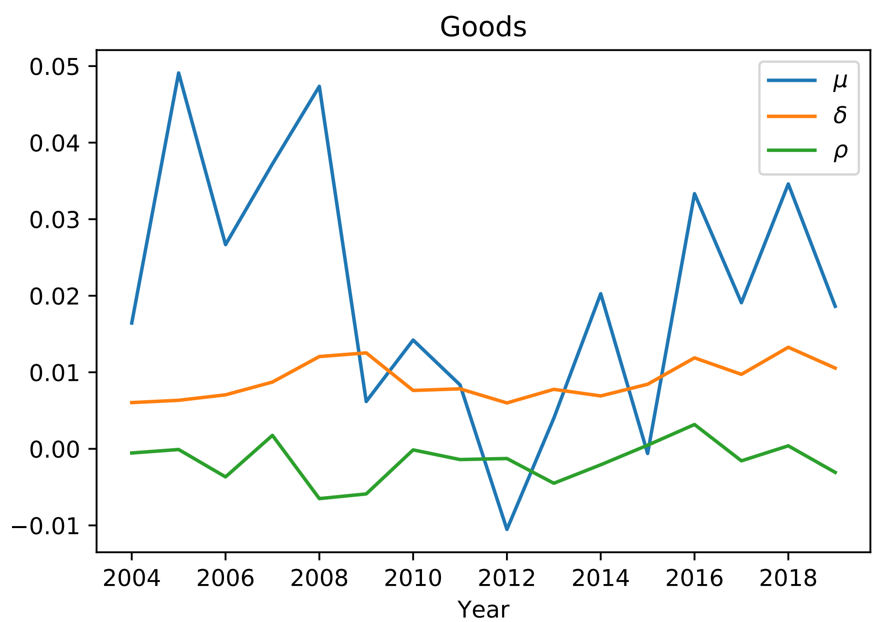

Figure 26: Average moments over time for goods. μ denotes the product group mean of item log price growth factors, δ the product group variance of item log price growth factors, and ρ the product group covariance between item log price growth factors and item log sales shares in the base period. Only annual price growth factors are included.

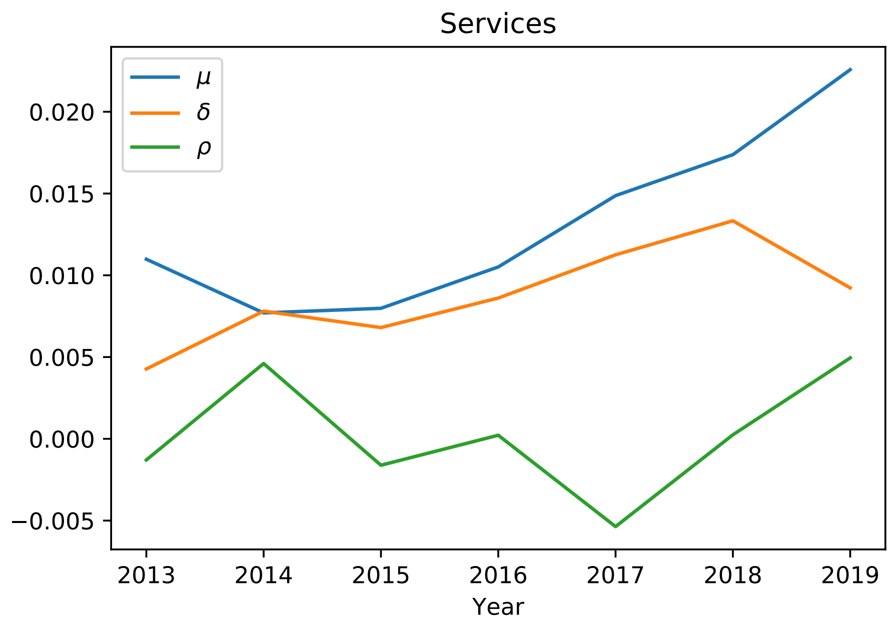

Figure 27: Average moments over time for services. μ denotes the product group mean of item log price growth factors, δ the product group variance of item log price growth factors, and ρ the product group covariance between item log price growth factors and item log sales shares in the base period. Only annual price growth factors are included.

B Additional data descriptions

B.1 PPI micro data in national product accounts

This section describes how Statistics Sweden measures real output and GDP for sectors not directly captured in the PPI micro data. The description relies on official documentation (e.g. SOU (2002)) and e-mail correspondence with statisticians working with the product accounts at Statistics Sweden. In general, constructing real production for various sectors of the economy is based on different sources, including the Consumer Price Index (CPI), the Real Estate Price Index, as well as production surveys (SOU, 2002). For example, production in some public services are measured by quantities directly, e.g., by the number of surgeries in the health care sector, or teaching hours in the education sector.

The PPI micro data include price information on domestic and foreign production of, for example, agriculture, materials, energy, manufacturing, transportation, information and communication, real estate, as well as professional and scientific services such as legal and management services, and personal services provided by e.g. restaurants and hairdressers. However, there are a couple of sectors that are not covered. These include most public services, i.e. health care, education, national defence, and social insurance. Also, price information for the construction sector and partly the finance sector is not collected. Although prices in these sectors are not directly measured, the PPI micro data are still used for measuring their real output and value added for the national product accounts. For example, to measure real production in public services, Statistics Sweden deflates input costs or measure real quantities of inputs directly, e.g. by hours worked. When deflating input costs, the PPI micro data is used. Although most of the finance sector is not covered in the PPI micro data, prices of one product group (SPIN5 code 64.190, which includes banking services except central banking) are measured. These prices are used to deflate the production of several product groups in the finance sector. For other product groups, deflators are constructed using e.g. wage indices for the finance industry and the GDP-deflator.

Other parts of the private sector that are not directly price measured in the PPI micro data are the wholesale and retail trade sectors. However, although there is formally no price information for product groups with codes (SPIN5) belonging to these sectors, the PPI micro data is used to deflate their output and value added. When measuring production in wholesale and retail trade, Statistics Sweden first divides these sectors into units that produces goods or services in the traditional sense, i.e. by adding value to the product before being sold, and into units that are pure traders, i.e. who sell the same goods and services that they buy in. Traditional production is deflated using prices in the PPI and CPI micro data for product groups related to the production in question. The sales of pure traders are deflated using more aggregate indices, which are also based on prices from both the PPI and CPI surveys.

In conclusion, measuring real output is an intricate process involving various data sources. Although deflators are constructed also using surveys on quantities and other price surveys such as the CPI, the PPI micro data is a fundamental part of measuring real output, as it is used to deflate nominal output in virtually all sectors of the economy.

B.2 Goods PPI micro data

The system for collecting producer prices has changed at least three times in 1992–2019. One such change resulted in that all item IDs was changed in January 2013. This means that we can not directly link items in January 2013 to items in December 2012 by merely comparing item IDs between these two months. To overcome this issue, we use a procedure to link the same items between December 2012 and January 2013 by matching an item's base price in January 2013 to an item's reported price (adjusted to SEK if needed) in December 2012. If these are identical for two item IDs within the same firm, we assume that they represent the same item. This procedure is identical to how we identify items that do not undergo quality or quantity adjustments in other periods, explained in Section 3.2. Hence we can only match item IDs for cases where the item did not undergo a product substitution. With this method, we match 65 % of the items, i.e. unweighted sample, in January 2013 to the items in December 2012, which can be compared to the average share of continuing non-substituted items between December and January for the rest of the years, 1992–2019 excluding 2013, which is 81 %. Although it is possible that the share of continuing items in January was lower in 2013 than other years due to the major changes made at the time, it is presumably the case that our method drops too many items. For the price change frequencies and sizes, which includes year 2013, it turns out that our results are robust to excluding 2013. For example, Figure 28 shows that 2013 does not stand out as compared to the monthly mean and median weighted price change frequencies in other years.

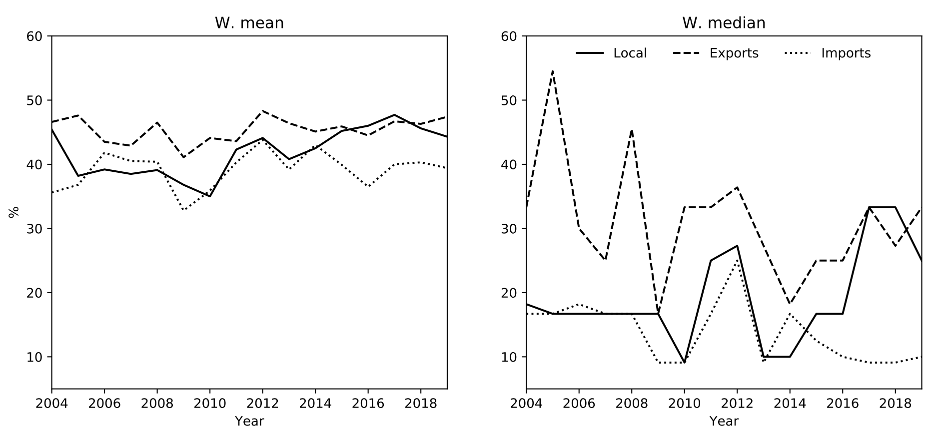

Figure 28: Weighted mean and median of monthly price change frequencies in each year between 2004–2019, for goods sold and produced domestically (*Local*), exports and imports, respectively.

The approach described above cannot distinguish between newly introduced items and existing items subject to quality change as measured by a change in the base price. For this reason, we drop the year 2013 in Figure 18, reducing the sample to 523,331. However, year 2013 is included in the rest of the analysis in the main text.

Another change in the data is the adding of a variable that flags when a new item is introduced due to rotation of the panel. This enables an additional split of new items into those that are rotated and not rotated. Figure 29 reports the shasre of different item types in 2015–2019 when the flag for rotations is available.

