Time to build and the real-options channel of residential investment*

Hyunseung Oh† and Chamna Yoon‡

April 30, 2017

Abstract

A standard real-options model predicts that time-to-build investment decisions could be delayed by uncertainty over future revenue. This paper examines the first-order importance of this mechanism by looking into the micro-data for residential construction during the 2003–2009 boom and bust. We first demonstrate empirically the large time-to-build variations in this period. We then develop and estimate a model of sequential irreversible investment with stochastic bottlenecks. In the estimated model, the boom period increase in time-to-build is due to frequent construction bottlenecks whereas the fall in house prices and the rise in uncertainty during the bust generated further delays in construction.

Keywords: Investment; Housing; Real options.
JEL Classification Numbers: E22, E32, R31.

*We thank Rick Bond, Jeff Campbell, Yongsung Chang, Mario Crucini, Morris Davis, Thorsten Drautzburg, Matteo Iacoviello, Yannis Ioannides, James Kahn, Oleksiy Kryvtsov, Ricardo Reis, Peter Rousseau, Stephanie Schmitt-Grohe, Holger Sieg, Jon Steinsson, Gianluca Violante, Yildiray Yildirim, and seminar participants at FRB Philadelphia for their helpful comments. All errors are ours. First draft: December 17, 2015.
†Vanderbilt University. E-mail: hyunseung.oh@vanderbilt.edu.
‡Corresponding author, Sungkyunkwan University. Email: chamnayoon@skku.edu

1 Introduction

Many investment projects are irreversible, their future payoffs are uncertain, and the timing of investment is flexible. The real-options theory of investment considers optimal investment decisions under such conditions, as articulated in Dixit and Pindyck (1994). The models predict that the option value of waiting for new information influences the optimal investment decisions.

This paper investigates residential investment in the 2003–2009 housing cycle through the lens of the real-options mechanism. The novelty of this paper is that we look into an investment margin that has not been discussed in the recent housing cycle: the *time-to-build* of incomplete houses. Based on the US census micro-data on residential construction, we document new facts about investment in the recent housing cycle and study their implications for the general housing supply using a real-options model of time-to-build investment.

The real-options mechanism has been a focus of recent business-cycle literature. When uncertainty is high, the option value of waiting becomes high and investment might often be postponed in a standard real-option model. As shown in Jurado et al. (2015), measures of macroeconomic uncertainty spiked in the 2007–2009 Great Recession; as detailed in Bloom (2009), high uncertainty generates sizable slumps in investment and employment through the real-options channel. The rise in uncertainty and the induced wait-and-see effect have been suggested as drivers of several macroeconomic variables in the Great Recession (Bloom, 2014).

Although the significant cyclicality and volatility of the housing market were well-known before the Great Recession (Davis and Heathcote, 2005), the recent housing boom-bust cycle was unprecedented in size.[1] At the same time, house prices have been volatile, and measures of uncertainty in the housing market have been high during the bust period.[2] These facts motivate us to look further into the implications of uncertainty on residential investment dynamics in the recent housing market.

Residential investment is composed of both an extensive margin, new housing starts, and an intensive margin: the sequential construction of incomplete houses.[3]

[1] In fact, the decrease in real residential investment between the 4th quarter of 2007 and the 2nd quarter of 2009 accounts for about 33 percent of the decline in real GDP in the same period, based on the national income and product accounts (NIPA) table "Real Gross Domestic Product, Chained Dollars" from the Bureau of Economic Analysis.

[2] The beginning of section 4 describes these measures.

[3] New construction is by far the largest component of residential investment data in the GDP.

Most papers studying the recent housing market focus on new housing starts.[4] However, we argue that time-to-build investment - the intensive margin of residential construction - has also been crucial in shaping housing dynamics in the recent housing bust for the following three reasons.

First, the stock of incomplete houses is large. Building a house takes a certain amount of time. Even after building permits are issued, the average single-family house takes 6 months from start (i.e., excavation) to finish. Multi-family houses take around 10 months to build. As a result, for each housing unit started in a given month, 8.3 housing units are under construction.[5] With such a large stock of incomplete houses relative to new housing starts, even small variations in construction intensity can significantly affect the movement of residential investment.

Second, construction intensity affects housing start decisions. When forward-looking homebuilders decide to build a house, they take into account their expected investment intensity during the entire construction process.[6] Therefore, any shift in construction intensity should also affect new housing starts, and thereby residential investment.

Third and most important, time-to-build shifted significantly during the recent housing boom-bust period. The average time from start to finish for both single-family and multi-family houses was stable between 1984 and 2003. However, from 2003 to 2009, those times shifted upward by 2 and 4 months, respectively. This translates to a sizable $25-40\%$ slowdown in average construction intensity per house under construction.[7]

That the completion of a project could be timed is a main element underlying

---

Residential investment also includes improvements on existing houses and brokers' commissions and other ownership transfer costs, which are not the focus of this paper.

[4]Haughwout et al. (2013) give a comprehensive review of this literature and provide some data analysis.

[5]Data from December 1969 to December 2014.

[6]These forward-looking aspects, such as gestation lags or flow adjustment costs, have been widely used in the business-cycle literature to explain various dynamic responses observed in the data. For example, Christiano et al. (2005) introduce investment adjustment costs to generate hump-shaped investment responses to monetary shocks, and Uribe (1997) introduces gestation lags and convex adjustment costs to capital accumulation to generate the observed slow convergence of inflation between nontradables and tradables in an experiment of an exchange-rate-based stabilization plan. More recently, Arezki et al. (2017) also use gestation lags for an application of the effects of news shocks in an open economy.

[7]Figure 1 in section 2 plots the time series of time-to-build for both single-family and multi-family houses.

the sequential irreversible investment literature. Majd and Pindyck (1987) extend a one-period irreversible investment model to multiple periods and analyze investment behavior up to the completion of a project. In this model, investors can stop investment that has been initiated and time the completion date of a project. As uncertainty about the future rises, the tendency to delay a started project also increases. The observed decrease in construction intensity is consistent with the sequential irreversible investment model.

This model mechanism also captures the following narrative aspect of the single-family housing market in the recent bust. Single-family houses are smaller than multi-family houses and they typically take less than a year to finish. Therefore, when starting single-family houses, developers are likely to secure credit to cover the total construction costs. Nevertheless, developers might delay construction if they believe that property values will soon bottom out and demand for their houses will gradually pick up as the supply of new houses slows down in the market. Future housing market outcomes are highly uncertain because of the unpredictable nature of those turning points, and some developers could delay construction to take advantage.[8]

In the data section, we study the time-to-build (TTB) dynamics during the recent boom and bust. In section 2, we look into the time series of average TTB of residential buildings in the US. In section 3, we use census micro-data to examine the time series of the *distribution* of TTB for single-family residential buildings across the US. We find that average TTB increased during both the boom and bust but with different distributional properties. The mean TTB increase in the boom period is caused by an overall shift in the distribution. However, the further increase in the bust period is due to long deferrals for a few houses rather than an overall shift. For built-for-sale houses in particular, we find that long deferrals occurred only for houses that were unsold before the start of construction.

In section 4, we examine our data from a real-options perspective by developing a TTB model with irreversible investment (Majd and Pindyck, 1987) and stochastic bottlenecks. In section 5, we apply the model to the recent housing boom-bust cycle. Estimating the model using the observed TTB distribution and house price dynamics indicates that the real-options channel is an important driver for the dynamics of TTB during the recent bust. In addition to the movement of the median house price,

[8]This narrative is centered on the New York Times Economix blog "What's stalling New York's skyscrapers?" by Edward Glaeser (March 31, 2009) and the ensuing discussions.

the model suggests that uncertainty built up in the recent boom-bust period and bottleneck effects were the dominant force only during the boom period.

Based on the model mechanisms that account for the observed TTB dynamics, we study several other housing supply implications in section 6. In particular, we find that, when the intensive margin of investment falls, residential investment does not lag housing starts; its initial movement is dominated by the intensive margin rather than the extensive margin. Compared to a constant TTB model, the real-option TTB model predicts a deeper investment slump due to an overhang of incomplete housing inventory. We find that our model mechanism accounts for about 50 percent of the fall in housing starts and residential investment from 2003 to 2009, whereas the constant TTB model accounts for less than half of that. Section 7 concludes our argument.

Related literature The real-options channel of investment has been a major topic of interest in the literature of both macroeconomics and financial economics. The theoretical channels are well summarized in Dixit and Pindyck (1994), and their empirical applications are widespread across both fields, such as Leahy and Whited (1996), Moel and Tufano (2002), Bloom et al. (2007), Bulan et al. (2009), Bloom (2009), Bachmann et al. (2013), and Gilchrist et al. (2014).

However, extensions of the real-options insight to TTB investment decisions have been less explored. Following Kydland and Prescott (1982), TTB investment is typically assumed to be exogenous from business cycles.[9] Although this assumption looks innocuous most of the time, the Great Recession was a period in which TTB itself shifted a lot.

We are the first to document this using micro-data on residential TTB across the US. Because the TTB of each building is a result of dynamic investment decisions, a structural model is appropriate to study its further implications. Therefore, we develop a structural model of both endogenous and exogenous TTB investment based on Majd and Pindyck (1987) and discipline the model using empirical moments constructed from the micro-data. From the estimated model, we conduct counterfactual experiments to quantify the causes of the slowdown in TTB investment during the drastic fall in residential investment in the recent housing bust. These findings contribute to the real-options literature by extending the insights of the uncertainty

[9]Campbell (1998), Lucca (2007) and Edge (2007) are good examples.

Figure 1: TTB for completed houses in the US

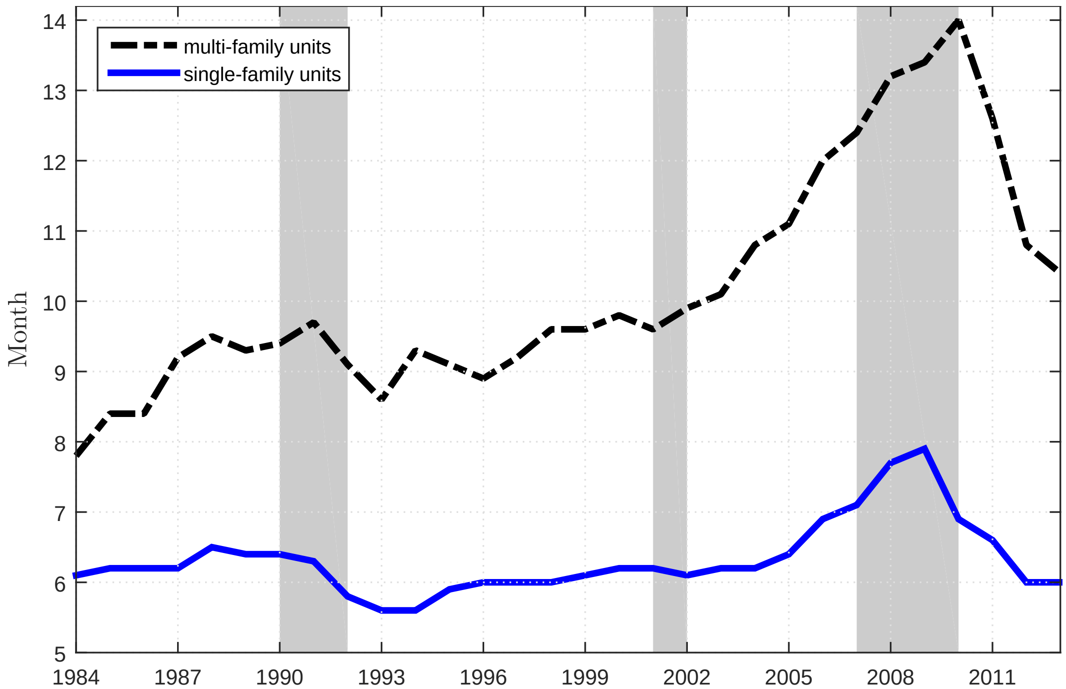

Note: Shaded areas are NBER recessions.

effects on investment into the intensive margin.[10]

By documenting new findings on residential construction and exploring their housing supply implications, this article also contributes to the housing investment literature, such as Topel and Rosen (1988), Iacoviello (2005), Glaeser et al. (2005), Leamer (2007), Glaeser et al. (2008), Saiz (2010), and Kydland et al. (2016).

