BANCA D'ITALIA
EUROSISTEMA

Questioni di Economia e Finanza

(Occasional Papers)

Mapping supply chains with micro data: the case of the Italian automotive industry

by Luca Brugnara, Gianmarco Cariola, Pasquale Recchia and Simone Zuccolalà

July 2026

Number 1041

BANCA D'ITALIA
EUROSISTEMA

Questioni di Economia e Finanza

(Occasional Papers)

Mapping supply chains with micro data: the case of the Italian automotive industry

by Luca Brugnara, Gianmarco Cariola, Pasquale Recchia and Simone Zuccolalà

Number 1041 – July 2026

The series Occasional Papers *presents studies and documents on issues pertaining to the institutional tasks of Banca d'Italia and the Eurosystem. The* Occasional Papers *appear alongside the* Working Papers *series which are specifically aimed at providing original contributions to economic research.*

The Occasional Papers *include studies conducted within Banca d'Italia, sometimes in cooperation with the Eurosystem or other institutions. The views expressed in the studies are those of the authors and do not involve the responsibility of the institutions to which they belong.*

The series is available online at www.bancaditalia.it .

ISSN 1972-6643 (online)

Designed by the Printing and Publishing Division of Banca d'Italia

MAPPING SUPPLY CHAINS WITH MICRO DATA: THE CASE OF THE ITALIAN AUTOMOTIVE INDUSTRY

by Luca Brugnara*, Gianmarco Cariola*, Pasquale Recchia* and Simone Zuccolalà*

Abstract

This paper develops a firm-level input–output framework to measure participation in production networks using transaction-level data and applies it to the automotive industry. Using a novel dataset of firm-to-firm transactions, we construct a firm-specific measure of supply chain exposure based on the Ghosh inverse, which is a sufficient statistic for the upstream propagation of demand shocks within a standard production network model. We use this framework to trace both direct and indirect sales linkages across firms. We find that around 8,000 Italian firms are significantly exposed to the EU automotive industry, with indirect linkages accounting for more than half of the supply chain. Exposure is highly heterogeneous across firms, sectors, and regions, and simulated demand shocks generate sizeable upstream propagation effects: a 1 per cent decline in final demand for cars would reduce value added among final producers by €63 million, with an additional €93 million propagating upstream through the network. The framework is scalable and can be extended to other downstream industries with clearly identifiable final producers, providing a flexible tool for analysing production networks and the upstream transmission of demand shocks.

JEL Classification: D57, L14, L62.

Keywords: production networks, firm-to-firm transaction data, shock propagation, automotive industry.

DOI: 10.32057/0.QEF.2026.1041

---

* Banca d'Italia.

1 Introduction[1]

The automotive industry is one of Europe's long-standing industrial pillars, combining large direct contributions to employment and value added with dense upstream and downstream input–output linkages and substantial R&D intensity (Draghi, 2024). Italy is no exception: beyond the firms directly classified in ATECO division 29, a wide network of suppliers and service providers is intertwined with vehicle production and sales (Orame et al., 2025). Over the last decade the European automotive system has been buffeted by major shocks that have reshaped its technological and organizational landscape. Alongside the reputational effects of the 2015 "Dieselgate" scandal (Pianeselli and Orame, 2023) and an increasingly stringent regulatory framework (the EU Fit for 55 package in 2021 and Regulation EU/2023/851 setting a 2035 phase-out for internal combustion engines), the sector has faced a series of technological shocks, including the rapid diffusion of electric powertrains, digital technologies and automation in production processes. At the same time, the ongoing internationalization of leading European carmakers and the increase in import competition from China have further contributed to redefining the structure and geography of the industry.

These developments have intensified competitive pressures and raised significant concerns about the resilience of upstream suppliers, especially in policy circles (Draghi, 2024). As vehicle production contracted across Italy and more broadly in Europe (ACEA, 2025), questions emerged about which firms would be most affected by the industry reconfiguration and where vulnerabilities may cluster across sectors and regions. Addressing these issues requires going beyond standard industry classifications and reconstructing the full network of firms that are linked directly or indirectly to carmakers.

Against this background, our contribution is threefold. First, we provide the first firm-level perimeter of the automotive supply chain, while existing studies are based on industry-level data, which necessarily overlook firm-level heterogeneity and do not allow a detailed geographical characterization of the network. Second, to construct this perimeter we adopt an input–output perspective and compute firm-level Ghosh inverse coefficients, using a recursive procedure to trace upstream sales shares starting from final producers, both Italian and EU. Third, we show that the resulting coefficients have a sufficient statistic interpretation: in a model à la Acemoglu et al. (2016), featuring a Cobb–Douglas production function, constant returns to scale and perfect competition, they measure the upstream propagation

[1]We are thankful to Matteo Alpino, Emidio Cocozza, Silvia Del Prete, Maura Francese, Elena Gentili, Andrea Linarello, Francesco Manaresi, Andrea Orame, Lorenzo Pesaresi and Patrick Zoi for their comments and suggestions. All remaining errors are our own. The opinions expressed in this paper do not necessarily reflect those of the Bank of Italy.

of a demand shock along the production network. This makes our metric well suited to capture the impact of short-run demand shocks (such as changes in consumer preferences or increased import competition) and, in the long-run, to measure firms' exposure to shocks coming from downstream producers.

Our study documents that the Italian participation in the European automotive supply chain is sizable. This perimeter includes final car producers located in Italy as well as Italian firms that supply, directly or indirectly, car producers located either in Italy or elsewhere in the EU. In 2022, this corresponded to around 8,000 Italian firms with a significant degree of exposure. Among limited-liability companies, the exposed value added amounted to €15.6 billion, corresponding to 1.6 percent of the value added of the non-financial and non-real-estate private sector. About 40 percent of this amount was attributable to car producers located in Italy, 47 percent to their domestic suppliers, and the remaining share to Italian suppliers of car producers located in other EU countries. Indirect relationships play a central role: around one third of total exposed value added originates from indirect linkages, underscoring the importance of identifying the full upstream network rather than relying solely on direct suppliers. Under the assumptions of our model, a 1 percent decline in final demand for cars produced in the EU would reduce value added by €63 million among final producers located in Italy, with an additional €93 million loss propagating to Italian suppliers through their direct and indirect links to car producers located both in Italy and elsewhere in the EU. Overall, this implies a network multiplier of about 2.5; restricting the analysis to direct suppliers only would reduce the multiplier to 1.7, thereby substantially understating the extent of upstream propagation.

From a sectoral perspective, the bulk of the supply chain is concentrated in industries traditionally upstream of vehicle manufacturing (motor-vehicle components, metals, rubber and plastics), although several service sectors also play a non-negligible role. Geographically, while most automotive-related activity is located in Italy's major northern regions, the strongest dependence on automotive demand is observed in areas surrounding vehicle production plants, pointing to pronounced geographic clustering. Additionally, while large firms account for the majority of value added related to car makers, smaller suppliers display the greatest reliance on the sector, reflecting their limited diversification.

To further investigate firm-level heterogeneity, we run a series of OLS regressions comparing firms within the automotive supply chain with other firms operating in the same sector and region, and with similar employment levels and capital intensity. We find that firms in the automotive supply chain tend to have higher sales, value added, productivity, wages, tangible and intangible capital but lower diversification of sales across clients. Additionally, if we restrict attention to direct linkages, we see that firms in the supply chain have a higher propensity to patent; their innovation activity is particularly concentrated in mechanical and electronic technologies.

Our paper relates to three strands of the literature. First, it contributes to the growing body of studies on the automotive industry. Much of this work consists of policy papers by governmental institutions (e.g. Dechezleprêtre et al., 2023, OECD, 2024 and European Commission, 2025), but a growing number of academic studies examines the structure and evolution of the this industry in greater detail (see for example Sturgeon et al., 2008, Pavlínek, 2020, Pianeselli and Orame, 2023 and Orame et al., 2025). Second, our contribution is linked to the literature that measures supply chain exposure using input–output tables and clarifies the concepts, metrics and interpretations of upstream and downstream linkages (see Baldwin and Lopez-Gonzalez, 2015, Baldwin and Freeman, 2022 and Baldwin et al.,2022). Our work builds on the insights introduced by these papers but departs from the standard framework by implementing the Ghosh inverse at the firm rather than sector level, thus markedly increasing the granularity of the analysis. Finally, our paper is related to the growing literature that uses firm-to-firm transaction data to study how shocks propagate through production networks (e.g. Dhyne et al., 2015, Boehm et al., 2019, Carvalho et al., 2021 and Bijnens et al., 2025). From a modeling perspective, although the literature on shock propagation in production networks is vast and rapidly growing (see for example Baqaee and Farhi, 2019, Baqaee and Rubbo, 2023, and Bertolotti et al., 2025 for an application tailored to the Italian economy), we build on Acemoglu et al. (2016), which provides a tractable framework that can be calibrated at the firm level and highlights the assumptions that we need in order to give the Ghosh inverse coefficients a sufficient statistic interpretation.

The paper is structured as follows. In Section 2 we describe the novel dataset and the procedures used to prepare the data for the empirical analysis. In Section 3 we present our methodology based on the estimation of Ghosh inverse coefficients, aimed at reconstructing the automotive supply chain. In Section 4 we show aggregate results from this mapping. Section 5 shows our regression analysis, which aims to analyze individual characteristics of the firms that are most exposed to the automotive industry. Finally, in Section 6 we present the results of a simulation of a car demand shock, using a simple theoretical framework based on Acemoglu et al. (2016).

2 Data

The bulk of our analysis relies on a novel dataset of electronic invoices provided by *Agenzia delle Entrate*, the Italian tax revenues agency. The dataset, which we refer to as "Fattel" throughout the paper, spans the period 2019–2023[2] and contains domestic transactions between sellers and buyers, aggregated at the quarterly level. Each observation corresponds to a seller–buyer–quarter triple and reports the aggregated gross transaction value, the corresponding value added tax (VAT) liability, and the number of exchanged invoices[3].

The buyers and sellers involved in the transactions can be of several types: (1) firms with legal personality (*persone giuridiche*); (2) firms treated as natural persons (*ditte individuali*); (3) public sector organizations; and (4) final consumers. Individual-level transactions are reported when both the buyer and seller are type (1) or (3) entities. For type (2) firms, transactions are aggregated by 2-digit ATECO sector. Transactions involving final consumers are aggregated across all consumers.

Overall, Fattel contains over 160 million observations per year, corresponding to approximately €2,800 billion in annual trade. For further details, see Astuti and Piras (2025).

To capture Italian firms' participation in the European automotive supply chain, we complement domestic transaction data with customs records from the Italian Customs and Monopolies Agency. These data provide yearly export flows at the level of Italian firm–foreign counterpart pairs, allowing us to observe the foreign buyers of Italian exporters and link them to firm-level information from Orbis. In this way, we are able to identify export flows directed to final car manufacturers, as described in subsection 3.b.

In addition, we merge Fattel and customs data with several complementary firm-level datasets using tax identifiers: (i) balance sheet data from Cerved, that covers the universe of Italian limited liability companies and provides information on firms' value added and financial structure, (ii) business registry data from the Italian Chamber of Commerce, containing firms' administrative information and the distribution of employees across municipalities, (iii) social security records from INPS and (iv) patent data from OrbisIP. Finally, we retrieve information on value added and turnover of the non-financial private sector at the territorial level from the SBS Frame dataset.

[2]In the empirical analysis we restrict the sample to the period 2019–2022, as additional data sources used in the paper (in particular export data) are not available for 2023.

[3]For fiscal reasons, firms belonging to a value added tax group (*gruppo IVA*) are treated as a single taxable entity. Intra-group transactions are therefore not recorded.

2.a Preparing the data for the empirical analysis

The raw dataset provided by the tax authority contains some inconsistencies and potential reporting anomalies. To address these issues, we implement a cleaning routine adapted from Dhyne et al. (2015). First, we drop transactions with zero sales, as they are not economically meaningful. We then perform the following steps:

1. *Misreporting of decimal points*: we detect several cases of incorrectly reported decimal points in sales declarations. To identify them, we compute the observed VAT rate as $\frac{VAT}{sales} \times 100$ and compare it with the standard VAT rates applied in Italy[4] (4%, 5%, 10% and 22%). In some cases, the observed VAT rate equals, for instance, 0.22%, 2.2% or 220%: these patterns indicate misreporting of decimal points. We assume that the misreporting occurs on the transaction value side, since a grossly incorrect VAT amount would likely be detected and corrected by the tax authority or the firm itself, as it directly affects tax payments. Accordingly, we correct the reported transaction value by dividing or multiplying it by 10, 100, or 1,000 whenever the observed VAT rate deviates from a standard rate by a factor of 10, 100, or 1,000, respectively.

2. *Swapped VAT amounts and transaction amounts*: there are cases were the paid VAT is larger than the transaction value. In these cases, we compute the observed VAT rate: if it is within the inverse of (21.5; 22.5), (9.5; 10.5) or (3.5; 4.5) intervals (i.e. it is compatible with standard VAT rates), we infer that transactions and VAT have been mistakenly inverted and thus we swap them[5].

3. *Clean inconsistencies over time*: since trade relationships tend to be stable over time, we interpret huge spikes in a single year as a misreporting. In this sense, we aggregate quarterly data at the year level: for a given seller-buyer couple, if we observe that in year t the trade amount is 1,000, 10,000 or 100,000 times larger (smaller) than in t-1 and in t+1, we divide (multiply) the yearly transaction value by 1,000, 10,000 or 100,000, respectively.

The cleaning operations have relatively light effects, as shown in Table 2.1, Panel A: less than 5,000 observations per year are cleaned and with a less than 1% decrease in total trade value between the pre- and post-cleaning routine.

[4]This procedure is applied only to seller–buyer–quarter triples for which a single invoice is observed. While Dhyne et al. (2015) rely on transaction-level data, our dataset aggregates invoices at the quarterly level; restricting the correction to cases with exactly one underlying transaction prevents misidentification of decimal-point errors.

[5]As in the previous point, this procedure is applied only to seller–buyer–quarter triples for which a single invoice is observed.

