



Written for the NBER International Seminar on Macroeconomics

Price discrimination within and across EMU markets: Evidence from French exporters[☆]

François Fontaine^a, Julien Martin^b, Isabelle Mejean^{c,*}

^a Paris School of Economics, Université Paris Panthéon Sorbonne, France

^b Université du Québec à Montréal, CEPR, Canada

^c CREST-Ecole Polytechnique, CEPR, France

ARTICLE INFO

Article history:

Received 26 August 2019

Received in revised form 13 December 2019

Accepted 17 January 2020

Available online 23 January 2020

Mendeley Data <https://data.mendeley.com/datasets/k8c3wmdj8z/draft?a=c32c2dae-08b8-4f66-b3a5-a6634401cc3c>

Keywords:

Price discrimination

Firm-to-firm trade

ABSTRACT

We study the cross-sectional dispersion of prices paid by EMU importers for French products. We document a significant level of dispersion in unit values both within product categories across exporters, and within exporters across buyers. This latter source of price discrepancies, which we call price discrimination, reflects the ability of exporters to sell similar or differentiated varieties of a given product at different prices to different buyers. Price discrimination (i) is substantial within the EU, within the euro area, and within EMU countries; (ii) has not decreased over the last two decades; (iii) is more prevalent among the largest firms and for more differentiated products; (iv) is lower among retailers and wholesalers; (v) is also observed within almost perfectly homogenous product categories, which suggests that a non-negligible share of price discrimination is partly triggered by heterogeneous markups rather than quality or composition effects. We then estimate a rich statistical decomposition of the variance of prices to shed light on exporters' pricing strategies.

© 2020 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

The pricing strategy of exporters is a central element of international macroeconomics. Whether exporters sell differentiated or homogenous products and charge the same or different prices across partners located in various destinations is key to understand deviations from the law of one price and to anticipate exporters' reaction to domestic and foreign shocks. In this paper we exploit firm-to-firm data to shed light on the pricing strategies of French exporters. The data cover both final and intermediate products sold to different buyers within and across EMU countries. These data enable us to quantify the extent of price discrimination, to analyze the heterogeneity in the pricing behavior across producers and sectors, and to discuss the mechanisms behind price discrimination.

We document a significant level of dispersion in the unit values charged by producers across their buyers within the European Monetary

[☆] We are grateful to Brent Neiman, Ida Hjortsoe, Charles Engel, Pierre-Olivier Gourinchas, Jordi Gali, and participants at the NBER-ISOM conference for insightful comments, as well as Thibault Cezanne for excellent research assistance. Fontaine gratefully acknowledges financial support from the Institut Universitaire de France. Mejean gratefully acknowledges support from a public grant overseen by the French National Research Agency (ANR) as part of the "Investissements d'Avenir" program (Idex Grant Agreement No. ANR-11-IDEX-0003-02, Labex ECODEC Grant Agreement No. ANR-11-LABEX-0047 and Equipex reference: ANR-10-EQPX-17 à l'Accès sécurisé aux données à CASD) as well as the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 714597).

* Corresponding author.

E-mail address: isabelle.mejean@polytechnique.edu (I. Mejean).

Union. Price discrimination is prevalent across buyers located in different countries, but also across buyers within a country. The data show a large heterogeneity across French exporters in their propensity to price discriminate. Large multiproduct exporters tend to adopt more discriminatory pricing strategies. Retailers and small exporters instead charge less dispersed and sometimes nearly uniform prices. Finally, price dispersion is stronger for more differentiated and more durable products.

These evidence rely on data on the unit prices charged by French exporters to their European buyers over 2002–2016 for each of the 9000 different product categories of the CN8 trade nomenclature. At this level of aggregation, the uncovered price dispersion can be driven by i) exporters charging different unit prices for differentiated varieties of the same product, or ii) exporters setting different markups across their customers of a given variety.¹ Whereas, we cannot distinguish between these two forms of discrimination, we show that price discrimination remains substantial for a subset of product categories that offer little ground for product differentiation. Focusing on the price per kilogram of 200 raw molecular substances, we indeed find a level of price discrimination in this subsample that is two third of the price discrimination measured in the whole dataset.

We conclude the paper with an examination of the extent to which exporters' ability to set high prices on their European partners, together

¹ Regarding driver i), differences in variety across buyers might reflect differences in the attributes of the products (eg. its quality), packaging differences, or differences in the mix of varieties purchased by the importer.

with importers' tendency to renegotiate prices on the match participate to the dispersion of prices within a firm. Estimates recovered from a high-dimensional fixed effect model show that prices tend to decrease over time, within a firm-to-firm relationship, while exporters set high prices on newly met partners. These behaviors suggest that some form of price bargaining occurs between the exporter and her foreign partner.

These new facts build on a unique dataset on firm-to-firm trade in the EU. For each of the 9000 products that the data cover, we observe a set of export transactions taking place in a given quarter between a particular French firm and one of its partners in the EU. The data enable us to compare the price strategy of two French exporters selling a similar product to a given EU destination as well as prices set by the same firm over different partners. At firm-level, any dispersion in the FOB unit values means exporters set different markups and/or supply differentiated varieties of the same product to buyers in their portfolio. This level of dispersion constitutes our measure of price discrimination.²

We start our analysis by quantifying how this source of price discrepancies influences the overall variance of prices observed in the data. To this aim, we construct a measure of price dispersion at the product level for each quarter and calculate the extent to which these price discrepancies come from different exporters serving European markets at different mean prices, a "between" component, versus individual exporters price discriminating between partners in their portfolio, a "within" component. The level of price dispersion recovered from these data is substantial. The mean coefficient of variation of prices in the EU is as high as 1.3. Two thirds of this dispersion is due to the between component, that is, exporters setting heterogeneous average prices to serve the same or different partners with potentially differentiated products. Still, a third of the cross-sectional variance in prices is attributable to the within-seller dimension, that is, exporters charging heterogeneous prices across their different clients. The rest of the analysis is dedicated to this specific source of price discrepancies, which we refer to as *price discrimination*.

Price discrimination is a common practice among French exporters. The median coefficient of variation of prices across buyers purchasing the same product from a given exporting firm in a specific quarter is as high as 30.5%. This average, however, hides a substantial amount of heterogeneity. In the limit, 14% of exporters have uniform pricing strategies in the EMU, yet these firms are relatively small and thus contribute little to aggregate exports.

Although the within-firm price dispersion implies systematic deviations from the LOP within the euro area, we also document that price dispersion at the firm-level is less severe within the EMU than in the overall EU. Mean differences across country samples within a firm are quantitatively important. Prices within the extended EU are, on average, 10% more dispersed than within the EU restricted to its 15 old members, whereas they are 14% less dispersed in the EMU than in the EU15. These differences are in part due to composition effects, within a firm, but we show the difference is still significant when we use firm-level randomization to compare prices within and outside of the EMU. This finding confirms that sharing a common currency causes greater market integration. The level of price dispersion has, however, increased over time, especially for relatively small firms. The coefficient of variation of prices recovered within a firm was 25% higher in the 2010s than in the 2000s, a result that is robust to composition effects. This result goes against the view that both the increasing integration of European markets and new communication technologies should enable consumers to arbitrage across goods, which is expected to force the

² One may argue that LOP should be considered at the level of consumer prices, thus including transportation costs. If arbitrage is strong enough, exporters may be forced to absorb trade costs, which would transmit into heterogeneous *FOB* prices but homogenous *CIF* prices. Consistent with existing evidence based on firm-to-destination export data (e.g., see Manova and Zhang, 2012; Martin, 2012, for Chinese and French data), our firm-to-firm *FOB* prices are increasing in distance. This finding suggests that, if anything, the corresponding *CIF* prices should be more dispersed than the *FOB* prices we study.

convergence of prices. Instead, the increasing dispersion of prices observed within an exporter over time suggests small exporters in our sample manage to maintain high price discrepancies, potentially thanks to product differentiation.

In a second step, we study how firm and product heterogeneity is related to the degree of price discrimination. Among the characteristics that might explain why firms are unequally prone to price discriminating, we find a significant effect of the firm's size and profits. Large multi-product exporters and firms with a greater profitability within their sector of activity are found to price discriminate more intensively. Within a firm, the propensity to price discriminate is weaker over the firm's core product. Finally, we find evidence of heterogeneity in price discrepancies across sectors. Retailers and wholesalers charge less dispersed prices than manufacturing firms supplying the same type of products. Price dispersion is stronger for differentiated products, especially durable ones.³ Along the value chain, price discrimination is more stringent for more downstream products. This heterogeneity may explain by systematic differences in consumers' ability to arbitrage induced by product differentiation or the structure of competition in various markets.

The dispersion of prices within a firm is consistent with two potentially complementary mechanisms. First, exporters may price discriminate across their partners through product differentiation, for example, by customizing their product to their customers' needs. Such a strategy should be especially relevant for differentiated goods, thus the higher mean dispersion of prices observed for these products. Second, exporters may sell the same product variety to various buyers at differentiated prices, thus adjusting their markup to their buyers' valuation for the good. Although the data do not allow us to quantify the relative contribution of both factors to the observed dispersion of prices, we conclude the analysis with two exercises that are meant to dig deeper into the underlying mechanisms of price discrimination. In the first exercise, we focus on a sub-sample of roughly 200 chemical products that we argue offer very little ground for product differentiation, because they correspond to raw molecular substances. By comparing the level of price dispersion in this sample and in the rest of the dataset, we can provide some indicative elements regarding the role of product differentiation as a source of price discrepancies. In the sample of homogenous products, the mean coefficient of variation is about 10 percentage points lower than in the control group. The difference is significant, including when identified within firms selling both homogenous and differentiated chemical products, controlling for unobserved heterogeneity across firms. Extrapolating these results beyond the chemical industry suggests that about a quarter of the observed price dispersion is due to product differentiation within a firm.⁴

The second exercise digs deeper into the pricing strategies of French exporters, using a rich linear model to analyze the determinants of export price levels. Using insights from the labor literature (Abowd et al., 1999), we estimate a price equation with two-sided unobserved heterogeneity (seller and buyer) that allows us to characterize the dynamics of firm-to-firm prices, conditional on sellers' and buyers' unobserved heterogeneity. Results show that firm-to-firm prices tend to decrease with the age of the buyer-seller relationship, which is consistent with buyers renegotiating and increasing their share of the transaction's surplus as they increase their outside options.⁵ Despite downward price

³ This finding confirms the singular role played by durable goods in open-macroeconomics (see Engel and Wang, 2011; Levchenko et al., 2010).

⁴ Such extrapolation is arguably heroic because the average difference in the dispersion of prices between homogenous and heterogeneous products is estimated for firms in the chemical industry, which may not be representative of the average firm in the data.

⁵ Studying the dynamics of prices within a firm-to-firm relationship is insightful as it allows us to dampen composition effects that may drive price dispersion between but also within an exporter. Still, unit values recovered from individual transactions may suffer from measurement bias since buyers may adjust the composition of the bundle purchased from a seller over time. The systematic (downward) pattern observed in the corresponding series of firm-to-firm unit values suggests that i) either this composition effects evolve over time in the direction of firms purchasing more and more of the low price varieties or ii) prices are renegotiated downward on the match which is our preferred interpretation.

renegotiations, the mean price set by French exporters increases over time. The reason is that more experienced exporters manage to expand their portfolio of buyers and charge their new consumers relatively high prices. Interestingly, firms at the top of the distribution of sales in their sector are especially good at charging new consumers high prices while suffering from relatively less pronounced downward price renegotiation. We thus offer some insights on the way superstar firms exert their market power.

Our data enables us to observe unit values set by seller across different buyers. As such, our work contributes to the literature on firm-to-firm trade, and to the literature on the pricing strategy of exporters. The literature on firm-to-firm trade has not explored the price dimension yet, and most of the literature on the pricing strategy of exporters uses unit values without the buyer dimension. The main limitation of our data is that differences in unit values are subject to composition effects, and thus reflect both differences in prices and, possibly, in the (mix of) varieties sold to different buyers.⁶ For most product categories, the price dispersion we document comes both from markup differences and variety differences. Nevertheless, the high level of dispersion we document implies that a model where sellers sell the same variety and charge the same markup to all their buyers would not fit the data. Moreover, we also find a high level of dispersion in a sample of chemicals where composition biases are likely to be limited, suggesting that markup differences are probably pervasive overall.

Literature review. In addition to the works cited above, this paper pertains to different strands of the literature. Deviations from the LOP are often associated with market segmentation and border effects. Engel and John (1996) document systematic deviations from the LOP using disaggregated consumer price indices across Canadian and US cities. Using similar data across European cities, the authors do not find evidence of price convergence after the introduction of the euro (Engel and Rogers, 2004). Within the car industry, Goldberg and Verboven (2005) find a strong positive impact of the European integration on price convergence, and a weaker impact on the level of price dispersion. We focus here on the absolute version of the LOP. As in Engel and John (1996), we exploit the granularity of the data in the spatial dimension to compare the level of price discrepancies within the euro area and within countries of the euro area.

Part of the literature relates deviations from the LOP at the consumer level to the extent of local distribution costs (Crucini et al., 2005; Crucini and Shintani, 2008). According to Gopinath et al. (2011), these distribution costs are not the main source of price discrepancies, which are instead high upstream in the value chain, at the wholesale level. Our analysis confirms their result by documenting the large degree of price discrepancies at the producer level. The evidence documented in Gopinath et al. (2011) further suggests that the price differences we document are likely to translate into price discrepancies at the consumer price level.