2 Aggregate time-to-build data

In this paper, we use Survey of Construction data. It is available from the Census Bureau, which is a national sample survey (sampling rate: 1/50) of the builders and owners of new houses. The dataset contains information on the building and geographic characteristics of new houses across the US in each survey year, including the starting and completion months of the houses, sales prices, and the month in

[10]That extensive and intensive investment might deserve separate attention is also described in Jovanovic and Rousseau (2014) where they argue that cyclical investment behavior differs between established and new firms.

which each house was sold (if sold), along with square footage, number of rooms, and so on. Houses authorized by building permits but not started at the end of the year or under construction at the end of the year are also included.[11]

Based on this dataset, the census reports the aggregate series for "average length of time from start to completion," for both single-family and multi-family units. In each given year, this series thus reveals information on the average TTB of completed houses. In this section, we examine this aggregate series. In section 3, we discuss the underlying micro-data to get a deeper understanding of the findings in this section.

2.1 Aggregate facts

As described in the introduction, Figure 1 depicts the average time from start to completion of both single-family and multi-family houses from 1984 to 2013. The average TTB for single-family and multi-family houses in 1984–2013 was 6 and 9.7 months, respectively.

However, from 2003 to 2006, TTB for single-family and multi-family houses increased by 1 and 2 months, respectively. One possible contributor to this increase is the difficulty of hiring construction workers during the housing boom period.[12] In Figure 2, we plot several construction-sector time series. We observe that construction activity, such as housing starts and construction employment, surged in this period. We also plot two measures of bottlenecks in the construction sector: (i) construction sector unemployment rate and (ii) construction sector labor market tightness, which is the job openings to unemployment ratio.[13] The low unemployment rate and high labor market tightness support the view that hiring workers in the construction sector during the housing boom period was a bottleneck.[14]

Third, from 2006 to 2009, average TTB further increased for both single-family and multi-family houses, again by 1 and 2 months, respectively.[15] However, measures

[11]Houses for which construction was abandoned after permit issuance or after start are not included.

[12]For example, Kalouptsidi (2014) finds that higher investment activity lengthened the production lag of the capacity-constrained Greek ship bulking industry.

[13]Construction sector unemployment rate is reported in the Bureau of Labor Statistics (BLS) website. This series is based on the household survey where the number of unemployed people in the construction sector is based on unemployed households that report their previous job in the construction sector.

[14]This view is expressed in, for example, Green et al. (2005).

[15]The average TTB for multi-family houses peaked in 2010.

Figure 2: Construction sector variables during the housing boom and bust

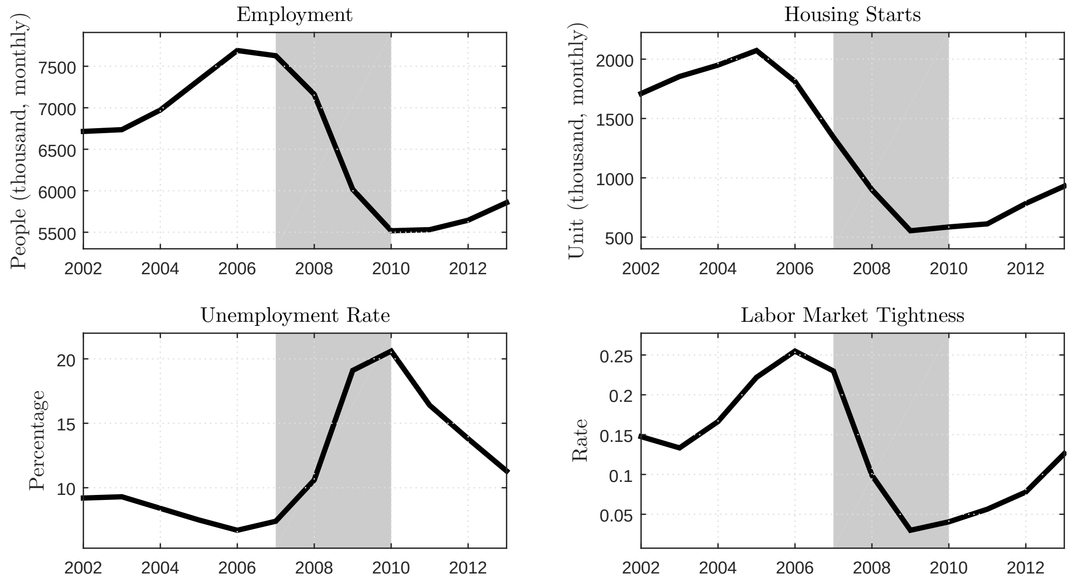

Note: Labor market tightness is the rate between construction job openings and unemployment with previous jobs from the construction sector. Shaded areas are NBER recessions.

of bottlenecks in the construction sector no longer support the view that bottlenecks were the main contributor to this further increase in TTB. In fact, as shown in Figure 2, the housing market was entering a bust period with a dramatic fall in both housing starts and construction employment, and all bottlenecks were resolved. This suggests the existence of a strong mechanism that overcame the negative bottleneck effects on TTB dynamics during the housing bust.

After 2010, average TTB returned to pre-2003 level.

2.2 Implications of aggregate facts

In a standard business-cycle model of investment, TTB is assumed to be constant. This is empirically supported in the period from 1984 to 2003. Although the economy experienced two NBER recessions during this period, TTB showed little variation.

That assumption no longer holds in the recent housing boom-bust cycle, during which average TTB increased by $25 - 40\%$ relative to the 1984–2003 average. Holding fixed the number of ongoing projects, this translates to a $25 - 40\%$ decrease in residential investment. The initial increase in the boom is consistent with bottleneck mechanisms, and we find that the labor market of the construction sector was indeed tight during this period.

However, a puzzling feature is the further increase in TTB during the bust. The data suggest that construction sector bottlenecks were resolved in the housing bust period, and that should create a decrease in average TTB. Therefore, the further increase in TTB during the bust period suggests that bottleneck theories were no longer the main action. Why did TTB increase during the housing bust? In the next section, we investigate this question using the underlying micro-data for TTB.

3 Micro-data analysis

Our goal in this section is to understand the dynamics of the *distribution* of TTB during the boom-bust period. To that end, we use the underlying micro-data for the census statistics on average TTB. Our analysis is based on the distribution of TTB for *single-family houses*, for which data are publicly available since 2000.[16] We first construct a measure of economic TTB by controlling for the geographic and building characteristics of each completed house, and then compare its cross-sectional distribution across time.

3.1 Economic time-to-build

All buildings are different, and the TTB for each house depends on various factors. For example, a large, difficult-to-build house will require lengthy construction time. Other factors for lengthy construction are houses built on locations with severe weather conditions and stringent regulations.

For the purpose of this section, we focus on the dynamics of TTB that are independent from these geographic and building characteristics.[17] Because the micro-data provide many of these features, we use them to construct a measure of economic TTB for each completed single-family building in the US. In particular, for geographical

[16]We use houses started since January 2000 for our sample.

[17]We understand that building and geographic characteristics can be correlated with economic conditions; for example, larger houses might be built during boom periods. Therefore, our estimate should be taken as a conservative measure. However, because our focus in this paper is the dynamics of TTB while controlling for housing start decisions, we find it best to remain silent on these possible correlations.

characteristics, we control for the 9 census divisions and whether the house is built in a metropolitan area, which is the finest level of geographical information available in the public data. For building characteristics, the list of control variables includes building purpose (owner built, contractor built, built for sale, built for rent), building method (site built, panelized, modular), and the square footage of the house.

We regress the log of TTB on those various control variables.[18] Table 1 reports the results from this regression using data from 2000 to 2013. Although the main focus is not on understanding the link between the control variables and TTB, several interesting results are worth mentioning. The first column summarizes the frequency of the sample. The division with the highest number of completed houses during the sample period is South Atlantic, accounting for 26.9% of the total sample. The smallest number of houses were built in New England (3.4%). For building purposes, built-for-sale houses compose the majority of the sample (74%), followed by contractor-built houses (14.2%) and owner-built houses (8.7%). For single-family units, built-for-rent houses are only a small portion (3.2%).

Looking into regression (a), notice first that the New England, Middle Atlantic, and Pacific divisions show longer TTB than the other divisions. In particular, TTB is on average 27% higher in the Middle Atlantic than it is in the West South Central division. Second, owner-built houses have longer TTB than contractor-built houses, which could reflect either the efficiency of contractors or the inclusion of various housing amenity preferences into owner-built houses. Built-for-sale houses take the least time to build. Third, site-built houses have longer TTB than panelized or modular houses, which could reflect exposure to bad weather conditions. Fourth, 2 and more story buildings take more TTB than 1 story buildings, and the square footage of a house has a positive relation with TTB.

In our subsequent analyses, we use the residual of regression (a) of Table 1 as our measure of economic TTB. Regression (b) adds the division-level unemployment rate as an additional variable to control for bottleneck effects. This regression will be discussed later.

[18] The qualitative discussions are similar even if we regress on the level of TTB instead of the log. Because a log regression provides estimates in percentage terms, we use this as a baseline. The level TTB regression is provided in the Appendix.

Table 1: Regression on TTB ($\log TTB$)

	Frequency	(a)		(b)	
New England	0.034	0.259	(0.00969)	0.242	(0.00966)
Middle Atlantic	0.071	0.270	(0.00736)	0.271	(0.00733)
East North Central	0.124	0.162	(0.00610)	0.183	(0.00609)
West North Central	0.076	0.130	(0.00718)	0.0957	(0.00719)
South Atlantic	0.269	0.0599	(0.00478)	0.0453	(0.00477)
East South Central	0.053	0.0813	(0.00640)	0.0985	(0.00641)
West South Central	0.144	-		-	
Mountain	0.098	0.0388	(0.00625)	0.0223	(0.00621)
Pacific	0.132	0.202	(0.00611)	0.233	(0.00608)
Built-for-sale	0.740	– 0.210	(0.00828)	– 0.220	(0.00828)
Contractor-built	0.142	– 0.0578	(0.00887)	– 0.0633	(0.00887)
Owner-built	0.087	0.166	(0.00982)	0.159	(0.00983)
Build-for-rent	0.032	-		-	
Modular	0.028	– 0.532	(0.0106)	– 0.539	(0.0107)
Panelized	0.024	– 0.144	(0.00804)	– 0.143	(0.00802)
Site built	0.948	-		-	
1 Story	0.430	-		-	
2+ Story	0.570	0.00628	(0.00271)	0.00531	(0.00269)
Square feet (×100)		0.00825	(0.000180)	0.00829	(0.000179)
Unemployment rate				– 0.0406	(0.000699)
Constant		2.025	(0.0129)	2.250	(0.0134)
Other controls:					
Metropolitan area		yes		yes	
Number of full bath		yes		yes	
Detached		yes		yes	
Deck		yes		yes	
Parking facility		yes		yes	
Foundation		yes		yes	
Material of wall		yes		yes	
Observations		261,184		261,184	

Note: Robust standard errors in parentheses.

3.2 The distribution of economic time-to-build

Using our measure of economic TTB, we compare their cross-sectional distribution of TTB across three different periods: steady state (2000–03), housing boom (2004–06), and the subsequent bust (2007–09). For a clearer exposition of our argument, we choose 2003 as our steady state, 2005 as the housing boom, and 2009 as the

Figure 3: Distribution of economic TTB (Total sample)

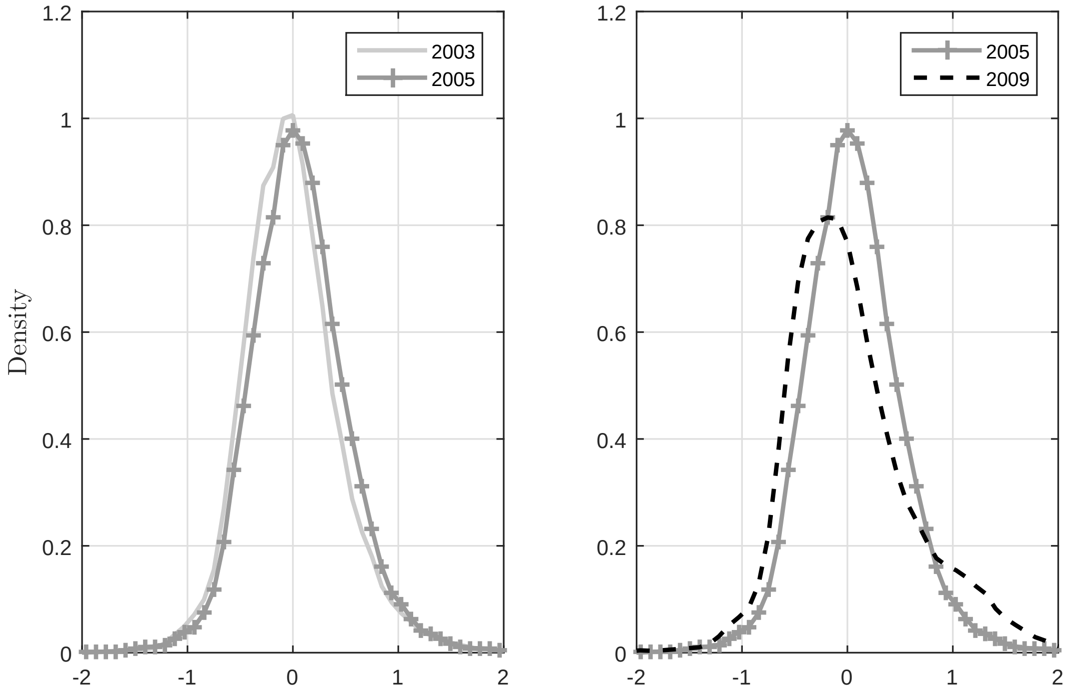

Note: Kernel density of economic TTB for all single-family houses in log values. Left compares 2003 and 2005, and right compares 2005 and 2009.

subsequent bust.