Table 2.1. Data cleaning and further exclusions: effects on total sales

	2019	2020	2021	2022
Panel A. Effects of data cleaning				
Pre-cleaning sales (bn. €)	2,706	2,284	2,771	3,490
1) Decimal rounding *(∼2,500 obs. per year)*	−0.02%	−0.03%	−0.02%	−0.02%
2) VAT-transaction swap *(∼1,500 obs. per year)*	+0.00%	+0.00%	+0.00%	+0.00%
3) Time inconsistencies *(∼500 obs. per year)*	+0.00%	−0.26%	−0.04%	−0.02%
Post-cleaning sales (bn. €)	2,705	2,277	2,769	3,489
Pre–post cleaning var.	−0.02%	−0.29%	−0.06%	−0.04%
Panel B. Further cleaning: weight on total sales				
Negative values	4.62%	3.27%	2.18%	2.74%
Self-invoices	1.27%	1.21%	0.72%	0.71%
Transactions < €1,000	0.43%	0.45%	0.40%	0.33%
Sellers/Buyers < €10,000	0.08%	0.09%	0.07%	0.06%
Final consumers	12.67%	10.47%	9.79%	8.83%

Notes: Panel A reports the variation in yearly total sales induced by each data cleaning step; in parentheses, the approximate number of affected observations. Panel B reports the weight of further excluded transactions on total sales by year.

Additionally, we apply some cleaning steps aimed at filtering out economically irrelevant transactions; in particular, we remove self-invoices and, in the case of negative transaction values, we invert the sign and swap buyer and seller. These steps have only a limited effect on total sales, as shown in as shown in Table 2.1, Panel B.

2.b Outliers management

In the Fattel dataset, some observations are flagged as outliers by the data provider, typically corresponding to unusually large sales values. The preliminary cleaning routine only marginally affects these cases (see also Benecchi et al., 2025). To address this, we implement an additional outlier-cleaning procedure based on a validation with external balance sheet data. Specifically, we focus on selling firms that can be matched to the Cerved dataset; for these companies, we reconstruct annual revenues by aggregating all domestic sales recorded in Fattel and adding yearly export values from customs data (when available). We compute this aggregate both including and excluding outlier observations, and then compare it with total revenues reported in Cerved balance sheets (sales plus disinvestment of capital goods)[6].

[6] Two relatively small groups of firms are excluded from this procedure for comparability reasons: (1) firms belonging to *gruppi IVA*, since intra-group transactions are not recorded in the dataset; and (2) firms reporting under International Accounting Standards (IAS).

For each firm-year, we drop outlier observations if their removal narrows the gap between the two aggregates, defined as:

$$\Delta(Revenues) = \left|Revenues_{Cerved} - Revenues_{Fattel+export}\right|$$

In the vast majority of cases, revenues reconstructed from Fattel plus exports fall short of those reported in balance sheets[7]. Consequently, our procedure tends to preserve outliers, making it relatively conservative. The results are summarized in Table 2.2: our cleaning affects approximately 1,000 firms per year, while leaving the (relatively small) median $\Delta(Revenues)$ unaffected and only marginally reducing the total volume of sales. However, by eliminating abnormally large transactions, the procedure lowers the average $\Delta(Revenues)$ by 5–15% and reduces the incidence of outliers on total sales by 1–4 percentage points annually.

Table 2.2. Outliers management

	2019	2020	2021	2022
Fattel sales (bn. €)				
Pre-cleaning	2,705	2,277	2,769	3,489
Post-cleaning	2,592	2,230	2,733	3,391
Var. %	-4.2%	-2.1%	-1.3%	-2.8%
N. *cleaned* firms	1,104	1,181	1,165	1,059
Outliers share				
Pre-cleaning	15%	13%	10%	13%
Post-cleaning	11%	11%	8%	11%
Median Δ(Revenues) (€)				
Pre-cleaning	36,342	30,038	35,764	40,914
Post-cleaning	36,256	30,000	35,684	40,822
Average Δ(Revenues) (€)				
Pre-cleaning	1,007,973	798,819	924,718	1,241,366
Post-cleaning	858,844	737,140	878,669	1,110,231
Var. %	-14.8%	-7.7%	-5.0%	-10.6%

Notes: *Cleaned* firms are those whose outlier observations are removed by our routine. Outliers incidence refers to the incidence of outliers on total Fattel sales. Δ(Sales) refers to the absolute difference between Revenues from Fattel plus export and Cerved.

[7] This gap arises mainly because not all transactions, especially those with final consumers, require electronic invoicing. It is particularly pronounced in industries with frequent direct sales to consumers, such as retail and wholesale trade.

3 Methodology

A natural starting point for measuring a firm's participation to the automotive supply chain is the share of its sales directed to car manufacturers. This direct sales share is intuitive and informative: it identifies firms that sell directly to the automotive industry and quantifies the intensity of this link. However, it remains incomplete. In real-world production networks, suppliers may be indirectly connected to final producers through their customers. Ignoring such indirect linkages would lead to an overly narrow characterization of the supply chain.

To illustrate this point, consider the simple production network depicted in Figure 3.1, where firms are arranged across three layers: final producers (automotive and non-automotive), their direct suppliers, and a second tier of upstream suppliers. Suppose that firm D sells €100 to a car manufacturer and €300 to firms unrelated to the automotive sector. According to the direct sales share, D's exposure to the automotive industry is 25 per cent, and D would be classified as part of the supply chain.

Figure 3.1. An example of a simple production network

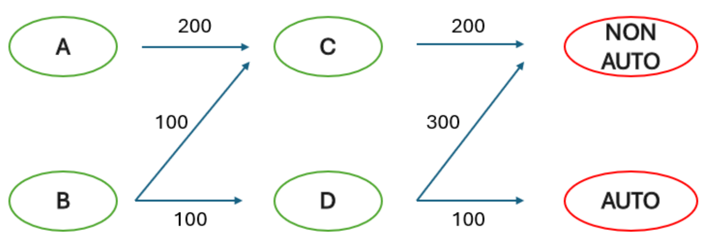

Notes: Example of a production network in which end nodes are a carmaker and a non-carmakers. Numbers on the directed arrows denote sales amounts.

Now consider firm B, which sells half of its output to D. Based solely on direct sales, B would appear to have no connection with the automotive sector. Yet this conclusion is counterintuitive: since half of B's sales go to firm D, and 25 per cent of D's sales go to carmakers, B is indirectly exposed to the automotive industry. By contrast, firm A, which supplies only the non-automotive part of the network, has no such dependence.

This example shows that a measure based exclusively on direct linkages correctly identifies first-tier suppliers, but fails to capture firms that participate indirectly through the production network: a comprehensive metric must therefore account for both direct and indirect relationships. In what follows, we formalize this idea by constructing a firm-level measure of participation in the automotive supply chain that exploits the structure of the production network.

3.a Measuring supply chain participation at the firm level

Let T_{ji} denote the value of sales from firm i (seller) to firm j (buyer)[8], and let y_i be firm i's gross output, which in our setting coincides with its total sales[9]. Following the supply-driven input–output approach initiated by Ghosh (1958) and recently extended by Baldwin et al. (2022), we normalize each firm's outflows by its total output and define the direct sales shares:

$$a_{ji} \equiv \frac{T_{ji}}{y_i}, \tag{1}$$

Collecting these coefficients yields the $N \times N$ matrix A that summarizes the structure of the production network: a positive a_{ji} coefficient indicates that firm i sells to firm j, and its magnitude measures the intensity of this link as a share of i's total sales. However, participation in a supply chain also involves indirect relationships. As we pointed out in the previous section, a firm may not sell directly to a car manufacturer, but may supply a customer that does, or a customer of a customer, and so on.

The higher-order connections implied by the network can be formalized in a simple way. The matrix A captures one-step (direct) buyer–supplier links. The matrix A^2 captures two-step links: its (j, i)-th element sums over all possible intermediate firms z the product $a_{jz}a_{zi}$, i.e. the share of i's sales that reaches j through exactly one intermediate customer z. More generally, A^k collects all paths of length k in the firm network, with each path weighted by the product of the corresponding direct sales shares.

In this spirit, we can write the matrix

$$H \;\equiv\; I + A + A^2 + A^3 + \dots, \tag{2}$$

whose (j, i)-th entry h_{ji} aggregates the contribution of all directed paths from firm i to firm j. Under the usual input–output regularity condition that the spectral radius of A is less than

[8] For computational purposes, when building T_{ji} we drop transactions with a value below €1,000 and sellers (or buyers) whose total annual sales (or purchases) are below €10,000 in a given year. These filters substantially reduce the number of observations and thus the computational burden of constructing the production network, but, as shown in Table 2.1, the combined weight of these excluded categories in terms of aggregate sales is extremely limited.

[9] In constructing total sales y_i, we use the maximum between firm-level turnover reported in Cerved and the sum of domestic sales from Fattel and exports from customs data, whenever both measures are available. This choice reflects the fact that Fattel does not capture certain components of firms' revenues, including parts of final demand not subject to electronic invoicing and exports of services. Moreover, for firms belonging to VAT groups, intra-group transactions are not observed in Fattel; in such cases, balance-sheet turnover provides a more comprehensive measure of total output.

one, this Neumann series converges and we can write the following:

$$H = (I - A)^{-1}, \tag{3}$$

which we refer to as the Ghosh inverse at the firm level.

In this framework, h_{ji} can be interpreted as the total (direct and indirect) share of firm i's sales that is ultimately absorbed by firm j through all layers of the production network. Accordingly, if we focus on the columns of H associated with a set of final car manufacturers (see Section 3.b), the corresponding coefficients provide a measure of firm i's participation in the automotive supply chain, taking into account both direct and indirect linkages. Empirically, we recover these coefficients by tracing upstream sales linkages recursively from the set of domestic and EU car manufacturers identified in Section 3.b, summing the weighted paths that connect each firm to those final producers. This is computationally equivalent to the fixed point iteration procedure adopted, among others, by Dhyne et al. (2021) and Bernard et al. (2022). Additional details are provided in Appendix A.

We will show in Section 6 that this measure can be naturally recast within a production network model with perfect competition and a Cobb–Douglas production function with constant returns to scale, thereby giving it a sufficient statistic interpretation.

3.b Selecting car manufacturers

The starting point for mapping the Italian perimeter of the European automotive supply chain is the set of final car producers located both in Italy and elsewhere in the EU, which constitute the terminal nodes of the production network. For Italy, the identification of car manufacturers relies on the classification adopted by Orame et al. (2025), which is based in turn on the list provided by the European Automobile Manufacturers' Association (ACEA).

Identifying car producers located abroad is inherently more complex. We restrict our scope to EU car manufacturers for both economic and data-related reasons: on the one hand, the European automotive industry has experienced a significant contraction in recent years and is central to the current policy debate (Draghi, 2024), making it particularly relevant to assess Italian firms' exposure to EU producers; more concretely, firm-level information on foreign counterparts in customs records is consistently available up to 2022 only for intra-EU transactions, which allows for a consistent and reliable identification and matching procedure[10].

[10] According to OECD Inter-Country Input–Output tables for 2022, around 10 percent of Italian sales to the

To identify European final producers, we proceed in three steps.

First, using customs data, we match foreign counterpart identifiers to firm-level information from Orbis.

Second, within the matched sample, we select firms classified under NACE Rev. 2 code 291 (Manufacture of motor vehicles). However, this classification is not sufficiently granular for our purposes: NACE 291 includes producers of trucks, tractors and other vehicles that do not correspond to complete passenger car manufacturers, as well as firms with heterogeneous or ambiguous activity descriptions. We therefore implement an additional screening procedure based on textual information contained in Orbis[11], which we submit to a Large Language Model (Mistral) through an API-based interface for firm-level classification. The model is instructed to classify firms according to whether they manufacture or assemble complete passenger automobiles, explicitly excluding firms involved in the production of trucks, buses, tractors, motorcycles, components, engines, services, leasing, or distribution activities.

We require the model to return a structured output indicating: (i) a binary classification variable (i.e. a dummy equal to 1 if and only if the firm produces automobiles), (ii) a confidence score (low, medium, high), and (iii) a short textual excerpt providing explicit evidence supporting the classification.

Finally, we perform manual validation checks on all firms classified with low confidence. In addition, we manually review all firms with more than 500 employees, irrespective of the model's confidence score. This combined automated and manual procedure allows us to construct a conservative and well-validated set of European car manufacturers, which we then treat as terminal nodes of the foreign part of the production network.

It is important to emphasize a key difference between the part of the supply chain linked to car producers located in Italy and the part linked to car producers located elsewhere in the EU. For the former, the availability of domestic firm-to-firm transaction data allows us to reconstruct the entire set of indirect relationships among Italian firms. In this case, the recursive algorithm traces all upstream linkages within the observed national production network. For the latter, instead, only direct export relationships between Italian firms and EU car manufacturers are observable. Indirect exposure to foreign automotive demand is therefore limited to the Italian suppliers of firms that export directly to those producers. By construction, we do not observe indirect relationships that unfold within foreign production

---

automotive sector (NACE 29) are directed to non-European destinations.

[11]The relevant variables are `trade_description_en`, `products_services`, `description_history`, and `overview_full_overview`

networks, nor cases in which an Italian firm sells to a foreign intermediary that subsequently supplies a foreign car manufacturer.

As a consequence, while the domestic part of the supply chain is fully mapped in terms of direct and indirect linkages, exposure to producers located elsewhere in the EU captures only the portion mediated through observable export relationships.

4 Mapping the Italian participation in the EU automotive supply chain

We now turn to the aggregate patterns that emerge from our firm-level measure of participation in the EU automotive supply chain. By scaling each firm's total value added[12] by its Ghosh inverse coefficient, we obtain the value added attributable—directly or indirectly—to final car manufacturers. This allows us to characterize the sectoral, regional and firm-level distribution of supply chain activity.

4.a Firm-level exposure

Table 4.1 reports descriptive statistics for the firm-level Ghosh inverse coefficients computed as described in Section 3. The distribution is highly skewed: the vast majority of firms have no direct or indirect trade links with car producers. Among firms with a non-zero Ghosh inverse coefficient, the average share of sales to final manufacturers is about 3%. However, this mean is driven by a small number of firms with exceptionally high participation in the supply chain: the median is relatively small, around 0.3%.

