POLITECNICO DI TORINO

Corso di Laurea in Ingegneria Gestionale

The Impact of Regulation on Short-Term Rentals

Evidence from the Italian Airbnb Market

Prova finale di

LUIGI BOCCIA

Relatore: Francesco Luigi Milone – A.A. 2025-2026

Contents

List of Figures

List of Tables

Chapter 1
Introduction

1.1 Airbnb and the sharing economy revolution

Airbnb, founded in 2007 in San Francisco by Brian Chesky and Joe Gebbia, is one of the most emblematic examples of how digital platforms have transformed traditional markets through technological disintermediation. Initially created as a solution for renting air mattresses during conferences in San Francisco, the platform quickly evolved into a global marketplace connecting accommodation owners with travellers seeking alternative accommodation solutions to traditional hotels.

The platform's growth has been exponential: from a few hundred listings in 2009, Airbnb reached over 8 million listings in more than 220 countries in 2024[1]. The company's Initial Public Offering (IPO) in December 2020, with an offering price of $68 per share and a valuation of $47 billion, saw shares more than double on the first day of trading, reaching a market capitalisation of approximately $86.5 billion[2], marking the consolidation of a business model that has redefined the global hospitality industry.

1.1.1 The business model: the two-sided market

From an economic perspective, Airbnb operates as an intermediary platform in a two-sided market, connecting two distinct groups of users: hosts (property owners offering accommodation) and guests (travellers seeking accommodation) (Rochet & Tirole, 2006; Zervas et al., 2017). The platform generates value by drastically reducing the search and transaction costs that characterised traditional tourist rental markets, while simultaneously offering:

- Hosts: access to a large base of potential customers, booking management tools, secure payment systems, and insurance coverage[3];
- Guests: a wide choice of accommodation with sophisticated search filters, host reviews and ratings, instant booking, and refund guarantees in case of problems.

[1] Q4 2024 Shareholder Letter
[2] CNBC - Airbnb skyrockets 112% in public market debut, giving it a market cap of $86.5 billion
[3] Airbnb Resource Center - "How Airbnb protects hosts"

The revenue model is based on commissions applied to both sides of the market. Traditionally, Airbnb has used a “split-fee” model with a 3% commission for hosts and a service fee for guests of between 14.1% and 16.5%[4], and a “single fee structure” where the entire 15.5% commission is deducted from the host's payout, with no separate fee for the guest[5]. The latter is mandatory for hotels and hosts using Property Management Software. This asymmetric pricing structure reflects the competitive dynamics of the market: the higher commission on the demand side keeps the offer attractive to hosts, who are the most critical side of the market as providers of the scarce resource (accommodation).

1.1.2 Network effects and disruption

Airbnb's success is linked to the network effects that characterise multi-sided digital platforms: the greater the number of hosts on the platform, the greater its attractiveness to guests, and vice versa, creating a virtuous circle of growth (Evans & Schmalensee, 2016). The increase in the number of listings available makes the platform more attractive to travellers, increasing traffic and bookings; this, in turn, encourages more owners to sign up, creating a virtuous circle of growth. This self-reinforcing dynamic creates significant barriers to entry for potential competitors and contributes to market concentration.

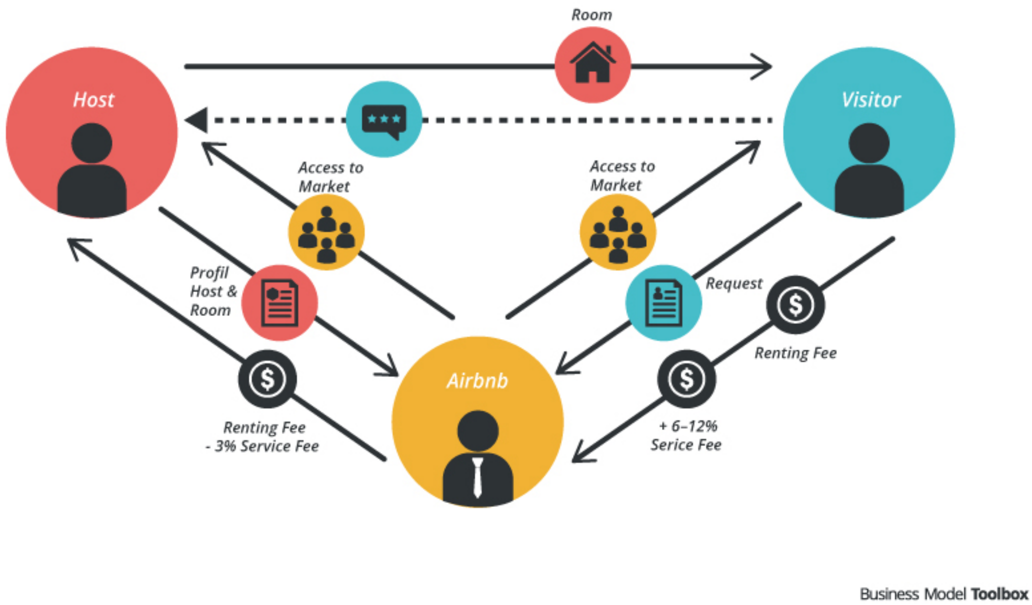

Figure 1. Airbnb Business Model. Source: bmtoolbox.net

[4] Airbnb Resource Center - "How much does Airbnb charge hosts?"
[5] Airbnb Help Center - "Airbnb service fees"

The platform has disintermediated traditional channels (travel agencies, local property portals), allowing hosts to reach end customers directly. This has reduced traditional intermediation margins but has also transferred a significant share of the value created through commissions to the platform itself. In addition, Airbnb provides services that traditional intermediaries did not offer, such as reputation systems based on bilateral reviews, dispute resolution mechanisms, and insurance guarantees.

1.1.3 Operational evolution: from sharing to professional management

In its early years, Airbnb embodied the pure sharing economy model: owners who occasionally rented out their primary residence or a spare room to supplement their household income. This peer-to-peer model was characterised by non-professional hosts, occasional activity, and a strong component of social interaction between hosts and guests.

Over the years, however, a gradual professionalisation of the market has been observed. According to industry data, the share of listings managed by hosts who own or manage multiple properties has grown significantly. According to one study, professional hosts, defined as those who manage multiple properties or rent full-time, generate approximately 71% of Airbnb's revenue in the twelve major markets (Xie & Kwok, 2021). Professional operators have emerged who manage dozens or hundreds of units, often through dedicated property management companies, transforming what began as an occasional peer-to-peer exchange into a structured business activity.

At the same time, self-check-in has become widespread, i.e. a method of handing over keys and accessing accommodation that does not require the physical presence of the host. This includes the use of key boxes (key boxes with codes), electronic locks with numerical codes, or instructions for retrieving hidden keys. The smart lock market is estimated to grow at a Compound Annual Growth Rate (CAGR) of 19.7% until 2030, indicating that their use has become an established practice in the short-term rental (STR) market[6]. Self-check-in allows for high turnover to be managed without the need for direct coordination with guests upon each arrival.

This evolution has created tensions with traditional regulatory frameworks, which were designed for an accommodation market where guests were necessarily identified through direct interaction at a hotel reception desk. The issue of guest identification in short-term

[6] Smart Lock Market Size & Outlook, 2030

rentals has thus become a critical point of regulatory debate in many countries, including Italy.

1.2 The STR market in Italy and the issue of guest identification

In Italy, the short-term rental market has developed rapidly since the second half of the 2010s, following global trends but with some specific characteristics linked to the country's strong tourist appeal. According to 2020 data, Italy ranked second in the world in terms of the number of Airbnb listings, with approximately 180,000 listings, after the United States[7]. The concentration of listings is particularly high in major urban centres: Rome leads the way, followed by Milan and Florence, confirming that the short-term rental market is strongly polarised in large cities and areas with greater tourist appeal and potential demand. The Italian regulatory framework has addressed the phenomenon of short-term rentals in a fragmented manner. At the fiscal level, Decree-Law 50/2017[8] (converted into Law 96/2017) introduced in Article 4, paragraph 5, the obligation for intermediaries to apply a 21% withholding tax on rents at the time of payment to the landlord. The same Article 4, paragraph 2, also provided for the possibility of opting for a flat-rate tax initially set at 21%, which was subsequently modified, with platforms operating as withholding agents. At the local level, many municipalities have introduced tourist taxes and registration requirements on municipal portals, creating a heterogeneous regulatory mosaic across the country.

1.2.1 Identification requirements: the Consolidated Law on Public Security

A key aspect, often discussed in a biased manner in the debate on short-term rentals, concerns the obligations to identify and report guests as provided for in the Consolidated Law on Public Security (TULPS), Royal Decree No. 773 of 18 June 1931. In particular, Article 109 requires accommodation managers to communicate the details of guests to the public security authorities within 24 hours of arrival, with a reduced deadline of 6 hours for stays of less than 24 hours. The subjective scope of the provision has been extended by way of interpretation by the legislator, clarifying that these obligations also apply to landlords or sublandlords who lease properties, or parts thereof, with contracts of less than thirty days' duration. Failure to comply with the obligations under the TULPS constitutes a misdemeanour, punishable by up to three months' imprisonment or a fine of up to €206.

[7] Radical Storage (2025) - "Airbnb Statistics: Revenue and Users"
[8] DL 50/2017, art. 4, comma 5, convertito con L. 96/2017

The regulation, originally designed for hotel accommodation, does not specify the technical methods by which the guest's identity must be verified. In traditional practice, identification is carried out by directly checking the document upon arrival. With the spread of short-term rentals managed through self-check-in, an important question of interpretation has arisen in terms of compliance: is it sufficient to acquire and transmit the data from the document provided by the guest, or is it necessary to verify the document?

1.2.2 The Circular of the Ministry of the Interior of 18 November 2024

The issue was addressed by the Ministry of the Interior in the circular issued by the Department of Public Security, ref. no. 38138 of 18 November 2024, concerning the identification of persons staying in accommodation facilities and tourist rentals. Without introducing any new provisions, the circular reiterated a restrictive interpretation of the obligations already provided for in Article 109 of the TULPS, stating that identification must be carried out by means of a visual check, i.e. through a direct verification of the correspondence between the identity document and the person. In this context, check-in procedures that do not allow for such verification (e.g. delivery of keys via key boxes or similar systems without identity verification) are non-compliant; the verification may also be carried out by a person delegated by the host.

In terms of case law, the circular has been the subject of litigation: with judgment no. 10210 of 27 May 2025, the Lazio Regional Administrative Court ordered its annulment; subsequently, in judgment no. 9101 of 21 November 2025, the Council of State reformed this decision, reaffirming the legitimacy of the obligation of “de visu” identification and specifying that it can also be satisfied through synchronous (real-time) video identification, provided that it is suitable for guaranteeing the certainty of identity before access.

From a penalty perspective, violation of the obligations under Article 109 of the TULPS is punishable under Article 17 of the TULPS; liability falls on the person providing accommodation (manager/host), not on the intermediary platforms.

1.2.3 Economic implications of regulation

From an economic perspective, the introduction of the requirement for visual identification can be interpreted as a shock to hosts' variable operating costs. Unlike regulations that introduce fixed compliance costs (e.g., the requirement for safety equipment or one-off administrative registration), the requirement for physical presence at each check-in generates a cost that varies proportionally to the number of arrivals handled.

For an occasional host who rents out their home a few times a year, the requirement for physical presence may be relatively inexpensive, as the number of check-ins is limited. For a professional host who manages multiple properties with high guest turnover, however, the organisational cost can be significant: it requires time coordination (availability at variable times), possible outsourcing of the service to third parties (co-hosts, check-in agencies), or investment in compliant video identification technologies.

The direction of the effect on several types of hosts is not clear-cut a priori. On the one hand, professional hosts face a higher volume of check-ins and therefore higher overall operating costs; on the other hand, they can benefit from economies of scale in the organisation of the service (dedicated staff, standardised processes), while occasional hosts may not find it economically viable to outsource the service for a few arrivals per year. Empirical evidence is needed to understand which mechanism prevails.

1.3 Rationale for the study

The Ministry of the Interior Circular of 18 November 2024 is a particularly suitable case study for investigating how public regulation influences the functioning of digital markets, for at least three reasons.

Firstly, the intervention affects a central operational component of the short-term rental business model: arrival management. The possibility of self-check-in has been one of the main factors enabling the growth and professionalisation of the sector, as it allows the business to scale up to multiple units by reducing the constraint of physical presence. The introduction of an "in-person" identification requirement alters this balance, increasing organisational constraints and potentially changing the structure of supply and the composition of operators.

Secondly, Italian regulations offer the opportunity to analyse a type of intervention that is still relatively unexplored in the literature on short-term rentals. Much of the available empirical evidence concerns quantitative measures (e.g., limits on the number of days that can be rented or the number of properties that can be managed) or fixed compliance costs (registrations, requirements, or mandatory equipment). In contrast, the requirement to be present at each check-in represents a shock to variable costs and the operational intensity of the activity, with theoretical and empirical implications that differ from more traditional regulations.

Thirdly, the issue has immediate policy relevance. The measure's ability to achieve its stated objective, strengthening public safety through reliable guest identification, depends on host compliance and market adaptation strategies. If, for example, the intervention leads to a selective exit of operators or a reconfiguration of management methods that does not guarantee an effective improvement in identification, the regulation could prove ineffective or produce undesirable side effects. From this perspective, an empirical analysis can provide useful evidence for an ex-post evaluation of the measure's effectiveness and trade-offs.

1.4 Research question

The central research question of this thesis is as follows: What is the impact of the introduction of the obligation to identify guests in person (Circular of the Ministry of the Interior, 18 November 2024) on the Italian short-term rental market in terms of supply, operational intensity, and prices?

The empirical analysis aims to estimate the average effect of the regulation and, at the same time, to clarify through which channels the intervention is transmitted to operators and which market segments are most exposed. The main question is therefore divided into the following sub-questions.

Q1: Effect on supply and operational intensity (extensive vs intensive margin).

Has regulation reduced the presence of listings on the platform, measured by the share of active listings and the evolution of operations? Does the adjustment occur mainly on the intensive margin (reduction in days of activity/availability and capacity actually placed on the market) or on the extensive margin, understood as a prolonged cessation of activity observed on the platform?

Q2: Heterogeneity by type of host (professional vs. occasional).

Are the effects different between professional hosts (multi-property) and occasional hosts? To what extent is the response compatible with a channel linked to the variable operating costs associated with check-in management (more relevant for those managing high volumes) compared to a channel of organisational capacity (greater possibility of adaptation for structured operators)?

Q3: Heterogeneity by property type (entire home vs. other types).

Does the impact differ between entire home listings and room listings (private/shared)? This dimension is relevant because different formats may reflect different pre-policy operating modes and, in particular, a different propensity to adopt self-check-in.

Q4: Heterogeneity by pre-policy operational intensity (turnover proxy).

Do listings with higher pre-regulation intensity, measured by observable proxies (e.g., days booked/active or calendar usage indicators), show a more pronounced response than sporadic listings? Such evidence would help isolate the role of the variable cost channel, which grows with the number of check-ins managed.

Q5: Effects on prices (asking and, when observable, realised).

Does regulation affect the prices charged by hosts? Any increase in prices could indicate an attempt to pass on higher operating costs to customers; conversely, price reductions or rigidity could signal strategies to maintain occupancy or limits to pass-through in the presence of competitive constraints. Where data allow, the analysis distinguishes between asking prices and realised prices, taking into account the selection of the observable sample for the latter.

1.5 Contribution of the thesis

This thesis aims to contribute to the economic literature on digital market regulations and, more specifically, on short-term rentals, in three ways.

Firstly, it offers empirical evidence on a relatively unexplored type of intervention. The existing literature has mainly analysed quantitative regulations (e.g. limits on the number of days that can be rented) and fixed compliance costs (registrations, requirements, or mandatory equipment). Referring to Chapter 2 for a systematic discussion, this study focuses instead on an operational constraint that translates into variable costs linked to intensity of use: the obligation to identify oneself in person at each check-in. This feature makes the context particularly informative for assessing how several types of cost shocks, fixed vs. variable, affect operators' supply decisions and operations.

Secondly, the analysis exploits an institutional framework in which regulation is uniform at the national level, while operators remain highly heterogeneous in terms of relevant characteristics (degree of professionalisation, scale of operations, type of ownership). This allows us to study the heterogeneity of responses not primarily in geographical terms, as is the case in many studies on municipal regulations, but rather in terms of business model and organisational capacity, which are potentially more generalisable and central to understanding the dynamics of adaptation in platform markets.

Thirdly, the study has direct policy relevance. The effectiveness of interventions aimed at strengthening safety and traceability depends crucially on compliance and behavioural adjustments by operators. Analysing whether the requirement for visual identification leads to a reduction in supply, a downsizing of operational intensity, or a price adjustment, and for which market segments, can contribute to an ex-post evaluation of the intervention and provide useful insights into possible refinements to the regulatory design.

1.6 Theoretical framework: expected mechanisms

Without going into formal detail, which will be developed in subsequent chapters, it is useful to briefly outline the main economic mechanisms through which the obligation of visual identification can influence the behaviour of hosts.

