Shock or Opportunity? Structural Change and Employment Dynamics Around Amazon Facilities

Giulia Antonacci

December 17, 2025

Abstract

This thesis investigates how the opening of Amazon Fulfillment Center reshapes local economic structures and labor market dynamics within Italian provinces. While the existing literature is mainly based on United States analysis and is divided between positive employment effects and potential crowding out pressures, far less is known about the mechanisms through which large logistics investments operate in European economies, which are more rigid and heterogeneous. To adress this gap, this thesis applies Augmented Synthetic Control Method to a high-frequency quarterly panel combining Movimprese and ISTAT data for the period 2013-2022. The analysis focuses on five provinces in which an FC opened between 2017 and 2020 (Rome, Turin, Vercelli, Rieti and Rovigo) and evaluates both macroeconomic and microeconomic adjustments following the opening of an FC. Four outcomes emerge as robust: the relative weight of digital intensive and real estate activities and wage levels for apprentices and blu collar workers. The results reveal a consistent pattern: Amazon generates meaningful and persistent effects only where the local economy is capable of functioning as a logistics and technological hub.

Overall, the finding suggest that Amazon's impact is neither uniform nor automatic across host territories. Instead, it critically depends on the sectoral composition and structural elasticity of each host province. Based on this, Amazon's FCs become transformative forces or remain a isolated presences within the local economy.

0.1 Introduction

Amazon is nowadays one of the most recognizable companies in the world, employing 1.5 million people in more than fifty countries and serving an estimated 300 million active clients, with around 12 million daily orders. In Europe alone, there are more than 250 logistics infrastructures employing 150 000 permanent workers in 21 States. Today, it is engaged in e-commerce, cloud computing, online advertising, digital streaming and artificial intelligence and it is one of the most complex and vertically integrated global companies in terms of logistics and digital operations. While its diversification is often considered as the key driver of Amazon's success, its true competitive core lies in its exceptional ability to design, operate and optimize its logistics network.[1]

To deliver a fast and reliable service to its customer, which is the key promise of Amazon, geographical diffusion is crucial. Through an extensive network of hubs located in strategic areas, Amazon supports and sustains the whole value chain, from the handling of warehousing to the final delivery to the client, significantly more efficient than the traditional retailers. Fulfillment centers (FCs) represent the backbone of this system: they are large-scale, multi-process facilities designed to handle inbound inventory, storage, picking, packing and outbound shipping. There are many different types of FC that may have different building size, automation level and product type. Typically, it covers a surface area between 40 000 and 100 000 square meters, has a storing capacity up to 20 million individual items and employs between 500 and 3000 workers, depending on automation levels and regional logistics demand. FCs the biggest storing facilities, and are supported by sortation centers and last-mile delivery stations for the customer distribution. These support facilities were first introduced with the launch of 'Prime' services promising 1-day-delivery and are now complementary in the logistics pipeline.[2]

[1] Houde, Newberry, and Seim, *Nexus Tax Laws and Economies of Density in E-Commerce*.
[2] Henaway, "Amazon's Distribution Space."

However, while it is clearly an advantage for Amazon to expand further and faster across the territory, many researchers have questioned the local impact that it has and the implications for the host communities. A central question in the literature is whether the opening of an FC generates net local job creation, by increasing labor demand, raising wages, and possibly attracting new workers into the area, or whether it mainly produces crowding-out effects, displacing employment in traditional retail, warehousing and small logistics firms that can't compete with Amazon's scale and efficiency.[3] In other words, does the FC add new economic activity to the local economy or does it mostly reorganize and centralize activity that was already present but distributed across many small operators?

From this perspective, the arrival of an FC can be seen as a trade-off. On the one hand, is is expected to increase employment opportunities for low and middle skilled workers in the area, stimulate the creation of complementary jobs in transport, business services and facility management and potentially stimulate a migration of workers attracted by new job positions and more stable contracts.[4] On the other hand, the same process may put pressure on the cost of living, particularly on housing and local services, so that nominal wage gains are partly or fully offset in real terms. At the same time, intensified competitive pressure may force some small local firms to contract or exit the market and workers and households who are unable to afford higher rents or who do not find a place in the new employment structure may be pushed to leave.

The research question of this thesis is therefore not only whether employment increases after an FC opening, but how this adjustment takes place, which sectors expand or shrink, which types of firms survive or exit, which groups of workers gain or lose and whether the local economy becomes more dynamic and diversified or mode dependent on a single dominant employer.

[3] Pathania and Netessine, *The Impact of Amazon Facilities on Local Economies*.
[4] Cunningham, *Local Labor Market Effects of Amazon* (2024).

0.2 Existing literature

One of the most influential empirical contributions is 'The impact of Amazon Facilities on Local Economies' by Pathania and Netessine,[5] which employs a difference-in-difference design combined with nearest-neighbour matching to address selection concerns. The paper is mainly focused on mid-sized US counties in a time range of 3 years (2014-2017) and its results show that an opening of an Amazon FC leads, on average, to a 1.46% increase in employment to population ratio, a 2.7% reduction in poverty and a 2.3% rise in median household income in the first three years after entry. These results suggest a positive labor market spillover and improvements in socioeconomic conditions. Other studies supporting this theory are Cunningham's, published in 2024 [6] and 2025.[7] He studied Amazon's effects over employment using the staggered roll out of Amazon's FCs across US metropolitan areas and counties. Its results show that in the area where the FC opened there is a raise in employment rate by about 1%, which interests particularly non-college and younger workers. Cunningham also found evidence of a wage increase of 0.7% that mainly affects warehousing and transportation fields, with an earning distribution shifting modestly upward for non-college workers.

However, another interesting finding of Cunnigham's papers is that the rise in employment is not the whole picture. Beyond aggregate labor market effects, it is also important to understand how Amazon reshapes the internal composition of local economies, reallocating employment across sectors connected, directly or indirectly, to the logistics pipeline. Cunningham showed how Amazon's entry systematically shifted the local industrial structure: the share of employment in transportation and warehousing increases across all demographic groups, while retail and wholesale trade contract, especially among younger and non-college workers.[8]

Economic Policy Institute (EPI) also supported this results, arguing that Amazon's FCs employment successes are actually offset by the reduction in their local industries, particularly brick-and-mortar retail and traditional storage/transport firms.[9] This line of research supports

[5] Pathania and Netessine, *The Impact of Amazon Facilities on Local Economies*.
[6] Cunningham, *Local Labor Market Effects of Amazon* (2024).
[7] Cunningham, *The Local Economic Impacts of Amazon* (2025).
[8] Cunningham, *The Local Economic Impacts of Amazon* (2025).
[9] Chava et al., *Labor Market Effects*.

a crowding-out hypothesis, with Amazon's superior productivity and scale overcoming local competitors and destroying previous traditional companies that used to operate in logistics and retail. They analyzed county-level data on employment between 2000 and 2015 using synthetic controls and they found no statistically significant increase in overall county employment, but an important decline in small-scale logistics operators, which indicates that Amazon did not create new jobs, but employed the exact same number of people that lost their jobs because of its entrance in the market.

These findings are consistent with Chava et al. (2023),[10] who document a 2.1% decline in employment and a 2.4% decline in total earnings in proximate brick-and-mortar retail establishments following the opening of FCs, indicating that part of Amazon's labour demand is met by reallocating workers out of traditional retail. The easier explanation is that Amazon's conditions can be better than local smaller companies who cannot offer the same wages and are forced to exit the market.

At the same time, Amazon's presence tends to stimulate growth in sectors complementary to logistics. Henaway (2019)[11] and Hassel Sieker (2022)[12] show that FC opening increases local demand for subcontracted transport companies, temporary work agencies, security, maintenance, ICT services and technologically specialized suppliers. Taken together, these studies suggest that the opening of an FC is not only an employment shock, but can also stimulate a profound reshaping of the local labor market.

Moreover, an important aspect to cover to really understand the changes of the territory is the significant increase in the housing market and the rising cost of living: as measured by Cunningham (2025),[13] Amazon's entry brings in average an increase in rents of about 1.1%, home values of 5.6% and utility costs 6%. This means that it is true that blue-collar workers have a significant increase in their wages, but it is not clear whether it is off-set by the increase in these living costs.

[10] Chava et al., *Labor Market Effects*.
[11] Henaway, "Amazon's Distribution Space."
[12] Hassel and Sieker, "The Platform Effect."
[13] Cunningham, *The Local Economic Impacts of Amazon* (2025).

All these existing studies focus almost exclusively on US, where the labor market structures, firm demographics and regulatory frameworks differ substantially from EU countries and for this reason evidence from US cannot be directly extrapolated to the European context and cannot predict the heterogeneous territorial adjustments that may occur in other countries. Italy, in particular, represents a compelling case to study for several reasons.

