Land development and frictions to housing supply over the business cycle*

Hyunseung Oh† Choongryul Yang‡ Chamna Yoon§

February 2, 2024

Abstract

Using a novel data set of U.S. residential land developments, we document that the average time to develop residential properties—which includes both the time spent preparing land infrastructures and construction—is about three years, consistent with sizable lags in housing investment projects. We show that the time to develop is highly dispersed across locations, a finding that helps quantify the housing supply elasticity that is relevant for assessing local housing variations over the business cycle. We also show that incorporating long and dispersed time to develop into an otherwise standard housing investment model helps rationalize some empirical facts on the housing market. Our model implies that policies to boost housing supply are less effective in immediately stabilizing house prices for regions where land development takes a long time.

Keywords: Housing supply; house price dynamics; residential investment.
JEL Classification Numbers: E22, E32, R31.

*We thank Harun Alp, Etienne Gagnon, Daniel Garcia, Greg Howard, Matteo Iacoviello, Ryan Kim, Emi Nakamura, Joseph Nichols, Jón Steinsson, Luis Torres, and numerous seminar and conference participants for helpful comments and suggestions. The views expressed are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System.
†Federal Reserve Board. E-mail: hyunseung.oh@frb.gov.
‡Federal Reserve Board. E-mail: choongryul.yang@frb.gov.
§Seoul National University. E-mail: chamna.yoon@snu.ac.kr.

1 Introduction

Researchers have argued that significant adjustment costs are needed in standard investment models to account for the empirically slower response of investment to economic shocks. Because ships and factories cannot be built in a day, these adjustment costs are typically motivated as stand-ins for the time it takes to produce new capital and the difficulties in adjusting investment plans once they are in train. A strand of literature on the housing market also highlights inelastic short-run housing supply as a sensible feature to explain the dynamic properties of the housing market, such as the difference in the short- and long-run housing market reactions to COVID-19 (Howard, Liebersohn and Ozimek, 2023).[1] Quantitatively, little is known, however, about how long it takes to develop land and what that implies on the adjustment of housing supply over time.

In this paper, we document the time it takes to build a house from undeveloped piece of land across regions using a new data set and study its aggregate and cross-regional implications using a model of investment dynamics. A desirable feature of our data set is that we observe the time it takes not only to build structures on a developed lot but also to develop land infrastructure on a vacant, undeveloped piece of land. Accordingly, we are able to analyze the *comprehensive* process of residential construction across major U.S. regions.

Empirically, we document two stylized facts on land development. First, residential land development is indeed a lengthy process that takes more than three years, on average, after receiving a preliminary approval of the site plan from the local government, including more than a year, on average, to develop raw land with a subdivision map into a lot on which structures could be built. Second, the time it takes to develop land is highly dispersed across locations, even after controlling for an extensive list of variables that are likely to affect local construction demand. In turn, we find that a county's median time to develop (TTD) is associated not only with a measure of its long-run housing supply elasticity, but also with adverse local weather conditions that hinder construction activity in the short run. This information suggests that local factors that determine TTD are not fully aligned with the local factors that determine housing supply in the long run.

We then use the local variations in TTD to quantify the local housing supply elasticity, a measure that has taken center stage in the macro-housing literature. As housing wealth is about one-third of the total net worth of U.S. households, house price changes have significant spillover effects to the broader economy and estimates of the housing supply elasticity are

[1] Many models that account for residential investment and house price dynamics rely on the assumption of fixed land supply—for example, Davis and Heathcote (2005); Kiyotaki, Michaelides and Nikolov (2011); and Kaplan, Mitman and Violante (2020).

frequently used to decipher the causal effect of house price changes to economic activity.[2] While the existing estimates of the housing supply elasticity mainly focus on the long-run determinants of housing supply, we show that TTD helps quantify the housing supply elasticity more relevant for the evolution of housing supply within the next five years. Towards that, we elaborate a housing investment model with TTD and derive analytical expressions that relate the short-run (up to five years) housing supply elasticity to TTD and the long-run housing supply elasticity. Combining the model and the data, we find that short-run housing supply elasticities vary significantly across counties and are indeed smaller than, and distinct from, corresponding long-run elasticities.

We show that the consideration of relatively long and dispersed TTD in an otherwise standard local general equilibrium model of housing investment helps rationalize a number of empirical facts. First, when the TTD friction is included, our otherwise standard model no longer predicts a counterfactual, strongly negative correlation between a region's relative price and relative quantity in response to a common demand shock across regions. Intuitively, as lengthy TTD lowers the short-run supply elasticity in all regions, developers in the region with a higher long-run supply elasticity are more likely to substitute short-run supply with long-run supply as they internalize the larger gap between the short- and long-run supply elasticities. Second, when the autocorrelation of the common demand shock is not highly persistent, our model predicts that a region's house price growth is better explained by the short-run supply elasticity than by the long-run supply elasticity. Using county-level data, we regress a county's house price change (relative to the national house price change) on the short- to long-run housing supply elasticities. Since the early 2010s, we find that our short-run supply elasticity outperforms the long-run supply elasticity in accounting for the observed cross-county variation in house prices.

Finally, we draw a policy implication by conducting a counterfactual exercise where the government aims to stabilize house prices by a discretionary housing supply policy. When TTD is present, we find that government incentives to boost the housing supply affect house prices through the expectations channel of future housing supply. Therefore, the policy could be somewhat less effective in immediately stabilizing house prices for regions where land development takes a long time, but it could be more effective in stabilizing medium-run house prices for those regions.

[2] According to the 2022 financial accounts of the United States from the Federal Reserve Board's flow of funds statistics, the total real estate at market value for households and nonprofit organizations is 47.1 trillion dollars and the net worth of households and nonprofit organizations is 143.7 trillion dollars.

Related literature. Studies on housing supply mainly focus on estimating its long-run determinants (Saiz, 2010; Baum-Snow and Han, 2022; Lutz and Sand, 2022), and these estimates are typically used to identify regional variations to economic shocks (Mian, Rao and Sufi, 2013; Mian and Sufi, 2014; Davis and Haltiwanger, 2019; Bhattarai, Schwartzman and Yang, 2021). This approach could be problematic when the economic shock of interest does not persist in the long run and when the short-run determinants of housing supply differ from the long-run estimates. For example, Guren, McKay, Nakamura and Steinsson (2020) suggest that a puzzling feature of the cross-regional housing price and quantity correlation discussed in Davidoff (2016) could be potentially resolved by assuming a lower short-run housing supply elasticity in all regions. Not much is known, however, about the determinants of housing supply elasticity in a business cycle frequency, with the notable exception of Topel and Rosen (1988), who estimate the short- and long-run elasticities of housing supply and find that most of the long-run response occurs within a year. We contribute to this literature by (i) quantifying frictions to housing supply at a business cycle frequency based on the observed duration of land development and (ii) studying the implications of our quantified frictions on the housing market through an equilibrium model of housing investment. By investigating the link between new housing supply data and the elasticity of housing supply, we also complement the literature that studies the sensitivity of local economic activity to house prices (Guren, McKay, Nakamura and Steinsson, 2021; Graham and Makridis, 2023). Research in this topic identifies plausibly exogenous house price variations by focusing on variables of the local economy or existing housing characteristics; we argue that data on the timing of new housing supply also capture an important source of local variation in house prices that could be used to estimate the sensitivity of local economic activity to house prices.

In the literature on business cycles, time to build has been noted as a key friction to investment dynamics at least since Kydland and Prescott (1982). Subsequently, several papers document the duration of building capital using newly available data or study its implications on investment behavior (Lucca, 2007; Kalouptsidi, 2014; Sarte, Schwartzman and Lubik, 2015; Millar, Oliner and Sichel, 2016; Kydland, Rupert and Šustek, 2016; Oh and Yoon, 2020; Meier, 2020; Charoenwong, Kiruma, Kwan and Tan, 2024). Our paper contributes to this line of work, both by providing new stylized facts on the *comprehensive* construction process from undeveloped land to the completion of new structures and by delivering a number of housing market implications of the new stylized facts. Relatedly, our work contributes to existing studies of housing investment (Mayer and Somerville, 2000; Green, Malpezzi and Mayo, 2005; Haughwout, Peach, Sporn and Tracy, 2013; Paciorek, 2013; Murphy, 2018; Nathanson and Zwick, 2018). More broadly, our findings could also shed light on the importance of

TTD frictions for nonresidential structures as both residential and nonresidential structures are likely to share some common hurdles from the site development stage.

As discussed earlier, the implications of these residential construction dynamics are less explored in the housing and macro literature. Most models that study the aggregate implications of the housing sector abstract from the dynamic aspect of land development (Davis and Heathcote, 2005; Iacoviello and Neri, 2010). Following the spirit of Glaeser, Gyourko and Saiz (2008), we explore the aggregate and regional implications of housing supply with a focus on the short-run dynamics.

Structure of the paper. In section 2, we present the land development data and estimate a TTD measure that is comparable across regions by controlling for various factors. In section 3, we develop a TTD model of housing investment and analytically derive housing supply elasticities in each horizon. In section 4, we quantify local housing supply elasticities in the short to medium run by using the model derivations and the empirical TTD estimates. In section 5, we study three implications of our model for house prices and housing quantity. Section 6 concludes. The Online Appendix provides additional details and sensitivity analyses of our theoretical and empirical results.

2 Duration of land development across regions

In this section, we use a unique data set that tracks development activities for residential properties in the U.S. to measure the duration of land development across regions. The desirable feature of our data set is that it includes the period of site development prior to building construction.

We show that the duration of land development that includes the period of site development is lengthy and varies widely across regions. These variations persist even after controlling for a number of observable regional demand factors.

2.1 Land development data and summary statistics

Our data set comes from Zonda, which provides data and analysis to the national residential home-building industry. The data set is constructed from Zonda's survey markets data, which cover many of the major metro areas with high residential construction activity in the U.S. The survey markets data put together a quarterly construction status survey in new home subdivisions, an area containing a large number of houses or apartments to be built close together at

Table 1: New housing completions between 2003 and 2019

(unit: $1,000$ housing)	Zonda	Census	Coverage
Total housing	$7,790$	$20,020$	39%
Single-family housing	$5,939$	$15,314$	39%

the same time. Large subdivisions are often broken down to multiple sections, each of which is typically built by a single-builder company. The data set displays the total number of housing units as well as other construction characteristics by sections. We have access to this data set from 2000 to 2021.

As shown in Table 1, our data set includes a large number of new housing supply across the U.S. Between 2003 and 2019, the data set contains 222,868 developed sections with a total of 7.8 million units of new housing. For reference, the Census Bureau reports a sum of 20 million new housing completions in the same period, implying a 39 percent Census coverage ratio of our data set. Our data set is not biased toward multi-unit housing development, as the Census coverage of single-unit housing completion is also around 39 percent.[3] These completions are distributed over 318 counties in 113 core-based statistical areas (CBSAs) that represent 48 percent of the U.S. population. The average population size of those CBSAs is 1,590,428, which is 4.7 times larger than that of the U.S. average CBSA.[4]

Besides the high coverage ratio, a desirable feature of the data set is that it continuously tracks the construction status of subdivisions and sections. Land development is a lengthy process, starting from the acquisition of land by developers and the design of a development plan. The plan is then submitted to the appropriate municipality for a preliminary review. The profile of the subdivision is first created and labeled as *future lots* in our data set when the municipality grants a preliminary approval as a first step in the process or, if the approval date is not available, after Zonda reviews and verifies the site plan submitted to the municipality. During each quarterly survey, the lot remains as future lots while there is ongoing land development, such as the site having survey stakes or equipment. It follows that the final site plan is submitted and approved, and the necessary permits are processed. Thereafter, the infrastructure for the land is developed and the lot is now labeled as *active*. At this stage, separate homebuilders enter for construction projects in the active lots not pursued by the developer.

[3] In our data set, single-family housing units comprise 82.6 percent of the completions, followed by townhouses (10.2 percent), condos (2.3 percent), and duplexes (1.2 percent). We show in the Appendix that the Census coverage does not significantly fluctuate across years.

[4] Of the top 20 CBSAs ranked by the 2020 Census population, only 2 CBSAs (Boston and Seattle) are not included in our data set.

When there is excavation activity with a slab or basement on these vacant developed lots, the units are classified as *under construction*, consistent with the Census Bureau's definition of housing starts. After the completion of home construction, each house is classified as either a *finished vacant* unit or an *occupied* unit, depending on its status of sales. Eventually, the subdivision is classified as *built out*, and it exits the data set when the number of occupied units equals its total units.