The left panel of Figure 3 compares the kernel density of economic TTB in 2003 and 2005. We observe that, between these years, there was an overall shift to the right of this distribution. Economic TTB increased for all types of houses, and bottleneck theories are consistent with this distributional shift.[19]

On the other hand, the right panel of Figure 3 compares the kernel density of economic TTB in 2005 and 2009. Two patterns are apparent. First, the mass of the distribution (including the mode) shifted back to the left. Second, a fat tail to the right appeared. The left shift in the mass of the distribution is consistent with resolution of bottlenecks. During the housing bust, the supply side have had less trouble finding available construction workers to build a house. Nevertheless, the fat tail to the right indicates that incomplete houses during the housing bust remained

[19]A two-sample Kolmogorov-Smirnov test rejects a common underlying distribution for the economic TTB in 2003 and 2005 (p-value < 0.001). In fact, all distributions in Figures 3, 4, and 5 are statistically different based on this test.

Figure 4: Distribution of economic TTB in 2005 and 2009

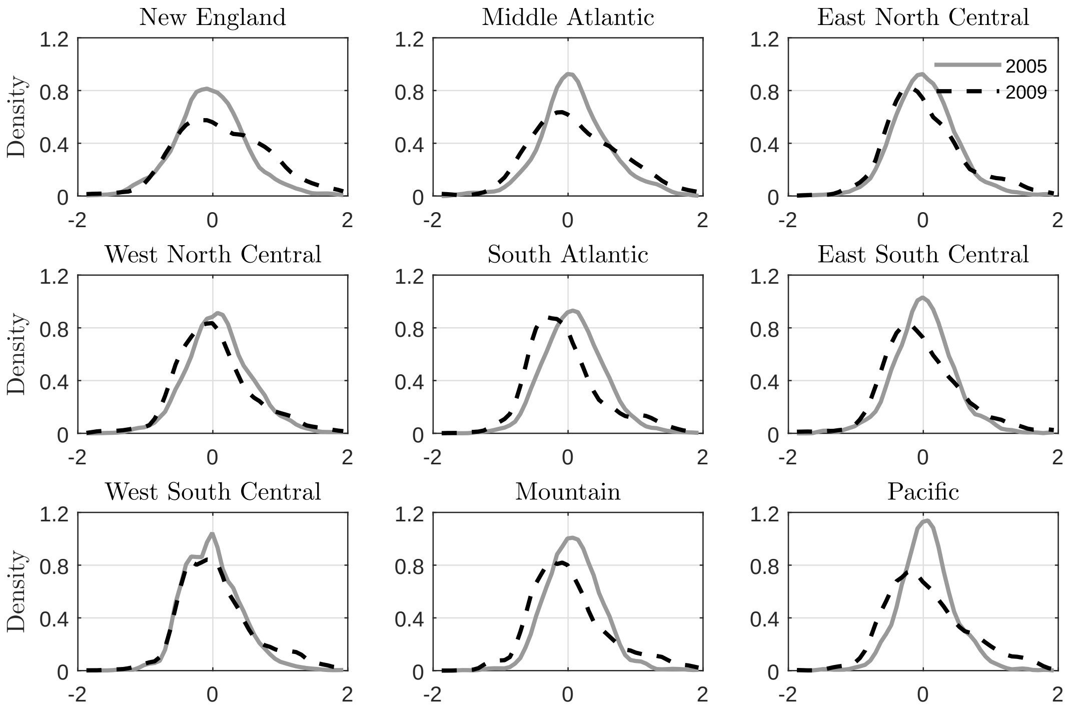

Note: Kernel density of economic TTB for all single-family houses in log values.

under construction for a long period. Linking this observation to the aggregate fact, the increase in TTB from this right fat tail dominated the negative bottleneck effect; hence, average TTB increased further in 2009 compared to 2005.

In Figure 4, we plot the kernel density of economic TTB in 2005 and 2009, across the 9 census divisions. The same two patterns - (i) mass of the distribution shifting to the left; (ii) fat tail to the right - hold in most divisions. Therefore, the dynamics of residential contruction lags between 2005 and 2009 appear to be a national phenomenon rather than a locally driven one.

3.3 Subsample evidence: Built-for-sale

From the overall distribution, we observe long deferrals in the construction of a portion of incomplete houses during the bust period. In this section, we narrow our focus to the built-for-sale subsample of the distribution, to better understand where these long deferrals disproportionately arise. We examine the built-for-sale subsample for two reasons: (i) built-for-sale houses contain information on the sales month and price that we can use to better infer the economic channels; (ii) built-for-sale houses compose the majority of single-family houses in the sample (74%) as shown in Table 1, when a significant fat tail to the right is observed for this category.

Based on the sales month information for these houses, we split this sample into two categories: houses sold before the start of construction and houses unsold before the start of construction. In Figure 5, only one pattern arises for houses sold before the start of construction: the distribution shifts back to the left. For this category of houses, bottleneck effects must have been the main force in the dynamics of TTB. On the other hand, for houses that were unsold before the start of construction, the fat tail to the right is pronounced in the 2009 distribution. Therefore, almost all of the increase in the right fat tail in this subsample comes from houses that were unsold before the start of construction.

For houses unsold before the start of construction, we break the data down further and compute the average economic TTB based on the time lag between the start date for construction and the sale date for 2005 and 2009, as shown in the right panel of Figure 5. In this figure, houses that were sold before the start of construction are all lumped into 0, and houses that were sold after 24 months or never sold are lumped into 24. Sales month lag and average TTB are positively correlated. That is, houses unsold for a long time tend to have lengthy TTB. In 2009, this pattern becomes more apparent. Comparing 2005 and 2009, houses that were sold during the early stage of construction were actually built faster in 2009. However, houses in the bust period that faced difficulty in selling early took a much longer time to complete. Long deferrals in 2009 were especially concentrated on houses that took a long time to sell, even compared to those in 2005.

3.4 Economic time-to-build and bottlenecks

To get a quantitative sense of the increase in economic TTB for built-for-sale houses, we compare the average TTB in the census data, our measure of economic TTB, and our measure of ex-bottleneck TTB that is explained below. The results are plotted in Figure 6.

First, considering the raw census data, we find that average TTB increased by 29% between 2003 and 2009. In particular, average TTB increased 12% by 2006, with

Figure 5: Distribution of economic TTB (built-for-sale)

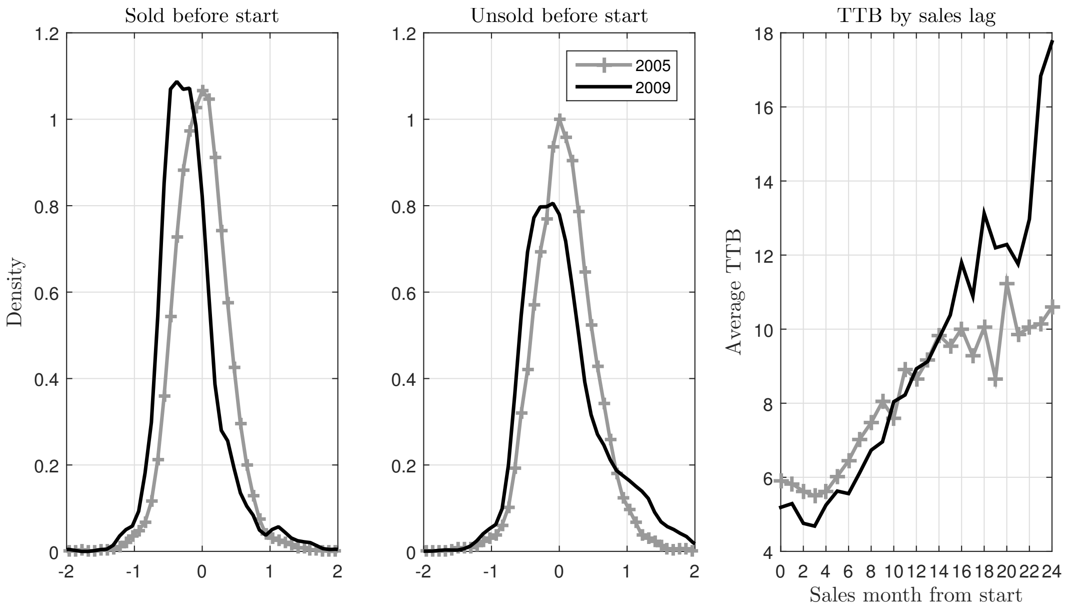

Note: Left two figures are kernel density of economic TTB for built-for-sale houses in log values, for 2005 and 2009. The right figure plots the average economic TTB by sales month from the start of construction for 2005 and 2009. Houses sold before the start of construction are lumped into 0, and 24 includes all houses that were sold after 24 months or not sold.

the remaining 17% increase occurring from 2006 to 2009. Therefore, for built-for-sale houses, the increase in average TTB was concentrated in the housing bust rather than in the boom.

Second, looking into the average of our economic TTB, we find that TTB increased by 22% between 2003 and 2009, which is smaller than the increase observed in the raw data. Because we control for geographic and building characteristics in our constructed measure, this result implies that some of the increase in TTB in the raw data originates from those characteristics. That is, houses built during the recent boom-bust era might be larger and of higher quality or built in regions where houses take longer to build. From 2003 to 2006, the increase was 12%, and the remaining 10% increase occurred from 2006 to 2009.

Third, because our subsequent analysis aims at understanding the increase in TTB during the bust, we also control for bottleneck effects in our economic TTB by regressing our series with bottleneck measures. Unfortunately, it is difficult to

Figure 6: Three measures of mean TTB for built-for-sale single-family houses

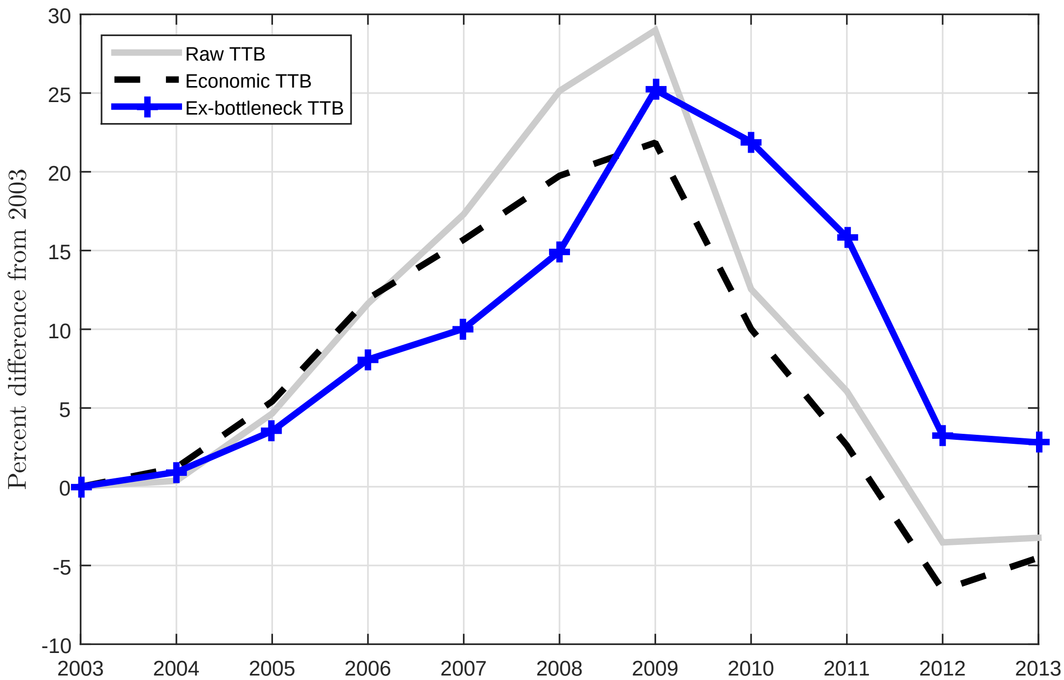

find direct measures of bottlenecks in the construction sector at the regional level. Our approach is to use the unemployment rate as a proxy for regional bottlenecks. Specifically, compared to our measure of economic TTB, we also control for the average division-level unemployment rate for the first three months from the starting month of construction and compute its residual, as shown in regression (b) of Table 1. Based on this ex-bottleneck measure of TTB, we find that total TTB increased by 25% from 2003 to 2009, and that during the boom, the increase was only 8%, whereas during the bust, the increase was 17%. That is, taking into account that bottlenecks were infrequent during the bust period, the increase in TTB unrelated to bottlenecks was large and quantitatively comparable to the overall increase in TTB from 2003 to 2009.

