

Information Shocks and Housing Market Dynamics: Evidence from TSMC's New Fab

Fang-Chang Kuo*, Tzu-Hsiang Hsu[†]

May 22, 2025

PRELIMINARY AND INCOMPLETE. COMMENTS ARE GREATLY APPRECIATED.

Abstract

This paper investigates the informational efficiency of the housing market by examining how property prices responded to a series of announcements related to TSMC's new fab in Kaohsiung. Using transaction-level data from Taiwan, we document an immediate 5% to 10% increase within three weeks. Event-study estimates further reveal gradual price adjustments over subsequent months, suggesting under-reaction to information shocks. We also observe significant declines in prices and transaction volume in response to negative updates, indicating limited downward price rigidity. Importantly, housing market responses appear to be driven by speculative activity, which is evidenced by a surge in transaction volumes immediately after the announcement and a subsequent increase in vacancy rates. We find that vacancy rates in Nanzi rose by 1.5 percentage points relative to similar districts. These results highlight both the informational responsiveness and speculative dynamics in local housing markets facing major industrial investment shocks.

JEL Codes: G14, R11, R21

Keywords: Information revelation; TSMC; Housing prices; Transaction volumes

*National Chung Cheng University, Department of Economics Email: fckuo@ccu.edu.tw

[†]Master's student at National Chung Cheng University

1 Introduction

The housing market represents a core component of the economy, accounting for a substantial share of household wealth and financial activity.¹ Yet, despite its central role, housing markets are often characterized by inefficiencies arising from information asymmetries, high transaction costs, and behavioral frictions. Understanding how efficiently housing markets incorporate new information is therefore crucial for researchers and policymakers alike. However, empirical investigations of informational efficiency have been limited by the scarcity of high-frequency data and the infrequent occurrence of clearly exogenous information shocks.

At the same time, the global strategic significance of semiconductors has increased substantially. These components are indispensable to a broad array of technologies, including consumer electronics, automobiles, defense systems, and artificial intelligence. As the world’s largest semiconductor foundry, Taiwan Semiconductor Manufacturing Company (TSMC) occupies a pivotal role not only in global supply chains and geopolitics, but also as a focal point in the study of industrial policy and technological strategy (Miao, 2024; Goldberg et al., 2024; Thurk, 2024). TSMC’s investment decisions—particularly regarding the siting of new fabrication facilities—constitute not only industrial milestones but also events capable of producing significant localized economic impacts, especially in real estate markets.

This paper exploits a sequence of announcements concerning TSMC’s new fabrication plant in Kaohsiung as a series of natural experiments to examine the informational efficiency of the housing market. Leveraging daily transaction-level data, we assess how quickly and thoroughly housing prices adjust in response to these exogenous information shocks. Our analysis begins with the initial news leak, which triggered a 5% increase in transaction prices within the first week and a 10% increase over the subsequent three weeks. To explore potential heterogeneity across housing segments, we compare cumulative density functions and conduct formal tests for first-order stochastic dominance. The results indicate that the price increase was broad-based and not limited to a particular segment of the market,

¹ As of the end of 2023, housing accounted for over 30% of household net worth in Taiwan. See: <https://eng.stat.gov.tw/News.aspx?n=2416&sms=10854>

suggesting widespread market responsiveness to new information.

To investigate long-term trends and responses to sequential information updates, we employ a difference-in-differences (DID) event-study methodology. This approach compares monthly price differentials between treated properties—those located in areas directly affected by TSMC’s investment—and untreated properties that serve as a control group. Our empirical findings yield several key insights. First, the housing market reacts not only to formal announcements but also to early news reports, suggesting that information begins to diffuse through informal channels before official confirmation. Second, price adjustments unfold gradually over time, indicative of a form of “under-reaction,” where market participants assimilate new information into housing prices in a staggered or delayed fashion. Third, we find limited evidence of downward price rigidity; housing prices declined in response to negative developments, such as the downsizing of TSMC’s investment plan. However, these negative shocks were also associated with a noticeable drop in transaction volumes, reflecting increased market uncertainty or reduced willingness to trade. Finally, the market appears to place a premium on technological advancement: when TSMC confirmed the adoption of the advanced 2nm process, housing prices rose further, signaling the capitalization of anticipated economic benefits tied to cutting-edge manufacturing capabilities.

In addition to examining price responses, understanding the underlying forces driving housing market dynamics is essential. To this end, we begin by analyzing short-term transaction volumes in the Nanzi District. We observe a sharp surge in transaction activity following TSMC’s initial announcement, which suggests that the observed price increases are unlikely to be driven purely by local residential demand. Instead, the spike in volume is more consistent with speculative behavior, where investors rapidly enter the market in anticipation of future capital gains.

To gain a more comprehensive view of longer-term dynamics, we turn to vacancy-related indicators—specifically, the number of low-use houses and the vacancy rate—to assess whether price surges coincide with an increase in underutilized properties. Our analysis reveals that, following the initial announcement, vacancy rates in Nanzi rose significantly, narrowing the

gap between Nanzi and other districts in Kaohsiung. This pattern further supports the hypothesis that a substantial share of post-announcement transactions may have been driven by speculation rather than end-user demand.

We adopt a difference-in-differences (DID) event-study framework to formally test the speculation conjecture. This specification includes a control for the total number of tax-eligible residential units, allowing us to account for changes in the housing supply over time. Controlling for housing stock is essential, as it helps distinguish increases in vacancy rates driven by speculative holding behavior from those caused by structural growth in newly built but unoccupied housing. This approach enables a more precise estimation of the impact of informational shocks on speculative activity in the housing market. Our findings indicate that, following the announcement, vacancy rates in the Nanzi District gradually increased, reaching approximately 1.5 percentage points above pre-announcement levels by mid-2024.

Our study sheds light on the informational efficiency of the housing market and contributes to several strands of literature. First, our findings relate to seminal work in finance on how markets respond to information shocks (Busse and Green, 2002; Chan, 2003). Second, we contribute to a growing body of research examining informational frictions in housing markets (Larsen and Weum, 2008; Hjalmarsson and Hjalmarsson, 2009; Anundsen and Røed Larsen, 2018; Broxterman and Zhou, 2023). A key innovation of our analysis lies in the use of daily transaction-level data to capture short-term price dynamics. To our knowledge, this is the first study to show that housing markets can incorporate new information within a matter of weeks. Moreover, given the highly differentiated nature of residential real estate, the observed shift in the entire distribution of transaction prices suggests that the effects of the information shock were broad-based, impacting all segments of the market.