1.6.1 Variable costs and intensive margin

The requirement for physical presence at each check-in introduces an operating cost that increases with the number of arrivals. This profile distinguishes it from regulations that mainly impose fixed costs (e.g. the purchase of mandatory equipment or one-off administrative requirements). In the presence of a variable cost linked to the activity, a host can react along two margins.

- Intensive margin: reduce the intensity of use of the listing by decreasing the days on which the property is made available or actually active for bookings, while keeping the listing on the platform.
- Extensive margin: discontinue activity on the market by making the listing persistently inactive or removing it from the platform.

The choice between these two adjustments depends on the residual profitability of the business after the introduction of the new constraint. Operators with high expected revenues per booking (or with greater capacity to organise check-ins) may find it convenient to absorb the cost and maintain operations, possibly downsizing on the intensive margin. Conversely, for hosts with tighter margins, the additional cost may make the business no longer profitable, increasing the likelihood of exit on the extensive margin.

1.6.2 Heterogeneity between professional and occasional hosts

The effect of regulation may differ between professional (multi-property) and occasional hosts through two potentially opposing channels.

- Volume channel: professional hosts manage a higher number of arrivals on average; consequently, a constraint that introduces a cost per check-in may result in a higher increase in overall operating costs. This mechanism suggests a more pronounced response, in the form of a reduction in operational intensity or a greater likelihood of cessation of activity observed on the platform.
- Organisational capacity channel: professional hosts may have resources and procedures in place that mitigate the impact of the constraint (e.g. dedicated staff, outsourced reception services, or more efficient organisation of arrival management). In this case, the response may be less pronounced than that of occasional hosts, for whom physical presence at each check-in can be particularly burdensome in terms of time, coordination, and rigidity.

The net direction of the effect is therefore ambiguous a priori and requires empirical verification. The analysis of heterogeneity by host type allows us to assess which of the two mechanisms prevails and to interpret the observed adjustment between the intensive and extensive margins.

1.6.3 Qualitative predictions

Based on the mechanisms outlined, it is possible to formulate some qualitative predictions, to be understood as hypotheses to be subjected to empirical verification and not as advance conclusions:

- Prevailing adjustments on the intensive margin: since the obligation directly affects check-in management, it is plausible to expect that part of the response will manifest itself through a reduction in operational intensity; for example, by decreasing the days of activity or availability of listings rather than through an immediate and generalised contraction of the overall supply observed on the platform.
- Heterogeneity based on pre-policy intensity: listings characterised by high activity prior to the intervention (proxy: higher turnover, active or booked days) could

react more strongly, in line with a variable cost channel that increases with the number of arrivals.
- Ambiguous heterogeneity by host and property type: as discussed, the direction of the effect for professional hosts compared to occasional hosts and for entire properties compared to other types is not determined a priori, depending on the balance between exposure to check-in volume and organisational capacity to absorb or reorganise the constraint.

These predictions will be compared with the empirical evidence presented in Chapter 5, including through heterogeneity analysis and decomposition of adjustments between intensive and extensive margins.

1.7 Empirical approach

The empirical analysis is based on a monthly panel at the listing level (property × month) and aims to measure how the outcomes of interest evolve following the introduction of the regulation. This approach allows us to compare trends before and after the intervention, distinguishing between groups of operators and listings that, based on observable characteristics, may differ in terms of organisational capacity and operational intensity (e.g., professional vs. occasional hosts, type of property, and pre-policy activity levels).

1.7.1 Identification Strategy

The core idea is to compare the pre/post evolution of outcomes between a group plausibly more exposed to regulation (e.g., professional hosts) and a comparison group (occasional hosts), estimating a Difference-in-Differences (DiD) model in a fixed effects specification on monthly panel data.

The baseline model includes:

- Property fixed effects: control for unobserved and time-invariant heterogeneity of each listing (e.g., quality, location, and structural characteristics).
- Month fixed effects: absorb common shocks over time (seasonality, macroeconomic trends, and other aggregate events).
- Standard errors clustered at the listing level: correct for potential serial correlation of errors within the same listing over time.

This Two-Way Fixed Effects (TWFE) specification allows for identification of the average treatment effect of regulation on the treated group through the treated × post interaction (while the "treated" and "post" components are absorbed by fixed effects). Identification relies on the parallel trends assumption: in the absence of intervention, the evolution of outcomes would have been similar between the treated and control groups. Operational details, alternative specifications, and empirical verification of the assumption are presented in a dedicated methodology chapter.

1.7.2 Dynamic Analysis and Event Study

Alongside the main estimation, the approach includes a dynamic analysis through event study, which allows:

- Assessment of the plausibility of the parallel trends assumption, verifying that in the months preceding the intervention no systematic differences emerge in outcome trends between the treated and control groups.
- Description of the temporal dynamics of the effect, observing whether the impact manifests immediately or gradually and whether it proves persistent or transitory in the post-intervention period.
- Identification of possible anticipatory effects, i.e., potential behavioural adjustments preceding the formal policy implementation, should operators have reacted in advance following expectations or available information.

1.7.3 Dependent Variables and Operating Mechanisms

The analysis considers several types of outcomes, with the objective of capturing both the overall effect of regulation and the main transmission channels:

- Activity/supply measures (intensive margin): active days and available days. These variables measure adjustments in the quantity of capacity actually brought to market by hosts.
- Utilization measures: booked days and occupancy rate. These indicators help distinguish reductions due to supply-side choices (lower availability) from demand-side variations (lower capacity to fill available days).
- Price measures: average price per night. Price estimates are interpreted with caution, since the "realized" component may be subject to selection: changes in activity or bookings can alter the composition of the sample for which price is observable.

1.7.4 Heterogeneity and Robustness

The analysis systematically explores effect heterogeneity along multiple dimensions, consistent with the theoretical channels discussed:

- Host type: professional hosts (operationally defined as hosts managing at least 2 properties in the dataset) vs. non-professional hosts.
- Property type: entire home listings vs. other types (private room, shared room).
- Pre-policy operational intensity: listings with high vs. low intensity in the pre-regulation period, measured through observable proxies of turnover/activity (e.g., booked days, occupancy rate, or calendar utilization indicators).

Targeted robustness checks are also conducted:

- Alternative time windows: restricting the sample to periods closer to the policy date to verify the sensitivity of estimates to the inclusion of temporally distant observations.
- Placebo tests: estimating the effect on fictitious dates in the pre-intervention period, in order to verify the absence of spurious differences and pre-existing dynamics erroneously attributed to regulation.

1.7.5 Limitations of the Approach

It is important to explicitly acknowledge some methodological limitations of the analysis.

- Identification. The Difference-in-Differences approach relies on the parallel trends assumption, which is not directly testable. It is only possible to verify that in the pre-intervention period no systematic differential trends emerge between the treated and control groups. The causal credibility of the estimates therefore depends on the plausibility of this assumption and on the supporting verification provided in the analysis.
- Control group definition. The distinction between "more exposed" and "less exposed" units (e.g., professional vs. non-professional) is constructed on observable characteristics and on an operational definition. These proxies may not perfectly capture the actual heterogeneity in sensitivity to regulation, nor any unobserved differences that vary over time.
- Time horizon. The analysis focuses primarily on short-term effects around the policy introduction. Long-term effects, such as organisational adjustments,

activity reorganisations, or structural market changes, may manifest beyond the observed window and therefore not be fully captured.

- Selection in price outcomes. Price measures, especially when based on observed transactions, may be subject to selection. Changes in activity or bookings can modify the composition of the sample for which price is observable, requiring interpretive caution.

Despite these limitations, the adopted approach represents an established standard in the empirical policy evaluation literature and allows production of informative evidence on the effects of regulation, provided it is interpreted with appropriate methodological cautions and supported by robustness analysis.

1.8 Thesis Structure

The thesis is organized in a sequence of chapters that proceeds from the theoretical and institutional framework to empirical evidence. Chapter 2 offers a critical review of the literature on short-term rentals, with particular attention to the effects of platforms on the real estate market, on the hotel sector, and to the role of regulations in determining market equilibrium. Chapter 3 reconstructs the regulatory context of short-term rentals, including a comparison between US and European approaches and the framing of Italian regulatory evolution, with focus on the Ministry of Interior Circular of November 18, 2024, and its operational implications. Chapter 4 describes the data used and develops a descriptive analysis of the Italian market, documenting price, supply, and booking dynamics, as well as synthetic performance indicators, with the objective of providing preliminary evidence and defining relevant outcomes for causal analysis. Chapter 5 presents the econometric strategy and results, quantifying the impact of regulation on outcomes of interest and discussing heterogeneity and robustness checks. Finally, the conclusions synthesize the main findings and discuss their implications for the literature and for policy.

Chapter 2
Literature Review

This chapter outlines the theoretical and empirical framework upon which the academic debate on short-term rentals takes place, providing a critical interpretation for understanding the effects of the Airbnb platform on contemporary urban dynamics. International literature has analysed the phenomenon from multiple angles, recognising Airbnb as a disruptive innovation (Guttentag, 2015) which, thanks to the extreme flexibility of its business model and the initial regulatory vacuum, has profoundly altered the balance of the property market, the hospitality sector, and the social composition of historic city centres. However, a careful analysis of these issues presents significant methodological challenges, first and foremost the extreme heterogeneity of the operating contexts. As highlighted by Nieuwland and van Melik (2020) and Guttentag (2015), the problem of external validity represents a structural limitation to the generalisation of the conclusions of statistical studies: the dynamics observed may differ from city to city, making it impossible to find a homogeneous solution that is objectively valid; the results obtained in American cities are not comparable with medium-sized tourist cities, let alone European cities, which have a completely different social and urban context.

Furthermore, it is necessary to consider the complexity of the causal links between the supply of short-term accommodation and the attractiveness of the areas. Many studies suffer from potential selection bias, as Airbnb tends to develop more in areas already characterised by strong tourism, making it difficult to isolate the net impact of the platform from pre-existing trends (Barron et al., 2021). Added to this is the problem of reverse causality: while the platform exacerbates overtourism, growing tourist demand further encourages Airbnb's penetration in the area, having a negative effect on hotels (Barone et al., 2019). To overcome these critical issues, the most recent literature has adopted quasi-experimental approaches and instrumental variables aimed at isolating the causal effect of regulations (Garcia-López et al., 2020; Koster et al., 2021; Valentin, 2021).

2.1 The impact of Airbnb on the property market

The academic debate on the relationship between STR platforms and the property market revolves around the economic concept of rent gap. The advent of Airbnb has introduced a new frontier in urban capital appreciation: property owners are incentivised to withdraw units from the long-term rental market and allocate them to tourist rentals, driven by the prospect of significantly higher returns and greater contractual flexibility (Wachsmuth & Weisler, 2018). This shift in residential supply towards tourist accommodation generates a negative shock on the housing stock available to residents, triggering a structural increase in prices.

In terms of causal identification, the study by Barron et al. (2021), conducted across the entire United States, represents the methodological benchmark. Using an instrumental variable approach based on changes in tourist demand, measured using an index created based on Google searches, the authors show that a 10% increase in Airbnb listings translates into an average increase of 0.18% in rental prices and 0.26% in house sale prices. The transmission mechanism identified is the reduction in the residential vacancy rate, which increases competition among local tenants for the remaining housing units. Analysing the New York City market using a spatial equilibrium model, it has been shown that the shift of housing units from the residential to the tourist market generates a net loss of welfare for the city's inhabitants. Specifically, there is a loss of welfare for renters, which translates into a transfer of $2.7bn to property owners (Calder-Wang, 2021).

Evidence gathered in European and international urban contexts confirms the regressive nature of this phenomenon. Garcia-López et al. (2020), analysing the case of Barcelona using a Difference-in-Differences (DiD) model, highlighted how the effect is heterogeneous and spatially concentrated: in neighbourhoods with higher tourist density, the pressure exerted by the platform has increased rents by 1.9%, transaction prices by 4.6% and advertised prices by 3.7%, making housing affordability a critical challenge for the local population. Similarly, the study by Koster et al. (2021) on Los Angeles confirmed, using spatial regression-discontinuity design (RDD) combined with DiD, that regulatory restrictions introduced to limit Airbnb's activity are able to partially correct these distortions, leading to a reduction in property prices in regulated areas.

2.2 The Impact on the Hotel industry

Airbnb's entry into the global market is described as a disruptive innovation that has profoundly altered the competitive balance of traditional hospitality (Lee et al., 2023). The literature documents a clear substitution effect. Specifically, in Texas, it has been documented that a 10% increase in the supply of short-term rentals leads to a 0.39% reduction in quarterly hotel revenues (Zervas et al., 2017). However, this impact is not uniform across different categories of accommodation, as it affects budget establishments and those without businessoriented services to a significantly greater extent, while luxury hotels show greater resilience (Zervas et al., 2017). In the Italian context, this heterogeneity is confirmed by the evidence that the impact of Airbnb is mainly concentrated on hotels with three stars or less, identified as direct substitutes for the platform's offering, unlike higher-end establishments, which are less affected by the competitive pressure of peer-to-peer services (Barone et al., 2019). In terms of overall profitability, it is estimated that average hotel profits would be 3.69% higher in the absence of competition from home-sharing platforms (Farronato & Fradkin, 2022).

A key mechanism through which this erosion manifests itself is the loss of market power during so-called “compression nights”, i.e. periods of peak demand linked to holidays or major events. In these circumstances, Airbnb's elastic supply, expanding rapidly when hotel occupancy is close to saturation, prevents traditional establishments from applying monopoly rates, drastically reducing their seasonal profit margins (Farronato & Fradkin, 2022). The expansion of the short-term rental market also exerts constant downward pressure on key hotel performance indicators, negatively affecting both occupancy rates, ADR (Average Daily Rate) and RevPAL (Revenue Per Active Listing) (Dogru et al., 2020). This dynamic forces hoteliers to revise their pricing strategies downwards in order to maintain occupancy levels, consolidating a transfer of value to consumers who benefit from a greater variety of choice and lower average prices during peak seasons (Farronato & Fradkin, 2022).

2.3 The impact of regulations on the short-term rental market

Economic literature has devoted increasing attention to the effectiveness of public policies in mitigating the negative externalities generated by short-term rental platforms. The introduction of restrictive regulations acts primarily by increasing operating costs and legal risks for landlords, influencing their decision to remain in the market. For example, the registration requirement and technological enforcement of licences in San Francisco caused a 20% to 27% decline in the total supply of listings, while also leading to an increase in average prices for units that remained active due to reduced internal competition (Bibler et al., 2025). At the same time, the exploitation of administrative boundaries in Los Angeles County has highlighted how the introduction of restrictive ordinances leads to a 50% decrease in Airbnb supply in regulated areas compared to neighbouring exempt areas, resulting in a downward correction in house prices of around 2% (Koster et al., 2021). However, the effectiveness of these interventions depends on the monitoring capacity of local authorities: in the case of Barcelona, the freeze on tourist licences was successful in curbing the growth of listings primarily in areas characterized by stronger enforcement, while having limited effects elsewhere, suggesting that formal regulations alone are insufficient to alter market outcomes without effective control mechanisms (Garcia-López et al., 2020).

A central aspect of the regulation concerns the distinction between occasional hosts and commercial operators. The effectiveness of the rules aimed at targeting commercial operators, so-called multi-hosts, is linked to the introduction of restrictions that make professional activity incompatible with the requirements for use of the property. Restrictions aimed at limiting the activity of professional players are not only effective in reducing pressure on the property market, but can paradoxically favour the supply of non-professional hosts, who see a reduction in internal competitive pressure within the platform (Chen et al., 2023). This evidence suggests that targeted regulation, which addresses the operational barriers of multi-hosts, is the most appropriate tool for preserving the core concept on which the sharing economy is based, without fuelling speculative phenomena.

The effectiveness of public intervention tools has also been analysed with respect to the introduction of time limits on rental activity, commonly referred to as day caps. Limiting the number of days of annual operation significantly reduces the supply of accommodation in the short-term rental market and can foster a partial reallocation of housing units toward the long-term rental sector (Gauß et al., 2024). These restrictions generate a redistribution of benefits among stakeholders: while the hotel sector may benefit from reduced competitive pressure, the effects on residents' and hosts' welfare critically depend on the stringency of the imposed cap (Gauß et al., 2024). The analysis further highlights potential trade-offs, as overly stringent limits may increase incentives for non-compliance, whereas well-calibrated day caps emerge as an effective instrument to mitigate the negative externalities of the sharing economy while preserving the local residential fabric (Gauß et al., 2024).

The literature suggests that non-price regulatory requirements, such as mandatory registration schemes, introduce fixed compliance costs that disproportionately induce the exit of less structured hosts. As a result, even when the total number of listings declines, the remaining supply becomes relatively more professionalised, increasing the relative weight of commercial operators (Bibler et al., 2025).