Based on this classification, we define a TTD measure for each new development section. TTD is defined as the duration between the quarter when the municipality likely approves a preliminary site plan and the quarter when the number of finished units (vacant or occupied) reaches at least half of the total number of units. The unique feature of our data set is that it captures the earliest stage of a completed development with a plan that is as concrete as a preliminary map submitted to the municipality. Our definition for the beginning quarter of TTD fully takes this feature on board; in the Appendix, we present results using an alternative definition for the beginning quarter of TTD.[5] The definition for the end quarter of TTD is consistent with the Census Bureau's definition of the completion of a multi-unit building, as it classifies the construction of a multi-unit building as complete when half of the units are finished. It is worth noting that the Census Bureau tracks construction time per building, whereas we can only track construction time per section. Therefore, our measure of the section's construction time could be longer than the construction time of an average building in that section if a developer decides to build structures sequentially rather than simultaneously. In the Appendix, we study the sensitivity of our empirical results when the end quarter of TTD is defined earlier than our baseline—that is, the date at which the number of finished units reaches a quarter of the total number of units in the section.

For the remaining analysis, we adopt the following sample selection criterion. Between 2003 and 2019, 222,868 sections were completed. We dropped 102,575 observations without information on TTD (for example, missing start dates), resulting in 120,293 observations.[6] We further dropped 16,097 observations without lot size information or demand controls, leaving us with 104,196 observations.

[5]Our baseline definition is also driven by data availability. In the data set, the preliminary approval date is missing for many sections, which limits our sample size substantially. In the Appendix, we provide details on how we measure the beginning quarter of development.

[6]Specifically, we define the start date as the first quarter when the total number of planned units is equal to the total number of future lots, based on Zonda's quarterly review of newly submitted maps at the municipality. We find that our defined start date is close to the municipality approval date of the preliminary site plan, when the latter date is available in the data set. We dropped sections where their first observation already had positive active lots, as land development on these sections likely started (according to our definition) before they entered the data set.

Table 2: Section TTD statistics

(unit: days)	Site TTD	Building TTD	Total TTD
Mean	569	760	$1,329$
Std. dev.	760	792	$1,077$
IQR	458	548	$1,006$
P10	91	182	366
P25	181	275	638
P50	275	458	$1,004$
P75	639	823	$1,644$
P90	$1,278$	$1,734$	$2,922$
Observations	$104,196$	$104,196$	$104,196$

Note: Each observation is a subdivision or a section of a subdivision when there are multiple sections in a subdivision. IQR stands for the interquartile range (P75−P25). Five different percentiles of each TTD distribution are shown—for example, P50 referring to the median (50th percentile) of the distribution.

2.2 Stylized facts on the duration of housing development

The total time it takes for housing development (TTD) comprises two parts: time to develop infrastructure at the site (site TTD) and time spent on the construction of buildings (building TTD). Just by looking into the raw measures of TTD based on Table 2, we find two stylized facts on the duration of housing development that stand out.

First, housing development is a lengthy process with significant time spent on land development. As shown in the first row of the table, housing development takes a total of $1,329$ days on average. While less emphasized in the literature because of limited data availability, we find that the duration between the land development plan approval and the finishing of site development is substantially long, averaging 569 days. The mean construction time of buildings on these developed sites is 760 days.

Second, there is substantial heterogeneity in the duration of housing development. The standard deviation and the interquartile range of total TTD are both around three years ($1,077$ days and $1,006$ days, respectively). The significant heterogeneity is also pronounced in site TTD, as its standard deviation is more than two years.

Note that the distribution of TTD is skewed to the right, as the mean is larger than the median in all TTD measures. This finding is also evident from the lengthy TTD at the 90th percentile.

2.3 Controlled measures of TTD

The lengthy and highly dispersed TTD across sections documented above could be driven by various factors. Our goal is to measure the developers' TTD constraint for housing development that is comparable across locations. Toward that goal, we need to first control for differences in construction characteristics that are likely to affect TTD. For each development section, the data set includes some of that information, such as the number of housing units, the average lot size, the type of housing, and the builder(s) of each section. Indeed, these construction characteristics have substantial variations. For example, there are an average of 42 housing units per each section, and the standard deviation is also 42 units.

The first column of Table 3 shows the regression result when the log of TTD is regressed on several construction characteristics in our sample. Builder fixed effects are included for the top 15 builders in our sample. Completion year fixed effects are also included to abstract from time variations in TTD. We find that both the number of units and the (average) lot size per housing unit are positively associated with TTD. One percent increases in the number of units and in the lot size per unit imply 0.131 percent and 0.136 percent increases in TTD, respectively. These results are highly significant and consistent with the findings in Oh and Yoon (2020), where the square footage of a single-family house under construction is shown to be positively associated with its time to build. In terms of the type of housing, townhouses and condos take a longer time to complete relative to single-family developments.

Even after we control for construction characteristics, TTD could also be driven by local economic factors that are linked to the developers' incentives in those locations. Because our model does not feature the developer's location choice, we would also need to control for these factors. The second column of Table 3 adds a number of local controls potentially relevant for housing supply—such as a Bartik-type variable that measures the local demand pressure, indicators of sand and coastal states, population shares of immigrants and college-educated adults, and population density—as suggested in Davidoff (2016). We also include the annual county-level GDP in the regression to control for any country-level time-varying economic factors. The regression results show that several of these local economic factors are associated with TTD in a statistically significant manner. Counties with higher Bartik demand pressure, immigrant share, and population density, or counties not in sand states and especially in coastal states, experience longer duration in development, as shown in the Appendix. Even after we control for these local economic factors, however, the R-squared shows limited improvement over the regression in the first column, and the regression coefficients for construction characteristics remain robust. These results suggest that local economic fac-

Table 3: Section TTD regression results

Variables	(1)	(2)	(3)
Log(number of units)	0.131***	0.140***	0.138***
	(0.00427)	(0.00436)	(0.00417)
Log(lot size)	0.136***	0.141***	0.140***
	(0.00480)	(0.00492)	(0.00477)
Single family	–	–	–
Townhouse	0.208***	0.198***	0.187***
	(0.0111)	(0.0113)	(0.0112)
Condo	0.191***	0.227***	0.235***
	(0.0417)	(0.0414)	(0.0384)
Duplex	0.0436	0.0341	0.0396
	(0.0309)	(0.0310)	(0.0293)
Etc.	0.0233	0.0461*	0.0444*
	(0.0263)	(0.0261)	(0.0246)
Builder fixed effect	✓	✓	✓
Year fixed effect	✓	✓	✓
Local controls		✓	
Local controls × Year			✓
Constant	4.505***	4.432***	5.284***
	(0.0505)	(0.0883)	(0.184)
Observations	104,196	104,196	104,196
R-squared	0.282	0.289	0.330

Note: Regression with log(TTD) as the dependent variable. Local control variables include Bartik-type predicted industry employment growth, indicators for sand state and coastal state, population share of immigrants, population share of college educated, population density, and county real GDP. Robust standard errors are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

tors might play a limited independent role after taking into account the developer's choice of construction characteristics.

The third column of Table 3 allows for additional flexibility in the time-varying response of TTD to local economic factors by interacting the time-invariant local controls with year fixed effects. While the R-squared moderately improves to 33 percent, the regression coefficients on construction characteristics remain relatively robust across all three specifications.

Table 4: County-level TTD statistics

(unit: days)	Raw TTD	Reg. (1)	Reg. (2)	Reg. (3)
Mean	$1,044$	973	969	971
Std. dev.	426	317	317	287
IQR	365	396	371	319
P10	640	605	605	631
P25	822	764	784	803
P50	$1,005$	954	959	964
P75	$1,187$	$1,160$	$1,155$	$1,123$
P90	$1,369$	$1,369$	$1,358$	$1,348$
Observations	267	267	267	267

Note: Each observation is a county's median TTD. We use counties with at least 10 completed sections observed. IQR stands for the interquartile range (P75−P25). Five different percentiles of each TTD distribution are shown—for example, P50 referring to the median (50th percentile) of the distribution.

2.4 County-level TTD

Using the regression results in Table 3, we now construct the county-level TTD for the representative housing. That is, we normalize TTD for each section by assuming that the controlled observables are at their national average values and add the fitted residuals. We then take the median of this value for each county of interest. Note that we use the median instead of the mean, as the distribution of TTD is skewed to the right. Moreover, the median TTD should be relatively insensitive to any remaining time-varying demand factors of TTD, such as the fat right tail in construction time during the Great Recession (Oh and Yoon, 2020).

Table 4 presents some cross-county moments of each county's measure of TTD. In the first column ("Raw TTD"), we observe that the average of the county-level TTD using raw section TTD data is almost the same as the average of the section TTD itself. The standard deviation and the interquartile range are sizable at 426 days and 365 days, respectively. The second to fourth columns present the same statistics using the controlled TTD estimates in Table 3. Controlling for construction characteristics and focusing on a nationally representative housing development, we find that the cross-county mean TTD is 973 days, just less than three years. The standard deviation and the interquartile range are also sizable at around one year. Results for the third and fourth columns are similar to those of the second column, consistent with the R-squared results in Table 3, suggesting that after one controls for construction characteristics, the marginal contribution of local demand factors is limited.

2.5 The geographic determinants of the land development process

Land development is a major topic of interest in civil engineering, as each construction site poses unique engineering challenges based on soil characteristics, topography, weather, and other physical features (Kone, 2006). As such, developers create not only a master plan design that conceptualizes their new development at the location of interest, but also a site engineering plan that adapts the master plan design to the physical properties of the site. These site engineering plans include (i) a grading plan that shows the elevations of grounds and buildings, (ii) a storm water management plan that shows the volume and rate of storm water runoff, and (iii) an erosion and sediment control plan that shows the erosion control barriers and materials at the site. The local government regulation on development varies based on its transitory and permanent environmental effects, and this factor also plays an important role in shaping the engineering plans of each site.

As shown in Table 4, our county-level TTD measures exhibit significant degrees of cross-regional variation even after controlling for construction characteristics and local economic factors. Based on the described process of land development, we then ask whether any observable geographic differences in engineering challenges could account for that variation. Some of the geographic differences could be highly correlated with factors that determine long-run housing supply described in Saiz (2010), but other geographic differences that do not play a large role in accounting for the long-run housing supply could nevertheless matter for the duration of land development. For example, each location is exposed to different weather conditions that might not materially affect the decision to develop land but might still matter for the cost and duration of land development. That is, locations with extreme storm and heat conditions could still be desirable for new construction, but severe weather will occasionally delay on-site construction activity and the developer's building design to tolerate extreme weather conditions might further lengthen the development process.

Table 5 displays the regression result using our controlled TTD measures and the observed geographic determinants described in this section. Both the long-run housing supply elasticity and rainfall intensity affect TTD in a statistically significant sense. That is, the duration of land development is longer for locations where (i) the long-run housing supply is limited and (ii) rainfalls are more intense, more frequent, or both. Heat also delays TTD, but its statistical significance is less pronounced. A careful study of these engineering challenges is beyond the scope of this paper, but we think that the significance of these geographic determinants in accounting for the variation in TTD across locations suggests the potential of our controlled TTD measure to reflect plausibly exogenous variations to economic shocks.

Table 5: County-level TTD regression results

Variables	Reg. (1)	Reg. (2)	Reg. (3)
Saiz elasticity	-0.162^{***}	-0.150^{***}	-0.134^{***}
	(0.047)	(0.048)	(0.041)
Rainfall intensity	0.125^{***}	0.099^{***}	0.099^{***}
	(0.022)	(0.022)	(0.021)
Heat	0.030	0.049^{**}	0.036^{*}
	(0.022)	(0.022)	(0.021)
Observations	222	222	222
R-squared	0.206	0.192	0.190

Note: We use counties with at least 10 completed sections observed. "Rainfall intensity" measures the rainfall inches per hour on a storm of one-hour duration and a 100-year return period for each county (Data source: National Oceanic and Atmospheric Administration's Atlas 14 Precipitation Frequency Estimates). "Heat"—that is, cooling degree days—is a measure of the year's temperature hotness, calculated as the difference between the daily temperature mean (the sum of the high and low temperatures divided by two) and 65 degrees Fahrenheit, multiplied by the number of days with a positive value of this difference in a given year (Data source: National Centers for Environmental Information's Annual Climatological Data). Robust standard errors are reported in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

3 Time-to-develop model of housing investment

This section outlays the theoretical framework that ties our TTD data to the housing supply elasticity, which is a key measure of interest in the housing literature. In the model, the developer makes housing investment decisions under a TTD constraint. We analytically derive and characterize the short- to long-run housing supply elasticities as a function of several parameters, including TTD. As our focus is to analytically formulate housing supply elasticity based on TTD, it is important to note that our framework abstracts from the endogeneity of TTD or other potential short-run determinants of housing supply studied in the literature.