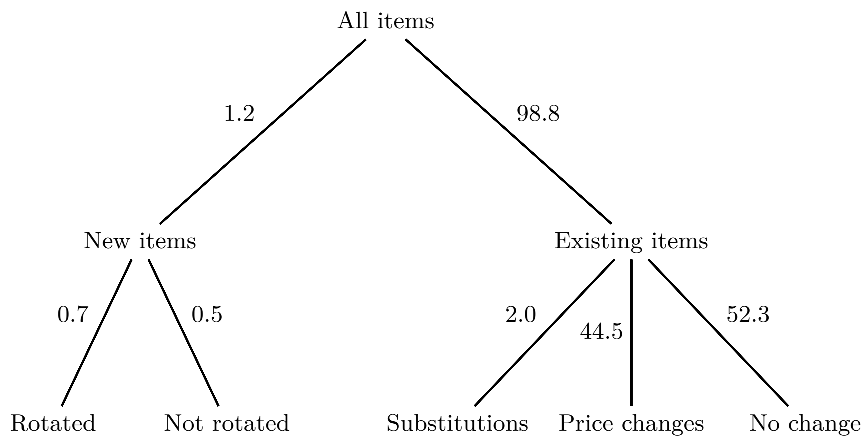

Figure 29: Shares of item adjustments on a monthly frequency in the goods PPI micro data in 2015–2019. The number on the edge above each node shows the weighted share of item-month observations of that type in relation to all item-month observations. *Substitutions* are non-comparable product substitutions that experience a quantity or quality adjustment in prices, *Price changes* are price changes that occur for the same item without a simultaneous quantity or quality change.

Finally, it should be noted that sales in a year a in the PPI micro data are not necessarily equal to actual sales due to the time it takes to collect firm accounts data. In the end of year a, Statistics Sweden only has information on last year's $(a-1)$ firms sales. This is used to construct sales for next year's $(a+1)$ PPI by inflating the sales in $a-1$ with narrowly defined price indices, often at the item level. So all sales observed in the micro data are sales two years back, inflated to better represent current nominal values.

B.3 Services PPI micro data

Figure 30 reports the share of different types of items for 2015–2019, when the flag for rotations is available.

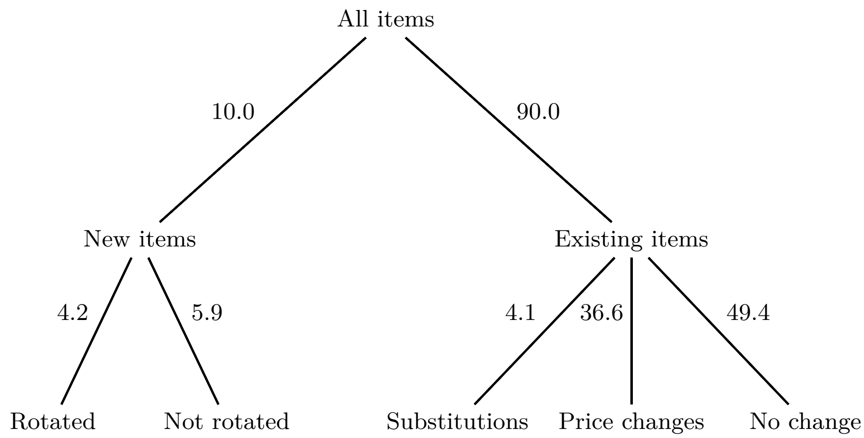

Figure 30: Shares of item adjustments on a quarterly frequency in the services PPI micro data in 2015–2019. See Figure 29 for a detailed description.

B.4 Firm reporting to the PPI micro data

Most firms reports their price information to Statistics Sweden by filling an electronic form, but they can also report by e-mail or mail. Respondents that fill in the electronic form first go to the PPI webpage where they can log in to report their prices.[17] Having logged in, they fill in the electronic form shown in Figure 31. A translated version in english is provided in Figure 32.

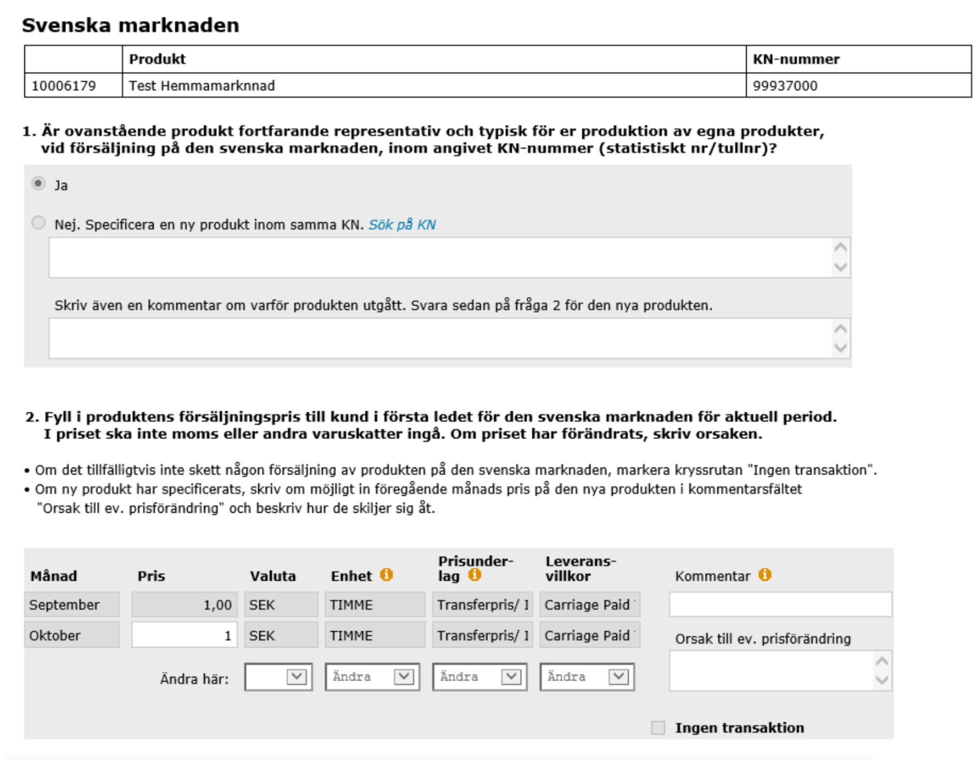

Svenska marknaden

	Produkt	KN-nummer
10006179	Test Hemmamarknnad	99937000

1. Är ovanstående produkt fortfarande representativ och typisk för er produktion av egna produkter, vid försäljning på den svenska marknaden, inom angivet KN-nummer (statistiskt nr/tullnr)?

Ja

Nej. Specificera en ny produkt inom samma KN. Sök på KN

Skriv även en kommentar om varför produkten utgått. Svara sedan på fråga 2 för den nya produkten.

2. Fyll i produktens försäljningspris till kund i första ledet för den svenska marknaden för aktuell period. I priset ska inte moms eller andra varuskatter ingå. Om priset har förändrats, skriv orsaken.

- Om det tillfälligtvis inte skett någon försäljning av produkten på den svenska marknaden, markera kryssrutan "Ingen transaktion".
- Om ny produkt har specificerats, skriv om möjligt in föregående månads pris på den nya produkten i kommentarsfältet "Orsak till ev. prisförändring" och beskriv hur de skiljer sig åt.

Månad	Pris	Valuta	Enhet	Prisunder-lag	Leverans-villkor	Kommentar
September	1,00	SEK	TIMME	Transferpris/ I	Carriage Paid	
Oktober	1	SEK	TIMME	Transferpris/ I	Carriage Paid	Orsak till ev. prisförändring
Ändra här:			Ändra	Ändra	Ändra	

Ingen transaktion

Figure 31: Electronic form filled in by respondents to the domestically sold goods PPI.

[17] https://www.scb.se/lamna-uppgifter/undersokningar/Prisindex-i-producent--och-importled/

Item ID	Product description	Product code (CN)
10006179	Test product	99937000

1. Is the product representative for the firm's own production sold in Sweden within the CN code?

☑ Yes
☐ No

- Specify a new product within the same CN code:
- Comment on why the earlier product is no longer available:

2. Provide the transaction price for the given period. The price should not include VAT or other taxes. If the price has changed, then also provide the reason behind the price change.

- If the product was temporarily not sold, mark the "No transaction" box.
- If a new product was specified, also provide, if possible, its price in the previous month under "Reason for price change" and describe how the product differs from the previous one.

Month	Price	Currency	Unit	Terms of delivery
September	1.00	SEK	HOURS	Carriage paid
October		SEK	HOURS	Carriage paid

- Comment:
- Reason for price change:

☐ No transaction

Figure 32: Electronic form filled in by respondents to the domestically sold goods PPI. Translated from the original form in Swedish.

B.5 Construction of sales in the PPI micro data

B.5.1 Item sales

This section provides a description of how sales are constructed in the goods and services PPI micro data. It relies on official documentation, such as Statistics Sweden (2019*b*), as well as e-mail correspondence with statisticians working with the PPI at Statistics Sweden. Assigning sales is an intricate process relying on information from different data sources. We therefore also briefly describe the data sets used for PPI sales construction in subsequent sections.

As mentioned in the main text, the unit of observation in the PPI micro data is an item, which is on a more detailed level than the 8-digit product code used in the goods data (CN8) and than the 7-digit product group code used in the services data (SPIN7). In particular, for a given item i, we observe its product code k, the firm f that produced the item as well as which 5-digit product group (*Standard för svensk produktindelning efter näringsgren*, SPIN5) g it belongs to. Each item is assigned a value representing its size, which is later used for aggregating item price changes into an aggregate index. Importantly, the size of this value is derived from an algorithm that depends on the item's sales share in relation to similar products, but is not always equal to true sales, which we denote as x_i. We refer to this value used in the PPI as the item *sales*. We observe item sales s_i as well as sales of the product group $\sum_{i \in g} s_i$. A *stratum* is a collection of products, potentially defined on a higher level of aggregation than a product group. Hence a stratum h is either represented by one or more product groups. One product group g is always assigned to one stratum h, so the same product group does not show up in several strata.