Because our data cover both manufacturing firms and wholesalers and retailers, for a wide range of different products, we can also compare the propensity to price discriminate at different points of the value chain. Although price discrepancies are large on average in all sectors, we do find some evidence of the propensity to price discriminate being smaller in the retail sector, within a product. The lower level of price discrimination by retailers is consistent with results in Cavallo et al. (2014) on the LOP within the EMU. The paper documents the importance of uniform pricing across euro countries for products sold online by four large retailers. To our knowledge, this paper is the first to

document uniform pricing across different countries. Although we find retailers (and non-durable goods) have a lower price dispersion in our data, the prevalence of uniform pricing is not striking. This behavior concerns about 14% of product varieties accounting for 2% of the value of trade.⁷

The literature has also examined price discrepancies in a national context. Most papers focus on specific industries and get quite different pictures. Using barcode data, DellaVigna and Gentzkow (2017) show that the vast majority of large US retailers charge uniform or nearly uniform prices across their stores.⁸ Cavallo (2018) shows the degree of uniform pricing of the largest US retailers across US locations has increased over the last 10 years, partly driven by on-line competition. By contrast, Adams and Williams (2019) focus on price dispersion in the home-improvement industry. They find substantial price dispersion in this sector and document the granularity of zone pricing. They further show that big players in this industry adopt different pricing strategies. Our work is also related to Kaplan and Menzio (2015), who describe the distribution of prices at which identical consumer goods are sold within a market. They find substantial dispersion in consumer prices, within narrowly defined products. As discussed above, we also document a substantial heterogeneity in the pricing practices of French exporters across sectors.

Our work also contributes to a literature that uses increasingly disaggregated data to understand the microeconomic underpinnings of incomplete exchange-rate pass-through and pricing-to-market⁹ (e.g., Berman et al., 2012; Amiti et al., 2019). The closest papers are Devereux et al. (2017) and Goldberg and Tille (2016) who use transaction-level data to discuss the role of market power on both sides of the trade relationship. Our estimates are consistent with exporters and importers sharing the surplus of the transaction. Because we can observe repeated transactions within a relationship, we can further discuss how this sharing evolves over time, and provide evidence of downward price renegotiation “on-the-match”. Moreover, we are able to document the extent to which market segmentation affects the dispersion of prices not only across countries but also within a destination, across the exporter's partners.

Finally, the paper is related to the emerging empirical literature on firm-to-firm trade. Papers have mostly focused on the value and growth of firm-to-firm trade flows (see Bernard and Moxnes, 2018, for a review). Our descriptive analysis capitalizes on two features of the data. First, we have information on firm-to-firm trade at the product level which allows us to compare sellers and buyers network and prices within detailed product categories. Second, we have transaction-level information on the value and quantity of exports, which allows us to compute unit values, and to provide the first descriptive evidence of price dispersion in this type of data.

The rest of the paper is organized as follows. Section 2 describes the data used to document the extent of price discrepancies in French

⁶ Examples of top product categories in French exports are: Aeroplanes and other powered aircraft of an unladen weight > 15,000 kg (excl. helicopters and dirigibles), Motor cars and other motor vehicles principally designed for the transport of persons (other than those of heading No 8702), incl. Station wagons and racing cars, with spark-ignition internal combustion reciprocating piston engine, of a cylinder capacity >1.500 cm³ but ≤ 3,000 cm³, new (excl. 8703.10–10 and 8703.23.11), Motor spirit, with a lead content ≤ 0,013 g/l, with a research octane number “RON” of <95, or Champagne of actual alcoholic strength of ≥ 8,5% vol.

⁷ Our results are not directly comparable with Cavallo et al. (2014) though, who use actual price data collected from four giant retailers. Their data do not suffer from the measurement bias inherent to using FOB unit values, which does drive a significant share of the dispersion we observe. Moreover, they solely cover final consumption goods while our data are recovered from business-to-business relationships, and include products used for intermediate consumption purposes. Such relationships display a lot more stickiness, as shown by the fact we observe repeated transactions involving the same partners over time. These long-term relationships give rise to bargaining, as suggested by our empirical evidence. Bargaining is itself a source of price discrepancies, in the cross-section.

⁸ Our paper differ along three dimensions from this paper. First, we work with unit values rather than barcode prices and thus the dispersion under study mixes composition effects and actual price differences. Second, we examine the price of intermediate products rather than final consumer prices. Third, the object of interest is different in the two papers. DellaVigna and Gentzkow (2017) examine the price dispersion of a given barcode product within chains, across stores (eg. the price of a can of Coke across different Walmart stores). We instead focus on the price dispersion of a seller across different firms (eg. the price of a can of Coke set by Coca-Cola to different retailers).

⁹ Pricing to market refers to situations in which a firm charges different prices when selling the same good to different markets segmented by different currencies (see, e.g., Atkeson and Burstein, 2008; Fitzgerald and Haller, 2014)

exports. Stylized facts on price discrepancies are then presented in three steps. Section 3 discusses the extent to which deviations from the LOP within a firm contribute to the overall dispersion of prices observed in the data. In Section 4, we study heterogeneity in firms' propensity to price discriminate over space, over time, and across firms. Section 5 digs deeper into exporters' pricing strategies to discuss the underlying mechanisms at the root of observed price discrepancies. Section 6 concludes.

2. Data and summary statistics

Throughout the analysis, we rely on export data provided to us by the French customs and covering the universe of export transactions from France to the rest of the EU. A full description of the data can be found in Bergounhon et al. (2018). Details on the construction of the variables used in the analysis can be found in Appendix A. The data allow us to identify both parties involved in a transaction, namely, the French exporter, identified by its Siren number, and the European importing firm, identified by its (anonymized) VAT number.¹⁰ This firm-to-firm dimension is useful because it allows us to compare pricing strategies across producers serving the same market and eventually the same buyer with the same product as well as prices offered by a given exporter to different partners located in the same or in different European markets. We exploit the cross-sectional richness in the analysis.

On top of the identity of both firms involved in the export flow, transactions recorded in the dataset are characterized by a date, at the monthly frequency, a product category at the 8-digit level of the combined nomenclature, the value of the transaction, and the physical quantity being traded.¹¹ Although the data are exhaustive, small exporters are allowed to complete a simplified form that does not request information on the product category or the physical quantity exported. Because these variables are key in the analysis, we neglect this population of firms. Between 2001 and 2006, the simplified regime concerned exporters whose annual export turnover in the EU was below 100,000 euros. The declaration threshold was increased to 150,000 euros in 2007 and to 460,000 euros in 2011. Therefore, our working sample is censored to the left of the distribution of exporters' size and the censoring increases over time. Censored observations, on average, represent 36% of exporters accounting for 13% of the value of trade during the main period of analysis, 2002–2006. We also present some results based on the 2012–2016 period, when the simplified regime represents 63% of exporters and 18% of the value of trade, on average.

The analysis mostly focuses on the cross-sectional dispersion of prices, within a given product category and a given quarter. But we also want to study how this cross-sectional dispersion evolves over time. Doing so requires identifying time-consistent product categories, which is cumbersome when working with the combined nomenclature because it continuously evolves over time. We follow Behrens et al. (2019) and harmonize product categories by nesting into broader clusters products that are connected through nomenclature updates. Because this methodology can produce relatively large clusters of products when applied over long horizons, we decided to restrict our attention to two five-year periods, 2002–2006 and 2012–2016. These subperiods are not affected by major revisions of the harmonized system at the root of the combined nomenclature. Working on relatively short periods limits the number of product categories that are grouped

¹⁰ These data are collected for VAT purposes and solely cover trade between firms. We thus do not include direct exports by a firm to a final consumer in the rest of the analysis. This restriction represents less than 1% of the value of exports in overall customs data. Unfortunately, the Customs do not report information on the nature of transactions (arm's length *vs* intrafirm).

¹¹ Although the raw data are available at the monthly frequency, we aggregate transactions within a quarter to compute statistics over the cross-sectional dispersion of prices in sections 3 and 4. Doing so allows us the benefit of high frequency while slightly increasing the dimensionality used to recover information on price discrepancies.

together through the harmonization algorithm.¹² But this also means that product categories are not fully comparable across sub-periods. Whenever price strategies are compared over long periods, within a firm, the analysis is restricted to product categories that are the same in both subperiods. This restricted sample represents 6896 products, out of the 9402 categories observed over 2002–2006.

For each transaction, we recover a price proxy, defined as the unit value:

$$p_{sb(c)pt} \equiv \frac{Value_{sb(c)pt}}{Quantity_{sb(c)pt}}$$

where the s , $b(c)$, p , and t subscripts, respectively, refer to the identity of the seller, the buyer (which is further identified by its origin country c), the product being exported, and the time of the transaction.¹³ The value of the transaction, $Value_{sb(c)pt}$, is measured in euros and is fob. The analysis excludes transactions below 100 euros, because of rounding issues. The quantity, $Quantity_{sb(c)pt}$ is either measured in kilograms or in physical units for some specific product categories. Therefore, unit values are not necessarily comparable across products but they are within a product category, the focus of the analysis.

The structure of the dataset, which enables us to compute unit values for each trade transaction, helps mitigate composition effects that have been argued to reduce the quality of unit values as a proxy for prices. Because unit values can still suffer from measurement issues when either the value or the quantity is misreported, we trim the data and remove price quotes that deviate from the median price set by the firm for this product over the considered year by more than 200%.¹⁴ The remaining differences in transaction-level unit values observed across and within an exporter for a given product and period imply the same quantity is sold at different prices. In theory, these price discrepancies can be attributable to heterogeneity in mark-ups, heterogeneity in marginal costs, and/or the vertical differentiation of the good that can take the form of different packagings for the same product, the addition of various optional characteristics to an existing good or the production of two differentiated varieties within a product category. Most of the analysis is agnostic about the origin of observed price discrepancies, the discussion of the mechanisms at the root of price discrimination being delayed to Section 5. Our approach consists of gradually reducing the potential for cost and product differentiation by first focusing on the dispersion within a product, then on price discrepancies within a particular exporter of this product and finally on the variance of prices over time within a firm-to-firm relationship. Even at this level of analysis though, unit values can be affected by composition effects, if the buyer purchases a mix of two differentiated varieties which composition varies over time. Section 5 discusses the extent to which product differentiation is likely to explain results in Sections 3 and 4.

Over 2002–2006, our dataset is composed of more than 37.7 million observations involving 70,649 exporters, 1.1 million importers located in 24 European countries, and 9400 (harmonized) product categories. Table 1 provides detailed statistics over the structure of the dataset, by destination country. Note that the period encompasses the entry of 10 Eastern European countries into the EU and thus into the dataset. For

¹² Within France's most important products in value share, one can cite several categories of transport equipments, a number of medicaments and champagne, i.e. products that are strongly differentiated and often supplied in many differentiated varieties within an exporter. Other important product categories include goods that are more homogenous, e.g. electrical energy or a number of raw molecules produced by the chemical industry.

¹³ We take the average price of a seller-buyer relationship within a quarter to increase the number of observations used to compute the coefficient of variation. The dispersion we measure is thus the dispersion in the quarterly price of a given product.

¹⁴ This price range may still appear large. However, Adams and Williams (2019) document that the price of Home Depot's 4' x 8' x 1/2" mold-resistant drywall ranges from 7.65 to 23.71 USD across locations. Kaplan and Menzio (2015) show that the price of a 36-oz plastic bottle of Heinz ketchup ranges from 0.5 to 2.99 USD. This restriction concerns less than 2% of the transactions in the data.

Table 1
Dimensionality of the data.

	Number of			
	Transactions	Exporters	Importers	Relationships
$sb(c)pt$	s	$b(c)$	$sb(c)p$	
(1)	(2)	(3)	(4)	
All	37,796,239	70,649	1,103,275	8,626,857
Austria	893,889	14,924	31,638	232,339
Belgium	6,329,954	46,765	128,592	1,397,788
Cyprus	63,891	3061	2271	19,906
Czech Republic	261,788	8601	8101	50,723
Denmark	697,829	15,106	17,968	159,829
Estonia	53,873	2283	1617	12,498
Finland	404,946	9287	9978	83,978
Germany	6,661,428	40,437	228,985	1,414,047
Greece	825,919	13,514	25,577	235,831
Hungary	203,617	6873	5884	40,803
Ireland	532,835	11,297	12,898	138,614
Italy	5,134,450	34,992	188,556	1,290,050
Latvia	53,164	2546	1796	14,281
Lithuania	60,250	3420	2342	16,187
Luxembourg	941,590	19,289	18,226	254,588
Malta	44,014	2395	1279	12,454
Netherlands	2,286,535	28,684	63,231	506,606
Poland	431,354	11,956	16,664	95,659
Portugal	1,717,826	20,974	42,307	394,948
Slovak Republic	79,645	4008	2913	18,491
Slovenia	110,763	3548	2760	20,896
Spain	5,355,890	36,395	164,399	1,230,907
Sweden	767,925	13,547	19,947	156,392
United Kingdom	3,882,864	32,049	105,346	829,042

Notes: Column (1) is the number of transactions recorded over 2002–2006. Columns (2) and (3), respectively, report the number of French exporters and European importers involved in these transactions. Finally, column (4) is the number of exporter×importer×product triplets. The ratio of column (4) to (1) allows us to recover the mean number of transactions observed over time for a particular firm-to-firm relationship and a particular product.

this reason, we compute a number of statistics on a sample restricted to the 15 “old” members of the EU.¹⁵ Likewise, EMU members are those that were already part of the euro-area in 2002. For each observation, we observe a transaction-specific price quote. We use the high dimensionality to condition the statistics regarding the variance of prices on a particular position in the network. Namely, we start by computing the variance of prices conditional on a particular product×quarter. Then, we further focus on price discrepancies within a firm. The remaining dimensions in the data can then be used to discuss how the dispersion of prices varies across firms, over space, and over time.

3. From price dispersion to price discrimination

In this section, we establish that a substantial share of the dispersion in export prices in the data is driven by exporters selling a given product at different prices to different buyers. We call such individual price behavior “price discrimination” and study it in more details in sections 4 and 5.

3.1. Price dispersion in the EU

We start by documenting that the product-level prices of French exports to the EU are highly dispersed. A third of this dispersion is driven by individual sellers charging their buyers different prices, whereas two thirds of the dispersion is attributable to average price differences across sellers.

¹⁵ When working on the later 2012–2016 period, we also neglect transactions involving importers in Romania, Bulgaria, and Croatia because these countries joined the EU only recently.

