Residential construction lags and the real-options channel of housing supply*

Hyunseung Oh[†] and Chamna Yoon[‡]

April 27, 2016

Abstract

Housing supply decisions consist of both an extensive margin (new housing starts) and an intensive margin (construction intensity of incomplete houses). While it is well known that housing starts have declined dramatically during the 2006–2009 housing bust, the intensive margin of residential investment has not been studied in the literature. In this paper, we document that construction intensity of incomplete houses has also fallen significantly during the bust. Using the Census micro data for construction lags of single-family houses across the US, we show that average construction lags for completed houses increased during the bust, and that this increase comes from long deferrals of several houses under construction, especially those that were unsold at the early stage of construction. Motivated by these new facts, we study a time-to-build model of residential construction where investment in each stage is irreversible. The model predicts that as the level of uncertainty increases, the "wait-and-see" channel becomes relatively more important for the intensive margin than for housing starts. Calibrated and estimated to match the house price dynamics and the construction lag distribution during the recent recession, the model suggests that the real-options mechanism played an important role in the dynamics of residential investment during the recent bust. Several housing supply implications based on the model follow.

Keywords: Housing; Real options; Investment; Time to build; Adjustment costs.

JEL Classification Numbers: E13, E22, R31.

*We thank Rick Bond, Jeff Campbell, Mario Crucini, Thorsten Drautzburg, Matteo Iacoviello (discussant), James Kahn, Oleksiy Kryvtsov, Ricardo Reis, Peter Rousseau, Stephanie Schmitt-Grohe, Holger Sieg, Jon Steinsson, Gianluca Violante, Federal Reserve Bank of Philadelphia, and participants at the Midwest Macroeconomics Meetings (Fall 2015) for precious comments. All errors are ours.

[†]Vanderbilt University. E-mail: hyunseung.oh@vanderbilt.edu.

[‡]Baruch College, City University of New York. E-mail: chamna.yoon@baruch.cuny.edu.

1 Introduction

Since the Great Recession, understanding housing dynamics has become a main topic of interest for business-cycle studies.[1] While the significant cyclicality and volatility of the housing market were well-known facts even before (Davis and Heathcote, 2005), the recent housing boom-bust cycle has been nevertheless unprecedented in its size. New empirical research finds that the recent housing cycle had a large impact on the macroeconomy and contributed to the severity of the recession.[2]

Our focus of this paper is on understanding how the supply side reacted to the recent housing boom and bust. There are indeed several papers that look into each component of residential investment in the GDP data to study the various investment decisions in the housing market.[3] Like any investment variable, residential investment is also composed of both the extensive margin, new housing starts, and the intensive margins: (i) improvements on existing houses and (ii) construction of incomplete houses.[4] Most papers focus on the supply side determinants of either new housing starts – the extensive margin – or improvements on existing houses – the first intensive margin. Our paper focuses on the *other* intensive margin of residential investment that is often not discussed: the construction intensity of incomplete houses. In detail, we ask how construction intensity responded to the housing boom and bust.

Construction intensity of incomplete houses is an important variable to understand the dynamics of residential investment for the following two reasons.

First, the stock of incomplete houses is large. Building a house takes a significant amount of time. Even after building permits are issued, the average single-family house takes 6 months from start (i.e. excavation) to completion. Multi-family houses take 10 months to build. This implies that for each unit of house started in a given month, 8.3 units of houses are under construction.[5] As a result, even small variations in construction intensity should play a large role for the dynamics of residential investment.

[1] The second handbook of macroeconomics includes a chapter on housing and macroeconomics, which was not the case in the first version published in 1999.

[2] Recent papers in this literature include Leamer (2007), Mian and Sufi (2011) and Saiz (2010).

[3] Haughwout, Peach, Sporn, and Tracy (2013) give a comprehensive review of this literature and provide some data analysis.

[4] Residential investment also includes brokers' commissions and other ownership transfer costs.

[5] Data from December 1969 to December 2014.

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 have an impact on new housing starts, and hence on residential investment.

In the data, construction intensity is not directly observed. However, construction lags are observed in detail, and they have a sharp inverse relation with construction intensity. In the data sections, we study the dynamics of construction lags during the recent boom and bust. In section 2, we look into the time series of average construction lags of residential buildings across the US. In section 3, we use the Census micro data to look into the time series of the *distribution* of construction lags for single-family residential buildings across the US. We find that average construction lags increased during both the boom and bust, and that the increase in the bust period is due to long deferrals of several houses rather than an overall shift in the distribution. For "built for sale" houses, we find that long deferrals only occured for houses that were unsold before start.

Existing intuition on the time series behavior of construction lags assume either that construction lags are stable across time (Kydland and Prescott, 1982), or that construction lags would increase if bottlenecks exist for certain type of inputs (Kalouptsidi, 2014). Our finding that construction lags increased in the housing bust era, especially for several unsold houses, is not consistent with either of these assumptions. In section 4, we develop a model of time to build that nests both the standard Kydland and Prescott (1982) and the real options mechanism of Majd and Pindyck (1987), to study the first-order importance of the real options mechanism. In section 5, we apply the model to the recent housing bust. With the high level of uncertainty observed in this period, we find that the real options mechanism significantly accounts for the dynamics of construction lags. Combining with the fall in prices during the bust, we find that our model captures most of the construciton lag dynamics during the recession. Based on the dominant role that house price dynamics and our model mechanisms played in accounting for the observed TTB dynamics, we study several housing supply implications in section 6. In particular, we find that when the

[6] These forward-looking aspects, such as gestation lags or flow adjustment costs, have been widely used in the business-cycle literature to deliver 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, gestation lags are also used in Arezki et al. (2015) for an application of the effects of news shocks in an open economy.

intensive margin of investment falls, residential investment does not lag housing supply and its initial movement is dominated by the intensive margin rather than the extensive margin. Section 7 concludes.

Related literature Our paper connects to four broad strands in the literature. First, we connect to the business-cycle literature on time to build and investment dynamics, including Kydland and Prescott (1982), Campbell (1998), Lucca (2007) and Edge (2007). These papers all assume that time-to-build investment is exogenous, and look into the initial investment decisions for time-to-build projects. Our findings suggest that time-to-build investment itself has also been an important margin in the recent recession.

Second, we link to the real options literature of investment, as in Dixit and Pindyck (1994), Leahy and Whited (1996), Bloom et al. (2007), Bloom (2009), Bachmann et al. (2013) and Gilchrist et al. (2014). To the best of our knowledge, we are the first to look into the effects of uncertainty shocks for time-to-build investment decisions. Morever, we contribute to the recent interest on the quantitative importance of uncertainty shocks. We suggest evidence that uncertainty shocks have played a significant role in the intensive margin of residential investment.

Third, we connect to the housing investment literature such as Topel and Rosen (1988), Iacoviello (2005), Green et al. (2005), Glaeser and Gyourko (2005), Glaeser et al. (2005), Glaeser et al. (2008), Saiz (2010), and Kydland et al. (2012). In this literature, we provide new stylized facts on the distribution of construction lags across time.

Lastly, our paper is related to the recent interest in understanding the distinction between the extensive and intensive margin of investment dynamics in Jovanovic and Rousseau (2014).

2 Data and stylized facts: Average construction lags

In this paper we use the "Survey of Construction (SOC)" data available from the Census Bureau, which is a national sample survey (sampling rate: 1/50) on 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 completed month of houses, sales price and the month in which the house was sold if sold, square footage, number of rooms and so on. Houses authorized by building permits but not started at the end of the year, under construction at the end of the year, or for sale at the end of the year are also included.[7]

Based on this dataset, the Census reports the aggregate series for "average length of time from start to completion," for both single-family and muilt-family units. In each given year, this series reveals information on the average construction lags of completed houses. In this section, we look into this aggregate series. In the next section, we will move on to the underlying micro data to get a detailed understanding of the findings of this section.

2.1 Aggregate facts

Figure 1 depicts the average length of time from start to completion of both single-family and multi-family houses from 1984 to 2013. Notice first that construction lags are relatively constant until 2002. Average construction lags for single-family and multi-family houses during this period were 6 and 9.7 months, respectively.

Second, from 2002 to 2006, construction lags for single-family and multi-family houses increased by 1 and 2 months, respectively. One aspect that may have contributed to this increase is the shortage of construction workers relative to the number of construction projects during this housing boom period.[8] 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 unemployed ratio.[9] The low unemployment rate and the high labor market tightness support the view that there was a shortage in available workers in the construction sector during the housing boom period.[10]

Third, from 2006 to 2009, construction lags further increased for both single-family and multi-family houses, each again by 1 and 2 months, respectively.[11] However, measures of bottlenecks in the construction sector no longer support the view that bottlenecks were the main contributor to

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

[8]For example, Kalouptsidi (2014) finds that higher investment activity lengthened the production lag of the Greece ship bulking industry, which is capacity constrainted.

[9]Construction sector unemployment rate is reported in the 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.

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

[11]Average construction lag for multi-family houses peaked in 2010.

Figure 1: Construction lags for completed houses in the US

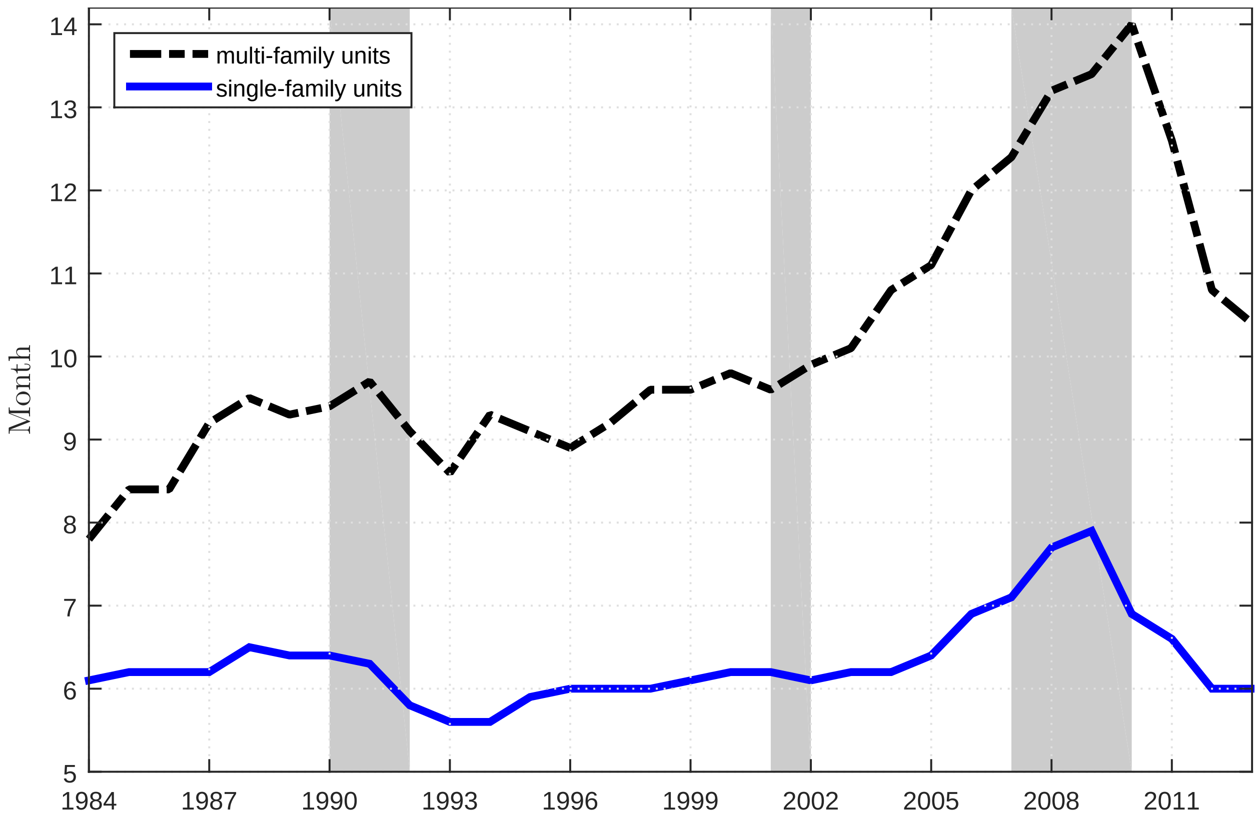

Note: Shaded areas are NBER recessions.

this further increase in construction lags. 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 bottlenecks were all resolved. This suggests that there must have been a strong mechanism that overturned the negative bottleneck effects for construction lag dynamics during the housing bust.

Finally, after 2010, construction lags recovered back to the pre-2002 level.

2.2 Implications of aggregate facts

In a standard time-to-build investment model, construction lags are assumed to be constant. This is empirically supported in the period from 1984 to 2002. Although the economy went over two NBER recessions during this period, construction lags showed little variation.