To construct item sales s, actual sales x are in a first step assigned to a firm-product *pair*, defined as a combination of a firm and products as defined by the CN8 code: $j = \{f, k\}$. This relies on combining information from rich data sources, such as the Industrial production survey (IVP) for goods and registry data on firm income statements and balance sheets (FEK) for services. To allocate sales to transactions occurring at the domestic market and as exports, surveys on foreign trade och goods and services (UHV and UHT) are also used.

After a measure of sales s has been assigned to each pair, Statistics Sweden then divides pair sales across the items in the same pair. The division across items either make use of additional information on the production or sales structure of the firm-product in question or, in lack of such information, splits pair sales evenly across items. We now focus on how item sales s are assigned in the first step, i.e. we focus on the algorithm that Statistics Sweden uses to assign sales to each pair in a stratum.

For a given stratum h, Statistics Sweden observes sales for each pair, x_j, where $j \in h$, from other data sources, such as IVP and UHV. Let N_h denote the number of pairs that are observed. Neyman allocation is then used to decide how many pairs, n_h, for which prices should be collected in the PPI micro data, where $n_h \leq N_h$.

Assigning sales s_j to each pair $j = \{1, 2, ..., n_h\}$ then proceeds as follows. First order all pairs in the stratum in terms of actual sales

$$x_1 \geq x_2 \geq ... \geq x_{N_h}.$$

Then use the following steps:

1. If

$$x_1 \geq \frac{x_1 + x_2 + ... + x_{N_h}}{n^h}$$

then let sales for the first pair represent its true sales, i.e. $s_1 = x_1$. Then proceed to the next step. If the above condition is not met, then instead let the sales of all pairs represent the sample average of total sales in the stratum:

$$s_j = \frac{x_1 + x_2 + ... + x_{N_h}}{n^h}, \quad \forall j \in \{1, 2, ..., n_h\}$$

and stop the algorithm.

2. If

$$x_2 \geq \frac{x_2 + ... + x_{N_s}}{n^h - 1}$$

then let sales for the second pair represent its true sales, i.e. $s_2 = x_2$. Then proceed to the next step. If the above condition is not met, then instead let the sales of all pairs from pair 2 represent the sample average of total sales in the stratum, less sales of pair 1:

$$s_j = \frac{x_2 + ... + x_{N_h}}{n^h - 1}, \quad \forall j \in \{2, ..., n_h\}$$

and stop the algorithm.

$\vdots$

Proceed with the above algorithm until the stopping rule applies.

B.5.2 Inflation of sales

Statistics Sweden creates the sample for the PPI in each year a at the end of year $a - 1$. At that point in time, the latest publications for the underlying data sources for sales are for year $a - 2$. So the weights in PPI in year a are constructed using sales in year $a - 2$. These weights are then inflated by price growth factors on the product level to capture the nominal changes in prices between $a - 2$ and a (Δ_P), so what is observed in the PPI micro data is actually an inflated measure of the weights: $\tilde{w} = \Delta_P \times w$.

B.6 Data sources for constructing sales

The main data sources on quantities and sales used to aggregate production are the industrial goods production survey (Industrins varuproduktion, IVP) and surveys and registry data on firm income statements and balance sheets (Företagens Ekonomi, FEK), as well as surveys on foreign trade of goods (Utrikeshandel med varor, UHV) and services (Utrikeshandel med tjänster, UHT).

The main data sources used to construct sales in the goods PPI micro data are IVP and UHV. We have access to the micro data of these surveys and use them directly to estimate elasticities of substitution across goods, as explained in the section below.

The firm income statements and balance sheets data, FEK, contains firm-level information on sales for all private firms in Sweden.[18] As a supplement to this dataset, Statistics Sweden also conducts a survey covering about 500 of the biggest firms as well as a sample of about 15,000 other firms. In this survey, firms are asked to specify the sources of sales. Thus, the survey is used to get sales on a firm-product level, where product is defined as a detailed 7-digit product group, SPIN7. We do not access this survey and therefore use the sales observed in the services PPI micro data to estimate elasticities of substitution. This is further described in the section below.

[18] FEK does not include the financial sector. In the services PPI micro data, sales for this sector are computed by using other sources.

C Estimating elasticities of substitution

This section lays out the empirical strategy for estimating elasticities of substitution between products within product groups, as well as the elasticity of substitution between product groups. First, we describe the estimation strategy. Second, we describe the data sets used and present results for goods ans services, respectively.

C.1 Estimation

C.1.1 Lower level of aggregation

To obtain estimates for the elasticities of substitution at the the product group level, we emulate the approach described in Feenstra (1994), Broda and Weinstein (2006), and Hottman et al. (2016). In particular, we adopt the same strategy as Hottman et al. (2016) for Universal Product Codes (UPCs). Broadly speaking, this approach addresses the simultaneity bias by differencing equations of supply and demand with respect to time and a reference product in order to exclude shocks affecting both price and quantity, such as changes in taste preferences. An important assumption is that these shocks are specific to product groups, and not to individual items within product groups.

Starting with demand, rewrite Equation (3) by taking logs and differencing across years a and $a-1$.

$$w_{i,g,a} = \left(\frac{p_{i,g,a}}{P_{g,a}}\right)^{1-\sigma_g} \gamma_{i,g,a} \tag{12}$$

$$\ln w_{i,g,a} = (1-\sigma_g)\Big[\ln p_{i,g,a} - \ln P_{g,a}\Big] + \ln \gamma_{i,g,a} \tag{13}$$

$$\Delta \ln w_{i,g,a} = (1-\sigma_g)\Big[\Delta \ln p_{i,g,a} - \Delta \ln P_{g,a}\Big] + \Delta \ln \gamma_{i,g,a}. \tag{14}$$

Next, assume supply to take the following form

$$p_{i,g,a} = \exp(\nu_{i,g,a}) y_{i,g,a}^{\omega_g}, \tag{15}$$

where $\nu_{i,g,a}$ represents a technology term and $\omega_g \geq 0$ is an inverse supply elasticity assumed to be constant across items in the same product group and over time. Proceed by writing the time-differenced supply curve in logs, and substitute quantities, $y_{i,g,a}$ for sales shares $w_{i,g,a}$ as follows:

$$\Delta \ln p_{i,g,a} = \omega_g \Delta \ln y_{i,g,a} + \Delta \nu_{i,g,a}, \tag{16}$$

$$y_{i,g,a} := w_{i,g,a} S_{g,a} / p_{i,g,a} \tag{17}$$

$$\Delta \ln p_{i,g,a} = \omega_g \Delta \ln(w_{i,g,a} S_{g,a} / p_{i,g,a}) + \Delta \nu_{i,g,a} \tag{18}$$

$$\Delta \ln p_{i,g,a} = \omega_g \Big[\Delta \ln(w_{i,g,a}) + \Delta \ln(S_{g,a}) - \Delta \ln(p_{i,g,a}) \Big] + \Delta \nu_{i,g,a} \tag{19}$$

Where $S_{g,a}$ denotes total sales of products in group g in year a. As an intermediary step to reduce noise, we constrain the sample to the middle 99 percentiles of the variables $\Delta \ln w_{i,g,a}$ and $\Delta \ln p_{i,g,a}$.

While Feenstra (1994) and Broda and Weinstein (2006) use data on products sold in different countries, Hottman et al. (2016) exploits variation across products sold by the same firm. We instead use accounts of national production and estimate the elasticities across different products within the same product group. We define a base product within each group g in year a as a product that exists both in year a and the previous year $a-1$, and has the largest sum of sales over these two years. Consequently, the base product is allowed to change from one year to the next.

We then proceed by differencing Equations (14) and (19) with respect to the identified base product, denoted by $\underline{i}$, thereby eliminating product group variation, in addition to changes due to time.

$$\Delta^{\underline{i},a} \ln w_{i,g,a} = (1 - \sigma_g) \Delta^{\underline{i},a} \ln p_{i,g,a} + \varepsilon_{i,g,a} \tag{20}$$

$$\Delta^{\underline{i},a} \ln p_{i,g,a} = \frac{\omega_g}{1 + \omega_g} \Delta^{\underline{i},a} \ln w_{i,g,a} + \delta_{i,g,a} \tag{21}$$

Where $\varepsilon_{i,g,a} = \Delta^{\underline{i},a} \ln \gamma_{i,g,a}$, $\delta_{i,g,a} = \frac{1}{1+\omega_g} \Delta^{\underline{i},a} \nu_{i,g,a}$, and $\Delta^{\underline{i},a} \ln x_{i,a} = \Delta \ln x_{i,a} - \Delta \ln x_{\underline{i},a}$ for any variable x. The base products are excluded from the sample when performing the estimation.

The identifying assumption is that the error terms from the double-differenced expressions for demand (20) and supply (21) equations are independent, $\mathbb{E}(\varepsilon_{i,g,a} \delta_{i,g,a}) = 0$. If independence holds, we can combine (20) and (21) by multiplying the respective residuals and divide by $(\sigma_g - 1)$ to form:

$$
\begin{aligned}
Y_{i,g,a} &= \theta_{1,g} X_{1,i,g,a} + \theta_{2,g} X_{2,i,g,a} + u_{i,g,a}, \text{ where:} && (22)\\
Y_{i,g,a} &= (\Delta^{\underline{i},a} \ln p_{i,g,a})^2 && (23)\\
X_{1,i,g,a} &= (\Delta^{\underline{i},a} \ln w_{i,g,a})^2 && (24)\\
X_{2,i,g,a} &= \Delta^{\underline{i},a} \ln p_{i,g,a} \Delta^{\underline{i},a} \ln w_{i,g,a} && (25)\\
\theta_{1,g} &= \frac{\omega_g}{(1+\omega_g)(\sigma_g - 1)}, && (26)\\
\theta_{2,g} &= \frac{\omega_g(\sigma_g - 2) - 1}{(1+\omega_g)(\sigma_g - 1)}, && (27)\\
u_{i,g,a} &= \varepsilon_{i,g,a} \delta_{i,g,a} / (\sigma_g - 1). && (28)
\end{aligned}
$$

Note that this system requires that σ_g is not equal to 1, but allows for other positive values and, in particular, values arbitrarily close to 1. Although theory restricts σ_g and ω_g to be non-negative, we do not make such a restriction when performing the estimation.