3.5 Summary

We summarize our findings in this section as follows. First, the increase in TTB during the boom stemmed from a shift to the right of the overall distribution. Second, the further increase in TTB during the bust was a combination of two counteracting forces: a shift to the left of the mass of the distribution, and a fat tail to the right. The right fat tail was the dominating force and causing the mean TTB to increase further. To study the potential channels of this increase in TTB, we focus on houses built-for-sale and find the same pattern. Within this sample, we compare houses sold and unsold before the start of construction. We find that the fat tail to the right only appears for houses that were unsold before the start of construction. In fact, during the housing bust, average TTB decreased for houses that were sold before the start of construction. Controlling for geographic and building characteristics, as well as for bottleneck effects, mean TTB increased by 25% from 2003 to 2009.

4 Model of time-to-build

Based on those empirical findings, our next task is to understand the economic forces that could deliver the observed increase in TTB especially during the recent bust. The micro-level data for residential TTB have been unexplored in the literature, and the housing supply variables that we have recovered so far involve complicated dynamics. Therefore, the first step to analyze our findings is to filter them through a baseline dynamic model and devise a structural interpretation of the observed distributional facts.

To that end, we illustrate a TTB model of residential investment that incorporates bottlenecks and the real-options mechanism and quantify the fall in construction intensity during the recent housing boom and bust. The section starts with a discussion of the real-options mechanism in the data and then illustrates the model.

4.1 Why the real-options mechanism?

As discussed in the Introduction, the real-options model is a natural fit for analyzing residential construction in the recent cycle for two broad reasons. First, housing supply decisions generally involve significant irreversible costs, such as resources spent on acquiring permits and building foundations and the time spent on construction. The irreversible resources and time required by these investments introduce a significant option value not only at the beginning of construction, but also at the continuation stage. As shown in Table 2, homebuilders report a large amount of spending in the later stages of construction. In fact, only 16.3% of the total con-

Table 2: Single-family house construction cost breakdown (2013)

Stage of Construction	Cumulative
1. Site Work (Permit, Inspections, Architecture)	6.8%
2. Foundation (Excavation, Concrete)	16.3%
3. Framing (Roof, Metal, Steel)	35.4%
4. Exterior Finishes (Wall, Windows, Doors)	49.8%
5. Major Systems Rough-in (Plumbing, Electrical, HVAC)	63.2%
6. Interior Finishing (Insulation, Painting, Lighting, Flooring)	92.5%
7. Final Steps (Landscaping, Outdoor Structures, Clean Up)	99.1%
8. Other	100.0%

Note: Survey data available from the National Association of Home Builders

struction cost has been spent when the foundation is completed; hence, continuing the project from that stage requires significant resources and time.

Second, the dynamics of the housing market in the recent bust period show features consistent with those seen when the real-options mechanism operates. House prices fell significantly, and the measured uncertainty surrounding the housing market increased signficantly. For completeness, we use three datasets to confirm this statement: (i) the monthly purchase-only house price index published by the Federal Housing Finance Agency, (ii) the monthly S&P/Case-Shiller national home price index, and (iii) the daily Philadelphia Stock Exchange housing sector index. On the left panel of Figure 7, the two home price indices are plotted, each divided by the monthly consumer price index. The fall in both house-price indices is steep and large during the recent bust. On the right panel, both the previous 6-month moving-average standard deviation of the monthly growth rate of house prices (the blue solid line) and the standard deviation of daily stock returns for each month (the black dashed line) are plotted. During the recession period, both the house-price growth-rate volatility and stock-return volatility increased significantly. Those large movements in house price and housing market uncertainty motivate us to study the first-order importance of the real-options mechanism in housing supply decisions.[20]

[20] However, it is important to note here that our focus will be on new construction rather than the overall housing market. The overall housing market condition depends disproportionately on transactions of pre-existing houses.

Figure 7: Price and uncertainty measures of the overall housing market

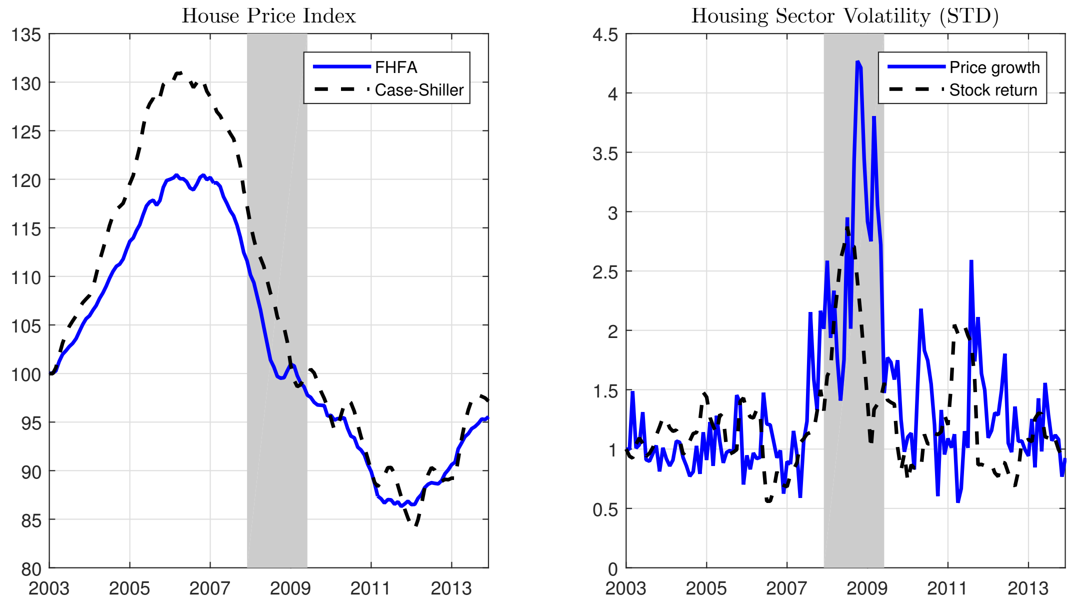

Note: For the house price index, we use both the monthly Federal Housing Finance Agency (FHFA) purchase-only index and the monthly S&P/Case-Shiller national home price index. We also compute two volatility measures. First, using the FHFA index to compute the monthly home price growth rate, we plot the standard deviation for the previous 6 months. For the stock price index, we use the daily Philadelphia Stock Exchange housing sector index and calculate the standard deviation of daily stock returns for each month. All measures are normalized at their January 2003 values and divided by the monthly consumer price index. Shaded areas are NBER recessions.

4.2 Model details

The model we lay out incorporates both the standard fixed TTB assumption as in the business-cycle literature and the real-options channel for TTB investment. Through the model, we investigate the first-order importance of the real-options view of TTB investment. The model considers a case in which a project takes time to complete with multiple irreversible stages. Payoff occurs only after the project is completed; hence, irreversible investment decisions are made sequentially at each stage with uncertainty about the future payoff.

The model blends three key elements that are widely discussed in the investment literature. First, a real-options channel is introduced by illustrating a discrete-time version of the sequential irreversible investment model of Majd and Pindyck (1987) and Dixit and Pindyck (1994). Second, we study the effect of uncertainty shifts on investment following the work of Bloom (2009). Third, the model incorporates stochastic bottleneck probabilities in investment to control for TTB originating from either input (labor and material) bottlenecks or weather effects.

In the extension in section 6, we also study the investment implications relative to the fixed TTB assumption of Kydland and Prescott (1982).

TTB and investment. A house takes a total real investment of $\bar{K}$. In each period, the maximum level of investment is κ. If $\kappa \geq \bar{K}$, then it is possible to complete the project in one period. However, if $\kappa < \bar{K}$, then the investment takes time to build, with physical TTB being $\bar{K}/\kappa$.

There is also a bottleneck probability for ongoing projects that follows a Poisson process. In each period, bottlenecks occur with probability p_c. Investment is delayed when bottlenecks occur, and the expected minimum TTB with bottlenecks is

$$1 + \frac{\bar{K}/\kappa - 1}{1 - p_c}.$$

Note that, without bottlenecks ($p_c = 0$), the expected minimum TTB is equal to the physical TTB.

Price dynamics. The value of a completed house i in period t is denoted P_{it}. This value is determined by both a macro price factor P^M and a construction unit idiosyncratic factor P^U_{it}:

$$P_{it} = P^M \times P^U_{it}.$$

The stochastic process for P^U_{it} is

$$log(P^U_{it}) = log(P^U_{it-1}) - \frac{\sigma^2}{2} + \sigma W_{it}, \quad W_{it} \sim N(0,1), \tag{1}$$

where σ is the level of uncertainty in this price process, and W_{it} is the idiosyncratic price innovation term. Note that, in this process, the mean growth rate of the idiosyncratic price factor is zero regardless of the level of uncertainty. For construction that has not started, we normalize the previous value ($P^U_{it-1} = 1$).

In the price process, P^M and σ are assumed to be constant. Therefore, builders in the model face house-price movements driven by time-varying idiosyncratic price factors rather than the macro price or uncertainty factors. Although this assumption might be reasonable in the short run, macro-level movements are also relevant in the medium run, especially during the recent boom-bust cycle. Later in the simulation, we allow the macro price and uncertainty factors to be time-varying across regimes and study the role of these factors in the housing boom and bust.

Bellman equation. The three state/choice variables for builders are K_{it}, I_{it} and B_{it}. First, K_{it} denotes the total remaining capital for completion of house i in period t. When $K_{it} = \bar{K}$, the house is yet to be started and the investment decision is on the extensive margin. On the other hand, when $K_{it} < \bar{K}$, then the homebuilder decides on the construction intensity of an existing project. Second, I_{it} is the flow cost of investment for house i in period t. The upper bound of this value is κ, which is our TTB constraint. Third, B_{it} is an indicator function for a construction bottleneck of house i in period t.

With these variables, the value $V(\cdot)$ of an incomplete house i with a current construction unit price factor P^U_{it}, $K_{it} > 0$ remaining capital to completion, previous investment of I_{it-1}, and bottleneck indicator B_{it} is

$$V(P^U_{it}, K_{it}, I_{it-1}, B_{it}; \Lambda) = \max_{I_{it}\in\{0,\kappa\}} \Big\{ -(1 - B_{it})I_{it} - \gamma 1_{\{K_{it}=\bar{K} \ \& \ I_{it}>0\}} + \left(\frac{1}{1+r}\right) \mathbb{E}V(P^U_{it+1}, K_{it+1}, I_{it}, B_{it+1}; \Lambda)\Big\},$$

where $\Lambda = \{p_c, P^M, \sigma\}$. Without bottlenecks ($B_{it} = 0$), the cost of real investment in each stage is I_{it}. Additionally, when building a new house ($K_{it} = \bar{K}$), the builder pays a fixed cost γ to start construction ($I_{it} > 0$). This cost parameter incorporates the various sunk costs in the decision to start a house, such as obtaining building permits. TTB investment decisions are discrete (either invest κ or not).[21]

[21] In a continuous-time model without adjustment costs, this bang-bang type of TTB investment turns out to be the solution of the model even when a continuous range is considered (Dixit and Pindyck, 1994). In general, this discrete investment decision assumption could be relaxed by solving the model with a higher frequency and aggregating across time.

The evolution of the state variables is

$$K_{it+1} = K_{it} - (1 - B_{it})I_{it},$$

$$B_{it} = \begin{cases} 1 & \text{with prob} \quad p_c \\ 0 & \text{with prob} \quad 1 - p_c. \end{cases}$$

Without bottlenecks ($B_{it} = 0$), the benefit of real investment is the reduction of the remaining capital to completion. In each period, bottlenecks occur with probability p_c. When there is a bottleneck, the construction of capital is deferred.

Lastly, the value of a completed house i ($K_{it} = 0$) is

$$V(P_{it}^U, K_{it}, I_{it-1}, B_{it}; \Lambda) = P^M P_{it}^U.$$

Hence, when a house is finished, the builder earns the market price of the house.

5 Model solution and simulation

The model is calibrated monthly. After calibrating some parameters of the model, we discuss the optimal investment decision of builders under different levels of uncertainty. Afterward, we simulate the model and estimate the remaining parameters to match the simulated TTB distribution to the empirical TTB distribution given in section 3.