Given this strong skewness, many firms display only minimal indirect connections to car manufacturers, often reflecting marginal or incidental relationships. For the purposes of the subsequent analysis—particularly Sections 4.e and 5 where we need a discrete classification of supply chain firms— we consider within the automotive supply chain those firms that display a Ghosh inverse coefficient larger than 0.1, in order to avoid treating firms with negligible exposure as meaningful participants to the automotive network; only a small fraction of firms meet this criterion—roughly 5% of those with a positive coefficient (see Table 4.1)[13].

[12] Firm-level value added is available only for limited liability companies, for which balance-sheet information is recorded in the Cerved archives.

[13] We assess the robustness of this definition by considering two alternative thresholds: a more restrictive cutoff at 30% and a fully inclusive 0% cutoff that incorporates all firms with any positive direct or indirect sales to car makers. As reported in Table 4.2, these alternative definitions substantially alter the number of exposed

Table 4.1. Distribution of participation shares in the automotive supply chain (2022)

Percentiles		Summary stats	
(a) Total sample:			
1%	0.0000	Observations	1,298,939
5%	0.0000	Mean	0.0030
10%	0.0000	Std. dev.	0.0315
25%	0.0000	Variance	0.0010
50% (Median)	0.0000	Skewness	20.02
75%	0.0000	Kurtosis	484.99
90%	0.0005	Min	0
95%	0.0046	Max	1
99%	0.0556		
(b) Auto sales > 0:			
1%	0.0001	Observations	147,396
5%	0.0002	Mean	0.0263
10%	0.0004	Std. dev.	0.0901
25%	0.0011	Variance	0.0081
50% (Median)	0.0035	Skewness	6.73
75%	0.0125	Kurtosis	56.30
90%	0.0470	Min	0
95%	0.1134	Max	1
99%	0.5148		

Notes: Summary statistics for firms' share of direct and indirect sales to Italian and EU car producers on total domestic and foreign sales (Ghosh inverse coefficients). Panel a includes the full sample of firms; panel b only includes firms whose contribution to automotive is larger than 0.

According to our estimates (Table 4.2), around 147,000 firms with available balance-sheet information were directly or indirectly exposed to car producers in 2022. Among them, roughly 8,000 exhibited a participation share in the automotive supply chain exceeding 10 percent. In terms of economic magnitude, the value added generated by Italian suppliers of car producers amounted to about 9.3 billion euros. Most of this amount, 7.4 billion euros, was associated with supply relationships with car producers located in Italy, while the remaining 1.9 billion euros reflected linkages to producers located in other EU countries. Adding the value added generated by car producers located in Italy, equal to roughly 6.3 billion euros, brings the overall perimeter of the supply chain to 15.6 billion euros. This corresponds to 1.6 percent of the value added of the Italian non-financial and non-real-estate private sector, or 0.8 percent of Italian GDP[14].

firms, while the associated value added declines much less sharply, confirming that many of the firms included under the most inclusive definition are only marginally exposed to the automotive sector.

[14] The relatively small magnitude of this figure partly reflects the conservative nature of our estimation

Table 4.2. Italian perimeter of the European automotive supply chain at various thresholds

	2019		2022	
	N. of firms	VA	N. of firms	VA
Car producers located in Italy	10	4,165	9	6,315
Direct & indirect suppliers of Italian & EU car producers				
No threshold (0%)	151,817	11,407	147,387	9,333
Threshold 10%	10,695	7,982	8,163	6,256
Threshold 30%	4,076	6,270	2,943	4,714
Direct suppliers of Italian & EU car producers				
No threshold (0%)	6,386	4,753	5,933	4,497
Threshold 10%	1,630	4,152	1,455	3,831
Threshold 30%	817	3,097	743	2,854
Direct & indirect suppliers of Italian car producers				
No threshold (0%)	147,341	8,937	143,539	7,429
Threshold 10%	8,913	5,898	6,776	4,723
Threshold 30%	3,352	4,589	2,412	3,435

Note: The first row reports final car producers located in Italy. The three panels below refer to Italian suppliers. The first panel reports suppliers with direct or indirect exposure to car producers located either in Italy or elsewhere in the EU, based on the Ghosh inverse. The second panel restricts the measure to direct supplier relationships with those producers. The third panel restricts exposure to car producers located in Italy (both direct & indirect). "Threshold" denotes the minimum share of sales to car producers required for a firm to be included in the supply chain. VA represents exposed value added, expressed in million euros at current prices. All figures refer to the universe of limited-liability companies covered by Cerved.

Approximately half of the exposed value added generated by suppliers (4.5 billion euros) reflects direct relationships with car producers, while the remaining half originates from indirect upstream linkages. This pattern supports the relevance of accounting for indirect connections, as focusing solely on direct suppliers would capture only a fraction of the total supply chain value added.

4.b Sectoral composition

Panel A of Figure 4.1 reports the sectoral composition of the automotive supply chain, excluding final car producers and considering Italian suppliers of car producers located both in Italy and elsewhere in the EU. The manufacture of motor vehicles emerges as the leading industry, accounting for about 20% of total automotive supply chain value added. This primarily reflects the direct supply of vehicle parts and accessories[15]. The manufacture of fabricated metal products and that of basic metals follow as the second- and third-largest contributors (jointly 22% of aggregate supply chain value added), typically providing metallic components used in engines, body structures, and interior assemblies. Other relevant contributors include the manufacture of rubber and plastic products, covering both interior components and tires, and the manufacture of machinery and equipment.

Notably, several service sectors also display a non-negligible participation in the upstream automotive supply chain. ICT activities account for about 5% of total value added, reflecting the provision of software development and digital solutions that support both production processes and the increasing technological complexity of modern vehicles. Wholesale trade also represents a significant share of supply chain activity, largely reflecting intermediaries involved in the import and distribution of specialized inputs such as electronic components, mechanical parts, and electrical systems. Finally, transportation and logistics services represent another important pillar of the automotive value-chain.

Panel B of Figure 4.1 shows that the sectors contributing the most to the automotive supply chain are also those in which the supply chain is relatively more specialized. In these industries, the share of value added generated within the automotive supply chain is substantially larger than their weight in the overall economy, with only a few limited exceptions.

In Appendix B, Panel B of Figure B.1 isolates the portion of the value added in the supply chain that originates from direct and indirect connections with car producers located abroad

methodology, as discussed in Section 4.e.

[15] The baseline computation excludes final car producers in order to focus on the upstream supplier component of the supply chain. When final car producers located in Italy are included, the share of motor vehicles rises to 53%, as shown in Panel A of Figure B.1.

in the EU. The sectoral ranking remains broadly similar to the baseline, but this component is more concentrated in manufacturing industries, which typically have a higher propensity to export.

Figure 4.1. Sectoral composition of the automotive supply chain (2022)

(a) Composition **(b)** Specialization

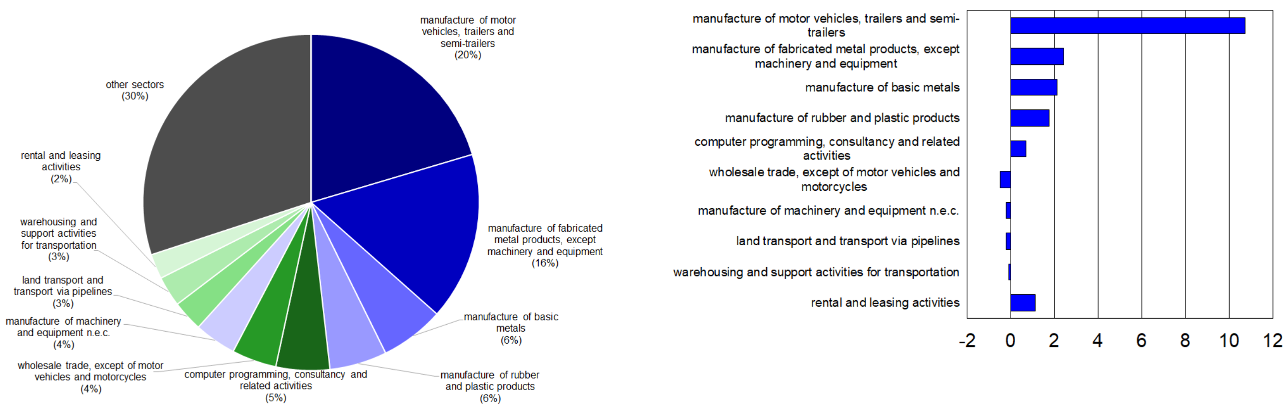

Notes: Top 10 sectors (ATECO 2-digit) by value added related to the automotive supply chain. Automotive-related value added is obtained by scaling all firms' total value added by their Ghosh inverse coefficient, and then aggregated across 2-digit sectors. Final carmakers are excluded. In panel a, sectoral automotive-related value added is expressed as a ratio of total automotive-related value added (in parentheses); manufacturing industries are in shades of blue, services in shades of green. In panel b, we report the ratio between one sector's contribution to automotive value added and its total contribution to total value added in the economy, minus 1; positive values indicate relative specialization.

4.c Regional composition

When it comes to the regional composition of the supply chain, a well-known limitation of Italian firm-level data is that location is recorded at the headquarters level, which may overlook production plants or offices situated elsewhere. We address this issue through a territorialization procedure based on business registry information: as mentioned in Section 2, for multi-plant firms, the Chamber of Commerce dataset reports the number of employees at the municipal level. We allocate each firm's value added across regions according to this distribution of employment, assuming that the spatial distribution of value added mirrors that of the workforce.

Panel A of Figure 4.2 shows the regional distribution of automotive-related value added, excluding final car producers. Automotive activity is strongly concentrated in the northern regions of Piedmont and Lombardy, which together account for more than half of the automotive supply chain. Other large industrial regions—such as Emilia-Romagna, Lazio, Campania and Veneto—also contribute substantially, though to a lesser extent. In some cases, these sizable contributions primarily reflect the overall scale of the regional economy; in others, such as Piedmont and Campania, they are also related to the presence of major car manufacturing plants[16]. This is also the case for some relatively small regions such as Abruzzo, where important production facilities are located: Panel b of Figure 4.2 shows that Abruzzo's contribution to automotive value added is disproportionately large relative to its share in the national economy.

In Appendix B, Panel B of Figure B.2 focuses on the value added attributable to observable links with car producers located abroad in the EU. The geographic distribution of this component remains strongly concentrated in Northern Italy. Lombardy plays a particularly prominent role, accounting for 35 percent of the associated value added, consistent with its position as one of Italy's leading exporting regions.

Figure 4.2. Regional composition of the automotive supply chain (2022)

(a) Composition

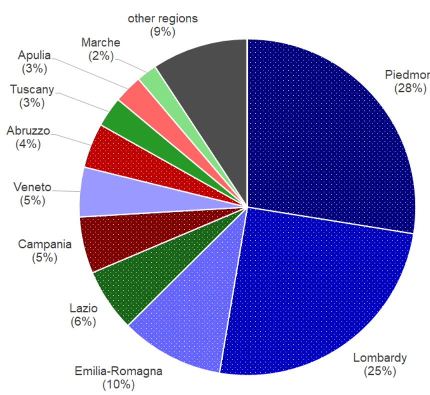

(b) Specialization

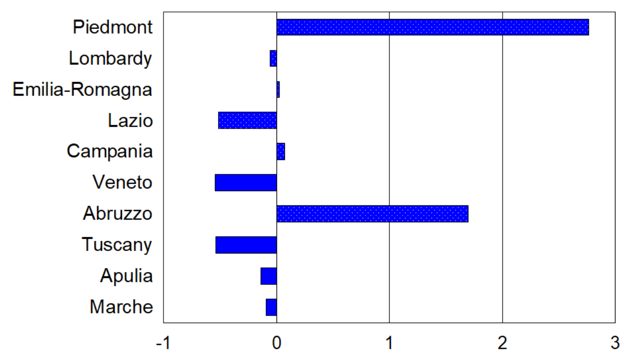

Notes: Top 10 NUTS-2 regions by value added related to the automotive supply chain. Automotive-related value added is obtained by scaling all firms' total value added by their Ghosh inverse coefficient, and then aggregated across regions. Final carmakers are excluded. In panel a, regional automotive-related value added is expressed as a ratio of total automotive-related value added (in parentheses); Northern regions are in shades of blue, Central in green and Southern in red. In panel b, we report the ratio between one region's contribution to automotive value added and its total contribution to total value added in the national economy, minus 1; positive values indicate relative specialization. In both panels, regions where final car production plants are localized are dotted.

4.d Dynamics of automotive sales over time

Beyond the static picture provided in the previous sections, an informative perspective comes from examining how supply chain activity evolved over time (Figure 4.3). During the COVID-19 pandemic, automotive-related real value added experienced a contraction stronger than that observed for the overall national economy. The subsequent rebound was more moderate and insufficient to fully restore pre-2020 levels.

To better understand the sources of the downturn in supply chain activity, we further decom-

[16]This result is not mechanical, as final car producers are excluded from these figures. When they are included, as in Panel A of Figure B.2, we observe a more pronounced concentration of the supply chain in regions hosting car manufacturing plants, with Emilia-Romagna emerging as the leading region due to the presence of major carmakers.

pose the total change into intensive and extensive margins. This decomposition shows that the intensive margin, reflecting incumbent firms that remain part of the automotive supply chain throughout the entire period, contributes positively to overall value added growth. In other words, conditional on maintaining a non-zero presence in the supply chain over the whole period, firms increased their output destined to the automotive sector, thereby partly offsetting the aggregate decline. By contrast, the extensive margin contributes negatively, reinforcing the downward trend. We distinguish two forces within the extensive margin:

1. *Entry–exit from the supply chain*: firms that either begin or cease trading with the automotive sector after 2019 while remaining active throughout the entire period.
2. *Entry–exit from the market*: firms entering or leaving the Fattel dataset altogether, capturing dynamics related to business demography rather than changes in supply chain participation alone.