That assumption no longer holds in the recent housing boom-bust cycle, where average con-

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

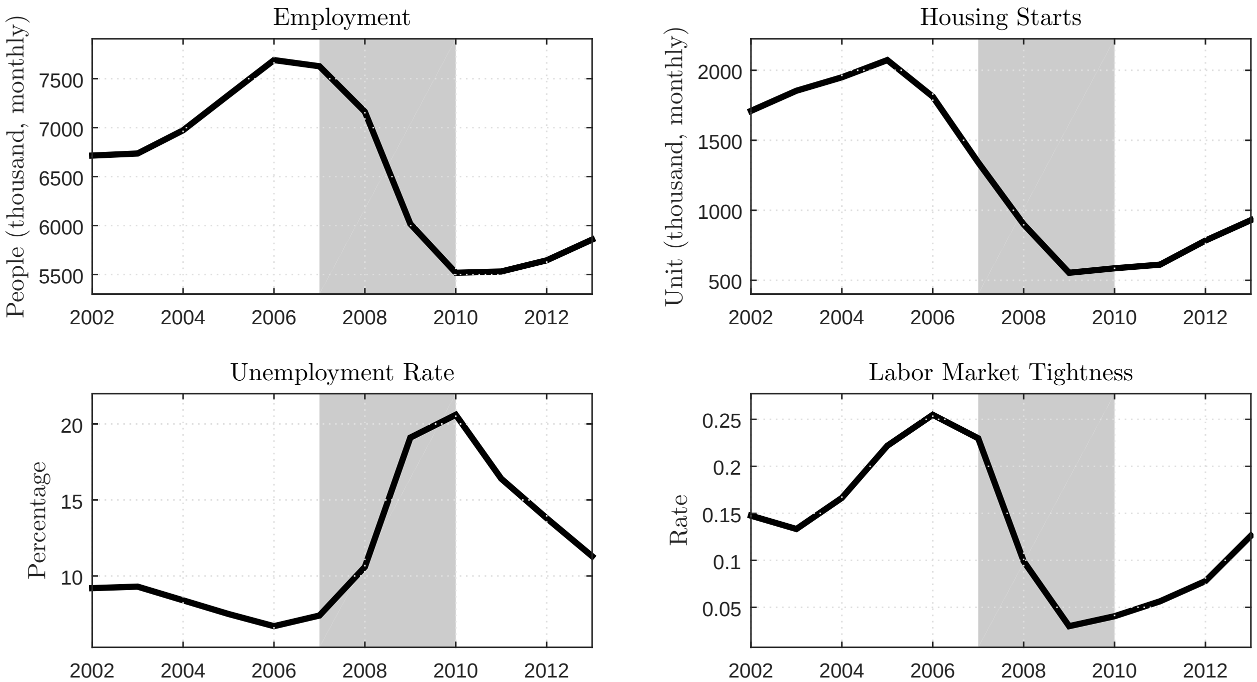

Note: Labor market tightness is the rate between construction job openings and unemployed with previous jobs from the construction sector. Shaded areas are NBER recessions.

struction lags increased by $30\% - 40\%$ relative to their 1984–2002 average. The initial increase in the boom is consistent with bottleneck mechanisms and we find that the labor market of the construction sector was tight during this period.

However, a puzzling feature is the further increase in construction lags during the bust. The data evidence shows that bottlenecks were resolved in the housing bust, and this should imply a fall in the average construction lags. Therefore, the further increase in construction lags during the bust period suggests that bottleneck theories were no longer the main action. Then why did construction lags increase during the housing bust? In the next section, we investigate this using the underlying micro data for construction lags.

3 Micro data analysis: Distribution of construction lags

Our goal in this section is to understand the dynamics of the *distribution* of construction lags during the boom-bust period. Towards this, we use the underlying micro data for the Census statistics on average construction lags. Our analysis will be based on the distribution of construction lags for *single-family houses*, of which the data is publically available since 2000.[12] We first construct a measure of "economic" construction lag by controlling for geographic and building characteristics of each completed house, and then compare its cross-sectional distribution across time.

3.1 "Economic" construction lags

All buildings are different and construction lags for each house depend on various factors. For example, a larger and more-difficult-to-build house will have lengthy construction lags. Another factor for lengthy construction periods are houses that build on severe weather conditions or stringent regulations.

Our goal in section is to focus on the dynamics of construction lags that are independent from geographic and building characteristics.[13] Since the micro data provides many of these features, we construct a measure of "economic" construction lags for each completed single-family building in the US by controlling for them. In particular, for geographical characteristics, we control for the 9 Census divisions, and whether or not 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 include building purpose (owner built, contractor built, built for sale, built for rent), building method (site built, panelized, modular), and square footage of the house.

We regress the log of time-to-build on the various control variables listed above. Table 1 reports the result from this regression, using the data from 2000 to 2013. While our main focus is not on understanding the link between the control variables and construction lags, we do find several interesting results that 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, which consists of 26.9% of the total sample. The least number of houses are built in New England (3.4%). For building purposes, built for sale houses consist 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 New England, Middle Atlantic, and Pacific divisions

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

[13]We understand that building and geographic characteristics may also be correlated with economic conditions; for example, larger houses may have been built during the boom period. In such case, our estimate should be taken as a conservative measure. However, since our focus of this paper is on the dynamics of construction lags controlling for housing start decisions, we find it best to remain silent on their possible correlations.

show lengthy construction lags compared to other divisions. In particular, construction lags are on average 27% higher in Middle Atlantic relative to that of the West South Central division. Second, owner built houses have longer TTB relative to contractor built houses, which may reflect either the efficiency of contractors or the selection into owner built houses for various housing preferences. Built for sale houses on average take the shortest time to build. Third, site built houses have longer TTB than panelized or modular houses, which may reflect the exposure to bad weather conditions for site built houses. Lastly, 2 and more story buildings take longer time than 1 story buildings, and square footage of the building also has a positive relation with construction lags.

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

3.2 The distribution of economic construction lags

Using our measure of economic construction lags, 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 subsequent bust. The annual time series of this distribution is plotted in the appendix.

The left panel of figure 3 compares the kernel density of TTB in 2003 and 2005. We observe that during this period, 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.

On the other hand, the right panel of figure 3 compares the kernel density of TTB in 2005 and 2009. We observe two facts. First, the mass of the distribution including the mode has shifted back to the left. Second, a fat tail to the right has appeared. The left shift in the mass of the distribution is consistent with bottlenecks being resolved. During the housing bust, the supply side may have had less issues with finding available construction workers to build a house. Nevertheless, the fat tail to the right indicates that several incomplete houses during the housing bust remained underconstructed for a long period. Linking this observation to the aggregate fact, the increase in TTB coming from this fat tail to the right dominated the negative bottleneck effect, and hence

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

	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 foot ($\times$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 bathrooms		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 are in parentheses. Regression (b) also with the unemployment rate.

average construction lags increased further in 2009 compared to 2005.

In figure 4, we plot the kernel density of TTB in 2005 and 2009, across 9 Census divisions. Two patterns that we emphasize – (1) mass of the distribution shifting to the left; (2) fat tail to the right – hold in most divisions. Therefore, the dynamics of residential contruction lags between 2005 and 2009 are not locally driven but rather a national phenomenon.

Figure 3: Distribution of economic construction lags (Total sample)

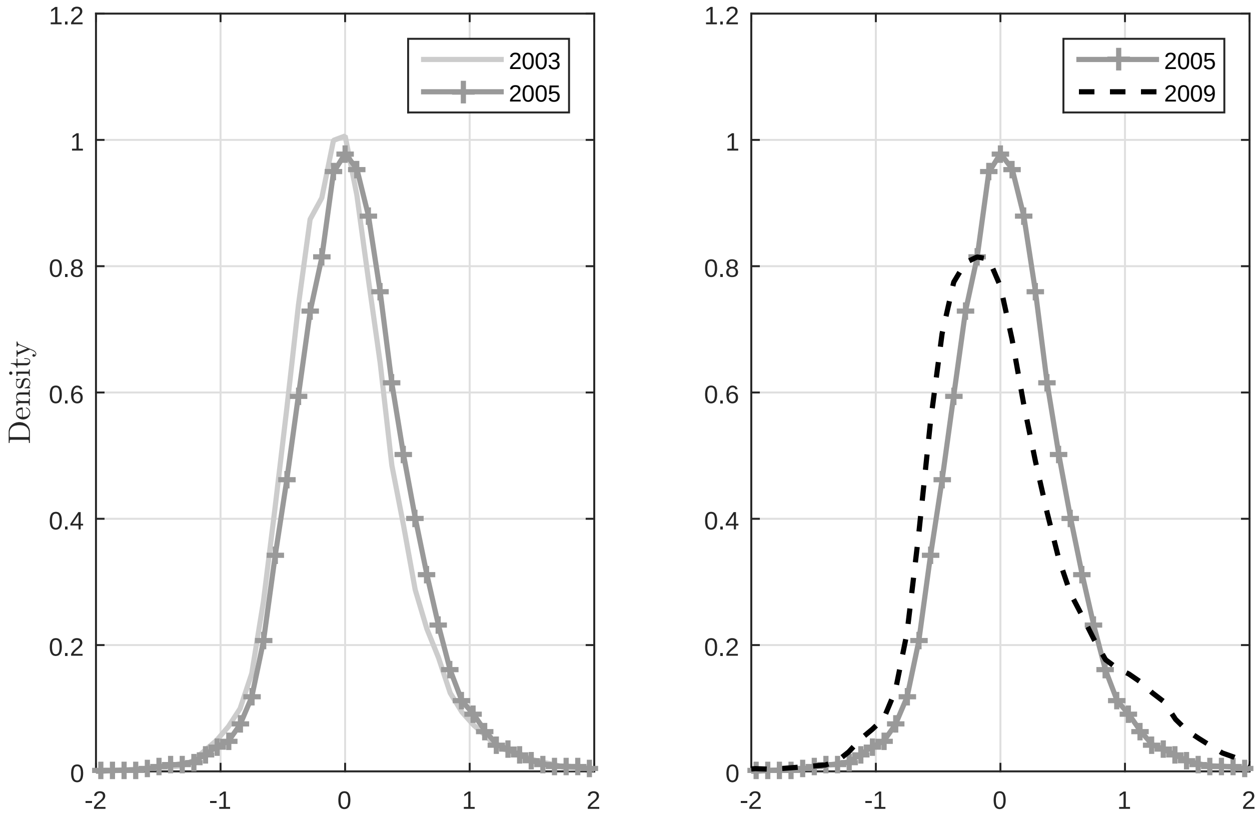

Note: Kernel density of economic TTB for total single-family houses in logs. Left compares 2003 and 2005, and right compares 2005 and 2009.

3.3 Subsample evidence: Built for sale

From the overall distribution, we observe that there were long deferrals of construction into several incomplete houses during the bust period. In this section, we narrow our focus into a subsample of the distribution, "built for sale" houses, to better understand where these long deferrals are disproportionately coming from. We look into 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 comprise the majority of single-family houses in the sample (74%) as shown in table 1, at the same time a significant fat tail to the right is observed for this category during the recent bust.

Based on the sales month information for these houses, we split this sample into two categories: houses sold before start, and unsold before start. In figure 5, we observe that for houses that were sold before start, we only observe one pattern: the distribution shifting back to the left.

Figure 4: Distribution of economic construction lags in 2005 and 2009

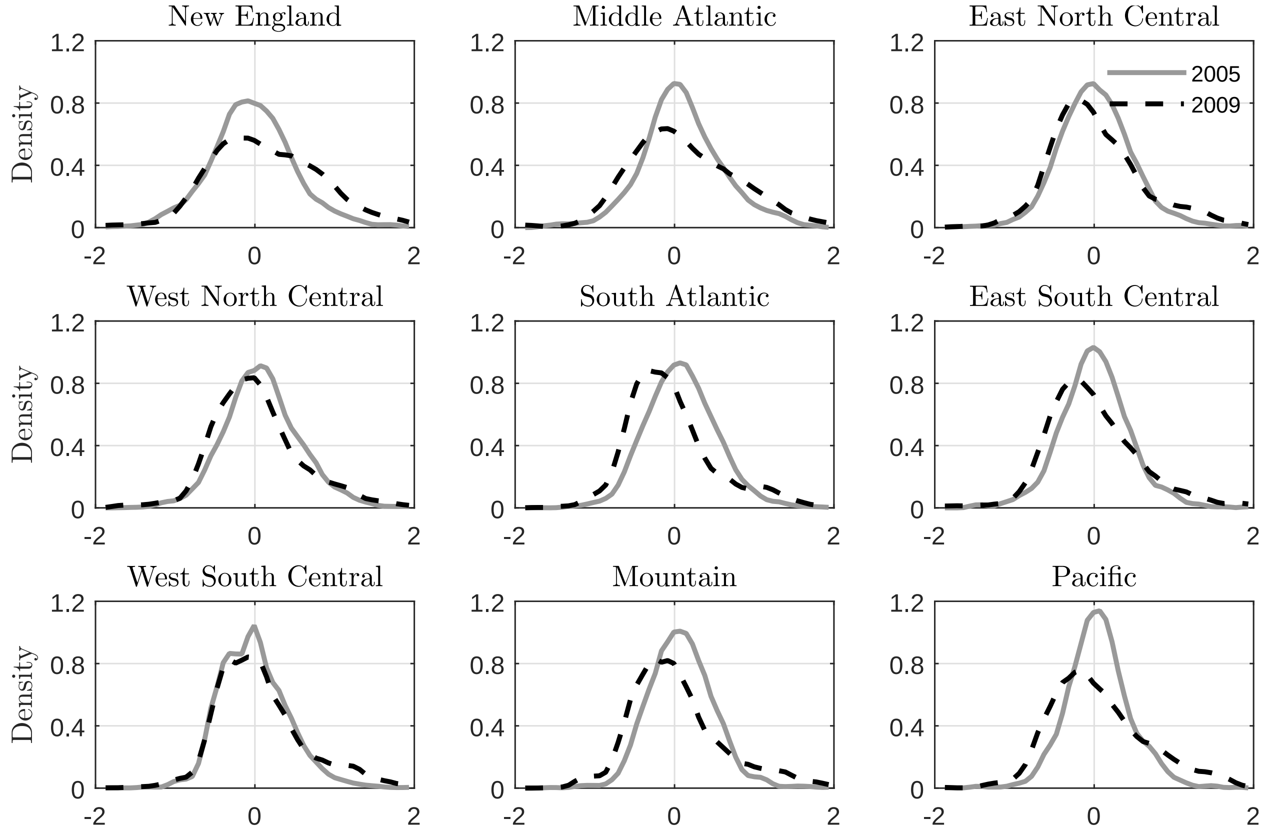

Note: Kernel density of economic TTB for total single-family houses in logs.