This paper also contributes to the literature on housing market responses to various exogenous shocks, including the siting of industrial facilities (De Vor and De Groot, 2011), entry of large retailers such as Walmart (Pope and Pope, 2015), corporate headquarters announcements (Chen et al., 2024; Hu et al., 2025), technological innovation (Beracha et al., 2022), and the release of risk-related information (Bernstein et al., 2019; Yi and Choi, 2020;

Cheng et al., 2022). Notably, while TSMC’s fabrication plant is fundamentally an industrial facility—typically associated with negative externalities and downward pressure on nearby housing prices—our results show that high-tech facilities of this nature can exert a positive influence on real estate values.² This finding suggests that the perceived economic benefits and positive externalities associated with advanced semiconductor manufacturing may outweigh the usual drawbacks linked to traditional industrial sites.

Finally, this paper contributes to the growing literature that leverages recent large-scale natural experiments—such as Amazon’s HQ2 selection—to evaluate market efficiency (Chen et al., 2024) and consumer expectations (Famiglietti et al., 2023). Our study is particularly closely related to Chen et al. (2024), as both papers investigate housing market responses to initial announcements and subsequent information updates. However, unlike their findings, which do not attribute market movements to speculation, our results suggest that price responses are at least partially driven by speculative behavior. This interpretation is supported by a sharp increase in transaction volumes and a subsequent rise in vacancy rates following the announcements, indicating that buyers may have been motivated by expected capital gains rather than long-term residential use.

The remainder of the paper is structured as follows. Section 2 outlines the background, data sources, and sample construction. Section 3 details the empirical strategy and presents the main results. Section 4 concludes.

2 Background and Data

Taiwan Semiconductor Manufacturing Company (TSMC) is the world’s largest dedicated semiconductor foundry and a central player in the global electronics supply chain. Founded in 1987 and headquartered in Hsinchu, Taiwan, TSMC produces high-performance chips for a wide range of applications, powering devices such as smartphones, laptops, game consoles, autonomous vehicles, and data centers. As the primary contract manufacturer for industry

²Following the COVID-19 pandemic, TSMC has announced major fab expansions in countries including the United States, Germany, and Japan. See: <https://www.reuters.com/technology/tsmc-prizes-japans-chips-skills-after-us-stumbles-sources-2023-09-12/>

giants including Apple, AMD, and Nvidia, TSMC's technological leadership and production capacity are critical not only to the global technology sector but also to national security and economic policy considerations across advanced economies.

In recent years, TSMC has expanded its fabrication footprint in response to growing geopolitical uncertainty, supply chain vulnerabilities, and surging global demand for semiconductors. As part of its diversification and capacity expansion strategy, TSMC has committed to building advanced fabrication facilities (fabs) not only in Taiwan but also in the United States, Japan, and Germany.³ These investment decisions are shaped by a combination of industrial policy incentives, customer proximity, and long-term capacity planning.

In Taiwan, a key decision was made in November 2021, when TSMC announced plans to construct a new fabrication facility in the Nanzi District of Kaohsiung. The project was initially set to include production lines for both the advanced 7-nanometer process and the mature 28-nanometer process, with completion targeted for 2024. This announcement marked a significant shift in the geographic distribution of Taiwan's semiconductor manufacturing, which has traditionally been concentrated in Hsinchu and Tainan. The decision to invest in southern Taiwan—particularly in Kaohsiung, a city with a robust industrial base but limited prior involvement in high-end semiconductor fabrication—came as a complete surprise to many observers. Nonetheless, it was widely regarded as a potential catalyst for regional economic development and industrial upgrading.

However, global macroeconomic headwinds and softening demand for chips led to strategic adjustments. In October 2022, TSMC announced a delay in the 7-nanometer production timeline due to tool delivery issues and capital expenditure constraints. Six months later, in April 2023, the company suspended the 28-nanometer expansion altogether, citing demand-side concerns. In August 2023, TSMC confirmed that the fab would instead focus on producing chips using its cutting-edge 2-nanometer process, reaffirming its long-term commitment to maintaining global technological leadership.

This evolving investment timeline—marked by public announcements, delays, and technological upgrades—has had significant implications for local markets, particularly in the

³ See: <https://www.reuters.com/technology/tsmc-prizes-japans-chips-skills-after-us-stumbles-sources-2023-11-11/>




Figure 1: Google Search Intensity

Notes: This figure plots weekly search intensities of “TSMC Kaohsiung” in Chinese, on Google search from January 2019 to December 2023. The search intensity ranges from 0 to 100. The intensities, ranging from 0 to 100, do not reflect actual numbers of searches as Google normalizes the highest search number of a keyword to 100 during the sample periods. Source: <https://trends.google.com.tw/trends>.

real estate sector. The prospect of hosting a next-generation TSMC fab has influenced not only infrastructure and labor planning in Kaohsiung but also expectations and behaviors in asset markets, most notably in housing. As a result, this setting offers a rare and valuable opportunity to study how local economic agents respond to a sequence of information shocks stemming from a major industrial investment decision.

Figure 1 plots the weekly Google search intensity for the term “TSMC Kaohsiung” in Chinese. Notably, a substantial spike occurred in September 2021, prior to the official announcement in November. This suggests that local news coverage and informal information channels had already begun to signal TSMC’s potential siting decision—pointing to a case of information leakage. Additionally, further increases in search intensity can be observed before October 2022, coinciding with emerging reports about possible delays and adjustments

Table 1: Timeline of TSMC’s Fab in Kaohsiung

Date	News and Announcements
2021-09-07	Supply chain confirmed that 7nm fab will be located in Nanzi, Kaohsiung (News)
2021-11-09	TSMC to build new chip factory in Nanzi for both 7nm and 28nm
2022-08-01	Delays, and site issues (News)
2022-10-13	TSMC cuts capex on tool delays, 7nm delayed
2023-04-20	28nm expansion suspended
2023-08-08	Advanced 2nm fab confirmed, no 7nm and 28nm

to the construction plan.