Figure 4.3. Value added dynamics (2019=100)

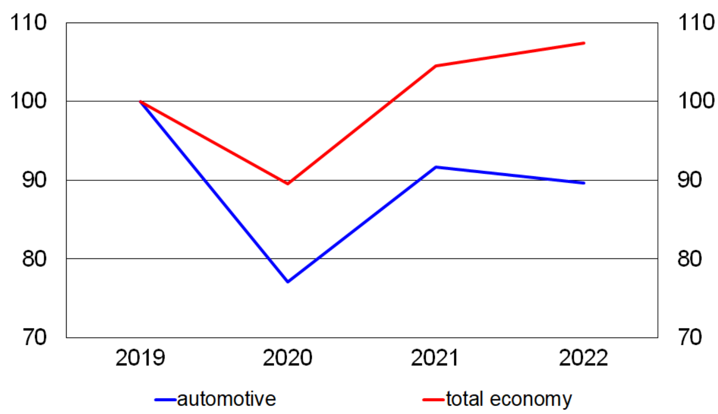

Notes: Automotive-related value added is obtained by scaling all firms' total value added by their Ghosh inverse coefficient. "Total economy" refers to firms' total value added from Cerved. Monetary aggregates are deflated using sector-specific producer price indices by Istat.

The former component displays a sizable reduction in automotive-related value added, while the latter shows only a negligible increase. Overall, this decomposition indicates that the contraction in automotive supply chain activity is not driven by underperformance among incumbent firms. The decline is mainly driven by firms that remain active in the economy but discontinue their relationships with car producers, suggesting a shift of activity away from the automotive sector. Entry of new suppliers only partly offsets this exit from the supply chain.

4.e Comparison with previous studies

Several estimates of the Italian automotive supply chain exist in the literature, differing substantially in methodology, data sources and resulting figures. Table 4.3 summarizes the main findings of such studies, whose estimates vary widely and are generally larger than those reported in our paper. This variation reflects differences in the perimeter being measured, while the gap with our estimates mainly reflects the structure of our analysis, which differs from previous work along three key dimensions.

First, our methodology focuses exclusively on the upstream supply chain: we trace sales linkages from Italian firms to car manufacturers, thereby excluding downstream activities such as dealerships, after-sales services, and the distribution of vehicles produced abroad.

Second, we quantify participation not only along the extensive margin, but also in terms of its intensity, as captured by each firm's Ghosh inverse coefficient. This distinction is empirically important: although many firms have some direct or indirect link to car producers, for most of them the corresponding sales share is very small. In particular, Table 4.2 shows that almost 150,000 firms display some exposure when no threshold is imposed, whereas this number falls to roughly 8,000 once we restrict attention to firms whose participation share exceeds 10 percent. Weighting firms' value added by their exposure coefficient is therefore essential to obtain a realistic measure of the supply chain; otherwise, treating all connected firms as fully part of the automotive perimeter would mechanically overstate its economic size.

Third, we rely on microdata capturing actual firm-to-firm transactions, derived from electronic invoices, rather than self-reported survey responses or sector-level aggregates.

Previous studies relying on firm-level data tend to produce larger estimates than those in this section. This is either because they only measure the extensive margin, or because they derive participation and turnover from self-reported questionnaires administered to a pre-selected sample of firms deemed relevant to the automotive sector, including downstream firms. By contrast, estimates based on sector-level data, such as Orame et al. (2025), assign exposure uniformly to all firms within a given ATECO division, implicitly assuming that every firm in, e.g., metal manufacturing contributes proportionally to automotive production.

Table 4.3. Main mappings of the automotive production network

Source	Year	Supply chain identification	Firms	Value Added	Turnover	Employees
Panel A. Our estimates of the automotive production network (1)						
	2022	Identification based on firm-to-firm production network linkages, capturing both direct and indirect sales to domestic and EU car manufacturers through the Ghosh inverse	8,172	15.6	95.3	-
Panel B. Existing mappings from official and industry sources						
ISTAT Permanent Census of Enterprises (2)	2021	Companies reporting participation in the "Rubber transport vehicles" supply chain in a specific question of the 2022 Permanent Census of Enterprises	102,198	110.3	-	1,704,103
		- *of which*: estimate of values exclusively attributable to the automotive supply chain	-	59.7	-	975,803
A. Orame, G. Cariola and G. Viggiano (2025) (3)	2022	Total direct and indirect contributions to global demand of ATECO 29, calculated on total non-financial private sector	99,284	45.2	204.2	627,782
TEA Report on Transformations of the Italian Automotive Ecosystem (TEA)	2021	ATECO 29 companies and another 40 divisions deemed relevant to the automotive sector and electric mobility; discretionary inclusion of additional companies, based on information collected also through interviews	2,400	-	100.0	255,000
ANFIA Study Center	2022	Production companies (45), component companies and other unidentified companies	5,528	-	86.2	273,600
		- *of which*: Automotive component supply chain	2,167	-	55.9	166,813
ISTAT Frame	2022	ATECO 29	2,356	15.7	77.2	168,730

Source: elaborations on Istat, ANFIA and TEA data.
Monetary aggregates (value added and turnover) are expressed in billions of euros. The number of firms and employees is reported in units.
Notes: (1) The value added and turnover figures include both direct and indirect suppliers as well as final producers; the number of firms referes to producers and suppliers with a significant exposure to the automotive industry (see section 4.a). (2) Our elaborations on aggregate data distributed by ISTAT. – (3) Identification, through input-output matrices, of the contribution of various sectors (selected ex-ante) to the production of ATECO 29. The calculation based on direct contribution considers, in addition to the total of ATECO 29 code values, the quantity of outputs from ATECO sectors directly used as inputs in automotive activities; the "direct and first-degree indirect method" also includes the contribution attributable to sales of goods to sectors that in turn are suppliers of intermediate goods to the automotive sector, while the total contribution considers each direct and indirect contribution of any degree. Calculations are made using 2022 data, while input-output matrices are updated to 2018.

A characterizing difference shared by previous studies is that they capture exposure to the automotive sector broadly defined, rather than participation in the upstream supply chain specifically. In fact, previous work includes leasing companies, and after-sales service providers—activities that are commercially related to automobiles but do not contribute to domestic vehicle production. Similarly, input-output methods trace linkages to final demand for ATECO 29 products regardless of whether that demand originates from domestic or foreign manufacturers. Our transaction-based approach, by contrast, identifies only firms whose sales—directly or through intermediate customers—ultimately reach domestic and EU car producers, thereby providing a precise, albeit smaller, mapping of the supply chain.

These methodological distinctions explain why our estimates are substantially smaller than those reported by previous studies. When it comes to sectors, the gap is largest for trade and services, where previous studies include downstream commercial activities that fall outside our upstream-focused perimeter. For manufacturing, estimates converge more closely, suggesting that differences are concentrated in the treatment of non-manufacturing activities rather than in the identification of the industrial core of the supply chain.

5 Characteristics of the supply chain

In this section, we characterize firms in the automotive supply chain along four dimensions: size, structure of sales, productivity and innovation, and labor outcomes. For each outcome, we estimate an «automotive premium» by comparing supply chain firms with other firms operating in the same sector and region, and with similar employment and capital intensity.

In order to detect if companies belonging to the automotive industry significantly differ from the rest along on some dimensions, we exploit the richness of our firm-level dataset to run OLS regression models of the following type[17]:

$$y_{i,t} = \alpha + \beta \times \mathbf{1}_{i,t}(\text{Automotive}) + \gamma \text{Empl}_{i,t-1} + \rho \frac{\text{Capital}_{i,t-1}}{\text{Empl}_{i,t-1}} + \delta_t + \delta_k + \delta_r + \delta_{rt} + \varepsilon_i$$

where $\mathbf{1}_{it}$(Automotive) is an indicator variable for firms participating in the supply chain (that is, with a Ghosh coefficient exceeding 10%, as illustrated in Section 4.a). Thus, the coefficient β should estimate an «automotive premium»: when the outcome variable y is expressed in logarithm, β expresses the percentage difference in the outcome variable between firms participating in the supply chain and those that do not. Since there can be significant size differences between the two groups, we control for employment (except for the regressions with employment as outcome) and for capital intensity, both at $t-1$; furthermore, we include year, NUTS-2 region, industry and industry $\times$ year effects.

Firm size. Table 5.1 examines differences in firm size. On average, firms participating in the automotive supply chain are about 16% larger in terms of total sales (Fattel plus exports).

[17] This regression framework is similar to the one employed in Cariola et al. (2025).

They also exhibit a sizeable revenue premium in Cerved data (5%), as well as higher value added (9%) and greater stocks of tangible and intangible assets. In addition, these firms are, on average, approximately 3% younger than non-automotive firms.

Table 5.1. Characteristics of automotive firms: firm size

Dependent variable:	Sales (log)	Revenues (log)	Value added (log)	Tangible assets (log)	Intangible assets (log)	Age (log)
	(1)	(2)	(3)	(4)	(5)	(6)
$\mathbf{1}_{it}$(Automotive)	0.160***	0.052***	0.094***	0.028**	0.045**	-0.034***
	(0.010)	(0.009)	(0.007)	(0.013)	(0.022)	(0.007)
$\mathbf{1}_{it}$(Direct supplier)	0.377***	0.307***	0.231***	0.040*	0.174***	0.018
	(0.017)	(0.015)	(0.011)	(0.022)	(0.038)	(0.013)
Employment$_{t-1}$ (log)	X	X	X	X	X	X
Capital intensity$_{t-1}$ (log)	X	X	X	X	X	X
Year effects	X	X	X	X	X	X
Industry effects	X	X	X	X	X	X
NUTS-2 effects	X	X	X	X	X	X
Industry × year effects	X	X	X	X	X	X
Adjusted R^2	0.60	0.66	0.73	0.65	0.34	0.31
Observations	1,919,545	1,915,473	1,820,172	1,828,418	1,210,715	1,919,166

Notes: Standard errors in parentheses are clustered at the firm level. Sales refer to total Fattel sales plus exports; the remaining dependent variables are from Cerved. Each coefficient is estimated in a separate regression. In the upper panel, the treatment variable is an indicator for firms belonging to the automotive supply chain, defined as firms with a Ghosh inverse coefficient above 10%. In the lower panel, the treatment variable is an indicator for direct suppliers, defined as firms belonging to the automotive supply chain that have positive direct sales to final carmakers. Controls, when indicated in the lower panel, include lagged employment, lagged capital intensity, year, industry, NUTS-2 region, and industry-by-year fixed effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

We then restrict attention to the subset of firms that sell directly to final carmakers. For these direct suppliers, size premia are even more pronounced across all measures. In terms of firm age, however, direct suppliers do not differ from other firms on average.

Structure of sales. Table 5.2 examines the structure of firms' sales. The sales premium associated with participation in the automotive supply chain is concentrated in the domestic market. Although supply chain firms exhibit a slightly higher export propensity than non-automotive firms, conditional on exporting they display 23% lower foreign sales. They also serve a smaller number of clients, resulting in a higher concentration of sales as measured by the Herfindahl–Hirschman Index (HHI). On the input side, these firms exhibit a similar concentration of suppliers relative to other firms.

When we restrict attention to direct automotive suppliers, we find a slightly different pattern. These firms display a significant premium in export propensity and a positive estimate in the level of foreign sales. They also exhibit high sales concentration, although to a lesser extent

Table 5.2. Characteristics of automotive firms: structure of sales

Dependent variable:	$\mathbb{1}$(Export >0)	Export (log)	N. of clients (log)	HHI Clients (log)	N. of suppliers (log)	HHI Suppliers (log)
	(1)	(2)	(3)	(4)	(5)	(6)
$\mathbf{1}_{it}$(Automotive)	0.013***	-0.228***	-0.420***	0.221***	-0.005	0.009
	(0.004)	(0.041)	(0.015)	(0.011)	(0.007)	(0.008)
$\mathbf{1}_{it}$(Direct supplier)	0.139***	0.122**	-0.343***	0.191***	0.144***	-0.081***
	(0.007)	(0.055)	(0.025)	(0.019)	(0.010)	(0.014)
Employment$_{t-1}$ (log)	X	X	X	X	X	X
Capital intensity$_{t-1}$ (log)	X	X	X	X	X	X
Year effects	X	X	X	X	X	X
Industry effects	X	X	X	X	X	X
NUTS-2 effects	X	X	X	X	X	X
Industry × year effects	X	X	X	X	X	X
Adjusted R^2	0.33	0.36	0.37	0.22	0.53	0.17
Observations	1,919,686	309,880	1,908,285	1,908,054	1,919,352	1,919,347

Notes: Standard errors in parentheses are clustered at the firm level. Each coefficient is estimated in a separate regression. In the upper panel, the treatment variable is an indicator for firms belonging to the automotive supply chain, defined as firms with a Ghosh inverse coefficient above 10%. In the lower panel, the treatment variable is an indicator for direct suppliers, defined as firms belonging to the automotive supply chain that have positive direct sales to final carmakers. Controls, when indicated in the lower panel, include lagged employment, lagged capital intensity, year, industry, NUTS-2 region, and industry-by-year fixed effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

than the broader group of supply chain firms, and maintain a larger pool of suppliers.

Productivity and innovation. Table 5.3 reports the results on productivity and innovation. Overall, firms in the automotive supply chain display significantly higher labor and total factor productivity (TFP), suggesting a more efficient combination of inputs. They also exhibit a markedly stronger propensity to invest in intangible assets, both in absolute terms and relative to total investment. The magnitude of this pattern remains unchanged when we focus on direct suppliers.

We further examine innovative activity directly. While the broader group of automotive firms exhibits a statistically significant - albeit small in magnitude - higher probability of holding patents, this correlation becomes more relevant once we restrict attention to direct suppliers: these firms are substantially more likely to engage in formal innovative activity, as captured by patent applications. This finding suggests that closer integration into the core of the supply chain may incentivize—or require—more systematic innovation efforts. Among firms that innovate (i.e., those holding at least one patent), the number of patents per firm does not differ significantly between automotive and non-automotive firms.

To characterize innovative activities in greater detail, we draw on the WIPO IPC technology concordance, allocating patents fractionally when they span multiple technological fields.