First, the Italian productive system is characterized by a dense network of small and medium-sized enterprises, often family-owned and, therefore, structurally more vulnerable to competitive shocks from large multinational firms. Second, Italy displays strong territorial heterogeneity, with northern regions more industrialized, export-oriented and integrated into EU value chain, while southern regions are characterized by a weaker productive base, higher unemployment and more fragile entrepreneurial ecosystem, which may lead to different impacts based on the region.[14] Lastly, Italy is characterized by strong labor market rigidities, lower geographical mobility and historically slower wage dynamics, which may amplify or dampen the channels identified in previous literature, generating adjustment patterns that diverge from US data. For all these reasons, studying Italy can be interesting to assess the external validity of the US evidence and therefore it is essential to understand the scale and nature of Amazon's presence in Italy.

According to Amazon's official assessments,[15] the company has invested over €25 billion in Italy since 2010, including both capital expenditure for logistic centers, offices and cloud infrastructure and operating costs, as wages and benefits for workers. Moreover, Amazon reports the creation of over 19 000 permanent jobs in Italy by 2025 across more than sixty facilities between operations, corporate offices, data centers and customer service centers. Independent analysis confirm this picture: a study by The European House Ambrosetti TEHA ranks Amazon as the first among all companies operating in Italy for net job creation between 2014 and 2023, the 20th for total number of employees and among foreign firms it is the third for total investments made in Italy during 2021-2023.

[14] Li, *E-Commerce and Regional Inequality*.
[15] Amazon EU, *Amazon EU Economic Impact Report 2023*.

Despite these information, there is still no documentation related to the actual local effects of the FC opening within the province. This study aims to fill that void and dive deeper in the main changes that are observed in the provinces that hosted a new FC. The next chapter presents the dataset used and outlines the variables that allow to track both the aggregate employment effects with other macro characteristics and the potential sectoral reshaping that may affect the local labor market.

0.3 Data

The main source is Movimprese, an official report that provides detailed and harmonized information on the entrepreneurial structure of Italian companies for each province quarterly.[16] Movimprese is produced by InfoCamere, on behalf of the public body UnionCamere, the coordinating institution of all Italian Chambers of Commerce. InfoCamere compiles, integrates and maintains the official business registry databases, ensuring temporal consistency and territorial comparability. Through these data, it is possible to analyze local market dynamics and structural changes in firm composition over time.

The original structure of the dataset has three core information blocks: field of expertise, juridical structure, activity indicators.

The field of expertise is defined through the ATECO classification, which is the Italian implementation of the European NACE framework.[17] Each company is assigned to a specific sector and division according to its main business activity. Studying the number of firms that enter and remain active in each field for each province over time is highly informative to see how labor market is changing and to assess which impact Amazon FC may have and if there are common patterns after its opening. This can give a clear answer to the question: is the FC stimulating the local economy through supply-chain complementarities and creating logistics hub or is it crowding out local firms that were previously working in the field and become unable to compete with Amazon's scale and efficiency? While warehousing and logistics activities are the

[16] InfoCamere and UnionCamere, Movimprese database.
[17] ATECO classification based on: ISTAT, "Classificazione delle attività economiche ATECO".

most directly affected, the impact may also extend to more knowledge-intensive sectors such as ICT, R&D, or professional services.

0.3.1 ATECO Clusters

To evaluate these mechanisms, ATECO codes have been aggregated into clusters designed to capture distinct transmission channels.

- **Logistics_core**: H52–H53 (warehousing, storage, courier and postal activities). This is the most directly affected sector and FCs operate exactly in this market as competitor or sometimes investors for local companies already working in the field.[18]

- **Logistics_broad**: H49–H50–H51 (road, rail, water and air transport). These activities supply upstream and downstream transport services, potentially positively stimulated by FC operations. They represent the broader transport infrastructure that supports the movement of goods across territories. Even if they are less tightly linked to warehouse inner operations, they are essential for enabling inbound and outbound freights flows and may adjust to accomodate changes in logistics demand.

- **Digital_extended**: K64–K65–K66–M71–M73–M74–N77–S95 (finance, insurance, scientific/technical services, advertising, leasing, IT repair). These sectors provide the complementary business services that typically expand alongside the digitalization of supply chain, contributing to coordination, risk management consulting, software maintenance and other high-value tasks that sustain a modern logistics ecosystem.

- **Information**: J58–J59–J60–J61–J62–J63 (publishing, audiovisual, telecommunications, ICT). These activities are part of the digital infrastructure that sustains large logistics platforms that typically rely intensively on data-flows, real-time communication and IT-intensive coordination mechanisms. Growth or restructuring in this cluster may reflect spillovers and adjustments in the technological fields in support to logistics operations [19]

[18] Cunningham, *The Local Economic Impacts of Amazon* (2025).
[19] Henaway, "Amazon's Distribution Space" (2023).

- **Research**: M72 (scientific R&D). Although not directly tied to Amazon's operational processes, it serves as a proxy for the knowledge-intensive component of the local economy. Changes here could signal broader technological spillovers in local innovation capacity.

- **Professional_support**: M69–M70–M82 (legal, accounting, management consulting, office support). These firms provide the administrative, regulatory and managerial support necessary for big comapnies, especially the ones entering in a new territory, and may respond to changes in the corporate organization or compliance needs generated by a major investments, in Amazon's specific case related to logistics.

- **Commerce_proxy**: G45–G46–G47 (vehicle trade, wholesale, retail). These sectors represent the traditional fabric of the territory and may be exposed to competitive adjustments triggered by Amazon's presence, particularly wholesale and retail operators that interact more closely with consumer markets and has proven to have suffered the opening of an FC in other studies, as Cunninghan's.

- **Realestate_construction**: L68–F41–F42–F43 (real estate, construction of buildings, infrastructure works). These sectors typically react during the investment phase, reflecting construction activities, but also may vary in the long-term period because of territorial changes and adaptations, such as the building of new accomodations to answer to the entrance of new workers.

Overall, the opening of an Amazon FC is expected to generate heterogeneous sectoral adjustments. Logistics_core is expected to be the most directly positively affected field, with warehousing, storage and courier services boosted by the new FC's demand and followed by a moderate expansion of broader transportation services (logistics_broad). Complementary service sectors (professional_support, digital_extended, information) are also likely to grow as firms reorganize around the operational and administrative needs induced by the new facility. Realestate_construction typically experience a sharp increase during the pre-opening period, corresponding to the investment phase of the FC, and then stabilize at elevated levels in the following quarters. Lastly, retail activities are expected to face competitive pressure from Amazon's online marketplace, potentially reducing firm entry and the share of traditional brick and

mortar shops.

These patterns are consistent with findings from previous research, mainly located in U.S. (e.g., Cunningham 2025;[20] Henaway 2023;[21] Hassel & Sieker 2022[22]). The key point of this study is to understand if in Italy the pattern will be the same or if there are any changes in this structure.

0.3.2 Juridical Forms

Firms are categorized into four macro-classes: sole proprietorships, partnerships, corporations, other forms. The distribution across juridical forms provides insights into the complexity and capital intensity of the local business environment. For example, increases in corporations (società di capitali) may suggest new suppliers or more structured logistics actors entering the market, attracted by the demand created by the FC. Small partnerships and sole proprietorships, instead, may be more vulnerable to displacement due to competitive pressure.

0.3.3 Activity Indicators

The dataset additionally reports: number of firms registered, number of active firms, number of firm closures. These indicators are the key values for the research and tracking their variations in ateco clusters and juridical forms allows to construct measures of business dynamism and potential crowding-out effects around the FC opening.

0.3.4 Complementary Data from ISTAT

Movimprese covers firm dynamics, but to have a complete picture it is necessary to study local socio-economic conditions and changes. The key data source is ISTAT (Istituto Nazionale di Statistica, the official national statistical institution).[23] These variables are crucial for understanding the social and aggregate labor market effects of the FC and are all calculated yearly. To have the quarterly data, the variables were divided into two categories: stock variables, which are considered slow-moving variables constant within a year, and flow variables, which are

[20] Cunningham, *The Local Economic Impacts of Amazon* (2025).
[21] Henaway, "Amazon's Distribution Space" (2023).
[22] Hassel and Sieker, "The Platform Effect" (2022).
[23] ISTAT, "Bilanci Demografici" and "Conti Economici Territoriali".

converted to quarterly frequency. All stock variables were duplicated over the quarters of the corresponding year, while flow variables were divided by four. Stock variables:

- **Resident population**: total number of residents, captures the demographic scale of each province and makes it possible to interpret firm-level indicators relative to the size of local community. It is mainly used for normalization and in synthetic control procedures.
- **Median wages of workers**: this measure is reported separately for white-collar, blue-collar and apprentices and represents a key labor market outcome. Since FCs primarily employ blue-collar workers, wage trends are informative on potential monopsony effects or competitive adjustments. At the same time, median wages reflect living standars and reveal how different occupational groups may experience the impact of new labor demand.
- **Labour input intensity**: expressed in full-time equivalent units (ULA) per resident, it summarizes the quantity of labor used in the provincial economy and corresponds to the work performed by one person employed full-time for the entire year. It incorporates both employment levels and typical working hours, offering an interesting indicator of economic activity and job offer of a territory.