3.1.1. Methodology

Hereafter, the object of interest is the dispersion of prices, within a 8-digit level product category, which we measure as:

$$Var_{pt}^{scb(c)}(p_{sb(c)pt}) = \frac{1}{N_{pt}-1} \sum_s \sum_c \sum_{b(c)} (p_{sb(c)pt} - \bar{p}_{pt}^{scb(c)})^2$$

where \bar{p} and $Var(p)$, respectively, refer to the first and second moments of the cross-section of prices and N_{pt} is the number of price quotes in the corresponding cross-section. Subscripts refer to the dimensionality of the corresponding variable, whereas superscripts denote the dimension in which the corresponding moment is calculated. $Var_{pt}^{scb(c)}(p_{sb(c)pt})$ thus denotes the variance of prices computed across sellers, buyers, and countries, for a particular product×period. Because the variance is increasing in the average level of prices, we use a normalized measure of dispersion, namely, the coefficient of variation:

$$CV_{pt}^{scb(c)}(p_{sb(c)pt}) = \frac{\sqrt{Var_{pt}^{scb(c)}(p_{sb(c)pt})}}{\bar{p}_{pt}^{scb(c)}}$$

In examining what share of these price discrepancies is attributable to different exporters selling a given product at different prices versus exporters price discriminating their partners in the EU, we further decompose the dispersion of prices into a *within* and a *between* components. The *within* component is a weighted average of the variance of prices within an exporter s , and the *between* component measures variations in mean prices between exporters of the same good. Formally,

$$Var_{pt}^{scb(c)}(p_{sb(c)pt}) = \underbrace{\sum_s \frac{N_{spt}-1}{N_{pt}-1} Var_{spt}^{b(c)}(p_{sb(c)pt})}_{Within} + \underbrace{wVar_{pt}^s(\bar{p}_{spt}^{b(c)})}_{Between} \quad (1)$$

where N_{spt} is the number of buyers connected to seller s , $Var_{spt}^{b(c)}(p_{sb(c)pt})$ is the variance of prices that this exporter sets on transactions with different partners,

$$Var_{spt}^{cb(c)}(p_{sb(c)pt}) = \frac{1}{N_{spt}-1} \sum_c \sum_{b(c)} (p_{sb(c)pt} - \bar{p}_{spt}^{cb(c)})^2$$

and

$$wVar_{pt}^s(\bar{p}_{spt}^{b(c)}) = \sum_s \frac{N_{spt}-1}{N_{pt}-1} (\bar{p}_{spt}^{b(c)} - \bar{p}_{pt}^{scb(c)})^2$$

is the variance of exporter-specific average prices. The ratio of the *within* component over the overall variance of prices is thus a measure of how much the cross-sectional dispersion of prices within a product is attributable to price discrepancies within a seller.

3.1.2. Results

Fig. 1 shows the evolution over time of the average coefficient of variation, using various country samples, namely, the EU25, the subset of countries that were already members of the EU in 2002 (EU15), the 11 original EMU countries, and the three members of the EU that do not participate in the common currency. The top panel corresponds to the 2002–2006 period, and the bottom one is for 2012–2016. The dispersion of prices is relatively stable over each subperiod, but price dispersion is higher in 2012–2016 than over 2002–2006. As expected, the coefficient of price variations is, on average, lower in the EMU subsample than in the whole EU. But the most striking difference is observed in the subsample of non-EMU members, in which the recovered coefficient of variation is an order of magnitude lower. This result is in part mechanical because this sample consists of three countries (the UK, Denmark and Sweden) that are not the most popular destinations for French exports. We further dig into this result in section 4.1, when

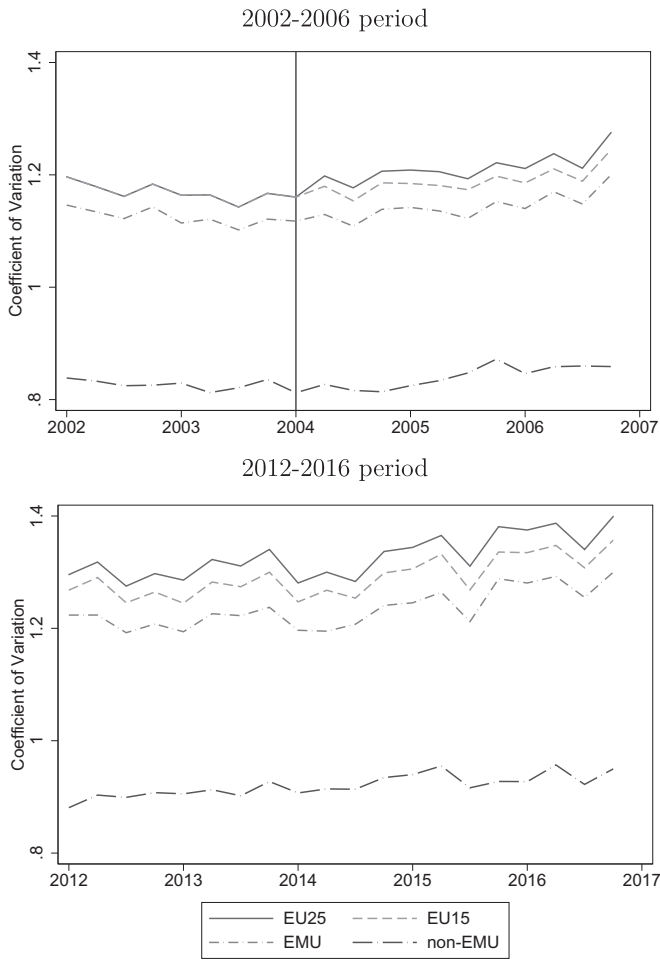


Fig. 1. Evolution of the mean dispersion of prices, over time. Notes: This figure plots the evolution of the mean coefficient of variation of prices, computed for each product and quarter, that is, the mean across products of: $CV_{pt}^{cb(c)}(p_{sb(c)pt}) = \frac{\sqrt{Var_{pt}^{cb(c)}(p_{sb(c)pt})}}{\bar{p}_{pt}^{cb(c)}}$ using the notations in Eq. (1). Coefficients of variation are computed across exporters, countries, and importers, in the whole sample (“EU25”) and in three subsamples restricted to EU15, EMU, and non-EMU EU15 member countries. The vertical line in the top panel corresponds to the EU enlargement.

the analysis is restricted to the dispersion of prices within a seller and we can compare price dispersion across geographical areas, conditional on a number of partners.

Table 2 provides further details on the distribution of the product- and period-specific coefficients of variation recovered from data covering the 2002–2006 period.¹⁶ The level of price dispersion varies substantially among products and quarters. The mean coefficient of variation is thus equal to 1.3 in the EU25, but the median is substantially lower, at 0.8, which indicates the distribution in the dispersion is skewed to the right. Some product quarters display much higher price dispersion than the median. This heterogeneity is to a large extent driven by the product dimension, because 8-digit product categories cover a large array of different types of goods, some much more differentiated than others, thus offering more ground for price discrepancies.

Unobserved heterogeneity across sellers (and the varieties they produce), the *between* component in Eq. (1), is a key driver of the dispersion of prices in French export markets. But the dispersion of prices *within* a particular seller, our measure of price discrimination, is also

¹⁶ Statistics are based on the subsample of coefficients of variation computed on at least five price quotes. Results are qualitatively similar over 2012–2016.

substantial in the bottom panel of Table 2. The contribution of the within-seller component to the overall price dispersion is around 30%, on average, regardless of the sample of countries considered in columns (1)–(4).

Note that these figures tend to underestimate the extent of price dispersion within sellers across buyers. Indeed, at the 8-digit product and quarter level, almost 50% of sellers (\times period) display zero within price dispersion, because they serve a single client in the EU. Although we later drop these firms from the analysis of price discrimination within a firm, they contribute to the overall dispersion of prices, within a product. As such, we include them in the statistics of Table 2. By definition, they solely contribute to the between-firm component, thus mechanically inflating its contribution to the overall dispersion.

3.2. Price discrimination within the EU

This section focuses on sellers' propensity to price discriminate, as measured by the dispersion of prices they charge their buyers for a given product and period. Price discrimination may arise due to firms charging different markups for the exact same variety sold to different buyers, or it may be a consequence of sellers selling different varieties of a given product or different variety bundles to their buyers. For now, we refer to both strategies as price discrimination.¹⁷ We come back on the origin of such price discrimination in section 5. Price discrimination is a common practice among French exporters. The coefficient of variation of prices computed within a seller based on prices set to buyers located in the EMU is equal to 35.7%, on average. Despite the within-exporter dispersion of prices being substantial, on average, we do observe a subsample of firms charging close to uniform prices in the EMU. These firms are relatively small, on average, and thus do not contribute much to aggregate exports.

3.2.1. Methodology

To study price discrimination, we now focus on sellers connected to multiple partners in a given period. The coefficient of variation of prices, measured within a seller, product, and time period, reads

$$CV_{spt}^{cb(c)} \equiv \frac{\sqrt{Var_{spt}^{cb(c)}(p_{sb(c)pt})}}{\bar{p}_{spt}^{cb(c)}} \quad (2)$$

$$\text{with } Var_{spt}^{cb(c)}(p_{sb(c)pt}) = \frac{1}{N_{spt}-1} \sum_c \sum_{b(c)} (p_{sb(c)pt} - \bar{p}_{spt}^{cb(c)})^2$$

where $Var_{spt}^{cb(c)}(p_{sb(c)pt})$ is the variance of prices set by seller s , computed across all transactions with partners located in various countries, N_{spt} is the number of partners the firm is connected to in this particular time period, and $\bar{p}_{spt}^{cb(c)}$ is the mean price of its export transactions. These statistics are defined for the 50% of French exporters that interact with at least two European importers within a given quarter. In what follows, the analysis is further restricted to firms serving at least five partners in a given quarter, around 40% of the overall distribution. We remove this restriction once the analysis can control for the number of partners involved into the corresponding cross-section.

3.2.2. Results

Fig. 2 represents the distribution of coefficients of variation. The distribution is bi-modal. About 10% of product-seller pairs have a coefficient of variation below 1%. The corresponding firms do not discriminate across partners, within a product. We examine this extreme form of “uniform pricing” in more details in the next paragraph. The rest of the distribution exhibits a substantial level of price

¹⁷ Firms may indeed offer a menu of packages or bundles (in terms of prices, quantity, and/or quality) to consumers to price discriminate. This is often referred to as second-degree price discrimination (Tirole, 1988).

Table 2
Summary statistics on the coefficient of variation, within a product, and quarter.

	EU25	EU15	EMU	non-EMU
	(1)	(2)	(3)	(4)
	Coefficient of variation $CV_{pt}^{scb(c)}(p_{sb(c)pt})$			
Mean	1.293	1.280	1.236	0.998
Median	0.790	0.785	0.774	0.722
10th percentile	0.331	0.330	0.325	0.312
90th percentile	2.753	2.720	2.595	2.007
	Contribution dispersion within s			
Mean	0.312	0.312	0.304	0.286
Median	0.265	0.264	0.253	0.203
10th percentile	0.006	0.006	0.005	0.002
90th percentile	0.704	0.705	0.696	0.725
Count obs	142,266	141,521	138,354	75,669

Notes: The table reports summary statistics on the dispersion of prices within a product and period, in various country sub-samples. The first panel reports statistics on the distribution of coefficients of variations:

$$CV_{pt}^{scb(c)}(p_{sb(c)pt}) = \frac{\sqrt{\text{Var}_{pt}^{scb(c)}(p_{sb(c)pt})}}{\bar{p}_{pt}^{scb(c)}}$$

using the notations in Eq. (1). The second panel corresponds to the contribution of the within component also described in this equation. Statistics are based on the 2002–2006 period and are restricted to coefficients of variation recovered from at least five points.

discrimination, with a mode around 25% and a distribution that is skewed to the right.

Table 3 provides additional summary statistics over the distribution of the coefficients of variation of prices, measured within a seller, product, and time period for different samples of countries. Column (1) is based on the whole country sample, as is Fig. 2. Columns (2), (3), and (4) are then restricted to importers located in the EU15, the EMU, and the non-EMU members of the EU, respectively. As expected, the level of price dispersion is lower once we focus on the within-exporter dimension. However, price discrepancies are still quantitatively important, the standard deviation being slightly above 35% of the mean level of prices, on average. This level of dispersion is in the range of what Kaplan and Menzio (2015) find for consumer goods sold in various US stores. Restricting the country sample to increasingly integrated markets as we do from column (1) to column (3) implies a distribution




Fig. 2. Distribution of coefficients of variations, across exporters, product and time. Notes: This figure plots the distribution of variation coefficients, computed for each exporter×product×period according to Eq. (2):

$$CV_{spt}^{cb(c)}(p_{sb(c)pt}) = \frac{\sqrt{\text{Var}_{spt}^{cb(c)}(p_{sb(c)pt})}}{\bar{p}_{spt}^{cb(c)}}$$

The analysis is restricted to statistics computed on buyers in the EU25, over 2002–2006 and based on at least five points.

that is slightly shifted to the left. This finding is consistent with the view that price discrepancies are reduced in more integrated markets. Instead, the distribution recovered from the non-EMU countries is slightly shifted to the right, thus suggesting that firms exporting to these destinations tend to set more dispersed prices.¹⁸

3.2.3. Uniform pricing

Uniform pricing has recently attracted new scrutiny as several papers have documented that large retailers tend to adopt such strategy in the US market, and in the euro area (Cavallo et al., 2014; DellaVigna and Gentzkow, 2017). The mass around zero in Fig. 2 shows a subsample of exporters that do adopt such pricing strategies. We now study this population in more details.¹⁹

We follow DellaVigna and Gentzkow (2017) and compute a measure of nearly uniform pricing based on close to zero coefficients of variation. Namely, we define a dummy variable that is equal to 1 for firms adopting near uniform pricing strategies:

$$NUP_{spt} = 1 [CV_{spt}^{cb(c)} < .01] \tag{3}$$

Using that measure, we examine the prevalence of near uniform pricing within the euro area and within euro countries. We report the results in Fig. 3. Over the period 2002–2006, about 14% of the products exported by French firms to their euro-area buyers are priced nearly uniformly (solid line, top panel, left-hand side) and almost 10% are priced nearly uniformly while exported to at least two EMU destinations. This finding is a significant result that has not been documented using trade-price data to our knowledge. The prevalence of near uniform pricing is twice as low, around 7%, but still significant over 2012–2016 (top panel, right-hand side).