Table 5.3. Characteristics of automotive firms: productivity, investments, and innovation

Dependent variable:	VA per worker (log)	TFP (log)	Intangible invest. (log)	% intangible invest. (log)	$\mathbb{1}$(N. patents>0)	N. patents (log)
	(1)	(2)	(3)	(4)	(5)	(6)
$\mathbf{1}_{it}$(Automotive)	0.066***	0.015***	0.117***	0.068***	0.003***	0.065
	(0.006)	(0.003)	(0.023)	(0.019)	(0.001)	(0.065)
$\mathbf{1}_{it}$(Direct supplier)	0.155***	0.012***	0.220***	0.061**	0.017***	0.083
	(0.011)	(0.004)	(0.039)	(0.031)	(0.003)	(0.077)
Employment$_{t-1}$ (log)	X	X	X	X	X	X
Capital intensity$_{t-1}$ (log)	X	X	X	X	X	X
Year effects	X	X	X	X	X	X
Industry effects	X	X	X	X	X	X
NUTS-2 effects	X	X	X	X	X	X
Industry × year effects	X	X	X	X	X	X
Adjusted R^2	0.29	0.24	0.24	0.15	0.03	0.20
Observations	1,788,031	1,869,742	590,716	590,716	1,919,686	7,209

Notes: Standard errors in parentheses are clustered at the firm level. Each coefficient is estimated in a separate regression. In the upper panel, the treatment variable is an indicator for firms belonging to the automotive supply chain, defined as firms with a Ghosh inverse coefficient above 10%. In the lower panel, the treatment variable is an indicator for direct suppliers, defined as firms belonging to the automotive supply chain that have positive direct sales to final carmakers. Controls, when indicated in the lower panel, include lagged employment, lagged capital intensity, year, industry, NUTS-2 region, and industry-by-year fixed effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 5.1 shows that direct suppliers concentrate approximately 67% of their patents in Transport, Mechanical elements, Electrical machinery, as well as Engines, pumps and turbines. Relative to other innovative firms, they exhibit strong specialization in these four technological classes, as well as a comparatively higher degree of specialization in Environmental Technologies.

Figure 5.1. Distribution of patents by technological field

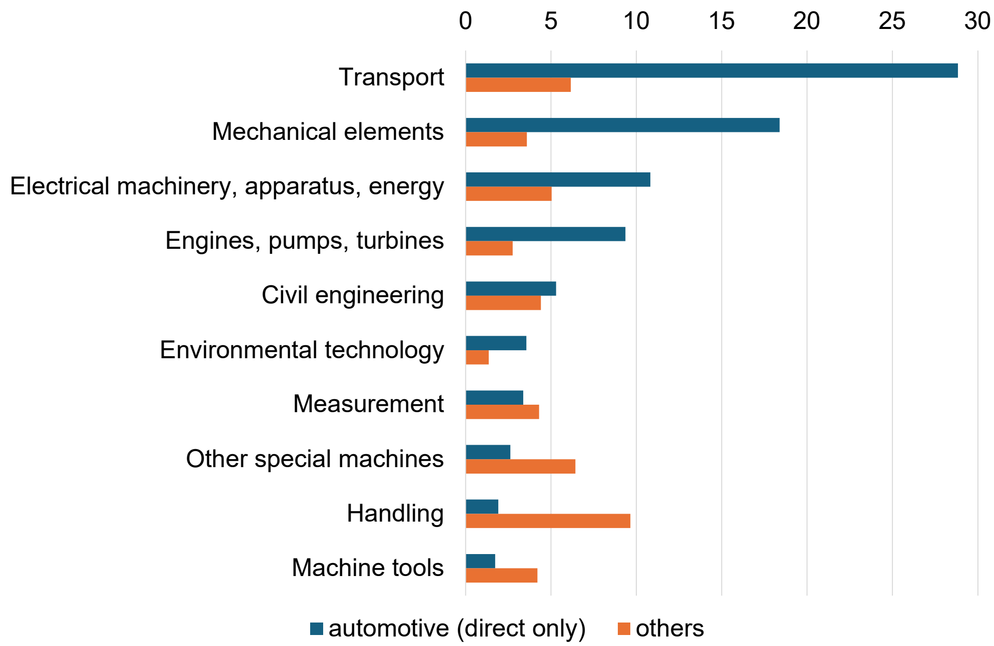

Notes: Distribution of patents by technological field, for automotive direct suppliers and for other firms. Patents are allocated fractionally when they span multiple technological fields. Top 10 fields for automotive suppliers.

Labor. Finally, Table 5.4 documents substantial differences in labor outcomes. Automotive firms employ roughly 51% more workers than comparable non-automotive firms and pay higher average wages, while also hiring a larger share of young employees. They additionally make more intensive use of short-time work schemes (*Cassa Integrazione Guadagni*, CIG), consistent with their greater exposure to sector-specific fluctuations. For direct suppliers, the employment and wage premia are even stronger.

Taken together, these regressions indicate that automotive supply chain firms tend to be larger, more productive, more reliant on tangible and intangible capital and associated with higher-wage employment than other firms, yet operate with a more concentrated client base and depend disproportionately on domestic demand. Size and productivity premia are more pronounced for direct suppliers, which may reflect either a selection mechanism—whereby supplying large and highly productive car manufacturers requires above-average size and efficiency, in a Melitz-type setting (Melitz, 2003)—or a learning mechanism, whereby proximity to large and productive buyers enhances firms' own performance (De Loecker, 2013).

However, disentangling these channels is beyond the scope of this paper.

Table 5.4. Characteristics of automotive firms: labor outcomes

Dependent variable:	Employment (log)	Under-35 empl. (log)	% under-35 (log)	Average wage (log)	Hours of CIG (log)
	(1)	(2)	(3)	(4)	(5)
$\mathbf{1}_{it}$(Automotive)	0.507***	0.054***	0.028***	0.065***	0.295***
	(0.016)	(0.010)	(0.009)	(0.004)	(0.016)
$\mathbf{1}_{it}$(Direct supplier)	1.225***	0.121***	0.045***	0.106***	0.363***
	(0.028)	(0.017)	(0.015)	(0.006)	(0.027)
Employment$_{t-1}$ (log)		X	X	X	X
Capital intensity$_{t-1}$ (log)	X	X	X	X	X
Year effects	X	X	X	X	X
Industry effects	X	X	X	X	X
NUTS-2 effects	X	X	X	X	X
Industry × year effects	X	X	X	X	X
Adjusted R^2	0.13	0.55	0.20	0.40	0.59
Observations	1,873,739	1,223,231	1,223,225	1,639,219	514,637

Notes: Standard errors in parentheses are clustered at the firm level. Each coefficient is estimated in a separate regression. In the upper panel, the treatment variable is an indicator for firms belonging to the automotive supply chain, defined as firms with a Ghosh inverse coefficient above 10%. In the lower panel, the treatment variable is an indicator for direct suppliers, defined as firms belonging to the automotive supply chain that have positive direct sales to final carmakers. Controls, when indicated in the lower panel, include lagged employment, lagged capital intensity, year, industry, NUTS-2 region, and industry-by-year fixed effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

6 The potential impact of a demand shock to car makers

In this Section we draw on the production network frameworks of Acemoglu et al. (2012) and Acemoglu et al. (2016) to assess how a demand shock to car producers propagates upstream along the supply chain. We adopt a deliberately simple setup in which firms combine labor and intermediate inputs taking their input–output linkages as predetermined and operate under constant returns to scale and perfect competition. This parsimonious structure can be calibrated at the firm level with our transaction data and delivers a transparent characterization of how a reduction in final demand is mechanically transmitted to suppliers through the network.

A central implication of the model is that the Ghosh inverse coefficients computed in Section 3 summarize the response of each upstream firm's output and value added to demand shocks affecting car makers, thus acting as a sufficient statistic for evaluating the short-run impact of such shocks on the automotive supply chain.

6.a A theoretical model

We consider a network economy where I firms are linked through input-output relationships. Each firm i produces output y_i using a Cobb-Douglas production function:

$$y_i = z_i l_i^{\alpha} \prod_{j=1}^{I} x_{ij}^{a_{ij}}, \tag{4}$$

where y_i is the output of firm i, z_i is a Hicks-neutral productivity term, l_i is labor input, x_{ij} is the quantity of input j used by firm i, and α and a_{ij} are the Cobb-Douglas exponents, with the constraint $\alpha + \sum_j a_{ij} = 1$ to ensure constant returns to scale. For the sake of simplicity, we assume $z_i = 1$ throughout the derivation. Given the Cobb-Douglas production function, value added at the firm level can be expressed as:

$$VA_i = (1 - \sum_j a_{ij}) y_i. \tag{5}$$

meaning that a log change in y_i will induce exactly the same log change in VA_i, i.e. $\frac{dVA_i}{VA_i} = \frac{dy_i}{y_i}$.

The market-clearing condition for firm i is:

$$y_i = \sum_j x_{ji} + c_i, \tag{6}$$

where c_i represents consumers' demand for firm i's output and x_{ji} represents the sales of intermediates of firm i to firm j.

We define the utility of the representative consumer as follows:

$$u(c_1, c_2, ...c_I, I) = (1 - l)^{\gamma} \prod_{i=1}^{I} c_i^{\beta_i} \tag{7}$$

where the coefficients β_i represent the consumption share of each good, with $\sum_{i=1}^{I} \beta_i = 1$, and γ governs the disutility of labor. Maximizing utility subject to the budget constraint $\sum_{i=1}^{I} p_i c_i = wl$ implies that, in the absence of taxation, labor supply is constant and equal to $\frac{1}{1+\gamma}$. Taking the wage as the numeraire, this also corresponds to the consumer's total income. Given these assumptions, Acemoglu et al. (2016) establish that equilibrium prices are pinned down by the coefficients of the production function; hence, shocks to final demand leave the price structure unchanged. Additionally, shock propagation is asymmetric, meaning that demand shocks propagate upstream while supply shocks propagate downstream (see proposition 1 and Appendix A in Acemoglu et al., 2016).

To derive the upstream propagation of a shock to final demand in our framework, we assume that the consumer experiences a shock to its preferences, i.e. a vector of changes $d\beta_i$ such that $\sum_i d\beta_i = 0$[18]. This corresponds to a reshuffling of expenditure shares across goods: consumers reduce the share of income allocated to some goods and increase it for others. A realistic example in our context is a reduction in the expenditure share on cars produced within the EU in favor of cars produced in foreign plants whose supply chain linkages with Italian firms are negligible. In such cases, the upward shift in demand for foreign goods does not generate any upstream domestic propagation, whereas the reduced expenditure on EU cars triggers an adjustment that travels through the domestic production network.

An alternative way to model a decline in final demand would be to introduce public demand as in Acemoglu et al. (2016) and consider a reduction in government purchases of cars (which could be interpreted as a decrease in subsidies for the purchase of new vehicles, a common policy in many countries). Under the assumptions of our model, the mechanism of upstream propagation would be essentially unchanged. However, formalizing this extension would require introducing either taxation, with additional effects operating through the government budget constraint, or public debt, which would call for an intertemporal framework. To keep the analysis parsimonious and focused on the propagation mechanism, we therefore model the shock as a shift in private expenditure shares.

Because of Cobb-Douglass preferences, expenditure on each good is given by

$$p_i c_i = \frac{\beta_i}{1+\gamma}, \tag{8}$$

so that the shock $d\beta_i$ induces

$$d(p_i c_i) = \frac{1}{1+\gamma}\, d\beta_i. \tag{9}$$

[18]In the firm-to-firm data, car sales to final users are often mediated by wholesalers and dealers rather than recorded as direct transactions between manufacturers and households. For the purposes of our framework, we nonetheless treat car producers as the relevant final-demand nodes. This is consistent with standard input–output accounting, where sales are attributed to the producing industry and wholesale/retail activities are typically captured through trade margins converting basic prices into purchasers' prices; see, for example, Eurostat (2026). Another difference relative to standard aggregate input–output tables is that firm-to-firm transactions also include investment goods.

Let us now examine how such a perturbation to final demand propagates upstream through the input-output network. Starting from the market-clearing condition and multiplying by p_i, total differentiation gives

$$\frac{d(p_i y_i)}{p_i y_i} \;=\; \frac{d(p_i c_i)}{p_i y_i} \;+\; \sum_j a_{ji} \frac{d(p_j y_j)}{p_i y_i}, \tag{10}$$

where Cobb–Douglas technology implies $a_{ji} = \frac{p_i x_{ji}}{p_j y_j}$. Defining instead

$$\hat{a}_{ji} = \frac{p_i x_{ji}}{p_i y_i}, \tag{11}$$

we can rewrite $a_{ji} = \hat{a}_{ji} \cdot \frac{p_i y_i}{p_j y_j}$ and obtain

$$\frac{d(p_i y_i)}{p_i y_i} \;=\; \frac{d(p_i c_i)}{p_i y_i} \;+\; \sum_j \hat{a}_{ji} \frac{d(p_j y_j)}{p_j y_j}. \tag{12}$$

Collecting terms in vector form yields

$$d \ln y = \hat{A}^T d \ln y + dC_y, \tag{13}$$

where dC_y is the vector whose generic element is equal to $\frac{1}{p_i y_i} \frac{1}{1+\gamma}\, d\beta_i$, i.e. the demand shock to firm i scaled down by its output, and $\hat{A}$ contains the normalized sales shares $\hat{a}_{ji}$. Solving for $d \ln y$ gives

$$d \ln y = \hat{H}\, dC_y, \qquad \hat{H} = (I - \hat{A}^T)^{-1}. \tag{14}$$

where $\hat{H}$ is the Ghosh inverse of the economy. As pointed out in Section 3, each element h_{ji} corresponds to an infinite sequence of direct and indirect sales shares, i.e. the direct and indirect sales of firm i to firm j normalized by the total value of i's sales[19].