3.1 Model description

In period t, the developer produces housing units, I_t, using inputs built in current and previous periods, $\{U_{t-p|t}\}_{p=0,1,\cdots,P}$, based on the following TTD construction function:

$$I_t = \left(\sum_{p=0}^{P} U_{t-p|t}^{\frac{\theta-1}{\theta}} \right)^{\frac{\theta}{\theta-1}}, \quad \theta > 0. \tag{3.1}$$

The parameter P is the number of periods it takes to complete a project from the beginning, and θ governs the substitutability of the different stages of construction. This generalized TTD specification follows Sarte et al. (2015) and nests the investment assumption in Kydland and Prescott (1982) as a special case when $\theta \to 0$.[7]

In turn, the developer builds construction inputs $U_{t|t+p}$ at a lot where housing completions are scheduled in period $t+p$ for each $p \in \{0, 1, \cdots, P\}$. These inputs are built based on a Cobb-Douglas production function:

$$U_{t|t+p} = M^{1-\alpha}_{t+p-P|t+p} N^{\alpha}_{t|t+p}, \quad \forall p \in \{0, 1, \cdots, P\}, \tag{3.2}$$

where $M_{t+p-P|t+p}$ is the housing permit in the beginning period of development $t+p-P$ for the lot where new housing is scheduled to be completed in period $t+p$ and $N_{t|t+p}$ is the amount of variable construction input at that lot.

In each period, the developer purchases building permits, $M_{t|t+P}$, from the local government at a price, $q_{M,t}$, for a lot that is at the beginning stage of development. Moreover, the developer hires variable construction inputs for each lot under development at a competitive cost, w_t. When a lot is fully developed, its completed new houses, I_t, are sold at a unit price, q_t. The developer's profit in period t, Φ_t, is

$$\begin{aligned} &\Phi_t = q_t I_t - q_{M,t} M_{t|t+P} - w_t N_t, \\ &\text{where} \quad N_t = \sum_{p=0}^{P} N_{t|t+p}. \end{aligned} \tag{3.3}$$

The developer builds new houses at multiple lots by purchasing building permits and utilizing variable construction inputs to maximize its discounted sum of profits:

$$\max_{\{I_t, M_{t|t+P}, N_t, \{N_{t|t+p}, U_{t|t+p}\}_{p=0}^{P}\}} \quad \mathbb{E}_0 \sum_{t=0}^{\infty} \Lambda_{0|t} (q_t I_t - q_{M,t} M_{t|t+P} - w_t N_t),$$

subject to the TTD equations (3.1), (3.2), and (3.3). The variable $\Lambda_{0|t}$ is the stochastic discount factor between periods 0 and t, and $\mathbb{E}_t$ is the expectations operator conditional on information available in period t.

[7] Consistent with our TTD construction function when $\theta \to 1$, we assume that $I_t = \prod_{p=0}^{P} U_{t-p|t}$ when $\theta = 1$. To simplify the analysis without loss of generality, we assume $\theta \neq 1$ in this section.

Finally, permits are supplied by the local government and are elastic to the house price:

$$M_{t|t+P} = \bar{M} q_t^{\gamma}. \tag{3.4}$$

This assumption is common in the literature and is consistent with the long-run housing supply elasticity as a function of γ (Guren, McKay, Nakamura and Steinsson, 2020).

3.2 Developer's housing supply

We denote $\mu_{t|t+p}$ as the period-t Lagrange multiplier of equation (3.2) for each p and express the respective optimality conditions of (i) construction at each stage, (ii) variable inputs at each stage, and (iii) building permits at the beginning stage of development as follows:

$$\begin{aligned}
\mu_{t|t+p} &= \mathbb{E}_t \left[\Lambda_{t|t+p} q_{t+p} \left(\frac{I_{t+p}}{U_{t|t+p}} \right)^{\frac{1}{\theta}} \right] \quad \text{for } p = 0, 1, \cdots, P, \\
w_t &= \alpha \mu_{t|t+p} M^{1-\alpha}_{t+p-P|t+p} N^{\alpha-1}_{t|t+p} \quad \text{for } p = 0, 1, \cdots, P \\
q_{M,t} &= (1-\alpha)\, \mathbb{E}_t \left[\sum_{p=0}^{P} \Lambda_{t|t+p} \mu_{t+p|t+P} \frac{U_{t+p|t+P}}{M_{t|t+P}} \right].
\end{aligned}$$

The first condition shows that construction at each stage is chosen such that its shadow value, $\mu_{t|t+p}$, is equal to its marginal contribution to the expected discounted value of the completed house in the future. The second and third conditions equate the costs of variable inputs and the building permit to the respective marginal products.

After log-linearizing the previous optimality conditions as well as equations (3.1) through (3.4), we derive the following lemma that summarizes the model's housing supply conditions.

Lemma 1 (log-linearized dynamic housing supply equilibrium) *Let each hatted variable be the log deviation from its steady-state value. The following log-linearized equilibrium conditions summarize the TTD model of housing investment:*

$$\begin{aligned}
\hat{I}_t &= \frac{1}{B(P)} \sum_{p=0}^{P} \left(\tilde{\beta}^{\alpha(\theta-1)/\theta} \right)^p \hat{U}_{t-p|t}, \\
\left(\frac{1-\alpha}{\alpha} + \frac{1}{\theta} \right) \hat{U}_{t|t+p} &= \frac{1}{\theta} \mathbb{E}_t \hat{I}_{t+p} + \mathbb{E}_t \hat{q}_{t+p} + \frac{1-\alpha}{\alpha} \hat{M}_{t+p-P|t+p} - \hat{w}_t + \mathbb{E}_t(\hat{\lambda}_{t+p} - \hat{\lambda}_t), \\
\hat{U}_{t|t+p} &= (1-\alpha) \hat{M}_{t+p-P|t+p} + \alpha \hat{N}_{t|t+p},
\end{aligned}$$

$$\hat{N}_t = \left(\frac{1-\tilde{\beta}}{1-\tilde{\beta}^{1+P}}\right)\sum_{p=0}^{P}\tilde{\beta}^p \hat{N}_{t|t+p},$$

$$\hat{M}_{t|t+P} = \gamma \hat{q}_t,$$

where

$$\Lambda_{t|t+p} = \beta^p \frac{\lambda_{t+p}}{\lambda_t},\ \tilde{\beta} = \beta^{\frac{\theta}{\theta+\alpha(1-\theta)}},\ \textit{and}\ \ B(t) = \frac{\tilde{\beta}^{(\alpha(\theta-1)/\theta)(1+t)}-1}{\tilde{\beta}^{\alpha(\theta-1)/\theta}-1}.$$

Note that in Lemma 1, we introduced the deterministic per-period discount factor parameter $\beta < 1$. As such, the stochastic discount factor between period t and $t+p$, $\Lambda_{t|t+p}$, is decomposed into the deterministic discount factor β^p and the net stochastic discount factor λ_{t+p}/λ_t.

Using Lemma 1, the following proposition derives the developer's period-by-period housing supply choice as a function of house prices and other general equilibrium forces.

Proposition 2 (period-by-period housing supply curve) *Based on the lemma and assuming a steady-state equilibrium before period* 0*, we derive the following period-by-period housing supply curve:*

$$\hat{I}_t = \begin{cases} \Upsilon_t(P)\hat{q}_t + \mathbf{GE}_t(P), & \forall t \in [0, P), \\ \frac{\alpha}{1-\alpha}\hat{q}_t + \gamma\hat{q}_{t-P} + \widetilde{\mathbf{GE}}_t(P), & \forall t \in [P, \infty), \end{cases}$$

where

$$\Upsilon_t(P) = \frac{B(t)}{\left(\frac{1-\alpha}{\alpha}+\frac{1}{\theta}\right)B(P) - \frac{1}{\theta}B(t)},$$

$$\mathbf{GE}_t(P) = -\frac{\Upsilon_t(P)}{B(t)}\sum_{j=0}^{t}(\tilde{\beta}^{\alpha(\theta-1)/\theta})^j\left(\hat{w}_{t-j} + (\hat{\lambda}_{t-j} - \hat{\lambda}_t)\right),$$

$$\widetilde{\mathbf{GE}}_t(P) = -\frac{\Upsilon_P(P)}{B(P)}\sum_{j=0}^{P}(\tilde{\beta}^{\alpha(\theta-1)/\theta})^j\left(\hat{w}_{t-j} + (\hat{\lambda}_{t-j} - \hat{\lambda}_t)\right).$$

All the proofs are available in the Appendix. Proposition 2 decomposes housing supply into its partial equilibrium and general equilibrium components. The general equilibrium component, denoted by $\mathbf{GE}_t(P)$ and $\widetilde{\mathbf{GE}}_t(P)$, depends on the current and past histories of construction wages and the stochastic discount factors. Our object of interest in this section is housing supply in partial equilibrium, and the general equilibrium forces that depend on the setup of the overall economy will be studied in section 5. The partial equilibrium component

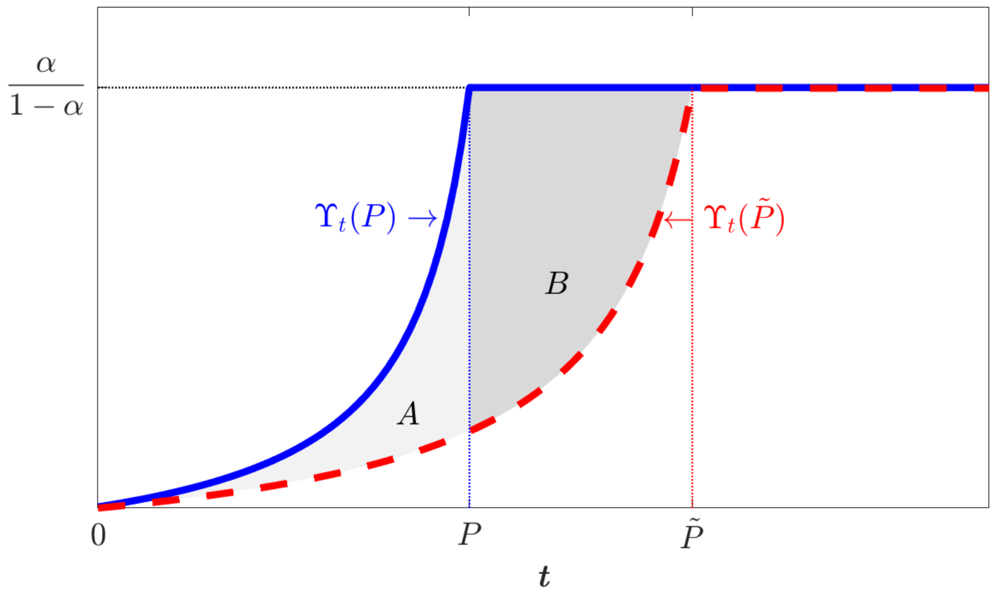

Figure 1: Theoretical static housing supply elasticities

consists of the current housing price before the TTD constraint ($t < P$) and the P-period lagged house price after the TTD constraint ($t \geq P$). The partial equilibrium component has a well-defined static housing supply elasticity, which is $\Upsilon_t(P)$ when $t < P$ and $\alpha/(1-\alpha)$ when $t \geq P$. The following corollary provides us some useful comparative statics with regard to the derived static housing supply elasticities.