Flow variables:

- **Migration flows**: this measure is reported separately for inward and outward flows and can be useful for detecting attraction effects of the FC. Increases in in-migration may indicate that new job opportunities are attracting workers, while higher out-migration may reflect the arise of new difficulties for local workers.
- **GDP**: this measure offers the broadest measure of economic performance and summarizes the value of goods and services produced locally.

All stock variables are duplicated across the four quarters of the corresponding year.

Combining Movimprese with ISTAT data provides a comprehensive picture of the potential impact of Amazon FC openings on: firm dynamics across sectors and juridical forms, labor market conditions and wages, demographic adjustments such as migration, local economic performance. This integrated dataset allows us to test whether Amazon FCs in Italy act as catalysts

for local economic development or crowd out existing operators, mirroring or diverging from patterns observed internationally.

As previously mentioned, Amazon has opened many FCs in Italy over the last 15 years. I selected only the ones opened between 2017 and 2020, which are respectively collocated in the following provinces: Rieti (2017), Vercelli (2017), Turin (2019), Rome (2020) and Rovigo (2020). It is important to signal that these provinces are not equally spread in the Italian territory and differ in demographic size and labor market structure, which is essential to interpret heterogeneous treatment effects.

Rome and Rieti are located in the south-central part of Italy. The province of Rome is a large metropolitan area that hosts an highly diversified labor market and represents one of the country's major economic hubs for public administration, advances services, ICT activities and a dense logistics network, connecting the north to the south. The FC opened in 2020 was not the first one: today there are 2 FCs, 1 delivery station and other 4 small Amazon facilities operating in the province. By contrast, the province of Rieti is characterized by a more traditional and diversified economy, with a labor market centered on local services, small manufacturing firms and agriculture. Rieti is much smaller and the FC opened in 2017 is the only Amazon's facility in the territory.

The other three provinces are located in the north of the country, specifically Turin and Vercelli in Piemonte and Rovigo in Veneto. Turin is a large metropolitan province, committed in a new industrial transition and technological innovation and traditionally highly developed in automotive and manufacturing. It hosts two Amazon's FC and the one opened in 2019 is now employing around 1700 workers.

As for Vercelli, it is a small province with a labor market specialized in agriculture (particularly rice production), textile manufacturing and a limited digital services segment. Lastly, Rovigo is a mid-sized province with a labor market that is recently growing in logistic sector. It hosts small manufacturing firms and agro-industrial activities. Amazon has 2 facilities in the province today.

All these information will be necessary for understanding the effect of the opening of the new FC.

0.4 Data Merging and Harmonization

To guarantee econometric reliability all datasets were merged into one unique panel.
First, because the core analysis is made over quarters, all the non-quarter datasets needed to be transformed. As mentioned before, this transformation was made based on the basic distinction between stock and flow variables.
Second, geographical harmonization is necessary to have a complete dataset. The merging process was conducted using standardized provincial identifiers to ensure a one-to-one correspondence across sources. All provinces were first standardized and then matched through all datasets. Ten provinces were dropped because they were not consistently available across all datasets. In particular, the entire region of Sardinia is excluded from the final sample due to incomplete coverage of GDP and migration variables in several sub-periods.

For each province and quarter, the panel presents measures of sectoral composition, wages, population, migration flows, GDP, labor input and juridical forms. All variables are expressed in comparable formats, such as levels, per-capita values or share, depending on their nature. This standardization makes it possible to jointly analyze sectoral, legal form and macroeconomic dynamics and to use all series as predictors or outcomes within the ASCM procedure. The final panel has 101 provinces and 36 quarters from 2013 to 2022.

0.5 Methodolody

To assess whether the opening of FCs has a real causal impact on the local economic structure and labor market dynamics across Italian provinces, I adopt the Synthetic Control Method (SCM), originally developed by Abadie, Diamond, and Hainmueller (2010)[24], and its more modern and robust extension, the Augmented Synthetic Control Method (ASCM), introduced by Ben-Michael, Feller, and Rothstein (2021).[25]

[24] Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. "Synthetic Control Methods for Comparative Case Studies." *Journal of the American Statistical Association*, 2010.

[25] Ben-Michael, Eli, Avi Feller, and Jesse Rothstein. "The Augmented Synthetic Control Method." *Journal of the American Statistical Association*, 2021.

SCM is designed for settings with one treated unit and many potential control units observed over time. The method constructs a synthetic control as a weighted average of non-treated provinces using only pre-treatment observations (which constitute the donor pool), selecting the combination that best reproduces the trajectory of the treated province. SCM relies on the assumption that units exposed to similar observed and unobserved components will have similar trajectories in the pre-treatment period (Linear Factor Model)[26]. Therefore, the synthetic control is built to closely reproduce the pre-treatment path of the treated unit, implicitly matching both observed and unobserved determinants of the outcome. The closest the synthetic is to the actual treated path, the more reliable their values will be in the post-treatment period. The causal effects of the treatment will be measured as the difference between the trajectory of the treated unit's outcome and the counterfactual one.

However, to reproduce the treated path SCM looks for linear combinations of the donor pool's observations, and therefore the main constraints are that weights must be non-negative and sum to one, so that treated units lie within the convex hull of the donors. When this condition fails, pre-treatment fit deteriorates and causal inference becomes unreliable.

This limitation motivates the transition to ASCM. ASCM relaxes the non-negativity constraint and augments the synthetic control with a penalized regression adjustment—typically ridge regression—estimated on the donor pool.[27] If SCM already achieves good pre-treatment fit, ASCM will have similar results; otherwise, ASCM improves balance and reduces bias.[28]. ASCM can be written as:

$$\widehat{Y}_{1T}^{ASCM}(0) = \underbrace{\sum_j w_j^{SC} Y_{jT}}_{\text{Synthetic Control prediction}} + \underbrace{\left[\widehat{m}(X_1) - \sum_j w_j^{SC}\widehat{m}(X_j)\right]}_{\text{Bias correction term}}. \tag{1}$$

$$\hat{Y}_{1t}^{\text{ASCM}}(0) = \hat{Y}_{1t}^{\text{SC}}(0) + \widehat{\text{bias}}_t.$$

[26] Abadie, Diamond, and Hainmueller (2010), see also Abadie, Diamond, and Hainmueller (2015), "Comparative Politics and the Synthetic Control Method."(12)

[27] Ben-Michael, Feller, and Rothstein (2021).

[28] See also Doudchenko and Imbens (2017), "Balancing, Regression, Difference-in-Differences and Synthetic Control Methods."(14)

In the Amazon application, several adaptations are implemented. First, the donor pool is restricted to provinces that never receive treatment and are located at least 100 km from the treated unit, in order to minimize spillover effects.[29] Second, outcomes and predictors are standardized before estimating the penalized regression, ensuring comparability across variables with different scales. Third, Principal Component Analysis (PCA) is applied to the donor matrix to reduce dimensionality and increase model stability, in line with recommendations in the recent SCM literature.[30] Fourth, the ridge penalty is selected through cross-validation in the pre-treatment period, following the placebo-split approach described in Ben-Michael et al. (2021). Finally, macroeconomic indicators exhibit smoother dynamics and stronger common trends than ATECO or juridical-form outcomes; therefore, a more restrictive PCA and stronger penalization are adopted.

To evaluate credibility, an in-time placebo test is conducted.[31] By simulating alternative treatment dates in the pre-treatment period and re-estimating the model, it becomes possible to assess whether the estimated treatment effect is unusually large relative to the distribution of placebo effects.

To identify the most reliable results, all outcomes were processed through a structured selection procedure. The only outcomes retained were the ones showing for at least four out of the five treated provinces coherent and consistent results, displaying statistical significance in the in-time placebo test and with no sign of overfitting or underfitting. In this way, the post-treatment divergence could not be attributed to noise and the results can be considered reliable.

A more detailed and structured description of the methodology is in Appendix 1, where all the models and econometrics' topics are properly explained.

[29] For the importance of donor pool restrictions see Abadie, Diamond, and Hainmueller (2015).

[30] See Xu (2017), "Generalized Synthetic Control Method", for discussion of dimensionality reduction in panel causal inference.(15)

[31] Placebo strategies are discussed extensively in Abadie, Diamond, and Hainmueller (2010; 2015).

0.6 Selected Outcomes and Results

This section dives deeper into the selected outcomes and the results they show. The metrics used to evaluate the causal effects estimated through the ASCM are mainly the following.