In 2023, the search intensity data reveal another prominent spike, followed by two smaller ones. The first major spike aligns with TSMC’s suspension of the 28nm expansion, an event that introduced uncertainty regarding the project’s future direction. The two smaller spikes are likely associated with increasing speculation that TSMC would shift its investment toward the more advanced 2nm process in Kaohsiung. Taken together, these fluctuations in search behavior correspond closely with the timeline of announcements and revisions, supporting our identification of these events as sequential, exogenous information shocks to the local housing market. Table 1 summarizes the key announcements and their corresponding dates.

The primary dataset used in this study is derived from Taiwan’s Actual Price Registration System (APRS), which provides comprehensive, government-mandated housing transaction records. Launched in August 2012, the APRS captures the universe of residential real estate sales across the country. Each transaction record includes detailed information such as the date of sale, final transaction price, and various housing characteristics.⁴ Given the comprehensive nature of the dataset—capturing all legally registered transactions—it is possible to compute transaction volumes and price trends across different geographic regions and time periods. This level of coverage and detail makes the APRS an ideal source for evaluating

⁴ Housing characteristics include building age, floor level, property type, number of rooms and bathrooms, and the full address. This study restricts its analysis to residential property transactions.

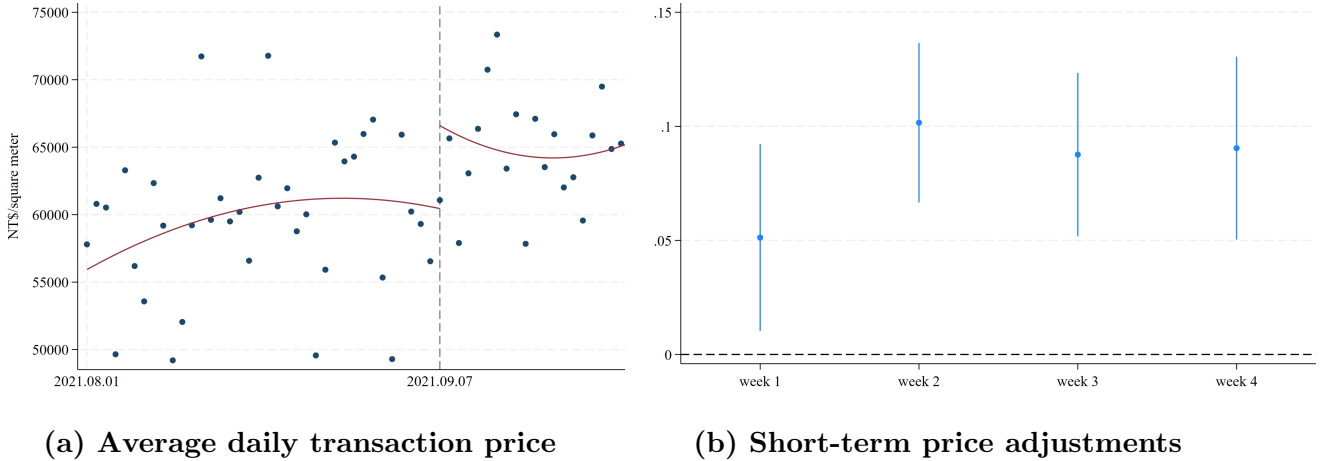


Figure 2: Short-term Dynamics of Housing Prices

market-level responses to policy or economic shocks, such as those examined in this study.

3 Empirical Analysis

3.1 Immediate Price Reactions

In this section, we evaluate the speed at which housing prices responded to the initial announcement concerning TSMC’s planned fab in the Nanzi District. To do so, we analyze transaction data specifically from August to September 2021. Figure 2a presents the unconditional average transaction prices per square meter for residential properties in Nanzi during this two-month period. The vertical dashed line in the figure indicates the date (September 7, 2021) when news reports first revealed TSMC’s intention to construct the new fab. The quadratic fitted lines separated by the vertical dashed line clearly illustrate a significant upward shift in average transaction prices immediately following the news disclosure. Prior to the announcement, daily average transaction prices predominantly remained below 60,000 NT\$ per square meter, whereas post-announcement prices consistently surpassed this threshold.

To quantify the precise timing and magnitude of price adjustments, we estimate the

following hedonic price model:

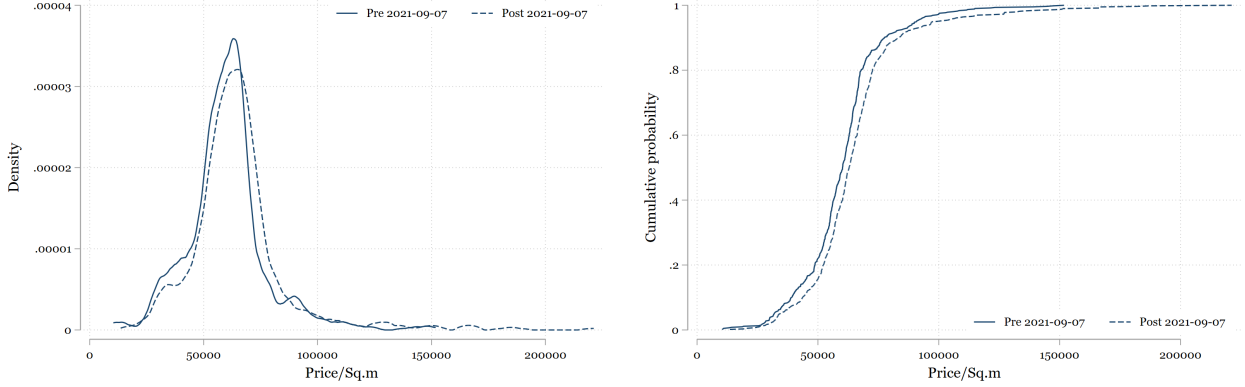
$$\log(p_{it}) = \sum_{w=1}^4 \eta_w \cdot D_w + \mathbf{x}_{it}\boldsymbol{\beta} + \epsilon_{it}$$

Here, $\log(p_{it})$ denotes the natural logarithm of the transaction price for property i at time. The matrix \mathbf{x}_{it} includes property-specific characteristics such as age, number of rooms, building type, and size. Our key explanatory variables are weekly indicator variables (D_w), marking each of the four weeks immediately following the initial news report on September 7, 2021. The parameters of interest, η_w , measure the average weekly price response.