5.1 Calibration

We start by calibrating the net monthly interest rate r and the physical TTB $\bar{K}/\kappa$.[22] As a benchmark, we set r such that the annual interest rate is 10 percent (Bloom et al., 2007).[23] The physical TTB is set at 4 months because, in the empirical TTB distribution, fewer than 10 percent of houses are built within 3 months.

For the fixed cost of starting construction, we set $\gamma = 0.073\bar{K}$, which implies that

[22] The overall construction cost $\bar{K}$ is normalized at 1.

[23] The model has also been solved and estimated with different interest rates in the range of $5 - 10\%$ and the main lessons are robust. In the Appendix, we also present the estimation result of the model under a 6.5 percent annual interest rate (Bloom, 2009).

the initial sunk cost is 6.8 percent of total construction spending (Table 2).

The three remaining parameters are p_c, P^M, and σ. To recall, p_c is the bottleneck probability, P^M is the macro price factor, and σ is the uncertainty faced by builders. These three parameters will be estimated using several key moments in the empirical TTB distribution after we discuss the solution of the model.

5.2 Model solution and the effect of uncertainty

The solution is characterized by a threshold price above which the builder invests. For each K_i stage of construction, the optimal threshold price, $P^*(K_i;\Lambda)$, is solved. The investment decision of the builder is

$$I_{it} = \begin{cases} \kappa & \text{if } P_{it}^U \geq P^*(K_i;\Lambda), \\ 0 & \text{if } P_{it}^U < P^*(K_i;\Lambda). \end{cases} \tag{2}$$

where $\Lambda = \{p_c, P^M, \sigma\}$. Notice that $P^*(1;\Lambda)$ indicates the threshold price for housing start decisions, and $P^*(K_i;\Lambda)$ for $K < 1$ refers to the threshold price for TTB investment decisions with remaining construction of K_i. Recall that for a housing start decision, the idiosyncratic price is drawn from a log-normal distribution with its mean at the macro price factor. Therefore, the builder is not forced to start construction at a pre-constructed structure or an area with a history of bad prices. On the other hand, for the TTB investment decision on incomplete construction, the idiosyncratic price is drawn based on its previous price. That is, once a housing start decision is made and foundations are laid, the building is locked into its own history of shocks.

Before moving on, we note that builders are not being forced to sell the house immediately after completion. Even if we allow the builders in the model to hold an option to sell the completed house later, they will not exercise that option. Instead, the builders in our model are optimally choosing to sell a completed project immediately.[24] The two views lead to different interpretations of the data. That is, builders in our model are extending TTB not because sales are forced, but rather because the irreversibility option is lost for a completed house.

[24]Technically, the expected future price equals the current observed price in our model. At completion, since the irreversibility option no longer exists and prices are not projected to be better in the future, the builder prefers to sell today rather than pay the opportunity cost of waiting, which is the interest payment.

Numerical example. The three parameters $\{p_c, P^M, \sigma\}$ will be estimated in the next section. For illustrative purposes, in this part we numerically solve the model by assuming zero bottlenecks ($p_c = 0$) and $P^M = 2$. For annual σ, the following 3 values are considered: 0.2, 0.4, 0.6. We also momentarily set the physical TTB at 5 months instead of 4 months to better understand TTB investment decisions.

The first observation is that, when uncertainty is higher, both housing starts (for $K = 1$) and TTB investment (for $K < 1$) require a higher threshold price:

$$P^*(K; \sigma = 0.6) > P^*(K; \sigma = 0.4) > P^*(K; \sigma = 0.2), \qquad \forall K.$$

Consistent with the literature on irreversible investment and uncertainty, the option value increases with the level of uncertainty and housing start decisions require a higher price.

Second, conditional on starting a house, the TTB investment threshold price falls as the construction process nears completion:

$$P^*\left(\frac{\bar{K} - \kappa}{\bar{K}}; \sigma\right) > P^*\left(\frac{\bar{K} - 2\kappa}{\bar{K}}; \sigma\right) > \cdots > P^*\left(\frac{2\kappa}{\bar{K}}; \sigma\right) > P^*\left(\frac{\kappa}{\bar{K}}; \sigma\right), \qquad \forall \sigma.$$

Intuitively, as TTB investment occurs, the future payoff period becomes closer to realization, which lowers the option value.

Discussion. These two results are also consistent with the findings of the sequential irreversible investment models outlined in Dixit and Pindyck (1994). However, these two observations do not give us a clear answer for how TTB investment decisions differ across different levels of uncertainty.

On the one hand, a homebuilder facing high uncertainty could be cautious in new construction and set a very high threshold price to start. Because the builder enters only with a price observation that is high enough, it is unlikely that the builder will delay construction. Therefore, for buildings started during highly uncertain periods, TTB might still remain relatively constant.

On the other hand, a homebuilder could also take an optimistic view in starting construction under high uncertainty, since a highly profitable outcome is possible. When that outcome occurs, it might be too late for the builder to reap that profit when the builder has not already started the project, because completion takes time.

In that case, the builder might want to bet on a good realization and engage in wait-and-see behavior after building at least the foundation of the house. Therefore, TTB could increase for buildings started after observing high uncertainty.

Uncertainty and TTB. To address those possibilities, we now examine whether TTB investment is sensitive to uncertainty after controlling for housing starts. Using our model, we ask whether TTB investment decisions become more cautious than housing start decisions during periods of high uncertainty.

In Figure 8, we plot the log of the threshold price of investment in each stage of construction relative to its threshold price of investment when the project started, scaled by the level of uncertainty:

$$W_K^* - W_1^* = \frac{1}{\sigma}\big[\log P^*(K;\sigma) - \log P^*(1;\sigma)\big]. \tag{3}$$

The left-hand side variables W_K^* and W_1^* are the threshold price innovation values for TTB investment at stage K and housing starts, respectively, in terms of the standard normal distribution assumed in (1). Because the scale of these threshold price innovation values is uncertainty-invariant, the above measure allows us to compare the threshold price of investment in each stage relative to the threshold price for starts, across different levels of uncertainty.

In the figure with low uncertainty ($\sigma = 0.2$), the relative threshold price of investment falls steeply as construction approaches completion. At the last stage of construction, the builder needs to observe a 9-standard-deviation decrease in the price innovation relative to the initial price innovation to defer the project. Therefore, once a project starts, unless the price innovation falls by a significant amount, the builder continues investing in the project until completion.

On the other hand, with medium uncertainty ($\sigma = 0.4$), the fall in the relative threshold price of investment is only gradual. Even at the last stage of construction, the builder defers the project when observing a 3-standard-deviation decrease in the price innovation relative to the initial price innovation. In fact, with high uncertainty ($\sigma = 0.6$), the threshold price of investment for first-stage TTB investment becomes even higher than the threshold price for starts.[25] Therefore, when uncertainty is

[25] This non-monotone response is due to the additional loss in option for initial construction decisions. Although the unit prices are based on the aggregate price level for new construction, for

Figure 8: Uncertainty-invariant threshold price of investment relative to that at the start of construction

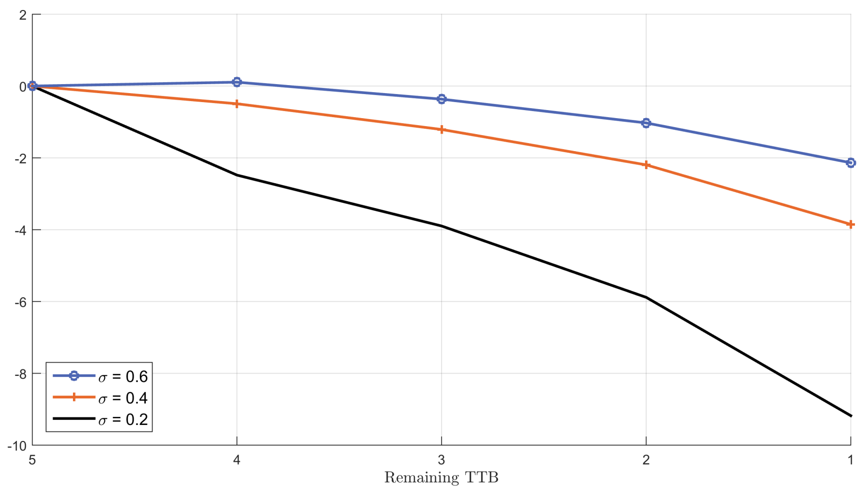

Note: The plotted measure is $(1/\sigma)[\log P^{*}(K;\sigma) - \log P^{*}(1;\sigma)]$, and the x-axis is $5 \times K$. For this plot, we assume physical TTB to be 5 months. The value of σ is annualized.

high, the same change in the uncertainty-invariant price innovation leads to deferrals of TTB investment, which is not the case under low uncertainty.

To summarize, we stress three results from the model solution: with higher uncertainty, (i) the threshold price of investment for housing start increases, (ii) the threshold price of investment for incomplete construction increases, and (iii) the relative threshold price of investment for incomplete construction to housing start increases. That is, even controlling for housing start decisions, TTB investment decisions are disproportionately affected as the level of uncertainty grows. In particular, the option value becomes relatively more sensitive to the degree of uncertainty as the project nears completion. At the last stage of construction ($K = \kappa/\bar{K}$), our model becomes a standard 1-stage irreversible investment model, for which the uncertainty effects are extensively studied in Bloom (2009). We now move on to the simulation to quantify the effects of price, uncertainty, and bottlenecks during the housing boom-bust period.

incomplete projects, they are based on their own previous price path. Intuitively, once foundations are set, the builders are locked into that location/structure, which amounts to giving up an option to build in a new location or to build a new structure.

5.3 Model simulation

In this section, we simulate and estimate the model to study the dynamics of TTB investment with regard to house price dynamics.

5.3.1 Simulation details

The economy consists of 20,000 builders. In each period, a builder might have an incomplete building under construction or not. Builders without an ongoing project decide whether to start a new building. Builders who have an incomplete project, on the other hand, make intensive investment decisions. The completion value of a house for the builder is the market price.[26] After a project is completed, the builder is available for a new construction. She gets a fresh price draw based on the aggregate price factor and decides whether to initiate new construction.

The measured TTB time series in the simulation is the average duration of projects completed in each period, consistent with the reported census data statistic. One point to make is that, if the duration distribution is extremely skewed to the right, then the mean could be driven by a few outliers. In the simulation, some building projects could experience low price realizations for a long time. When these projects eventually pick up, they would significantly drive up the average and potentially overestimate the dynamics of TTB. For data guidance, we compute extreme values for the raw TTB of completed projects between 2000 and 2013. The maximum observed TTB is 89 months, the 99.9 quantile is 46 months, and the 99.0 quantile is 27 months. We use this information to discipline our simulation. The maximum duration of a project is allowed to be 90 months, and the project is abandoned if it is incomplete by then.

5.3.2 Estimation details

The goal in the simulation exercise is to understand the model-based drivers of TTB during the housing boom-bust and recovery period. Therefore, matching the distributional properties of TTB is useful in understanding the main drivers of TTB dynamics that the mean TTB itself does not reveal. In particular, the model uses three key governing parameters; the bottleneck effect (p_c), the price effect (P^M), and

[26]That is, while under construction, the builder only optimizes on the completion value of the house under construction and not on future investment opportunities.

the uncertainty effect (σ). We use our micro-data to decompose the three structural channels that the model contains and to understand whether uncertainty effects played a distinctive role during the housing cycle.

For many houses, the completion year is different from the start year. That is, the TTB of a given year typically reflects economic conditions in both the current and the previous year. Therefore, we estimate the parameters for each odd year in our data (2003, 2005, 2007, 2009, 2011, 2013). The estimated parameters for each odd year should be interpreted as the two-year average housing market conditions faced by the developers.

Because the median house price for *new* construction is available in the data, we also use this information in our estimation. As a baseline, we construct the real new house price data by deflating the "median sales price of new homes sold in the US" (Census) by the consumer price index (BLS). In the simulation, we set the two-year average new house price as the baseline price that homebuilders face.

Simulated Method of Moments. For each odd year t between 2003 and 2013, the three values $\{p_{ct}, P_t^M, \sigma_t\}$ should be specified. We estimate the initial value P_{2003}^M. The remaining prices $\{P_t^M | t > 2003\}$ are set by the observed growth rate of new house prices data described above. The bottleneck parameters $\{p_{ct}\}$, and the uncertainty parameters $\{\sigma_t\}$ do not have an apparent empirical counterpart and are thus estimated. We use the simulated method of moments (SMM) to determine each parameter. The moments for economic TTB that we use for estimation in each year are (1) mean TTB, (2) $Pr(TTB \leq 6)$, (3) $Pr(6 < TTB \leq 9)$, and (4) $Pr(9 < TTB \leq 12)$, where the TTB data are our economic TTB. These moments well identify the various channels of the model. For instance, given a certain mean TTB, bottleneck effects are likely to govern the exogenous movements of TTB across the overall distribution, whereas uncertainty effects could govern the tail distribution.