As pointed out in Feenstra (1994), the term $u_{i,g,a}$ is correlated with the regressands $X_{1,i,g,a}, X_{2,i,g,a}$ because prices and sales shares are correlated with the errors $\varepsilon_{i,g,a}$ and $\delta_{i,g,a}$. However, with panel data, this can be overcome by assuming that demand and supply disturbances are constant over varieties of the same product. Hence we average over time and obtain

$$\bar{Y}_{i,g} = \theta_{1,g} \bar{X}_{1,i,g} + \theta_{2,g} \bar{X}_{2,i,g} + \bar{u}_{i,g}, \tag{29}$$

which we use to define moment conditions resting on the assumption of independence of the unobserved demand and supply disturbances for each other product in the same group over time. Thus, the moment condition for each product i in group g is

$$\mathcal{G}(\beta_g) = \mathbb{E}_{\mathbb{A}}(u_{i,g,a}(\beta_g)) = 0 \tag{30}$$

where $\beta_g = \begin{pmatrix} \sigma_g \\ \omega_g \end{pmatrix}$ and $\mathbb{E}_{\mathbb{A}}$ is the expectation operator over time. We stack all moment conditions to form the GMM objective function and obtain

$$\hat{\beta}_g = \arg\min_{\beta_g} \mathcal{G}^*(\beta_g)' \mathcal{W} \mathcal{G}^*(\beta_g) \tag{31}$$

where $\mathcal{G}^*(\beta_g)$ is the sample analog of $\mathcal{G}(\beta_g)$ and $\mathcal{W}$ is a positive definite weighting matrix. We weight each product i by its average sales share within product group g, implying that double-differenced products that have larger average sales shares over time get a higher weight. In order to obtain estimates for σ_g from these equations, we use the non-linear Levenberg Marquardt algorithm on product group subsets of the data. For some product groups, this estimation procedure results in imprecise point estimates, e.g. due to few observations in a product group. If that is the case, we instead estimate the elasticity across products within a more aggregated grouping. Our procedure to estimate parameters using data on different levels of aggregation is done as follows. First, we estimate the elasticity of substitution across products using the entire data set to obtain a reference value of the elasticity of substitution at the highest level, σ^0. This value is the elasticity of substitution between all products, i.e. assuming that $\sigma_g = \sigma^0 \; \forall g$. Note however that the double differenced variables are still differenced with respect to a product belonging to its respective product group and that the base product varies over the product groups. The estimated σ^0 is then used as a starting value when running the same regression for each SPIN1 group using products within the respective SPIN1 group. In this step, elasticities are assumed to be constant for all product groups that belong to the same SPIN1. The process is then repeated, but with starting values stemming from the SPIN1 estimated value, if it is defined, to estimate at the SPIN2 level, and so on until the regression has been run once for each product group, assuming heterogeneous elasticities across all groups at the SPIN5 level for goods and the SPIN7 level for services.

Inevitably, not all estimates at any given level, such as SPIN5, will be successful, in the sense that the point estimates are imprecise. We characterize a failed estimation as one when the point estimate of the elasticity of substitution is smaller than its standard error. To obtain estimated values for product groups where the estimation failed, we impute the value with the elasticity estimated obtained when estimating elasticities for the higher aggregated groups, starting with SPIN4, or if this value was also unsuccessful, using the SPIN3 value, and so on until only σ^0 remains as a feasible candidate. With this approach, all product groups end up with an estimated value for σ_g. Table 8 shows a summary over the origin of estimated σ_g.

C.1.2 Upper level of aggregation

The strategy for estimating the upper level elasticity across product groups corresponds to the approach used for estimation on the firm level in Hottman et al. (2016). The estimates of the product elasticities of substitution $\{\sigma_g\}$ from the section above are used to solve for demand shocks at the product level, $\{\gamma_{i,g,a}\}$, and compute the firm price indices $P_{g,a}$, using Equation (3) and the price index

$$P_{g,a} = \left(\sum_{i \in I_{g,a}} \gamma_{i,g,a} p_{i,g,a}^{1-\sigma_g} \right)^{\frac{1}{1-\sigma_g}} . \tag{32}$$

The sales share of product group g in relation to all other products is given by

$$W_{g,a} = \gamma_{g,a} \left(\frac{P_{g,a}}{P_a} \right)^{1-\sigma} . \tag{33}$$

where P_a is the aggregate price index in year a

$$P_a = \left(\sum_{g \in G} \gamma_{g,a} P_{g,a}^{1-\sigma} \right)^{\frac{1}{1-\sigma}} \tag{34}$$

Taking logs and double-differencing (33) across years and in relation to a reference product group, $\underline{g}$, yields

$$\Delta^{\underline{g},a} \ln W_{g,a} = (1-\sigma) \Delta^{\underline{g},a} \ln P_{g,a} + \kappa_{g,a} \tag{35}$$

where $\kappa_{g,a} = \Delta^{\underline{g},a} \ln \gamma_{g,a}$. As pointed out in Hottman et al. (2016), estimating equation (35) using OLS is potentially problematic because changes in product group price indices could be correlated with changes in product group demand shocks. To find a suitable instrument for changes in product group price indices, we use the structure of the model to write changes in product group price indices in terms of the underlying product characteristics of the product group. To this end, denote the geometric mean of a variable by a tilde above that variable, i.e. for a given variable x,

$$\tilde{x}_{g,a} = \exp \left\{ \frac{1}{|I_{g,a}|} \sum_{i \in I_{g,a}} \ln x_{i,g,a} \right\} \tag{36}$$

where $I_{g,a}$ is the set of products in product group g in year a, and $|I_{g,a}|$ is its cardinality. Using the optimal sales share for a product i in product group g, as given by equation (3), then yields

$$\frac{w_{i,g,a}}{\tilde{w}_{g,a}} = \gamma_{i,g,a} \left(\frac{p_{i,g,a}}{\tilde{p}_{g,a}} \right)^{1-\sigma_g} \tag{37}$$

where the geometric mean of product demand shocks within a product group is assumed to equal 1: $\tilde{\gamma}_{g,a} = 1$.

Substituting the implied expression of product demand shocks $\gamma_{i,g,a}$ from Equation (37) into the product group price index in Equation (32), taking logs and double-differencing with respect to time and the reference group yields

$$\Delta^{\underline{g},a} \ln P_{g,a} = \Delta^{\underline{g},a} \ln \tilde{p}_{g,a} + z_{g,a} \tag{38}$$

where the last term is given by

$$z_{g,a} = \frac{1}{1-\sigma_g} \Delta \ln \left(\sum_{i \in I_{g,a}} \frac{w_{i,g,a}}{\tilde{w}_{g,a}} \right) - \frac{1}{1-\sigma_{\underline{g}}} \Delta \ln \left(\sum_{i \in I_{\underline{g},a}} \frac{w_{i,\underline{g},a}}{\tilde{w}_{\underline{g},a}} \right). \tag{39}$$

Equation (38) is used as a first stage, instrumenting for the double-differenced product group price index in Equation (35), which is used in the second stage.

C.2 Data

C.2.1 IVP micro data

The data used for estimating the elasticities of subsitution across goods is the Swedish industrial goods production survey, IVP, an encompassing database containing information on yearly production value and quantities of goods reported by a sample of Swedish production-units, active last year in the tax records. We have access to this data set between the years 1996 and 2017. In total, we observe 292,244 observations where an observation is a good reported yearly by a *production unit*. A production unit corresponds to a firm if the firm is small and a subdivision of a firm if the firm is large. The unit at which goods are reported follows a variant of the combined nomenclature (CN) down to an eight or nine digit level. While Statistics Sweden points to potential reporting errors of the exact CN value, this is believed to be a negligible issue for our purposes; the CN classification is used merely to differentiate between goods and for identifying more aggregated product groups, at the SPIN5 level, which are more likely to be reported correctly due to the aggregation. Finally, the time variable is year, which indicates the fiscal year of production. Fiscal years that do not correspond to the calendar years are reported as the later year if the fiscal year ends after the first of May, otherwise it is reported as the previous year. While good, production unit, and year jointly identify an observation in most cases, there exists observations where the production unit reports different values for a given good in a given year. These cases are removed from the data set to ensure the integrity of the data's primary key and to accurately track observations over time.

While the exact sampling of firms has been subject to revisions over the years, the sampling method used in 2019 samples all active firms registered in Företagsdatabasen having more than 20 employees or a turnover exceeding 75 million SEK. In addition, firms with less than 20 employees, but more than 10 employees were sampled for the industries, *Other extraction of minerals; Production of fertilizers and other nitrogen products; Production of cement, lime, and plaster; Production of concrete-, cement-, and plaster-based goods; Logging; Shaping and final processing of stone and manufacture of abrasives and other non-mineral goods.* Instead of sampling values for small firms that do not meet the cut-off criteria, production values are estimated using a model and then imputed. This works for the purpose of IVP, namely estimate overall industry production values, but means we have to restrict our sample to the larger firms for which survey data exists. Changes in the sampling method over time has mainly affected the sample status of smaller firms. To give some examples, between 1996 and 2002, all firms with 10 employees or more were sampled, and during the period 2003 to 2006, some firms were sampled only every other year to ease reporting burden.