Figure 2b illustrates these estimated weekly coefficients. The results indicate an immediate 5% increase in transaction prices within the first week after the initial announcement. Subsequent weeks exhibit a larger and relatively stable price premium, ranging between 9% and 10%. This pattern suggests the market rapidly incorporates the new information, with prices stabilizing after approximately two weeks.

In addition to the average price responses within a month, transaction data also allow us to analyze shifts in the price distribution resulting from the announcement. Figure 3 presents both the probability density and cumulative distribution functions (CDFs) of transaction prices per square meter in the Nanzi District before and after TSMC’s fab announcement. Both density and cumulative distributions visibly shift to the right post-announcement, indicating higher prices across the entire distribution. This indicates that the rise in transaction price following TSMC’s decision is widespread, benefiting sellers across the spectrum of residential properties.

To statistically validate the observed shift, we will conduct formal Kolmogorov-Smirnov tests for first-order dominance. Table 2 reports the p -values from one-sided Kolmogorov-Smirnov (KS) tests for first-order stochastic dominance (FOSD) between the pre- and post-announcement distributions of transaction prices per square meter. The test of the null hypothesis $H_0 : F(P_{pre}) \leq F(P_{post})$ yields a p -value of 0.000, strongly rejecting the null in favor of the alternative that the post-announcement distribution is stochastically greater than the pre-announcement distribution. Conversely, the reverse test $H_0 : F(P_{post}) \leq F(P_{pre})$



(a) Density

(b) Distribution

Figure 3: Short-term Pre/Post Price Distribution

Notes: This figure presents the estimated (a) probability density functions and (b) cumulative distribution functions of housing transaction prices per square meter in the Nanzi District, comparing the periods before and after the initial announcement of TSMC’s fab. The sample period covers transactions from August 1, 2021, to September 30, 2021.

Table 2: Stochastic Dominance Tests

	$H_0 : F(P_{pre}) \leq F(P_{post})$	$H_0 : F(P_{post}) \leq F(P_{pre})$
<i>p</i> -value	0.000	1.000

Notes: The table reports *p*-values from one-sided Kolmogorov–Smirnov tests for first-order stochastic dominance (FOSD). For further details on the implementation of the tests, see [Lee and Whang \(2024\)](#).

produces a *p*-value of 1.000, indicating no evidence that the pre-announcement distribution dominates the post-announcement distribution.

Taken together, these results provide formal statistical confirmation that the distribution of transaction prices shifted significantly upward following the TSMC announcement. The post-announcement distribution first-order stochastically dominates the pre-announcement distribution, implying that the observed price increase was broad-based and not confined to specific segments of the market. This finding aligns with the visual evidence presented in Figure 3b, demonstrating that the information shock resulted in widespread gains for residential property owners across the Nanzi District.

3.2 Tracing Housing Price Dynamics

To estimate how the housing market responds dynamically to a series of information shocks, we employ a difference-in-differences (DID) event-study approach. Unlike a conventional DID specification, which typically summarizes the treatment effect with a single coefficient associated with the interaction between treatment status and post-treatment period, the event-study framework explicitly traces the evolution of treatment effects over time. This allows for capturing differential price movements between treated and untreated groups both before and after the event, providing insights into pre-treatment trends and the temporal pattern of responses following the announcement. Specifically, we estimate the following econometric model:

$$\log(p_{ijkt}) = \sum_{\substack{\tau=-q \\ \tau \neq -1}}^m \beta_{\tau} \cdot D_k \cdot \mathbf{1}(t - t_j^* = \tau) + \mathbf{x}_{ijkt} \boldsymbol{\beta} + \gamma_{jk} + \delta_t + \epsilon_{ijkt} \quad (1)$$

where $\log(p_{ijkt})$ denotes the natural logarithm of the house price for transaction i , property type j , located in township k at time t . The matrix \mathbf{x}_{ijkt} includes observable property characteristics such as age, number of rooms, and size. The primary independent variables are interaction terms between D_j , a binary indicator equal to 1 for properties in the treated district (Nanzi), and the event-time indicators, the event-time indicators $\mathbf{1}(t - t_j^* = \tau)$. The variable τ the number of months relative to the event month t_j^* , which is September 2021.

To account for unobserved heterogeneity, we include fixed effects at various levels. Specifically, γ_{jk} time-invariant differences across markets and property types, while δ_t controls for common shocks affecting all properties at a given time, such as changes in interest rates, regulatory updates, or macroeconomic conditions. The error term ϵ_{ijkt} captures unobserved factors influencing prices. We use robust standard errors clustered at the market and property-type levels to account for potential heteroskedasticity and within-cluster correlations.

The coefficients of interest, β_{τ} , capture differentials in the log prices between treated and untreated properties relative to the omitted reference month ($\tau = -1$). Estimates for β_{τ}

when $\tau < -1$ allow us to assess the parallel trends assumption by examining pre-treatment price patterns.

For empirical robustness, we initially exclude adjacent districts surrounding Nanzi to mitigate possible spillover effects that could bias estimates and use all other districts in Taiwan as the primary control group.⁵ Subsequently, we refine our control group selection by including only districts with average household incomes similar to Nanzi. Specifically, districts with average household incomes between 640 and 680 thousand NT\$ in 2019 are included, resulting in an alternative control group comprising 26 districts after excluding adjacent districts.⁶

To explore the price dynamics in responses to information shocks about TSMC’s new fab, Figure 4 plots the estimated event-study coefficients (β_τ) along with their 95% confidence intervals. The vertical lines correspond to key dates detailed in Table 1. Prior to the first news leak on September 7, 2021, regarding the confirmation of the 7nm fab location in the Nanzi District, we observe stable pre-trends, supporting the validity of our parallel trends assumption. Following this initial announcement, there is a noticeable upward trend in the price differentials, gradually reaching around 5% within three months. Although this initial reaction appears moderate, it suggests early market optimism regarding future economic benefits brought by TSMC’s planned investment.