Corollary 3 (comparative statics) *The static housing supply elasticity when $t < P$, $\Upsilon_t(P)$, has two properties. First, $\Upsilon_t(P)$ is positive and an increasing function of time, with an upper bound of $\alpha/(1-\alpha)$:*

$$0 < \Upsilon_{t-1}(P) < \Upsilon_t(P) < \frac{\alpha}{1-\alpha}.$$

Second, $\Upsilon_t(P)$ is larger when the TTD constraint P is shorter:

$$\Upsilon_t(P) > \Upsilon_t(\tilde{P}) \quad \textit{when } P < \tilde{P}.$$

Corollary 3 is visualized in Figure 1. As observed, the static housing supply elasticity is an increasing function up to the TTD constraint. Afterward, the static housing supply elasticity is determined by the parameter α, which represents the production elasticity to variable construction inputs. Comparing the housing supply elasticity between two regions with different

TTD constraints, P and $\tilde{P}$, we find that the region with a shorter TTD constraint has a higher static housing supply elasticity for two reasons. First, housing supply is more flexible during periods under construction, represented by area A in the figure. Second, housing supply determined at the beginning period is completed earlier, represented by area B in the figure. In turn, area $A + B$ is the cumulative housing supply elasticity difference in the two regions. As noted earlier, our model nests the TTD assumption in Kydland and Prescott (1982) as a special case when $\theta \to 0$. In this case, the static housing supply elasticity becomes a step function: 0 before the TTD constraint is reached and $\alpha/(1-\alpha)$ afterward. As such, the difference in the static housing supply elasticity across regions in Kydland and Prescott (1982) arises only after the TTD constraint is reached in the more flexible region, which is area B.

3.3 The short- and long-run empirical housing supply elasticities

Using the proposition, we define the T-horizon housing supply elasticity that is consistent with existing empirical measures of the housing supply elasticity.

Definition 4 (T-horizon housing supply elasticity) *The T-horizon housing supply elasticity is defined as the average of the theoretical partial equilibrium period-by-period housing supply elasticities between periods 0 and T. Using Proposition 2 and assuming $\Upsilon_t(P) = \alpha/(1-\alpha)$ when $t \geq P$ for simplicity of notation, we define the T-horizon housing supply elasticity with P-period TTD, $\mathcal{E}_T(P)$, as*

$$\mathcal{E}_T(P) \equiv \frac{\Delta_{t=0}^{T} \hat{I}_t}{\Delta_{t=0}^{T} \hat{q}_t} \equiv \frac{1}{T+1} \sum_{t=0}^{T} \left[\Upsilon_t(P) + \gamma(T-P+1) \times \mathbf{1}_{\{T \geq P\}} \right], \tag{3.5}$$

where $\mathbf{1}_{\{T \geq P\}}$ is an indicator function equal to 1 when $T \geq P$. Moreover, the long-run housing supply elasticity, $\mathcal{E}_\infty$, is defined as

$$\mathcal{E}_\infty \equiv \lim_{T \to \infty} \mathcal{E}_T(P) = \frac{\alpha}{1-\alpha} + \gamma.$$

By taking an average of the period-by-period housing supply elasticities, our T-horizon housing supply elasticity summarizes the evolution of housing supply over time based on the supply-side behavior. In the short run when $T < P$, housing supply elasticity is not a function of γ but a function of TTD and other parameters of the housing construction function. In the long run, housing supply elasticity is purely a function of γ and α; TTD is no longer relevant. For $P \leq T < \infty$, housing supply elasticity is a weighted average of TTD and the long-run elasticity, where the latter matters more as $T \to \infty$.

Of note, this definition is not the only way to characterize the T-horizon housing supply elasticity. Depending on the endogenous forces that drive house prices, different weights on the period-by-period housing supply elasticities might better characterize the average housing supply elasticity over the horizon of interest. Indeed, our unweighted average of the period-by-period housing supply elasticities might be viewed as an agnostic measure to the various driving forces of the house price throughout the horizon.

4 Quantifying housing supply elasticities in each horizon

Using the duration of land development statistics in section 2 and our theoretical result in section 3, this section presents regional housing supply elasticities in each horizon. We ask whether the significant TTD variations we find in the data translate to significant variations in the $T-$horizon housing supply elasticities by comparing those with the counterpart long-run elasticities. In this section, our main focus is on quantifying housing supply elasticities at different horizons; their relevance in accounting for housing market dynamics will be studied in section 5.

4.1 Parameterization

Recall county i's T-horizon housing supply elasticity from Definition 4. In equation (3.5), the elasticity $\mathcal{E}_T^i(P_i)$ requires five inputs: P_i, γ_i, α, β, and θ.[8] We begin by discussing the calibration of the key county-level parameters P_i and γ_i. First, P_i is calibrated by the county median TTD estimate in the previous section. Note that the TTD estimate is the time span between the approval of the preliminary site plan and the completion of the project. As such, it is inclusive of the county's average time span between submitting a project for final approval and receiving a decision, documented in Gyourko, Hartley and Krimmel (2021). Next, the permit supply elasticity parameter (γ_i) is derived from Saiz's long-run housing supply elasticity together with α described below.

For the remaining parameters, we set the time discount factor, β, at 0.99, as the model is calibrated at a quarterly frequency. The construction labor share, α, is set at 0.385 which implies that a county with the smallest Saiz's supply elasticity has a permit elasticity γ_i at its lower bound of zero. This value of α is consistent with our estimate of the construction labor income of 37 percent in the KLEMS account when we assume that overhead labor costs

[8] Due to data availability, we assume that the time discount factor (β), the construction labor share (α), and the elasticity of substitution across construction stages (θ) are identical across all counties.

are about 10 percent of the total labor cost.[9] The elasticity of substitution across construction stages, θ, is set at 0.5 in our baseline. The calibration of θ is not straightforward, but we think a complementarity assumption ($\theta < 1$) is intuitive, as most construction activities are conducted sequentially on site. Another piece of evidence that tells us that θ is small is the minor difference in our data between the number of housing units planned at the beginning of development and the number of housing units completed at the end of development. For example, the total number of housing units submitted at the beginning of development does not change at the end of development for about 60 percent of the sections in our sample. In the Appendix, we show that the regional variation of our supply elasticity measures is less sensitive to this parameter.

4.2 The $T-$horizon housing supply elasticity

In Figure 2, we quantify the county-level house price elasticity for each horizon. The left panel compares our measure of horizon-specific elasticities with the Saiz elasticity. The takeaway from this panel is that the shorter-run housing supply elasticities are not simply a monotonically smaller version of the long-run supply elasticities. While one-year ($T = 4$) housing supply elasticities are close to zero and show little variability, five-year ($T = 20$) housing supply elasticities show significant variations. As T increases, the $T-$horizon housing supply elasticities increase and converge to the 45 degree line that equates the long-run elasticity.

The right panel shows the standard deviation of the $T-$horizon elasticities for each T normalized by the standard deviation of the long-run elasticities. The blue solid line plots the results when we use the county-level TTD measure used in the regression (3) column in Table 4. For example, when $T = 9$, our $T-$horizon supply elasticity is half as variable in standard deviation as the long-run elasticity. The red dashed line plots the results when we assume that all counties have the same TTD at the national median of $T = 11$. In this case, the cross-county variation in housing supply elasticity obviously disappears in the short run. Therefore, the difference between the blue solid line and the red dashed line quantifies the share of TTD variations in accounting for the cross-county variations of housing supply elasticities for each T. We find that the difference between the blue solid line and the red dashed line remains sizable especially for lower values of T, suggesting that TTD variations play a significant role in accounting for total variations in each of the $T-$horizon housing supply elasticity. For example, at $T = 16$, TTD contributes to about 50 percent of the total

[9] KLEMS stands for K-capital, L-labor, E-energy, M-materials, and S-purchased services; the term refers to broad categories of intermediate inputs consumed by industries in their production of goods and services.

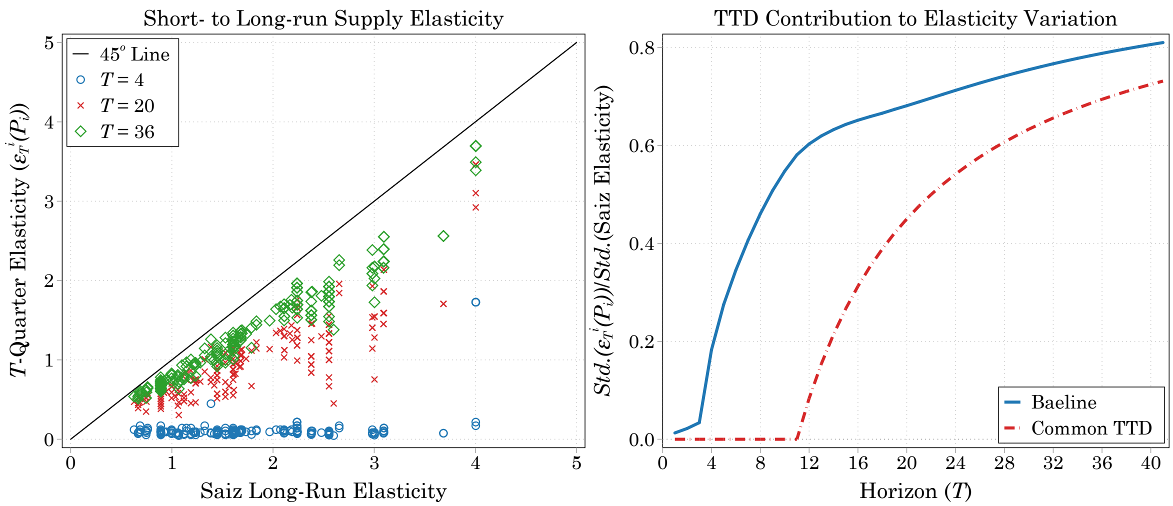

Figure 2: Supply elasticities in each horizon

Note: The left panel shows the scatter plot comparing the Saiz supply elasticity and the $T-$horizon supply elasticities with $T = 4, 20, 36$ quarters. The right panel shows the cross-county variation of the $T-$horizon supply elasticities when we use (i) the TTD measure constructed by using regression (3) in Table 3 (blue solid line) and (ii) the median county TTD measure of 11 quarters (red dashed line).

variation in the $T-$horizon housing supply elasticity.

In sum, we find that the large variations of TTD documented in the data translate to quantitatively sizable variations in the shorter-run housing supply elasticities. This suggests that our shorter-run housing supply elasticities could provide a new perspective compared to the long-run housing supply elasticity in the study of housing dynamics.

Nevertheless, our focus so far on housing supply elasticities has limitations because the equilibrium housing market behavior is a combination of both housing supply and demand and it is unclear whether the measured variations in supply elasticities would imply significant variations in housing market variables that we observe through the data. That is the topic of the next section, where we study the equilibrium effects of short-run housing supply elasticities on house prices and housing investment.

5 Theory and application of short- and long-run elasticities

We documented that TTD is lengthy and dispersed across counties. Mapping the empirical TTD into our model's analytical expression for housing supply elasticity, we find that the short-run housing supply elasticity is also highly dispersed across counties. As the housing supply elasticity is derived from a partial equilibrium model of housing developers, our results are not specific to the nature of economic shocks or the endogenous response of housing demand. However, housing supply elasticity does not directly inform us about the dynamics of house prices and quantities that are of interest in the literature. To understand its implications on house prices and quantities, one needs to specify the demand side of housing.

In this section, we study the theoretical implications of our short- and long-run housing supply elasticities on house prices and quantities using a local general equilibrium model of housing investment. That is, we extend the partial equilibrium model of developers in section 3 by incorporating the household and the nondurable goods–producing sectors. After specifying the local general equilibrium model, we hit the economy with a common housing demand shock and study the differential local house price and quantity responses. We show that our calibrated model predicts two salient housing market features that can be tested with house price and quantity data. We also investigate the effectiveness of a government's discretionary housing supply policy when TTD is taken into account.

5.1 Local general equilibrium TTD model

The local economy consists of housing developers, households, nondurable goods producers, and a local government. Since the national central bank sets the interest rate, we assume that the local economy takes it as given. Therefore, the bond and nondurable goods markets do not clear locally, analogous to the assumptions in small open economy models in the international macro literature. Housing developers follow the same specification and notation as in section 3. Below, we describe the households, the rest of the economy, and the equilibrium of the model.

5.1.1 Households

The representative local household's expected lifetime utility is

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t U(C_t, H_t, N_{n,t}, N_t; \varphi_t), \tag{5.1}$$

where C_t is the household consumption of nondurable goods, H_t is the service flow of housing, $N_{n,t}$ is the labor supply for the nondurable goods sector, and φ_t is an exogenous process for housing demand. The household's one-period subjective discount factor, β, is consistent with the housing developers' deterministic discount factor specified in section 3.