Chapter 3
Short-Term Rental Regulations

The previous chapter analysed the main research conducted in literature on the subject of Airbnb. In this chapter, the focus is on the regulations implemented by legislators in major cities internationally to try to alleviate the negative externalities produced. The advent and rapid spread of the sharing economy in the tourism sector have radically transformed the global hospitality market over the last decade. What began as a phenomenon of occasional sharing of unused spaces has evolved into a professional industry, capable of competing directly with the traditional hotel sector and profoundly changing the socio-economic context of tourist destinations. The main challenge that legislators had to face was to understand the socio-economic context of the city and develop *ad hoc* regulations in order to preserve the positive aspects of the sharing economy. In most cases there was initially a *laissez-faire* approach, so a regulatory gap had to be filled. However, the regulatory response has not been uniform (Nieuwland & van Melik, 2020): different jurisdictions have adopted different strategies, ranging from simple bureaucratic registration to partial or total bans, often in response to specific local tourism pressure.

In order to systematise the varied international case studies that will be presented, it is useful to first adopt a taxonomy of regulatory instruments, as identified in the literature. Regulatory interventions can be divided into four main macro-categories (Jefferson-Jones, 2015), often combined by local administrations in a heterogeneous policy mix. The first category concerns time restrictions (*Day Caps*), which impose a maximum number of nights per year for rentals effectively distinguishing between occasional home sharing and business activities. The second lever is geographical restrictions (*Zoning*), used to prohibit or limit the issuance of new licences in saturated urban areas to relieve tourist pressure in specific parts of the city (Gurran & Phibbs, 2017). A third type, known as the Primary Residence Requirement, is the most severe measure against professionalisation, allowing STR only in the host's main residence. Finally, there are administrative and transparency requirements (*Licensing & Registration*), which subject the activity to obtaining an identification code and complying with safety and tax standards. The latter category aims to increase compliance costs and reduce regulatory asymmetries with the traditional hotel sector (Guttentag, 2015).

3.1 American overview

The US landscape is the starting point for regulatory analysis, as the United States is the birthplace of the sharing economy and, consequently, the first to have addressed its legal challenges. The response of US cities has been characterised by greater heterogeneity and, in some cases, unprecedented regulatory severity. Below, some regulations implemented in major American cities are analysed by way of example.

The analysis of the most relevant case studies begins in California, where San Francisco has pioneered the application of the Primary Residence Requirement. As the birthplace of Airbnb and one of the areas with the most acute housing tension in the country, the administration introduced the "One Host, One Home" principle. This regulation allows STR only in the residential unit in which the host permanently resides for at least 275 days per year and imposes a maximum cap of 90 days per calendar year for un-hosted STR, i.e. rentals in which the owner does not reside. This legal framework is explicitly aimed at making the activity of multi-property investors unviable, with the aim of preserving the residential stock and preventing the removal of housing from the long-term market. In addition to the above, other measures include the Ellis Act, zoning restrictions, and registration requirements.

An even stricter approach, based on a total ban on unattended rentals, has been adopted by Santa Monica. The city's Home-Sharing Ordinance has established a clear and insurmountable distinction between "Home Sharing", which is considered legal, and "Vacation Rental", which is prohibited throughout the municipal area. Specifically, it is strictly prohibited to rent an entire home for periods of less than 30 days, allowing only the rental of a spare room or portion of the primary home. As discussed in the literature by Wachsmuth and Weisler (2018), this measure is designed to radically eliminate the supply of "entire homes" from the tourist market, forcing them back into the residential circuit through proactive enforcement and severe penalties.

A prime example of geographical restrictions, i.e. zoning, is New Orleans. STR of properties, rental for less than 30 days, were illegal when Airbnb was founded. The city, historically subject to intense tourist pressure, adopted a differentiated approach in 2017, mapping the territory into permitted and prohibited zones. The key element of the legislation is a total ban on STR in the French Quarter (Vieux Carré), the city's oldest and most iconic neighbourhood. The administration determined that preserving the architectural integrity and residential fabric of this historic area was incompatible with the commercial nature of STR. Outside this "red zone" regulations vary, limited permits, often tied to the presence of the owner, are allowed in residential areas, while the rules are more permissive in commercial areas. This selective use of zoning aims to confine the impact of tourism, protecting districts of greater cultural value from overtourism and gentrification.

Regulatory developments reached their peak with the case of New York City and the introduction of Local Law 18 in 2023. This regulation represents a turning point in the global landscape, combining the residency requirement with an unprecedented system of platform liability. The law not only reiterates the ban on renting entire flats for less than 30 days but also introduces a financial blocking mechanism at source: digital platforms are prohibited from processing transactions for listings that do not have a registration number verified by the Office of Special Enforcement (OSE). The legislator conceived this shift in the burden of control to overcome the historical difficulties of monitoring by local administrations, making it technically impossible to publish non-compliant listings.

3.2 European overview

Shifting the analysis to the European continent, a regulatory approach that differs from the radicalism of some US bans is highlighted. Although the European Union has recently promoted greater transparency through data sharing (DAC7 Directive[9]), the management of tourist flows has remained firmly in the hands of local administrations. Major European cities have developed intervention models aimed at balancing openness to tourism with the protection of residential areas, mainly using time-based quota systems (*Day Caps*).

The most widespread model in Northern Europe is based on time limits, designed to distinguish occasional home sharing from commercial activity. A prime example is London, where the Deregulation Act of 2015 introduced the “90-day rule”. The rule allows owners to rent out their entire residence for up to 90 nights per year without requiring specific permits; above this threshold, the activity is reclassified as a change of use, requiring planning permission, which local authorities rarely grant in order to preserve the housing stock. The regulation does not apply if a host guests in a room while he is also present.

As regards Berlin, it has structured its regulatory response around the strict protection of housing stock, operating through the Zweckentfremdungsverbot-Gesetz[10]. This legislation, introduced in 2013, is based on the principle that housing should not be diverted from its

[9] DAC7 (Directive (EU) 2021/514) is a European directive that strengthens the tax reporting obligations of digital platforms.

[10] German Law on the Prohibition of Misuse of Housing.

primary social function. Initially configured as a *de facto ban*, the law was subsequently amended to allow for more flexible but controlled regulation. Currently, in order to rent a property on a short-term basis, it is mandatory to obtain a registration number, while the rental of a main residence is permitted under favourable conditions. For second homes, there is a maximum limit of 90 days per year, subject to a specific permit.

Even more complex is the case of Amsterdam, which adopts a hybrid system characterised by a continuous tension between restrictive political will and judicial examination. The general rule imposes a maximum limit of 30 nights per year, accompanied by the obligation to notify the municipality of each individual stay. In an attempt to curb overtourism, in July 2020 the administration introduced a total ban on STR in three districts of the historic centre. However, this measure was deemed excessive by the Council of State (Raad van State), which overturned the ban on the grounds that it was disproportionate to the objective pursued.

Finally, Barcelona represents the most structured example of zoning in Europe. Through the Special Urban Plan for Tourist Accommodation (PEUAT[11]), the city has been divided into concentric zones with different levels of restriction. In the central areas, defined as saturated, a total ban on the issuance of new licences has been imposed, allowing activity only in peripheral areas and under strict conditions (Ajuntament de Barcelona, 2017). Recently, the city's policy has shifted towards an even more restrictive approach, with the announcement of its intention not to renew existing licences upon their expiration in 2028, potentially paving the way for the future elimination of residential tourist accommodation, in line with restrictive approaches observed in some American cities.

3.3 Italian overview

Regulatory developments in Italy have followed a path of progressive centralisation. While the matter was initially fragmented by various local regulations, state intervention has gradually become more incisive in order to respond to three fundamental needs: tax compliance, public safety, and the distinction between occasional and entrepreneurial activities.

In the first phase of expansion of the phenomenon (2008–2016), in the absence of a national framework law, responsibility fell almost exclusively on the regions, which legislated to distinguish tourist rentals from traditional accommodation facilities. However, this fragmentation led to asymmetries and difficulties in monitoring. The turning point came in the

[11] Barcelona’s urban zoning plan regulating tourist accommodation.

two-year period 2017-2019, with a focus on tax revenue recovery and security. Decree Law (D.L.) 50/2017 introduced a 21% flat tax regime (*cedolare secca*) specifically for STR, establishing the role of platforms as withholding agents, required to operate tax withholding at source. At the same time, from a security perspective, D.L. 113/2018 definitively clarified the applicability of Article 109 of the TULPS[12] to private individuals, making it mandatory to communicate guests' personal details to the police headquarters via the "Alloggiati Web" portal within 24 hours of arrival. A first attempt at administrative order was made with D.L. 34/2019 (*Decreto crescita*), which legitimised the dissemination of Regional Identification Codes (CIR, IUN, CIS), anticipating the need for a national database. The current regulatory framework was outlined by D.L. 145/2023, converted into Law 191/2023, and by the subsequent 2024 Budget Law. This comprehensive reform is based on three main pillars: the establishment of the National Identification Code (CIN) linked to the National Database of Accommodation Facilities (BDSR) to combat the underground economy, the introduction of new plant safety requirements that require the installation of gas detectors and portable fire extinguishers, and an increase in the tax burden, with the flat-rate tax rate rising to 26% from the second rented property onwards.

However, the regulation covered by this study is the circular issued by the Ministry of the Interior on 18 November 2024, which provides clarification on the procedure for remote identification of guests staying in short-term accommodation facilities. The circular refers to the obligations set out in Article 109 of the TULPS, establishing the obligation to identify guests in person and effectively abolishing the possibility of remote identification. The main rationale for introducing this measure lies in the need to reduce the risks associated with the illegal occupation of STR by potentially dangerous individuals, particularly during major political and cultural events. Although the circular is primarily aimed at public safety, the potential economic effects of its application cannot be ruled out. Before the circular was issued, remote check-in was permitted, without the need for the host to be physically present at the property. This aspect is crucial to the analysis conducted in this study, as the new regulations appear to have a significant impact on multi-hosts, who, managing multiple properties, may find it impossible to visit each unit simultaneously to identify guests. Another category affected is that of individuals who make their entire property available, as the

[12] TULPS (Testo Unico delle Leggi di Pubblica Sicurezza) is the Italian consolidated law governing public order and public security.

property may be located far from the host's actual residence, making it impractical to comply with the circular.

Operationally, the circular has asymmetrical effects depending on how check-ins were managed prior to its introduction. For hosts who already conducted check-ins in person, it does not materially alter existing practices or generate additional costs. For hosts relying on self-access procedures such as key boxes, electronic locks, or remote instructions, it instead requires a reorganisation of the arrival process. The practical implications then vary sharply with scale. A non-professional host managing a single property near their residence can typically comply with limited additional time, whereas a professional operator managing many properties across various locations faces a scalability constraint: physical presence cannot be replicated across simultaneous arrivals, making external staff necessary. This implies additional operating costs that increase with the number of check-ins. Importantly, these costs are predominantly variable, proportional to listing activity rather than to portfolio size, which distinguishes the circular from measures such as mandatory registration that generate largely fixed costs per property. The requirement also complicates late-night arrivals, further increasing reliance on third parties such as greeters. Hosts may respond by passing these costs on to guests through higher platform prices, absorbing them through lower marginal revenue, or, if they become excessive, reallocating supply toward the long-term rental market. Consequently, although introduced under a public-safety rationale, the circular can also be interpreted as indirectly curbing professionalised STR operations and shifting the market toward a model closer to the original sharing-economy logic.

Chapter 4
Data and Descriptive Analysis

To analyse the impact that regulation has had on the Airbnb market, a dataset containing 396,364,643 rows provided by AirDNA was used. The data provided refers to the entire Italian market and is on a daily basis with a reference period from 01-01-2023 to 31-08-2025. The dataset contains a number of variables that are useful for the purposes of the analysis. A brief description of these variables is provided below:

- *Property ID*: unique code identifying each property in the dataset.
- *Date*: daily reference indicating the date of the stay.
- *Status*: variable indicating whether the property is reserved (Status = "R") or available (Status = "A") on that day, i.e. not booked on that specific day. There is also a value indicating whether the property has been removed by the host on that day (Status = "U")[13].
- *Price (USD)*: price associated with the transaction in dollars.
- *Booked Date*: date on which the booking is made on the platform.
- *Reservation ID*: unique code associated with a booking.

4.1 Descriptive statistics

Table 1 shows a series of statistics aimed at providing an overall description of the Airbnb market in Italy. These statistics are compiled by aggregating data from January 2023 to August 2025 on a monthly basis. These results outline the main characteristics of the market in terms of prices, size, and performance, providing an overview of the short-term rental sector before moving on to empirical analysis.

Panel A focuses on prices. The average monthly price per transaction is $171, with a median value of $168, suggesting a relatively symmetrical distribution of prices. The standard deviation of approximately $22 highlights a significant variability in prices over time, consistent with what is expected from the market, which is strongly affected by seasonality and variations in the composition of demand. The minimum and maximum values, $133 and $214 respectively, confirm the market's tendency to fluctuate. A similar picture emerges

[13] This analysis has excluded rows with this value as they represent a negligible number of the entire dataset, 0.015% of the total.

when looking at the monthly median price, which is lower than the average price in terms of both average levels and variability, indicating that the latter is influenced by high-end listings.

Panel B focuses on market size, specifically analysing three variables: active listings, the average monthly number of listings on the platform, regardless of whether they are available or booked on a given day; available listings, the average monthly number of listings actively available for booking, excluding those already booked; and booked days, the total monthly number of days booked, which measures the overall volume of transactions actually carried out in the market. On average, there are approximately 394,000 distinct listings active each month, regardless of whether they are booked or not, while unbooked listings alone number approximately 260,000 per month. The gap suggests that a sizeable share of active listings is booked for at least part of the month. The wide range between the minimum and maximum values for active listings, between 180,000 and 490,000, indicates significant variability in the number of monthly listings within the period observed. The number of days booked averages around 4 million per month and, like the other two variables, shows high dispersion, reinforcing the idea that the short-term rental market is strongly influenced by the seasonality of tourist demand and possible structural changes in supply.

Panel C presents some summary performance indicators for assessing market performance. The average occupancy rate, calculated as the total number of days booked in the month divided by the total number of active listing days per month, is 0.32; this indicates that a listing is booked for approximately one-third of the days it is active in a month. Both the range between the minimum and maximum values and the standard deviation are high, highlighting strong seasonal differences in the intensity of listing usage. Consistently, the average number of days booked per active listing per month is approximately 10 days, suggesting an average usage of about one-third of the month. Finally, the analysis examines the Revenue proxy Per Active Listing (RevPAL), defined as the product of the average transaction price and the average number of days booked, which shows an average value of approximately $1,750 per month. Although this does not represent actual revenue at the individual listing level, this indicator effectively summarises the interaction between prices and demand intensity, also highlighting considerable heterogeneity over time.

Table 1. Summary statistics

Panel A. Prices (USD)					
Variable	**Mean**	**Median**	**SD**	**Min**	**Max**
Monthly mean transacted price	171.36	168.19	21.6	133.12	214.22
Monthly median transacted price	118.27	114	14.07	96	146
Panel B. Market size (monthly)					
Variable	**Mean**	**Median**	**SD**	**Min**	**Max**
Active listings (avg per month, A+R)	394,373	398,071	45,386	182,166	492,139
Available listings (avg per month, A)	263,434	270,384	44,001	148,047	321,415
Booked days (monthly, R)	3,997,183	3,987,280	1,740,414	1,023,572	7,803,955
Panel C. Market performance (monthly)					
Variable	**Mean**	**Median**	**SD**	**Min**	**Max**
Occupancy rate (booked/active listing-days)	0.323	0.331	0.12	0.173	0.552
Booked days per active listing (avg per month)	9.85	10.09	3.74	5.36	17.11
RevPAL (USD per month)	1,752.96	1,670.43	855.08	752.05	3,549.00

In addition, some supporting graphical analyses were conducted; the aim of these is to assess, at a preliminary stage, possible recurring trends or anomalous behaviour following regulation. The focus in this assessment phase was mainly on a number of key aspects, namely: analysis of price dynamics, market structure, booking behaviour and, finally, market performance.

Price dynamics are analysed using representations that allow us to evaluate the trend of prices observed on the platform, as well as those actually traded.

Figure 2 shows the average price trend over time, calculated based on all listings observed (A+R) and smoothed using a 7-day moving average, with the aim of mitigating the volatility that can be found in weekly trends, such as higher prices at the weekend. Consistently with the preceding evidence, the series confirms the strong seasonality of the sector, with peaks in the summer and contractions in the winter but also shows a gradual increase in price levels over the period under analysis. Close to the regulation, marked by the vertical red line, it can be seen that the contraction in the winter period appears to be more dramatic than in the previous year, with January 2025 settling at lower average prices than the same month of the previous year, but from the following months onwards, the rise to the summer peak is steeper, reaching a peak of over $220, which represents the all-time high for the market.

Figure 2. Average price trend

The second analysis conducted examines the issue of average prices in greater depth, as shown in Figure 3, focusing exclusively on listings that are actually booked (R) and comparing the average price with the median price. When evaluated together with the previous analysis, this representation allows us to isolate the behaviour of the prices actually paid and assess the shape of their distribution. Throughout the period analysed, there is a systematic deviation between the average and median lines, highlighting the presence of an asymmetric distribution with a tail towards listings with higher prices. After regulation, both measures show a growing trend, with the average growing at a relatively higher rate than the median, suggesting that the higher price ranges may more drive the increase in prices.