Whereas a significant share of products is priced uniformly, they account for a more modest share of trade in value terms. The dotted lines in Fig. 3 show that NUP weights about 2% of French exports toward euro trade partners in 2002–2006 and a lower 1.7% in the more recent period.²⁰

We further consider the possibility that firms choose to price uniformly within a market but not across destinations, which would be consistent with zone pricing (Adams and Williams, 2019).²¹ The prevalence of NUP within EMU countries is summarized in the bottom panels of Fig. 3. Over 2002–2006, 13% of varieties exported by French firms in a given destination are priced uniformly. This amount is about the same magnitude as the prevalence observed in the whole euro area. However, the economic weight of NUP within a destination is about twice as large as in the overall euro area, at almost 4%. Here as well, the prevalence of NUP seems to decrease over time, in terms of both frequency and exported value.

¹⁸ We have also computed the same statistics based on the EMU less Belgium and Netherlands. Trade with these countries is somewhat biased by the presence of major trade ports in Anvers and Rotterdam. A substantial share of trade between the EU and the rest of the world indeed transits through these ports. In principle, these trade flows have been excluded from the dataset as we dropped trade flows that are reported to be exported to destinations outside of the EU but exit France through another EU country. However, a significant number of trade flows could still be intermediated by firms in these countries, in which case the flow is recorded as an intra-EU transaction. The comparison of EMU results with numbers recovered from this restricted sample were very similar. We concluded from this finding that exports to these countries do not bias our results.

¹⁹ Remember that our price data are not especially well-suited to study uniform pricing since the differentiation of the underlying varieties cannot be ruled out. This can explain that the prevalence of uniform pricing in our data is an order of magnitude smaller than in Cavallo et al. (2014) and DellaVigna and Gentzkow (2017).

²⁰ Part of the discrepancy is explained by the largest firms being relatively less likely to adopt near uniform pricing. When we perform the analysis conditional on the number of partners served in the EMU, we observe that NUP is relatively more prevalent among firms serving two to four buyers in the EMU, which are not the largest ones.

²¹ Such zone pricing has also been documented to some extent by Cavallo et al. (2014), who show that Zara has a pricing strategy specific to Spain and Portugal on the one hand, and other euro countries on the other hand.

Table 3
Summary statistics on the coefficient of variation, within a seller, product, and quarter.

	EU25	EU15	EMU	non-EMU
	(1)	(2)	(3)	(4)
	Coefficient of variation $CV_{spt}^{cb(c)}(p_{sb(c)pt})$			
Mean	0.364	0.362	0.357	0.365
Median	0.314	0.311	0.305	0.307
10th percentile	0.027	0.026	0.027	0.058
90th percentile	0.761	0.758	0.749	0.760
	Contribution dispersion within c			
Mean	0.506	0.526	0.580	0.839
Median	0.551	0.581	0.656	0.941
10th percentile	0.000	0.001	0.020	0.505
90th percentile	0.950	0.957	0.979	1.000
Count Obs	863,275	835,386	716,780	104,410

Notes: The table reports summary statistics on the dispersion of prices within an exporter, product, and period, in various country samples. Price dispersion is measured as:

$$CV_{spt}^{cb(c)}(p_{sb(c)pt}) = \frac{\sqrt{\text{Var}_{spt}^{cb(c)}(p_{sb(c)pt})}}{\bar{p}_{spt}^{cb(c)}}$$

using the notations in Eq. (2). The second panel corresponds to the contribution of the within component described in Eq. (4). The period of analysis is 2002–2006. Statistics are computed on the distribution of variation coefficients recovered from at least five points.

4. Heterogeneity in the level of price discrimination

This section investigates the heterogeneity in the level of price discrimination across markets, sectors, firms, and over time. It shows the level of price discrimination is substantial within the EMU - and within EMU destinations - but remains lower than outside the EMU. We further show the level of price discrimination has increased over time, mostly driven by the behavior of small firms. Large firms and firms active in more differentiated sectors are more likely to price discriminate, whereas retailers and wholesalers charge less dispersed prices.

4.1. Price discrimination across markets

In this section, we study the extent of price discrimination within a firm and across various geographical areas. In particular, we study the extent to which price discrimination is lower in the EMU than in the rest of the EU. One of the expected benefits of the monetary union is indeed the convergence of prices, through arbitrage. Such arbitrage should limit firms' ability to price discriminate. Our data also enables us to compare the extent of price discrepancies within a country and across countries. Here as well, the comparison is insightful inasmuch as we think of countries as relatively well-integrated geographical areas that should thus display less dispersion in prices than larger geographical units.

4.1.1. Price discrimination within vs outside the EMU

We start by documenting that price discrimination is lower within the common currency area. To do so, we construct a panel of coefficients of variation in which each observation is identified by the exporting firm, the product being exported, and the period of analysis (the *spt* triplet), and the country sample over which price discrepancies are recovered (either EU25, EU15, or EMU). We regress these measures of dispersion on dummies indicating the geographical area considered and exporter×product×period fixed effects. The coefficients on the dummies thus measure the extent to which price discrimination varies within a firm, across various geographical areas. Results are presented in Table 4.

First, consider columns (1) and (2), which compare the mean dispersion of prices across country sub-samples. Consistent with Table 3, results show that price discrepancies are, on average, larger in the complete sample than in the sample restricted to the 15 historical members of the EU, whereas they are lower in the EMU than in the EU15. This

observation is true in both periods, although the difference between the EU25 and the EU15 is not statistically different from zero over 2002–2006.²² Mean differences across country samples within a firm are quantitatively important because prices within the EU25 are, on average, 10% more dispersed than within the EU15, whereas they are 14% less dispersed in the EMU than in the EU15. By construction, the coefficients of variation are computed using a larger number of observations for the EU samples than for the EMU one. Columns (3) and (4) include the number of buyers as a control, to ensure that the differences across geographic areas are not mechanically driven by such differences in the dimensionality of the underlying variables. Results show they are not because the coefficient on the EMU dummy continues to be significantly negative and of the same order of magnitude once we control for the number of buyers.

To further assess the robustness of this result to potential composition effects, we do an additional exercise. The idea is to fix the number of buyers per seller, and compare the level of discrimination among EMU buyers relative to the level of discrimination between buyers located both within and outside the EMU. More specifically, we restrict the sample to firms (×product×period) serving at least three partners in the EMU and at least one partner outside of the EMU (but within the EU15).²³ We then compute a measure of price dispersion within the EMU based on a random sample of three price quotes, drawn from the firm-specific portfolio of EMU partners. This statistics is then compared with a measure of price dispersion in the EU15, recovered from two random draws from the firm's EMU partners and one random draw from its non-EMU partners. This exercise amounts to comparing the dispersion of prices within a firm (×product×period), within and outside of the EMU, conditional on a fixed number of partners. Fig. 4 presents the distribution of the coefficients of variation recovered from the exercise. As expected, the distribution recovered from EMU buyers is denser at low levels of the coefficient of variation. This finding is consistent with price discrimination being lower within a firm, and within rather than outside of the EMU. Note, however, the difference is not very pronounced, with the difference in means being equal to two percentage points.

4.1.2. Price discrimination within and across countries

We have shown that price discrimination is lower within the EMU than in the EU15, but remains substantial at 35.7%, on average. The data offer a unique opportunity to dig deeper into the importance of administrative borders for price discrimination because they allow us to compare prices set by the same exporting firm over different partners located in the same country. Using the same strategy as in the previous paragraph, we now compare the extent of price discrimination within a country and across countries, within the EMU.

Following the same logic as in Eq. (1), price discrepancies within an exporter can further be decomposed into *within* and *between* components according to:

$$\text{Var}_{spt}^{cb(c)}(p_{sb(c)pt}) = \underbrace{\sum_c \frac{N_{sptc}-1}{N_{spt}-1} \text{Var}_{spt}^{b(c)}(p_{sb(c)pt})}_{\text{Within}} + \underbrace{w \text{Var}_c^c(\bar{p}_{spt}^{b(c)})}_{\text{Between}} \quad (4)$$

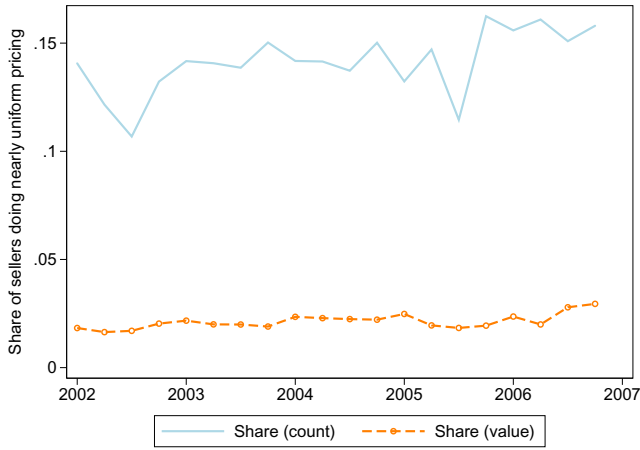
where $\text{Var}_{spt}^{b(c)}(p_{sb(c)pt})$ is the variance of prices that this exporter sets on transactions with these buyers, N_{sptc} is the number of buyers connected to seller *s* in country *c*, and

²² One possible reason for the lack of significance of the EU25 dummy in column (1) is the size of the sample used to identify the coefficient. Our dataset does not cover bilateral data prior to countries' entry into the EU. Therefore, the coefficient on the EU25 dummy for the 2002–2006 period is de facto identified over observations recovered from 2004 to 2006 data.

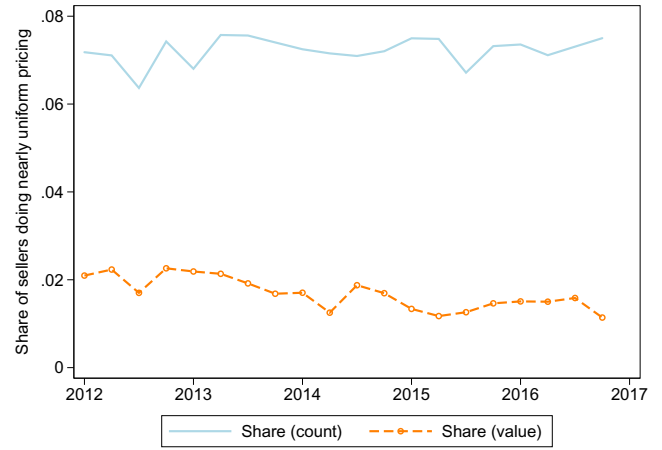
²³ This restriction reveals itself to be quite demanding, because it reduces the population of firms to 34% of the overall sample.

Within the EMU

2002-2006



2012-2016



Within EMU countries

2002-2006



2012-2016




Fig. 3. Near uniform pricing within the EMU. Notes: This figure reports the share of near uniform pricing within the euro area. Near uniform pricing is defined in Eq. (3). The top panels report the prevalence of NUP within the EMU; two thirds of the firms doing NUP are selling their product to more than one destination within the EMU. The bottom panels report the prevalence of NUP within EMU destinations.

$$wVar_{spt}^c(\bar{p}_{spt}^{b(c)}) = \sum_c \frac{N_{spt} - 1}{N_{spt} - 1} (\bar{p}_{spt}^{b(c)} - \bar{p}_{spt}^{cb(c)})^2$$

is the variance of mean prices set by seller s , across destination countries. The *within* component in Eq. (4) thus captures what is attributable to the seller price discriminating across buyers within a destination country. The *between* component instead measures discrepancies in mean prices across destinations, that is the PTM component. The decomposition is calculated for each firm serving at least two partners in the EU, the *within* (respectively *between*) component being mechanically equal to zero if the firm serves a single buyer within each destination (respectively, a single destination, but at least two partners there).

The second panel in Table 3 provides statistics over the contribution of the *within* component to the overall dispersion of prices set by an exporter. On average, in the EU25, half of the price dispersion is

attributable to exporters setting different prices on their different partners located in the same EU country. The remaining 50% of the dispersion is due to the firm applying different mean prices across destinations, and in particular across EMU and non-EMU destinations. Note the contribution of the *within* component naturally increases when the analysis is restricted to smaller country samples, but this is just the consequence of the *between* term being computed over a smaller cross-section. Within the EMU, 58% of the variance of seller-specific prices is observed within a country.

Although these numbers indicate the average contribution of the *within* and *between* components of price discrimination in the data, they hide a substantial amount of heterogeneity. In particular, differences between firms that mostly export to a single destination and firms that serve few buyers in many different destinations can mechanically induce a substantial dispersion. To control for this heterogeneity, we again rely on randomization. Namely, we retrain the

Table 4
Dispersion of prices within a seller across EU markets.