The official confirmation by TSMC on November 9, 2021, explicitly stating the establishment of fabs for both 7nm and 28nm manufacturing processes, is associated with a substantial and immediate increase in price differentials, exceeding 10%. The pronounced jump suggests that the official announcement provided a stronger, credible signal to the market, bolstering investor confidence and significantly impacting local housing valuations.

Throughout the first half of 2022, housing prices continued to exhibit incremental growth, gradually peaking at approximately 15% above pre-announcement levels. This steady appreciation can be interpreted as a market adjustment to optimistic expectations about employ-

⁵ The TSMC shocks may affect nearby districts due to spillover effects, potentially invalidating them as suitable control areas.

⁶ Adjacent districts remain excluded to maintain consistency in our control group definition.




Figure 4: Event-study results

Notes: This figure plots estimated coefficient β_τ along with their 95% confidence intervals for Equation (1). Outcome variable is the log of sale prices. Standard errors are robust and clustered at township and building type levels.

ment opportunities, economic spillovers, and potential infrastructure improvements linked to the new TSMC facilities. However, the rising trend began to weaken in August 2022, coinciding with media reports of potential delays and logistical issues regarding the fab site.

The subsequent negative announcement on October 13, 2022, regarding TSMC’s significant capital expenditure cuts and indefinite delay of the 7nm process, had a clearly negative and immediate impact on housing market dynamics. As indicated by Figure 4, price differentials sharply decreased following this announcement. Additionally, the uncertainty introduced by the negative news was reflected in wider confidence intervals, consistent with reduced transaction volumes and heightened market volatility during periods of adverse information shocks.

Entering 2023, the announcement on April 20 of the suspension of the 28nm expansion

plans added further negative sentiment, extending downward pressure on price differentials. Market uncertainty about the future of TSMC’s investment persisted until a dramatic reversal occurred on August 8, 2023. On this date, TSMC officially confirmed plans for the construction of a highly advanced 2nm fab, substantially elevating investor expectations regarding future regional economic growth. This renewed enthusiasm is clearly reflected in the immediate and robust increase in housing prices, driving the differentials to their highest levels since the original announcement in September 2021. This strong reaction underscores the housing market’s sensitivity to major positive news, particularly involving high-tech and advanced manufacturing investments.

Overall, these dynamic responses indicate a housing market that rapidly incorporates new information but also reveals prolonged periods of adjustment consistent with the presence of significant transaction costs, speculative activity, and uncertainty regarding the full implications of major industrial announcements.

Figure 5 displays the monthly number of residential transactions in the Nanzi District from 2018 through 2023. The dashed vertical lines mark key dates corresponding to major announcements in TSMC’s fab investment timeline, beginning with the September 2021 news leak confirming the siting of a 7nm facility. Prior to this event, transaction volumes typically ranged between 300 and 400 units per month. In response to the initial news, volume surged to over 700 units in September 2021, more than double the pre-announcement average, suggesting a strong and immediate market reaction.

However, transaction volumes dropped sharply following subsequent negative developments—such as the October 2022 announcement delaying the 7nm fab and the April 2023 suspension of the 28nm expansion—falling well below historical levels. This decline may reflect reduced investor enthusiasm, market uncertainty, or tightening supply constraints. A gradual recovery is observed in 2023 and peaked at August 2023, coinciding with the confirmation that the fab would focus on the advanced 2nm process. These volume dynamics complement the observed price responses and reinforce the interpretation that the housing market reacts sensitively to both positive and negative informational shocks.




Figure 5: Monthly transaction volume

Notes:

3.3 Mechanisms

Our empirical analysis thus far has established robust evidence of significant housing market responses to informational shocks related to TSMC’s new fab announcements. In this section, we investigate whether these observed price dynamics reflect speculative behavior—defined as investors purchasing properties primarily to capitalize on expected price appreciation rather than for long-term occupancy. Identifying speculation as the driving force behind these market responses carries important economic implications, as speculative activities can exacerbate housing market cycles. Specifically, speculation tends to amplify price increases during boom periods, potentially leading to sharper price corrections and more severe economic contractions in subsequent downturns (Gao et al., 2020).

To capture speculative activity in the housing market, we use the prevalence of *low-use*

houses as an empirical proxy. This approach is motivated by Gao et al. (2020), who measure speculation using the fraction of non-owner-occupied homes among all mortgage-financed home purchases, based on the notion that such properties yield fewer housing services and are therefore more likely acquired for investment rather than residential purposes. In our context, while our data do not directly record ownership status, we instead utilize low-use houses numbers as reported by Taiwan’s Ministry of the Interior. A residential property is classified as *low-use* if its average electricity consumption does not exceed 60 kilowatt-hours per month during designated monitoring periods—May and June for the first half of the year, and November and December for the second half—based on data from Taiwan Power Company. This definition effectively identifies identifies homes that are infrequently occupied, left vacant for extended periods, or used only occasionally. Since such properties provide minimal residential utility to their owners, an elevated share of low-use houses serves as a practical and informative proxy for speculative housing demand. It reflects investor behavior driven by expected capital gains rather than immediate housing needs.

Figure 6 presents both the number of low-use houses in the Nanzi District and the corresponding vacancy rates for Nanzi and the broader Kaohsiung area over the period from 2020 to the first half of 2024.⁷ The left vertical axis shows the total number of low-use houses in Nanzi, depicted by grey bars, while the right vertical axis shows the percentage vacancy rates for both Nanzi (red line) and the rest of Kaohsiung (green line).

Before the announcement of TSMC’s fab project (marked by the dashed vertical line at the end of 2021), Nanzi exhibited a relatively stable and moderate vacancy rate, consistently lower than the Kaohsiung average. The number of low-use houses also remained steady, fluctuating slightly around 6,500 to 6,800 units. However, in the first half of 2022, we observe a sharp upward trend in both the number of low-use houses and the vacancy rate in Nanzi.

This increase is particularly notable in 2022h1 and 2023h1, with the number of low-use houses rising by nearly 2,000 units over two years, and the vacancy rate in Nanzi converging

⁷ Vacancy rates refer to the share of residential units classified as low-use, calculated by dividing the number of low-use houses by the total number of tax-eligible residential properties.