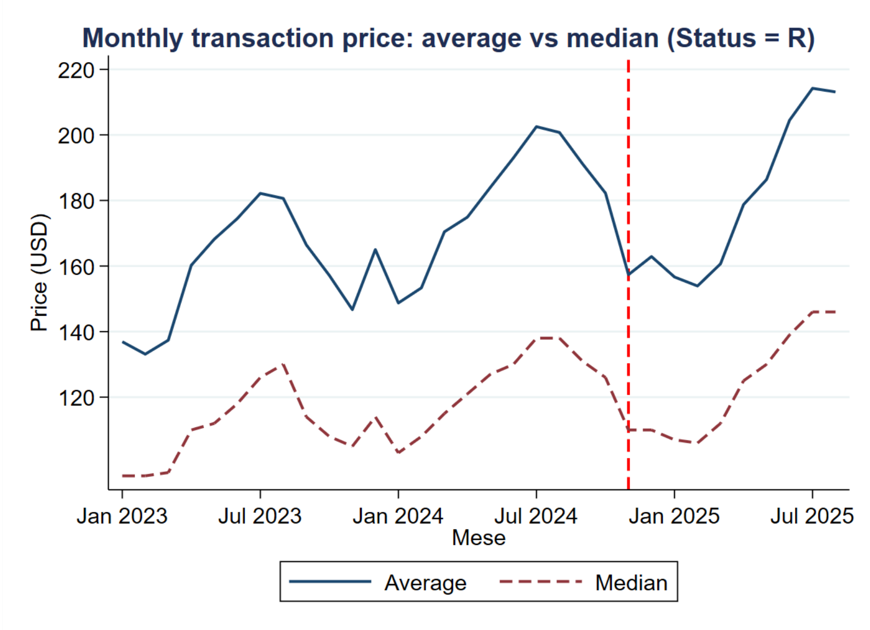

Figure 3. Transaction price

The joint analysis of the dynamics, shown in Figure 4, reveals a complex picture of the post-regulatory price equilibrium. In the winter period immediately following the reform (January 2025), the average transacted price reached a seasonal low of around $155, which is broadly in line with the approximately $150 recorded in the same period of the previous year.

However, a key pattern emerges in the evolution of the spread relative to the listing price. While in 2024 the gap between the asking price and the paid price was more contained, in the post-regulatory period (vertical dotted line) there is a wider divergence. With the listing price showing strong rigidity, remaining above $180 even in the low season, demand did not validate these levels, with transactions finalised at significantly lower prices. The widening of this "gap" is consistent with a sorting pattern, with consumers systematically allocating themselves to the most competitive units, leaving unsold listings with excessive economic demands, which in fact raise the average listing price but not the transacted price. This gap tends to be reabsorbed only with the arrival of the summer demand shock when supply saturation forces transaction prices to align with list prices.

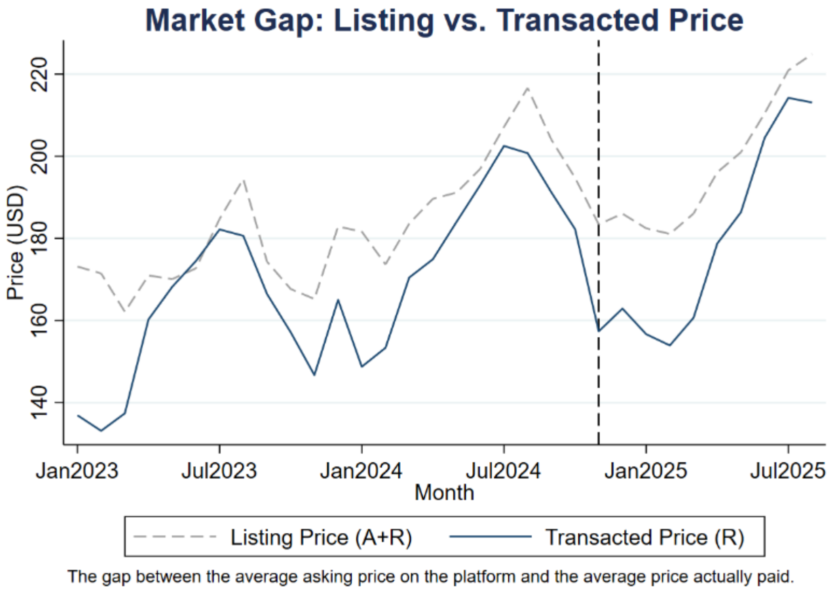

Figure 4. Asking Price vs. Realized Price

Moving on, the analysis examines the structure of the market and its evolution with the advent of regulation, focusing on both its size and its internal composition.

Figure 5 suggests that, after a period of sustained growth between 2023 and the first half of 2024, supply peaks in the summer season of the same year. As the regulation comes into force, there is a net reduction in the number of listings, which continues until the summer months of 2025, when there is a slight recovery, but one that does not reach the peaks seen in the previous year, suggesting the emergence of a new equilibrium.

Figure 5. Total number of listings

In order to examine this change in greater depth, it was decided to compare active listings (A+R) with those actually booked (R), measured on the basis of monthly listing days. Figure 6 confirms that following the regulation there has been a reduction in active supply; however, the line representing booked listings shows a less pronounced decline and a faster recovery. This indicates that, despite there being fewer active listings, demand continues to be similar to previous years, suggesting an increase in the intensity of use of active listings.

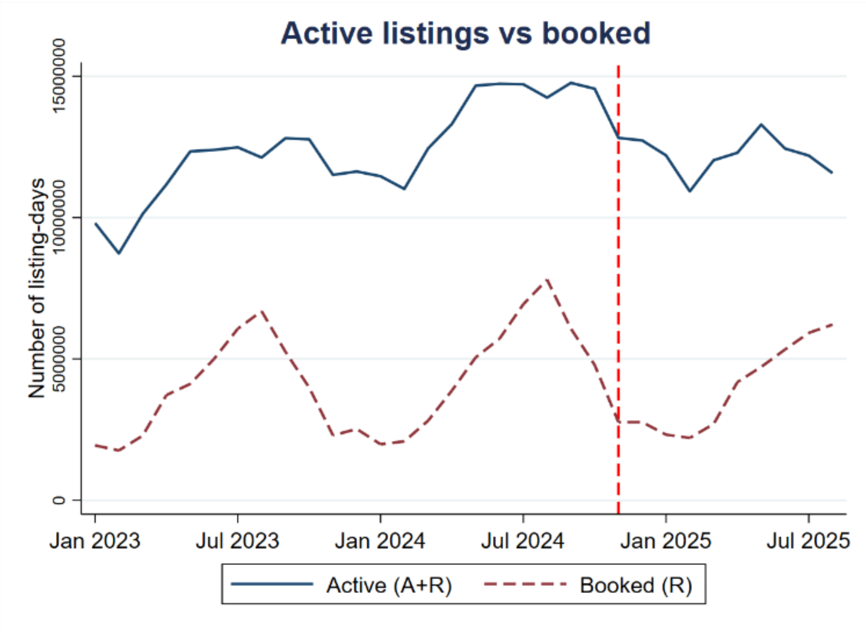

Figure 6. Supply activity

Next, the composition of the supply is analysed, distinguishing between price ranges (budget, mid-range and luxury). In order to assign each listing to its respective category, the distribution of historical prices in the pre-regulation period (January 2023 – November 2024) was analysed. This classification was kept fixed in the post-reform period in order to observe the number of listings that remained on the market, purging the analysis of any reclassification effects due to price increases or the entry of new operators. Figure 7 suggests that the mid-range segment is the most prevalent on the market and is also the one that has suffered the greatest decline in the number of listings on the market since regulation. Analysis of the other two segments, budget and luxury, suggests that the former remained stable in terms of the number of active listings, while the latter seems to follow the trend of the intermediate segment, with a sharp decline around the time of the reform and a slight hint of recovery in the summer season, due in part to the seasonality of the market.

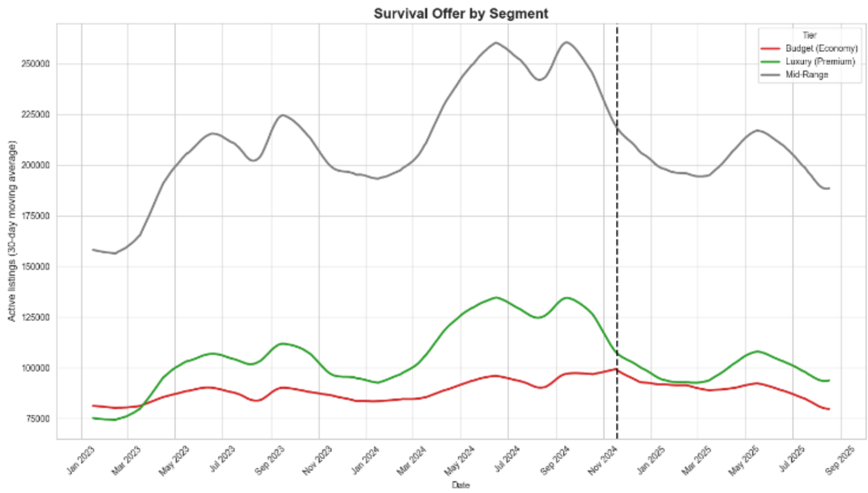

Figure 7. Active supply by segment

Finally, the issue of turnover was examined in depth in a structural manner, distinguishing permanent changes from simple seasonal breaks. For each PropertyID, a birth (entry) was defined as the first appearance in the dataset, and a death (exit) as the last observed appearance. This approach makes it possible to distinguish periods of temporary inactivity from definitive exits from the market.

Both graphs should be analysed with caution, as the approach is influenced by the edges of the sample, which affect the two measures asymmetrically. Regarding births, the first few months appearing in the dataset show very high numbers because the PropertyID appears for the first time in the dataset, simply because no previous periods are observed. Conversely, deaths are concentrated in the last months available, as the algorithm identifies the last month available in which the listing is observed as the exit. To mitigate statistical artefacts that would overestimate exits in the latest period, the analysis of market exits was conducted by truncating the observed series and excluding the last two months of data.

Despite this technical limitation, if the initial and final tails are excluded, it can be seen an upward trend for deaths and a downward trend for births around the time of regulation. Overall, this descriptive evidence provides useful context, but requires further investigation in the future, ideally with larger time series datasets, in order to assess more reliably whether, and to what extent, the changes observed are real and persistent over time.

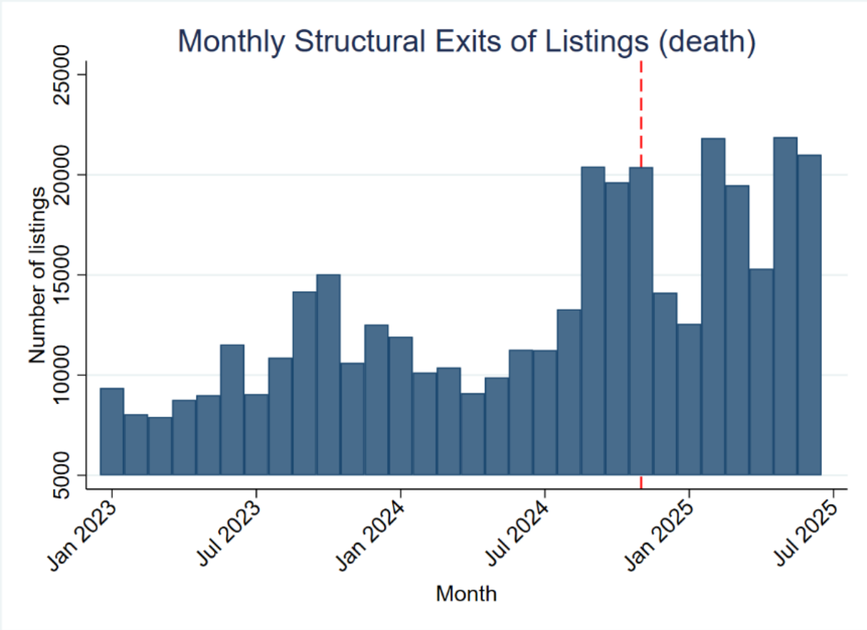

Figure 8. Listing deaths

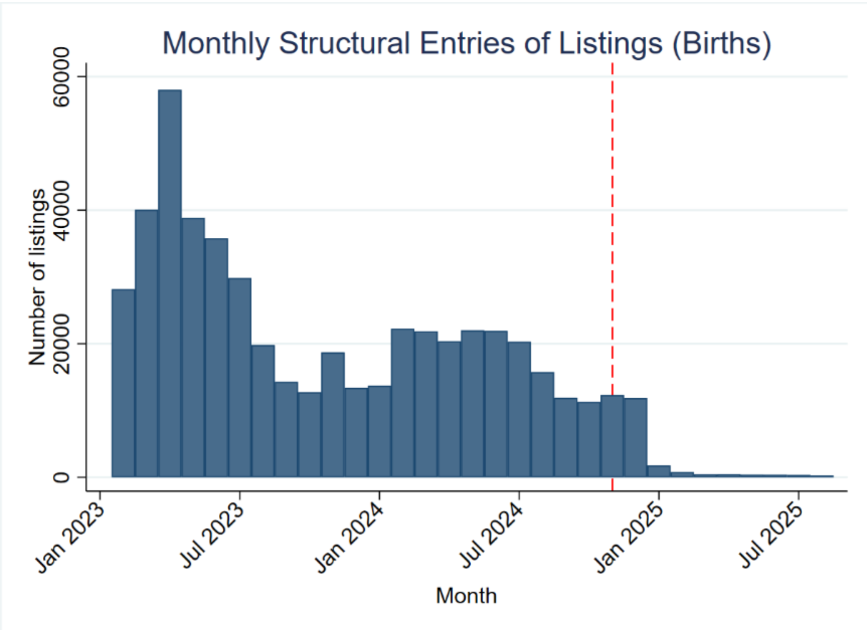

Figure 9. Listing births

Following with the descriptive analysis, after examining price dynamics and supply structure, booking behaviour is evaluated with the aim of assessing how it evolves both in terms of intensity of use of available capacity and in terms of booking times.

The monthly occupancy rate was calculated, on a monthly basis, as the proportion of days booked out of the total number of active days:

$$Occupancy_m = \frac{R_m}{A_m + R_m}$$

Where R_m indicates booked listing days and $A_m + R_m$ indicates active listing days. Figure 10 shows a clear and recurring seasonal pattern: the occupancy rate grows progressively from spring until it peaks in the summer months (reaching between 50-55%), then declines in the following months towards winter lows.

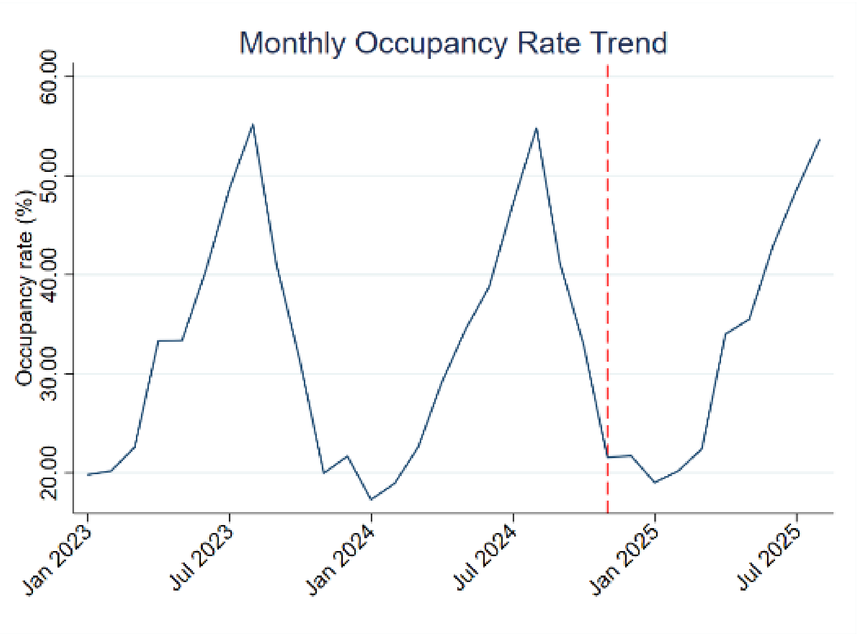

Figure 10. Monthly occupancy rate

However, joint analysis with supply volumes (Figure 6) highlights a structural change in the post-regulatory period. Despite the sharp decline in active listings, the occupancy rate did not simply replicate past trends but recorded higher capacity absorption in seasonal lows. Specifically, the seasonal lows in winter 2025 were higher than in the previous year, above 20% compared to less than 18% in 2024, avoiding the sharp decline typical of the past. This increase in occupancy, mathematically resulting from a more marked decline in supply than in demand, confirms that the market has become tighter; tourist demand, which appears broadly comparable year-on-year in the same months, has concentrated on the remaining housing stock, saturating its capacity to a greater extent than in the past.