	Dep. Var: \ln Coefficient of variation $CV_{szpt}^{b(c)}(p_{sb(c)pt})$			
	2002–2006	2012–2016	2002–2006	2012–2016
	(1)	(2)	(3)	(4)
EU25 dummy	0.082 (0.049)	0.104 ^c (0.035)	0.091 (0.047)	0.112 ^c (0.035)
EMU dummy	-0.140 ^a (0.013)	-0.141 ^b (0.033)	-0.151 ^a (0.012)	-0.150 ^b (0.033)
Count partners			-0.003 ^b (0.000)	-0.002 ^b (0.000)
# Observations	2,063,813	2,742,399	2,063,813	2,742,399
Fixed Effects	Product×period×seller			
# FE	961,105	1,270,168	961,105	1,270,168
Adjusted R ² (overall)	0.930	0.874	0.930	0.875

Notes: The LHS variable is the log of the coefficient of variation, across buyers within a seller (for each product and quarter), calculated for a specific geographical area z , the EU25, the EU15, or the EMU. The sample thus has a maximum of three observations per firm×product×period. We further impose a minimum of two observations; that is, the sample is restricted to firms that serve different partners in at least two zones in a given period. “Count partners” is the number of buyers served by the firm in the corresponding area.

dataset to the 10% of firms that export to at least three partners in three different destinations. In this subsample, we randomize to recover three price quotes per destination and three price quotes in three different destinations. Based on this subsample, we can compute measures of price dispersion within and across destinations that are fully comparable within a firm, because they are based on the same number of price quotes, with the location of the partners being the only source of variation. The distributions recovered from the two samples are reported in Fig. 5.

Within a firm (×product×period), the dispersion of prices charged on three random partners located in a given destination is 22% lower than the corresponding statistics based on three random partners located in three different countries, on average. Market segmentation is thus a significant source of price discrepancies, including within the EMU. However, within-country price discrepancies are also substantial, in this sub-sample as in the overall population of French exporters. For




Fig. 4. Price discrimination within and outside the EMU: randomized sample. Notes: This figure shows the distribution of the coefficients of variation in Eq. (2) recovered from the randomization of prices within and outside of the EMU. These statistics are based on the sub-sample of firms exporting to at least three partners in the EMU and at least one partner outside of the EMU (but within the EU15) in a given month over 2002–2006. Coefficients of variation plotted in the “EMU” dispersion are calculated from three random price quotes within the EMU in the seller’s portfolio. The “non-EMU” dispersion is based on two random draws in the EMU and one random draw outside of the EMU. Horizontal lines materialize the means.




Fig. 5. Price discrimination within and across countries: Randomized sample. Notes: This figure shows the distribution of the coefficients of variation in Eq. (2) recovered from the randomization of prices within and across EMU destinations. These statistics are based on the subsample of firms exporting to at least three partners in three EMU countries in a given month over 2002–2006. Coefficients of variation plotted in the “Between” dispersion are calculated from three random price quotes in three different countries in the seller’s portfolio. The “within” dispersion is based on three random price quotes in each destination. Horizontal lines represent the means.

the mean firm in the considered population, the within-country dispersion recovered from the randomization is indeed as high as 25% of the mean unit value, only somewhat below the coefficient of 36% found for the average firm in the overall EMU. This finding suggests price discrimination is an important feature of the data, including within a destination market.

4.2. Price discrimination over time

After having analyzed the geography of price discrimination within and outside the EMU, this section focuses on the time dimension. We show price discrimination has increased over time within the EMU. This trend has been mainly driven by the pricing behavior of small exporters. This result goes against the common view that both the increasing integration of European markets and new communication technologies should allow consumers to arbitrage more easily across goods, which is expected to force the convergence of prices. Instead, the increasing dispersion observed within an exporter, over time, suggests exporters in our sample manage to maintain high price discrepancies, potentially thanks to product differentiation.

These results are recovered using a panel of coefficients of variations, computed for each seller×product×period across buyers in the EMU. We restrict the analysis to the subsample of firm-product pairs observed at least once over 2002–2006 and once over 2012–2016. This sample represents 20% of the overall population of exporters. Based on this panel, measuring the evolution of price discrimination over time, *within a particular seller*, is possible.²⁴ Results are presented in Table 5, top panel. On top of the fixed effects, the list of regressors includes a post-2012 dummy. The coefficient estimated on this variable measures the average difference in the level of price dispersion, between the first and the second periods. In column (1), the coefficient is positive but not precisely measured. Column (2) further controls for the number of

²⁴ Importantly, the exercise is conducted within a firm. Doing so implies that results are immune from composition effects related to the change in the declaration threshold, between the first and second periods of analysis. As explained in Section 2, changes in customs procedures imply that the population of firms covered by the 2012–2016 period is smaller than the population of exporters covered in the 2002–2006 sample. Results in this paragraph are based on the population of firms whose sales are above both declaration thresholds.

Table 5
Dispersion of EMU prices, within a seller: over time.

	All buyers				3 random buyers		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Dep. Var: In Coefficient of variation $CV_{spt}^{b(c)}(p_{sb(c)pt})$						
Post 2012 dummy	0.254 (0.055)	0.262 ^c (0.028)	0.286 ^c (0.029)	0.345 ^c (0.033)	0.318 ^b (0.021)	0.356 ^b (0.020)	0.392 ^b (0.025)
Post 2012 ×High Mket Power			-0.054 ^b (0.003)			-0.121 ^b (0.005)	
Post 2012 ×large firm				-0.360 ^b (0.019)			-0.308 ^b (0.022)
In count partners		0.368 ^b (0.002)	0.374 ^a (0.006)	0.370 ^a (0.006)			
# Observations	1,345,036	1,345,036	1,151,259	1,297,404	977,945	836,402	942,559
Fixed effects	Seller×product						
# FE	64,501	64,501	54,945	62,125	46,608	39,698	44,871
Adjusted R ²	0.420	0.435	0.434	0.435	0.730	0.726	0.728
	Dep. Var: In Coefficient of variation $CV_{spt}^{b(c)}(p_{sb(c)pt})$						
Post 2012 dummy	0.225 (0.055)	0.238 ^c (0.029)	0.258 ^c (0.030)	0.290 ^c (0.032)	0.280 ^b (0.017)	0.311 ^b (0.022)	0.322 ^b (0.023)
Post 2012 ×High Mket Power			-0.040 ^b (0.002)			-0.102 ^c (0.008)	
Post 2012 ×Large Firm				-0.254 ^b (0.013)			-0.191 ^c (0.020)
In Count Partners		0.355 ^b (0.009)	0.357 ^b (0.010)	0.359 ^b (0.009)			
# Observations	1,444,353	1,444,353	1,256,483	1,390,426	817,549	717,462	787,302
Fixed Effects	seller×country×product						
# FE	86,421	86,421	74,508	83,081	48,465	42,201	46,655
Adjusted R ² (overall)	0.391	0.403	0.404	0.402	0.737	0.733	0.734

Notes: The LHS variable is the log of the coefficient of variation, across buyers within a seller (for each product and quarter) in the EMU (top panel) or in a particular country of the EMU (bottom panel). The panel is restricted to coefficients of variation that can be compared in both subperiods, within a particular individual, that is, within a seller×product in the top panel and within a seller×product×country in the bottom panel. The first four columns are based on the overall sample. The last three columns report results obtained from the randomization of three buyers per individual and quarter. Standard errors are clustered by sub-period.

buyers underlying the measure of price dispersion. With this additional control, the coefficient on the post-2012 dummy turns significant at the 10% level and is positive, suggesting the level of price discrimination has increased over time.

The weak precision of the estimates suggests possible heterogeneity in the evolution of price discrimination in the data. Regressions in columns (3) and (4) include an interaction between the Post 2012 dummy and measures of firms' size and profitability. Size is measured by the total value of the firm's sales (domestic and export) sales. Profitability is recovered from the ratio of the firm's gross operating surplus over its value added. Both dummy variables are equal to 1 for firms belonging to the top of the distribution, within their sector of activity.²⁵ Coefficients on both interactions are negative and significant, suggesting the increase in within-firm price discrimination has been relatively less pronounced for large and high profitability firms. Price discrimination is actually stable over time for the top four firms in each sector (column (4)).

To rule out the possibility that results are driven by composition effects, we again rely on randomization. Regressions in columns (2)–(4) are reproduced based on coefficients of variation computed for each firm×product×quarter, using a random sample of three EMU partners. Results reported in columns (5)–(7) are very similar, both qualitatively and quantitatively. In particular, we confirm the heterogeneity along the distribution of firms, with the average increase in the level of price discrimination being mostly driven by relatively small firms.

Finally, the bottom panel of Table 5 reproduces the same regressions, using the dispersion of prices within a particular EMU country as left-

hand-side variable, and identifying the variation over time within a firm×country×product. All results are qualitatively the same. This finding implies the overall increase in the level of price dispersion observed within an exporter is not triggered by more dispersion across EMU destinations but indeed by a stronger degree of price discrimination, including across buyers located in a given country.

4.3. Individual determinants of price discrimination

In the previous sections, we documented some heterogeneity in price discrimination across markets and over time, within a firm. This section takes another perspective and examines the individual characteristics driving the heterogeneity in price discrimination across sellers. We show the degree of price discrimination varies substantially across firms and products. Large, multi-product firms and producers of differentiated goods exhibit a higher degree of price discrimination. Price dispersion is instead lower for retailers and wholesalers.

4.3.1. Methodology

To measure the characteristics at the root of heterogeneity in firms' propensity to price discriminate, we rely on the following linear decomposition:

$$\ln CV_{spt}^{b(c)} = \beta X_{spt} + FE_s + FE_{pt} + e_{spt} \tag{5}$$

FE_s and FE_{pt} refer to seller-specific and product×period fixed effects, respectively. X_{spt} is a vector of time-varying controls. In a second step, we use the estimated fixed effects to study the determinants of price discrepancies that are (i) specific to an exporting firm and (ii) product-specific. Results of the first and second stages are presented in Tables 6–8. Details on the construction of the different explanatory variables can be found in Appendix A.

²⁵ The "High Profit Rate" dummy is equal to 1 for firms in the fourth quarter of the sectoral distribution of profit rates. The "Large Firm" dummy is equal to 1 for the four largest firms in their sector. Results are very similar, though more difficult to interpret quantitatively, when both size and profit rate are measured using a continuous variable, namely, the level of the variable for the firm, in relative terms with respect to the median firm in its sector.

Table 6
Determinants of the dispersion of EMU prices, within a seller: Seller-product determinants.

	Dep. Var: In Coef of var. $CV_{spt}^{b(c)}(p_{sb(c)pt})$			
	(1)	(2)	(3)	(4)
In count partners	0.321 ^a (0.002)	0.390 ^a (0.002)	0.417 ^a (0.002)	0.318 ^a (0.002)
In Size	0.059 ^a (0.001)			0.062 ^a (0.001)
Experience (product)		0.030 ^a (0.000)		0.027 ^a (0.000)
Core product			-0.081 ^a (0.004)	-0.169 ^a (0.004)
# Observations	1,945,787	1,931,140	1,945,787	1,931,140
Fixed effects	Product×period, Seller			
# pt	126,124	125,798	126,124	125,798
# s	42,614	42,241	42,614	42,241
Adjusted R ² (overall)	0.559	0.557	0.558	0.558
Within R ²	0.040	0.041	0.039	0.043

Notes: The LHS variable is the log of the coefficient of variation, across EMU buyers within a seller (for each product and quarter) as in Eq. (5). "In Count Partners" is the log of the number of price quotes used to compute the variance of prices. "In Size" is the log of the value of the seller's overall exports in the EU, during this particular period and for this particular product. "Experience (product)" is the number of years since the firm began exporting the product in the EMU. "Core Product" is a dummy equal to 1 if the product under consideration is the exporter's main source of export revenues. Standard deviations are clustered in the seller dimension. In column (4), product×period fixed effects explain 10% of the overall variance, and the contribution of seller fixed effects is equal to 44%.

Table 6 presents results of the first stage estimation. The model in Eq. (5) explains about 56% of the variance in the data, most of it being attributable to the fixed effects. Namely, heterogeneity in coefficients of variation across product×periods explains 10% of the variance, whereas as much as 44% is attributable to unobserved heterogeneity across sellers. This finding confirms the high degree of heterogeneity across exporting firms in terms of their propensity to price discriminate. Once we control for this heterogeneity, we find price dispersion is larger over large trade flows, even conditional on the number of partners served, but relatively less pronounced for the firm's core product. Finally, price dispersion tends to increase with the firm's experience as an exporter of the product. We dig deeper into this result in section 5.2.

Table 7
Determinants of the dispersion of EMU prices, within a seller: seller-specific determinants.

	Dep. var: seller fixed effect $\hat{F}E_s$				
	(1)	(2)	(3)	(4)	(5)
In Relative Sales	0.328 ^a (0.006)				0.259 ^a (0.008)
In Relative Market Power		0.065 ^a (0.015)			0.079 ^a (0.014)
Wholesaler			-0.067 ^a (0.024)		-0.098 ^a (0.028)
Retailer			-0.191 ^a (0.045)		-0.266 ^a (0.051)
In Count products				0.508 ^a (0.009)	0.387 ^a (0.011)
# Observations	35,094	28,044	39,231	42,244	27,851
Adjusted R ²	0.069	0.001	0.001	0.076	0.115

Notes: The LHS variable is the estimated seller fixed effect recovered from the estimation of Eq. (5). "In Relative Sales" is the (log of) the seller's turnover in 2006, normalized by the median firm's sales in the same sector. "In Relative Profit Rate" is a measure of the seller's relative profit rate, in comparison with the median firm in its sector, where a firm's profit rate is proxied by the ratio of gross operating surplus over value added, in 2006. "Wholesaler" and "Retailer" are dummy variables for sellers belonging to the wholesaling and retailing sectors, respectively. "In Count products" is the (log of) the number of products that the firm exports in the EMU.

Table 8
Determinants of the dispersion of EMU prices, within a seller: Product-specific determinants.