The household's service flow of housing is proportional to its housing stock. For simplicity of notation, the housing stock is also denoted as H_t. The housing stock evolves over time by

$$H_t = (1-\delta)H_{t-1} + I_t, \tag{5.2}$$

where δ is the depreciation rate of the housing stock.

The household flow budget constraint is given by

$$C_t + q_t I_t + \frac{B_{t+1}}{R_t} + \frac{\psi_b}{2}B_{t+1}^2 = w_{n,t}N_{n,t} + w_t N_t + B_t + \Phi_t + T_t, \tag{5.3}$$

where B_{t+1} is the household's one-period bond holdings that mature in period $t+1$, R_t is the gross bond interest rate between periods t and $t+1$, $w_{n,t}$ is the real wage for working in the nondurable goods sector, Φ_t is the period-t profit of developers because households are the final owners of the developers, and T_t is transfers from the local government. As is standard in small open economy models, the household is subject to the bond portfolio adjustment cost $\psi_b B_{t+1}^2/2$. The parameter ψ_b is calibrated to be positive for stability in solving the model but small enough to not materially affect the model dynamics.

5.1.2 The rest of the economy

As we will discuss next, the rest of the economy consists of the nondurable goods producers, the local government, and the market-clearing conditions. We also specify the exogenous process for housing demand that we use for later applications.

Nondurable goods producers. The representative nondurable goods producer operates with a linear production technology, $Y_t = \bar{Z}N_{n,t}$, where Y_t is the output of the nondurable good and $\bar{Z}$ captures its productivity. The profit of the producer is $Y_t - w_{n,t}N_{n,t}$, where both the input and output markets are perfectly competitive. The nondurable goods are tradable to other regions.

Local government. As specified in equation (3.4), the supply of housing permits is determined by its local government, which in turn is elastic to the region's equilibrium house price. For each housing permit, the local government collects a fee, $q_{M,t}$, from developers. The local government also collects the bond portfolio adjustment cost from households. The local government follows a balanced budget by rebating back its revenue to the households in the form of transfers T_t:

$$T_t = q_{M,t} M_t + \frac{\psi_b}{2} B_{t+1}^2. \tag{5.4}$$

Market clearing. The labor markets for the nondurable goods sector and the construction sector clear by equating the supply and demand of labor in each sector. The permit market clears by equating permit supply to permit demand. The market for new housing clears by equating the supply and demand of new housing investment. The bond and nondurable goods markets do not clear locally as we assume that the interest rate is exogenously determined by the national central bank. Finally, the following resource constraint of the local economy needs to be satisfied:

$$C_t + \frac{B_{t+1}}{R_t} = w_{n,t} N_{n,t} + B_t. \tag{5.5}$$

Exogenous housing demand. The exogenous component of housing demand, φ_t, is attached to the household's preference over the housing stock H_t in the utility function (5.1) and follows a first-order autoregressive process in logs:

$$\log \varphi_t = (1 - \rho_\varphi)\bar{\varphi} + \rho_\varphi \log \varphi_{t-1} + \epsilon_{\varphi,t}, \tag{5.6}$$

where $\epsilon_{\varphi,t}$ is the exogenous housing demand shock and ρ_φ is the persistence of exogenous housing demand from its mean $\bar{\varphi}$.

5.1.3 Equilibrium

The local general equilibrium is a set of variables $\{U_{t|t+p}, N_{t|t+p}, \mu_{t|t+p}\}_{p=0}^{P}$, $M_{t|t+P}$, N_t, I_t, H_t, C_t, Y_t, $N_{n,t}$, B_{t+1}, w_t, $w_{n,t}$, q_t, R_t for $t \geq 0$ such that taking as given the endogenous prices w_t, $w_{n,t}$, and q_t, the exogenous processes R_t and φ_t, and the initial conditions B_0 and H_{-1}, the following conditions hold:

1. Housing developers maximize their profit subject to (3.1) through (3.3).
2. Households maximize their lifetime utility (5.1) subject to (5.2) and (5.3).
3. Nondurable goods producers maximize their profit.
4. The local government supplies permits and balances its budget according to (5.4).

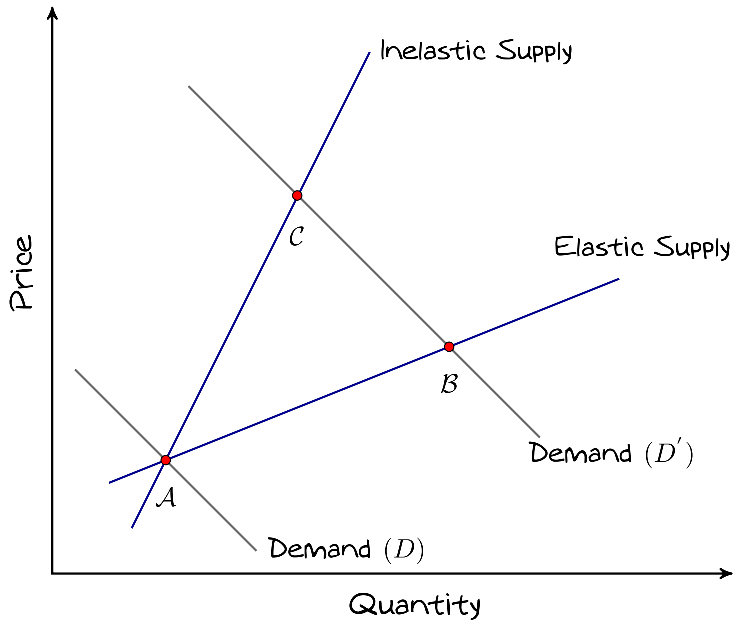

Figure 3: Housing supply and demand curves

5. Markets clear for nondurable goods labor, construction labor, permits, and housing investment, and the resource constraint (5.5) is satisfied.

We assume that the interest rate is fixed by the national central bank. The exogenous component of housing demand follows (5.6).

5.2 Relative price and relative quantity correlation in the housing market

Consider two regions with different housing supply elasticities. In Figure 3, the housing supply curves of the two regions are denoted as inelastic supply and elastic supply. Assume that the initial housing market equilibrium for both regions is at A, where the demand curve is denoted as D. When there is positive housing demand that shifts the demand curve from D to D', the equilibrium price and quantity responses are different for the two regions. In the inelastic supply region, the equilibrium is formed at C, where prices increase by a lot and quantities increase by little. By contrast, in the elastic supply region, the equilibrium is formed at B, where prices increase by little and quantities increase by a lot. This outcome implies that conditional on a demand shock, the cross-regional correlation between the relative price and the relative quantity should be negative.

In the data, however, the cross-county correlation of the relative price and the relative quantity is mildly positive. Using the county house price indexes from the Federal Housing Finance Agency and the annual estimates of the number of housing units from the Census Bureau, we find that the cross-county correlation of the annual growth rates of prices and quantities is positive in 15 years between 2001 and 2019 with an average correlation of 0.13. This discrepancy between theory and data poses a challenge to the view that housing supply elasticities play an important role in the cross-regional dynamics of the housing market (Davidoff, 2016).

We provide an explanation for this discrepancy through a calibrated local general equilibrium model, as specified above. When a positive demand shock hits the economy, the local governments supply new permits, and the region with elastic long-run housing supply observes a larger number of new permits supplied by the local government compared to the region with inelastic long-run housing supply, as shown in the left panel of Figure 4. This response is the standard *level effect* of the supply elasticity—that is, housing supply is higher in locations with higher long-run elasticity.

With TTD, however, our model also generates a *slope effect* of the supply elasticity as TTD generates a positive slope between the short- and long-run supply elasticity.[10] As developers in each region deal with multiple projects under different stages of construction, they make the most out of this process by allocating more variable inputs to the project with a large number of permits at the expense of lowering inputs of some other projects with fewer permits. The middle panel of Figure 4 shows that this substitution channel is stronger in an elastic region relative to an inelastic region, because developers in regions with high long-run supply elasticity face a steeper slope of the elasticity path and hence a stronger intertemporal substitution motive. As a consequence, the right panel of Figure 4 shows that new housing quantities in the elastic region are lower than those in the inelastic region until the period when the TTD constraint is met.

Quantifying the county-level TTD and long-run elasticity data, Figure 5 presents the conditional relative price and relative quantity rank correlations at each horizon under a common housing demand shock. In the baseline, the correlation turns positive in the short run, peaks at around six quarters, and then gradually trends down until it moves into negative territory as the TTD friction is lifted in many counties. We also plot the correlation in a counterfactual model in which TTD for all counties is set to be at the national median of 11 quarters. In this case, the only source of cross-county variation is the long-run elasticity. This model predicts that the conditional cross-county correlation is either 1 in the short run or -1 in the longer run when the TTD constraint is lifted. In both models, the long-run correlations are -1, as

[10] For simplicity of discussion, we assume that the TTD constraint is the same 11 quarters for both regions.

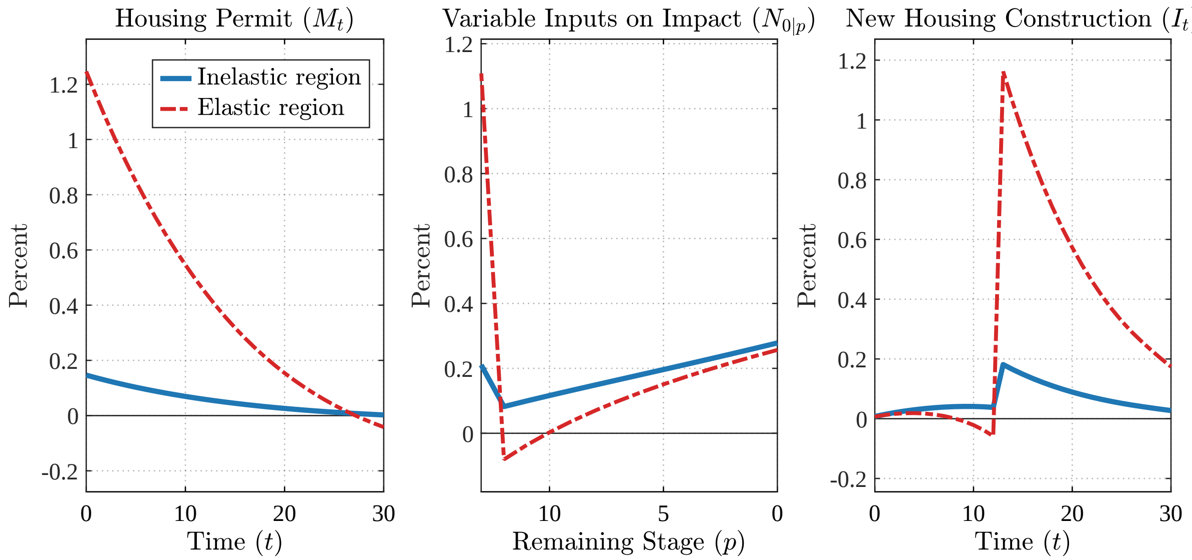

Figure 4: Impulse responses to a common demand shock

Note: This figure shows the impulse response of the housing permit (left panel), time-0 responses of each p-stage construction labor (middle panel), and the impulse response of new housing construction (right panel), to a common demand shock. Blue solid lines show the responses in the region with inelastic housing supply ($\gamma = 0.9$) and red dashed lines show the responses in the region with elastic housing supply ($\gamma = 3.0$). The common demand shock is assumed to have a persistence of 0.95.

housing supply in the long run is determined only by the long-run elasticity.

We gauge the empirical relevance of the theoretical channel by simulating our model with common demand shocks and computing the unconditional cross-county correlations of the annual growth rates of prices and quantities in the same manner as our empirical estimates above.[11] We evaluate the importance of the positive short-run relative price and relative quantity correlation in generating the relative price and relative quantity correlation observed in the data by comparing our results with and without TTD.

Figure 6 indicates that the strong negative cross-county correlation between the annual growth rates of prices and quantities disappears as the short-run elasticity is assumed to be smaller than the long-run elasticity, although the correlation does not become mildly positive like in the data. This result suggests that while the discrepancy is much smaller than in the model without any TTD, other shocks are also needed to fully account for the correlation we observe in the data, such as the location-specific housing demand shocks studied extensively in Howard and Liebersohn (2021, 2023). The figure also illustrates that what matters for the model's cross-county correlation is a lengthy median TTD rather than the variation of TTD across counties. That is, there is little difference in the correlation between the baseline case

[11] We present our model calibration in the Appendix.