An important feature of the IVP data is that we observe the production value and the respective quantity, which is reported for each good. There are two issues related to observing production values instead of sales values when estimating the elasticities of substitution. First, it is not clear that everything produced end up being sold so production might be higher than sales. Second, another consequence of observing production data instead of sales data is that we cannot immediately distinguish between domestic consumption and exports. Nevertheless, a measure of domestic consumption can be obtained by making use of trade data (UHV), which contains monthly information of imports and exports at the firm and CN8 level. We can therefore join the trade data with IVP to construct a fraction of production at the CN8 level of a firm that was exporting in a given year. There are some complications associated with this procedure since when the data in UHV is aggregated up to the yearly level, that value represents production value during the calendar year production that may or may not correspond to a given firm's fiscal year. While this cannot be addressed without further information, we constrain the estimate of domestic production to be non-negative and assume that goods for which there is no match in the trade data are only sold in the domestic market.

Before using the IVP data, some observations had to be dropped in several consecutive steps. First, the data is filtered by removing observations with imputed production values leaving a total of 244,296 observations. After this step, we remove the remaining 2,071 observations that are not uniquely identified by the primary key. We thereafter remove observations with missing production value or quantity, or that have negative values for these variables, leaving a total of 178,660 observations. A final filtering step is to remove observations where the unit of the quantity is different from its value last year, if a value was reported. This leaves us with 177,886 observations.

C.2.2 Services PPI micro data

The services PPI micro data contains quarterly price data on an item level between year 2013–2019. Restricting the sample to domestic sales and exports, and dropping observations that are imported indices and hence do not represent items but groups of items, it contains 99,977 item-quarter observations. Since the unit of observation in the services PPI micro data is an item-quarter, we aggregate the data in two dimensions before using it for estimation. First, as described in Section B.5, firms with large market shares in a given market-SPIN7 group can have their sales split up between several items, each with a distinct item price. If Statistics Sweden does not have information on item specific sales shares, they simply do an even split, leading to several items in the same firm-market-SPIN7 with the same reported sales. We summarize these items into one consolidated item with sales given by the sum of sales across items and with a price given by the arithmetic average across items. Specifically, if we observe a firm-market-SPIN7-quarter combination with several items given by the set $\tilde{I}$ with identical sales $s_i = \tilde{s}\ \forall i \in \tilde{I}$, we construct one consolidated item $\tilde{i}$ with sales given by $s_{\tilde{i}} = \sum_{i \in \tilde{I}} s_i$ and price given by $p_{\tilde{i}} = \frac{1}{|\tilde{I}|} \sum_{i \in \tilde{I}} p_i$, where $|\tilde{I}|$ is the number of items in $\tilde{I}$. After performing this aggregation, the number of item-quarter observations becomes 58,653.

Second, since sales only vary at an annual frequency, we aggregate the data to annual observations and let the annual price and sales be given by the arithmetic average across quarters

$$p_{i,g,a} = \frac{1}{|Q|} \sum_{q \in Q} p_{i,g,q,a} \tag{40}$$

$$s_{i,g,a} = \frac{1}{|Q|} \sum_{q \in Q} s_{i,g,q,a} \tag{41}$$

where Q is the set of quarters that item i exist in the data for year a. The new data set with observations on an item-year frequency contains 16,104 observations.

To control for quality and quantity changes, we use the ratio of observed prices and their base prices, which are determined by Statistics Sweden to account for quality and quantity adjustments. An alternative approach would be to identify items that experience substitution and drop them from the sample. In principle, both approaches could be used. However, we choose to use adjusted prices in our baseline estimation since we then do not need to drop observations from our sample. In the end of this section, we also report the results when using the alternative approach by dropping substitutions.

C.3 Results

C.3.1 Goods

Median and mean product elasticities within 2-digit product groups: Goods

2-digit id	Group name	Median	Mean	Min. level of est.	N. of product groups
01	Farming	3.26	3.26	SPIN5	24
02	Forestry	6.97	6.97	SPIN5	4
03	Fishing	6.00	6.00	SPIN1	1
05	Crude coal	6.00	6.00	SPIN1	2
06	Crude petr., natural gas	6.00	6.00	SPIN1	2
07	Metal ore	16.34	15.22	SPIN5	2
08	Misc. mineral extraction	3.73	3.42	SPIN5	6
10	Foods	4.65	8.98	SPIN5	30
11	Beverages	7.64	10.14	SPIN5	7
12	Tobacco	6.19	6.19	SPIN1	1
13	Textiles	11.16	8.86	SPIN5	11
14	Apparel	8.25	7.50	SPIN5	7
15	Hides, skins, leather	1.75	2.81	SPIN5	3
16	Lumber & wood	6.19	5.80	SPIN5	13
17	Paper	6.19	5.69	SPIN4	14
18	Printing services & recording	6.19	6.19	SPIN1	6
19	Coal & refined petr.	2.97	2.98	SPIN5	2
20	Chemicals	3.91	4.10	SPIN5	16
21	Pharmaceuticals	2.79	2.79	SPIN5	2
22	Rubber & plastic	9.09	7.81	SPIN5	6
23	Misc. non-metallic mineral	3.44	4.07	SPIN5	26
24	Metals	43.11	29.59	SPIN5	14
25	Misc. metal products	11.92	9.43	SPIN5	18
26	Computers, electronics, optics	10.95	11.91	SPIN5	10
27	Electric devices	3.47	5.22	SPIN5	10
28	Misc. machinery	4.36	3.74	SPIN5	21
29	Motor vehicles	5.61	11.82	SPIN5	5
30	Misc. transportation	5.27	5.22	SPIN5	8
31	Furniture	3.02	5.32	SPIN5	7
32	Misc. manufactured goods	2.08	2.06	SPIN5	10
33	Repair/install for machi. & dev.	3.02	3.02	SPIN1	10
35	Electr., gas, heating, cooling	4.38	4.41	SPIN5	5
36	Water	3.02	3.02	SPIN1	2
38	Waste disposal	3.01	3.01	SPIN5	6
39	Decontamination	3.02	3.02	SPIN1	1
58	Printing & publishing	5.50	4.91	SPIN5	4
59	Film/video/TV/audio services	12.51	11.73	SPIN5	2

Table 29: *N. of product groups* shows the number of unique 5-digit product groups within the given 2-digit group. *Min. level of est.* states the minimum level of aggregation when estimating elasticities for the 5-digit product groups within the given 2-digit group. For example, if the lowest level of aggregation is SPIN5, at least one elasticity within the 2-digit group was estimated only using products within its 5-digit product group. If the lowest level of aggregation is SPIN0, all elasticities were imputed from the estimation performed on all products in the sample.

Table 30 shows the estimated value for $1 - \hat{\sigma}$ from the two-stage least squares regression, implying an elasticity of substitution of $\hat{\sigma} = 1.237$ across 5-digit product groups.

Estimated elasticity across product groups: Goods

$1 - \hat{\sigma}$	-0.237^{***}
	(0.016)
Constant	-0.023
	(0.015)
Observations	4,504
R^2	0.049
Adjusted R^2	0.049
Residual Std. Error	1.028 (df = 4502)

Table 30: Estimated $1 - \sigma$ from two-stage least squares regression. The double-differenced group price index $\Delta^{g,a}_{-} \ln P_{g,a}$ is first regressed on its instruments, including the geometric average of prices in the group and a variant of a Theil index consisting of product sales shares relative to the geometric mean of sales shares in a given group. In the second stage, the fitted values are regressed on the product group sales shares $\Delta^{g,a}_{-} \ln W_{g,a}$. When performing an F-test on the instruments in the first stage, the null hypothesis that the instruments are weak is strongly rejected, with a p-value below 0.01. Asterisks denote the following p-values *p<0.1; **p<0.05; ***p<0.01

C.3.2 Services

Median and mean product elasticities within 2-digit product groups: Services

2-digit id	Group	Median	Mean	Min. level of est.	N. of product groups
49	Land & pipeline transportation	1.37	1.37	SPIN7	11
50	Water transportation	5.51	5.51	SPIN7	13
51	Air transportation	6.60	6.60	SPIN1	9
52	Warehousing	6.60	6.60	SPIN1	18
53	Courier & postal	30.46	21.72	SPIN7	8
55	Hotels & accommodation	3.73	3.87	SPIN7	4
56	Bars & restaurants	6.60	6.60	SPIN1	4
58	Publishing	6.60	6.60	SPIN1	31
59	Audio & video	6.60	6.60	SPIN1	13
60	Broadcasting (& scheduling)	2.56	2.56	SPIN0	4
61	Telecommunication	2.56	6.99	SPIN7	18
62	Computer programming	1.30	2.05	SPIN7	5
63	Information	2.56	2.56	SPIN0	4
64	Banking	1.34	1.40	SPIN7	2
68	Real estate	2.56	2.14	SPIN7	10
69	Legal & accounting	2.56	2.59	SPIN7	18
70	Business consultancy	4.16	4.22	SPIN7	11
71	Architectural & technical	7.94	16.92	SPIN7	30
73	Advertising & marketing	2.22	4.57	SPIN7	14
74	Misc. legal, business	2.22	2.22	SPIN1	12
77	Rental & leasing	1.78	1.77	SPIN3	24
78	Employment & staffing	2.22	2.22	SPIN1	7
79	Travel	2.22	2.22	SPIN1	4
80	Security	4.09	4.09	SPIN7	3
81	Property maintenance	2.03	4.85	SPIN7	7
82	Office & administration	3.99	6.05	SPIN7	9
93	Sports & recreational	27.04	27.04	SPIN1	1
95	Repair (electronics, instruments, etc.)	32.86	32.86	SPIN7	1
96	Misc. consumer services	27.04	27.04	SPIN1	3

Table 31: *N. of product groups* shows the number of unique 7-digit product groups within the given 2-digit group. *Min. level of est.* states the minimum level of aggregation when estimating elasticities for the 7-digit product groups within the given 2-digit group. For example, if the lowest level of aggregation is SPIN7, at least one elasticity within the 2-digit group was estimated only using products within its 7-digit product group. If the lowest level of aggregation is SPIN0, all elasticities were imputed from the estimation performed on all produscts in the sample.