As noted, we estimate the initial price level P_{2003}^M. Intuitively, when this aggregate price level is low, projects could start with a good idiosyncratic price draw because the value of waiting for a better price draw is small, but builders might not continue to invest afterward if the idiosyncratic price draw does not increase enough to compensate for the low aggregate price level. Therefore, the completion rate of initiated projects is helpful in identifying this channel. However, the completion rate cannot be computed for each year from the census data. Fortunately, the census does include

Figure 9: TTB distribution (SMM)

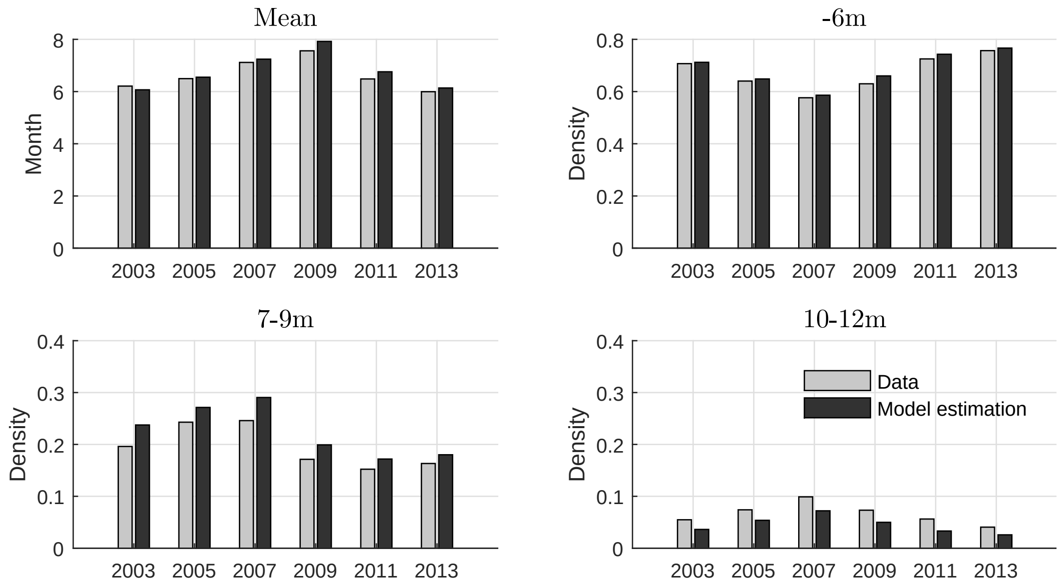

Note: The 4 moments used in the estimation of parameters (mean TTB, fraction of TTB in 4–6 months, 7–9 months, and 10–12 months) for each odd year are plotted in the panels above. The grey bars are the empirical moments and the black bars are the moments estimated using SMM.

the abandon rate of single-family houses in 2002, which was 0.5%. Therefore, in 2003, we also impose a penalty function on the distance between this data abandon rate of 0.5% and the simulated abandon rate.

5.3.3 Simulation results

Figure 9 reports both the data and the simulated moments for the obtained values. The model fit of the estimated values is decent across the distributional moments in each year. In particular, the estimated model is consistent with the dynamics of the TTB distribution for each of the four depicted moments. In Figure 10a, given the time series of new house price data, the estimated values of bottlenecks and uncertainty across time are plotted. The estimated values for the three values in 2003 are $P^M_{2003} = 1.984$, $p_{c,2003} = 0.379$, $\sigma_{\text{annual},2003} = 0.395$.[27] The bottleneck parameter increases to

[27] An alternative and direct empirical measure to infer house price uncertainty in the model would be the standard deviation of the homebuilders' perceived house price growth rate before starting a project. Unfortunately, in the data, prices are observed only when they are sold, for houses that end up being sold. Moreover, data for the unit-level growth rate of new construction sales prices are not

Figure 10: SMM for the baseline model

(a) Model parameters

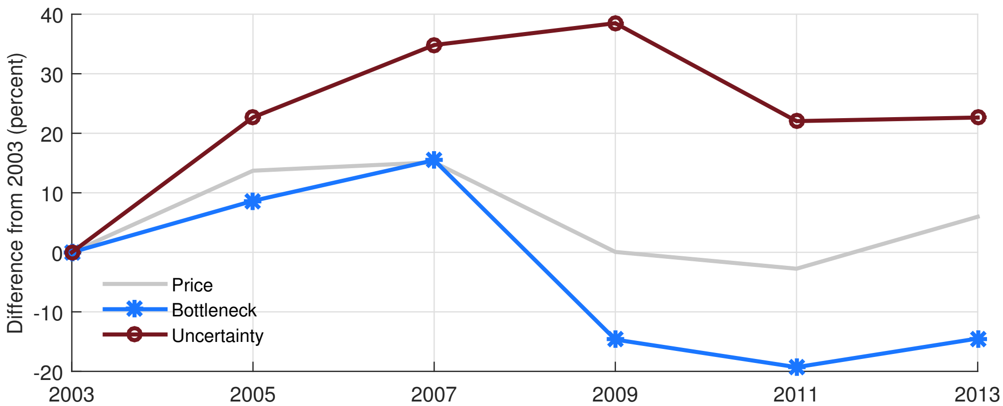

(b) Historical decomposition of TTB

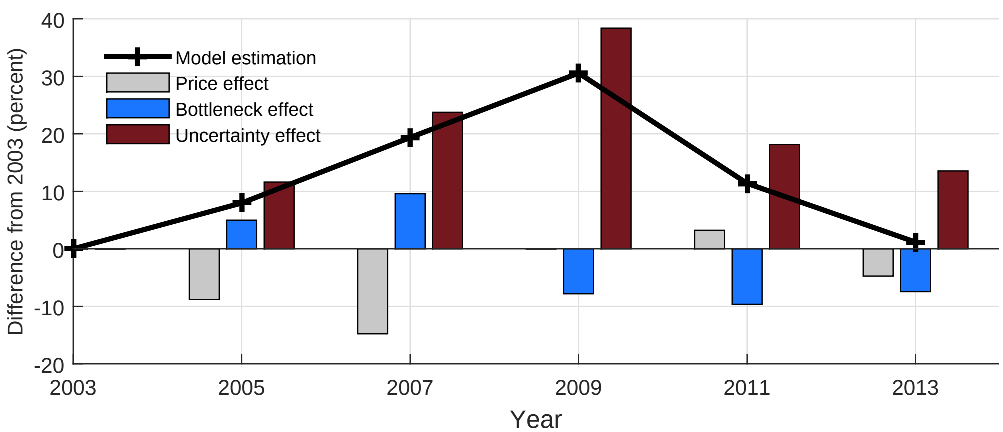

Note: Panel (a) plots the estimated parameters (bottleneck and uncertainty) relative to the first year in percentage terms. The price levels relative to the first year are set based on the two-year average median sales price of new homes sold (census) deflated by the consumer price index (BLS).

15 percent until 2007 and plummets during the housing bust and afterward. On the other hand, the uncertainty parameter remains high in 2009, about 38 percent higher

readily available.

than its 2003 value.[28]

To understand the estimated channels of TTB dynamics, we compute the historical decomposition of TTB. Because the solution of the model is nonlinear, there are two different ways to gauge the effect of each channel. For example, when computing the price effect, we could fix p_c and σ at their respective 2003 values and plot the mean TTB only when prices move. On the other hand, we could fix P^M at its 2003 value and plot the implied mean TTB as the other values change across time. Subtracting this counterfactual value from the model-estimated mean TTB would provide the net price effects. Figure 10b reports the average of these two methods of computing the historical decomposition.

Several things are worth noting. First, the increases in TTB in 2005 and 2007 are driven by a combination of bottleneck and uncertainty effects.[29] However, the large rise in prices during these periods counter this increase in TTB. Second, the further increase in TTB in 2009 is not driven by bottleneck effects. In fact, the model estimates suggest that bottlenecks are resolved during this period and with small price effects; uncertainty effects dominate during this period. The negative bottleneck effects during this period are consistent with the pattern of ex-bottleneck TTB depicted in Figure 6. Third, after 2009, the fall in uncertainty drives the fall in mean TTB.

5.4 Other aspects

In this section, we address several issues relevant to applying our data to the residential TTB model.

5.4.1 Construction cost and alternative simulation

One driver of the increase in the new house price during the recent housing boom period might be a rise in building costs. Moreover, if building costs did not fall as quickly as house prices during the bust period, this difference could better explain

[28] For example, although our paper does not measure uncertainty at the macroeconomic level, the estimated increase during the housing boom and bust period is less than those measured in Jurado et al. (2015) and Shin and Zhong (2016).

[29] The estimated increase in uncertainty during the boom is consistent with the increase in the overall house price dispersion documented in Van Nieuwerburgh and Weill (2010).

Figure 11: SMM for the alternative model (price deflated by building cost index)

(a) Model parameters

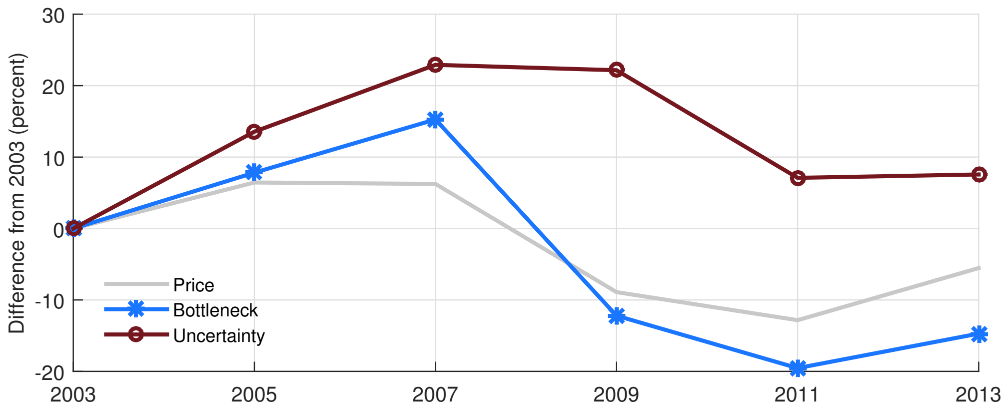

(b) Historical decomposition of TTB

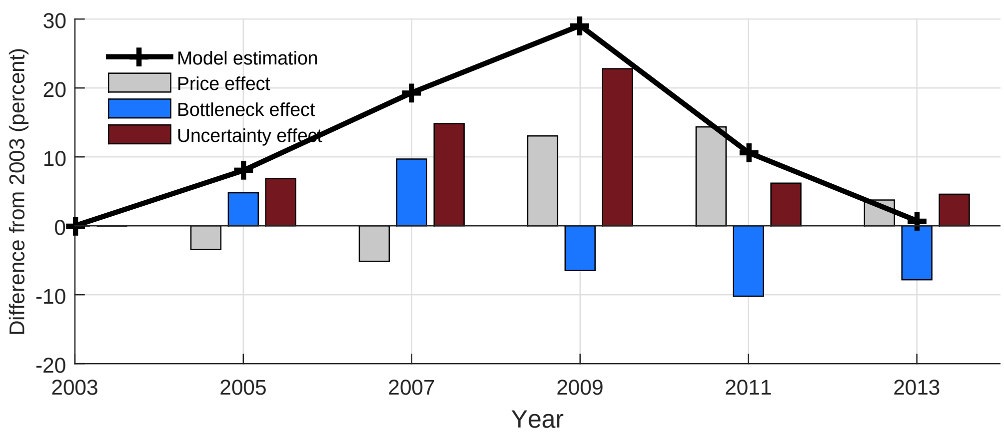

Note: Panel (a) plots the estimated parameters (bottleneck and uncertainty) relative to the first year in percentage terms. The price levels relative to the first year are set based on the two-year average median sales price of new homes sold (census) deflated by both the consumer price index (BLS) and the building cost index (Engineering News Record).

TTB dynamics during the bust period than the rise in uncertainty. To gauge the first-order importance of the uncertainty effect while taking building costs into account, we use the building cost index constructed by the "Engineering News Record."[30] That is, the model is re-estimated using new house price data deflated by the building cost index.

Figure 11a plots the new estimation with the newly imposed house price index. For the estimated values, the bottleneck effects remain similar in size until 2007, but with the smaller increase in prices, the uncertainty effect is also estimated to be smaller. Also, the bottleneck value in 2009 becomes smaller than its 2003 value, whereas uncertainty remains more than 20 percent above its 2003 value.