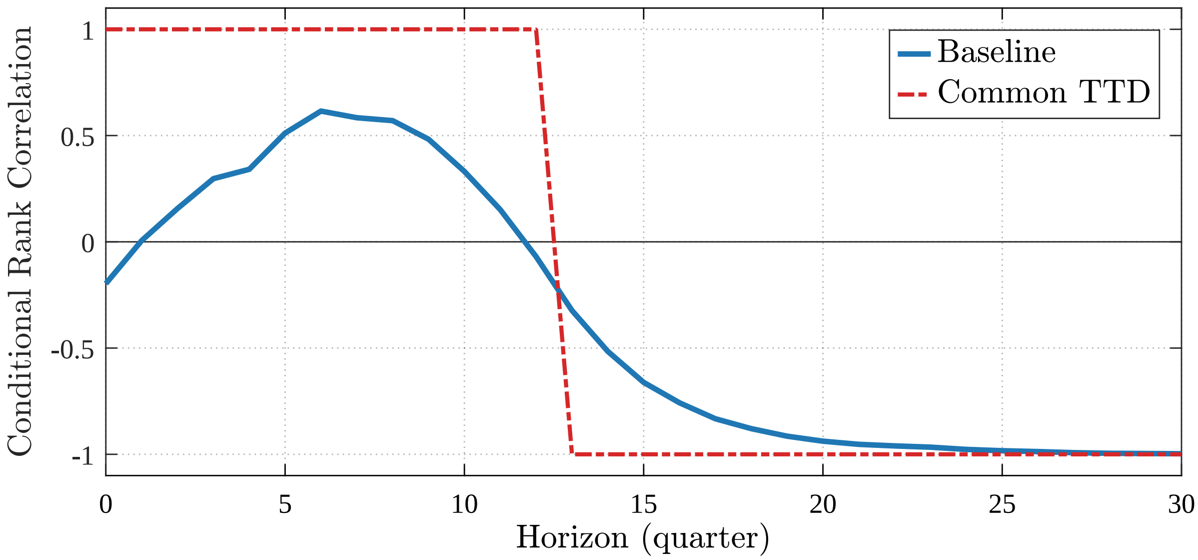

Figure 5: The relative price and relative quantity correlation under a common demand shock

Note: This figure shows the Spearman's rank correlation between (cumulative) house price response and (cumulative) new construction response to a common demand shock. Blue line shows the correlation in the baseline model with cross-county variations in both TTD and the long-run elasticities. Red line shows the correlation in the model with common TTD for all counties.

and the case in which counties have a common TTD.

Overall, our argument supports the view that housing supply elasticities could play an important role in accounting for housing market dynamics if short-run housing supply frictions in the form of TTD are also taken into account, related to a point raised in Guren et al. (2020).

5.3 House price dynamics in the short and long run

Our local general equilibrium model predicts that house prices react more sensitively to a housing demand shock in locations with inelastic housing supply conditions. With TTD constraints, there is a wedge between the short- and long-run housing supply elasticities, and section 4 finds that cross-county variations in the short-run housing supply elasticity are large and distinct from the variations in the long-run elasticity. As such, we now study the quantitative importance of these measured short-run housing supply elasticities in accounting for house price dynamics both in our model and in the data.

Model investigation. To investigate the degree to which our short-run housing supply elasticities should influence house prices in theory, we solve our model and compute each county's house price response to a common housing demand shock.

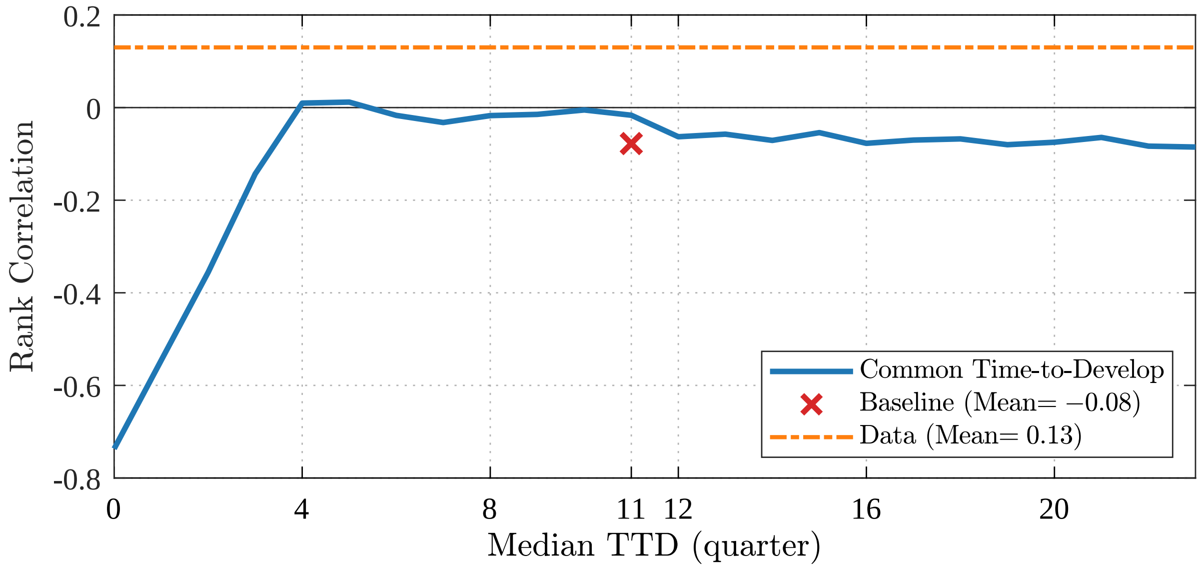

Figure 6: Unconditional relative price and relative quantity correlation: Model versus data

Note: This figure shows the average of Spearman's rank correlation between house price and new construction in response to a common demand shock. Blue solid line shows the average correlation of the simulated data in the common TTD models as a function of TTD (P). Red cross shows the average correlation in the baseline model where both TTD and long-run elasticity are heterogeneous (median TTD is 11 quarters). Lastly, orange dashed line shows the average correlation in the 2001-2019 US county data.

In the left panel of Figure 7, we plot the rank correlation for each horizon T between a county's T-horizon housing supply elasticity, $\mathcal{E}_T$, and its impact house price response to a common housing demand shock, q_0. For each persistence parameter of interest for the housing demand shock, the rank correlation is negative as house prices tend to increase more in counties where housing supply elasticities are lower. The rank correlations are not exactly -1, however, because the county rankings of the T-horizon housing supply elasticity are not the same across the horizons, as discussed in section 4. Moreover, the non-monotone rankings of the various horizons of housing supply elasticities suggest that the rank correlation should depend on the persistence of the housing demand shock. In the figure, we find that when the persistence is relatively high ($\rho_\varphi = 0.92$), the rank correlation is close to -1 for housing supply elasticities at a horizon of 7 or more years. This near-perfect, negative alignment between the longer-run housing supply elasticities and the impact house price response no longer holds when the persistence is relatively low ($\rho_\varphi = 0.88$), as the lowest rank correlation is around -0.9 at the four-year housing supply elasticity and the correlation is around -0.7 at the long-run housing supply elasticity.

These results suggest that our short-run housing supply elasticities have implications on the richer dynamics of house prices not captured by the long-run housing supply elasticity. That is,

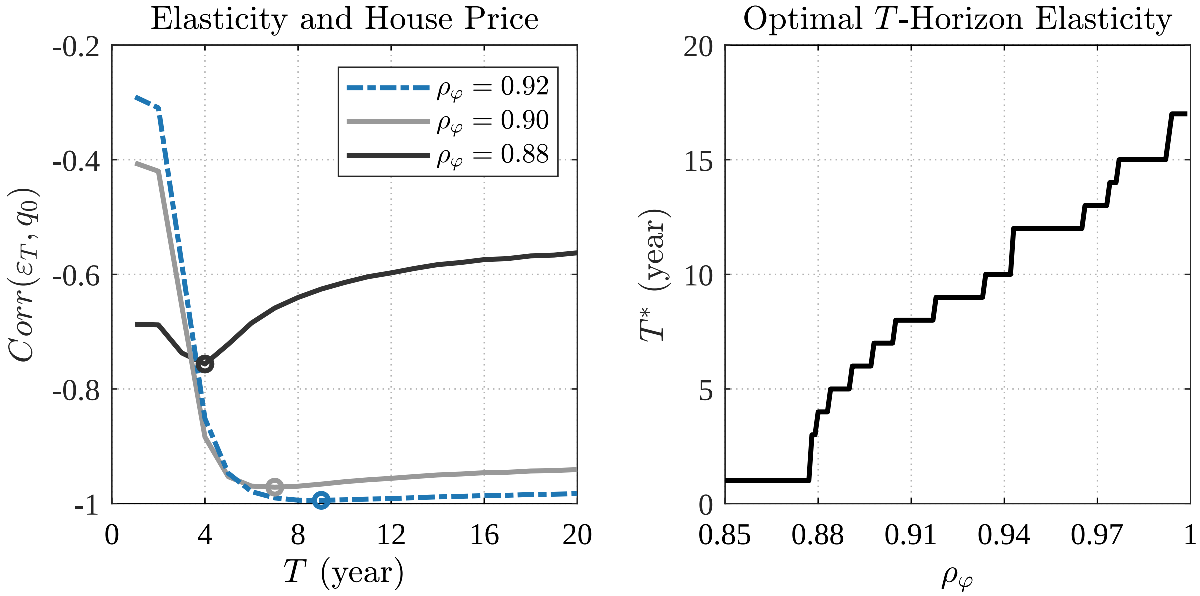

Figure 7: Short-run housing supply elasticities and house prices across counties

Note: The left panel plots the Spearman's rank correlation coefficient (at each T) between a county's T-horizon housing supply elasticity and the size of its impact house price response, when each county is hit by a common housing demand shock subject to the three calibrated persistence parameters. The circle marker of each line indicates the lowest correlation coefficient for the given persistence parameter. The right panel plots the T value that is consistent with the lowest correlation between the T-horizon housing supply elasticity and the size of the impact house price response for each persistence parameter $\rho_\varphi \in (0.85, 1)$ of the common housing demand shock.

when the persistence of the common housing demand shock is relatively low, our shorter-run housing supply elasticities are more relevant in accounting for the sensitivity of house prices in terms of exhibiting the lowest rank correlation. In the right panel of Figure 7, we show the optimal horizon T^* at which the Spearman's rank correlation between the T-horizon housing supply elasticity and the impact house response is minimized, for each persistence parameter ρ_φ of the housing demand shock:

$$T^*(\rho_\varphi) = \arg \min_{T \in [1,\infty)} Corr(\mathcal{E}_T, q_0(\rho_\varphi)).$$

We find that the most relevant housing supply elasticity depends on the persistence of the housing demand shock.[12] For example, when the housing demand shock of interest is highly persistent at or above 0.98, the optimal T-horizon housing supply elasticity is 15 years or higher, suggesting that the Saiz long-run elasticity is a relatively good benchmark to study

[12]We only plot the case where the persistence parameter implies that the half-life of the shock is more than a year.

house price responses. However, when the persistence of the housing demand shock of interest is less than 0.90, the optimal T-horizon housing supply elasticity is 7 years or lower and the relevance of the long-run elasticity is diminished.

In sum, we find that when the housing demand shock of interest has a relatively low persistence, housing supply elasticities with short horizons are more relevant than the long-run elasticity in accounting for the relative house price dynamics across counties. The performance of the long-run elasticity is weak because shocks with lower persistence do not last long enough to affect housing supply in the long run. We also verify that when the housing demand shock of interest is highly persistent, the long-run housing supply elasticity dominates the short-run elasticity as variations in TTD matters less in this case.

These findings hold even when we allow for cross-regional spillovers. In the Appendix, we study a two-region general equilibrium model with asymmetric housing supply conditions. Specifically, we assume that the short-run housing supply elasticity is larger in region one but that the long-run housing supply elasticity is larger in region two. Conditional on a positive common housing demand shock, we find that the impact house price response is larger in region two, but the response reverses eventually, and the medium-run house price response is larger in region two.