For this category of houses, bottleneck effects must have been the main force in the dynamics of construction lags. On the other hand, for houses that were unsold before start, the fat tail to the right is pronounced for the 2009 distribution. Therefore, almost all of the increase in the fat tail to the right in this subsample comes from houses that were unsold before start.

Within houses that were unsold before start, we break this down further and compute the average economic TTB based on the time lag between starting date of construction and the sales occuring date. We plot this for years 2005 and 2009 in the right panel of figure 5. In this figure, houses that were sold before start 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 that are unsold for a long time tend to have lengthy construction lags. In 2009, this pattern has become more apparent. Comparing 2005 and 2009, houses that were sold at the early stage of construction were actually built faster in 2009. However, houses in the bust period that

Figure 5: Distribution of economic construction lags (Built for sale)

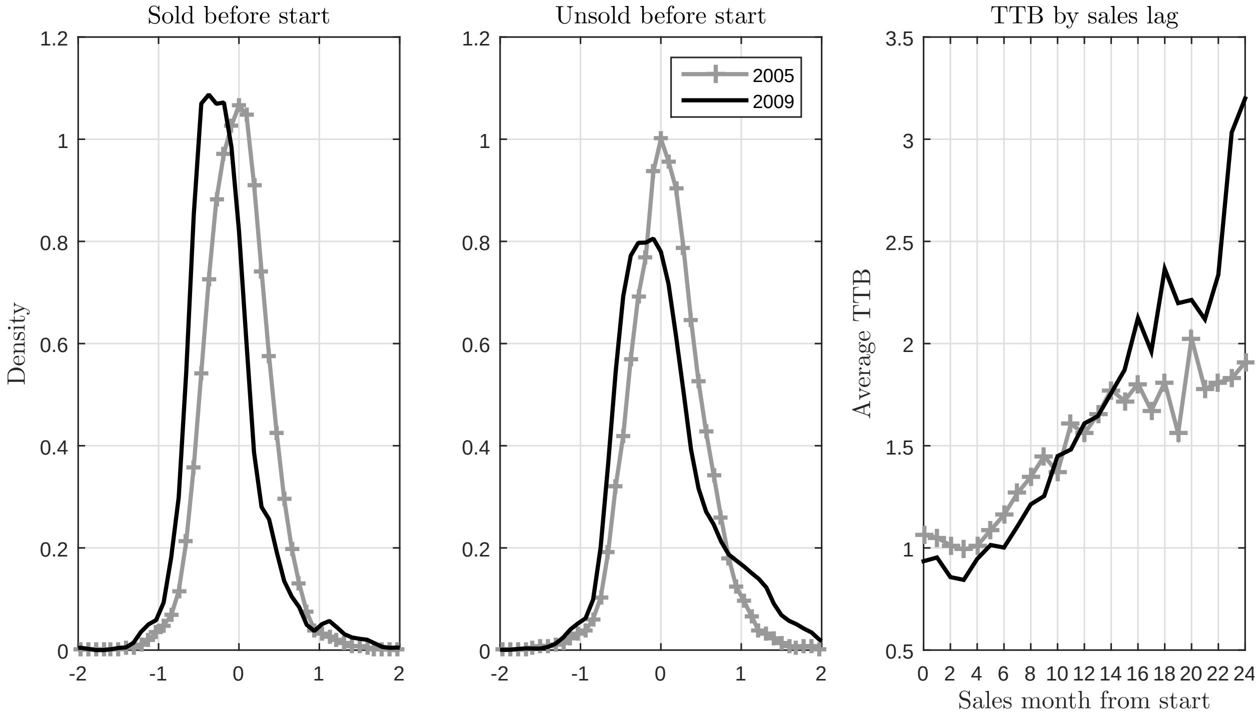

Note: Left two figures are kernel density of economic TTB for built-for-sale houses in logs, for 2005 and 2009. The right figure plots the average economic construction lags by sales month from start, for 2005 and 2009. Houses that were sold before start are lumped into 0, and 24 includes all houses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

faced difficulty in selling early took a much longer time to complete. From this figure, we find that long deferrals in 2009 were especially concentrated on houses that took a long period to sell, even compared to those in 2005.

3.4 Economic construction lags not explained by bottlenecks

To get a quantitative sense of the increase in economic construction lags 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 we explain below. The results are plotted in figure 6.

First, looking into the Census raw data, we find that average TTB increased by 29% between 2003 and 2009. In particular, average TTB increased 12% by 2006 and the remaining 17% increased from 2006 to 2009. Therefore, for "built for sale" houses, the increase in average TTB were 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. Since we control for geographic and building characteristics in our constructed measure, this implies that some of the increase in TTB in the raw data originates from these characteristics. That is, houses that were built during the recent boom-bust era might be larger and higher quality, or built in regions where houses take a longer time to build. From 2003 to 2006, the increase is 12%, and the remaining 10% increased from 2006 to 2009.

Third, since 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 find direct measures of bottlenecks in the construction sector at the regional level. Our approach is to take 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, controlling for the lax bottlenecks during the bust period, the increase in TTB not related to bottlenecks were large and quantitatively comparable to the overall increase in TTB from 2003 to 2009. We conclude that even quantitatively, there is a sizable variation in TTB in the recent bust period that cannot be explained by geographic and building characteristics, as well as the bottleneck effects.

3.5 Summary

We summarize our findings in this section as follows. First, the increase in TTB during the boom was 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 fat tail to the right was the dominating force and average TTB increased. To study the potential channels of this increase in TTB, we focus on houses that are "built for sale" where we find the same pattern. Within this sample, we compare houses that were sold and unsold before start. We find that the fat tail to the right only appears for

Figure 6: Three measures of average construction lags for single-family houses

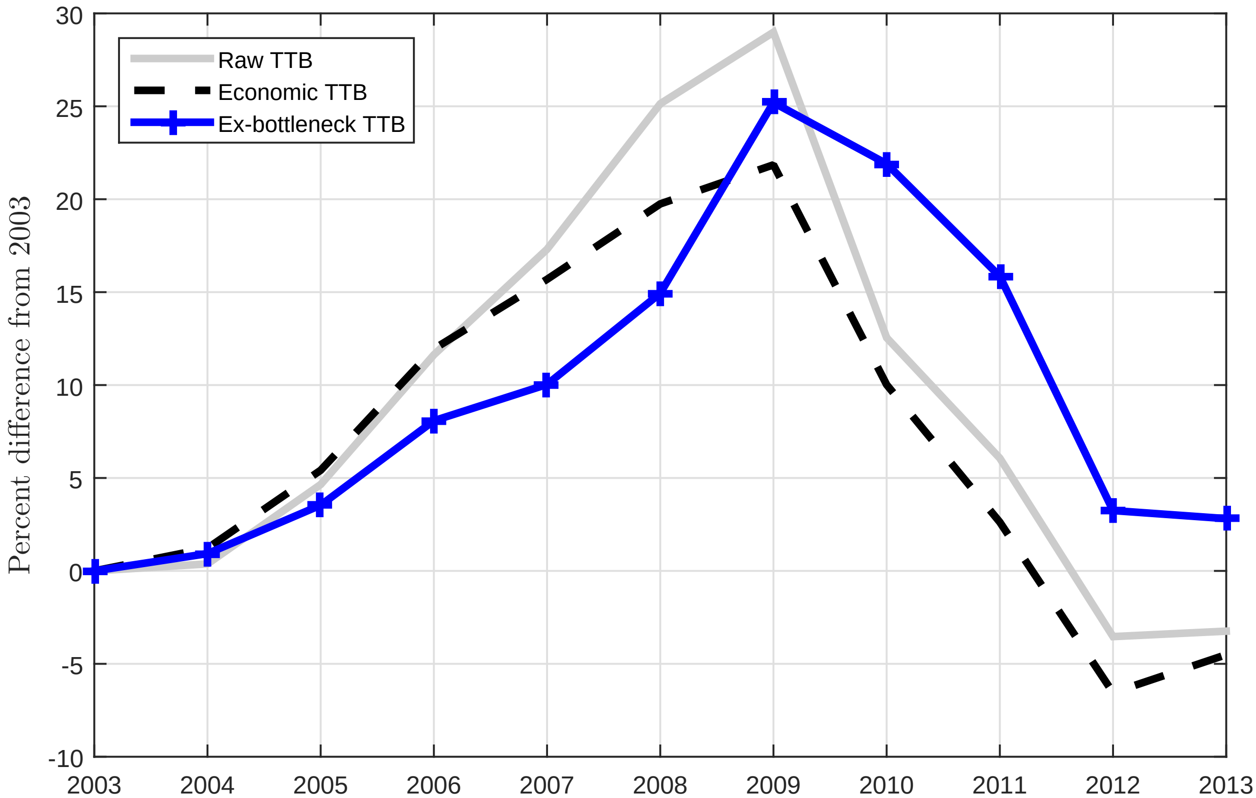

houses that were unsold before start. In fact, during the housing bust, average TTB decreased for houses that were sold before start. We quantify the effects and find that controlling for geographic and building characteristics, as well as for bottleneck effects, average TTB increased by 17% from 2003 to 2009.

4 Model of construction lags

Based on the empirical facts, our next task is to understand the economic forces that deliver the observed increase in TTB during the recent bust. In this section, we ask whether the "wait-and-see" channel is of first-order importance to account for the observed TTB dynamics. Towards this, we illustrate a TTB model of residential investment that incorporates the real options mechanism and quantify the fall in construction intensity during the recent housing bust. We start with a discussion of the real options mechanism in our data and then illustrate our model.

Table 2: Single-Family House Construction Cost Breakdown (2013)

Stage of Construction	Cumulative
1. Site Work (Permit, Inspections, Architecture)	6.8%
2. Foundations (Excavation, Concrete)	16.3%
3. Framing (Roof, Metal, Steel)	35.4%
4. Exterior Finishes (Wall, Windows, Doors)	49.8%
5. Major Systems Rough-ins (Plumbing, Electrical, HVAC)	63.2%
6. Interior Finishes (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 National Association of Home Builders

4.1 Why the real options mechanism?

In a classical investment model, the decision to invest in a project depends on the the net present value of expected cash flows.[14] However, for projects with large sunk costs, investment also entails a significant opportunity cost by making a commitment and giving up the option of waiting. The real options model of investment extensively explored in Dixit and Pindyck (1994) demonstrates that this opportunity cost could be a major factor for investment decisions under plausible parameter values.

We argue that the real options model is a natural fit for our residential construction data for two broad reasons. First, housing supply decisions involve significant costs that are irreversible, such as resources spent on both acquiring permits and building foundations, and the time spent on construction. The irreversible resources and time required on 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 occurring at the later stage of construction. In fact, only 16.3% of the total construction cost is spent when foundations are completed, hence continuing the project from that stage still requires a significant amount of resources and time.

Second, the dynamics of the housing market in the recent bust period also share many features that are consistent with the real options mechanism. Our finding that investment deferrals were especially concentrated on unsold houses during the recent housing bust is consistent with the "wait-and-see" behavior implied by real options models. Moreover, house prices fell a lot (first

[14] Hall and Jorgenson (1967)

moment), and volatility (or uncertainty) surrounding the housing market increasing signficantly (second moment). 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, we plot the two home price indices, divided by the monthly consumer price index. We find that the fall in both house price indices is steep and large during the recent bust. On the right panel, we plot 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). During the recession period, both the house price growth rate volatility and the stock return volatility have increased significantly. These large movements in house prices and housing market uncertainty motivate us to study the first-order importance of the real options mechanism in the housing supply decisions.