First, the average treatment effect on the treated (ATT) measures the post-treatment difference between the treated province and its synthetic counterfactual. A positive ATT indicates that the outcome increased more than expected and shows an upward divergence with the counterfactual. The SMD is calculated rescailing ATT by pre-treatment variability for the outcome to garantee comparison across different outcomes: values around 0.5 denote medium effects, values above 1 suggest large, economically meaningful changes that can't be attributed to the noise.

The other essential metric, also used by Abadie, is the mean squared prediction error, which reflects how the synthetic control managed to reproduce the treated unit's trajectory in the pre-treatment period. A low MSPE_pre signals strong credibility of the counterfactual. The ratio between post and pre-treatment MSPE is indicated with ratio_safe and is analyzed with placebo test value to address for the overall significance of the results.

0.6.1 Digital sector: `digital_extended active/share`

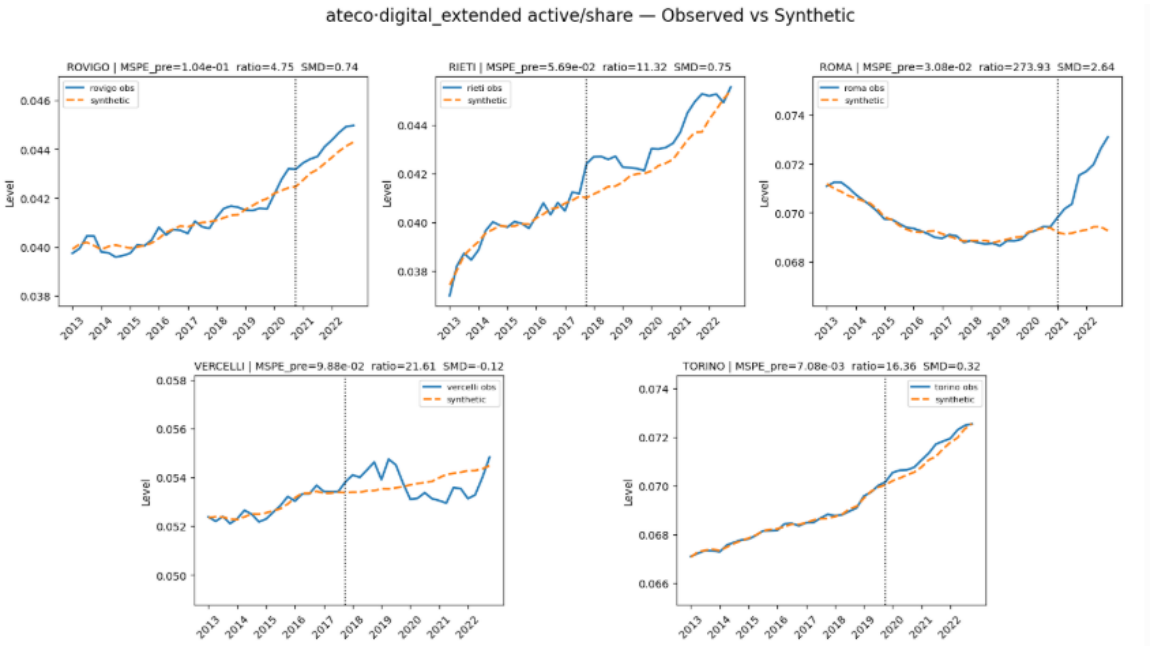

Figure 1: *

(a) Treated vs. synthetic and ATT.

--- Outcome: ateco·digital_extended active/share ---

	Province	ATT_mean_post	SMD	MSPE_pre	ratio_safe	p_cond_prefit
0	rovigo	0.000687	0.743189	0.104304	4.747505	0.206897
1	rieti	0.000817	0.751797	0.056895	11.323025	0.062500
2	roma	0.002146	2.638430	0.030763	273.926641	0.033333
3	vercelli	-0.000065	-0.124625	0.098789	21.611901	0.187500
4	torino	0.000245	0.315002	0.007083	16.359684	0.000000

Figure 2: *

(b) Summary statistics by province.

In the pre-treatment period, the share of digital-extended firms, as finance, insurance, scientific/technical services, advertising, leasing and IT repair, displays a clear upward trend in all provinces. With the entrance of the new FC, there are major changes in the environment. In fact, large logistic platforms like Amazon tend to rely on a constellation of digital services, as automation software, IT support, data processing, tracking technologies, and the increase of demand of such services reinforces the relative importance of digital activities in areas already

endowed with tech-intensive capabilities. Therefore, the opening of the FC should boost the digital sector, with the clearest pattern emerging in Rome, where the treated series diverges sharply upward from its synthetic counterpart. This positive result is significant for Rieti, Rome and Turin, indicating that such a divergence is uncommon among their respective donor pools and that digital firms grew at a faster pace than other sectors.

Focusing on Rome and Turin, these results are coherent with the inner structure of the two provinces, being not only the largest between the five treated provinces considered, but were also the most innovative and digitalized (y axes ranging between 0.066 and 0.074 for both of them). Rome is one of the most digitally advances territories in Italy according to Unioncamere rankings and hosts multiple logistics facilities, such as centers in Colleferro, Ardea, Pomezia and Rome Settecamini, embedded in a broader program for a regional support in logistics growth (Zona Logistica Semplificata).

Also Turin is currently investing in programs that support automation and mobility technologies, with innovation programs as S3 Piemonte.

These conditions provided for both provinces a fertile ground for complementarities: the arrival of the FC made the digital hub stronger and digital suppliers, software providers, data-processing firms and ICT logistics services were boosted and incentivated to grow and kept intensify their presence supporting the increase in demand of such services.

A plausible interpretation is that Amazon's presence alone is insufficient to trigger digital upgrading in territories lacking pre-existing innovation ecosystem or a critical mass of tech-intensive firms.

Also Rovigo and Rieti show a smaller but still positive shift. As for Rovigo, the increase suggests a mild complementary activity growth, possibly supporting logistics operations, but not a large scale realignment of the local economy. This is true also for Rieti and can be interpreted both as digital firms not expanding faster than the rest of the economy or that these provinces lack the ecosystem necessary to translate Amazon's presence into a meaningful repositioning toward digital sectors.

Vercelli's digital share is the only negative result between the treated provinces. While not strongly significant, this result can be analyzed as the insufficient strenght of the FC to boost

the growth of the digital sector in a mainly manufacturing environment as Vercelli. However, as shown in Figure 1, the two paths keep crossing in the years after the opening of the FC and these results are considered unreliable and not significant. This is consistent with the idea that less diversified local economies have more stable sectoral shares and are less reactive to large individual entrants, especially in sectors not particularly developed.

0.6.2 Real Estate: `realestate_construction active/share`

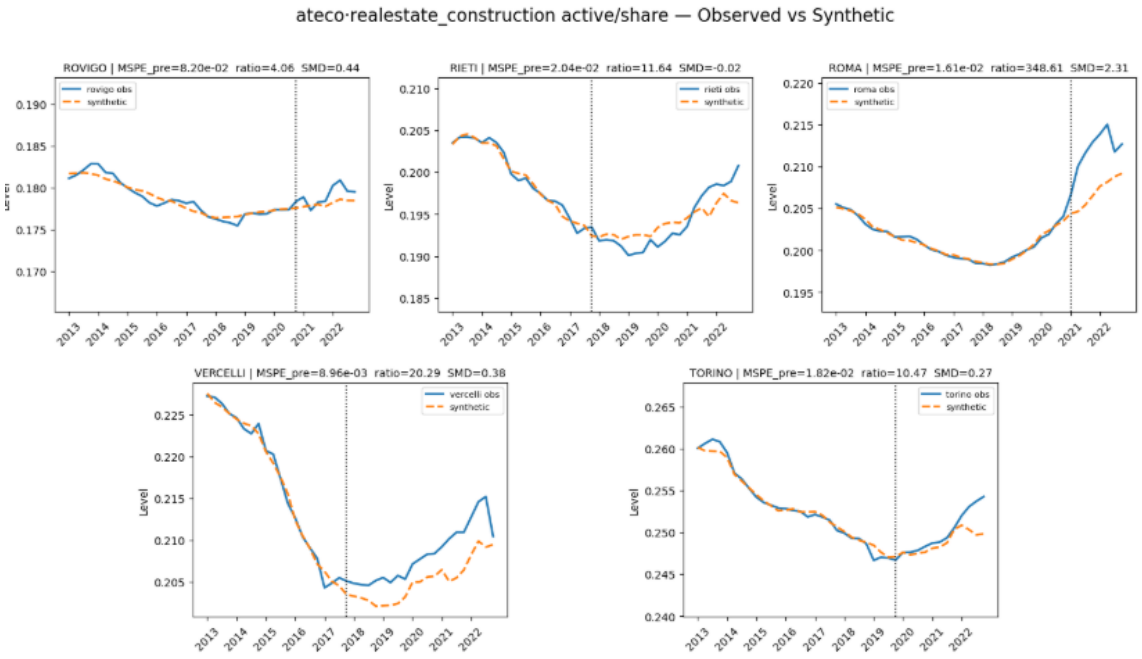

Figure 3: *

(a) Treated vs. synthetic and ATT.