Empirical exercise. To study the empirical relevance of our short-run housing supply elasticity in accounting for house price dynamics, we focus on four episodes of the recent housing cycle: (1) the 2000s housing boom (2002–06), (2) the 2000s housing bust (2006–09), (3) the 2010s housing recovery (2012–19), and (4) the COVID housing boom (2020–22). In each of these episodes, we estimate the following relative house price regression for each horizon $T \in \{1, 2, 3, \cdots\}$:

$$\Delta \log \left(P_i/P_N\right) = \kappa_T \tilde{\mathcal{E}}_T^i + \Omega \mathbf{X_i} + u_i. \tag{5.7}$$

The left hand side variable is the log change of county i's house price index, P_i, relative to the national house price index, P_N, between the years of interest. The variable $\tilde{\mathcal{E}}_T^i$ is county i's $T-$horizon housing supply elasticity standardized to have the same cross-county variation across T. We allow for a set of control variables (including a constant), $\mathbf{X_i}$, as well as a residual term, u_i, that captures other unmodeled drivers of the relative house price. Because a higher value of the housing supply elasticity should imply a lower sensitivity of the relative house price to shocks, a simple test to validate the empirical relevance of the $T-$horizon housing supply elasticity is to show that the estimated coefficient κ_T is significantly negative.

Figure 8 displays the estimation results of the relative house price regression in each of

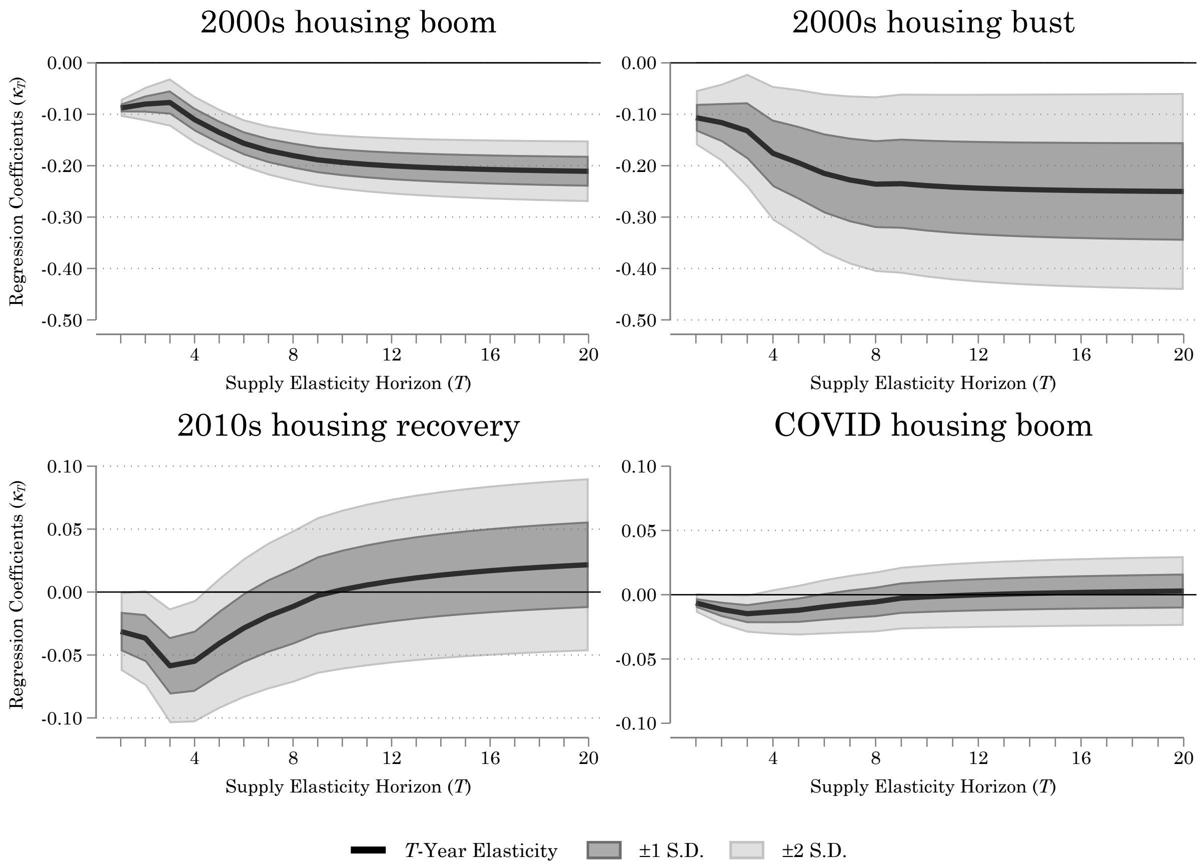

Figure 8: Relative house price regression coefficients

Note: Each figure shows the estimated coefficients for κ_T as a function of the standardized supply elasticity horizon T in equation (5.7). The confidence intervals for the 1 and 2 standard deviations of the estimates are shown as the dark gray and light gray areas, respectively. The four figures present the results using data for 2002–06 (top left panel), 2006–09 (top right panel), 2012–19 (bottom left panel), and 2020–22 (bottom right panel). Local control variables include Bartik-type predicted industry employment growth, indicators for sand state and coastal state, population share of immigrants, population share of college educated, population density, and county real GDP growth.

the four episodes. During the 2000s housing boom and bust periods, we find that the long-run elasticity is more relevant than the short-run elasticity in accounting for changes in a county's relative house price, in the sense that the estimated coefficients for the long-run elasticities are more negative than the coefficients for the short-run elasticities. In both the 2010s housing recovery and the COVID housing boom periods, however, the short-run elasticity outperforms the long-run elasticity. The estimated coefficients for the elasticities below the 10-year horizon are negative and the two standard deviation confidence intervals around the 3- and 4-year elasticities remain below zero. On the other hand, the estimated coefficients for longer horizon elasticities are positive and not distinguishable from zero.

Through the lens of our model, the results imply that both the positive housing demand shock in the 2000s boom and the negative housing demand shock in the 2000s bust were perceived as highly persistent, leading to the higher relevance of the long-run elasticity in accounting for house prices. After the 2000s housing cycle, however, agents might have expected the 2010s recovery to be less persistent, possibly reflecting on the recent housing boom experience, which then made the short-run elasticity more relevant to account for house prices. That is, even in a location where the long-run elasticity is high and buildable land is plentiful, residential developers in the 2010s might have thought that if TTD is too lengthy, pursuing new development is not as worthy as before due to concerns that the positive demand could quickly reverse course. The same goes with the COVID boom, where developers might have continued to believe that the higher housing demand induced by the greater flexibility of work-from-home would not last once the virus was under control.

Finally, we note some caveats to the empirical analysis. First, the estimated regression coefficients since the 2010s are smaller in absolute value compared to the 2000s. It is indeed likely that national housing shocks played a limited role since the 2010s amid location-specific shocks in the housing market. Second, TTD could have shifted especially during the COVID boom when there were known bottlenecks to construction, such as the shortage in lumber. While we think these bottlenecks were widespread across the country and did not meaningfully affect the *relative* TTD across locations, if TTD was disproportionately shifted in several locations to the extent that the overall TTD rankings were significantly changed, then our results should be taken with more caution.

The housing wealth effect. Following the discussion that our short-run housing supply elasticity is as relevant as the long-run housing supply elasticity in accounting for relative house price dynamics across locations, we now estimate the housing wealth effect implied by our elasticity. Following Guren et al. (2020), the housing wealth effect we estimate is the elasticity of retail employment per capita to real house prices at the CBSA level. We construct an instrument for the real house price change by interacting our 5-year housing supply elasticity with the national real house price change in the same period. We follow the analysis in Guren et al. (2020) and display the 10-year rolling window estimates of the housing wealth effect in Figure 9. For example, the point estimate for 1990q1 is based on the regression between 1985q1 and 1995q1. For comparison, we also present the same sample estimates using both the sensitivity instrument of Guren et al. (2020) and the housing supply elasticity of Saiz (2010). We find that our results are consistent with those based on the other two instruments especially after the early 2000s. As such, we confirm that the magnitude of the housing wealth

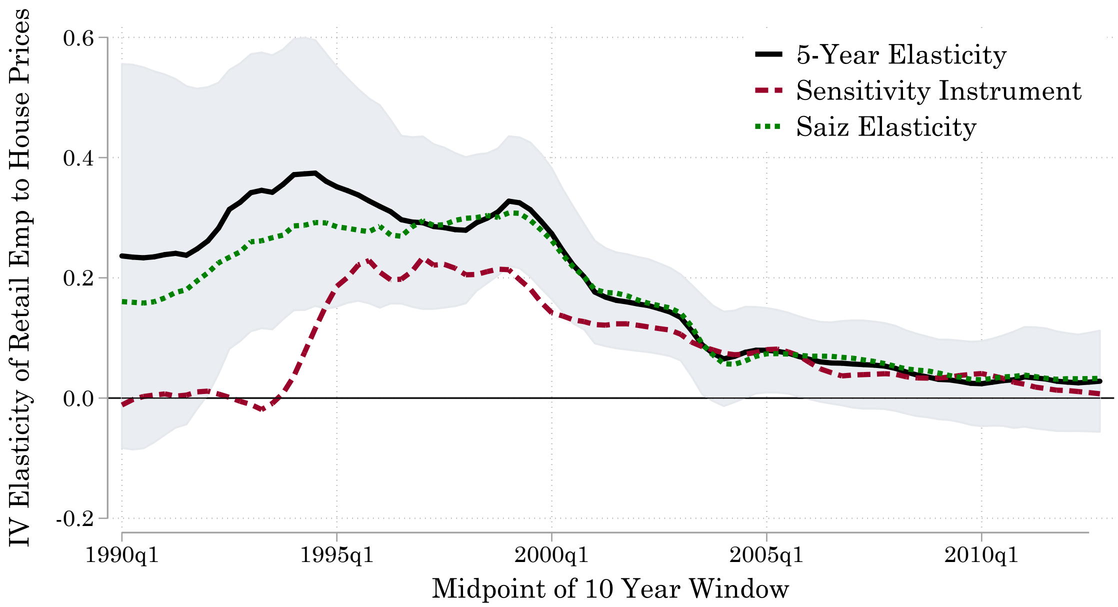

Figure 9: The elasticity of retail employment per capita to house prices over 10-year windows

Note: The figure plots the elasticity of retail employment per capita to real house prices at the CBSA level for rolling 10-year sample periods for three different methods. Each point indicates the elasticity for a 10-year sample period with its mid-point in the quarter stated on the horizontal axis. The black solid line uses an instrument that interacts the 5-year housing supply elasticity derived from equation (3.5) with the national annual log change in house prices. The red dashed line uses the sensitivity instrumental variable estimator in Guren et al. (2020). The green dotted line uses an instrument that interacts the estimated housing supply elasticity from Saiz (2010) with the national annual log change in house prices. All three specifications use the same controls and CBSA fixed effects as described in Guren et al. (2020). The sensitivity instrument also includes region-time fixed effects, while the 5-year elasticity and the Saiz elasticity only use time fixed effects. The gray area shows 95% confidence intervals for the 5-year elasticity regression where the standard errors are constructed using two-way clustering by CBSA and time. All cases use the same sample of 79 CBSAs where the TTD data are available.

effect discussed in Guren et al. (2020) is also sensible when using our elasticity that is based on a different source of variation from the existing approaches.

5.4 The effectiveness of housing supply policy in stabilizing house prices

A rapid increase in house prices raises concerns of policymakers, as these developments could subsequently lead to outsized drops in those prices that amplify stress in the financial system and the broader economy. As such, stabilizing house prices is a key objective of policymakers, and various measures are discussed and implemented in practice. In this part, we study the effectiveness of the government's discretionary housing supply policy as a tool for house price stabilization when TTD is taken into account.

Before the analysis, we clarify what we mean by discretionary housing supply policy. As summarized in Glaeser and Gyourko (2008), new construction in the U.S. is regulated in terms of building codes and land-use rules. In particular, there are numerous examples of land-use regulations that directly limit housing supply across regions, such as minimum lot size requirements, height restrictions, or growth-control policies. The discretionary housing supply policy we have in mind is a temporary relaxation of these existing land-use regulations, as in practice, new development could receive a waiver to some of the regulations.

To be specific, we modify the local government's permit supply assumption in equation (3.4) to the following:

$$M_{t|t+P} = v_t \bar{M} q_t^{\gamma}, \qquad \log v_t = \rho_v \log v_{t-1} + \varepsilon_t^v,$$

where the variable v_t indicates the government's discretionary housing supply policy that follows a first-order autoregressive process in logs.