Figure 6: House Vacancy

Notes: This figure plots the number of low-use houses in Nanzi as well as the semi-annually vacancy rates from 2020 to 2022.

with and eventually surpassing the Kaohsiung average. By 2024h1, the vacancy rate in Nanzi reaches its highest point in the sample period, exceeding 11%, compared to around 10% in Kaohsiung overall. These patterns are not observed prior to the announcement and therefore do not appear to be part of an ongoing trend.

The post-announcement divergence in vacancy behavior between Nanzi and the rest of Kaohsiung is consistent with speculative housing activity. If the increase in transactions were primarily driven by genuine residential demand—such as households relocating to the area in anticipation of job creation associated with TSMC’s investment—we would expect vacancy rates to decline as newly purchased units become occupied. Instead, the data reveal a rising share of properties remaining unused, suggesting that many recent purchases are being held as investment assets rather than for immediate residential use. This pattern implies that buyers are motivated by expectations of future capital gains rather than the intention to reside in the properties.

The empirical pattern observed in Figure 6 reinforces the interpretation that speculative motives—rather than shifts in fundamental residential demand—play a central role in shaping post-announcement housing dynamics in Nanzi. However, this analysis is descriptive and does not establish a causal relationship. In particular, changes in vacancy rates may also be driven by an increase in the supply of newly built housing units that have yet to be occupied. Such additions to the housing stock can mechanically raise vacancy rates, even in the absence of speculative activity.

To more rigorously test the speculation hypothesis, we adopt a difference-in-differences (DID) event-study framework similar to the one used in our earlier price analysis. In this specification, we include controls for the total number of tax-eligible residential properties to account for changes in housing supply over time. This adjustment allows us to distinguish between increased vacancy arising from speculative holding behavior and that driven by structural growth in available housing units, thereby providing a cleaner estimate of the effect of informational shocks on speculative market responses. Specifically, we estimate the following econometric model:

$$\text{vacancy rate}_{jt} = \sum_{\substack{\tau=-4 \\ \tau \neq -1}}^4 \eta_{\tau} \cdot D_k \cdot \mathbf{1}(t - t_j^* = \tau) + \alpha \log(\text{total houses})_{jt} + \theta_k + \delta_t + \mu_{jt} \quad (2)$$

where vacancy rate_{jt} denotes the proportion of low-use residential properties in township j at time t , measured semiannually. The main explanatory variables are a set of event-time dummies, $\mathbf{1}(t - t_j^* = \tau)$, interacted with a treatment indicator D_k , which equals one for the treated unit, Nanzi District, and zero otherwise. These interaction terms allow us to estimate the dynamic treatment effects η_{τ} relative to the omitted reference period, $\tau = -1$.

We also include $\log(\text{total houses})_{jt}$ to control for changes in the size of the housing stock over time, as increases in vacancy rates may be mechanically driven by the supply of newly constructed units that have yet to be occupied. Township fixed effects θ_k absorb time-invariant unobserved heterogeneity across locations, while time fixed effects δ_t control for

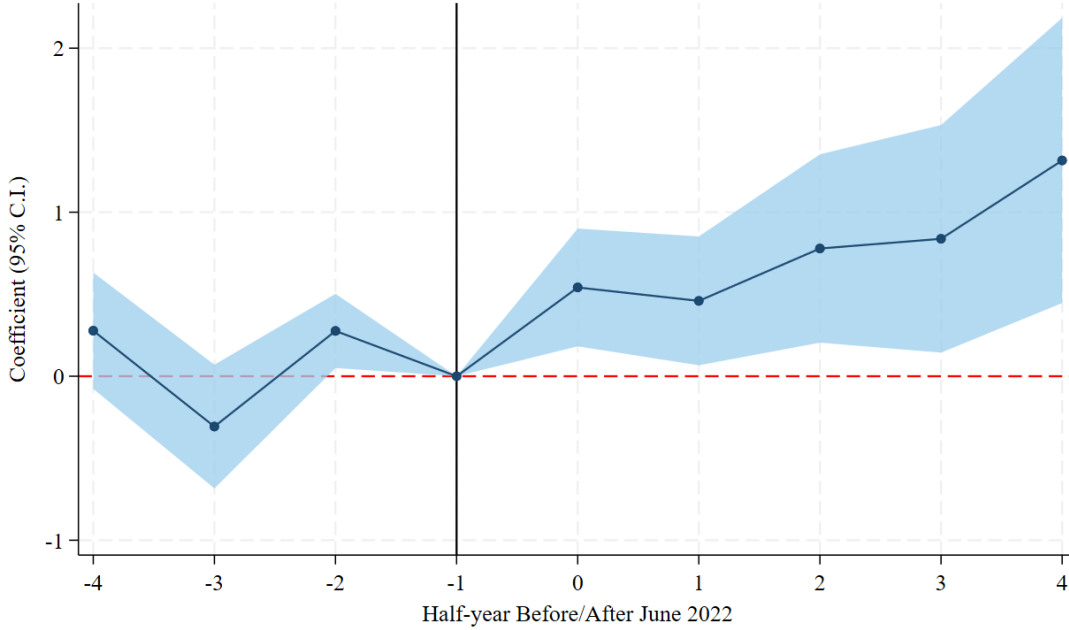


Figure 7: Changes in Vacancy Rate

Notes: This figure plots the estimated coefficients η_τ and their 95% confidence intervals from Equation (2), where the outcome variable is the vacancy rate (in percentage points) measured at the township level. Each time period corresponds to a six-month interval. Standard errors are robust and clustered at the township level. The control group consists of all districts within Kaohsiung City, excluding those adjacent to Nanzi District in order to mitigate potential spillover effects. The average vacancy rate in Nanzi during the sample period is 9.1%.

common shocks affecting all townships in a given period. In the estimation, the control group consists of all districts within Kaohsiung City, excluding those adjacent to Nanzi District in order to mitigate potential spillover effects.⁸ The error term μ_{jt} captures unobserved factors influencing vacancy rates. Standard errors are robust and clustered at the township level to account for heteroskedasticity and serial correlation.

Figure 7 presents the estimated event-time coefficients η_τ and their 95% confidence intervals from Equation (2). The horizontal axis measures event time in half-year intervals relative to June 2022, which marks the midpoint of the post-announcement period following

⁸ We focus exclusively on Kaohsiung City because local governments have the authority to influence long-term housing supply through housing taxes, zoning regulations, and land use policies, which can have broader implications for regional housing dynamics (Hsieh and Moretti, 2019; D’Amico et al., 2024; Bartik et al., 2025).