Estimated elasticity across product groups: Services

$1 - \hat{\sigma}$	-0.071
	(0.072)
Constant	-0.142
	(0.087)
Observations	371
R^2	0.003
Adjusted R^2	-0.0001
Residual Std. Error	1.679 (df = 369)

Table 32: Estimated $1 - \sigma$ from two-stage least squares regression. The double-differenced group price index $\Delta^{\underline{g},a} \ln P_{g,a}$ is first regressed on its instruments, including the geometric average of prices in the group and a variant of a Theil index consisting of product sales shares relative to the geometric mean of sales shares in a given group. In the second stage, the fitted values are regressed on the product group sales shares $\Delta^{\underline{g},a} \ln W_{g,a}$. When performing an F-test on the instruments in the first stage, the null hypothesis that the instruments are weak is strongly rejected, with a p-value below 0.01. Asterisks denote the following p-values *p<0.1; **p<0.05; ***p<0.01

C.3.3 Results for services when dropping substitutions

We here report the alternative approach of dealing with quality and quantity changes, instead of using the base prices constructed by Statistics Sweden. This approach instead uses observed prices and drop items that experience a substitution. If an item undergo a substitution in a quarter in a given year a, we remove all observations of that item in that year from the sample, resulting in two distinct spells of observations for that item. For producers with large market shares and identical item sales, we drop the consolidated item it at least one of the underlying items experience a substitution. Removing substitutions reduces the number of item-year observations to 15,292.

The distribution of elasticities is shown in the table below. Dropping substitutions leads to more extreme elasticities at the top and bottom of the distribution, and the median of 3.13 is below the median of 5.51 in the baseline estimation using quality-adjusted prices. However, the upper level estimation yields a higher estimate, since we get that $\hat{\sigma} = 1.26$ and is statistically different from the Cobb-Douglas case when $\sigma = 1$, which was the result under the baseline approach.

Distribution of estimated elasticities across products: Alternative approach (services)

Ranked Percentile	Estimate	Estimate (SPIN7 only)
99	1.12	1.14
95	2.06	1.34
90	2.14	1.38
75	2.22	2.22
50	3.13	4.81
25	5.94	8.50
10	17.56	17.71
5	24.85	38.37
1	55.15	52.64
N. estimates:	298	26

Table 33: The table shows percentiles of elasticities of substitution across products when using the alternative estimation approach that drops all substituted items. *Estimate* shows the percentiles using all elasticities, including those imputed from higher level groupings than the 7-digit product group. *Estimate (SPIN7 only)* excludes all imputations.

C.4 Bootstrapping the PPI and TPI with estimated elasticities

The main text, section 4.4, reports the PPI index with estimated elasticities of substitution. We show in this section that the uncertainty about the estimated elasticities of subsitution, σ_g and σ, are negligible relative to the difference between the baseline PPI index and the index with estimated elasticities. To obtain confidence bands for the PPI with estimated elasticities of substitution we bootstrap the index as follows. First, we draw a value for σ_g (for each product group) and σ from a normal distribution with the previously estimated mean and standard deviation. For σ_g this distribution is specific to each product group. Second, we compute the PPI index with the drawn elasticities. We repeat those steps 1000 times resulting in a total of 1000 indices. We then compute moments of the price index distribution. Figure 34 shows the results. The blue line indicates the baseline PPI index with elasticities equal to zero at both levels. The orange line shows the average index value with the corresponding 95% confidence band in dashed lines. The range of the confidence band is small compared to the difference between the baseline PPI index and the index with estimated elasticities. We conclude that uncertainty about the estimated elasticities of substitution matters less than differences in aggregation methods.

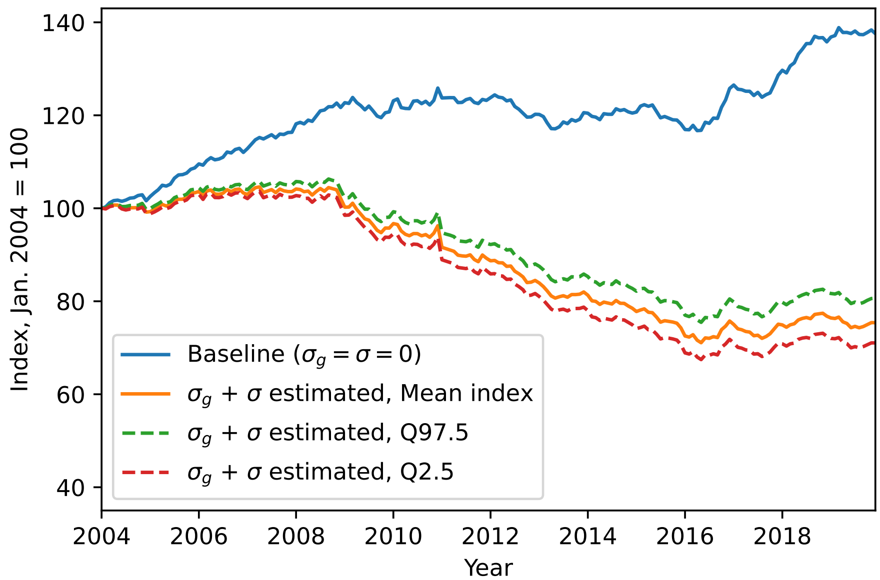

Figure 33: PPI index with bootstrapped elasticities of substitution at the lower and higher level, σ_g and σ.

We repeat the same exercise for the service price index, TPI. In line with the PPI, differences in price levels that arise from differences in aggregation methods exceed differences in price levels attributed to uncertainty about the estimated elasticities of substitution.

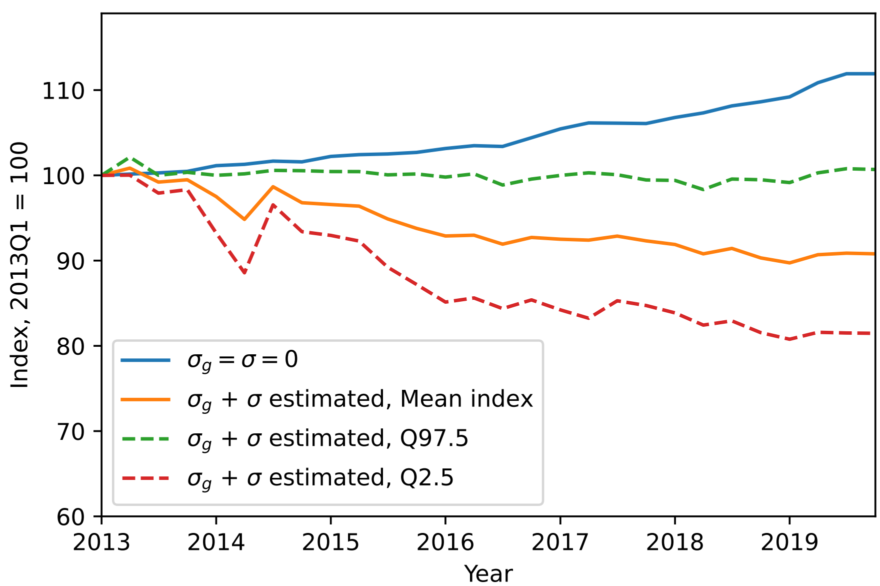

Figure 34: TPI index with bootstrapped elasticities of substitution at the lower and higher level, σ_g and σ.

D Proofs & derivations

D.1 Price index effect of a higher elasticity of substitution

To show that higher σ leads to (weakly) lower inflation, the following result needs to hold

$$\left(\sum_{i=1}^{N} w_i \left(\frac{p_{t,i}}{p_{0,i}}\right)^{1-\sigma_1}\right)^{\frac{1}{1-\sigma_1}} \geq \left(\sum_{i=1}^{N} w_i \left(\frac{p_{t,i}}{p_{0,i}}\right)^{1-\sigma_2}\right)^{\frac{1}{1-\sigma_2}}, \tag{42}$$

where $\sigma_1 < \sigma_2$. Consider three cases. First, with $\sigma_2 > 1$ the expression can be rearranged to

$$\left(\sum_{i=1}^{N} w_i \left(\frac{p_{t,i}}{p_{0,i}}\right)^{1-\sigma_1}\right)^{\frac{1-\sigma_2}{1-\sigma_1}} \leq \sum_{i=1}^{N} w_i \left(\frac{p_{t,i}}{p_{0,i}}\right)^{1-\sigma_2}. \tag{43}$$

Jensen's inequality states:

$$\varphi\left(\sum_i w_i x_i\right) \leq \sum_i w_i \varphi(x_i),$$

for a convex function φ, with the inequality reversed for a concave function. Setting $\varphi(\circ) = (\circ)^{\frac{1-\sigma_2}{1-\sigma_1}}$ and $x_i = \left(\frac{p_{t,i}}{p_{0,i}}\right)^{1-\sigma_1}$, Jensen's inequality turns into equation (43). Second, if instead $\sigma_2 < 1$, equation (42) can be rearranged equivalently such that Jensen's inequality for a concave function applies (with $\varphi(\circ) = (\circ)^{\frac{1-\sigma_2}{1-\sigma_1}}$ and $x_i = \left(\frac{p_{t,i}}{p_{0,i}}\right)^{1-\sigma_1}$). Third, if $\sigma_2 = 1$ equation (42) turns into

$$\left(\sum_{i=1}^{N} w_i \left(\frac{p_{t,i}}{p_{0,i}}\right)^{1-\sigma_1}\right)^{\frac{1}{1-\sigma_1}} \geq \prod_{i=1}^{N} \left(\frac{p_{t,i}}{p_{0,i}}\right)^{w_i}$$

Using Jensen's inequality for a concave function, setting $\varphi(\circ) = \frac{1}{1-\sigma} \ln(\circ)$ and $x_i = \left(\frac{p_{t,i}}{p_{0,i}}\right)^{1-\sigma_1}$ shows that the above inequality (after taking logs) holds.