	Dep. Var: Product fixed effect $\hat{F}E_{pt}$					
	(1)	(2)	(3)	(4)	(5)	(6)
Durables dummy	0.631 ^a (0.009)					0.024 ^a (0.009)
Share of differentiated inputs		1.818 ^a (0.014)				1.273 ^a (0.020)
Upstreamness			-0.299 ^a (0.003)			-0.159 ^a (0.005)
Product complexity				0.156 ^a (0.003)		0.135 ^a (0.004)
Relationship stickiness					0.100 ^a (0.008)	0.015 ^c (0.009)
# Observations	125,798	108,696	124,342	112,431	125,787	102,173
Fixed effects	Period					
# t	20	20	20	20	20	20
Adjusted R ²	0.044	0.147	0.061	0.023	0.005	0.157

Notes: The LHS variable is the estimated product×period fixed effect recovered from the estimation of Eq. (5). The "Durables dummy" is an indicator variable that takes the value 1 for products classified as durables in the BEC classification (i.e., durable consumption goods and capital goods). "Share of differentiated inputs" is the percentage share of inputs used to produce the corresponding product that are classified as "differentiated" according to Rauch (1999) classification, which corresponds to the measure of "Input specificity" in Nunn (2007). "Upstreamness" measures the product's average position in value chains and is taken from Antras et al. (2012). "Product complexity" is a measure of the complexity of the product as measured by Hausmann and Hidalgo (2014). Finally, "Relationship stickiness" is the level of stickiness of relationships estimated for each HS6 product category in Martin et al. (2019). All regressions include a period fixed effect so that coefficients are identified across products, within a quarter.

4.3.2. Seller characteristics and price discrimination

Heterogeneity across exporters is investigated in Table 7, where the estimated seller fixed effects recovered from Eq. (5) are regressed on various firm-level variables. In column (1), the estimated fixed effects are regressed against the seller's sales (normalized by the median firm's sales in the industry). The (normalized) profit rate of the firm is included in column (2). The variable is measured by the ratio of gross operating surplus over value added, relative to the median firm in the sector. Column (3) introduces dummy variables, for Wholesalers and Retailers. The control group is the rest of the economy, mostly firms in manufacturing sectors. To control for heterogeneity between single- and multi-product firms, column (4) correlates the fixed effects with the number of products exported by the firm. Finally, column (5) corresponds to the multivariate specification.

In terms of explanatory power, the most important variables are the firm's size and the number of products it exports. Relatively large and multi-product firms tend to price discriminate more, conditional on the number of buyers served and the product being exported, because we use these variables as controls in the first stage (Table 6). Each of these variables explains about 7% of the variance in the fixed effects, and they are not fully redundant because the adjusted R-squared of the multivariate regression in column (5) is above 11%.

Less important in terms of explanatory power but also highly significant is the heterogeneity across firms with different degrees of profitability (column (2)). High profit rate firms are found to display larger price discrepancies, on average. Our results also point to heterogeneity across sectors with retailers and wholesalers being relatively less prone to price discrimination, on average (column (3)).²⁶ Here as well, the R-squared of the regression is low, meaning heterogeneity in pricing strategies is strong, even within these sectors. Not all retailers adopt near uniform pricing strategies. All of

²⁶ Retailers and wholesalers, respectively, represent 35% and 7% of firms in the sample.

these results remain the same in the multivariate regression of column (5).

4.3.3. Product characteristics and price discrimination

In Table 8 and Figs. 6 and 7, we characterize the heterogeneity across products, in the degree of price discrimination. Histograms in Figs. 6 and 7 report the mean level of price dispersion, across industries and broad economic categories, respectively.²⁷ In both cases, the ranking of products appears consistent with expectations. We find that prices are relatively less dispersed in industries producing relatively homogenous goods such as petroleum, food products, minerals, and some chemicals. The largest average levels of price discrepancies are instead found within highly differentiated industries, for example, Machineries and Professional equipments. We find the same ranking between primary and processed goods when products are classified according to the BEC classification in Fig. 7. The least dispersed prices are found for primary goods. At the other side of the distribution, the highest average coefficients of variation are obtained for durables and capital goods. This result is consistent with the view that the differentiation of products facilitates price discrimination. Engel and Wang (2011) argue that trade in durable goods is key to understanding the volatility and comovement of exports and imports in open macroeconomic models. Our findings further highlight the central role of durable goods for the level of price discrimination associated with exports.

This intuition is broadly confirmed by the multivariate regressions in Table 8. In columns (1) to (5), the product fixed effects estimated in Eq. (5) are regressed against measures of product durability, product differentiation, upstreamness, product complexity, and relationship stickiness. Column (1) confirms price discrimination is more pronounced in sectors classified as durables in the BEC classification. Column (2) also confirms prices are more dispersed across buyers for more differentiated products as measured by Nunn (2007). Column (3) shows less discrimination among products with a more upstream position in value chains, with upstreamness defined as in Antras et al. (2012).²⁸ Column (4) shows that firms selling more complex products (as defined by Hausmann and Hidalgo (2014)) tend to have a higher level of price discrimination. And column (5) implies a positive correlation between price discrimination and the level of stickiness in relationships (estimated by Martin et al. (2019) across HS6 product categories). These results are still valid in a multivariate regression including all the variables (column (6)). Together, these variables explain almost 16% of the dispersion in the level of price discrimination across HS6 product categories. Most of the dispersion is explained by Nunn's measure of input specificity confirming the role of product differentiation in firms' pricing strategies.

These results show not all firms and products are equally prone to price discrimination. Depending on their size, profitability, and the type of products they sell, French exporters appear strongly heterogeneous with respect to the dispersion of prices set on various buyers in the EMU. A corollary is that the mechanisms at the root of such price discrimination are heterogeneous. Price discrimination may sometimes reflect buyers' heterogeneity in terms of their willingness to pay, or come from differences in sellers' ability to differentiate products and serve

²⁷ Statistics in Figs. 6 and 7 are recovered from the following second-stage regression: $\hat{F}E_{pt} = Class_p + FE_t + e_{pt}$ where $Class_p$ is a full set of industry (Fig. 6) or BEC categories (Fig. 7). Because the left-hand side variable of the first stage is a log, numbers can be interpreted in percentage change from the omitted category in each figure.

²⁸ This observation may seem to contradict results in column (3) of Table 7 that price dispersion is less pronounced for retailers. By definition, retailers tend to sell goods at the end of value chains, which should thus display less dispersed prices according to column (3) in Table 8. The apparent contradiction comes from the reference used to interpret both coefficients. In Table 7, we study heterogeneity across sellers, conditionally on a product being exported. The low propensity of retailers to price discriminate must thus be understood in relative terms with respect to manufacturing firms selling the same product. Instead, the coefficient on upstreamness identified in Table 8 is identified from the comparison of different products that are positioned at various points of the value chain.

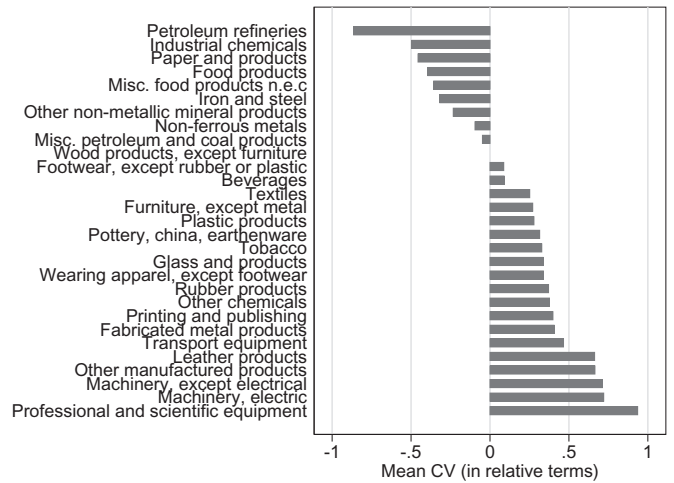


Fig. 6. Mean dispersion of prices, across industries. Notes: This figure reports the mean coefficient of variation per broad industry, in relative terms with respect to the sector producing wood products. These statistics are recovered by regressing the product×period fixed effects of Eq. (5) on a set of time and sector dummies. The sector dummies are reported on the graph. Since the left-hand side variable of Eq. (5) is in log, the y-axis can be interpreted in percentage terms.

different buyers with different products, or may be related to heterogeneity in the way the seller and its partner share the surplus. Until now, we have been agnostic about the underlying causes for such price discrimination. In Section 5, we provide some indicative evidence of the relative importance of these factors.

5. An exploration of the sources of price discrimination

This section explores into more details the mechanisms at the root of observed price discrimination strategies. Price discrimination occurs whenever consumers of a firm's product are heterogeneous in terms of their product valuation, or even with identical consumers in the presence of information frictions on the consumer or producer side (Stigler, 1961; Varian, 1980; Tirole, 1988). In such environments, sellers benefit from the fact that consumers cannot perfectly arbitrage to set prices between the competitive price and a maximum price where the buyer's




Fig. 7. Mean dispersion of prices, across BEC categories. Notes: This figure reports the mean coefficient of variation per category of the BEC classification. Results are relative to non-durable consumption goods. These statistics are recovered by regressing the product×period fixed effects of Eq. (5) on a set of time and BEC category dummies. The estimated coefficients on the BEC categories are reported on the graph. Because, the left-hand side variable of Eq. (5) is in log, the y-axis can be interpreted in percentage terms.

surplus is zero. The share of the total surplus that goes to a seller is a measure of its market power. Firms may price discriminate by offering different prices to different consumers for the same variety, or by offering different varieties or similar varieties with different packages (e.g., different volumes or different qualities) at different prices to all consumers.

Although our data do not allow us to measure heterogeneity in preferences, nor the amount of information available to firms while deciding on their prices, nor the exact packages sold by the firm, we propose two exercises to make progress on these questions. First, we restrict the analysis to a subsample of products that we argue are highly homogenous, so that we can confidently exclude the possibility that observed price discrepancies are explained by vertical differentiation. Second, we rely on a rich statistical decomposition of the variance of prices to discriminate between various models of price discrimination.

5.1. Price discrimination for homogenous products

In this section, we provide an indirect assessment of the importance of product differentiation as a source of price discrimination. To do so, we compute the level of price discrimination for a subset of products for which we expect very little room for vertical differentiation. This approach allows us to rule out that, in this sample, price discrepancies reflect a differentiation of products sold by the exporter to its partners in the EMU, and instead interpret the dispersion in terms of heterogeneous markups.

5.1.1. Methodology

We used a two-step method to select the sample of homogenous products. We started from 2446 product codes of sections V, VI, and VII of the combined nomenclature (Mineral Products, Products of the Chemical or Allied Industries, Plastics and Articles Thereof; Rubber and Articles Thereof). Among sections V and VII, we kept mineral, plastic, and rubber in their primary forms. Among chemical products, we excluded specialty chemicals such as pharmaceutical products, fertilizers, tannins, or pigments because these products are usually produced in low-volume and customized for the clients. We then excluded all product categories that explicitly mention that several varieties can enter in the description (e.g., Chlorides and chloride oxides) and those that include derivatives or esters of chemical compounds (e.g., Derivatives of acyclic hydrocarbons), leaving us with 402 CN8 product categories that we thought were potentially quite homogenous.

We then interviewed a chemist that went through the 402 product categories and classified them into eight groups: elements entering the Mendeleiev classification, molecular organic substances, molecular inorganic substances, essential oils, polymers, industrial inorganic products, industrial organic products, and petroleum products. Based on his assessment, we decided to exclude from the group of homogenous products essential oils and petroleum products, whose quality usually depends on the concentration of the product in the substance, as well as polymers and industrial inorganic and organic products that, contrary to the raw organic and inorganic substances, can be produced in various ways and qualities, leaving us with 276 product categories classified as homogenous. Examples of homogenous products include Fluorine (CN828,013,010), Calcium (CN8 28,051,200), Sulphuric acid (CN8 28,070,010), Solid potassium hydroxide (CN8 28,152,010), Trichloroethylene (CN8 29,032,200), or Octan-2-ol (CN8 29,051,620). The complete list of these products is available upon request.²⁹

Once the products were identified, we merged the list with the customs data to select the subsample of export transactions for these particular homogenous products. This selection leaves us with 7441 observations covering 276 product categories exported by 395 sellers.

For these products, unit values are defined in euros per kilogram. Price discrimination thus reflects differences in the price per kilogram of a given product, by a given seller across buyers.

5.1.2. Results

Fig. 8 compares the distribution of coefficients of variation recovered for firms (\times product \times quarter) in the overall sample (solid line) and in the restricted sample of homogenous products (dashed line). Compared with the overall distribution, the distribution based on homogenous products is shifted to the left. This finding is consistent with the view that homogenous products are less prone to price dispersion within a seller, because firms cannot price discriminate by differentiating the products they offer to their different partners. The mean dispersion of prices in the sample of homogenous goods is still substantial though, at 27% of the mean price.

Table 9 digs further into this difference by systematically comparing the level of price discrimination across different samples. Column (1) is the analog of Fig. 7. The coefficient of variation is explained by a dummy that is equal to 1 if the exported product is a homogenous chemical product, the control group being composed of the rest of the dataset. Price discrimination is 5.6 percentage points lower in the selected sample. Because all homogenous products belong to the chemical industry, the difference in means may be entirely explained by composition effects across industries. Column (2) is thus restricted to products in the chemical and allied industries. This sample includes all homogenous products but also other chemical products such as pharmaceuticals, specialty chemicals, or any chemical substance that was not sufficiently homogenous to pass the selection process. In this sample as well, we find prices of homogenous products significantly less dispersed than more differentiated products, with a difference of 16 percentage points. Finally, columns (3) and (4) in Table 9 further restrict the analysis to the sub-sample of firms that do export at least one homogenous product. Here as well, the coefficient estimated on the homogenous product dummy is significantly negative, at -10 percentage points, whether identified across or within a firm. Note the within specification in column (4) identifies the coefficient of interest solely on multi-product firms that export homogenous and non-homogenous chemicals. Controlling for unobserved heterogeneity across sellers, we can still confirm that prices of homogenous products are significantly less dispersed.

Although the external validity of these results is questionable, these results are indicative of composition effects within a seller being a factor at the root of some of the price discrepancies discussed in this paper. In the most demanding specification that allows us to compare




Fig. 8. Price discrimination among homogenous products. Notes: This figure reports the distribution of the coefficients of variation in Eq. (2), in the full sample and in the sample restricted to homogenous products. Horizontal lines represent that means.

²⁹ We are thankful to Luc Mejean for his valuable help at this stage of the procedure.