--- Outcome: ateco·realestate_construction active/share ---

	Province	ATT_mean_post	SMD	MSPE_pre	ratio_safe	p_cond_prefit
0	rovigo	0.000968	0.442378	0.081961	4.064677	0.038462
1	rieti	-0.000065	-0.016176	0.020446	11.641455	1.000000
2	roma	0.004971	2.306637	0.016066	348.611762	0.032258
3	vercelli	0.003116	0.379371	0.008957	20.292610	0.562500
4	torino	0.001174	0.269605	0.018224	10.471508	0.200000

Figure 4: *
(b) Summary statistics by province.

As for real estate activities, in the pre-treatment period all provinces showed a negative trend and a reduction of the importance of this sector in the local economy. After the opening of the FC, Rome and Rovigo display the most significant and positive results. As for Rovigo, it is important to emphasize that in a province with 227 000 residents, 50 000 of which living in the main city, the opening of a new facility employing around 1000 new stable workers and many seasonal ones can have an interesting impact not only for the actual construction of the facility, but also for the accomodation services of the surrounding area, which can be boosted by the increase of demand.

This can be also true for Rome, in which FC, as previously mentioned, are currently working more than 2000 employers with full time position. An expansion of all the services that take care of the need of fast, easy accomodations is possibly plausible considering the logistic hub that is placed in the same province.

Also Turin and Vercelli show positive results, but not as significant as the previous provinces and therefore it is not possible to assess if the FC had an impact on the growth of the sector. In particular, Vercelli confirms what shown in digital services share and the entrance of the FC does not modify drastically the composition of the local labor market which remains stable. This explanation can be applied also to Rieti's results, being a much smaller province than the others and so closer to Vercelli's characteristics.

Overall, these results support the idea that Amazon helps real estate sector not only in the first part of the construction of the facility, but also in the the medium term because of the new

need for accomodations. However, because the results are not omogeneously significant, it is impossible to consider this as a pattern for local opening of an FC.

0.6.3 Apprentice Wages: `apprentice wage_mean/level`

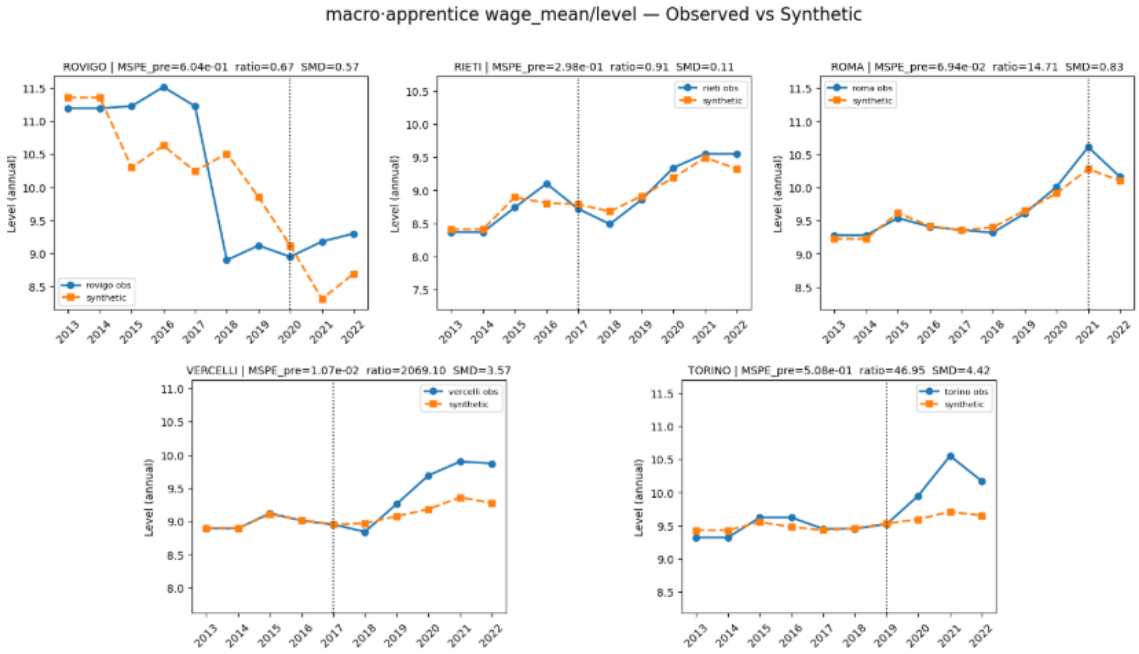

Figure 5: *

(a) Treated vs. synthetic and ATT.

`--- Outcome: macro·apprentice wage_mean/level ---`

	Province	ATT_mean_post	SMD	MSPE_pre	ratio_safe	p_cond_prefit
0	rovigo	0.634664	0.571089	0.604387	0.667816	0.740741
1	rieti	0.032722	0.113986	0.297865	0.912339	0.750000
2	roma	0.194249	0.828341	0.069426	14.705767	0.033333
3	vercelli	0.324693	3.573421	0.010700	2069.102506	0.400000
4	torino	0.527793	4.420196	0.508303	46.951402	0.040000

Figure 6: *

(b) Summary statistics by province.

The average apprentice wage show a positive increase for all the provinces, thus being significant only for two of them. Even if apprenticeships span multiple sectors, Amazon's hiring

practice could still make a difference: its job offer includes a variety of possibilities in IT, logistics and corporate positions, all governed by the National Logistics Contract.
Among all treated units, Rome displays the strongest and most credible increase in apprentice wages: the treated trajectory rises above the synthetic control immediately after the opening. Turin also exhibits a significant positive impact, suggesting that larger labor markets with more entry-level positions are more sensitive to the presence of a major employer offering wages above local standards. In Rome this can be also boosted by the presence of multiple Amazon's facilities integrated into a broader regional logistics ecosystem.

In Rovigo and Vercelli apprentice wages show positive, though more irregular divergence from the synthetic counterfactual. These movements suggest that Amazon's hiring practices may influence the lower tail of the wage distribution, but the magnitude and stability of the effect depend heavily on local absorptive capacity and competition among companies offering entry level positions. Also in this outcome Rieti shows almost no discernible apprentice wage effect that can be related to the opening of the FC.

0.6.4 Blue-Collar Wages: `blu_collar wage_mean/level`

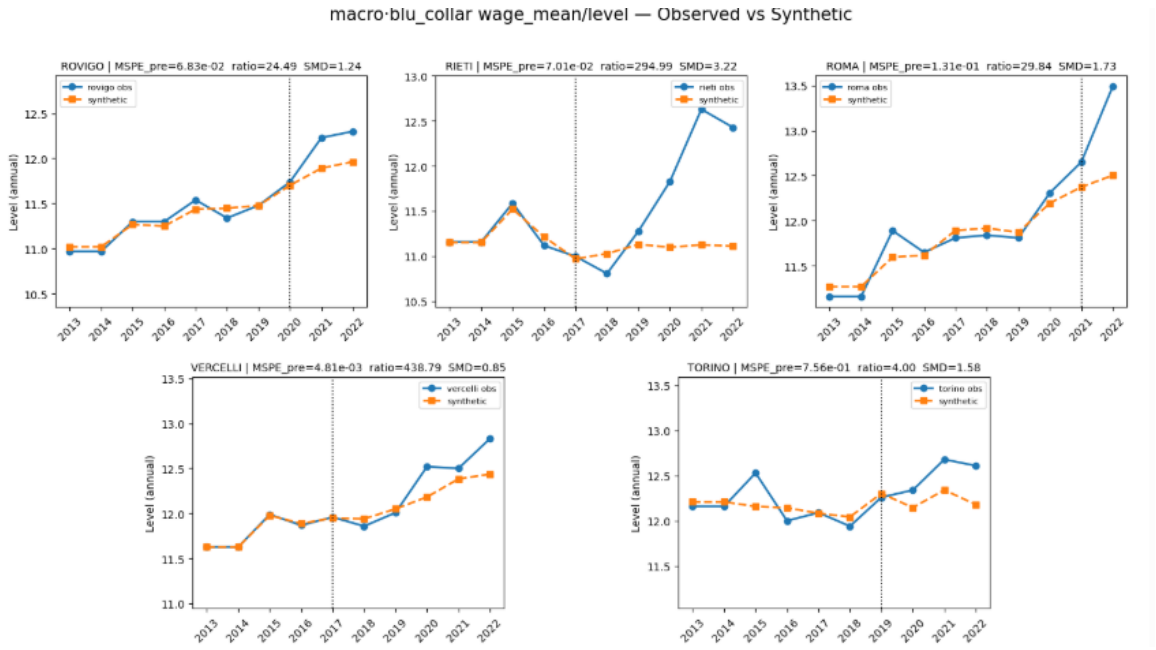

Figure 7: *

(a) Treated vs. synthetic and ATT.