In addition to the intensity of use, the timing of bookings is analysed through the booking lead time, calculated as the distance between the day of stay and the day of booking, for reserved listings only.

$$LT_i = Date_i - BookedDate_i; \quad \overline{LT}_m = \frac{1}{N_m} \sum_{i \in R_m} (Date_i - BookedDate_i)$$

Where $Date_i$ is the date of stay; $BookedDate_i$ is the day on which the booking was made; N_m is the number of reserved observations in month m and R_m is the set of observations with Status = "R" that fall in month m.

Here too, as shown in Figure 11, the trend shows clear seasonality, with lead time increasing in the high season and decreasing in subsequent periods. As with the other analyses, regulation falls at a time when the measure is typically lower, while in the following months there is an increase emphasised by the approach of the summer season. There are no obvious discontinuities that interrupt the cyclicality.

In summary, the indicators analysed paint a consistent picture in which occupancy should be assessed as a measure of capacity absorption rather than demand, while the booking lead time shows the profile of how demand is organised. Booking behaviour appears to be mainly driven by the seasonal cycle, with no obvious break around the regulation date in the lead-time series.

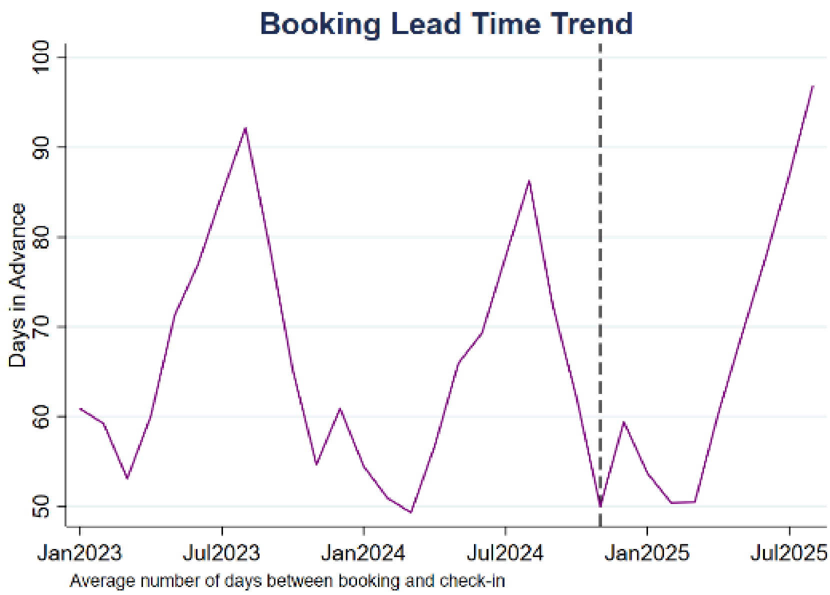

Figure 11. Booking lead time

Continuing with the descriptive, a synthetic market performance indicator is calculated from the hosts' perspective, which is based on a measure of average potential revenue per active listing. Specifically, the indicator is constructed as a combination of (i) a quantity indicator, i.e. booked days per active listing, and (ii) a price indicator, i.e. the average transaction price calculated on the basis of booked days. Operationally, it is defined as follows:

$$RevPAL_m = \{P\}_m^{(R)} \times \left(\frac{R_m}{Active\ listings_m}\right)$$

where $\{P\}_m^{(R)}$ is the average price on days booked in month m, R_m is the number of booked listing days, and Active $listings_m$ is the average monthly number of active listings. This measure allows us to evaluate both price levels and intensity of use.

As with the other analyses conducted, strong, wide, and regular fluctuations are confirmed; there are sudden increases in the middle months of the year, in the summer months, followed by sharp declines in the winter months, when demand is weaker. At the end of the series, in the summer of 2025, there is a new strong increase, consistent with what has been observed in previous years. In the months when demand is highest, a strong multiplier effect is observed, as both the intensity of use and the average transaction price increase, as also observed in previous analyses, confirming the strong seasonality on which the short-term rental market depends.

The vertical line associated with the regulation, as noted above, falls within a low-demand winter period, but there is no immediate break in the series around the regulation date when comparing the same season across years. In descriptive terms, the post-regulatory path appears broadly consistent with past seasonal dynamics, although it unfolds in a context of changing supply volumes. The graph suggests that the combination of booking intensity and price levels are the main drivers of the measure, with stronger dynamics in peak months and weaker dynamics in the low season.

Looking at the previous analyses, the two measures that make up the indicator separately can be evaluated and interpret the trend more clearly. In the post-regulatory period, both measures contribute to the peak in the summer season of 2025, with usage intensity increasing as there is a reduction in active listings while booked listing-days remain broadly comparable year-on-year in the same months. Meanwhile, for the price component, the trend in average prices transacted on booked listings alone is increasing in the same time segment, further strengthening the growth of the RevPAL.

In summary, this measure provides an overview of the average monetisation capacity of the supply, which is useful for assessing the overall behaviour of the sector when the quantities sold and the price associated with those sales vary simultaneously.

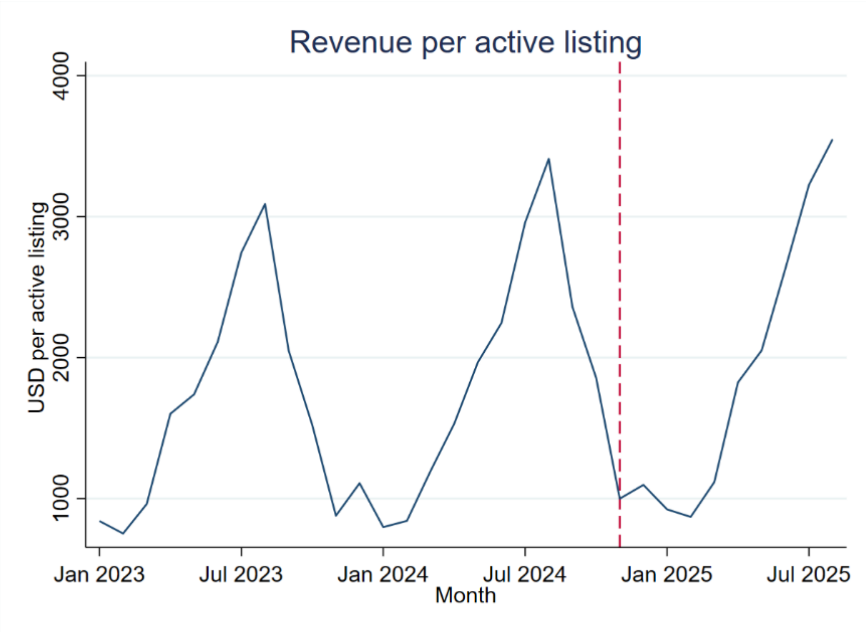

Figure 12. Revenue per active listing

In summary, the descriptive analysis outlines a market undergoing structural adjustment in conjunction with the introduction of regulation. Descriptive evidence suggests that the negative shock to supply, manifested through a sharp contraction in active listings, interacted with demand that appears resilient in year-on-year seasonal comparisons. This imbalance generated an increase in market tightness, observable both in the improvement in occupancy rates during seasonal lows and in the increase in average monetisation capacity (RevPAL).

However, the price adjustment was not uniform: the emergence of a significant spread between list prices and transaction prices indicates that, although the market moved towards higher values, demand reacted by activating self-selection mechanisms towards the most competitive units, not fully validating the hosts' expectations of price increases.

Although the timing of these patterns overlaps with the introduction of the regulation, descriptive analysis does not allow the causal effect of the policy to be isolated from other concomitant factors, such as macroeconomic cyclicality or industry trends. To overcome these limitations and quantify the net impact of the regulation, the following chapter adopts an econometric approach to assess the causal effect of the policy on the short-term rental market.

Chapter 5
Methodology and Empirical Analysis

This chapter illustrates the empirical strategy adopted to assess the consequences of regulation on the behaviour of professional hosts. Section 5.1 defines the research question; Section 5.2 describes the data and the construction of the estimation variables; Section 5.3 provides a descriptive analysis of the evolution of outcomes in the two groups in the pre-policy period, with both an informative and diagnostic function. Section 5.4 develops the identification strategy based on Difference-in-Differences (DiD) with Two-Way Fixed Effects (TWFE). The main results are presented in Section 5.5; Section 5.6 presents the event study as a tool for assessing parallel trends prior to the policy. Sections 5.7 and 5.8 collect the robustness tests and the analysis of economic mechanisms. Section 5.9 develops the heterogeneity extensions.

5.1 Research question and evaluation design

5.1.1 Research question

The central research question is to verify the extent to which the policy has changed the intensity and structure of the supply of hosts classified as professional, compared to that of non-professional hosts, in the post implementation period. The answer is articulated through four outcomes. The main outcome is the share of calendar activity, called *active share*, understood as a direct measure of the supply response. The occupancy rate, referred to as occupancy, serves as a complementary outcome and captures the dynamics of demand realized downstream of the available supply. The logarithm of the list price, referred to as asking price, and the logarithm of the realized price, referred to as realized price, are treated as secondary outcomes, subject to specific interpretative cautions illustrated in the dedicated sections.

5.1.2 Unit of analysis, time horizon, and panel structure

The basic unit of analysis is the listing–month, defined by the combination of a listing's unique identifier (*pid*) and the corresponding calendar month (*mdate*). The panel covers the period from January 2023 to August 2025, for a total of thirty-two months, with the policy month set at December 2024, as the policy came into force in the second half of November.

The structure of the panel is unbalanced: listings enter and leave the sample over time for reasons related to actual market dynamics and data availability. No minimum requirements for presence in the panel have been imposed, with the aim of preserving the representativeness of the sample with respect to the entire distribution of active listings. Table 2 reports the general sample statistics and Table 3 presents descriptive statistics for all outcome variables.

Table 2. General sample statistics

Observations	14,997,652
Unique listings	890,383
Time period	January 2023 — August 2025
Professional classified listings	74 %
Entire classified listings	81.6%

Table 3. Descriptive statistics of the main variables

Variable	**N**	**Mean**	**SD**
active share	14,997,652	0.868	0.247
occupancy	14,996,984	0.352	0.377
asking price	13,498,376	4.805	0.785
realized price	9,273,839	4.871	0.679
active days	14,997,652	26.42	7.54
available days	14,997,652	17.63	11.65
booked days	14,997,652	8.80	10.06

The TWFE model can also be estimated on unbalanced panels. However, identification requires that entries and exits from the sample do not generate treatment-related differential selection. If, in the post-policy period, professional listings exit the sample to a different extent than non-professional ones for reasons related to the policy itself, the DiD coefficient may partly reflect this selection. For this reason, the robustness tests on balanced windows also serve to assess the sensitivity of the results to the composition of the sample. Standard errors are clustered at the listing level in all specifications to account for within-listing serial correlation over time.

5.1.3 Definition of the treatment and logic of the DiD design

The treatment is defined based on the classification of hosts as professional, for which $treat_i = 1$ corresponds to $professional_i = 1$, or non-professional, for which $treat_i = 0$ corresponds to $professional_i = 0$. The temporal discontinuity is set at December 2024, so $post_t = 1$ if *mdate* is equal to or later than December 2024. The parameter of interest in DiD is the coefficient β on the interaction $treat_i \, x \, post_t$, which captures the differential variation in the outcome for professionals compared to non-professionals more than the common temporal variation. The design is non-staggered, as all professional listings share the same treatment date, which simplifies the analysis compared to cases of staggered adoption. In the latter, recent literature has documented problems of interpreting aggregate DiD in the presence of heterogeneous effects across units (Callaway & Sant'Anna, 2021; de Chaisemartin & D'Haultfoeuille, 2020).

It is essential to specify that the treatment is differential and not absolute. The control group is not necessarily immune to the policy, but it is assumed to be affected to a different extent or through different channels than professionals. In the presence of market spillovers, for example if the reduction in professional supply alters the competitive structure to the advantage of non-professional hosts, the DiD coefficient measures the differential gap between the two groups and not the absolute effect on the treated group.

5.1.4 Stratification variables: professional, entire, and host type

The *professional* variable classifies hosts as professional, to which the value 1 is assigned, and non-professional, to which the value 0 is assigned. To assign this variable, the existing *host type* variable is used, which differentiates hosts by the number of properties they own; hosts managing two or more properties are classified as professional. This is the primary criterion for defining the treated group. The underlying rationale is that hosts managing multiple properties are more likely to face larger negative effects from the policy, as they may need to handle several check-ins simultaneously and therefore rely more heavily on remote or automated self check-in procedures.

The *entire* variable indicates whether the listing is for an entire accommodation, in which case it takes the value 1, or a private or shared room, in which case it takes the value 0. It is used as the third dimension of stratification in the Triple Difference (DDD) analysis developed. The economic rationale lies in the fact that entire apartments and rooms occupy distinct market segments with different cost structures and demand profiles, so it is plausible that the policy will affect professional operators differently depending on the composition of their portfolio. In particular, hosts offering private rooms or shared accommodations are less likely to be negatively affected by self check-in restrictions, as they are more likely to be physically present on the premises and can facilitate in-person check-ins, whereas operators managing entire-home listings rely more heavily on remote access and automated check-in procedures.

5.2 Data, variables, and construction of the estimation sample

5.2.1 Structure of the dataset and panel construction pipeline

The estimation dataset is a panel of listings per month constructed from the platform's activity data. Each observation corresponds to a pair consisting of the unique identifier of the listing and the reference month. Time-invariant variables, including professional, entire, and host size category, are associated with each listing for the entire duration of the panel.

5.2.2 Operational definition of outcomes

Table 4 shows the operational definition of each outcome, the expected range, and the number of valid observations in the M4 sample. The divergence in the samples between outcomes is not random but reflects structurally different cleaning rules: *active share* and occupancy are excluded when the respective denominator is zero; asking price is excluded when is zero or missing; realized price is excluded when is zero, missing, or when the listing is not booked. The latter rule yields the smallest and most selectively composed sample, systematically determined by the characteristics of listings that generate bookings.

Table 4. Operational definitions of outcomes and sample sizes in the M4 model

Outcome	Operational definition	Range	N — M4 model
active share	active days / days_in_period	[0,1]	14,997,652
occupancy	booked days / available days	[0,1]	14,996,984

asking price	log of average monthly list price	log(price > 0)	13,498,376
realized price	log of the average monthly realized price	log(price > 0)	9,273,839

5.2.3 Mechanism variables and selection indicator

For the interpretation of the economic mechanisms developed in Section 5.8, three variables relating to the structure of the monthly calendar are used. The variable *active days* counts the days of the month on which the listing is active on the platform. The variable *available days* counts the days made available and not subject to booking. The variable *booked days* counts the days booked. There is a relationship between these three variables, whereby *active days* is given by the sum of *available days* and *booked days*. This structure allows us to distinguish between supply-side adjustments and changes in actual demand.

5.3 Preliminary descriptive evidence

5.3.1 Pre-trends by group: structure and dynamics

Before proceeding with econometric estimates, it is appropriate to examine the evolution of the main outcomes for the four groups defined by the intersection between the professional variable and the entire variable in the pre-policy period. For this purpose, reference is made to the following figures: Figure 13 for occupancy; Figure 14 for active share; Figure 15 for asking price; Figure 16 for realized price. This is purely descriptive evidence, which does not identify causal effects.

The occupancy graph shows marked seasonality common to all groups, with significantly higher summer peaks for professionals with entire accommodations and lower winter lows for non-professionals without entire accommodations. The seasonal structure appears visually similar across groups, which is descriptively consistent with the parallel trends assumption for this outcome, although the amplitudes of the fluctuations are heterogeneous. The graph relating to the number of listings reveals a highly asymmetrical composition of the sample: the group composed of professionals with full accommodation is by far the largest, followed by the non-pro/entire group and the pro/non-entire group in descending order. The non-pro/non-entire group is small, which makes the estimates for this subgroup less accurate.

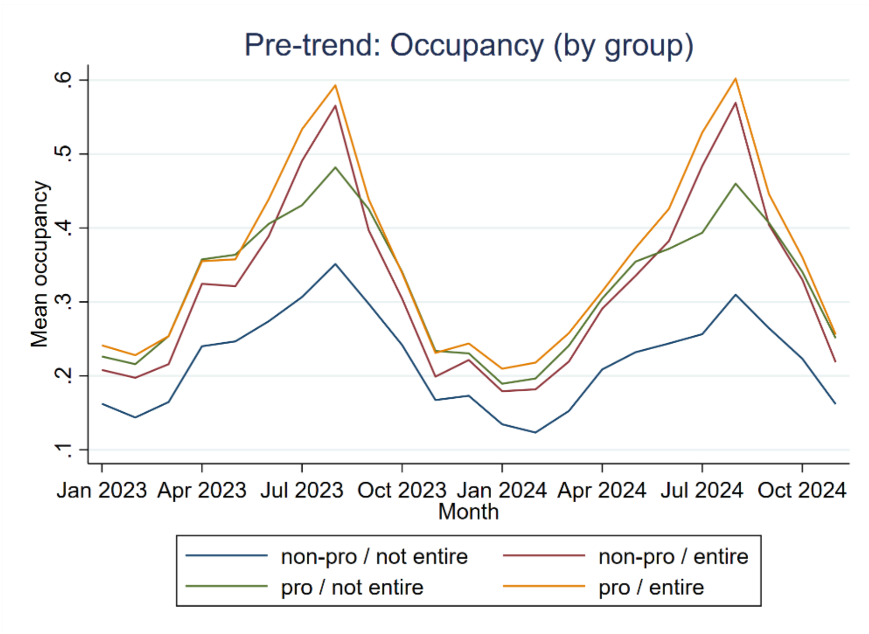

Figure 13. Occupancy by group

The active share pre trend plot suggests that the main source of heterogeneity is the entire dimension rather than professional status. Both not entire groups display high and relatively stable active share levels, with only minor fluctuations and a small gap between professionals and non-professionals. In contrast, both entire home groups show lower average active share and pronounced seasonality, with sharp summer troughs visible in 2023 and especially in 2024.