The previous proof shows that higher σ leads to weakly lower inflation with respect to the base period $t = 0$. For the PPI this is December and for the TPI it is the 4th quarter of the previous year. It can occur though that for the example of the TPI the $Q2_t$ to $Q3_t$ growth rate is increasing in σ due to mean reversion of item prices within product groups. Consider the following example: assume that there are two items with a price of 100 in a product group in Q4_{t-1}. The price of the first good increases to 150 in Q1_t whereas the price of the second remains at 100. In Q2_t the price of the second good increases to 150 as well. If the price index for this group stands at 100 in Q4_{t-1} a price index with unweighted arithmetic averaging of price growth rates with respect to Q4_{t-1} stands at 125 in Q1_t $(0.5 \times 1.5 + 0.5 \times 1)$ and at 150 in Q2_t $(0.5 \times 1.5 + 0.5 \times 1.5)$. Instead a price index where both goods are perfect substitutes stands at 100 in Q1_t (just the cheaper good is consumed) and at 150 in Q2_t (both goods are equally expensive again). In the case of perfect substitutes price growth with respect to Q4_{t-1} is weakly lower, however the Q1_t to Q2_t growth rate is higher (150/100 compared to 150/125).

D.2 Proof for Proposition 1

Proof. Start from Equation (4):

$$\frac{P_{t,g}}{P_{0,g}} = \frac{\left(\sum_{i=1}^{N} \gamma_{0,i} p_{t,i}^{1-\sigma_g}\right)^{\frac{1}{1-\sigma_g}}}{\left(\sum_{i=1}^{N} \gamma_{0,i} p_{0,i}^{1-\sigma_g}\right)^{\frac{1}{1-\sigma_g}}} \tag{44}$$

where $N = |I_g|$ is the number of items. Taking logs yields

$$\ln\left(\frac{P_{t,g}}{P_{0,g}}\right) = \frac{1}{1-\sigma_g}\left(\ln\left(\sum_{i=1}^{N} \gamma_{0,i} p_{t,i}^{1-\sigma_g}\right) - \ln\left(\sum_{i=1}^{N} \gamma_{0,i} p_{0,i}^{1-\sigma_g}\right)\right)$$

and adding and subtracting $\frac{1}{1-\sigma_g}\ln\left(\frac{1}{N}\right)$ allows us to rewrite the expression as

$$\ln\left(\frac{P_{t,g}}{P_{0,g}}\right) = \frac{1}{1-\sigma_g}\left(\ln\left(\frac{1}{N}\sum_{i=1}^{N} \gamma_{0,i} p_{t,i}^{1-\sigma_g}\right) - \ln\left(\frac{1}{N}\sum_{i=1}^{N} \gamma_{0,i} p_{0,i}^{1-\sigma_g}\right)\right).$$

By the law of large numbers, taking the limit $N \to \infty$ on both sides yields

$$\ln\left(\frac{P_{t,g}}{P_{0,g}}\right) = \frac{1}{1-\sigma_g}\left(\ln\left(\mathbb{E}\left[\gamma_{0,i} p_{t,i}^{1-\sigma_g}\right]\right) - \ln\left(\mathbb{E}\left[\gamma_{0,i} p_{0,i}^{1-\sigma_g}\right]\right)\right)$$

where we have $\gamma_{i,0} = w_{i,0}\left(\frac{p_{i,0}}{P_{g,0}}\right)^{\sigma_g - 1}$, as given by Equation (3). Plugging this into the above yields

$$\ln\left(\frac{P_{t,g}}{P_{0,g}}\right) = \frac{1}{1-\sigma_g}\left(\ln\left(\mathbb{E}\left[w_{0,i} P_{g,0}^{1-\sigma_g}\left(\frac{p_{t,i}}{p_{0,i}}\right)^{1-\sigma_g}\right]\right) - \ln\left(\mathbb{E}\left[w_{0,i} P_{0,g}^{1-\sigma_g}\right]\right)\right),$$

or

$$\ln\left(\frac{P_{t,g}}{P_{0,g}}\right) = \frac{1}{1-\sigma_g}\left(\ln\left(\mathbb{E}\left[w_{0,i}\left(\frac{p_{t,i}}{p_{0,i}}\right)^{1-\sigma_g}\right]\right) - \ln\left(\mathbb{E}\left[w_{0,i}\right]\right)\right). \tag{45}$$

We now want to find a simpler expression for

$$\mathbb{E}\left[w_{0,i}\left(\frac{p_{t,i}}{p_{0,i}}\right)^{1-\sigma_g}\right].$$

To this end, assume that $w_{i,0}$ and $\frac{p_{t,i}}{p_{0,i}}$ are jointly log-normally distributed. First note that $w_{0,i}\left(\frac{p_{t,i}}{p_{0,i}}\right)^{1-\sigma_g} = \exp\left\{\ln(w_{0,i}) + (1-\sigma_g)\ln\left(\frac{p_{t,i}}{p_{0,i}}\right)\right\}$. Since a linear combination of random variables with joint normal distribution also follows a normal distribution, we have that $\ln(w_{0,i}) + (1-\sigma_g)\ln\left(\frac{p_{t,i}}{p_{0,i}}\right)$ is normally distributed and hence that $\exp\left\{\ln(w_{0,i}) + (1-\sigma_g)\ln\left(\frac{p_{t,i}}{p_{0,i}}\right)\right\} = w_{0,i}\left(\frac{p_{t,i}}{p_{0,i}}\right)^{1-\sigma_g}$ has a log-normal distribution. Second, any variable z with a log-normal distribution satisfy the property $\mathbb{E}[z^n] = \exp\left\{n\mathbb{E}[\ln(z)] + \frac{1}{2}n^2\mathbb{V}[\ln(z)]\right\}$ for any scalar n, where $\mathbb{V}$ denotes the variance. Thus, we have that

$$\begin{aligned}\mathbb{E}\left[w_{0,i}\left(\frac{p_{t,i}}{p_{0,i}}\right)^{1-\sigma_g}\right] &= \exp\left\{\mathbb{E}\left[\ln\left(w_{0,i}\left(\frac{p_{t,i}}{p_{0,i}}\right)^{1-\sigma_g}\right)\right] + \frac{1}{2}\mathbb{V}\left[\ln\left(w_{0,i}\left(\frac{p_{t,i}}{p_{0,i}}\right)^{1-\sigma_g}\right)\right]\right\} \\ &= \exp\left\{\mathbb{E}\left[\ln(w_{0,i})\right] + (1-\sigma_g)\mathbb{E}\left[\ln\left(\frac{p_{t,i}}{p_{0,i}}\right)\right] + \frac{1}{2}\mathbb{V}\left[\ln(w_{0,i}) + (1-\sigma_g)\ln\left(\frac{p_{t,i}}{p_{0,i}}\right)\right]\right\}.\end{aligned}$$

Using this in combination with

$$\mathbb{V}\left[\ln\left(w_{0,i}\right)+(1-\sigma_g)\ln\left(\frac{p_{t,i}}{p_{0,i}}\right)\right]=\mathbb{V}\left[\ln\left(w_{0,i}\right)\right]+(1-\sigma_g)^2\mathbb{V}\left[\ln\left(\frac{p_{t,i}}{p_{0,i}}\right)\right]+2(1-\sigma_g)\mathbb{C}\left[\ln\left(w_{0,i}\right),\ln\left(\frac{p_{t,i}}{p_{0,i}}\right)\right]$$

(where $\mathbb{C}$ denotes the covariance) we get

$$\begin{aligned}\ln\left(\frac{P_{t,g}}{P_{0,g}}\right)=&\ \mathbb{E}\left[\ln\left(\frac{p_{t,i}}{p_{0,i}}\right)\right]+\frac{1-\sigma_g}{2}\mathbb{V}\left[\ln\left(\frac{p_{t,i}}{p_{0,i}}\right)\right]\\&+\frac{1}{1-\sigma_g}\Big(\mathbb{E}\left[\ln\left(w_{0,i}\right)\right]-\ln\left(\mathbb{E}\left[w_{0,i}\right]\right)\Big)\\&+\frac{1}{2(1-\sigma_g)}\mathbb{V}\left[\ln\left(w_{0,i}\right)\right]+\mathbb{C}\left[\ln\left(w_{0,i}\right),\ln\left(\frac{p_{t,i}}{p_{0,i}}\right)\right]\end{aligned}\tag{46}$$

Finally, use the property that the marginal distributions of jointly log-normal variables are univariate log-normal, i.e., since $w_{i,0}$ and $\frac{p_{t,i}}{p_{0,i}}$ are jointly log-normally distributed, we know that $w_{i,0}$ is log-normally distributed. Hence $\ln\left(\mathbb{E}\left[w_{i,0}\right]\right)=\mathbb{E}\left[\ln\left(w_{i,0}\right)\right]+\frac{1}{2}\mathbb{V}\left[\ln\left(w_{i,0}\right)\right]$ and (46) simplifies to

$$\ln\left(\frac{P_{t,g}}{P_{0,g}}\right)=\mathbb{E}\left[\ln\left(\frac{p_{t,i}}{p_{0,i}}\right)\right]+\frac{1-\sigma_g}{2}\mathbb{V}\left[\ln\left(\frac{p_{t,i}}{p_{0,i}}\right)\right]+\mathbb{C}\left[\ln\left(w_{0,i}\right),\ln\left(\frac{p_{t,i}}{p_{0,i}}\right)\right].\tag{47}$$

Finally, formula (10) in the main paper is given by taking the exponential on both sides and defining $\mu_p=\mathbb{E}\left[\ln\left(\frac{p_{t,i}}{p_{0,i}}\right)\right]$, $\rho_p^2=\mathbb{V}\left[\ln\left(\frac{p_{t,i}}{p_{0,i}}\right)\right]$, and $\rho_{p,w}=\mathbb{C}\left[\ln\left(w_{0,i}\right),\ln\left(\frac{p_{t,i}}{p_{0,i}}\right)\right]$.