`--- Outcome: macro·blu_collar wage_mean/level ---`

	Province	ATT_mean_post	SMD	MSPE_pre	ratio_safe	p_cond_prefit
0	rovigo	0.303536	1.235864	0.068338	24.490570	0.034483
1	rieti	0.663022	3.223700	0.070068	294.986676	0.562500
2	roma	0.631677	1.726370	0.131469	29.842600	0.032258
3	vercelli	0.138109	0.854101	0.004806	438.789548	0.533333
4	torino	0.292875	1.580151	0.755939	4.001999	0.208333

Figure 8: *

(b) Summary statistics by province.

The blu-collar wage outcomes present the clearest and most robust evidence of Amazon's influence on the local economy and wage-setting. This is actually predictable given that Amazon's FCs majority of workers are warehouse operators, commonly young and with no degree, whose job opportunities in Italy are often limited by a fragmented market dominated by small firms. In Rome and Rovigo wages consistently rise above their synthetic counterparts, with substantial

and stable post-treatment trajectories. These increases mirror Amazon's strategic positioning as a high-wage employer within the logistic sector: the company strongly supports paying levels above local market standards, with benefits supplements such as meal vouchers, structured progression paths and repeated wage revisions over the years. Rome and Rovigo, having a strong logistic infrastructure before the opening of the FC, had a stronger impact from the entrance of the new competitor and all the small comapnies operating in the same field had to adjust. In labor markets where local firms traditionally exercise monopsonistic power, especially in logistics, transport and warehousing, the arrival of a new employer offering systematically higher wager disrupts the pre-existing equilibrium, forcing competitors to adjust compensation levels in order to retain workers.

Even if also positive, Turin, Rieti and Vercelli's effects can't be considered reliable. This can be consistent with the idea that when Amazon operates in territories with sufficient density of logistics and manufacturing activity, its wage policies exert substantial competitive pressure, lifting wages across the operational segments of the local labor market.

0.7 Conclusion

This thesis has examined the local economic effects of Amazon Fulfillment Centers on Italian provinces using high-frequency administrative data and an Augmented Synthetic Control framework. The questions underneath my research were mainly: how does the opening of an Amazon FC reshape local economic structures, sectoral composition and labor market dynamics in Italian provinces? It questioned whether the FCs truly create new employment or merely displace existing firms; whether they stimulate complementary sectors or crowd out traditional ones; whether wage gains benefit workers or are offset by rising cost of living and whether the effects observed in the US literature can be expected to hold in a country as Italy, with a structurally different productive system and rigid labor market institutions.

Across all outcome dimensions, the most significant and interesting results showed major changes in the labor market, with digital and real estate companies increasing their relative importance in the local context, and in macroeconomic setting, with an important increase in

apprentice and blu-collar wages. However, the results were not the same for all the treated provinces, and therefore the underlying pattern discovered shows that Amazon generates the strongest and most persistent local impact when its facilities operate as part of a broader logistics hub rather than as isolated establishments.

In Rome and Turin, where the FCs are embedded within dense, diversified and technologically advances ecosystems, the opening of a new facility triggers significant structural adjustments. These provinces offer more dynamic environments and a diversified economy, with many activities able to interact productively with Amazon's operations (ICT services, transport, construction, business support, manufacturing), able to absorb and amplify Amazon's demand. In these cases, Amazon's FC does not act as an isolated shock, bus becomes a catalyst, accelerating transformations already underway. Digital-intensive firms expand in relative importance, real estate and construction regain weight and wages, especially for blu-collar workers, rise consistently above their counterfactual trajectories. The result is a structural reallocation of local economy, with complementary sectors growing faster than the others and reshaping the internal composition of firms.

In Rovigo and Vercelli the results show milder but still notable effects. In these provinces the FC stimulates some expansion in digital services and real estate activity and generates positive wage effects, but the magnitude and stability of these adjustments are more limited. This reflects lower structural elasticity, meaning that their economic fabric is less responsive to large external shocks. Their firm population is smaller, less diversified and less interconnected, so the capacity of local companies to scale in response to Amazon's demand is more limited.

Lastly, Rieti represents a particular case were the isolated and relatively small FC is unable to generate strong results with the exception of the increase of the share of digital companies in the local territory.

Overall, it is possible to affirm that Amazon does not reshape local economies uniformly. Its effects depend critically on whether the hosting territory is capable of functioning as a logistic and technological hub. Only provinces that possess a dense network of complementary activities and the structural elasticity needed to integrate a large employer as Amazon experi-

ence significant changes. However, there are not either proofs for a crowding out effects and a reallocation of workers over the territory, with migrations data not showing any significant results. This means that if it is proven that wages have an increase in the sectors interested by the FC, there are no elements to assess whether these improvements are offset by the increasing cost of life. It is true though that real estate increase may imply an higher cost of living, but this has to be proven in further analysis. Italian provinces seem to respond heterogeneously and US results are not completely seen in this context.

Bibliography

[1] Pathania, Vikram, and Serguei Netessine. *The Impact of Amazon Facilities on Local Economies*. Amazon Research Report, n.d.

[2] Cunningham, Evan R. *Local Labor Market Effects of Amazon*. Working Paper, University of Minnesota, October 31, 2024.

[3] Cunningham, Evan R. *The Local Economic Impacts of Amazon: Evidence from Fulfillment Centers and Lessons for Local Economic Development*. PhD diss., University of Minnesota, 2025.

[4] Houde, Jean-François, Peter Newberry, and Katja Seim. *Nexus Tax Laws and Economies of Density in E-Commerce: A Study of Amazon's Fulfillment Center Network*. NBER Working Paper No. 23361. Cambridge, MA: National Bureau of Economic Research, 2017. Revised April 2022.

[5] Henaway, Mostafa. "Amazon's Distribution Space: Constructing a 'Labour Fix' through Digital Taylorism and Corporate Keynesianism." *ZFW – Advances in Economic Geography* 67, no. 4 (2023): 202–216.

[6] Goldfarb, Avi, and Catherine Tucker. "Digital Economics." *Journal of Economic Literature* 57, no. 1 (2019): 3–43.

[7] Hassel, Anke, and Felix Sieker. "The Platform Effect: How Amazon Changed Work in Logistics in Germany, the United States and the United Kingdom." *European Journal of Industrial Relations* 28, no. 3 (2022): 363–382.

[8] Chava, Sudheer, et al. *Labor Market Effects*. Working Paper, 2023.

[9] Li, Elmer Zongyang. *E-Commerce and Regional Inequality: A Trade Framework and Evidence from Amazon's Expansion*. Working Paper, December 2024.

[10] Amazon EU. *Amazon EU Economic Impact Report 2023*. Amazon Corporate Publication, 2023.

[11] Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. "Synthetic Control Methods for Comparative Case Studies." *Journal of the American Statistical Association*, 2010.

[12] Abadie, Alberto, Alexis Diamond, and Jens Hainmueller. "Comparative Politics and the Synthetic Control Method." *American Journal of Political Science*, 2015.

[13] Ben-Michael, Eli, Avi Feller, and Jesse Rothstein. "The Augmented Synthetic Control Method." *Journal of the American Statistical Association*, 2021.

[14] Doudchenko, Nikolay, and Guido Imbens. "Balancing, Regression, Difference-in-Differences and Synthetic Control Methods: A Synthesis." 2017.

[15] Xu, Yiqing. "Generalized Synthetic Control Method: Causal Inference with Interactive Fixed Effects Models." *Political Analysis*, 2017.

Appendix

Methodological framework: from synthetic control to augmented synthetic control

This section develops the full theoretical grounding of these methods and goes through the key elements of each.

Classical Synthetic Control Method (SCM)

The key idea on which the classical Synthetic Control Method is based is that a combination of units often provides a better comparison for the treated unit than any single unit alone. The aim of this approach is to find for each unit exposed to the intervention a synthetic counterfactual

using a weighted average of control units, known as a synthetic control, that matches the treated unit's pretreatment outcomes and then predict the counterfactual's post-treatment trajectory and compare it with the one observed in the treated.

The counterfactual has to be built using only the non-treated pre-treatment observations. Weights on each non-treated province are chosen to reproduce the pre-treatment characteristics of the treated unit as closely as possible. The estimated impact is then the difference in post-treatment outcomes between the treated and the synthetic control.

Let $Y_{it}(0)$ denote the potential outcome for unit i at time t in the absence of treatment and $Y_{it}(1)$ the treated potential outcome that would be observed for unit i at time t once exposed to the intervention. The assumption is that before the treatment $Y_{it}(0) = Y_{it}(1)$ and there are no anticipation effects before T_0.