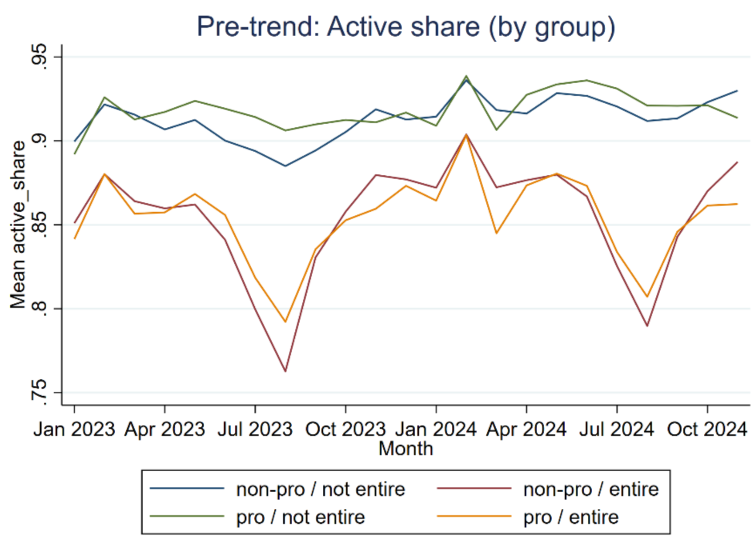

Figure 14. Active share by group

This pattern is consistent with greater intermittency among entire home listings, potentially related to owners' personal use or discretionary availability, while differences between professionals and non-professionals within the same accommodation type appear comparatively limited in the pre-period.

Log asking and log realized price figures show a stable ranking across the four groups over the entire sample period. Entire home listings are systematically priced above not entire listings, and professional entire homes exhibit the highest values throughout. Pre-policy dynamics are not fully reassuring, since the professional entire home group displays more pronounced seasonal peaks and a relatively stronger upward movement in the second half of 2024, most clearly for asking prices. This pattern suggests that parallel trends may be weaker for price outcomes and that DiD estimates for prices should be interpreted cautiously, particularly for realized prices given their selective availability.

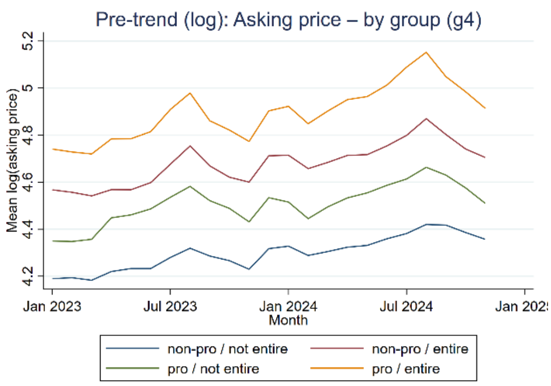

Figure 15. Asking price by group

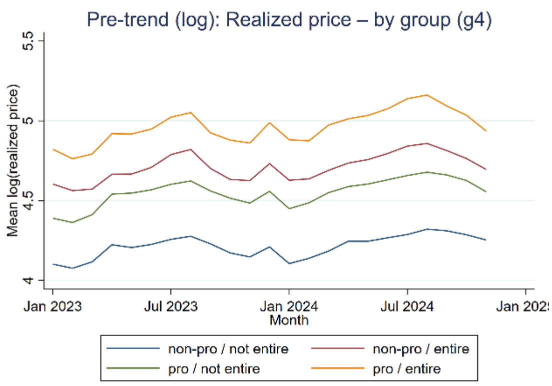

Figure 16. Realized price by group

5.4 Identification strategy: Difference-in-Differences with Two-Way Fixed Effects

5.4.1 Main econometric specification

The main specification adopted for each outcome is the Two-Way Fixed Effects model, referred to below as TWFE:

$$\boldsymbol{Y}_{\boldsymbol{it}} = \alpha_i + \delta_t + \beta(Pro_i \times Post_t) + \varepsilon_{it}$$

In this equation, Y_{it} denotes the value of the outcome for the listing$_i$ in the month$_t$. The term α_i represents the listing fixed effects (AirbnbFE), operationally implemented by absorbing the unit identifier: these fixed effects absorb all the time-invariant characteristics of the listing/host and, consequently, also include the constant component of the treatment variable $\text{Pro}_i = 1$, which is not identified separately in the model. The term δ_t represents the monthly Time FEs, which absorb all time shocks common to the entire sample and make $\text{Post}_t = 1$ unidentifiable separately. The parameter β is the DiD coefficient of interest and estimates the average differential effect on professionals compared to non-professionals in the post-period compared to the pre-period, through the interaction $\text{Pro}_i \times \text{Post}_t$. The term ε_{it} is the idiosyncratic residual; standard errors are clustered at the listing level.

As specified in Section 5.1.3, the control group is not immune to policy: if the reduction in professional supply generates competitive spillovers in favour of non-professionals, β captures the differential gap between the two groups and not the absolute effect on the treated group alone. In the presence of spillovers, the coefficient β does not identify the absolute causal effect on the treated group alone but measures the differential change between treated and controls in the post-period compared to the pre.

In all M1-M4 models, standard errors are clustered at the listing level to manage intra-listing autocorrelation. No weights are used: each observation has a unit weight. This convention is applied uniformly to all estimates reported in the chapter, including mechanisms, DDD, and heterogeneity analyses.

5.4.2 Progressive M1-M5 specifications

In order to assess the stability of the DiD coefficient with respect to the progressive inclusion of structural controls, five specifications are estimated. Model M1 is pooled Ordinary Least Squares (OLS) without fixed effects, included as a benchmark for the raw sign of the association. Model M2 adds monthly Time FEs, controlling for common time shocks but not for persistent differences between listings. Model M3 introduces only Airbnb FEs, absorbing the time-invariant component of listing characteristics but not time trends. Model M4, which includes both sets of fixed effects, is the benchmark estimate for all comparisons and robustness checks. Model M5 extends the TWFE specification by adding City by Time fixed effects.

For *active share*, the most informative change emerges in the transition from M2 to M3: in M2, the coefficient associated with $\text{Pro}_i \times \text{Post}_t$ is -0.0008, while in M3 it becomes -0.0058. The introduction of Airbnb FEs therefore increases the absolute value estimate by about seven times compared to M2, indicating that time-invariant heterogeneity at the listing level, absorbed by Airbnb FEs, is relevant for estimating the post-pre differential between professionals and non-professionals. In other words, the specification with only Time FEs

(M2) and that with only Airbnb FEs (M3) exploit different sources of variation and can produce different estimates when there are structural differences between groups that interact with temporal dynamics; this pattern is consistent with fixed listing characteristics playing a non-negligible role, without implying a definitive interpretation in terms of one specification being biased relative to the other. Furthermore, the estimate is very stable between M3 and M4: the addition of Time FE after controlling for Airbnb FE modifies the coefficient only marginally, suggesting that, once the fixed heterogeneity of listings is absorbed, common monthly shocks do not materially alter the estimate of the differential.

For *occupancy*, the dynamics in the M1–M4 specifications are qualitatively different: the TWFE (M4) estimate indicates a post–pre differential increase for professionals, but of limited magnitude. Unlike *active share*, which directly reflects the host's operational decisions, how many days the listing is active, occupancy is an outcome that is determined only after the offer is exposed and also depends on demand and seasonality; consequently, it is interpreted as a complementary measure of market equilibrium, while the main evidence remains anchored to measures of activity.

For prices, the TWFE (M4) specification returns negative coefficients for both *asking price* and *realized price*; however, the two outcomes differ in terms of the informative quality of the sample: *asking price* is observed on a broader sample, while *realized price* is only available when the realized is observed. Furthermore, the presence of a sign reversal in M3 for *realized price* highlights a greater sensitivity of the estimate to the specification, consistent with the selective nature of this outcome. For these reasons, both price measures are treated as secondary results.

Finally, Model M5 augments the benchmark TWFE specification by adding City by Time fixed effects. This control structure absorbs city-specific shocks that vary over time and provides a more conservative robustness check because identification relies on within-city-month comparisons between professional and non-professional listings. Reassuringly, the estimated coefficients remain close to the benchmark M4 results across all outcomes, suggesting that the main findings are not driven by unobserved city-level time-varying confounders. Given this close alignment between M4 and M5, the benchmark TWFE specification in Model M4 is retained as the reference model for the subsequent analyses to ensure internal consistency and comparability across extensions.

Table 5. DiD M1-M4 main results and M5 diagnostic

Outcome / Term	M1 Pooled OLS	M2 + Time FE	M3 + Listing FE	M4 TWFE main estimate	Robustness: City x Time FE
Panel A — active share					
treat x post	-0.0013*** *(0.000418)*	-0.0008** *(0.000417)*	-0.0058*** *(0.000416)*	-0.0054*** *(0.000416)*	-0.0059*** *(0.000417)*
N	14,997,652	14,997,652	14,997,652	14,997,652	14,968,071
Panel B — occupancy					
treat x post	0.0039*** *(0.000719)*	0.001 *(0.00071)*	0.0082*** *(0.000583)*	0.0042*** *(0.000571)*	0.0048*** *(0.000554)*
N	14,996,984	14,996,984	14,996,984	14,996,984	14,967,403
Panel C — asking price					
treat x post	-0.0021 *(0.00151)*	-0.0077*** *(0.001511)*	-0.0006 *(0.000797)*	-0.0089*** *(0.000818)*	-0.0103*** *(0.00082)*
N	13,498,376	13,498,376	13,498,376	13,498,376	13,476,125
Panel D — realized price					
treat x post	-0.0086*** *(0.001456)*	-0.0138*** *(0.001446)*	0.0014* *(0.000694)*	-0.0040*** *(0.000674)*	-0.0067*** *(0.000662)*
N	9,273,839	9,273,839	9,273,839	9,273,839	9,215,586
Control structure: Monthly FE time: M1 No | M2 Yes | M3 No | M4 Yes | M5 Yes AirbnbFE: M1 No | M2 No | M3 Yes | M4 Yes | M5 Yes City x Time FE: M1 No | M2 No | M3 No | M4 No | M5 Yes SE cluster at the pid level: Yes in all models					

*Note: OLS coefficient of the treat term per post. Standard errors in parentheses. Significance: *** p < 0.01; ** p < 0.05.*

5.4.3 The parallel trends assumption and tools for its evaluation

The causal interpretation of DiD TWFE requires that, in the absence of the policy, the difference in outcome between professional and non-professional listings would have followed a comparable evolution over time. Since the counterfactual is not observable, this condition cannot be verified directly and is assessed indirectly using complementary instruments. In that case, these checks provide less than fully favourable indications, so the coefficients are cautiously interpreted as differential variations between professionals and non-professionals consistent with the impact hypothesis, rather than as conclusive causal estimates.

The first tool is the event study developed in Section 5.6: the pre-policy coefficients provide a diagnostic indicator of the validity of the parallel trends assumption, as they measure the differential evolution between professionals and non-professionals before the policy came into force. Pre-policy coefficients that are systematically different from zero suggest the presence of pre-existing differential dynamics, making the estimate informative but not conclusive in causal terms. The second tool is the placebo sweep illustrated in Section 5.7.2: by replicating the specific TWFE (M4) and placing the policy date in fictitious months of the pre-period, it is verified whether the model produces statistically significant DiD effects even when the policy under study has not yet been implemented. As shown in the following sections, both instruments reveal identification weaknesses that require caution in interpretation; the estimates are therefore read as differential variations consistent with the impact hypothesis and the economic mechanisms observed, without attributing strong causal significance to them.

5.5 Main results

5.5.1 Active share

The DiD coefficient in the specific TWFE (M4) for *active share* is β = -0.0054. The effect is statistically significant in all four specifications: at 5% in M2 and at 1% in M1, M3, and M4. The estimated coefficients range from -0.0013 in M1 to -0.0058 in M3.

In economic terms, the negative sign indicates that in the post-policy period, professional hosts' listings recorded a reduction in the share of active days compared to non-professionals, consistent with an adjustment in the operating margin of supply. This interpretation is further supported by the mechanism variables analysed in Section 5.8.1. Finally, the substantial stability between M3 and M4, a difference of about 0.0004 points, suggests that, once the time-invariant heterogeneity at the listing level is absorbed through the Airbnb FE, the inclusion of the Time FE only marginally changes the estimate of the differential effect.

5.5.2 Occupancy

The DiD coefficient in the TWFE (M4) specification for occupancy is β = +0.0042. The positive sign indicates a differential increase in the occupancy rate for professional listings compared to non-professional listings in the post-policy period. A plausible interpretation is consistent with the structure of the mechanisms discussed in Section 5.8.1: since occupancy is a ratio of booked days to available days, a more marked reduction in availability than the change in bookings can generate an increase in the occupancy rate even in the absence of an increase in demand. From this perspective, the estimated effect on occupancy is interpreted as a result consistent with a mainly supply-side adjustment.

At the same time, occupancy is an equilibrium outcome shaped jointly by supply decisions and demand and seasonal conditions, so its interpretation requires caution. For this reason, the stability of the estimate is assessed explicitly through the robustness analyses reported in the following sections.

5.5.3 Listing prices and realized prices

The log of the list price in M4 records a coefficient of β = -0.0089. The result is also significant in M2, but not in M1 and M3: in particular, in M3 the coefficient is -0.0006. The difference between M3 and M4 suggests that, for *asking price*, the absorption of common monthly shocks through Time FE has a substantial impact on the estimation of the interaction $\mathrm{Pro}_i \times \mathrm{Post}_t$. In other words, the result on *asking price* is more sensitive to the specification, and the significance in M4 but not in M3 indicates that the evidence on list prices depends to a large extent on the combined effect of the two sets of fixed effects.

The log of the realized price presents a further critical issue. In M4, the coefficient is β = -0.0040, while in M3 it is β = +0.0014, thus with the opposite sign compared to the other models. This reversal signals a strong sensitivity of the estimate to the specification, consistent with the fact that *realized price* is observed on a selected sub-sample and that the composition of this sample may vary over time and between specifications.

5.6 Event study: temporal dynamics and diagnostic evaluation of parallel trends

5.6.1 Event study specification and role in the design

The event study plays a dual role in this design, which is important to keep conceptually distinct. The first is diagnostic: the coefficients β_k in the pre-period ($k < 0$) allow us to assess whether the differential between professional and non-professional listings was already evolving before the policy was introduced. Pre-policy coefficients that are systematically different from zero are a sign of possible fragility of the parallel trends assumption and require caution in interpretation. The second role is descriptive of the post-policy dynamics: the coefficients β_k in the post-period illustrate how the differential gap changed month by month after the policy, compared to a reference period.

The main estimate in the chapter remains the TWFE (M4) specification, which summarises the post-change in a single parameter β . However, in light of the pre-trend checks, this parameter is read as a differential variation in the post-period compared to the pre and not as a conclusive causal estimate; the event study does not replace the TWFE but qualifies its interpretability by providing evidence on the dynamics of the pre-trends. The event-study figures are presented using the City by Time specification (M5), which absorbs city-specific shocks that vary over time and therefore provides a more conservative diagnostic in a context characterized by strong seasonality. The resulting dynamic patterns are closely aligned with those obtained under the benchmark TWFE specification, so this choice serves primarily to enhance the transparency of the pre-trend assessment and the description of post-policy dynamics, without affecting the substantive conclusions drawn from the main TWFE estimates.

The event study disaggregates the treatment differential by relative time k, defined as $k = \text{mdate} - 2024m12$. It replaces the single interaction term $Pro_i \times Post_t$ with a set of interactions between Pro_i and relative time indicators, omitting one period as the reference category.

The event study specification is as follows:

$$\boldsymbol{Y_{it} = \alpha_i + \delta_{cxt} + \sum_{k \neq -1} \beta_k (Pro_i \times 1[t - t_0 = k]) + \varepsilon_{it}}$$

The period $k = -1$, corresponding to November 2024, is omitted as the baseline. All coefficients β_k therefore measure the differential between professionals and non-professionals relative to this reference period. The specification includes AirbnbFE (α_i) and City by Time FE (δ_{cxt}); standard errors are clustered at the listing level and estimates are made without weights.