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 this model, our goal is to investigate the first-order importance of the real-options view of TTB investment. The model considers the case where a project takes time to complete with multiple irreversible stages each requiring a certain cost. Payoff occurs only after the project is completed, hence irreversible investment decisions are made sequentially at each stage under uncertainty about its future payoff.

The model blends three key elements widely discussed in the investment literature. First, we introduce a real options channel by illustrating a discrete-time version of the sequential irreversible investment model of Majd and Pindyck (1987) and Dixit and Pindyck (1994). Second, we introduce investment adjustment cost parameters to be consistent with the fixed TTB assumption of Kydland and Prescott (1982). Third, we study the effect of uncertainty on TTB investment following the work of Bloom (2009).

Along these elements, the model also incorporates stochastic bottleneck probabilities in investment to control for construction lags originating from either input (labor and material) shortages or weather effects.

Figure 7: Price and uncertainty measures of the housing market

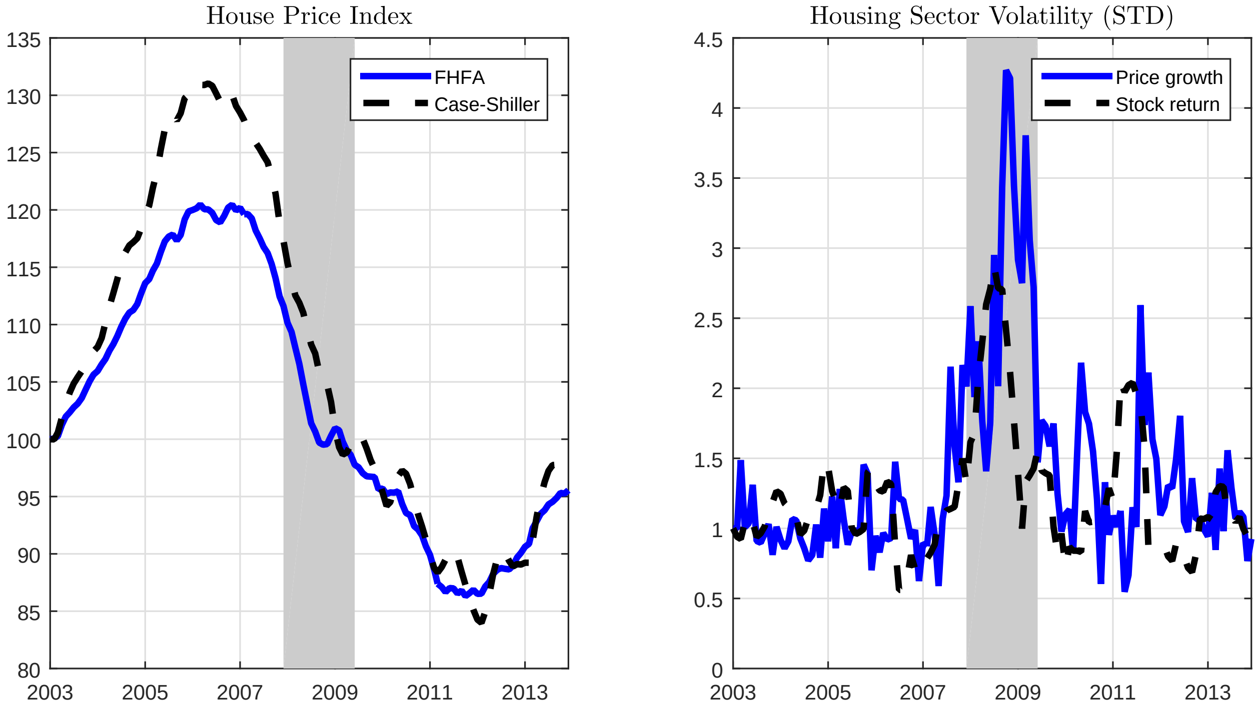

Note: For the house price index, we use both the monthly Federal Housing Finance Agency (FHFA) purchase-only index and the monthly S&P/Case-Shiller national home price index. We also compute two volatility measures. First, using the FHFA index to compute the monthly home price growth rate, we plot the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

TTB and investment Getting into the details, 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 investment takes time to build, with the physical TTB being $\bar{K}/\kappa$.

There is also a bottleneck probability for ongoing projects which follows a Poisson process. In each period, bottlenecks occur with probability p_c. Investment is delayed when bottlenecks occur, and 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 equates 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_{it}^U:

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

The stochastic process for P_{it}^U is

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

where σ is the level of uncertainty of 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 preserved to be zero regardless of the level of uncertainty. For construction that has not started, we normalize the previous value ($P_{it-1}^U = 1$).

In the price process, P^M and σ are assumed as constant. Therefore, builders in the model face house price movements that are driven by idiosyncratic price factors rather than macro price or uncertainty factors. While this might be reasonable in the short run, macro-level movements should also be relevant in the medium run, especially during the recent period. Later in the simulation, we allow the macro price and uncertainty factors to be time-varying as well across regimes, and study the role that these factors played 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 of house i in period t. The upper bound of this value is κ, which is our time-to-build constraint. Third, B_{it} is an indicator function for construction bottleneck of house i in period t.

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

$$
\begin{aligned}
V(P_{it}^U, K_{it}, I_{it-1}, B_{it}; \Lambda) = \max_{I_{it} \in \{0, \kappa\}} \Big\{ &- (1 - B_{it}) I_{it} - \phi(I_{it-1}, I_{it}, K_{it}) \\
&+ \left(\frac{1}{1+r}\right) \mathbb{E} V(P_{it+1}^U, K_{it+1}, I_{it}, B_{it+1}; \Lambda) \Big\},
\end{aligned}
$$

where $\Lambda = \{p_c, P^M, \sigma\}$. Without bottlenecks ($B_{it} = 0$), the cost of real investment is I_{it}. Additionally, the function $\phi(\cdot)$ summarizes the adjustment costs of the model which will be described below. TTB investment decisions are discrete (either invest κ or not).[15]

The evolution of the state variables are

$$
\begin{aligned}
K_{it+1} &= K_{it} - (1 - B_{it}) I_{it}, \\
B_{it} &= \begin{cases} 1 & \text{with prob} \quad p_c \\ 0 & \text{with prob} \quad 1 - p_c. \end{cases}
\end{aligned}
$$

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

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

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

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

Adjustment costs The adjustment cost function $\phi(\cdot)$ carries three variables: previous investment (I_{it-1}), current investment (I_{it}), and total remaining capital (K_{it}). Specifically, we assume that a construction project faces the following adjustment cost function:

$$
\phi(I_{it-1}, I_{it}, K_{it}) = \gamma_0 1_{\{K_{it} = \bar{K} \ \& \ I_{it} > 0\}} + \gamma_1 1_{\{I_{it-1} = 0 \ \& \ I_{it} > 0\}} + \gamma_2 1_{\{I_{it-1} > 0 \ \& \ I_{it} = 0\}}, \tag{2}
$$

[15] 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.

where $1_{\{\cdot\}}$ is an indicator function for the arguments within the parenthesis and $\{\gamma_0, \gamma_1, \gamma_2\}$ are nonnegative values. Three different types of adjustment costs are assumed by this functional form. First, the parameter γ_0 denotes the fixed cost to the extensive margin of construction. This cost parameter incorporates the various suck costs in the decision to start a house. Second, the parameter γ_1 denotes the fixed cost to an upward adjustment of investment for an existing construction project. That is, if previous investment is zero and current investment is positive, then reinitiating the project bears an adjustment cost of γ_1. Third, the parameter γ_2 denotes the fixed cost to a downward adjustment of investment for an existing project, which is exactly the opposite logic of the parameter γ_1.

By setting different values for γ_1 and γ_2, our model incorporates two views on TTB investment. On the one hand, when γ_1 and γ_2 are large, the builder of an existing project always continues to invest until completion as in Kydland and Prescott (1982). On the other hand, when $\gamma_1 = \gamma_2 = 0$, then the model is a discrete-time version of a pure real options model of TTB investment as in Majd and Pindyck (1987).

5 Model solution and simulation

The model is calibrated at monthly frequency. After calibrating the parameters of the model, the optimal investment decision of builders under different levels of uncertainty are discussed. Afterwards, we simulate the model and estimate the parameters to match the simulated TTB distribution to the empirical TTB distribution in section 3.

5.1 Calibration

We start by calibrating the net monthly interest rate r and the physical TTB $\bar{K}/\kappa$.[16] As in Bloom (2009), we set r such that the annual interest rate is 10 percent. In the TTB distribution, less than 10 percent of houses are built within 3 months. Therefore, the physical TTB is set at 4 months.

For the adjustment cost parameters, we set $\gamma_0 = 0.073\bar{K}$, which implies that the initial sunk cost is 6.8 percent of total construction spending (table 2). For the intensive margin adjustment

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

cost parameters of γ_1 and γ_2, we set them both to be zero to study the pure real-options mechanism. We will later set different values for these parameters to study their counterfactual implications.

The three remaining parameters are p_c, P^M, and σ. As discussed in the previous section, 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 cutoff price of investment for each K_i stage of construction, $P^*(K_i;\Lambda)$, such that

$$I_{it} = \begin{cases} \kappa & \text{if } P_{it}^U \geq P^*(K_i;\Lambda), \\ 0 & \text{if } P_{it}^U < P^*(K_i;\Lambda). \end{cases} \tag{3}$$

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

For illustrative purposes, we numerically solve the model by assuming zero bottlenecks ($p_c = 0$) and $P^M = 2$. For uncertainty, the following 3 values are considered: 0.2, 0.4, 0.6. The three parameters $\{p_c, P^M, \sigma\}$ will be estimated in the next section.

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

$$P^*(K;\sigma = 0.6) > P^*(K;\sigma = 0.4) > P^*(K;\sigma = 0.2), \qquad \forall K \in \left\{ \frac{\kappa}{\bar{K}}, \frac{2\kappa}{\bar{K}}, \cdots, \frac{\bar{K} - \kappa}{\bar{K}}, 1 \right\}.$$

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 cutoff price falls as the construction process nears completion:

$$P^*\left(\frac{\bar{K}-\kappa}{\bar{K}};\sigma\right) > P^*\left(\frac{\bar{K}-2\kappa}{\bar{K}};\sigma\right) > \cdots > P^*\left(\frac{2\kappa}{\bar{K}};\sigma\right) > P^*\left(\frac{\kappa}{\bar{K}};\sigma\right), \qquad \forall\sigma \in \{0.2, 0.4, 0.6\}.$$

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

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

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

On the other hand, a homebuilder could also look at the bright side in starting construction under high uncertainty, since now a highly profitable outcome is also possible. When that outcome realizes, it might be too late for the builder to reap that profit when the builder has not already set foot in the project, since completion takes time. In this case, the builder might want to take some bet on good realizations and rather engage in "wait-and-see" behavior after at least building the foundations of the house. Therefore, TTB may increase for buildings that started after observing this high uncertainty.

Uncertainty and TTB To answer the above, we now study whether TTB investment is sensitive to uncertainty even after controlling for housing starts. Using our model, we ask whether periods of high uncertainty are also when TTB investment decisions become even more cautious than housing start decisions.

In figure 8, we plot the log cutoff price of each stage relative to its initial cutoff price when the project started, scaled by the level of uncertainty:

$$W_K^* - W_1^* = \frac{1}{\sigma}\big[\log P^*(K;\sigma) - \log P^*(1;\sigma)\big]. \tag{4}$$

Figure 8: Uncertainty-invariant cutoff price relative to start cutoff price

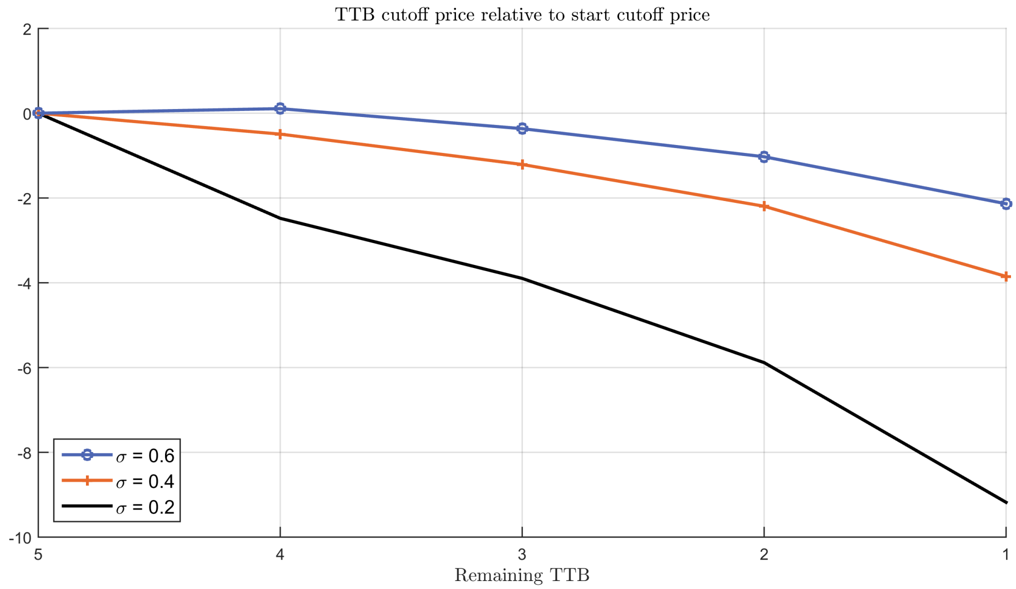

Note: The plotted measure is $(1/\sigma)[\log P^*(K;\sigma) - \log P^*(1;\sigma)]$, and the x-axis is $5 \times K$. For this plot, we assume physical TTB to be 5 months. For the results that follow, we set physical TTB at 4 months.