Table 9
Price discrimination: homogenous versus differentiated products.

	Dep. Var.: $CV_{sp}^{sb(c)}(p_{sb(c)pt})$			
	(1)	(2)	(3)	(4)
Average	0.323 ^a (0.000)	0.372 ^a (0.001)	0.369 ^a (0.001)	0.368 ^a (0.001)
Homogenous products	-0.056 ^a (0.003)	-0.105 ^a (0.001)	-0.101 ^a (0.004)	-0.097 ^a (0.004)
Sample	All	Chemicals & Allied indus.	Sellers exporting at least One homogenous product	
Fixed effects	no	no	no	seller×time
Observations	1,402,748	109,676	61,851	61,851
R ²	0.000	0.037	0.010	0.191

Notes: The table reports the estimated average coefficient of variation, in the overall sample and in the sub-sample of homogenous products (the coefficient being expressed in relative terms). The different columns correspond to various control groups: The rest of the economy in column (1), other chemical products in column (2), other products exported by firms that sell homogenous goods in columns (3)–(4).

homogenous and differentiated products sold by the same firm, price dispersion is 26% lower for homogenous products. Assuming that firms in the chemical industry are representative of the rest of the sample implies about a quarter of the price dispersion discussed in this paper is attributable to firms differentiating the products they offer to various partners in the EMU. Even within homogenous products such as molecular substances, price dispersion within a firm is found to be high, on average, at 27% of the mean price. Because this number is obtained within a firm, for a particularly homogenous product, it indicates exporters price discriminate across their partners by adjusting their markup.

5.2. Exporters' dynamic pricing strategies as a source of price dispersion

Interpreting firms' propensity to price discriminate requires us to dig deeper into their pricing strategies, in the cross-section and over time. In particular, we need to separate systematic differences *between* firms, coming from unobserved product differences, production technology, or market power, from differences *within* firms but *between* buyers if firms do not face the same marginal cost and/or do not apply the same mark-ups over the different partners served in a given period. For that purpose, we use a rich linear model, controlling for both seller and buyer fixed effects to analyze the determinants of price levels. Results are used to recover insights about the likely sources of the previously discussed price dispersion.

5.2.1. Methodology

The decomposition takes inspiration from the labor literature. Following Abowd et al. (1999), this literature has extensively used matched employer–employee data and high-dimensional fixed-effect estimators to identify the sources of the dispersion in wages observed in the data. Our dataset has the same bipartite graph structure, and we can thus rely on this methodology to decompose the observed variance of export prices.

The estimated model takes the following form:

$$\ln p_{sb(c)pt} = \beta \mathbf{X}_{sb(c)pt} + FE_s + FE_{b(c)} + e_{sb(c)pt} \tag{6}$$

where $\mathbf{X}_{sb(c)pt}$ is a set of control variables, FE_s is a (time-invariant) fixed effect for seller s , $FE_{b(c)}$ is a (time-invariant) fixed effect for buyer $b(c)$, and $e_{sb(c)pt}$ is a residual that captures the unexplained dispersion of prices within a seller–buyer match. As shown by Abowd et al. (1999), such an equation can be estimated on panel data to recover the contribution to the dispersion in prices of (i) unobserved heterogeneity across sellers absorbed into FE_s , (ii) unobserved heterogeneity across buyers absorbed into $FE_{b(c)}$, (iii) observable variables regarding the relationship involving seller s and buyer b $\mathbf{X}_{sb(c)pt}$, and (iv) a residual $e_{sb(c)pt}$ that is

specific to the seller–buyer relationship and the particular product and period under consideration.³⁰

Hence, the seller fixed effect captures both unobserved product heterogeneity between sellers and systematic differences in market power, with some sellers being able to set higher prices, on average. In the same way, the buyer fixed effect captures both the heterogeneity in good valuation among buyers, and differences in bargaining power, for example, because some buyers are new to the market and have limited information about the other sellers. Finally, the residual reflects any match-specific difference: a buyer experiencing particularly low or high prices with that seller or price fluctuations coming from idiosyncratic shocks.

The set of controls systematically includes a product×period effect that absorbs the mean price set by French firms in the product market for this particular quarter. These fixed effects control for the effect of inflation on prices and also absorb any difference induced by unit values being defined with respect to different physical quantities for some products. Identification of the buyer and seller fixed effects comes from the variance of prices across sellers and across buyers within a product×period. In comparison with similar statistical decompositions applied to matched employer–employee data, the parameters of this model are well identified. In a panel of almost 40 millions transactions, product×period fixed effects are identified, on average, on 226 price quotes (=37,470,412/165,730) while seller and buyer components are measured on, respectively, 30 and 2.3 price quotes per quarter, on average. This dimensionality renders the estimation of buyer and seller fixed effects more precise and alleviates the so-called limited mobility bias (Andrews et al., 2008).

5.2.2. Results

Table 10 reproduces the results based on 2002–2006, starting with the baseline equation that solely controls for fixed effects in column (1), before sequentially adding more controls. The simplest specification captures 89% of the variance of the data, with more than 65% solely attributable to the product×period fixed effects. The large contribution of the product×period fixed effects is not surprising given the diversity of products in the data that cover transactions over airplanes, wine bottles, car components, and so on. Besides the heterogeneity across products, the variance decomposition confirms the role of unobserved heterogeneity between exporters, which explains as much as 20% of the overall variance. In comparison, the contributions of the buyer unobserved heterogeneity and the residual are small, at 4.5 and 10.5%, respectively.

The bottom panel of the table reports the variance decomposition of the within-seller (×product×period) components of prices.³¹ These components are the ones that we are mostly interested in, because the objective is to understand price discrimination within a firm. In this dimension, the buyer fixed effect and the residual, respectively, account for 17% and 83% of the variance. The importance of buyers' unobserved heterogeneity as a determinant of sellers' pricing strategies can be interpreted in terms of the heterogeneity of buyers' preferences, which makes sellers optimally differentiate the product they sell and/or adjust their mark-up. Fig. 9 shows the correlation between the mean value of the buyer fixed effects and the GDP per capita of their country of origin. The correlation is strongly positive. This finding is consistent with the view that buyers in rich countries have higher

³⁰ This equation can be estimated whenever the underlying bipartite graph is connected, which is largely the case in our data. Namely, the largest connected component of the graph encompasses more than 99% of all observations. In the rest of the section, we neglect the remaining 1% of observations and estimate Eq. (6) on the largest connected component, which allows estimated fixed effects to be comparable.

³¹ In practice, the decomposition is based on the three-way fixed-effects estimation presented in Eq. (6). More specifically, the buyer and match components are regressed on normalized (log) prices. Prices are normalized in the product×period and seller dimensions. The estimated coefficients measure how much of the dispersion in prices within sellers and products is attributable to the buyer fixed effects and the residuals.

Table 10
Results of the fixed-effect decomposition of price dispersion.

	Dep.Var: ln price $p_{sb(c)pt}$					
	(1)	(2)	(3)	(4)	(5)	(6)
In Transaction size		0.036 ^a (0.000)	0.039 ^a (0.000)	0.036 ^a (0.000)	0.036 ^a (0.000)	0.036 ^a (0.000)
In Distance		0.008 ^a (0.001)	0.008 ^a (0.001)	0.009 ^a (0.001)	0.009 ^a (0.001)	0.016 ^a (0.002)
Seller's experience			0.002 ^a (0.000)	−0.000 (0.000)	0.000 (0.000)	−0.000 (0.000)
Age of the relationship			−0.009 ^a (0.000)	−0.010 ^a (0.000)	−0.009 ^a (0.000)	−0.009 ^a (0.000)
−× Large firm						0.003 ^a (0.000)
In count partners				0.008 ^a (0.000)	0.008 ^a (0.000)	0.008 ^a (0.000)
Dummy First time					0.002 ^a (0.000)	0.002 ^a (0.001)
−× Large firm						0.003 ^b (0.001)
Dummy One shot					0.008 ^a (0.000)	0.005 ^a (0.000)
−× Large firm						0.013 ^a (0.001)
# observations	37,470,412	35,143,089	35,143,089	35,143,089	35,143,089	33,490,640
Adj R ²	0.891	0.892	0.892	0.892	0.892	0.891
Within R ²	0.000	0.005	0.005	0.005	0.005	0.005
# Estimated FE						
Seller	62,497	55,471	55,471	55,471	55,471	49,957
Buyer	808,383	777,965	777,965	777,965	777,965	751,553
Product×Period	165,730	164,507	164,507	164,507	164,507	163,562
	Share of price dispersion explained by					
Observables		−0.002	−0.002	−0.001	−0.001	−0.001
Product×period FE	0.652	0.647	0.647	0.647	0.647	0.647
Seller FE	0.197	0.201	0.201	0.201	0.201	0.200
Buyer FE	0.045	0.048	0.048	0.048	0.048	0.048
Match residual	0.106	0.105	0.105	0.105	0.105	0.106
	Share of within-seller price dispersion explained by					
Observables		0.000	0.001	0.002	0.002	0.002
Buyer FE	0.169	0.175	0.175	0.175	0.175	0.176
Match residual	0.831	0.825	0.824	0.823	0.823	0.821

Notes: The table reports results of the estimation of Eq. (6), over 2002–2006. The last two panels are variance decompositions of observed price discrepancies into the components entering Eq. (6), in the whole sample and within seller×product×period triplets.

valuations, and/or demand higher-quality goods, which inflates the relative price that they pay.³²

Columns (2)–(6) in Table 10 augment the specification with additional observed variables to dig deeper into the origin of the match-specific component ($e_{sb(c)pt}$) entering Eq. (6). In column (2), we control for the size of the transaction (in euros) and the distance between the exporter and her partner.³³ These two variables explain another 1% of the variance of prices. Export prices are larger over large transactions and toward distant countries. These results are consistent with evidence based on more aggregated firm-level data (Manova and Zhang, 2012; Martin, 2012).

In column (3), we add two variables that capture the dynamics of prices, over time. Namely, the “Seller's experience” is the number of years since the seller started exporting the product to the destination. The “Age of the relationship” is the number of years since the seller first started serving the buyer with the product. The coefficient on the seller's experience is positive and significant, whereas the impact of the age of the relationship is significantly negative. In quantitative

³² We also computed the correlation of the fixed effects with the number of French exporters the firm has been connected to, over time. The correlation is negative but very small.

³³ The “Distance” separating the firm and the buyer is a population-weighted average of distances between the seller's commuting zone in France and the destination's main cities. This measure of distance is taken from Laboureau (2018). We also run a specification with an alternative measure of distance, namely, the driving time between the commuting zone and the destination's main cities. Results were qualitatively unchanged.

terms, one more year of experience implies 0.2% higher prices, on average. The downward renegotiation of prices within a firm-to-firm relationship amounts to 1% after one year.³⁴

Taken together, these coefficients are interpreted as follows. Within a firm-to-firm relationship, a downward renegotiation of prices occurs over time. Assuming the renegotiation takes place given a constant variety sold by the seller to the buyer, implies the buyer recovers a larger share of the surplus of the transaction when interacting with the seller repeatedly. This observation is true even though, on average, the seller increases her average price over time, when she acquires experience in the destination. The reason for these conflicting results is that sellers compensate for the decreasing prices recovered from their existing customers by acquiring new buyers and charging them with higher prices. This interpretation is consistent with results in columns (4) and (5). In column (4), we add a measure of the number of partners that the firm serves in the destination. The coefficient associated with this variable is identified within a firm, over time, and thus is interpreted as the marginal effect of the firm's expanding its portfolio of clients. The impact is positive and makes the coefficient on the firm's experience turn nil. This

³⁴ Note the downward trend of prices within a firm-to-firm relationship is recovered from all repeated transactions, including those that display full price rigidity. Indeed, some of the relationships we observe in the data are likely based on long-term contracts between the firm and its client. If these long-term contracts include pre-set prices, the following series of prices observed over the corresponding repeated transactions will display very little variance, if any. Because we cannot distinguish long-term contracts from spot transactions, there is nothing we cannot control for this possibility.

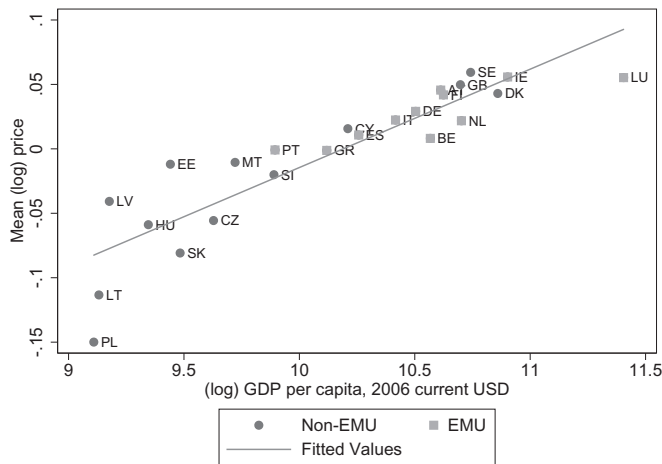


Fig. 9. Mean level of prices, across destinations countries. Notes: This figure correlates the mean level of prices set by French firms to a destination with the country's GDP per capita (in 2006, Source: World Bank). To account for composition effects, mean prices are recovered from the estimated buyer fixed effects entering Eq. (6).

finding is consistent with the interpretation that the positive impact of firms' experience is driven by the acquisition of new partners.³⁵ In column (5), we further augment the specification with two dummy variables, one for "One shot" relationships, that is, transactions involving a seller and a buyer who will never interact again in the future, and one for "First time" transactions, that is, the first transaction involving a particular pair of firms. Both coefficients are positive and highly significant. The outcome is that sellers charge higher prices, on average, across buyers that they have just met. In quantitative terms, "One shot" relationships are priced at 0.8% more than repeated transactions, whereas the first transaction of a repeated relationship is priced at 0.2% more than the next ones.