5.6.2 Active share

In the event study on active share, shown in Figure 17, the pre-policy coefficients are not centred around zero and display a systematic pattern, suggesting that professionals and non-professionals may have followed different dynamics already before the policy. In the post-policy period, coefficients indicate an initial reduction in the differential in the first months, followed by a gradual recovery thereafter. Coefficients also exhibit fluctuations that appear partly cyclical. Since the specification includes monthly time fixed effects, common market-wide seasonality is absorbed. Any recurring peaks in the β_k coefficients should therefore be interpreted as differential movements between the two groups, or as estimation noise, rather than as aggregate seasonality affecting all listings equally. Overall, evidence from the pre-policy leads weakens the parallel trends assumption, so the DiD estimates should be read as informative but not conclusive in a causal sense.

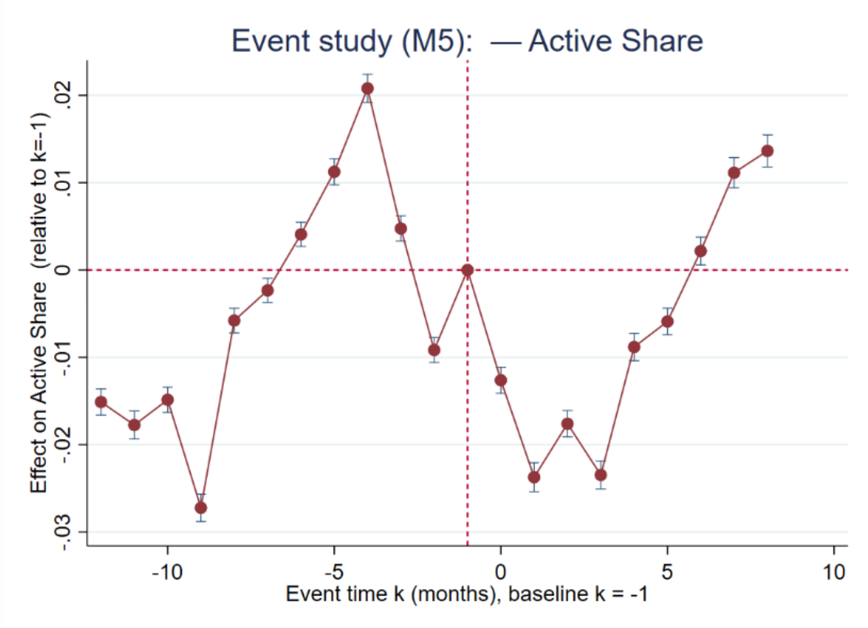

Figure 17. Event study on active share

5.6.3 Occupancy

The event study on occupancy, shown in Figure 18, exhibits a highly noisy profile in the pre-period, with pronounced oscillations in the differential between professionals and non-professionals. Since the specification includes monthly time fixed effects, common market wide seasonality is absorbed. Residual fluctuations in the β_k coefficients therefore reflect changes in the gap between groups, or estimation noise, rather than aggregate seasonality affecting all listings equally. In the immediate pre-period, coefficients vary frequently and often reverse sign, without a clearly identifiable monotonic trend. Relative to active share, whose pre-policy leads display a more systematic pattern, occupancy provides a less clean diagnostic signal. Its high variability makes it difficult to distinguish between a systematic departure from parallel trends and pre-existing differential fluctuations.

In the post-policy period, coefficients are initially below the base month and tend to increase in subsequent months, reaching higher values toward the end of the window. However, given that similar oscillations are already present in the pre-period, this pattern cannot be attributed with confidence to a causal effect of the policy and should be interpreted cautiously.

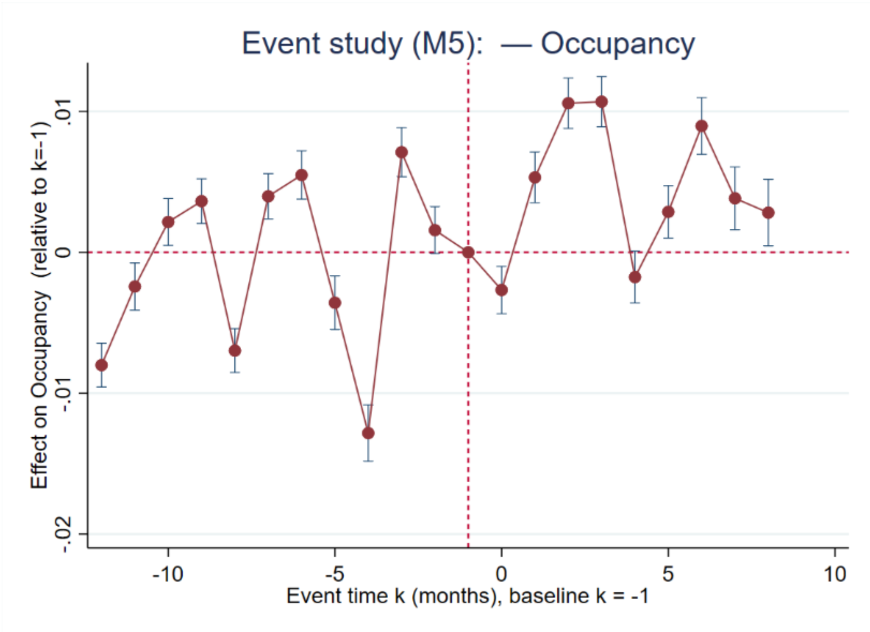

Figure 18. Event study on occupancy

5.6.4 Listing and realized prices

The event studies on prices, shown in Figure 19 and Figure 20, display a markedly cyclical pre-period profile, with a pronounced run-up in the professional–non-professional differential followed by a steep decline as the baseline month approaches. Under the City by Time specification, common month shocks and city-specific seasonal patterns are absorbed, so the remaining variation in the $\beta_k$coefficients reflect differential pricing dynamics between the two groups or residual noise rather than aggregate seasonality. In both series, the post-policy coefficients indicate an immediate negative adjustment in the differential, with a gradual partial rebound over subsequent months. The pattern is qualitatively similar for both the prices measures.

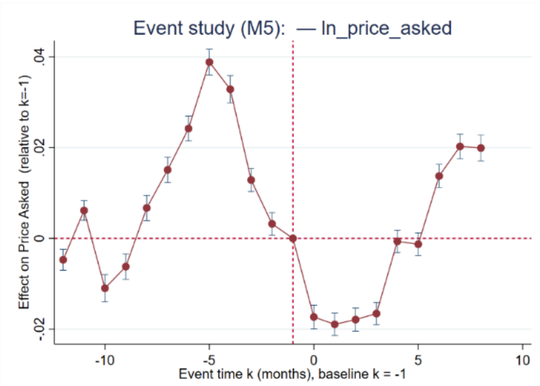

Figure 19. Event study on asking price

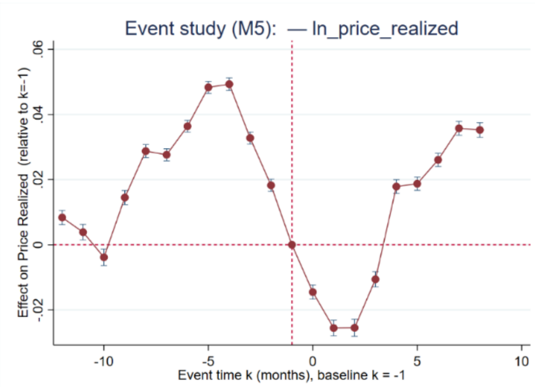

Figure 20. Event study on realized price

5.7 Robustness check

5.7.1 Alternative time windows

The first robustness analysis tests the sensitivity of the estimates to the choice of the estimation window, by restricting the sample around the starting month of the policy. The goal is to reduce reliance on potentially differential long run dynamics between groups, and to verify whether the results are supported within a comparable time frame on both sides of the policy. To this end, a balanced window is adopted, requiring that the number of months in the pre-policy period matches the number of months observed in the post-policy period. Since the post-policy data are available through August 2025, the estimation window runs from March 2024 to August 2025. This choice yields a smaller sample than the full panel and focuses the analysis on the period around the regulation.

Results in the balanced window only partially confirm those found in the complete sample. For *active share*, the estimated effect remains negative and is equal to -0.0069, consistent with the main estimate, and indicative of a relative reduction in the calendar activity of professionals. For *occupancy*, on the other hand, the effect in the balanced window is reduced to +0.0009 with a p-value of 0.142, which is not statistically significant; this result suggests that the evidence on occupancy is more sensitive to the temporal composition and is not robustly supported in the sample balanced around the policy. For prices, the estimates remain negative even in the balanced window, but the magnitude may differ significantly from the previous analysis; for example, for *realized price*, the estimated effect is -0.0195 compared to -0.0040 in the full sample. This amplification is consistent with the idea that price estimates are particularly sensitive to the window and, especially for the realized, to the selection of the observable sub-sample. Results are shown in Table 6.

Table 6. Balanced time window

Outcome	β	SE	t-stat	p-value	N
active share	-0.0069	0.000452	-15.19	< 0.001	8,421,507
occupancy	0.0009	0.000617	1.47	0.142	8,421,507
asking price	-0.0139	0.000813	-17.07	< 0.001	7,543,462
realized price	-0.0195	0.000680	-28.75	< 0.001	5,397,246

Note: M4 TWFE model estimated on the Mar24-Aug25 window.

5.7.2 Multi-date placebo tests

Falsification tests are a central element of DiD design validation because they verify whether the TWFE specification produces effect estimates even in the absence of regulatory intervention. To this end, a multi-date placebo test is implemented, replicating the M4 estimate and replacing the actual policy date with four dates in the pre-period and evaluating a symmetrical time window of ±6 months. The results are shown in Table 7.

Table 7. Multi-date placebo: TWFE estimates with placebo dates in the pre-period.

Placebo date	Outcome	β	SE	t-stat	p-value	N
2023m06	occupancy	0.0054	0.000756	7.09	< 0.001	5,270,891
2023m06	active share	0.0149	0.000619	24.07	< 0.001	5,271,559
2023m06	asking price	0.0047	0.001348	3.45	< 0.001	4,731,634
2023m06	realized price	0.0174	0.000912	19.06	< 0.001	3,252,158
2023m12	occupancy	-0.0045	0.000682	-6.63	< 0.001	6,186,730
2023m12	active share	-0.0161	0.000559	-28.75	< 0.001	6,187,121
2023m12	asking price	-0.0009	0.001178	-0.73	0.466	5,566,317
2023m12	realized price	0.0142	0.000892	15.87	< 0.001	3,821,473
2024m03	occupancy	0.0032	0.000726	4.43	< 0.001	6,429,344
2024m03	active share	0.0157	0.000531	29.58	< 0.001	6,429,431
2024m03	asking price	0.0297	0.000997	29.81	< 0.001	5,786,517
2024m03	realized price	0.0439	0.000899	48.79	< 0.001	3,993,709
2024m06	occupancy	0.0078	0.000731	10.67	< 0.001	6,033,496
2024m06	active share	0.0177	0.000545	32.39	< 0.001	6,033,496
2024m06	asking price	0.0161	0.000959	16.83	< 0.001	5,436,878
2024m06	realized price	0.0243	0.000799	30.37	< 0.001	3,732,562

Note: M4 TWFE model with fictitious data policy.

Overall, the multi-date placebo evidence points to significant critical issues for identification. Of sixteen estimated coefficients, fifteen are statistically significant at 1%. The only exception is *asking price* with a fictitious date of December 2023. The variability of the signs is also informative. For *active share*, positive coefficients are observed when the pseudo policy is set to June 2023, March 2024, and June 2024, while the coefficient becomes negative when the fictitious date is December 2023. This strong sensitivity to the choice of timing of the pseudo policy within the pre-period is consistent with the presence of differential dynamics between professionals and non-professionals even before the policy, which weaken the assumption of parallel trends.

5.8 Economic mechanisms

5.8.1 Adjustment of the supply calendar

To identify the channel through which the policy operated on active share, the TWFE estimates are replicated on the disaggregated variables of the monthly calendar. The rationale is that the DiD coefficient on *active share* is an estimate on a composite measure; the separate estimation of *active days*, *available days*, and *booked days* allows us to distinguish whether the adjustment occurs on the supply side of the calendar or on the demand side. The results are shown in Table 8.

Table 8. Calendar mechanisms: TWFE estimates.

Outcome	β	SE	t-stat	p-value	N
Active days	-0.165	0.013	-12.65	< 0.001	14,997,652
Available days	-0.178	0.01863	-9.53	< 0.001	14,997,652
Booked days	0.012	0.01517	0.81	0.418	14,997,652

The results on mechanisms indicate a contraction in the calendar activity of professionals in the post-policy period. In particular, *active days* decreases significantly, with $\beta = -0.165$ and $p < 0.001$; *available days* decreases similarly, with $\beta = -0.178$ and $p < 0.001$; *booked days*, on the other hand, shows no statistically significant changes, with $\beta = +0.012$ and $p = 0.418$.

A consistency check reinforces this interpretation: applying the M4 estimate for *active share*, with $|\beta| = 0.0054$, to an average month of 30 days, the expected reduction is $0.0054 \times 30 \approx 0.162$ days, a value close to that estimated on *active days* in the main DiD. The small difference is plausibly attributable to the different lengths of the months and the structure of the panel and does not indicate any substantial inconsistency.

Overall, the combination of *booked days*, which is essentially unchanged, and *available days* implies an increase in occupancy for professionals compared to non-professionals in the post-policy period. In summary, the evidence is consistent with a rationalisation of calendar supply: professionals appear to reduce the days exposed on the market without a proportional reduction in bookings.

5.9 Extensions: Triple Difference and Heterogeneity

5.9.1 Triple Difference: Heterogeneity between Entire and Non-Entire

Triple Difference (DDD) analysis was introduced for two complementary reasons. First, it allows us to verify whether the response to the policy shows heterogeneity along a structural dimension of the type of listing, distinguishing between entire and non-entire properties, which plausibly differ in terms of business model, demand elasticity, and operational constraints. Second, DDD represents an internal validation exercise of the DiD design: while it does not resolve any violations of the parallel trends assumption, the appearance of a pattern of heterogeneity consistent with an economically motivated channel reinforces the interpretation of the results as a behavioural response to the policy, rather than as an indistinct fluctuation between groups.

The entire variable is time invariant at the listing level. Each property id is associated with a single value, defined during the construction of the dataset, which does not vary in the panel. The intersection between entire and professional generates the four groups previously evaluated in the descriptive graphs.

The DDD model extends the TWFE specification by introducing the triple interaction $Pro_i \times Post_t \times Entire_i$, also including the two-way interactions $Pro_i \times Post_t$ and $Post_t \times Entire_i$, in addition to the usual fixed effects at the listing level and monthly fixed effects. As a result of the fixed effects at the listing level, all time-invariant components are absorbed and do not appear as separate regressors, in particular Pro_i, $Entire_i$ and $Pro_i \times Entire_i$. On the other hand, the terms that vary over time remain identifiable, namely $Pro_i \times Post_t$, $Post_t \times Entire_i$ and the triple interaction $Pro_i \times Post_t \times Entire_i$.

The coefficient δ on the triple interaction quantifies the difference in the DiD effect between entire and non-entire listings. With the adopted parameterisation, the effect for non-entire properties is β_{nonent}, while the effect for entire properties is $\beta_{entire} = \beta_{nonent} + \delta$. Consequently, if $\beta_{nonent} < 0$, a value $\delta > 0$ indicates that the reduction is less pronounced for entire properties, i.e. that the effect is attenuated for entire properties. The results are shown in Table 9.

Table 9. Triple Difference: Extended M4 TWFE model with triple interaction.

Outcome	β non-entire	delta	β entire	N
active share	-0.00868*** *(0.00117)*	0.00387*** *(0.00125)*	-0.00481*** *(0.000451)*	14,997,652
occupancy	0.00880*** *(0.00180)*	-0.00260 *(0.00190)*	0.00620*** *(0.000611)*	14,996,984
asking price	-0.01856*** *(0.00227)*	0.01533*** *(0.00244)*	-0.00323*** *(0.000897)*	13,498,376
realized price	-0.03210*** *(0.00208)*	0.03388*** *(0.00221)*	0.00179** *(0.000739)*	9,273,839

*Note: βnonent is the coefficient for non-entire; delta is the DDD differential; βentire is equal to the sum of βnonent plus delta. *** $p < 0.01$; ** $p < 0.05$*

For *active share*, the estimated effect for non-entire listings is $\beta_{nonent} = -0.00868$, indicating a marked reduction in the calendar activity of professionals in this category. The triple interaction coefficient is $\delta = +0.00387$, statistically significant and of opposite sign. This implies that, for entire properties, the reduction is attenuated compared to non-entire listings. Consistent with the DDD parameterisation, the total effect for entire properties is $\beta_{entire} = \beta_{nonent} + \delta = -0.00481$. In other words, both subgroups show a reduction in *active share*, but the adjustment is more intense among non-entire listings.

A possible interpretation is that entire properties have different operational constraints and demand conditions, which allow them to limit the contraction in activity. The evidence on list prices is also consistent with this picture. The DDD differential on *asking price* is $\delta = +0.01533$, indicating a smaller reduction in the list price for entire properties than for non-entire listings.