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

In the figure with low uncertainty ($\sigma = 0.2$), the scaled relative cutoff price steeply falls as construction approaches completion. At the last stage of construction, the builder needs to observe a 9 standard deviation fall 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 investment into the project until completion.

On the other hand, with medium uncertainty ($\sigma = 0.4$), the fall in the scaled relative cutoff price is only gradual. Even at the last stage of construction, the builder defers the project when observing a 3 standard deviation fall in the price innovation relative to the initial price innovation. In fact, with high uncertainty ($\sigma = 0.6$), the first-stage TTB investment cutoff price becomes even higher than the start cutoff price. Therefore, when uncertainty is high, the same change in the uncertainty-invariant price innovation would lead 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) housing start cutoff price increases, (ii) TTB investment cutoff prices increase, and (iii) TTB investment cutoff prices increase more than the start cutoff price. 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 bottleneck during the housing boom-bust period.

5.3 Model simulation

In this section, the model is simulated and estimated to study the dynamics of TTB investment with regards 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.[17] 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 decide whether or not to initiate a new construction.

The simulation is repeated 100 times for 45 years, with a burn-in period of 30 years. Our simulated TTB time series is the average duration of completed projects 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 may experience low price realizations for a long time. When these projects

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

eventually pick up, they would significantly drive up the average, and may 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, by dropping the project if it is incomplete by then.

5.3.2 Estimation details

The goal in the simulation exercise is to understand the drivers of TTB during the housing boom-bust and recovery period. Towards this, matching the distributional properties of TTB is useful in providing use a better understanding of TTB dynamics that the mean TTB dynamics does not reveal. In particular, our model allows for three key parameters that govern the bottleneck effect (p_c), the price effect (P^M), and the uncertainty effect (σ). We utilize our micro data to decompose the three structural channels that the model allows for, and to understand whether uncertainty effects played a distinctive role during the housing cycle.

We estimate the parameters for each odd year in our data (2003, 2005, 2007, 2009, 2011, 2013). Since the average 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 CPI (BLS).

Simulated Method of Moments For each year $t \in \{2003, 2005, 2007, 2009, 2011, 2013\}$, the three values $\{p_{ct}, P_t^M, \sigma_t\}$ should be pinned down. Given an initial value P_{2003}^M, the remaining price effects P_t^M are imposed by using the growth rate of the house price data described above. Therefore, there are three values to estimate in the initial year 2003, and two afterwards. We use simulated method of moments to pin down each of these. The moments we use in each year for estimation are (1) mean TTB, (2) $Pr(TTB \leq 6)$, (3) $Pr(6 < TTB \leq 9)$, (4) $Pr(10 < TTB \leq 12)$, where TTB data is our economic TTB. These distributional aspects well identify the two effects that we estimate. Given a certain mean TTB, bottleneck effects govern the exogenous movements of TTB across the overall distribution, while uncertainty effects govern the tail distribution. The 4 moments we use well characterize these distributional aspects.

Figure 9: Simulated method of moments: TTB distribution

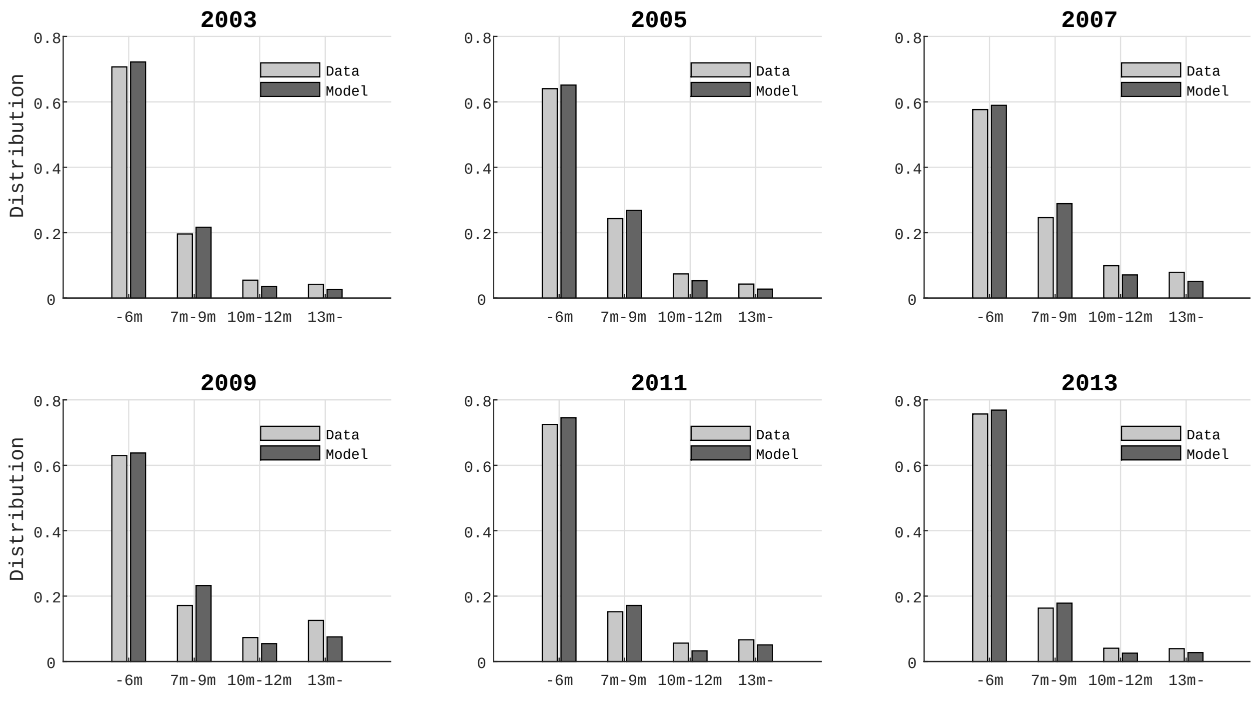

In year 2003, we additionally estimate the initial price level P^M_{2003}. Intuitively, when this aggregate price level is low, projects may start with a good idiosyncratic price draw since the value of waiting for a better price draw is small in that case, but not continue to invest afterwards if the idiosyncratic price draw does not increase enough to cover up the low aggregate price level. Therefore, the completion rate of initiated projects is helpful in identifying these channels. This moment cannot be computed using our data. Fortunately, the Census produced the abandon rate of single-family houses in 2002, which is 0.5%. Therefore, in year 2003, we use 5 moments (the four moments above and the abandon rate) to estimate the three values $\{p_{c2003}, P^M_{2003}, \sigma_{2003}\}$.[18]

5.3.3 Baseline simulation results

Figure 9 reports both the data and the simulated moments for the obtained estimated values. The model fit of the estimated values are decent across the distributional moments in each year. In figure 10, mean TTB as well as the estimated values of bottlenecks and uncertainty across time

[18]When comparing the results across time, it is important to fix the estimation methods. Our estimation procedure is different at the initial period with the abandon rate data. In the analysis below, we re-estimate the values p_{c2003} and σ_{2003} by imposing the the estimated value for P^M_{2003}, and dropping the abandon rate data. However, the discussion of the results remain even when we use the initially estimated p_{c2003} and σ_{2003}.

Figure 10: Simulated method of moments: TTB mean

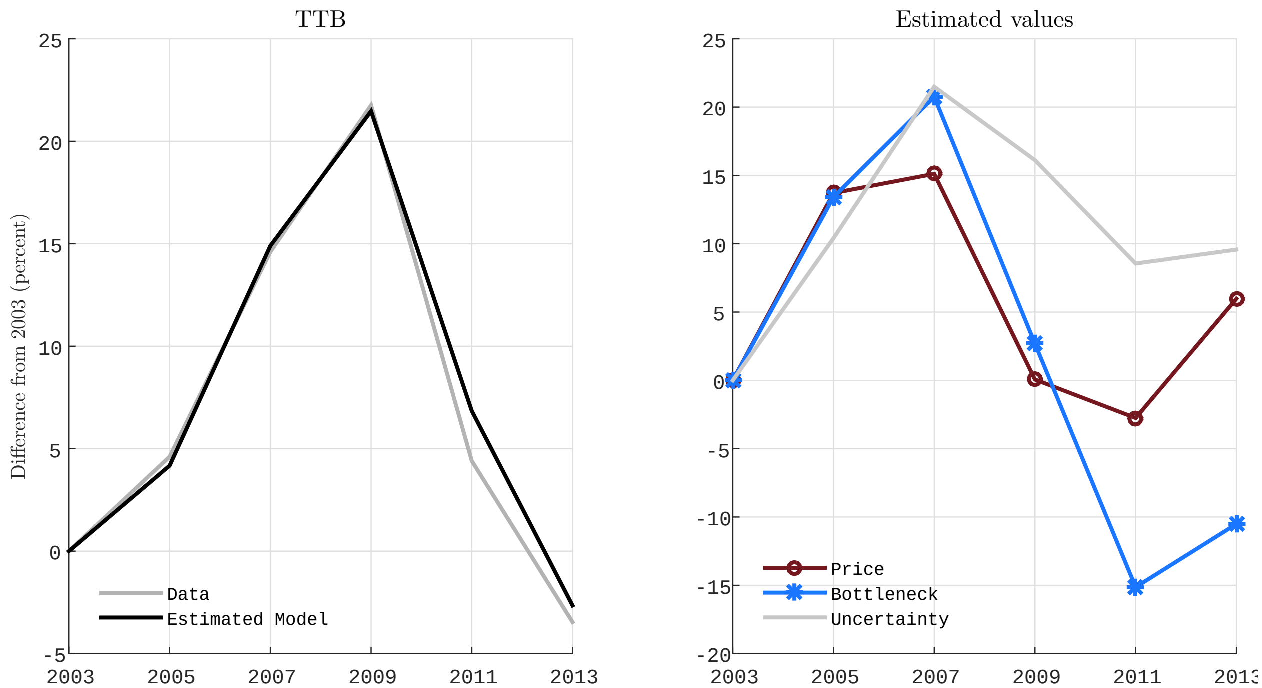

are plotted. The estimated values for the three values in 2003 are $P^M_{2003} = 1.935$, $p_{c2003} = 0.360$, $\sigma_{2003} = 0.431$. Again, the model fit of mean TTB is decent. Moreover, given the price path, both the bottleneck and uncertainty parameters increase beyond 20 percent of their respective values in 2007, and then fall afterwards. The fall in uncertainty is much more gradual than the fall in bottleneck probability.

Figure 11 plots the counterfactual TTB when bottleneck effects are shut down at its 2003 level. Consistent with our plot in figure 6, the estimated model favors bottleneck effects in driving TTB dynamics during the boom period but not afterwards. Therefore, we conclude that bottleneck effects are not the driving force of TTB dynamics during the housing bust period and afterwards.