TSMC’s major investment updates. The vertical black line at $\tau = -1$ denotes the omitted reference period, while the red dashed line at zero indicates the baseline against which all treatment effects are measured.

The estimates show that in the baseline periods, the coefficients are statistically indistinguishable from zero, supporting the parallel trends assumption. Beginning in the post-treatment period, the coefficients increase steadily, with statistically significant effects emerging at $\tau = 0$ and growing in magnitude over time. By $\tau = 4$, the estimated effect exceeds 1.5 percentage points. These results indicate a sustained rise in the vacancy rate in Nanzi district relative to the control group, after accounting for changes in housing stock. This pattern is consistent with speculative behavior, in which properties are acquired for investment rather than immediate use.

4 Concluding Remarks

This paper offers novel evidence on the informational efficiency and behavioral responses of the housing market using a unique natural experiment—TSMC’s phased announcements regarding its new fabrication plant in Kaohsiung. Leveraging high-frequency housing transaction data, we show that prices reacted quickly to early news reports and official announcements, with significant appreciation concentrated within the first few weeks. Price adjustments continued in response to subsequent updates, with observed under-reaction patterns suggesting that information diffusion was gradual and shaped by evolving market beliefs.

Our analysis also highlights the speculative component of the housing market’s response. Transaction volumes spiked immediately following the announcement, and the number of low-use houses and vacancy rates rose sharply in the treated district. These dynamics are consistent with investor-driven behavior rather than shifts in end-user demand. By incorporating a difference-in-differences event-study design and controlling for housing supply, we provide a clean estimate of the impact of informational shocks on speculative behavior.

Our findings contribute to the literature on housing market efficiency, investor behavior, and local economic impacts of industrial policy. In particular, we demonstrate that

high-tech manufacturing investments—while potentially introducing negative externalities typical of industrial facilities—can generate positive valuation effects in surrounding housing markets. Moreover, our evidence on speculation underscores the importance of monitoring behavioral responses to large-scale development projects, particularly when such projects generate information shocks that ripple through asset markets. Future work could extend this framework to investigate policy tools that mitigate speculative excess while preserving the broader economic benefits of strategic industrial investments.

Reference

- Anundsen, A. K. and Røed Larsen, E. (2018). Testing for micro-efficiency in the housing market. *International Economic Review*, 59(4):2133–2162.
- Armstrong, W. J., Cardella, L., and Sabah, N. (2021). Information shocks, disagreement, and drift. *Journal of Financial Economics*, 140(3):916–940.
- Barrett, G. F. and Donald, S. G. (2003). Consistent tests for stochastic dominance. *Econometrica*, 71(1):71–104.
- Bartik, A., Gupta, A., and Milo, D. (2025). The costs of housing regulation: Evidence from generative regulatory measurement. *Available at SSRN 4627587*.
- Ben-Shahar, D. and Golan, R. (2019). Improved information shock and price dispersion: A natural experiment in the housing market. *Journal of Urban Economics*, 112:70–84.
- Beracha, E., He, Z., Wintoki, M. B., and Xi, Y. (2022). On the relation between innovation and housing prices—a metro level analysis of the us market. *The Journal of Real Estate Finance and Economics*, pages 1–27.
- Bernstein, A., Gustafson, M. T., and Lewis, R. (2019). Disaster on the horizon: The price effect of sea level rise. *Journal of financial economics*, 134(2):253–272.
- Broxterman, D. and Zhou, T. (2023). Information frictions in real estate markets: Recent evidence and issues. *The Journal of Real Estate Finance and Economics*, 66(2):203–298.
- Busse, J. A. and Green, T. C. (2002). Market efficiency in real time. *Journal of Financial Economics*, 65(3):415–437.
- Chan, W. S. (2003). Stock price reaction to news and no-news: drift and reversal after headlines. *Journal of financial economics*, 70(2):223–260.
- Chang, S., Yang, T.-T., and Zhu, J.-D. (2024). The effect of capital gains taxation on housing sales: Evidence from taiwan. *Working Paper*.

- Chen, Y., Wilkoff, S., and Yoshida, J. (2024). Amazon is coming to town: Sequential information revelation in the housing market. *Real Estate Economics*, 52(2):277–323.
- Cheng, T.-C. F., Wang, T.-C., and Zhu, J.-D. (2022). The effect of risk information on housing prices in taiwan. *Academia Economic Papers*, 50(1):41–73.
- D’Amico, L., Glaeser, E. L., Gyourko, J., Kerr, W. R., and Ponzetto, G. A. (2024). Why has construction productivity stagnated? the role of land-use regulation. *Available at SSRN 4679195*.
- De Vor, F. and De Groot, H. L. (2011). The impact of industrial sites on residential property values: A hedonic pricing analysis from the netherlands. *Regional Studies*, 45(5):609–623.
- Famiglietti, M., Garriga, C., and Miravete, E. (2023). Amazon hq2: A tale of two shocks to housing price expectations. *Working Paper*.
- Gao, Z., Sockin, M., and Xiong, W. (2020). Economic consequences of housing speculation. *The Review of Financial Studies*, 33(11):5248–5287.
- Goldberg, P., Juhász, R., Lane, N., Lo Forte, G., and Thurk, J. (2024). Industrial policy in the global semiconductor sector. *NBER Working Paper*, (w32651).
- Hjalmarsson, E. and Hjalmarsson, R. (2009). Efficiency in housing markets: Which home buyers know how to discount? *Journal of Banking & Finance*, 33(11):2150–2163.
- Hsieh, C.-T. and Moretti, E. (2019). Housing constraints and spatial misallocation. *American economic journal: macroeconomics*, 11(2):1–39.
- Hu, M., Tsang, D., and Wan, W. X. (2025). Corporate relocation and housing market spillovers. *Management Science*, 71(5):4344–4376.
- Larsen, E. R. and Weum, S. (2008). Testing the efficiency of the norwegian housing market. *Journal of Urban Economics*, 64(2):510–517.