Figure 10 plots the impulse response functions of the house price, cumulative housing construction (as a percentage of the initial housing stock), and housing permits conditional on a discretionary permit supply shock that increases the 10-year (40-quarter) cumulative construction by 2 percent of the housing stock. Compared with the result when TTD is assumed to be zero, a positive discretionary housing supply policy is somewhat less effective in reducing house prices in the short run when TTD is set at the national median of 11 quarters. Note that the peak decline in house prices occurs at around two to three years with TTD, compared with the peak decline at around one year without TTD. This difference implies that a discretionary housing supply policy with TTD could be an effective tool to stabilize house prices in the medium run through its effects on forming expectations about future supply conditions. As the literature finds that house prices tend to show momentum in the short to medium run, these discretionary supply policies could be effective in countering that momentum by controlling the expectations of future supply under the TTD commitment.

In conclusion, setting aside the political constraints in implementing a discretionary housing supply policy, we find that lengthy TTD might also somewhat limit its effectiveness in stabilizing house prices in the short run. While a discretionary housing supply policy to stabilize house prices might not have been a discussion at the national level in the U.S., this policy was implemented in Korea to tackle surging house prices in early 2021.[13] Our analysis suggests the potential challenges of such a policy when TTD is lengthy. Of note, a nationwide

[13] See Cynthia Kim (2021), "S. Korea to Boost Seoul Housing Supply by 10% to Calm Buying Frenzy," *Reuters*, February 4, https://news.trust.org/item/20210204030650-r2wji.

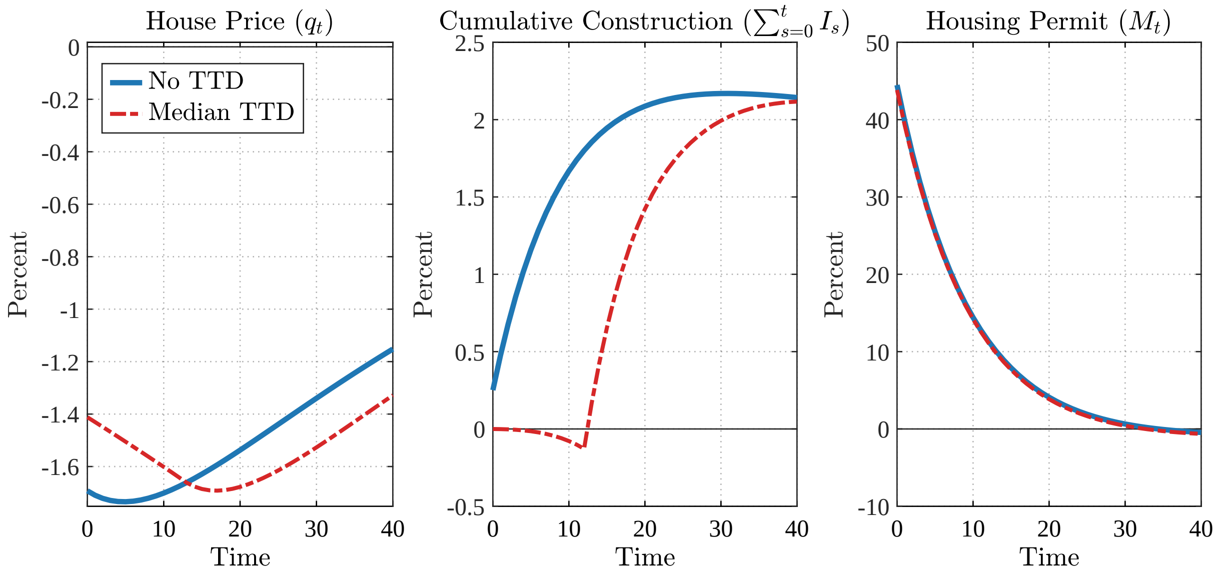

Figure 10: Model responses to a discretionary permit supply shock

Note: This figure shows the impulse responses of house price (left panel), cumulative housing construction as a percent of the initial housing stock (middle panel), and permit supply (right panel), to a discretionary permit supply shock with 0.9 persistence. We compare the model responses without TTD (blue solid lines) and those with the median TTD constraint (red dashed line). The size of the permit supply shocks in both models are scaled to have the same cumulative construction response at 40 quarters.

housing supply policy is likely to interact with the interest rate, which is not allowed in the above experiment using a local general equilibrium model.

6 Conclusion

In this paper, we use a TTD model of housing investment to formulate a link between short- and long-run housing supply elasticities and analyze TTD for residential development across the U.S. using a unique data set. We then quantify frictions to housing supply over the business cycle across major counties and draw their implications for housing market dynamics through a local general equilibrium model.

As we stated, a comprehensive process for land development takes about three years, on average, in the U.S. This feature alone introduces a large difference between the short- and long-run housing supply elasticities. In this paper, we adopt insights from the investment adjustment cost literature to shed light on the role that lengthy and dispersed TTD could play on housing market dynamics. Toward that objective, we abstract from several features of the data set that might be useful for future research. First, one could explore the time-varying nature of TTD, especially during the recent periods. While Oh and Yoon (2020) study the cyclical pattern of time-to-build in the context of the 2002–2011 housing boom-bust cycle, its lower frequency trend could also be explored in the context of understanding the half century decline in construction-sector productivity (Goolsbee and Syverson, 2023). Second, our TTD regression results suggest that geographic determinants could play a key role in construction activity. As most construction activity is still conducted on site, climate change and environmental regulation would also have a first-order effect on the construction sector. We hope that our modeling framework as well as our granular TTD data open up a new avenue of research along these lines.

References

Baum-Snow, Nathaniel and Lu Han, "The Microgeography of Housing Supply," *mimeo*, 2022.

Bhattarai, Saroj, Felipe Schwartzman, and Choongryul Yang, "Local Scars of the US Housing Crisis," *Journal of Monetary Economics*, 2021, *119*, 40–57.

Charoenwong, Ben, Yosuke Kiruma, Alan Kwan, and Eugene Tan, "Capital Budgeting, Uncertainty, and Misallocation," *Journal of Financial Economics*, 2024, *153*, 103779.

Davidoff, Thomas, "Supply Constraints Are Not Valid Instrumental Variables for Home Prices Because They Are Correlated With Many Demand Factors," *Critical Finance Review*, 2016, *5*, 177–206.

Davis, Morris and Jonathan Heathcote, "Housing and the Business Cycle," *International Economic Review*, 2005, *46* (3), 751–784.

Davis, Steven J and John C Haltiwanger, "Dynamism Diminished: The Role of Housing Markets and Credit Conditions," Working Paper 25466, National Bureau of Economic Research January 2019.

Glaeser, Edward L. and Joseph Gyourko, *Rethinking Federal Housing Policy*, The AEI Press, 2008.

Glaeser, Edward L, Joseph Gyourko, and Albert Saiz, "Housing Supply and Housing Bubbles," *Journal of Urban Economics*, 2008, *64* (2), 198–217.

Goolsbee, Austan and Chad Syverson, "The Strange and Awful Path of Productivity in the U.S. Construction Sector," *NBER Working Paper*, 2023, *30845.*

Graham, James and Christos A. Makridis, "House Prices and Consumption: A New Instrumental Variables Approach," *American Economic Journal: Macroeconomics*, January 2023, *15* (1), 411–43.

Green, Richard K, Stephen Malpezzi, and Stephen K Mayo, "Metropolitan-Specific Estimates of the Price Elasticity of Supply of Housing, and Their Sources," *American Economic Review*, 2005, *95* (2), 334–339.

Guren, Adam, Alisdair McKay, Emi Nakamura, and Jon Steinsson, "What Do We Learn From Cross-Regional Empirical Estimates in Macroeconomics," *NBER Macroeconomics Annual*, 2020, pp. 175–223.

— , — , — , and — , “Housing Wealth Effects: The Long View,” *Review of Economic Studies*, 2021, *88* (2), 669–707.

Gyourko, Joseph, Jonathan S. Hartley, and Jacob Krimmel, “The Local Residential Land Use Regulatory Environment across U.S. Housing Markets: Evidence from a New Wharton Index,” *Journal of Urban Economics*, July 2021, *124*, 103337.

Haughwout, Andrew, Richard W. Peach, John Sporn, and Joseph Tracy, “The Supply Side of the Housing Boom and Bust of the 2000s,” in Edward L. Glaeser and Todd Sinai, eds., *Housing and the Financial Crisis*, University of Chicago Press, 2013, chapter 2, pp. 69–104.

Howard, Greg and Jack Liebersohn, “Why is the Rent so Darn High? The Role of Growing Demand to Live in Housing-Supply-Inelastic Cities,” *Journal of Urban Economics*, 2021, *124*, 103369.

— and — , “Regional Divergence and House Prices,” *Review of Economic Dynamics*, 2023, *49*, 312–350.

— , — , and Adam Ozimek, “The Short- and Long-Run Effects of Remote Work on U.S. Housing Markets,” *Journal of Financial Economics*, 2023, *150* (1), 166–184.

Iacoviello, Matteo and Stefano Neri, “Housing Market Spillovers: Evidence from an Estimated DSGE Model,” *American Economic Journal: Macroeconomics*, April 2010, *2* (2), 125–64.

Kalouptsidi, Myrto, “Time to Build and Fluctuations in Bulk Shipping,” *American Economic Review*, 2014, *104* (2), 564–608.

Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante, “The Housing Boom and Bust: Model Meets Evidence,” *Journal of Political Economy*, 2020, *128* (9), 3285–3678.

Kiyotaki, Nobuhiro, Alexander Michaelides, and Kalin Nikolov, “Winners and Losers in Housing Markets,” *Journal of Money, Credit and Banking*, 2011, *43* (2–3), 255–296.

Kone, D. Linda, *Land Development*, 10 ed., BuilderBooks.com, 2006.

Kydland, Finn E. and Edward C. Prescott, “Time to Build and Aggregate Fluctuations,” *Econometrica*, 1982, *50* (6), 1345–1370.

— , Peter Rupert, and Roman Šustek, “Housing Dynamics over the Business Cycle,” *International Economic Review*, 2016, *57* (4), 1149–1177.

Lucca, David O., “Resuscitating Time-to-Build,” *mimeo*, 2007.

Lutz, Chandler and Ben Sand, "Highly Disaggregated Land Unavailability," *mimeo*, 2022.

Mayer, Christopher J. and C. Tsuriel Somerville, "Residential Construction: Using the Urban Growth Model to Estimate Housing Supply," *Journal of Urban Economics*, 2000, *48*, 85–109.

Meier, Matthias, "Supply Chain Disruptions, Time to Build, and the Business Cycle," *working paper*, 2020.

Mian, Atif and Amir Sufi, "What Explains the 2007-2009 Drop in Employment?," *Econometrica*, 2014, *82* (6), 2197–2223.

_ , Kamalesh Rao, and Amir Sufi, "Household Balance Sheets, Consumption, and the Economic Slump," *Quarterly Journal of Economics*, 2013, *128* (4), 1687–1726.

Millar, Jonathan N., Stephen D. Oliner, and Daniel E. Sichel, "Time-To-Plan Lags for Commercial Construction Projects," *Regional Science and Urban Economics*, 2016, *59*, 75–89.

Murphy, Alvin, "A Dynamic Model of Housing Supply," *American Economic Journal: Economic Policy*, November 2018, *10* (4), 243–267.

Nathanson, Charles G. and Eric Zwick, "Arrested Development: Theory and Evidence of Supply-Side Speculation in the Housing Market," *Journal of Finance*, December 2018, *73* (6), 2587–2633.

Oh, Hyunseung and Chamna Yoon, "Time to Build and the Real-Options Channel of Residential Investment," *Journal of Financial Economics*, 2020, *135* (1), 255–269.

Paciorek, Andrew, "Supply Constraints and Housing Market Dynamics," *Journal of Urban Economics*, 2013, *77*, 11–26.

Saiz, Albert, "The Geographic Determinants of Housing Supply," *Quarterly Journal of Economics*, 2010, *125* (3), 1253–1296.

Sarte, Pierre-Daniel, Felipe Schwartzman, and Thomas A. Lubik, "What Inventory Behavior Tells Us About How Business Cycles Have Changed," *Journal of Monetary Economics*, 2015, *76*, 264–283.

Topel, Robert and Sherwin Rosen, "Housing Investment in the United States," *Journal of Political Economy*, 1988, *96* (4), 718–740.