To understand the overall channels of TTB, we plot the historical decomposition of TTB with our three effects in figure 12. Since 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, one could fix p_c and σ at their respective 2003 values, and plot the mean TTB when only prices move. On the other hand, one could fix P^M at its 2003 value and plot the implied mean TTB when other values change across time. Subtracting this counterfactual value from the model estimated mean TTB, one could obtain the net price effects. Figure 12 reports the average of these two ways of

Figure 11: Simulated method of moments: Ex-bottleneck TTB

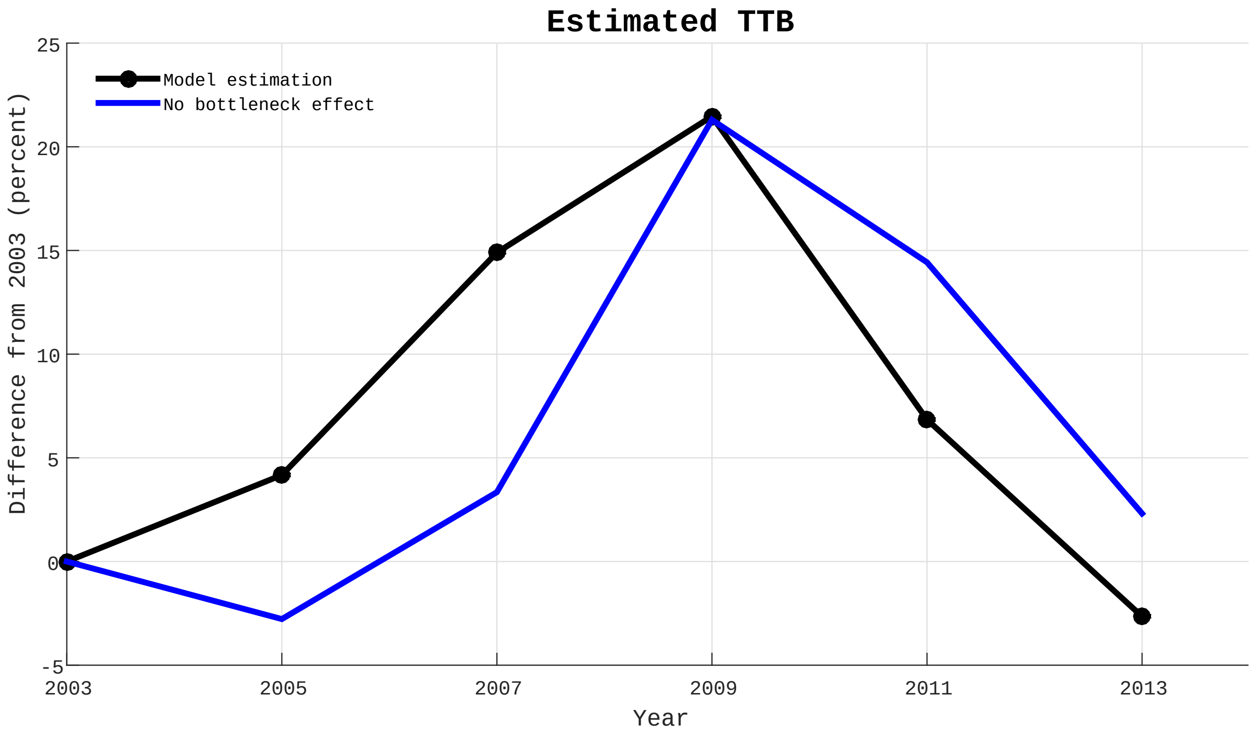

computing the historical decomposition.

Several things are worth noting. First, the increase in TTB in 2005 and 2007 are driven by a combination of bottleneck and uncertainty effects. On the other hand, the large rise in prices during these periods countered 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. Third, after 2009, the fall in uncertainty effects drive the fall in mean TTB.

To sum up, the model finds that after 2007, uncertainty effects dominate the dynamics of mean TTB.

5.3.4 Alternative simulation results

In the baseline simulation, we find that the observed house prices are not sufficient to account for the dynamics of TTB during the recent period. Importantly, bottleneck effects are important for the housing boom period, and uncertainty effects dominate the housing bust period. In this section, we ask whether this conclusion is robust to using other measures of house price data. One

Figure 12: Simulated method of moments: Historical decomposition of TTB

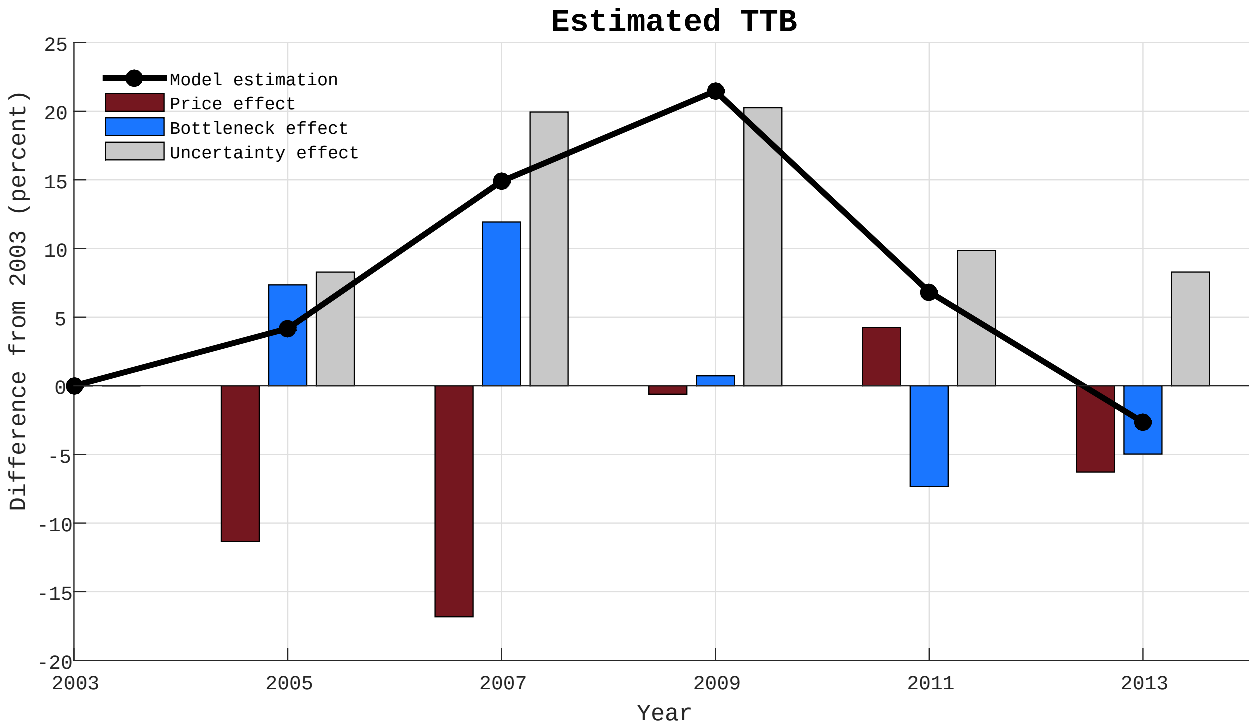

thing to note is that the movement of our price data in figure 10 is smaller than the Case-Shiller or FHFA price data in figure 7.[19] However, since our sample is biased towards locations with volumns of new construction, using these existing house price indices will overestimate the price movements of new houses, since these indices are driven by locations with low supply elasticity (Saiz, 2010).

Still, one driver of the increase in new house price during the recent housing boom period might be the rise in building costs. Moreover, if building costs did not fall as fast as house prices during the bust period, this difference may have contributed to TTB dynamics during the bust period instead of the rise in uncertainty. To gauge the first-order importance of the uncertainty effect under this building cost effect, we additionally use the building cost index constructed by the "Engineering News-Record."[20] Therefore, we deflate our housing price measure by this building cost index and re-estimate the model.

Figure 13 plots the new estimation with the newly imposed price index. Model fit remains

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

[20]As described in 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.*"

Figure 13: SMM: TTB mean (price deflated by building cost index)

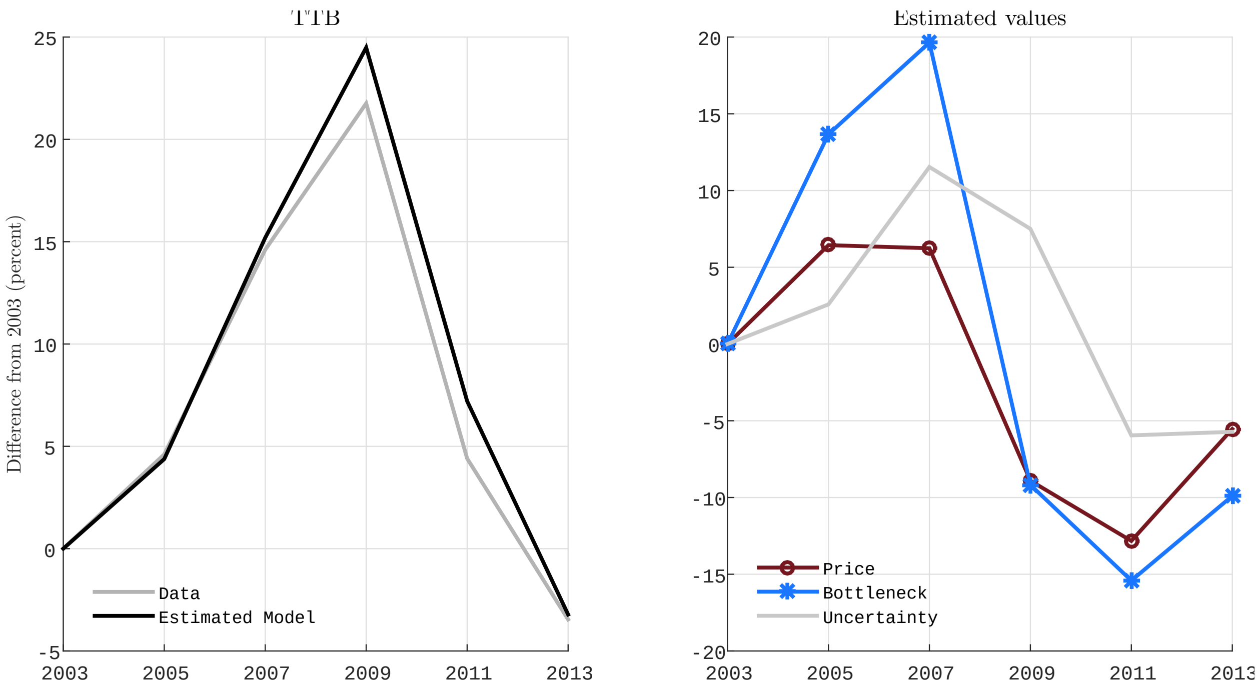

decent, although the 2009 increase in TTB is overestimated. For the estimated values, bottleneck effects remain the same in size until 2007, but with the small increase in prices, uncertainty effect is also estimated to be smaller. On the other hand, the bottleneck value becomes smaller than its 2003 value in 2009, while uncertainty still remains at around 7.5 percent above its 2003 value.

Figure 14 reports the historical decomposition of TTB with the newly imposed price index. With the alternative price index, price effects dominate the increase of TTB in 2009. However, with the fall in the bottleneck effect, uncertainty still remains to be an important channel in driving TTB dynamics. Importantly, in 2011, the fall in bottlenecks and uncertainty drive down the increase in TTB although prices effects are estimated to drive up TTB signficantly during this period. Overall, movements in TTB since 2009 are consistent with movements in the estimated uncertainty effects during this period.

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

Figure 14: SMM: Historical decomposition of TTB (price deflated by building cost index)

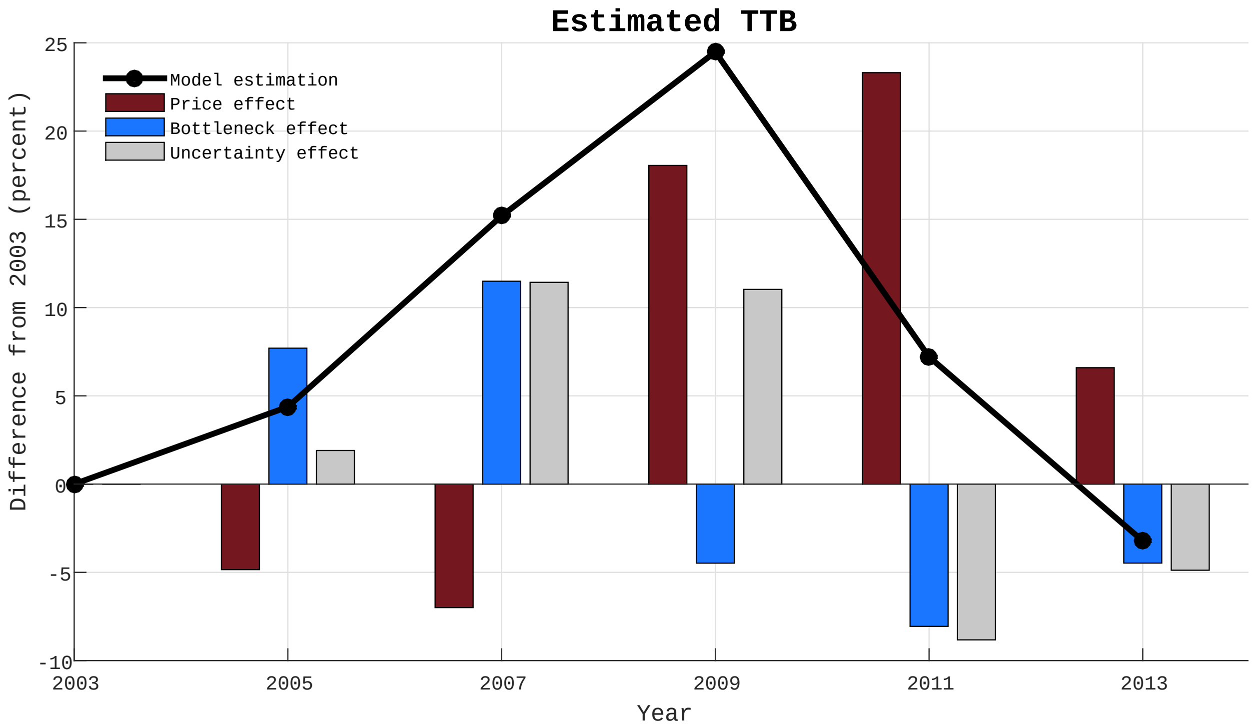

to look into several housing supply implications that follow from the channels studied so far. While understanding the dynamics of construction lags are interesting in itself, it is even more important by their implications on housing supply dynamics. We deliver two insights through our model. First, we look into its implications on 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 the results with a simulation under fixed construction intensity.