Combining equation (9) with the propagation mechanism implied by equation (14), we obtain:

$$\frac{d(p_i y_i)}{p_i y_i} = \sum_j h_{ji} \left(\frac{1}{p_j y_j} \cdot \frac{1}{1+\gamma} d\beta_j \right). \tag{15}$$

This expression shows that the log change in the output of an upstream supplier i is a

[19] This is evident if equation (14) is expanded through Newton's formula, i.e. $(I - \hat{A})^{-1} = I + \hat{A} + \hat{A}^2 + \hat{A}^3 + \cdots$

weighted sum of the shocks to downstream demand, where the weights h_{ji} capture the propagation mechanism through the production network and the term in parentheses represents the magnitude of the demand shock to firm j expressed as a fraction of its total sales.

6.b Interpretation of the model

The Ghosh inverse coefficients have implications that extend beyond their descriptive content. Under the modelling assumptions introduced in the previous section, they provide a sufficient statistic for the impact of demand disturbances on the domestic production network. The soundness of this interpretation, however, hinges on the validity of the assumptions of the model.

First, the static nature of the framework precludes the possibility that firms anticipate changes in final demand and adjust their sourcing or sales strategies before the shock materializes. In other words, the structure of the model rules out forward-looking behavior aimed at mitigating the impact of an expected contraction in downstream demand.

Second, the adoption of a Cobb–Douglas production function with constant returns implies fixed expenditure shares that coincide with the technological coefficients. As a consequence, following a demand shock suppliers do not adjust their input mix: intermediate-input shares remain predetermined.

Third, the assumption of perfect competition removes the possibility that firms adjust their markups in response to perturbations to final demand. Recent evidence by Conteduca and Panon (2025) indicates that firms tend to compress markups when hit by exogenous shocks, thereby partially absorbing the impact through prices rather than quantities. Our framework rules out this margin of adjustment and potentially overstates the extent of upstream propagation.

Despite these limitations, the framework is useful for our purposes because it combines a transparent propagation mechanism with a parsimonious structure that can be directly calibrated using firm-level input–output data. This allows us to quantify the upstream transmission of demand shocks and to document its heterogeneity across firms, sectors, and regions. Given the assumptions of the model, its structural interpretation is most appropriate for short-run, unexpected demand perturbations, when the production network can be reasonably taken as fixed and firms have limited scope to adjust their sourcing patterns, customer base, or production technologies. Over longer horizons, as firms adapt their customer base and reorganize their supply chains, the model is not designed to capture the resulting dynamics; its projections should therefore be viewed as upper bounds to the true adjustment.

6.c Simulation results

We now bring the model to the data to quantify the upstream propagation of a shock to the final demand for cars. To this end, we calibrate the framework to the 2022 firm-level production network and feed into the model a 1% demand shock to EU car manufacturers[20] [21]. Given equation (15), this shock maps directly into predicted percent changes in output and value added across all upstream suppliers through the firm-specific Ghosh inverse coefficients. Because the Cobb–Douglas technology implies proportionality between output and value added, we express all firm-level effects in terms of value added, which we observe from balance-sheet data for the universe of limited liability companies.

We first report the aggregate implications of the shock and then examine how its effects vary across territories, sectors and firm-size classes.

6.c.1 Aggregate results

According to our model calibrated to 2022 data, a 1% reduction in final demand for cars produced in the EU has, in the first place, first-order effects: it reduces the value added of Italian car manufacturers by approximately 63 million euros. Then the shock propagates upstream through the production network, generating an additional 93 million euros in losses among domestic direct and indirect suppliers. The total aggregate effect therefore amounts to a decline in value added of about 156 million euros. This implies a network multiplier of 2.48, defined as the ratio of total losses (156) to first-order losses borne by final car manufacturers (63). If we restrict attention to direct, first-tier suppliers only, we would capture an upstream loss of only about 45 million euros, implying a network multiplier of 1.71 and substantially understating the economy-wide exposure to a drop in car demand.

[20] As detailed in equation 14, the magnitude of the shock is scaled down by final producers' output (i.e. it should be interpreted as $\frac{dC}{Y}$ instead of $\frac{dC}{C}$); if all output of car manufacturers is absorbed by final demand, $\frac{dC}{Y} = \frac{dC}{C}$.

[21] In practice, demand changes in the automotive sector are substantially larger. According to ACEA data, over our sample period, EU car registrations declined by roughly 30%, while even in the pre-COVID years registrations typically fluctuated by about 5% per year on average. This decline was accompanied by a steady increase in imports of cars from outside the EU: the extra-EU import share over domestic car demand rose by about 5 percentage points between 2018 and 2023, exceeding 40% (De Santis et al., 2024). In our framework, given the linearity of the model, the effects of an n% shock can be obtained by multiplying the baseline results by n.

6.c.2 Projected impact by sector

Figure 6.1 reports the projected percentage decline in sectoral value added resulting from a 1% contraction in final demand for cars produced within the EU , restricting attention to upstream propagation (i.e., excluding the direct effect on car manufacturers). Unsurprisingly, the manufacture of automotive components exhibit by far the largest impact, with almost 0.3% of its value added ultimately affected by shocks to final car producers, which implies a 30% pass-through.

Figure 6.1. Projected impact of a 1% drop in car demand on the supply chain, by sector (% of sectoral value added)

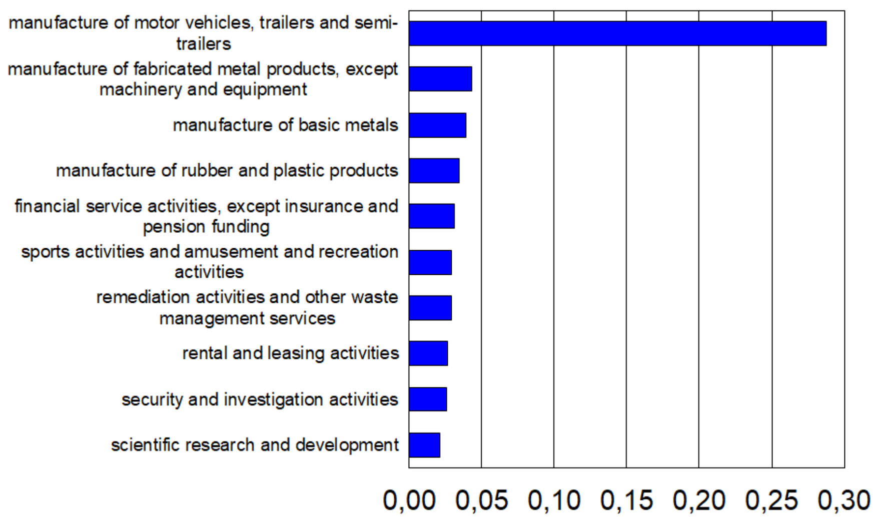

Notes: Sectoral (2-digit ATECO) impact is measured as the projected decline in value added induced by a 1% reduction in final demand for cars, expressed as a share of sectoral value added in 2022. For each sector, the impact is obtained by aggregating firm-level shocks proportional to the Ghosh inverse coefficients and 2022 value added. Final carmakers are excluded. Top 10 sectors by impact.

While sizeable, this share is far from encompassing the entire sector. Even within the ATECO 29 sector (manufacturing of motor vehicles), many firms sell a substantial portion of their output to customers outside the automotive industry. This pattern, which is also visible in the raw transaction data for domestic sales, suggests that many component producers operate as multiproduct suppliers serving several manufacturing sectors. This finding is relevant for the current policy debate, which often treats the whole motor vehicle manufacturing sector as fully dependent on final demand for cars. Our firm-level network approach shows instead that exposure to the demand of motor vehicles is substantial but far from prevalent even within the core sector of automotive manufacturing.

Other than that, three traditional industrial sectors—namely fabricated metal products, basic metals, and rubber and plastic products—rank among the main contributors to the supply chain and are likewise significantly affected, although to a smaller extent (a 1% shock to final demand leads to almost a 0.05% reduction in these industries' value added).

Among services, several activities also display non-negligible exposure to car demand shocks. Financial activities, in particular, appear significantly affected. Rental and leasing activities—which include agencies providing vehicle-leasing services on behalf of car makers—are, unsurprisingly, substantially impacted. At the same time, other smaller service sectors that are less obviously associated with automotive manufacturing also exhibit meaningful exposure, including R&D and security activities.

6.c.3 Projected impact by NUTS-3 region

When turning to the geographic heterogeneity of the projected effect, an interesting pattern emerges, as the magnitude of the projected decline in value added varies substantially across Italian NUTS-3 regions. The provinces experiencing the largest impact are those hosting major car-production facilities[22], indicating that upstream suppliers remain heavily concentrated around final assembly plants. This pattern is evident not only in Turin, Italy's traditional automotive hub, but also in several smaller and more peripheral areas, such as Potenza, Isernia, and other provinces in Southern and inland Italy, where large manufacturing plants or historically automotive-specialized clusters are located. In these territories, the projected impact of a 1% shock to final car demand exceeds 0.05% of total provincial value added of the non-financial and non-real-estate private sector, underscoring their pronounced dependence on automotive demand (Figure 6.2).

Taken together with the results from Section 4, where the largest overall contributions to the automotive supply chain originate from major industrial regions such as Piedmont, Lombardy and Emilia-Romagna, this evidence points to the coexistence of two distinct forces shaping the territorial organization of the supply chain.

On the one hand, scale economies and agglomeration forces foster the concentration of upstream suppliers in relatively larger regions. On the other, local idiosyncrasies and long-standing productive specializations generate clusters of high impact in specific provinces where automotive production plants are located, even when these areas are relatively small in economic size[23]. Together, these forces produce a markedly uneven geographic distribution of the supply chain, with both large industrial regions and smaller specialized territories playing key roles in the propagation of shocks to the automotive industry.

[22] Throughout this section, car producers are excluded from the analysis; hence, the estimated impact reflects solely the propagation of the shock through upstream suppliers and is not mechanically driven by the direct decline in car manufacturers' output.

[23] These two competing forces are also explored in the literature about the regional impact of international disruptions; see Borin et al. (2023) and Panon et al. (2024) for further details.

Figure 6.2. Projected impact of a 1% drop in car demand on the supply chain, by province (% of province value added)

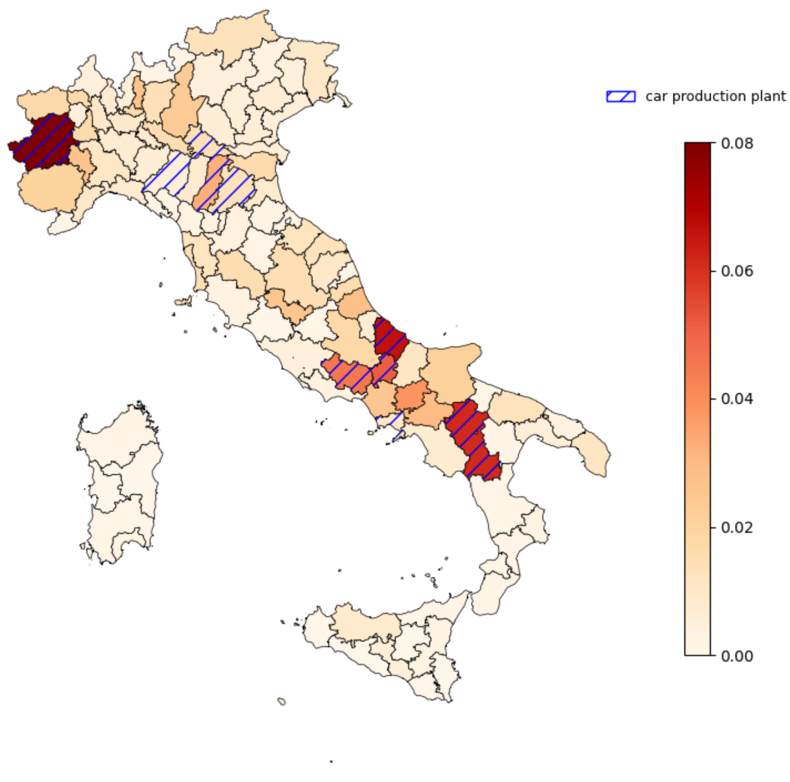

Notes: Provincial (NUTS-3) impact is measured as the projected decline in value added induced by a 1% reduction in final demand for cars, expressed as a share of provincial value added of the non-financial and non-real-estate private sector in 2022. For each province, the impact is obtained by aggregating firm-level shocks proportional to the Ghosh inverse coefficients and 2022 value added. Final carmakers are excluded. Regions hosting final car production plants are shaded.

6.c.4 Projected impact by firm size

We finally examine how the impact of the shock varies by firm size by aggregating the firm-level effect by quintiles of total revenues[24]. Panel a of Figure 6.3 shows that, following a 1% drop in car demand, firms in the first quintile experience a value added drop close to 0.2% of total value added, four to five times larger than that of firms in the top two quintiles, where the effect remains below 0.05%. This result indicates that smaller firms are considerably more exposed to disruptions in automotive demand, reflecting their greater reliance on a narrow set of downstream buyers. However, the greatest absolute losses are borne by large firms, which account almost the totality of value added affected by the shock (panel b of Figure 6.3). This stylized fact highlights a central feature of the automotive supply chain:

[24] As documented in Section 4.a, the vast majority of firms have zero impact; we therefore restrict the analysis to firms with a strictly positive Ghosh inverse coefficient.

large firms account for most of the aggregate exposure, whereas smaller suppliers exhibit greater vulnerability.