For occupancy, the DDD differential is $\delta = -0.00260$ with $p = 0.171$, which is not statistically significant. There is therefore no robust evidence that the adjustment in the occupancy rate differs systematically between entire and non-entire listings.

For *realized price*, the DDD differential is positive and significant, with $\delta = +0.03388$, implying an estimated effect for entire properties of $\beta_{entire} = +0.00179$. However, this result should be interpreted with caution considering the selection of the realised sample discussed in previous section.

5.9.2 Heterogeneity by operational scale within-professional

This section does not aim to estimate a causal effect in DiD form comparable to Table 5. The analysis is conducted exclusively within the professional group and does not use non-professionals as a counterfactual; consequently, it does not identify a difference-in-differences between groups but describes how outcomes evolve in the post-policy period among professionals with different operational scales. It follows that the coefficients in this table should not be interpreted as DiD estimates relative to the control group. The results are shown in Table 10.

Table 10. Heterogeneity by operational scale within professionals. Baseline: hosts with 2-5 listings.

Outcome	Term. Baseline: hosts with 2-5 listings	B	SE
Panel A — active share, N = 11,065,521			
active share	Post — baseline 2-5 listings	-0.0117***	0.0007
	6-20 vs. 2-5 listings differential	-0.005***	0.00049
	Differential 21+ vs 2-5 listings	-0.0163***	0.00054
Panel B — occupancy, N = 11,064,914			
occupancy	Post — baseline 2-5 listings	0.2566***	0.00089
	Difference between 6-20 and 2-5 listings	0.0033***	0.00072
	Differential 21+ vs. 2-5 listings	0.0155***	0.00075
Panel C — asking price, N = 9,978,029			
asking price	Post — baseline 2-5 listings	0.4286***	0.00127
	Difference between 6-20 and 2-5 listings	-0.0101***	0.00103
	Differential 21+ vs 2-5 listings	0.0802***	0.00141
Panel D — realized price, N = 7,019,566			
realized price	Post — baseline 2-5 listings	0.3679***	0.00113
	Difference between 6-20 and 2-5 listings	-0.0144***	0.00081
	Differential 21+ vs 2-5 listings	0.0345***	0.00125

*Note: *** $p < 0.01$.*

For *active share*, the post-policy change for the reference group, hosts with 2–5 listings, is −0.0117. Compared to this baseline, hosts with 21 or more listings show an additional differential of −0.0163, while the intermediate group, 6–20 listings, shows a differential of −0.0050. The evidence shows a clear dimensional gradient: the reduction in *active share* in the post-policy period increases with operational scale, with a particularly marked discontinuity in the 21+ category. The result is economically significant: while it cannot rule out the presence of economies of scale in the organisation of check-in, the data suggest that these economies are not sufficient to offset the greater exposure to variable costs associated with a high volume of arrivals. In other words, the observed heterogeneity is more consistent with a channel linked to operating volume than with a full effect of organisational capacity for larger operators. Consistently, large hosts seem to concentrate the adjustment through a selective reduction in activity, plausibly reallocating the portfolio towards the most profitable units, rather than through price increases. Since the analysis is conducted exclusively within the professional segment and is based on a pre-post variation without an external control group, these differentials should be interpreted as descriptive evidence of heterogeneous adjustment and not as DiD estimates relative to non-professionals.

For occupancy, the differentials across scale groups are positive and statistically significant, with hosts managing 21 or more listings showing the largest relative increase, equal to +0.0155. Even in this case the result suggests that the evolution of occupancy in the post-period systematically differs depending on the operational scale, while still subject to the same interpretative cautions related to the absence of an external counterfactual.

For prices, the relationship with scale does not follow a simple increasing or decreasing pattern. Relative to the baseline category of hosts managing 2 to 5 listings, the intermediate group of 6 to 20 listings shows a negative post-period differential for both asking prices and realized prices, equal to −0.010 and −0.014, respectively. The largest operators, those managing 21 or more listings, instead show a positive differential, equal to +0.080 for asking prices and +0.035 for realized prices. Overall, these results point to heterogeneous pricing dynamics across scale groups. Medium-scale operators may face tighter constraints in adjusting prices in the post-period, resulting in relative price compression, whereas the largest operators display higher differentials, potentially reflecting differences in portfolio composition or a stronger presence in higher-yield segments. The scale differentials nonetheless document substantial dispersion within the professional segment and motivate further investigation in future research.

5.9.3 Heterogeneity by pre-intervention intensity

To explore whether the response to the policy differs with pre-intervention listing intensity, a continuous measure equal to the average number of *booked days* in the pre-policy period is constructed for each listing. This measure is then centred and standardized so that a value of zero corresponds to an average listing in the pre-period, and a one-unit increase corresponds to one standard deviation higher pre-policy intensity. The coefficient β_{tp} captures the estimated effect for a listing with average pre-policy intensity, while β_{int} captures how this effect varies with pre-policy intensity, which is the heterogeneity component. The specification extends the TWFE model by interacting the treatment term with the standardized pre-policy intensity measure. Under this parameterization, the main coefficient describes the effect for a listing of average intensity, while the interaction coefficient indicates whether the effect becomes more or less pronounced as pre-policy bookings increase. A negative interaction coefficient implies that higher intensity listings in the pre-period experience a relatively larger reduction in the outcome in the post-period. Results are reported in Table 11.

Table 11. Heterogeneity by pre-intervention intensity: interaction with pre-policy booked days.

Outcome	β_tp	SE_tp	β_int	SE_int	N
active share	-0.00322***	0.000417	-0.01215***	0.000177	14,891,036
occupancy	0.01258***	0.000562	-0.04659***	0.000266	14,890,368
asking price	-0.00742***	0.000823	-0.01272***	0.000444	13,401,845
realized price	0.00416***	0.000725	-0.01330***	0.000401	9,205,689

*Note: M4 TWFE model with continuous interaction. β_tp is the treat-for-post coefficient for a listing with pre-intervention booked days equal to zero; β_int measures the differential change per unit increase in pre-policy booked days. *** p < 0.01.*

For *active share*, the interaction coefficient is $\beta_{int} = -0.01215$ with $p < 0.001$, indicating that the post-policy adjustment is more pronounced for professional listings with greater pre-intervention intensity. As pre-period bookings increase, the relative reduction in *active share* in the post-period becomes stronger. For *occupancy*, the interaction coefficient is $\beta_{int} = -0.04659$, suggesting that the relative increase in occupancy in the post-period is attenuated for listings with greater pre-policy intensity, consistent with the idea that already highly booked listings have less room for further increases in occupancy.

However, these results should be interpreted as exploratory evidence of heterogeneity. Pre-intervention intensity may be correlated with unobserved listing characteristics that affect both the response to the policy and outcome dynamics over time.

Chapter 6
Conclusion

The present work has analysed the possible effect of regulation on the short-term rental market in Italy. The legislation, which aims to preserve public safety by limiting remote check-in, represents an intervention that may indirectly affect the STR market, with potentially stronger implications for professional operators. In order to assess its effects, empirical analyses were conducted to measure the market response, distinguishing between professional and non-professional hosts.

Before presenting the results, it is useful to assess the extent to which the qualitative forecasts are supported by empirical evidence. The first prediction, concerning adjustments prevailing on the intensive margin, is supported by the results on calendar mechanisms; there is a significant reduction in active days and available days, while booked days does not show statistically significant variations, consistent with an adjustment of supply through calendar management rather than a proportional contraction in actual demand. The second prediction, concerning heterogeneity with respect to pre-policy intensity, is also consistent with the available evidence; estimates for pre-intervention intensity indicate an interaction coefficient $\beta_{int} = -0.012$ on active share, suggesting a relatively more pronounced impact for listings with higher pre-regulation intensity. The third prediction, relating to potentially ambiguous heterogeneity by host type and property type, finds less obvious but still informative confirmation; the within-professional dimensional gradient and the DDD differential between entire and non-entire suggest that channels related to operating volume and organisational capacity can both be influential, with sign and intensity varying across different dimensions.

The main result of the empirical analysis suggests that, for professional hosts, the active share is reduced by about 0.54 percentage points compared to non-professionals, corresponding to about 0.62% of its sample mean. This evidence, combined with the reduction in the number of days on which the listing is available and a substantially stable share of days booked, allows us to interpret the phenomenon as an adjustment on the supply side: professionals seem to reduce the availability, thereby optimizing the use of available capacity. Despite the identification weaknesses highlighted by event studies and placebo tests, the convergence of results on different but related measures makes this outcome particularly informative in describing this adjustment. In addition, the reduction in active share in the post-policy period is more pronounced as the size of the business increases, particularly for hosts with more than 21 listings, suggesting a size gradient in the adjustment. This pattern is consistent with the interpretation that larger operators, which are more exposed to regulation and have greater flexibility in managing their portfolios, concentrate the reduction in activity through a strategic reallocation of supply.

With regard to occupancy, the main analysis shows a slight increase of 0.42 percentage points, consistent with the evidence relating to booked and active days. However, the robustness analysis suggests that this result tends to fluctuate or lose statistical significance in the symmetrical and narrow time window around the policy, suggesting sensitivity to window choice and seasonal dynamics.

As regards prices, both asking and realized, the coefficients are statistically significant, despite the smaller sample relative to the other outcome variables, and show a counterintuitive trend. While one would expect, in response to a regulation that affects the variable costs of operators, an increase in prices mainly borne by professional hosts, both measures instead show a negative variation, a decline of about 0.9% in asking prices and 0.4% in realized prices for professionals relative to non-professionals. Three possible explanations can be put forward for this phenomenon. First, the post-policy period observed does not cover a full calendar year, so seasonal dynamics may influence the results. Second, the dataset does not include a variable that tracks additional service fees on top of the transaction price, which could introduce a distortion in the estimates. Third, it is plausible that some professional hosts have outsourced check-in management to third parties, with payment taking place directly between the parties without being intercepted by the platform. In this scenario, the additional cost of check-in is covered by the guest but is not incorporated into the price visible on the platform: this generates an iceberg effect in the observable data; the visible part, the price on the platform, remains stable or decreases due to competition while the hidden part, the cost of managing the arrival increases, shifting the cost from the host to the guest.

These considerations suggest caution in interpreting the negative price trend in the post-policy period.

The results obtained using Two-Way Fixed Effects (TWFE) should be interpreted with caution. As documented by event studies, the assumption of parallel trends is challenged for all outcome variables considered. The pre-treatment coefficients show a structural divergence between professional and non-professional hosts already prior to the entry into force of the legislation, with heterogeneity that appears to be related to seasonal market cycles. This differential divergence, i.e., the asymmetry with which the two groups respond to seasonality, and not seasonality itself, is not captured by time fixed effects and compromises the causal identification of the estimates. For price variables, this behaviour is consistent with evidence that professional hosts exercise greater market power during peak seasons, resulting in a systematically wider price differential than during low seasons (Casamatta et al., 2022). More generally, hosts of different sizes exhibit structurally heterogeneous seasonal patterns, suggesting that the adoption of non-professionals as a control group may not be sufficient to ensure pre-treatment comparability (Sainaghi & Baggio, 2021). The TWFE results should therefore be understood as descriptive estimates of the association between professionalism and policy response, and not as causal estimates.

Future research could further explore the analyses conducted in this study on several fronts. On a methodological level, the adoption of estimators robust to treatment heterogeneity would allow the causal effects of the policy to be isolated even in the presence of the pre-treatment structural divergences documented in this study. At the same time, the identification of a control group more homogeneous to professional hosts would further strengthen the credibility of causal identification. On a substantive level, a natural extension concerns the geographical heterogeneity of the effects: it is reasonable to expect that the impact of the regulation will differ significantly between markets with a high concentration of tourism, such as Rome, Florence and Venice, and smaller markets, a dimension that the aggregate analysis at the national level is unable to fully capture. Furthermore, the time window available at the time of the analysis covers only the first few months after the policy came into force: with more extensive data, it will be possible to assess whether the documented effects stabilize, attenuate as a result of strategic adaptation by hosts, or amplify in the medium term. Finally, future studies could examine the implications of regulation on local real estate markets and traditional accommodation supply, dimensions that are beyond the scope of this paper, but which represent the most relevant policy implications for regulators.

Bibliography

Ajuntament de Barcelona. (2017). *Pla Especial Urbanístic d'Allotjament Turístic (PEUAT).* Barcelona City Council.

Barone, G., Conti, L., & Papini, G. (2019). L'impatto economico di Airbnb: il caso italiano. In A. Petrella & R. Torrini (Coords.), *Turismo in Italia: numeri e potenziale di sviluppo* (pp. 22–23). Questioni di Economia e Finanza n. 505. Banca d'Italia.

Barron, K., Kung, E., & Proserpio, D. (2021). The effect of home-sharing on house prices and rents: Evidence from Airbnb. *Marketing Science,* 40(1), 23–47.

Bibler, A. J., Teltser, K. F., & Tremblay, M. J. (2025). How do short-term rental regulations affect market outcomes? Evidence from San Francisco's Airbnb registration requirement. *Real Estate Economics,* 53, 841–868.

Calder-Wang, S. (2021). The distributional impact of the sharing economy on the housing market (SSRN Scholarly Paper No. ID 3908062). Social Science Research Network.

Callaway, B., & Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics, 225*(2), 200–230.

Casamatta, G., Giannoni, S., Brunstein, D., & Jouve, J. (2022). Host type and pricing on Airbnb: Seasonality and perceived market power. *Tourism Management,* 88, 104433.

Chen, W., Wei, Z., & Xie, K. (2023). Regulating professional players in peer-to-peer markets: Evidence from Airbnb. *Management Science,* 69(5), 2893–2918.

de Chaisemartin, C., & D'Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review, 110*(9), 2964–2996.

Dogru, T., Hanks, L., Mody, M., Suess, C., & Sirakaya-Turk, E. (2020). The effects of Airbnb on hotel performance: evidence from cities beyond the United States. *Tourism Management,* 79, 104090.

Evans, D. S., & Schmalensee, R. (2016). *Matchmakers: The New Economics of Multisided Platforms.* Harvard Business Review Press.

Farronato, C., & Fradkin, A. (2022). The Welfare Effects of Peer Entry: The Case of Airbnb and the Accommodation Industry. *American Economic Review* 112(6), 1782–1817.

Garcia-López, M.-A., Jofre-Monseny, J., Martínez-Mazza, R., & Segú, M. (2020). Do short-term rental platforms affect housing markets? Evidence from Airbnb in Barcelona. *Journal of Urban Economics,* 119, 103278.

Gauß, P., Gensler, S., Kortenhaus, M., Riedel, N., & Schneider, A. (2024). Regulating the sharing economy: The effects of day caps on short- and long-term rental markets and stakeholder outcomes. *Journal of the Academy of Marketing Science.*

Gurran, N., & Phibbs, P. (2017). When tourists move in: How should urban planners respond to Airbnb? *Journal of the American Planning Association,* 83(1), 80–92. https://doi.org/10.1080/01944363.2016.1249011

Guttentag, D. (2015). Airbnb: disruptive innovation and the rise of an informal tourism accommodation sector. *Current Issues in Tourism,* 18(12), 1192–1217.

Jefferson-Jones, J. (2015). Airbnb and the housing segment of the modern "sharing economy": Are short-term rental restrictions an unconstitutional taking? *Hastings Constitutional Law Quarterly,* 42.

Koster, H. R., Van Ommeren, J., & Volkhausen, N. (2021). Short-term rentals and the housing market: Quasi-experimental evidence from Airbnb in Los Angeles. *Journal of Urban Economics,* 124, 103356.

Lee, M., Sisson, A., Costa, R., & Bai, B. (2023). Disruptive technologies and innovation in hospitality: A computer-assisted qualitative data analysis approach. *Journal of Hospitality & Tourism Research,* 47. https://doi.org/10.1177/10963480231156080

Nieuwland, S., & van Melik, R. (2020). Regulating Airbnb: how cities deal with perceived negative externalities of short-term rentals. *Current Issues in Tourism,* 23(7), 811–825.

Rochet, J.-C., & Tirole, J. (2006). Two-sided markets: A progress report. *The RAND Journal of Economics,* 37(3), 645–667.

Sainaghi, R., & Baggio, R. (2021). Airbnb host scaling, seasonal patterns, and competition. *Journal of Knowledge Management Practice,* 26(1).

Valentin, M. (2021). Regulating short-term rental housing: Evidence from New Orleans. *Real Estate Economics,* 49, 152–186.

Wachsmuth, D., & Weisler, A. (2018). Airbnb and the rent gap: Gentrification through the sharing economy. *Environment and Planning A: Economy and Space,* 50. https://doi.org/10.1177/0308518X18778038

Xie, K. L., & Kwok, L. (2021). Do professional hosts matter? Evidence from multi-listing and full-time hosts in Airbnb. *International Journal of Hospitality Management,* 96, 102950.

Zervas, G., Proserpio, D., & Byers, J. W. (2017). The rise of the sharing economy: Estimating the impact of Airbnb on the hotel industry. *Journal of Marketing Research,* 54(5), 687–705.