Figure 11b reports the historical decomposition of TTB with the newly imposed price index. With the alternative price index, price effects play a significant role in the increase in TTB in 2009. However, with the fall in the bottleneck effect, uncertainty remains the most important channel driving TTB dynamics. In 2011, the fall in bottlenecks and uncertainty drove down the increase in TTB even though price effects would be expected to drive TTB up significantly. Overall, movements in TTB since 2009 are consistent with movements in the estimated uncertainty effects.

5.4.2 New and existing house price

In the baseline simulation, the observed house price movements are insufficient to account for the TTB dynamics during the recent period. In fact, bottleneck effects are important for the housing boom period, and uncertainty effects are the dominant force for the housing bust period. One thing to note is that the movement of the new price data in Figure 10a is smaller than the Case-Shiller or FHFA price data in Figure 7.[31]

One reason for the gap between these two prices is the difference in housing supply elasticity across locations (Saiz, 2010). New construction sales price is weighted towards locations with large volumes of new construction, which have high housing supply elasticity. On the other hand, house price indices are based on repeated sales for existing houses, which are likely driven by locations with low housing supply elas-

[30] As described on their webpage, this index is constructed by computing the cost of "*68.38 hours of skilled labor at the 20-city average of bricklayers, carpenters and structural ironworkers rates, plus 25 cwt of standard structural steel shapes at the mill price prior to 1996 and the fabricated 20-city price from 1996, plus 1.128 tons of portland cement at the 20-city price, plus 1,088 board-ft of 2 × 4 lumber at the 20-city price.*"

[31] One caveat is that the price data we use in Figure 10a are the two year average (e.g. in 2003, the price data are the average of 2002 and 2003), but the movements are small even when computing at the same frequency.

ticity. The negative relation between house price elasticity and house prices could account for the muted movements in the new construction sales price.

Because our model is about new construction, we use new price data in the estimation. Using the existing house price index would overestimate the price effect.

5.4.3 Planning of construction

Our model does not include the planning process for construction, activities such as purchasing a lot and applying for building permits. The overall construction process takes longer than the construction activity itself. Our data do include permit issuance months for housing units that require building permits. In the raw data, the time from permit to the completion of a project was about 7 months before 2003 and 10 months in 2009. In the Appendix, we construct both economic and ex-bottleneck measures for the average length of time from permit to completion, following the same method described in section 3. Our TTB measures (start to complete) turn out to be conservative in illustrating the delay in construction between 2003 and 2009. That is, the time from permit to completion increased even more in percentage terms in 2009. This suggests that the delay in construction activity could also extend to the early planning stages of construction.

5.4.4 Demand and inventory of completed houses

Our model abstracts from inventory considerations of completed houses, since developers prefer to sell a completed unit immediately. In the data, sellers might not be able to find demand for their house immediately, and the sales frequency of completed houses fell in 2009.[32] In the Appendix, we extend the baseline model and allow for forced inventories of completed houses by assuming that not all houses meets demand immediately. Fitting the demand probabilities to the observed sales frequency of completed houses, the estimated model finds that the demand channel is limited in accounting for the increase in TTB in 2009 and also the decrease in TTB in 2011. The intuition is that when demand is low, forward-looking developers become more selective in starting a house, rather than extending TTB.

We note that our baseline model also has a demand channel: low prices indicate low demand. High uncertainty widens the gap between high and low price realizations,

[32]The statistics are provided in the Appendix.

and a high priced house faces a lower bound in TTB, whereas a low priced house will increase TTB to wait for a better future. Therefore, the uncertainty channel in our model incorporates a basic demand effect between prices and quantity. In the data, developers might also be waiting for demand above their desired price level rather than demand *per se*.

6 Housing supply implications

The results so far imply that the real-options mechanism potentially accounts for a significant portion of TTB dynamics in the recent housing cycle. In this section, we use our simulation to consider several housing supply implications that follow from the channels studied so far. Although understanding the dynamics of TTB is interesting in itself, their implications for housing supply dynamics are even more important. We offer two insights using our model. First, we consider the model's implications for residential investment under our baseline simulation. Second, we study how variable construction intensity driven by the observed house price dynamics affect housing start dynamics, by comparing our results with a simulation under fixed construction intensity.

6.1 Investment with variable construction intensity

Housing starts and residential investment are key data of interest to macroeconomists and policymakers because of their high volatility and stable lead-lag structure with GDP (Davis and Heathcote, 2005). In this section, we ask how our intensive margin affects the dynamics and lead-lag structure of residential investment.

To understand these short-run implications of our model, we study its transition dynamics using perfect foresight equilibrium with a regime change. Starting from the 2003 values, Figure 12 plots the transition of several variables when house prices fall by 5 percent. We find that housing starts fall by 5 percent on impact, and construction intensity for incomplete houses falls by less than that on impact. Therefore, in the short run, housing starts have a bigger effect on total investment relative to construction intensity with this price fall. Looking into the implications for TTB, we find a gradual increase in TTB for construction that started before the regime shift, and the TTB of completed houses gradually increases up to 4 percent higher than

Figure 12: Transition dynamics with a price fall

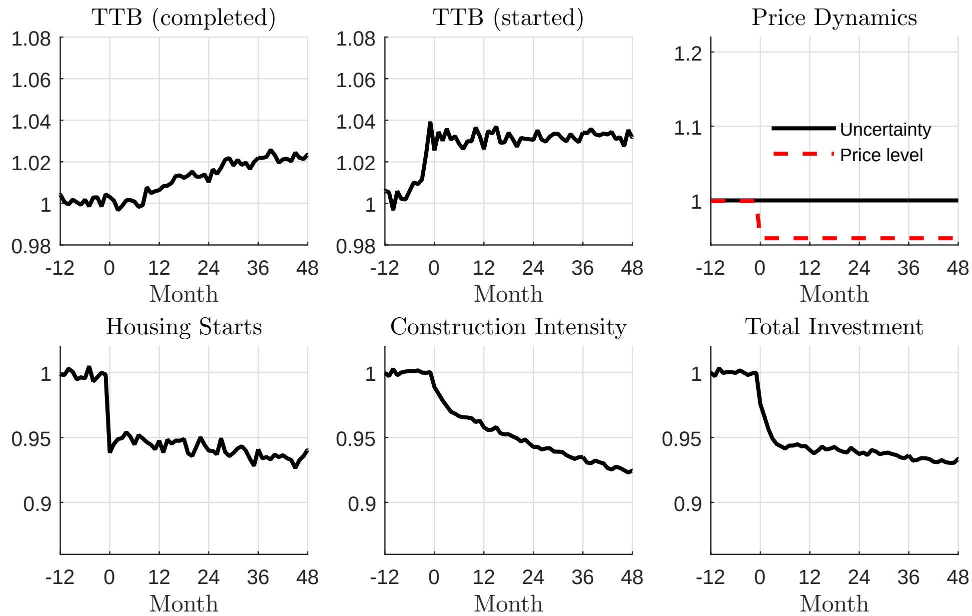

the previous level.

Figure 13 plots the same transition with a concurrent 20 percent increase in uncertainty. Compared to a pure price fall, we find that construction intensity falls even more than housing starts in less than a year. The fall in total investment is thus more affected by intensive investment than by extensive investment. Moreover, with an increase in uncertainty, TTB for houses started before the regime shift overshoots the new steady-state TTB. That is, builders who started their projects before the price and uncertainty change now find themselves in a different world and engage in wait-and-see behavior beyond that of builders starting new projects.

In Figure 14, we plot the within-year total housing starts and residential investment based on the price fall and uncertainty increase in Figure 13. To gauge how TTB dynamics (i.e., intensive investment) affect residential investment, we also plot a counterfactual time series for residential investment in which we shut down the TTB channel by assuming that there are no endogenous deferrals (making economic TTB and physical TTB the same).[33]

We conclude two points from this figure. First, for the initial periods after our experiment, construction intensity on incomplete houses drives the dynamics of residential investment rather than housing starts. For the first month, the counterfactual

[33]Constant stochastic bottleneck effects remain in action for both cases.

Figure 13: Transition dynamics with a price fall and uncertainty rise

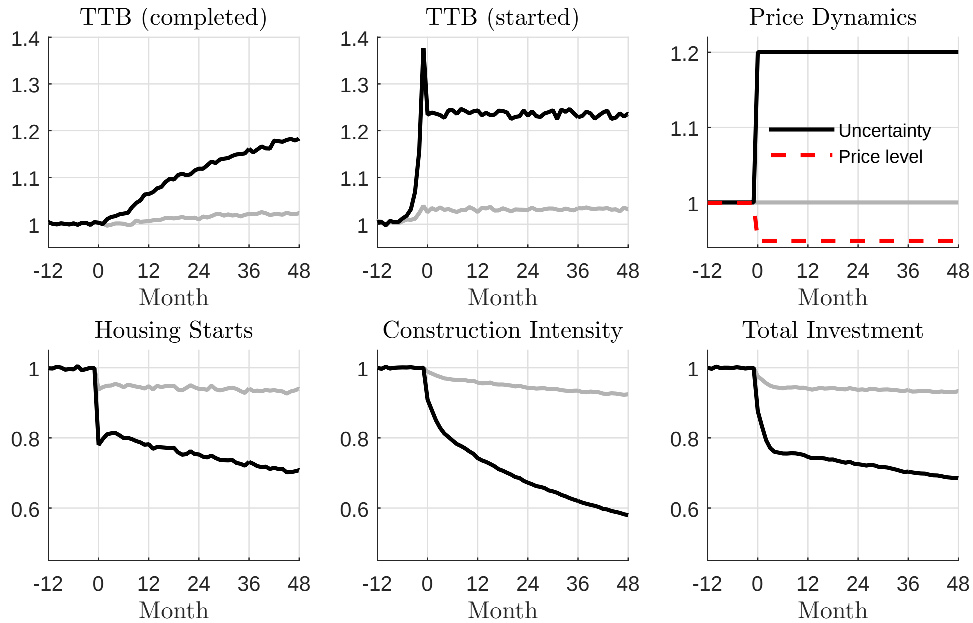

Note: Grey line depicts data from Figure 12, the pure price fall effect.

Figure 14: Housing starts and residential investment

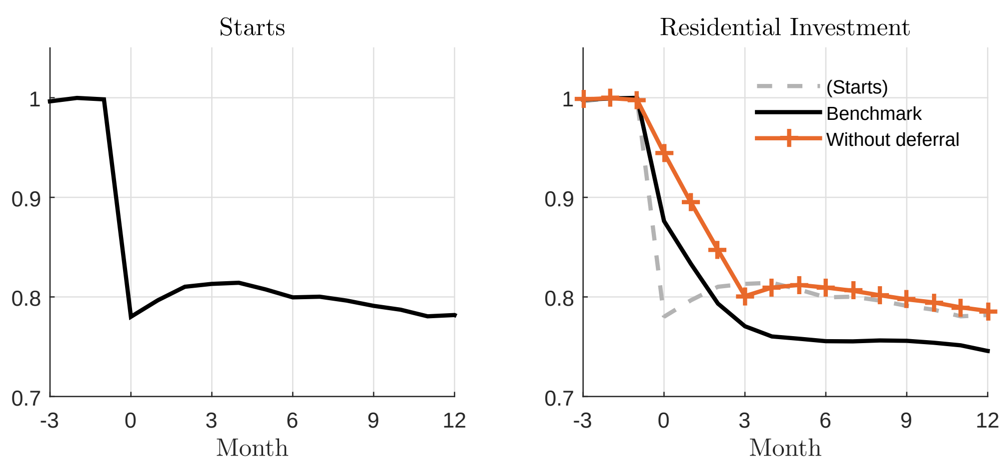

residential investment, which contains information only on housing starts, falls by only half of the actual simulated residential investment. The initial movements in residential investment are thus dominated by the intensive margin rather than the extensive margin. Second, whereas the counterfactual residential investment series lags housing starts, the actual residential investment responds immediately with deferrals, and the lag pattern becomes muted.

Figure 15: Housing dynamics with variable and fixed construction intensity

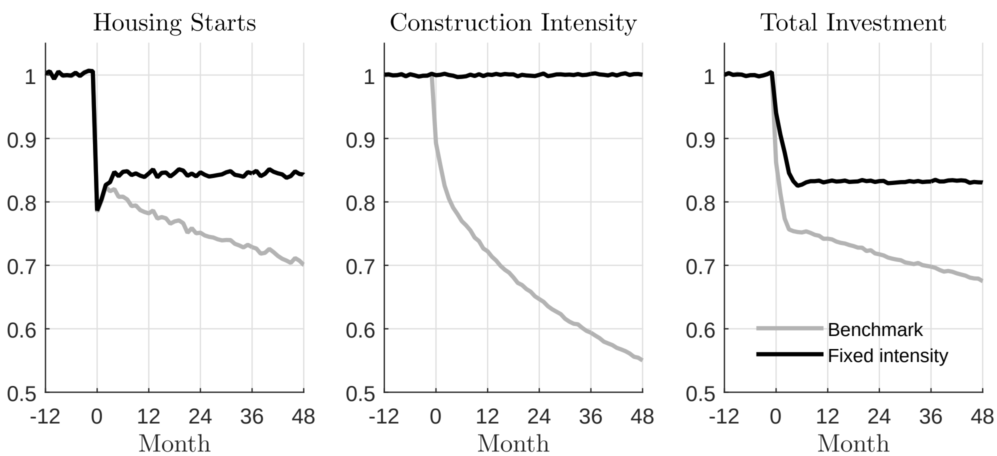

Therefore, neglecting the intensive margin of residential investment and forecasting residential investment based only on housing starts data could lead to incorrect short-term estimates, especially when the movement of house price dynamics is large, as in the recent recession.