Because only $Y_{it}(0)$ is observed, SCM relies on the key assumption that the evolution of outcomes across units can be described by a linear factor model (LFM). This data-generating process states that untreated potential outcomes can be defined as:

$$Y_{it}(0) = \delta_t + \theta_t' Z_i + \lambda_t' \mu_i + \varepsilon_{it},$$

where δ_t captures common time shocks affecting all units, Z_i are observed covariates with time-varying coefficients θ_t, $\lambda_t' \mu_i$ represents the interaction between time-varying unobserved factors λ_t and unit-specific loadings μ_i and ε_{it} is an idiosyncratic error term. In other words, LFM assumes that units exposed to similar observed and unobserved components (similar Z_i and μ_i) will show similar trajectories of $Y_{it}(0)$ in the pre-treatment period. Since the unobserved factor loadings μ_i cannot be directly measured, studying the patterns of pre-treatment outcomes can give valuable information of these latent components. The synthetic control is therefore built to closely reproduce the pre-treatment path of the treated unit, implicitly matching both observed and unobserved determinants of the outcome.

When building the counterfactual, SCM evaluates all possible combinations of the donor pool's units, assigning for each a possible weight. SCM's weights $W = (w_2, \ldots, w_{J+1})$ have

to satisfy:

$$w_j \geq 0 \qquad \text{and} \qquad \sum_j w_j = 1,$$

W^* is chosen such that the difference between pre-treatment characteristics of the treated unit, X_1, and those of the synthetic control, $X_0 W$, is minimized:

$$W^* = \arg\min_W \|X_1 - X_0 W\|_V,$$

where V is a positive semi-definite matrix reflecting the predictive relevance of each covariate. In other words, SCM actively searches for linear combination of the donor pool's observations. This implies that in order to successfully reconstruct the counterfactual trajectory of the treated unit, the treated unit's pre-treatment outcomes must be expressible as a convex combination of the donor pool outcome units. This requirement is known as the *convex hull condition*: the synthetic control can only reproduce the treated unit if its pre-treatment vector lies within the convex hull of the donor pool:

$$X_1 \in \text{Conv}\{X_2, \ldots, X_{J+1}\}.$$

Under the Linear Factor Model and when the convex hull condition is satisfied, matching pre-treatment outcomes ensures that the synthetic control also matches the relevant unobserved factors. Therefore, deviations between the treated and the synthetic control in post-treatment periods can be interpreted as causal effects of the intervention.

If the treated unit lies outside this region, there is no positive combination of donor units that can replicate its pre-treatment path, which means that it is not possible to approximate μ_1 using donor pool's observations. SCM cannot match the latent structure of the treated unit and therefore without a reliable analysis of the pre-treatment period, it is impossible to accurately predict the post-treatment trajectory of the synthetic. Unfortunately, this was exactly the case of my setting: none of my treated provinces could be replicated using a positive linear combination of the non-treated outcomes of their own donor pool. Therefore, the research dived deeper in more robust types of analysis.

Augmented Synthetic Control Method (ASCM)

The key idea behind the Augmented Synthetic Control Method (ASCM) is simple: it improves the traditional Synthetic Control (SC) whenever the synthetic unit does not perfectly reproduce the treated unit during the pre-treatment period. There are two main reasons why ASCM manages where SCM fails. First, in ASCM some of the typical restrictions of SCM are not valid: weights do not have to be always positive. Second, ASCM uses a ridge penalization to correct the bias created by an unreliable match between counterfactual and treated.

Let X_i denote the vector of pre-treatment outcomes for unit i. ASCM introduces an explicit outcome model regression for $Y_{it}(0)$ estimated using only donors:

$$\hat{m}(X_i) = \hat{\eta}_0 + X_i^\top \hat{\eta},$$

where $(\hat{\eta}_0, \hat{\eta})$ solve the ridge-penalized problem

$$(\hat{\eta}_0, \hat{\eta}) = \arg\min_{\eta_0, \eta} \left[\sum_{i \in \text{donors}} \left(Y_{iT}(0) - \eta_0 - X_i^\top \eta \right)^2 + \lambda_{\text{ridge}} \|\eta\|^2 \right].$$

The resulting bias term is

$$\widehat{\text{bias}}_t = \hat{m}(X_1) - \sum_{j=2}^{J+1} w_j^{\text{SC}} \hat{m}(X_j).$$

The ASCM counterfactual is

$$\widehat{Y}_{1T}^{ASCM}(0) = \underbrace{\sum_j w_j^{SC} Y_{jT}}_{\text{Synthetic Control prediction}} + \underbrace{\left[\widehat{m}(X_1) - \sum_j w_j^{SC} \widehat{m}(X_j) \right]}_{\text{Bias correction term}}. \tag{2}$$

$$\hat{Y}_{1t}^{\text{ASCM}}(0) = \hat{Y}_{1t}^{\text{SC}}(0) + \widehat{\text{bias}}_t.$$

It is clear that if the SC model perfectly fits in the pre-treatment there will not be bias and therefore the two models will coincide. However, if there is an unreliable match between counterfactual and treated, the ASCM will address for the bias and improve the model.

ASCM can also be written purely as a weighted average of donor outcomes:

$$\hat{Y}_{1t}^{\text{ASCM}}(0) = \sum_{j=2}^{J+1} w_j^{\text{ASCM}} Y_{jt},$$

These weights are not the original SCM weights, but they are adjusted weights that incorporate both the need for better pre-treatment balance and the desire to stay close to the original SCM solution. w^{ASCM} are defined as the solution to the following optimization problem:

$$w^{\text{ASCM}} = \arg \min_{w:\, \sum_j w_j = 1} \left\{ \frac{1}{2\lambda_{\text{ridge}}} \|X_1 - X_0 w\|^2 + \frac{1}{2} \|w - w^{\text{SC}}\|^2 \right\}.$$

The term $\|X_1 - X_0 w\|^2$ is the balance term: it measures the difference between the treated unit's pre-treatment features X_1 and the weighted average of donor ones $X_0 w$. The aim of the optimization is minimize this term prioritizing a good fit.

The term $\|w - w^{\text{SC}}\|^2$ is the regularization term: it penalizes weights that deviate too much from the original SC's weights, keeping the estimator from becomig unstable or relying too heavily on individual donors.

The constraint $w : \sum_j w_j = 1$ garantees that ASCM synthetic counterfactual remains a weighted average and therefore interpretable as a synthetic unit.

The parameter λ_{ridge} is the trade-off controller: if it is large, the balance term receives less weight and so the optimization stays close to w^{SC} and ASCM will be similar to SC; if λ_{ridge} is small the balance term becomes more important, so the weights are allowed to move further from w^{SC} to improve pre-treatment fit, resulting in a more flexible model.

ASCM Implementation in Amazon's FCs Case

To study Amazon's FC specific case, it is important to address to some peculiarities. First, the original donor pool is wide (101 non-treated provinces), which can easily generate synthetics that perfectly match the treated, creating an overfitting problem. Second, if for ateco and juridical form variables the values change for each quarter, macro variables have the same values over the year, and therefore it shrinks the pre-treatment period observation pool. Therefore, the model needed some adaptations to fully capture the effects of FC's on the territory.

(1) Donor pool construction and basic constraints

The construction of the correct donor pool for each treated province is the first step to guarantee a reliable counterfactual. Donor units must not be exposed to similar shocks or spillovers in order to have a credible model and therefore, for each treated province, the donor pool was restricted to provinces that have never been treated before T_0 and are beyond a minimum geographical distance threshold of 100km. In this way, possible spillover effects that arise from the FC are excluded and donor pool, ensuring that donors are temporally and geographically unaffected by the FC openings, which reduces the risk of contamination and strengthens the interpretation of the counterfactual. The result is a range of donor pool's observations between 82 and 96 for each treated province.

Moreover, each specification must include at least 16 pre-treatment observations and a minimum number of effective donors. This mirrors the identification guidance in Abadie 2010 and the factor-based arguments of Xu 2017, both showing that SC/ASCM relies on sufficiently long pre-treatment periods to estimate latent factors and stable relationships. Short pre-treatment windows or too few donors often lead to overfitting and unreliable post-treatment inference and must be excluded.

(2) Standardization of treated and donor series

Before fitting the ridge estimator, both the treated outcome and donor features are standardized. This step is necessary to allow the ridge penalty to apply equally to all predictors, creating a fair shrinkage. In other words, the penalized problem solved by the ridge estimator can be written as:

$$\min_{w} \ \|Y - Xw\|^2 + \lambda\|w\|^2.$$

Without standardization, two predictors in X with different scales might not be evaluated in the same way, giving more weight to the larger one even if the second is more informative. Standardization fixes this by forcing every column to have

$$\text{mean}(X_j) = 0, \quad \text{sd}(X_j) = 1.$$

After fitting the model, I transform predictions and residuals back to the original units to

preserve interpretability. This is consistent with the idea, in Ben-Michael et al. (2021) that controlling pre-treatment fit is meaningful only after suitable scaling and is also explicitly recommended by Feller Rothstein (2021) to Doudchenko, Imbens (2017) and Xu (2017), improving numerical stability and producing interpretable weight patterns.