6.1 Residential investment with variable construction intensity

Housing starts and residential investment are key data of interest to macroeconomists and policymakers, both because of their high volatility and their stable lead-lag structure with GDP (Davis and Heathcote, 2005). In this part, 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 the transition dynamics of our model by looking into the perfect foresight equilibrium with a regime change. Starting from

Figure 15: Transition dynamics with a price fall

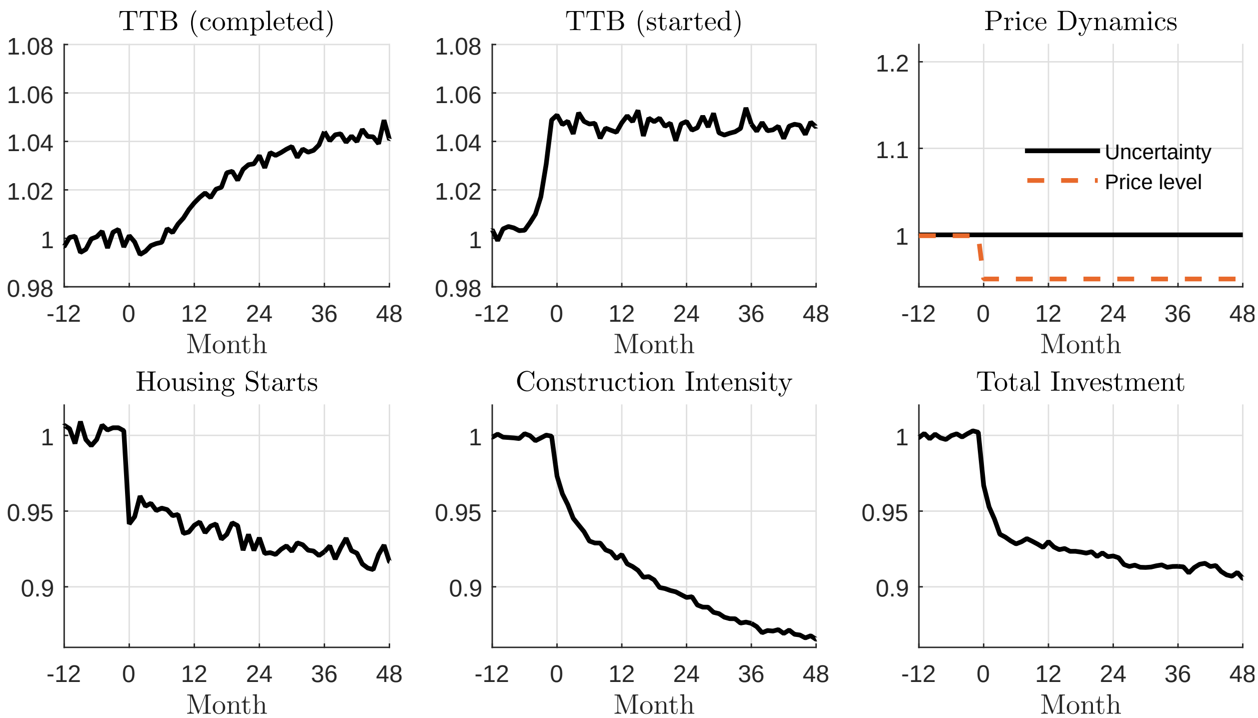

the 2003 values, figure 15 plots the transition of several variables when the house price falls by 5 percent. We find that housing starts fall by 5 percent on impact, and construction intensity of incomplete houses fall by less than that on impact. Therefore housing starts have a bigger impact on total investment relative to construction intensity with this price fall. Looking into implications on TTB, there is a gradual increase in TTB for construction that started before the regime shift, and TTB of completed houses gradually increase up to 4 percent higher than the previous level.

Figure 16 plots the same transition with also a 20 percent increase in uncertainty. Compared to a pure price fall, we find that construction intensity falls more than housing starts on impact. Therefore, the initial fall in total investment is dominated by intensive investment rather than extensive investment. Moreover, with an increase in uncertainty, TTB for houses that started before the regime shift overshoots the new steady state of TTB. That is, builders who started their projects before the price and uncertainty change now find themselves in a different world and engage in a "wait-and-see" behavior beyond that of newly started builders.

In figure 17, we plot total housing starts and residential investment based on the price fall and uncertainty increase in figure 16. To gauge the importance of TTB dynamics (i.e. intensive invest-

Figure 16: Transition dynamics with a price fall and uncertainty rise

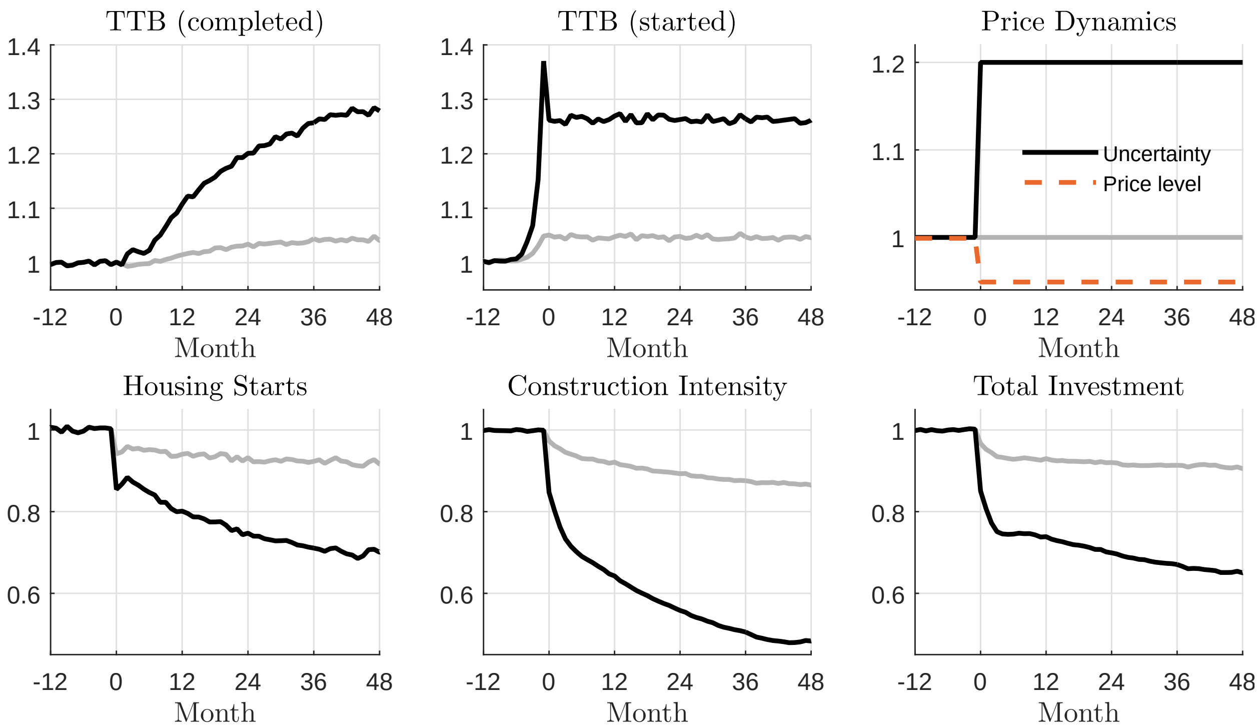

Note: Grey line depicts figure 15, the pure price fall effect.

ment) in affecting residential investment, we also plot a counterfactual time series for residential investment, where we shut down the TTB channel, by assuming that there are no deferrals and hence economic TTB and physical TTB are the same.

We make two points from this figure. First, for the initial periods after our experiment, construction intensity of incomplete houses drives the dynamics of residential investment rather than housing starts. For the first 6 months, the counterfactual residential investment, which only contains information on housing starts, falls by only a half of the actual simulated residential investment. The initial movements in residential investment are dominated by the intensive margin rather than the extensive margin. Second, while the counterfactual residential investment series lags housing starts, the actual residential investment responds immediately to the shock and hence does not lag housing starts.

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

Figure 17: Housing starts and residential investment

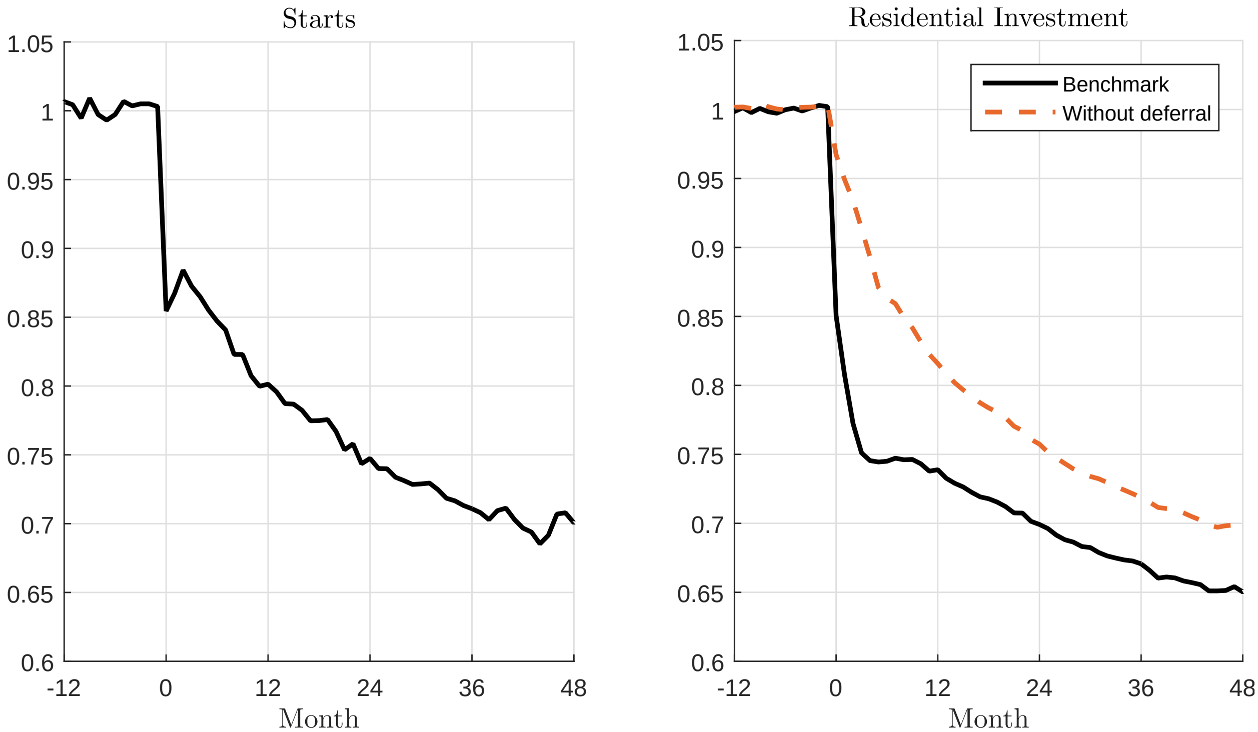

6.2 Housing starts with variable construction intensity

In this section, we ask through our model how variable construction intensity affects housing starts. In particular, we compare our result with that under the commonly used fixed TTB assumption. Figure 18 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 initial fall in housing starts is muted, but the fall is persistent and large, only slowly converging towards a lower steady state.[21]

In figure 19, we decompose the dynamics of housing starts in figure 18 into (i) the start rate per builder and (ii) the number of available builders. Looking into the first column, we find that the impact effect on the start rate is smoother with variable intensity than with fixed intensity. In the second column, the number of available builders under two scenarios are plotted. For both

[21] In our model, the average starting price of houses under fixed intensity is higher than that under variable intensity. Since existing projects are relatively "unhealthy" under variable intensity, with a lower price, these houses now more frequently defer during construction, which leads to a higher average economic TTB. Therefore, there is a lower flow of available builders in the new price regime, which leads to the lower level of starts.