6.2 Housing starts with variable construction intensity

In this section, we use our model to analyze how variable construction intensity affects housing starts. In particular, we compare our result with that under the commonly used fixed TTB assumption. Figure 15 compares the dynamics of housing starts for variable and fixed construction intensity. Under fixed construction intensity, the initial fall in housing starts is large and immediate, quickly approaching the new steady state. On the other hand, under variable construction intensity, the fall in housing starts is persistent and large, only slowly converging toward a lower steady state.

In Figure 16, the dynamics of housing starts in Figure 15 are decomposed into (i) the start rate per builder and (ii) the number of available builders. Looking at the first column, the effect on the start rate is similar in both cases. In the second column, the number of available builders is plotted under two scenarios. For both types of builders, available builders increase initially for a longer period because completed builders decide not to start new projects. However, the two types of builders diverge in the later period because, with variable construction intensity, pre-existing projects

Figure 16: Housing starts

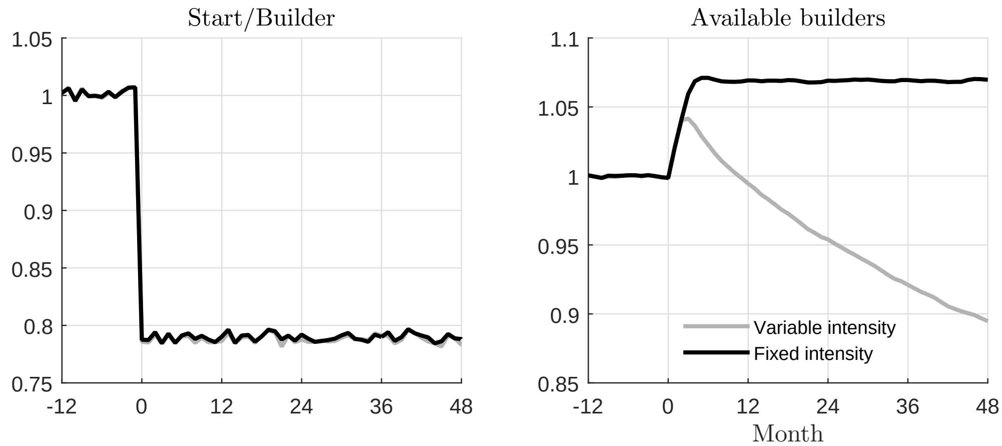

are all delayed and inventory overhang leads to fewer new housing projects. On the other hand, available builders do not fall below the steady state in the second column, because no delays occur with fixed construction intensity.

With fixed intensity, pre-existing projects must be completed on schedule even with the arrival of new house-price information. Therefore, the flow of builders completing projects is relatively stable. Those builders are quickly available to start a new house by drawing a new price. On the other hand, with variable intensity, pre-existing projects can now be deferred in response to the new pricing information. In the simulation, the completion rate falls, and the number of builders available for new projects remains low for a long time. This overhang of incomplete houses generates a deeper housing slump because the number of available builders shrinks.

6.3 Housing supply in the boom-bust period

Given the implications of our model for housing supply, we finally ask how our model channels account for the housing-start and residential-investment dynamics in the recent housing cycle. Using the estimated values for price, uncertainty, and bottleneck effects that match the TTB distribution in the data, Figure 17 plots the implied housing supply variables in each period: the implied housing starts and residential investment based on the baseline estimation, along with housing starts and residential investment data for single-unit houses.[34] Based on 2003 data, the model

[34]For residential investment, we use NIPA data: real private fixed investment - residential structures - permanent site - single family.

Figure 17: Housing supply data and model

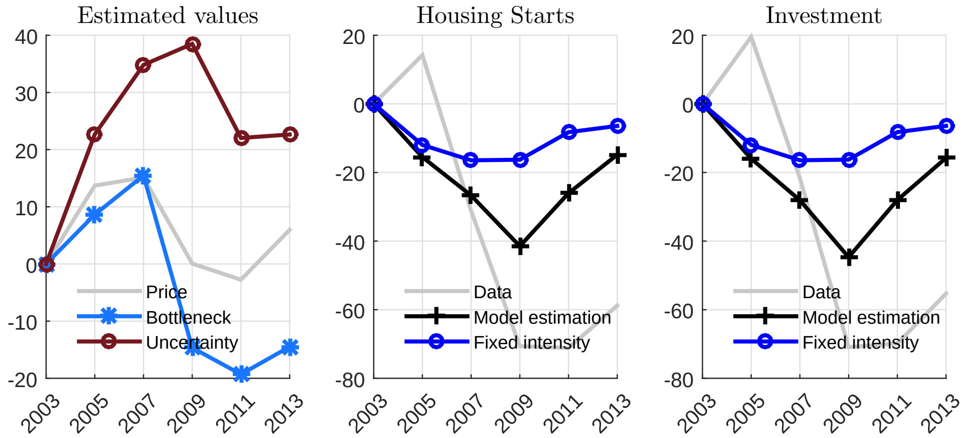

claims an overall 40 percent fall in housing starts and residential investment in 2009. On the other hand, beginning with the same estimated values but assuming a fixed intensity of construction, housing starts and residential investment fall by only 18 percent in 2009. In the data, housing starts and residential investment fell by around 70 percent. The intensive investment deferral channel generates a sizeable slump in housing starts and residential investment, closer to the data than the fixed intensity model. The variable intensity model accounts for more than half of the fall in housing starts in the data.

However, our model is also limited in two dimensions: (i) in 2009, housing starts fell by 70 percent instead of 40 percent as our model indicates, and (ii) in 2005, housing starts increased by 10 percent instead of a 17 percent fall as our model indicates. Similar arguments apply to residential investment.

In short, although our model is estimated to match the intensive investment distribution of residential construction, those channels have limited influence in generating the housing-start dynamics observed in the data. Obviously, the model is missing many variables that could also have affected housing starts in the boom-bust period. For example, the credit boom and bust would have affected the *initial* availability of construction loans, which could work disproportionately at the extensive margin, and the permit process might also changed during this period. Moreover, a micro-founded model of bottlenecks during the housing boom period would imply that bottlenecks increased because builders engaged in multiple projects at the same time, which our stochastic bottleneck assumption does not allow. Lastly, a general equilibrium model would imply that interest rates should also interact with prices and uncertainty. Interest rate movement could thus also be a potentially important channel affecting housing starts and investment decisions.[35]

7 Conclusion

We document new findings about the distribution of residential TTB across the US. The fall in economic activity in the recent housing bust is not limited to housing starts, but expands to TTB investment. Contrary to the notion that already started projects are costly to stop, we find that a significant portion of them were deferred during the housing bust. We study a model in which TTB investment responds to prices, uncertainty, and bottlenecks and simulate that model using the observed house price dynamics. The real-options mechanism can account for the drop in TTB investment during the housing bust.

Before concluding, it is important to note that, in this paper, we abstract from the searching and matching channels in housing demand, to focus on the supply-side decisions. We also leave aside the financial frictions channel in TTB investment projects. The construction sector is indeed a leveraged industry, and the recent housing boom-bust cycle is closely related to the availability of credit. Builders and lenders with different financing conditions and contracts would have behaved differently during the housing bust, and the overall financial constraints might have exacerbated the aggregate housing market collapse.

Although we find that the real-options mechanism can account for investment activity on incomplete projects, its potential interactions with housing demand or developers' financial frictions are interesting extensions that should be pursued further. In particular, as discussed in the previous section, there is room for improvement in accounting for the housing start dynamics during the recent housing cycle.

[35] Bloom et al. (2016) extend Bloom (2009) and study the general equilibrium channels of uncertainty shocks. The general equilibrium model can generate a sizable drop in GDP with a rise in uncertainty.

References

Arezki, R., V. R. Ramey, and L. Sheng (2017). News Shocks in Open Economies: Evidence from Giant Oil Discoveries. *Quarterly Journal of Economics 132*(1), 103–155.

Bachmann, R., S. Elstner, and E. R. Sims (2013). Uncertainty and Economic Activity: Evidence from Business Survey Data. *American Economic Journal: Macroeconomics 5*(2), 217–249.

Bloom, N. (2009). The Impact of Uncertainty Shocks. *Econometrica 77*(3), 623–685.

Bloom, N. (2014). Fluctuations in Uncertainty. *Journal of Economic Perspectives 28*(2), 153–176.

Bloom, N., S. Bond, and J. Van Reenen (2007). Uncertainty and Investment Dynamics. *Review of Economic Studies 74*, 391–415.

Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten, and S. J. Terry (2016). Really Uncertain Business Cycles. *mimeo*.

Bulan, L., C. Mayer, and C. T. Somerville (2009). Irreversible Investment, Real Options, and Competition: Evidence from Real Estate Development. *Journal of Urban Economics 65*, 237–251.

Campbell, J. R. (1998). Entry, Exit, Embodied Technology, and Business Cycles. *Review of Economic Dynamics 1*(2), 371–408.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (2005). Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy. *Journal of Political Economy 113*(1), 1–45.

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

Dixit, A. K. and R. S. Pindyck (1994). *Investment Under Uncertainty.* Princeton University Press.

Edge, R. M. (2007). Time-to-Build, Time-to-Plan, Habit-Persistence, and the Liquidity Effect. *Journal of Monetary Economics 54*(6), 1644–1669.

Gilchrist, S., J. W. Sim, and E. Zakrajšek (2014). Uncertainty, Financial Frictions, and Investment Dynamics. *NBER Working Paper 20038*.

Glaeser, E. L., J. Gyourko, and A. Saiz (2008). Housing Supply and Housing Bubbles. *Journal of Urban Economics 64*(2), 198–217.

Glaeser, E. L., J. Gyourko, and R. E. Saks (2005). Why Have Housing Prices Gone Up? *American Economic Review 95*(2), 329–333.

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

Haughwout, A., R. W. Peach, J. Sporn, and J. Tracy (2013). The Supply Side of the Housing Boom and Bust of the 2000s. In E. L. Glaeser and T. Sinai (Eds.), *Housing and the Financial Crisis*, Chapter 2, pp. 69–104. University of Chicago Press.

Iacoviello, M. (2005). Housing Prices, Borrowing Constraints and Monetary Policy in the Business Cycle. *American Economic Review 95*(3), 739–764.

Jovanovic, B. and P. L. Rousseau (2014). Extensive and Intensive Investment over the Business Cycle. *Journal of Political Economy 122*(4), 863–908.

Jurado, K., S. C. Ludvigson, and S. Ng (2015). Measuring Uncertainty. *American Economic Review 105*(3), 1177–1216.

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

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

Kydland, F. E., P. Rupert, and R. Sustek (2016). Housing Dynamics over the Business Cycle. *International Economic Review 57*(4), 1149–1177.

Leahy, J. V. and T. M. Whited (1996). The Effect of Uncertainty on Investment: Some Stylized Facts. *Journal of Money, Credit, and Banking 28*, 64–83.

Leamer, E. E. (2007). Housing IS the Business Cycle. *NBER Working Paper 13428*.

Lucca, D. O. (2007). Resuscitating Time-to-Build. *mimeo*.

Majd, S. and R. S. Pindyck (1987). Time to Build, Option Value, and Investment Decisions. *Journal of Financial Economics 18*, 7–27.

Moel, A. and P. Tufano (2002). When are Real Options Exercised? An Empirical Study of Mine Closings. *Review of Financial Studies 15*(1), 35–64.

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

Shin, M. and M. Zhong (2016). A New Approach to Identifying the Real Effects of Uncertainty Shocks. *Board of Governors of the Federal Reserve System Working Paper Series 040*.

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

Uribe, M. (1997). Exchange-rate-based Inflation Stabilization: The Initial Real Effects of Credible Plans. *Journal of Monetary Economics 39*, 197–221.

Van Nieuwerburgh, S. and P.-O. Weill (2010). Why Has House Price Dispersion Gone Up? *Review of Economic Studies 77*(4), 1567–1606.