□

E Log-normal approximation errors

This section first reports how the approximation errors of annual aggregate inflation vary over time. Then, it reports the distribution of approximation errors across group-years, with no restriction on the number of items in each group-year (in contrast to in the main text, where only group-years with at least four items were included).

Figure 35 shows how the approximation error for goods with $\sigma_g=0$ varies over time by plotting the annual inflation rates of the true index and the log-normal approximation as well as their difference. Three years—2005, 2008, and 2016—stand out in terms of having relatively large approximation errors of 0.6–0.7 percentage points of annual inflation.

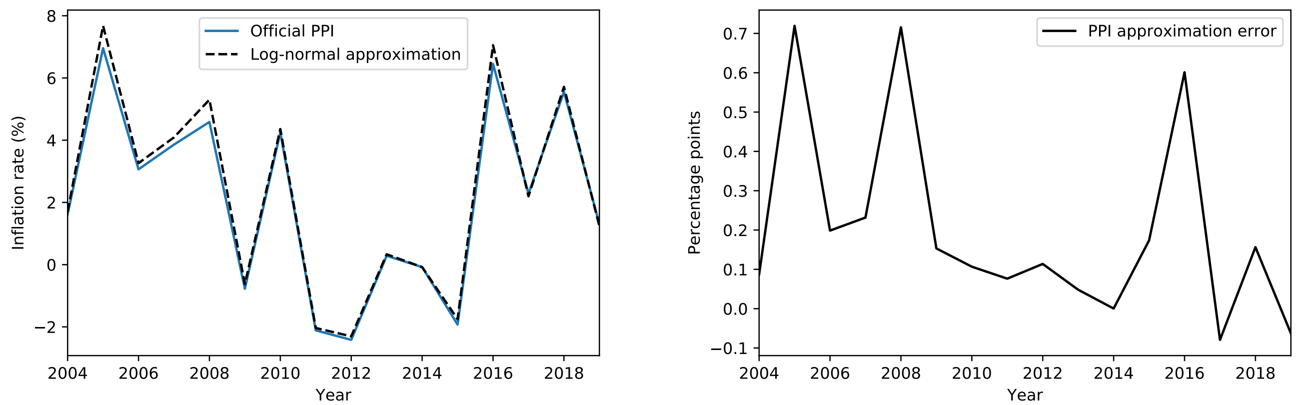

Figure 35: Annual inflation rates as measured by the baseline goods PPI and the log-normal approximation with $\sigma_g = 0$ and $\sigma = 0$ (left figure), and their difference (right figure).

Figure 36 shows the corresponding dynamics for services with $\sigma_s = 1$. The years 2017 and 2019 have relatively high approcimation errors of about 0.3 percentage points.

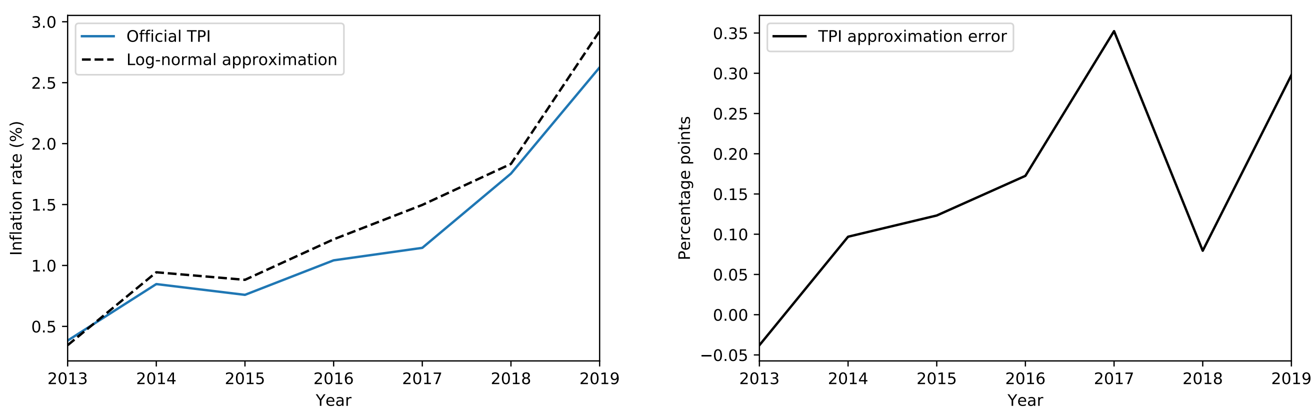

Figure 36: Annual inflation rates as measured by the baseline goods PPI and the log-normal approximation (left figure), and their difference (right figure).

We now show the distribution of approximation errors when including all group-year observations, and not only when keeping the group-years containing at least 4 items.

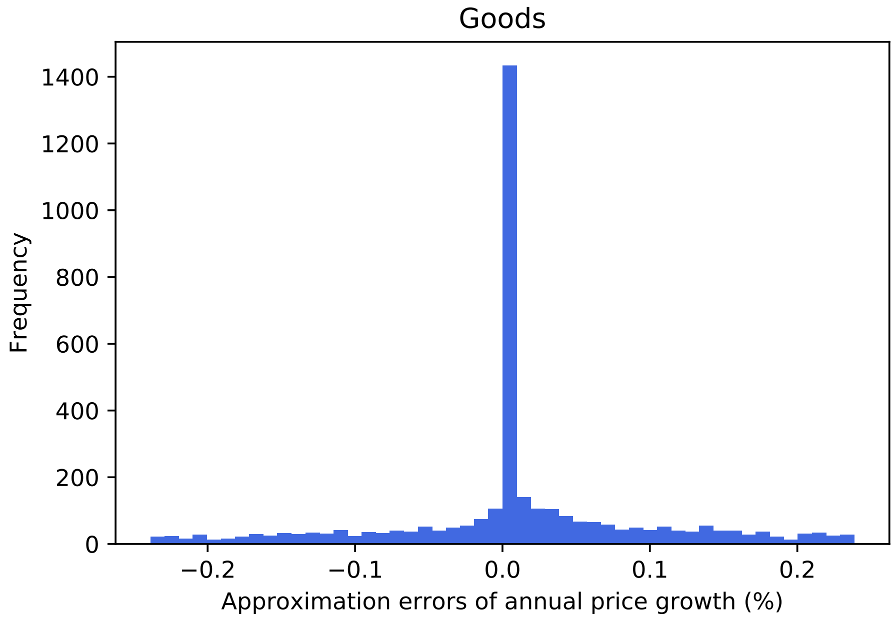

Figure 37: Distribution of group-level annual price growth approximation errors for goods. Each observation is given by a group-year pair, and all elasticities of substitution between products within product groups are set to $\sigma_g = 0$. The horizontal axis is cut off symmetrically around zero by the limits $\pm\max(|25^{\text{th}}$ percentile$|, |75^{\text{th}}$ percentile$|)$.

Group-level error distribution: Goods

	25^{th}	50^{th}	75^{th}	Mean
Approximation error (%)	-0.102	0.000	0.239	0.203

Table 34: Percentiles and mean of group-level annual price growth approximation errors for goods. Each observation is given by a group-year pair, and all elasticities of substitution between products within product groups are set to $\sigma_g = 0$.

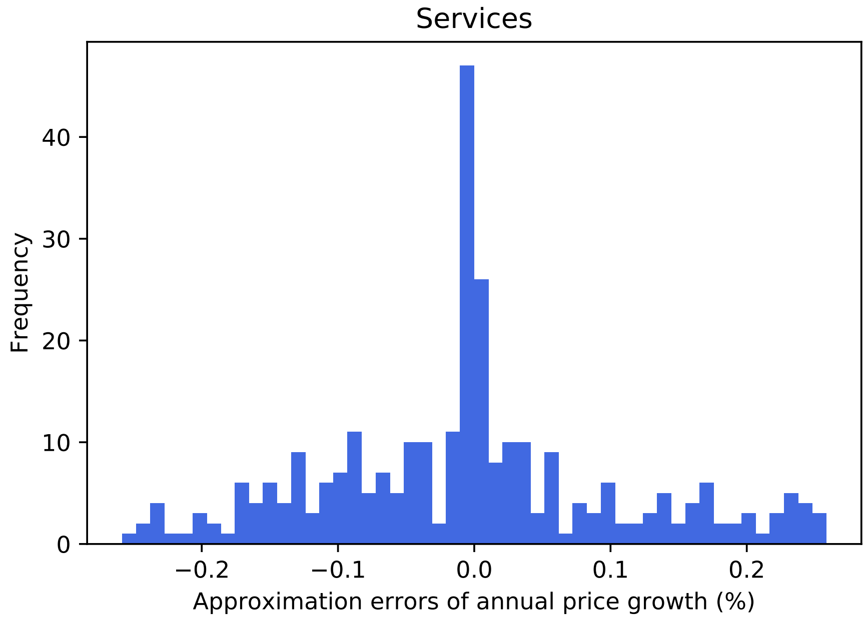

Figure 38: Distribution of group-level annual price growth approximation errors for services. Each observation is given by a group-year pair. The horizontal axis is cut off symmetrically around zero by the limits $\pm \max(|25^{\text{th}} \text{ percentile}|, |75^{\text{th}} \text{ percentile}|)$.

Group-level error distribution: Services

	25^{th}	50^{th}	75^{th}	Mean
Approximation error (%)	-0.200	-0.000	0.258	0.489

Table 35: Percentiles and mean of group-level annual price growth approximation errors for services. Each observation is given by a group-year pair, and all elasticities of substitution between products within product groups are set to $\sigma_g = 1$.