Figure 18: Housing dynamics with variable and fixed construction intensity

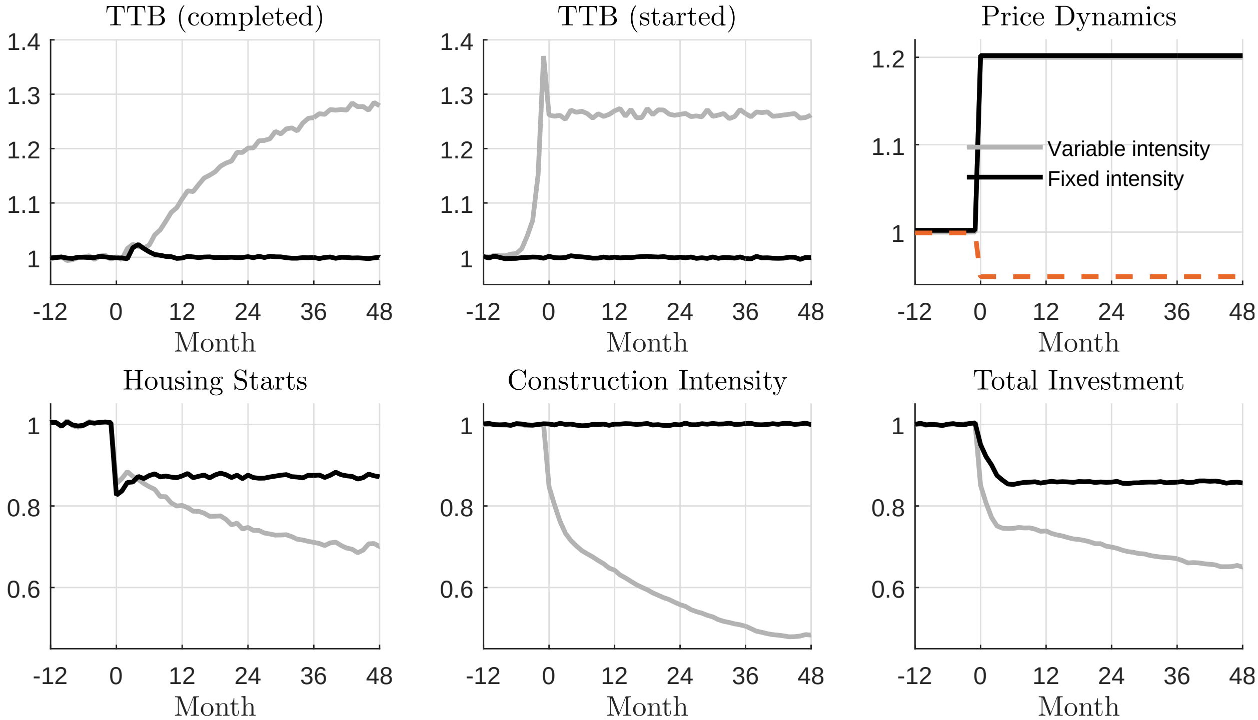

types of builders, available builders increase initially for a longer period. This is due to completed builders deciding not to start a new project. However, the two types of builders depart in the later period since with variable construction intensity, pre-existing projects are all delayed and hence inventory overhang leads to less new projects available. On the other hand, available builders do not fall below the steady state in the second column, since no delays occur with fixed construction intensity.

For the economic interpretation, we find that with fixed intensity, investment decisions are all loaded onto housing starts. Therefore, with the new information on price dynamics, housing starts take a big fall on impact. However, since pre-existing projects are committed to complete on schedule, there is a relatively stable flow of builders completing a project. These builders are quickly available to start a new house by drawing a new price. On the other hand, with variable intensity, investment decisions are spread out at each stage of construction. Therefore variable intensity builders that are considering new projects do not respond as sensitive as the fixed intensity builders. However, since pre-existing projects are now deferred with the new pricing information, completion rate falls and available builders for new projects remain lower for the long time. This

Figure 19: Housing starts

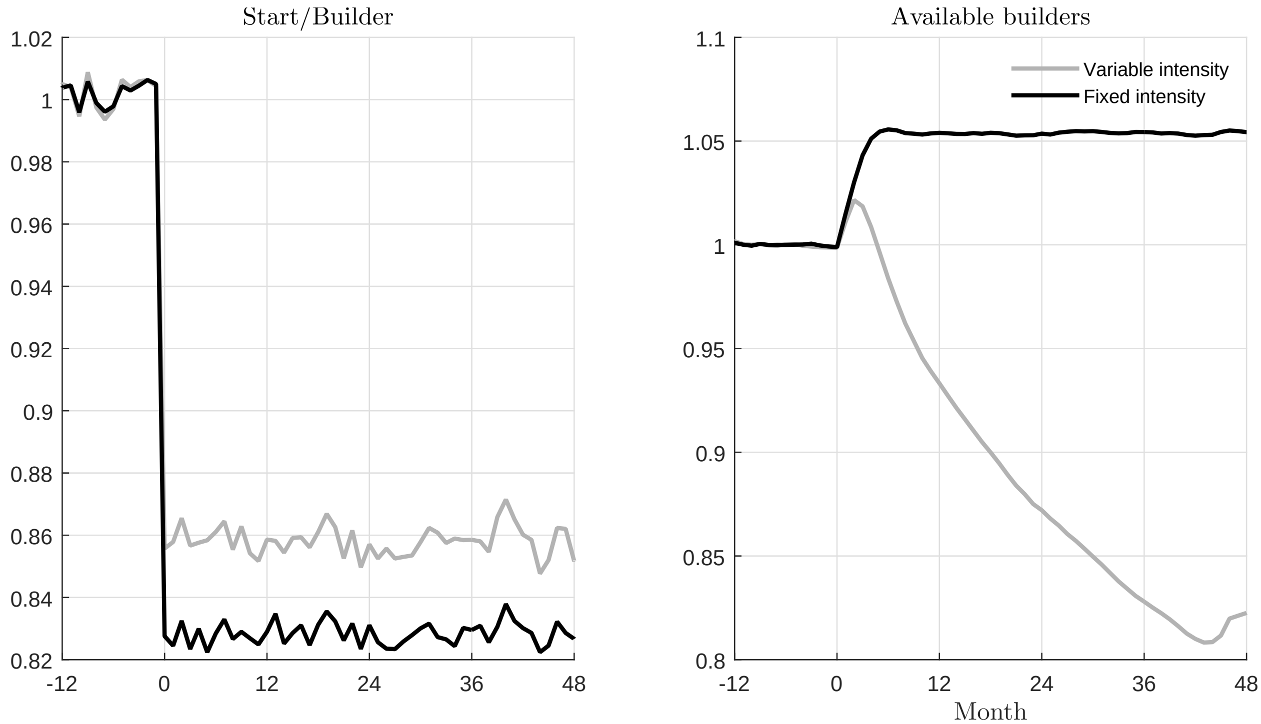

overhang of incomplete houses generates a housing slump since the number of availble builders shrink.

6.3 Housing starts in the boom-bust period

With the implications of our model on housing starts, we finally ask how much our model channels would account for the housing start dynamics in the recent housing cycle. Using the estimated values for price, uncertainty, and bottleneck effects that match the TTB distribution in the data, we plot the implied housing starts in each period. In figure 20, the implied housing starts based on the baseline estimation are plotted, along with housing starts of single-unit houses. Based on 2003, the model claims a 1 percent increase in housing starts in 2005 and a 20 percent fall in housing starts in 2009. In the data, housing starts increased by 10 percent in 2005 and fell by 70 percent in 2009. Therefore, our model does not explain the housing boom and accounts for 30 percent of the fall in housing starts in the housing bust period.

In short, while our model is estimated to match the intensive investment distribution of residential construction, these channels have limited influence in terms of generating the observed housing

Figure 20: Housing starts data and model

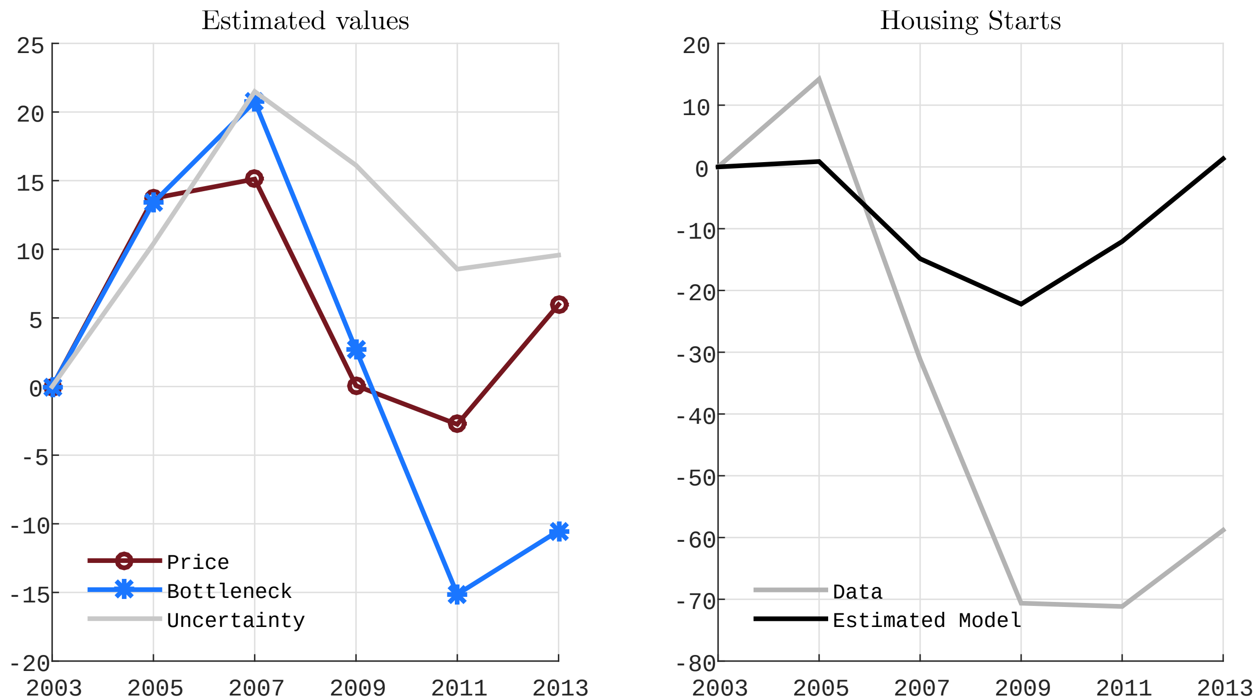

start dynamics in the data. Obviously, the model abstracts from many variables that affected housing starts in the housing boom-bust period. For example, the credit boom and bust would have affected the initial availability of construction loans that may affect disproportionately the extensive margin, and permit processes may have also changed in this period. Moreover, a micro-founded model of bottlenecks during the housing boom period would imply that builders engage in multiple projects at a time which would also lead to an increase in the number of buildings started in this period.

With these richer aspects on the housing start margin, the model may be able to account better for the movements of housing starts during the recent period. However, our intensive investment channel itself still accounts for a moderate fraction of the dynamics of housing starts in the data.

7 Conclusion

In this paper, we document some new facts on the distribution of residential construction lags across the US. Importantly, we emphasize that the fall in economic activity is not limited to extensive investment, but also expands to intensive investment. Contrary to the notion that time-to-build projects that have already started are costly to stop, we find that a significant portion of it has been deferred during the recent housing bust. Given the large movements in house prices during the era, we study a model where time-to-build investment responds solely to prices and uncertainty and simulate the model with the level of price and uncertainty movements observed in the recent recession. We find that our real-options mechanism is capable to account for all the drop in intensive investment during the housing bust. We argue that the real-options mechanism is of first-order importance for construction lags during the recent recession.

Before concluding, it is important to note that in this paper, we left aside the financial frictions channel in time-to-build investment projects. We are indeed aware that the construction sector is a levered industry, and that 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 to the housing bust, and the overall financial constraints may have exacerbated the aggregate housing market collapse. Our aim is rather on addressing the first-order importance of the real options mechanism in the recent housing cycle given the unprecedented magnitude of house price dynamics, than on providing a complete picture of the housing supply side behavior. As discussed in the previous section, there is still room for improvement in accounting for housing start dynamics in the recent housing cycle. While we find that the real-options mechanism is capable of accounting for most of the investment activity for projects under construction, its potential endogeneity with financial frictions is a topic of interest. We leave this out as a future research project.

References

Arezki, R., V. R. Ramey, and L. Sheng (2015). News Shocks in Open Economies: Evidence from Giant Oil Discoveries. *UC San Diego mimeo*.

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

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. and J. Gyourko (2005). Urban Decline and Durable Housing. *Journal of Political Economy 113*(2), 345–000.

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.

Hall, R. E. and D. Jorgenson (1967). Tax Policy and Investment Behavior. *American Economic Review 57*(3), 391–414.

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.

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 (2012). Housing Dynamics over the Business Cycle. *NBER Working Paper 18432*.

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.

Mian, A. and A. Sufi (2011). House Prices, Home Equity-Based Borrowing, and the U.S. Household Leverage Crisis. *American Economic Review 101*(5), 2132–2156.

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

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.