Finally, in column (6), we interact several variables with a dummy identifying the top four firms in their sector. The interaction is meant to take a focus on the "superstar firms" that have recently been argued to have gained market power (De Loecker and Eeckhout, 2017). Because superstar firms are, on average, more experienced in export markets, tend to serve more partners, and are involved in larger transactions, the coefficients estimated on the corresponding variables might to some extent reflect their pricing strategies. Part of the heterogeneity in pricing strategies between superstar firms and the rest of the distribution is, however, absorbed into the seller fixed effects entering Eq. (6). We confirm this result in Fig. 10, which shows the correlation of the seller fixed effects with the firms' relative sales. As expected, the correlation is positive, meaning large firms on average set higher prices. In column (6), we further interact the age of the transaction and the two dummies for "One shot" relationships and "First time" transactions with the dummy for large firms. The three coefficients on the interactions are found significantly positive. Therefore, price renegotiation on the match is relatively less pronounced for large firms, although these firms are also the ones that manage to set the highest prices across new consumers. This helps refine the evidence in De Loecker and De Loecker and Eeckhout (2017). According to our results, the increase in superstars' average markups is explained by their ability to attract new customers. High markups they charge these new customers allow sellers to compensate for downward renegotiations with their existing partners and maintain high average markups, over time.

³⁵ See Lenoir and Patault (2019) for a detailed discussion of the dynamics of exporters' buyers acquisition based on the same data.

Evidence in this section is thus consistent with the dispersion of within-seller prices being driven by two reinforcing factors. On the one hand, buyers are shown to be charged heterogeneous prices, which is consistent with price discrimination based on heterogeneity in consumers' valuation for the firm's product. On the other hand, the dynamics of prices charged by exporters, over time, suggests buyers that are charged high markups renegotiate over time the sharing of surplus which exerts downward pressures on export prices. To compensate for this downward pressure on their average markup, sellers must attract new buyers that they can charge high prices. This differentiation of prices between new and old customers can explain part of the cross-sectional dispersion of prices observed within a firm.

6. Conclusion

This paper exploits fine grained information on the unit price that French exporters charge their European buyers. We document a significant level of price dispersion both within product categories across exporters and within exporters across buyers. This latter source of price discrepancies is indicative of significant deviations from the LOP – even within the euro area. The median coefficient of variation of prices set by a French exporter over the different partners in its portfolio is 30%, but we show that pricing strategies are highly heterogeneous. Although a small fraction of exporters adopt a near uniform pricing strategy, most exporters charge different prices across buyers. We further document a substantial level of price discrimination among homogeneous chemical products, which suggests that a non negligible share of price discrepancies is triggered by differences in mark-ups rather than quality/composition effects.

Price discrimination is particularly prevalent among firms exporting differentiated products, in particular durable ones, and for large firms. Interestingly, although we provide evidence of downward price renegotiations as the buyer-seller relationship is maintained over time, our results suggest these large exporters maintain high average mark-up rates, by increasing their customer base and charging new buyers with high prices. These results shed new lights on some of the micromechanisms that could explain the observed increase in market power (De Loecker and Eeckhout, 2017) and how they potentially rely on informational frictions on the buyers' side.

Appendix A. Data Appendix

This section provides additional information regarding the various variables used as controls in the analysis.

A.1. Variables constructed using the Customs data

Count partners. Tables 4, 5, and 6 use as control the number of partners in the firm's portfolio, which is also equal to the number of observations used to recover the corresponding coefficient of variation. This variable is also used in Table 10 to understand the dynamics of average prices, within a seller.

Size. We use two measures of the "size" of trade relationships as controls in the analysis. Throughout the paper, the size of a trade relationship is measured as the value of exports recorded in the Customs database. In the raw data, the variable is available for each transaction involving a seller s , a buyer $b(c)$ over a particular product p , and for a specific period t . This is what is used as control in Table 10 and referred to as the "Transaction Size". Table 6 uses the value of trade as a measure of the size of the exporting firm. In that case, transactions are aggregated within an exporter \times product \times period, across all partners to which the firm is connected.

Age and experience. The duration of trade relationships is also constructed using the panel dimension of the data. The "Seller's experience" measures the experience of the firm as an exporter. In Table 10, it is




Fig. 10. Mean level of prices, across size bins. Notes: This figure correlates the mean level of prices set by French firms with their relative size in their sector. To account for composition effects, mean prices are recovered from the estimated seller fixed effects entering Eq. (6). The distribution of firms' sales in each sector is then discretized into 50 bins, and the mean fixed effect inside each bin is correlated with the relative sales of firms there.

measured as the number of years since the firm has started serving the destination.³⁶ In Table 6, the focus is on the overall dispersion of prices within a seller and product, and experience is thus measured as the number of years since the firm started exporting this particular product in the EMU. Finally, the “age” of a relationship used as control in Table 10 is measured relative to the date of the first transaction involving the seller and its foreign partner for a particular product.

Core product. In Table 6, we introduce a dummy variable identifying the firm's “Core” product. For each firm and product, we first aggregate exported values across all partners, all destinations, and all periods. The core product is then defined as the most important product generating at least 30% of the firm's overall sales. We also run a specification with the firm's main product (in value terms) without imposing that it represents at least 30% of the firm's sales but results were qualitatively the same because more than 75% of multi-product firms have their core product that represents at least a third of their exports. Finally, note the “Core” product dummy is normalized to zero for firms that export a single product.

Count products. In Table 7, column (4), we control for the number of products that the firm is exporting. Here as well, the variable is recovered from aggregated trade across all periods and destinations, within a firm.

First time and One-shot dummies. In Table 10, columns (5) and (6), we control for two dummy variables identifying firm-to-firm relationships displaying a single transaction (dummy “One shot”), as well as the first transaction involving a particular pair of firms (dummy “First time”). Note that the “First time” dummy is always equal to 1 when the “One shot” dummy is as well but the reverse is not true so that coefficients on these two dummy variables are identified.

A.2. Variables recovered from external sources

In Tables 7 and 10, we use various firm-level variables to document the heterogeneity in pricing strategies across French firms. We rely on the INSEE-Ficus dataset, which provides balance-sheet data covering the universe of French firms. We merge the dataset with the trade variable, using the French firm Siren identifier.

³⁶ The construction of this variable exploits the overall panel provided to us by the customs, that goes back to 1995.

Sector. Based on the balance-sheet data, we can recover information on the firm's sector of activity, as defined in the NAF nomenclature. We use this information to construct the “Wholesaler” and “Retailer” dummies introduced in column (3) of Table 7.

Relative sales. The “Relative Sales” variable in Table 7 is measured as the ratio of the firm's overall turnover divided by the median turnover of firms active in the same sector of activity. We get rid of the time dimension by calculating this ratio for a single cross-section, namely, 2006. The variable thus captures the relative size of the firm, in its sector. The “Large firm” dummy used in Tables 5 and 10 is constructed from this variable and is equal to 1 for the top four firms in each sector.

Relative profit rate. The “Relative Profit Rate” variable is defined based on information on the ratio of the firm's gross operating surplus over its value added, expressed in relative terms with respect to the median firm in the sector. Here as well, the variable has no time dimension and 2006 is used as reference. The “High Profit Rate” dummy used in Table 5 is constructed from this variable and is equal to 1 for the top 25% of firms in each sector.

Distance. The “Distance” variable used in Table 10 is recovered using the database constructed by Laboureau (2018). Because we are working with trade flows restricted to the EU, the mean distance from France is relatively low and only somewhat heterogeneous across countries. In such a restricted geographic area, the precise location of the firm in France becomes an important source of variation in distances to various destinations. For this reason, Laboureau's dataset dominates more standard datasets such as the CEPII's *distance* database which provides various measures of bilateral distances between all countries in the world. In Laboureau's dataset, distance is measured between a precise commuting zone and a given destination using a weighted average of bilateral distances with the country's most important cities. This dataset can be merged with the firm-level data using information on the firm's location provided by INSEE. The variable used in Table 10 is measured in kilometers, but we also estimated a specification using the driving time to the destination, recovered from the Google Map API.

References

- Varian, Hal, 1980. A model of sales. *Am. Econ. Rev.* 70 (4), 651–659.
- Laboureau, Thomas, 2018. *City gravidist*. Database.
- Abowd, John M., Kramarz, Francis, Margolis, David N., 1999. High wage workers and high wage firms. *Econometrica* 67 (2), 251–334 March.
- Adams, Brian, Williams, Kevin R., 2019. Zone pricing in retail oligopoly. *American Economic Journal: Microeconomics* 11 (1), 124–156 February.
- Amiti, Mary, Itskhoki, Oleg, Konings, Jozef, 2019. International Shocks and Domestic Prices: How Large Are Strategic Complementarities? *The Quarterly Journal of Economics*
- Andrews, M.J., Gill, L., Schank, T., Upward, R., 2008. High wage workers and low wage firms: negative assortative matching or limited mobility bias? *Journal of the Royal Statistical Society Series A* 171 (3), 673–697.
- Antras, Pol, Chor, Davin, Fally, Thibault, Hillberry, Russell, 2012. Measuring the upstreamness of production and trade flows. *Am. Econ. Rev.* 102 (3), 412–416.
- Atkeson, Andrew, Burstein, Ariel, December 2008. Pricing-to-market, trade costs, and international relative prices. *Am. Econ. Rev.* 98 (5), 1998–2031.
- Behrens, Kristian, Boualam, Brahim, Martin, Julien, 2019. Are clusters resilient? Evidence from Canadian textile industries. *J. Econ. Geogr.*
- Bergounhon, Flora, Lenoir, Clémence, Mejean, Isabelle, 2018. *A Guideline to French Firm-Level Trade Data*.
- Bernard, Andrew B., Moxnes, Andreas, August 2018. Networks and trade. *Annual Rev. Econ.* 10 (1), 65–85.
- Cavallo, Alberto, 2018. More Amazon effects: online competition and pricing behaviors. NBER Working Papers 25138. National Bureau of Economic Research, Inc October.
- Crucini, Mario J., Shintani, Mototsugu, April 2008. Persistence in law of one price deviations: evidence from micro-data. *J. Monet. Econ.* 55 (3), 629–644.
- Crucini, Mario J., Telmer, Chris I., Zachariadis, Marios, June 2005. Understanding European real exchange rates. *Am. Econ. Rev.* 95 (3), 724–738.
- DellaVigna, Stefano, Gentzkow, Matthew, 2017. Uniform pricing in US retail chains. NBER Working Papers 23996. National Bureau of Economic Research Inc November.
- Devereux, Michael B., Dong, Wei, Tomlin, Ben, 2017. Importers and exporters in exchange rate pass-through and currency invoicing. *J. Int. Econ.* 105 (C), 187–204.
- Jan, De Loecker, Eckhout, Jan, 2017. The Rise of Market Power and the Macroeconomic Implications. UCL Mimeo.
- Engel, Charles, Rogers, John H., July 2004. European product market integration after the euro. *Econ. Policy* 19 (39), 347–384.

- Engel, Charles, Wang, Jian, January 2011. International trade in durable goods: understanding volatility, cyclicalities, and elasticities. *J. Int. Econ.* 83 (1), 37–52.
- Fitzgerald, Doireann, Haller, Stefanie, 2014. Pricing-to-market: evidence from plant-level prices. *Review of Economic Studies* 81 (2), 761–786.
- Goldberg, Linda S., Tille, Cédric, 2016. Micro, macro, and strategic forces in international trade invoicing: synthesis and novel patterns. *J. Int. Econ.* 102 (C), 173–187.
- Goldberg, Pinelopi K., Verboven, Frank, January 2005. Market integration and convergence to the law of one price: evidence from the European car market. *J. Int. Econ.* 65 (1), 49–73.
- Gopinath, Gita, Gourinchas, Pierre-Olivier, Hsieh, Chang-Tai, Li, Nicholas, October 2011. International prices, costs, and markup differences. *Am. Econ. Rev.* 101 (6), 2450–2486.
- Hausmann, Ricardo, Hidalgo, Cesar, 2014. *The Atlas of Economic Complexity: Mapping Paths to Prosperity*. Vol. 1. MIT Press Books, The MIT Press January.
- Engel, Charles, John, H., 1996. Rogers, how wide is the border? *Am. Econ. Rev.* 86 (5), 1112–1125 December.
- Kaplan, Greg, Menzio, Guido, November 2015. The morphology of price dispersion. *Int. Econ. Rev.* 56, 1165–1206.
- Lenoir, Clemence, Patault, Berangere, 2019. Exporters Dynamics: Sales Managers and Firm's International Portfolios.
- Levchenko, Andrei A., Lewis, Logan T., Tesar, Linda L., 2010. The collapse of international trade during the 2008–2009 crisis: in search of the smoking gun. *IMF Economic Review* 58 (2), 214–253 December.
- Manova, Kalina, Zhang, Zhiwei, 2012. Export prices across firms and destinations. *Q. J. Econ.* 127 (1), 379–436.
- Martin, Julien, 2012. Markups, quality, and transport costs. *Eur. Econ. Rev.* 56 (4), 777–791.
- Berman, Nicolas, Martin, Philippe, Mayer, Thierry, 2012. How do different exporters react to exchange rate changes? *Q. J. Econ.* 127 (1), 437–492.
- Martin, Julien, Mejean, Isabelle, Parenti, Mathieu, 2019. Relationship Stickiness: Measurement and Applications to International Economics.
- Cavallo, Alberto, Neiman, Brent, Rigobon, Roberto, 2014. Currency unions, product introductions and the real exchange rate. *The Quarterly Journal of Economics* 129 (2), 529–595.
- Nunn, Nathan, 2007. Relationship-specificity, incomplete contracts, and the pattern of trade. *Q. J. Econ.* 122 (2), 569–600.
- Rauch, James E., June 1999. Networks versus markets in international trade. *J. Int. Econ.* 48 (1), 7–35.
- Stigler, George J., Jun., 1961. The economics of information. *J. Polit. Econ.* 69 (3), 213–225.
- Tirole, Jean, 1988. *The Theory of Industrial Organization*. Vol. 1. MIT Press Books, The MIT Press March.