Figure 6.3. Projected impact of a 1% drop in car demand on the supply chain, by firm-size quintile (2022)

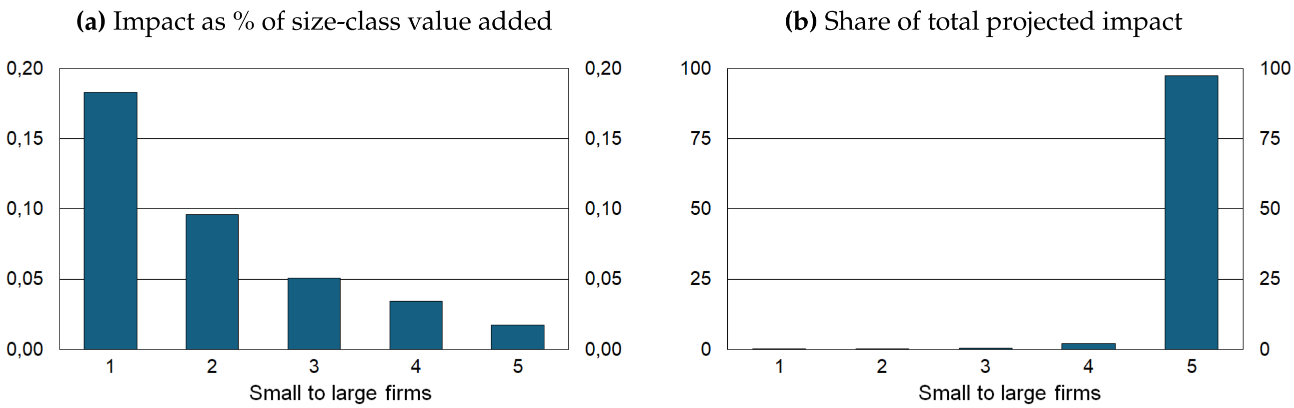

Notes: The figure reports the projected decline in value added induced by a 1% reduction in final demand for cars, by firm-size quintile. Panel a expresses the projected decline in value added as a percentage of total Cerved value added in that quintile. Panel b shows the distribution of the aggregate projected impact across size classes: each bar reports the share of the total decline accounted for by the corresponding quintile, with the total impact normalized to 100. For each size class, the projected decline in value added is computed by aggregating firm-level shocks proportional to the Ghosh inverse coefficients and firms' 2022 value added. Firms with no automotive-related value added and final car manufacturers are excluded.

7 Concluding remarks

In this paper, we use a novel dataset on firm-to-firm transactions to provide a detailed mapping of the automotive supply chain. The exceptional granularity of these data allows us to move beyond standard input–output aggregates and reconstruct supply chain participation at the level of individual firms. Our estimates indicate that in 2022 more than 8,000 Italian firms recorded a non-negligible share of direct or indirect sales to European car producers, and that over 15 billion euros of value added were directly or indirectly linked to the automotive sector.

Automotive-related transactions are concentrated in sectors traditionally involved in the production of vehicle parts and components—most notably motor-vehicle components, metals, rubber and plastics. At the same time, the supply chain extends beyond manufacturing, with several service activities accounting for a meaningful share of automotive-related value added; among these, ICT services play a relevant role, as the technological transformation of vehicles and production processes is likely to increase the importance of digital capabilities in the automotive industry. In absolute terms, the largest contribution originates from the major industrial regions of Central and Northern Italy, which host a substantial share of the country's manufacturing capacity.

Looking at relative dependence instead of absolute contributions reveals an equally important pattern. Alongside core manufacturing industries, several smaller sectors—particularly in services—display a relatively strong reliance on the automotive industry. Similarly, while large companies account for the majority of automotive-related value added, smaller suppliers tend to exhibit a higher degree of dependence on the sector, reflecting their more limited diversification across clients.

Our regression analysis further shows that, conditional on sector, region, employment and capital intensity, firms integrated into the automotive supply chain are on average larger, more productive, and more intensive in intangible investment than comparable firms outside the network, while also displaying a lower degree of diversification. In addition, direct suppliers of car manufacturers exhibit a significantly higher probability of patenting, particularly in traditional fields such as mechanical and electronic technologies.

Overall, our findings carry several policy implications. Given the challenges currently faced by the automotive industry, potential downturns in the demand for cars produced in the EU would likely have stronger effects in smaller, less diversified regional economies that depend heavily on the sector. The territorial patterns documented in the paper also point to the coexistence of two forces shaping the geography of the supply chain: agglomeration effects that concentrate upstream suppliers in large industrial regions, and localized clusters that emerge around major vehicle production plants. As a result, relatively small provinces hosting assembly facilities may display disproportionately high exposure to sectoral shocks.

Our results also suggest that sector-based classifications may provide an incomplete picture of supply chain exposure. Many firms classified in the motor vehicle sector sell a substantial share of their output to customers outside the automotive industry, while several firms outside the sector are indirectly connected to car manufacturers through intermediate buyers. Policies targeting industries defined by standard sectoral classifications may therefore not fully account for the complexity of supply chain structures and cross-sectoral linkages.

Finally, our analysis highlights the importance of accounting for indirect relationships when measuring supply chain exposure. Focusing exclusively on first-tier suppliers captures only a portion of the network and can substantially underestimate the magnitude of upstream linkages. Firm-level transaction data allows to recover these higher-order connections and provide a more accurate assessment of how shocks propagate through production networks.

References

ACEA (2025). The automobile industry pocket guide 2025/2026. Publication providing comprehensive facts and figures on the automotive industry in Europe.

Acemoglu, D., Akcigit, U., and Kerr, W. (2016). Networks and the macroeconomy: An empirical exploration. *NBER Macroeconomics Annual*, 30(1):273–335.

Acemoglu, D., Carvalho, V. M., Ozdaglar, A., and Tahbaz-Salehi, A. (2012). The network origins of aggregate fluctuations. *Econometrica*, 80(5):1977–2016.

Astuti, V. and Piras, D. (2025). Italian business-to-business invoicing data: A network analysis. *mimeo*.

Baldwin, R. and Freeman, R. (2022). Risks and global supply chains: What we know and what we need to know. *Annual Review of Economics*, 14(1):153–180.

Baldwin, R., Freeman, R., and Theodorakopoulos, A. (2022). Horses for courses: Measuring foreign supply chain exposure. NBER Working Paper 30525, National Bureau of Economic Research.

Baldwin, R. and Lopez-Gonzalez, J. (2015). Supply-chain trade: A portrait of global patterns and several testable hypotheses. *Journal of International Economics*, 98:1–13.

Baqaee, D. R. and Farhi, E. (2019). The macroeconomic impact of microeconomic shocks: Beyond hulten's theorem. *Econometrica*, 87(4):1155–1203.

Baqaee, D. R. and Rubbo, E. (2023). Nonlinear production networks with an application to the covid-19 shock. *Econometrica*, 91(6):2391–2436.

Benecchi, A., Borin, A., Conteduca, F. P., Leone, F., Mancini, M., Modugno, L., Mongardini, M., Papini, G., and Zoi, P. (2025). Us tariffs and italian firms: Mapping network exposure and early impact from survey data. Bank of Italy, mimeo.

Bernard, A. B., Dhyne, E., Magerman, G., Manova, K., and Moxnes, A. (2022). The origins of firm heterogeneity: A production network approach. *Journal of Political Economy*, 130(7):1765–1804.

Bertolotti, F., Linarello, A., and Zoi, P. (2025). Shock propagation and economic policies in the italian production network. Manuscript.

Bijnens, G., Duprez, C., Konings, J., and Van Cayseele, P. (2025). Propagation of supply and demand shocks in firm-to-firm networks. Working paper, National Bank of Belgium.

Boehm, C. E., Flaaen, A., and Pandalai-Nayar, N. (2019). Input linkages and the transmission of shocks: Firm-level evidence from the 2011 tōhoku earthquake. *American Economic Review*, 109(1):193–231.

Borin, A., Cariola, G., Gentili, E., Linarello, A., Mancini, M., Padellini, T., Panon, L., and Sette, E. (2023). Inputs in geopolitical distress: A risk assessment based on micro data. Technical Report 819, Bank of Italy.

Cariola, G., Carnevali, G. B., Linarello, A., Manaresi, F., Mirenda, L., Russo, E., Sartori, M., and Viggiano, G. (2025). Multinational enterprises in italy: Insights from firm-level data. Technical Report 894, Bank of Italy. SSRN working paper.

Carvalho, V. M., Nirei, M., Saito, Y. U., and Tahbaz-Salehi, A. (2021). Supply chain disruptions: Evidence from the great east japan earthquake. *Quarterly Journal of Economics*, 136(2):1255–1321.

Conteduca, F. P. and Panon, L. (2025). Natural disasters and markups. Temi di Discussione (Working Papers) 1496, Banca d'Italia.

De Loecker, J. (2013). Detecting learning by exporting. *American Economic Journal: Microeconomics*, 5(3):1–21.

De Santis, R. A., Di Nino, V., Furbach, N., Neumann, U., and Neves, P. (2024). Will the euro area car sector recover? *Economic Bulletin Boxes*, 4(None):None.

Dechezleprêtre, A., Díaz, L., Fadic, M., and Lalanne, G. (2023). How the green and digital transitions are reshaping the automotive ecosystem.

Dhyne, E., Kikkawa, A. K., Mogstad, M., and Tintelnot, F. (2021). Trade and domestic production networks. *Review of Economic Studies*, 88(2):643–668.

Dhyne, E., Magerman, G., and Rubinova, S. (2015). The belgian production network 2002-2012. Working Paper Research 288, National Bank of Belgium.

Draghi, M. (2024). The future of european competitiveness. Technical report, European Commission.

European Commission (2025). Industrial action plan for the european automotive sector. COM(2025) 95.

Eurostat (2026). Methodology – esa supply, use and input-output tables. `https://ec.europa.eu/eurostat/web/esa-supply-use-input-tables/methodology`. Accessed on March 19, 2026.

Ghosh, A. (1958). Input-output approach in an allocation system. *Economica*, 25(97):58–64.

Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6):1695–1725.

OECD (2024). The future of the automotive value chain: Implications for fdi-sme linkages.

Orame, A., Cariola, G., and Viggiano, G. (2025). The italian automotive sector in the green transition: empirical evidence and insights from industry experts. Questioni di Economia e Finanza (Occasional Papers) 911, Bank of Italy, Economic Research and International Relations Area.

Panon, L., Lebastard, L., Mancini, M., Borin, A., Caka, P., Cariola, G., Essers, D., Gentili, E., Linarello, A., Padellini, T., Requena, F., and Timini, J. (2024). Inputs in distress: geoeconomic fragmentation and firms' sourcing. Technical Report 861, Banca d'Italia.

Pavlínek, P. (2020). Restructuring and internationalization of the european automotive industry. *Journal of Economic Geography*, 20(2):509–541.

Pianeselli, D. and Orame, A. (2023). Thinking the green transition: Evidence from the automotive industry. Occasional Papers 767, Bank of Italy, Rome.

Sturgeon, T. J., Van Biesebroeck, J., and Gereffi, G. (2008). Value chains, networks and clusters: Reframing the global automotive industry. *Journal of Economic Geography*, 8(3):297–321.

Appendix A Computing firm-level exposure

This appendix provides additional details on the recursive procedure used to compute firm-level exposure coefficients, and reports robustness checks on two key implementation choices: the maximum number of iterations and the stopping threshold for negligible indirect trade flows.

Our empirical objective is not to recover the full matrix of firm-level Ghosh coefficients, but only the subset of coefficients that links firms to the selected set of downstream automotive producers. To this end, we exploit the recursive structure of the Ghosh inverse and trace upstream sales linkages starting from the set of domestic and EU car manufacturers identified in Section 3.b.

More specifically, for each final producer j, we first identify its direct suppliers i and compute the corresponding direct sales shares a_{ji}. We then move one step further upstream: for each supplier identified at a given step, we retrieve its own suppliers and the corresponding sales shares. Proceeding recursively in this way, we construct the upstream paths that connect firms to final car manufacturers. Along each path, the contribution of a firm is given by the product of the direct sales shares associated with the sequence of links. For any firm i, summing these weighted paths leading to producer j yields the corresponding firm-to-producer exposure coefficient. Aggregating these coefficients across all selected car manufacturers provides our measure of firm i's participation in the automotive supply chain, consistently with the Neumann-series representation of the Ghosh inverse (see equation 2). In practice, rather than computing the full inverse matrix, we recover only the coefficients that are relevant for the automotive application. When moving from iteration 9 to iteration 12, the cumulative increase in captured sales linkages is less than 1 percent of total direct sales to final producers. This suggests that extending the recursion further would add only marginal additional linkages.

In addition, to avoid unnecessary computational burden, we stop the recursion whenever the additional indirect trade flow identified at a given step is smaller than €1,000. To assess the sensitivity of the results to this choice, we raise the stopping threshold to €3,000 and compare the resulting exposure coefficients with the baseline. In the baseline specification, we cap the recursion at 12 iterations; that is, we allow sales linkages up to twelve degrees of distance from final producers. This choice is motivated by the fact that the additional amount of indirect sales captured in later iterations becomes very limited, as shown in Figure A.1.

Figure A.1. Indirect sales captured at each iteration, over direct sales (%)

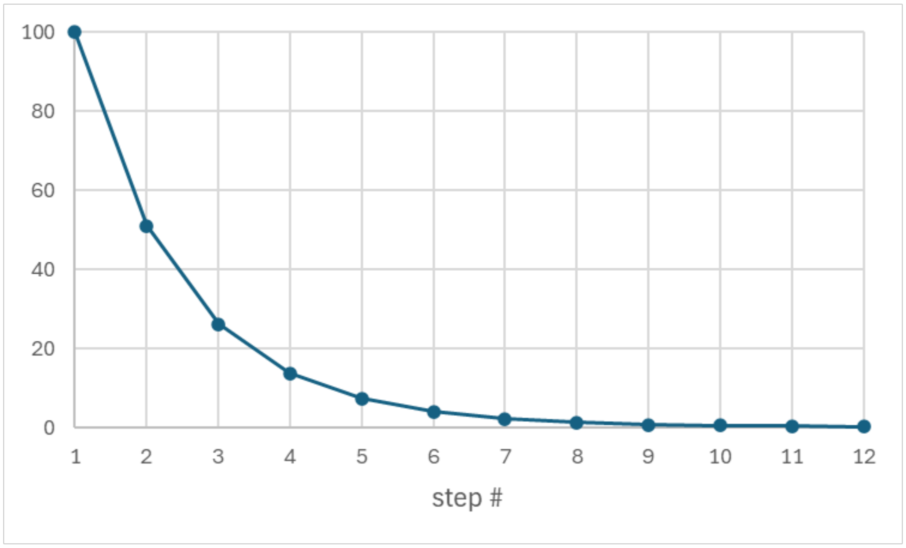

Notes: Additional sales linkages captured at each iteration, as a share of total direct sales to car manufacturers.

Figure A.2 shows that the results are only minimally affected. For firms with less than 1 percent exposure to the automotive sector, the higher threshold still captures about 95 percent of baseline automotive sales. For firms with greater exposure, the corresponding share is close to 100 percent.