- Lee, K. and Whang, Y.-J. (2024). Pysdtest: a python/stata package for stochastic dominance tests. *arXiv preprint arXiv:2307.10694*.
- Miao, W. (2024). Technology rivalry and resilience under trade disruptions: The case of semiconductor foundries. *Working Paper*.
- Pope, D. G. and Pope, J. C. (2015). When walmart comes to town: Always low housing prices? always? *Journal of Urban Economics*, 87:1–13.
- Pope, J. C. (2008). Buyer information and the hedonic: The impact of a seller disclosure on the implicit price for airport noise. *Journal of Urban Economics*, 63(2):498–516.
- Thurk, J. (2024). Outsourcing, firm innovation, and industry dynamics in the production of semiconductors. *Working Paper*.
- Votsis, A. and Perrels, A. (2016). Housing prices and the public disclosure of flood risk: a difference-in-differences analysis in finland. *The Journal of Real Estate Finance and Economics*, 53:450–471.
- Yi, D. and Choi, H. (2020). Housing market response to new flood risk information and the impact on poor tenant. *The Journal of Real Estate Finance and Economics*, 61:55–79.

Appendix

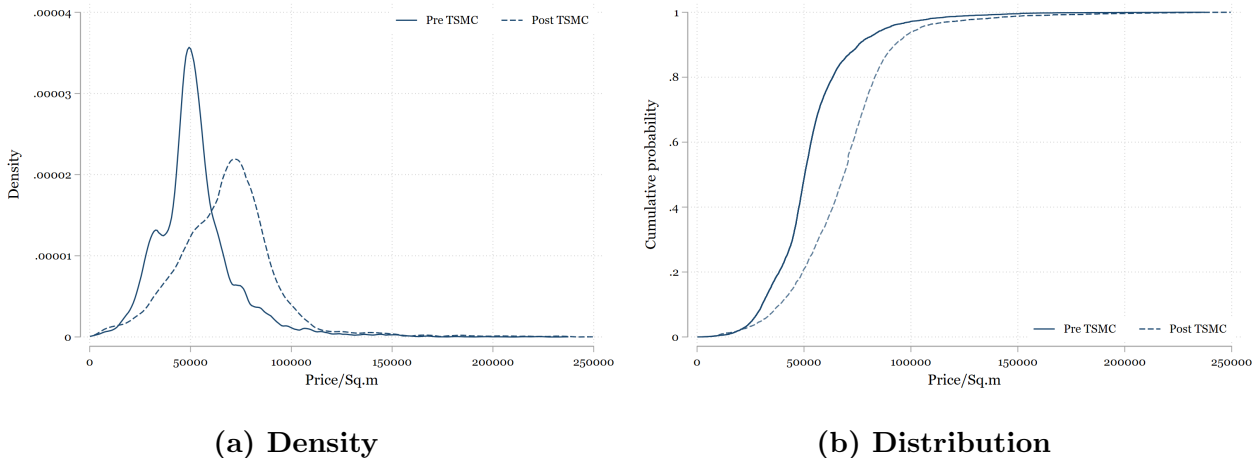


Figure A1: Long-term Pre/Post TSMC Price Distribution

Figure A1 presents analogous results to those in Figure 3, but based on an extended sample period from 2020 to 2023. The long-term comparison reveals a clear rightward shift in the distribution of residential transaction prices, indicating price increases across all quantiles.

Figure A2 show the binscatter plots for average monthly price of pre-owned and pre-sale housing transactions in Nanzi District. The average prices in pre-sale market is higher than those in pre-owned market because of higher construction costs, future expectations, or lack of wear and tear. Using September 2021 as the separating month, we find descriptive evidence that the average prices exhibit structural changes with increases in slopes and jumps at the event month.

Figure A3 shows differentials using alternative control groups with similar average household income. The evolution of point estimates are similar to those in Figure 4 despite the price differentials is over 20% at 2023-08.

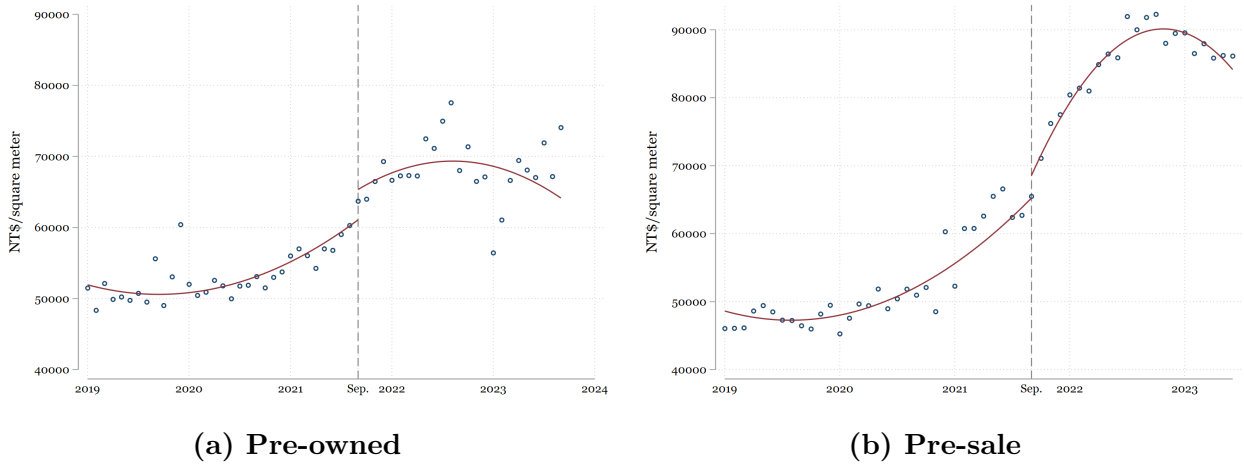


Figure A2: Average monthly price




Figure A3: Event-study results with alternative control group

Notes: This figure plots estimated coefficient β_τ along with their 95% confidence intervals for Equation (1). Outcome variable is the log of sale prices.

Table A1: Summary Statistics

Variable	Treated		Control		Difference	
	Mean	SD	Mean	SD	Diff	<i>p</i> -value
vacancy rate(%)	9.11	0.87	12.35	4.64	-3.23	0.04
log of total houses	11.30	0.03	9.37	1.40	1.93	0.00