(3) Dimensionality reduction via PCA

In the standard ASCM the donor weight matrix X_0 may directly contain many series, possibly including covariates and lags. Having more than 80 donors may mean that series become highly collinear or noisy. Therefore, I applied Principal Component Analysis (PCA) to the donor matrix before fitting the ridge regression. With X being the standardized donor matrix in the pre-treatment period, I used:

$$Z_{\text{pre}} = P^{\top} X_{\text{pre}},$$

where the columns P are the first k eigenvectors. Ridge ASCM is then run on Z_{pre} instead of the full X and the same PCA transformation is applied to the full sample before prediction. In this way, PCA keeps the common factors that explain donor variation with the direct effect of reducing the effective dimension of the problem and improving stability of ridge fit. In other words, the reduction in dimensionality should help avoiding overfitting while still capturing dominant latent factors driving variation among provinces.

(4) Placebo-based blocked cross-validation for the penalty parameter

One of the core parameters in the model is the Ridge penalty term λ and therefore it has to be chosen carefully. The optimal λ has to penalize the right parameters and avoid too strong extrapolation. Ben-Michael et al. (2021) proposed choosing the ridge parameter λ_{ridge} by cross-validation using only pre-treatment data, leveraging "placebo" splits of the pre-period into training and validation blocks. To each λ, I fit the ASCM on the training block and compute the Mean Squared Prediction Error on the validatation block. Then, I averaged MSPE across different cuts and chose the λ that minimizes this average MSPE.

Moreover, I impose a minimum value for λ to avoid "almost-OLS" solutions, which occur when regularization is too weak. Such solutions are known to generate unstable or extreme weights that extrapolate heavily outside the convex hull of the donors. Imposing a lower bound ensures

controlled extrapolation and increases robustness, stabilizing ASCM weights.

(7) Distinct tuning behaviour for macro variables

In the panel used, macroeconomic variables are significantly different from ateco and juridical ones: their values are constant over the year and typpicaly exhibit slower dynamics and strong common factors. Therefore, I adopted a more restrictive PCA, with maximum two components, reflecting the idea that macro series are driven by a few common trends, and I used a grid with higher values to choose λ to reduce the model's flexibility. This distinction follows the observation in Abadie 2010 and Xu 2017 that macro variables tend to share strong common trends, while micro outcomes often contain more idiosyncratic variation. Adapting the model to the nature of each outcome improves stability without sacrificing sensitivity.

In some extensions of SCM macro variables are used as additional regressors and controls. In Amazon's FC specific case, I tried different combinations of macros but the best outcome was the one that had none, because it preserved the model from overfitting.

Placebo test

Placebo tests are often used in SCM and ASCM models to assess the reliability of the effects estimated in the models. These tests answer to the simple question: applying the same method to units that were not treated, how often the effects would be as large as the ones estimated for the treated unit?

There are two classic versions: placebo in-space, which simulates the scenario were other non-treated units were treated, and placebo in-time, which pretends that the treated units were treated earlier in time. Both ideas were firstly introduced in the oridinal synthetic control papers (abadie, Diamond, Hainmueller 2010, 2015) and are now standard in ASCM literature (Ben-Michael, Feller, Rothstein 2021) as diagnostic and inference tools.

Having a range of possible donor between 82 and 96, I chose to use placebo in-time for my model. Practically, given a T_0 time of treatment, the tests selects a placebo time $\tilde{T_0}$ such as $\tilde{T_0} < T_0$ by assumption. Then, time series' are split at T_0 and recompute pre and post-MSPEs

as if the treatment started at $\tilde{T}_0$. The placebo ratio will be:

$$R(\tilde{T}_0) = \frac{\text{MSPE}_{\text{post}}(\tilde{T}_0)}{\text{MSPE}_{\text{pre}}(\tilde{T}_0)}.$$

Repeating this for multiple cutoff $\tilde{T}_0$ within the pre-period gives a distribution of MSPE ratios when no treatment actually occurred. If

$$R(T_0) = \frac{\text{MSPE}_{\text{post}}(T_0)}{\text{MSPE}_{\text{pre}}(T_0)} > R(\tilde{T}_0)$$

than it is possible to assume that the real intervention is unusually strong compared to what happens when T_0 is randomly assigned.

Selected Outcomes and Empirical Results

In this section I evaluate the most significant results. The ASCM and in-time placebo test were performed on all outcomes, but to select the most valuable ones there are three main controls. First, for all outcomes the post-treatment effects are evaluated across all treated provinces and retained only those outcomes where at least four provinces exhibit:

$$|\text{SMD}| > 0.15, \qquad \text{sign}(\text{ATT\_mean\_post}) = \text{sign}(\text{SMD}),$$

In other words, the outcomes are filtered to keep only the ones that have a non-negligible magnitude and are coherent between provinces.

Second, all the outcomes passing the coherence filter are then checked for significance: it is retained if at least two provinces satisfy

$$p_{\text{cond}} \leq 0.10,$$

It implies that the observed difference between treated and synthetic paths is statistically significant, relatively to placebo cuts inside pre-treatment period.

Finally, synthetic's outcomes showing paths perfectly identical to the treated or almost flat have to be eliminated, ensuring that selected results are not driven by instability or poor pre-treatment fit.

The key controls for overfitting or underfitting are MSPE, standard deviation and the λ selected for that particular outcome in the pre-treatment. A typical overfitting scenario would show an extremely low λ and MSPE: the first one meaning that there is almost no regularization and the model is free to adapt to the treated path in the pre-treatment; as for the null MSPE, it implies that there are no differences between the treated and the synthetic and therefore the model will not be able to generalize in the post-treatment period because it learned also the noise.
On the other hand, huge λ and too low standard deviation of the synthetic may be clues of an underfitting scenario: an extremely high value of λ implies that the regularization is so strong that the model risks to be flat by construction; the standard deviation check makes sure that the variation of the synthetic is not extremely lower than the treated one, meaning that it is able to follow the pattern.
If any of these checks fails, that particular outcome will not be considered.

To measure the effectiveness of the filters and the credibility od the synthetic control, I computed a set of global diagnostics based on the pre-treatment period for all estimated outcomes and then repeated the same diagnostics only on the final set of selected outcomes.

The first metric is the pre-treatment coefficient of determination, defined as $R^2_{\text{pre}} = 1 - \frac{\sum_{t \le T_0} (y_{it}^{\text{obs}} - \hat{y}_{it})^2}{\sum_{t \le T_0} (y_{it}^{\text{obs}} - \bar{y}_i)^2}$. It measures the proportion of the pre-treatmeent variance of the treated unit that is explained by its synthetic counterpart. To complement the average performance, I also record $R^2 < 0.3$, which identifies cases in which the synthetic control explains little of the pre-treatment variation, suggesting that any post-treatment deviation would be difficult to interpret causally.
A second diagnostic is the pre-treatment correlation, defined as $\text{corr}_{\text{pre}} = \text{Corr}\big(y_{it}^{\text{obs}}, \hat{y}_{it}\big)$, which focuses on the co-movement between the treated and synthetic series. It captures whether the synthetic control tracks quarter-to-quarter fluctuations, with low values of correlation indicating that the synthetic series fails to reproduce key temporal patterns even if the overall level may be

similar.

A third metric concerns the relative variability of treated and synthetic series. $\text{sd\_ratio}_{\text{pre}} < 0.25$, which compares the standard deviation of the synthetic control to the one of the treated unit before treatment. If this ratio is very small, the synthetic series is nearly flat relative to the treated series, as previously mentioned for underfitting scenarios.

=== Comparison: ALL vs SELECTED outcomes ===

	metric	value_all	share_of_total_all	value_selected	share_of_total_selected
0	n_total_results	290	1.000000	25	1.00
1	n_bad_fit_flag	38	0.131034	2	0.08
2	n_flat_flag	20	0.068966	1	0.04
3	n_overfit_flag	0	0.000000	0	0.00
4	n_R2<0.3	33	0.113793	2	0.08
5	n_corr<0.5	9	0.031034	1	0.04
6	n_sd_ratio<0.25	28	0.096552	1	0.04
7	n_sd_ratio>4.0	0	0.000000	0	0.00

Figure 9: Global Diagnostic.

Comparing global diagnostics before and after the filtering, they show a substantial improvement in the quality of synthetic control fit, with a fall of unreliable province-outcome combinations from 13.1% to 8% and a reduction both in R^2 and $corr_{pre}$.

Using this selection process, the following outcomes emerge as the most robust and informative.