Figure A.2. Trade linkages with a €3,000 threshold compared to baseline (%)

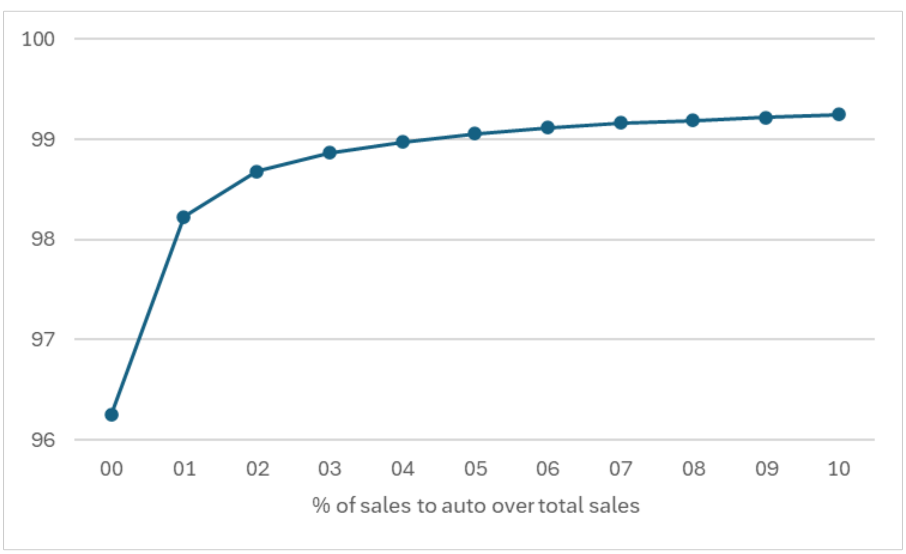

Notes: Sales linkages to car producers with a stopping threshold at €3,000 by automotive exposure, as a fraction of the baseline.

Appendix B Additional graphs on sectoral and regional composition

Figure B.1. Sectoral composition of the automotive supply chain (2022)

(a) Including final car producers **(b)** Suppliers of non-domestic EU producers

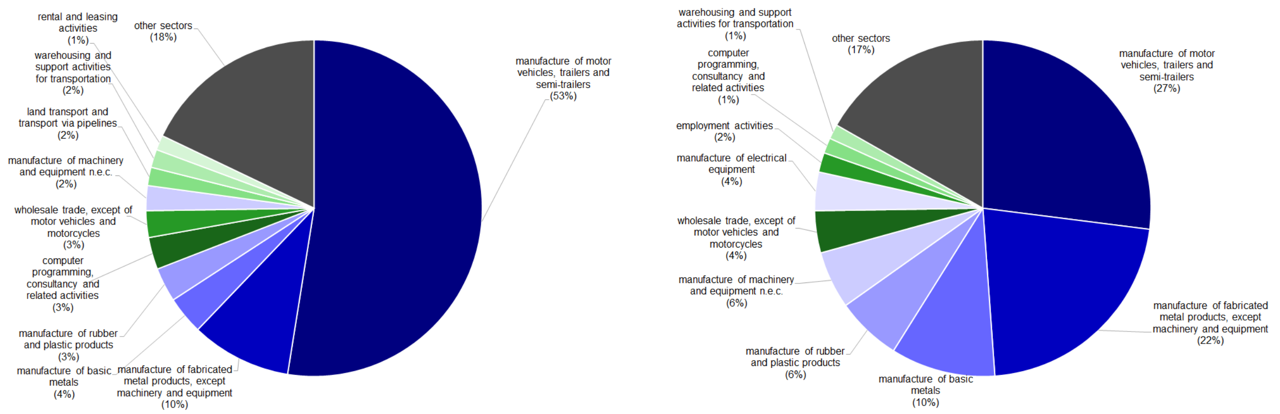

Notes: Panel a reports the top 10 sectors (ATECO 2-digit) by automotive-related value added, including final car producers located in Italy. Panel b reports the top 10 sectors among Italian suppliers linked to final carmakers located elsewhere in the EU. Automotive-related value added is obtained by scaling each firm's total value added by its Ghosh inverse coefficient and then aggregating across 2-digit sectors; in panel b, the coefficient is computed considering only exposure to non-domestic EU car producers. In both panels, manufacturing industries are shown in shades of blue and services in shades of green. Percentages in parentheses denote each sector's share of total automotive-related value added in the corresponding panel.

Figure B.2. Regional composition of the automotive supply chain (2022)

(a) Including final car producers **(b)** Suppliers non-domestic EU producers

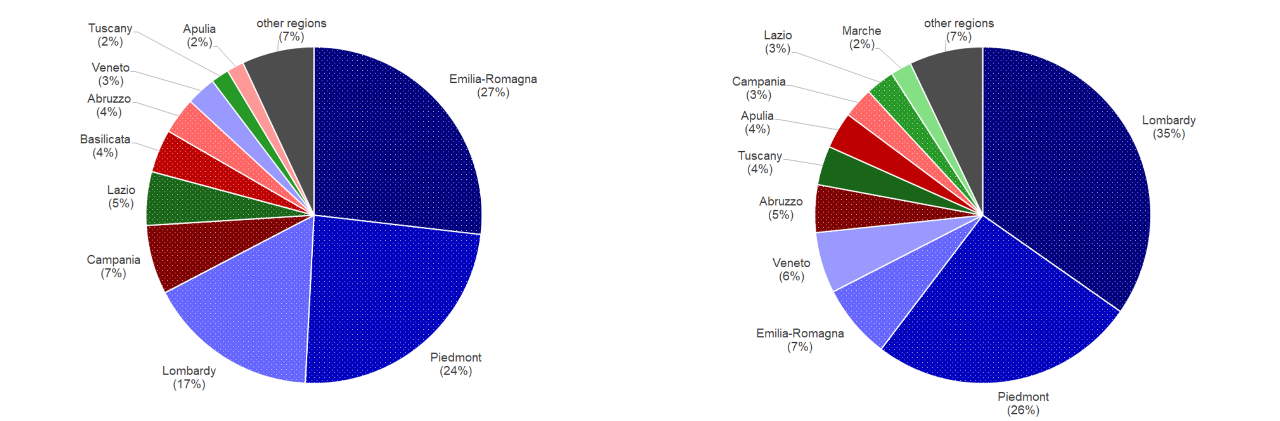

Notes: Panel a reports the top 10 NUTS-2 regions by automotive-related value added, including final car producers located in Italy. Panel b reports the top 10 regions among Italian suppliers linked to final carmakers located elsewhere in the EU. Automotive-related value added is obtained by scaling each firm's total value added by its Ghosh inverse coefficient and then aggregating across NUTS-2 regions; in panel b, the coefficient is computed considering only exposure to non-domestic EU car producers. In both panels, Northern regions are shown in shades of blue, Central regions in green and Southern regions in red; regions where final car production plants are located are dotted. Percentages in parentheses denote each region's share of total automotive-related value added in the corresponding panel.

Appendix C Alternative regression specification

Following comments received during the seminar and the refereeing process, this appendix reports a robustness exercise based on a more restrictive and time-consistent definition of participation in the automotive supply chain. In the baseline analysis, firms are classified as automotive if their Ghosh coefficient exceeds 10% in a given year. Here, instead, we define a firm as belonging to the automotive supply chain if and only if its Ghosh coefficient is above 10% in every year in which the firm is observed in the FATTEL data.

This alternative definition aims to capture firms with a stable and persistent exposure to the automotive industry, excluding firms with only sporadic or transitory links to the supply chain. We then re-estimate the same regression specifications presented in Section 5, using this more stringent indicator of automotive participation.

Overall, the results are very similar to those obtained under the baseline definition. Estimated automotive premia remain comparable in magnitude and statistical significance across all outcome dimensions—firm size, sales structure, productivity and innovation, and labor outcomes—confirming that our main findings are not driven by short-term fluctuations in firms' involvement in the automotive supply chain.

Table C.1. Characteristics of automotive firms: firm size (robustness)

Dependent variable:	Sales (log)	Revenues (log)	Value added (log)	Tangible assets (log)	Intangible assets (log)	Age (log)
	(1)	(2)	(3)	(4)	(5)	(6)
$\mathbf{1}_i$(Automotive)	0.150***	0.075***	0.115***	0.024	0.045	-0.031***
	(0.013)	(0.012)	(0.009)	(0.017)	(0.030)	(0.010)
$\mathbf{1}_i$(Direct supplier)	0.392***	0.345***	0.261***	0.037	0.185***	0.032*
	(0.022)	(0.020)	(0.014)	(0.027)	(0.049)	(0.017)
Employment$_{t-1}$ (log)	X	X	X	X	X	X
Capital intensity$_{t-1}$ (log)	X	X	X	X	X	X
Year effects	X	X	X	X	X	X
Industry effects	X	X	X	X	X	X
NUTS-2 effects	X	X	X	X	X	X
Industry × year effects	X	X	X	X	X	X
Adjusted R^2	0.60	0.66	0.73	0.65	0.34	0.31
Observations	1,919,545	1,915,473	1,820,172	1,828,418	1,210,715	1,919,166

Notes: Standard errors in parentheses are clustered at the firm level. The upper panel reports regressions where firms are classified as belonging to the automotive supply chain if their Ghosh inverse coefficient exceeds 10% for every year in which the firm is present in the Fattel archives. The lower panel focuses on the subset of automotive firms that have direct sales linkages to final carmaker. Each panel refers to a separate set of regressions. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.2. Characteristics of automotive firms: structure of sales (robustness)

Dependent variable:	$\mathbb{1}$(Export >0)	Export (log)	N. of clients (log)	HHI Clients (log)	N. of suppliers (log)	HHI Suppliers (log)
	(1)	(2)	(3)	(4)	(5)	(6)
$\mathbf{1}_i$(Automotive)	0.015***	-0.164***	-0.571***	0.345***	-0.033***	0.033***
	(0.006)	(0.052)	(0.020)	(0.015)	(0.009)	(0.010)
$\mathbf{1}_i$(Direct supplier)	0.146***	0.174**	-0.488***	0.322***	0.120***	-0.068***
	(0.009)	(0.068)	(0.032)	(0.024)	(0.013)	(0.018)
Employment$_{t-1}$ (log)	X	X	X	X	X	X
Capital intensity$_{t-1}$ (log)	X	X	X	X	X	X
Year effects	X	X	X	X	X	X
Industry effects	X	X	X	X	X	X
NUTS-2 effects	X	X	X	X	X	X
Industry $\times$ year effects	X	X	X	X	X	X
Adjusted R^2	0.33	0.36	0.37	0.22	0.53	0.17
Observations	1,919,686	309,880	1,908,285	1,908,054	1,919,352	1,919,347

Notes: Standard errors in parentheses are clustered at the firm level. The upper panel reports regressions where firms are classified as belonging to the automotive supply chain if their Ghosh inverse coefficient exceeds 10% for every year in which the firm is present in the Fattel archives. The lower panel focuses on the subset of automotive firms that have direct sales linkages to final carmaker. Each panel refers to a separate set of regressions. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.3. Characteristics of automotive firms: productivity, investments, and innovation (robustness)

Dependent variable:	VA per worker (log)	TFP (log)	Intangible invest. (log)	% intangible invest. (log)	$\mathbb{1}$(N. patents>0)	N. patents (log)
	(1)	(2)	(3)	(4)	(5)	(6)
$\mathbf{1}_i$(Automotive)	0.079***	0.027***	0.121***	0.082***	0.005***	0.097
	(0.008)	(0.004)	(0.030)	(0.025)	(0.002)	(0.080)
$\mathbf{1}_i$(Direct supplier)	0.174***	0.023***	0.241***	0.067*	0.021***	0.156*
	(0.013)	(0.005)	(0.048)	(0.038)	(0.004)	(0.093)
Employment$_{t-1}$ (log)	X	X	X	X	X	X
Capital intensity$_{t-1}$ (log)	X	X	X	X	X	X
Year effects	X	X	X	X	X	X
Industry effects	X	X	X	X	X	X
NUTS-2 effects	X	X	X	X	X	X
Industry $\times$ year effects	X	X	X	X	X	X
Adjusted R^2	0.29	0.24	0.24	0.15	0.03	0.20
Observations	1,788,031	1,869,742	590,716	590,716	1,919,686	7,209

Notes: Standard errors in parentheses are clustered at the firm level. The upper panel reports regressions where firms are classified as belonging to the automotive supply chain if their Ghosh inverse coefficient exceeds 10% for every year in which the firm is present in the Fattel archives. The lower panel focuses on the subset of automotive firms that have direct sales linkages to final carmaker. Each panel refers to a separate set of regressions. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table C.4. Characteristics of automotive firms: labor outcomes (robustness)

Dependent variable:	Employment (log)	Under-35 empl. (log)	% under-35 (log)	Average wage (log)	Hours of CIG (log)
	(1)	(2)	(3)	(4)	(5)
$\mathbf{1}_{it}$(Automotive)	0.679***	0.065***	0.028***	0.066***	0.333***
	(0.022)	(0.013)	(0.012)	(0.005)	(0.021)
$\mathbf{1}_{i}$(Direct supplier)	1.410***	0.147***	0.055***	0.103***	0.412***
	(0.037)	(0.021)	(0.020)	(0.008)	(0.033)
Employment$_{t-1}$ (log)		X	X	X	X
Capital intensity$_{t-1}$ (log)	X	X	X	X	X
Year effects	X	X	X	X	X
Industry effects	X	X	X	X	X
NUTS-2 effects	X	X	X	X	X
Industry × year effects	X	X	X	X	X
Adjusted R^2	0.13	0.55	0.20	0.40	0.59
Observations	1,873,739	1,223,231	1,223,225	1,639,219	514,637

Notes: Standard errors in parentheses are clustered at the firm level. The upper panel reports regressions where firms are classified as belonging to the automotive supply chain if their Ghosh inverse coefficient exceeds 10% for every year in which the firm is present in the Fattel archives. The lower panel focuses on the subset of automotive firms that have direct sales linkages to final carmaker. Each panel refers to a separate set of regressions. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.