

Global Research Unit

Working Paper #2021-024

Is housing still the business cycle? Perhaps not.

Richard K. Green, University of Southern California

© 2021 by Green. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Is housing still the business cycle? Perhaps not.

Richard K. Green

University of Southern California

May 24, 2021

Abstract

In 1998, I published a paper that showed that under a wide range of specifications, residential investment led GDP, while non-residential investment did not. That paper was followed by a number of others, including Coulson and Kim (2000), Davis and Heathcoate (2005) and Leamer (2007) that used more sophisticated techniques than my paper, but found the same outcome—that residential investment led GDP. Leamer famously announced that housing was the business cycle. But in light of the Great Financial Crisis, the subsequent crash in residential investment, and the fundamental changes in the mortgage market, I thought it worth revisiting housing as a leading indicator. I have found that it is a much weaker leading indicator than before, and that it is much less sensitive to Federal Reserve Policy—especially changes in the Federal Funds Rate—than before. It is possible that the increasing stringency of local land use policy had interfered with the ability of the Federal Reserve to use housing as an instrument of monetary policy.

I. Introduction

Twenty-three years ago, as of this writing, I published a paper in *Real Estate Economics* (Green 1998) on residential investment's remarkable track record for predicting subsequent GDP. The paper used a straightforward time-series technique—Granger causality—across a variety of specifications and found that while detrended residential investment Granger caused detrended GDP, non-residential investment did not. Moreover, residential investment was, at that time, orthogonal with respect to GDP—that is, adding GDP to a vector autoregression explaining residential investment did not improve the ability to forecast residential investment.

The fact that Granger causation went in one direction—from housing to GDP—was important to identify the forecasting parameters from housing to GDP. The paper made no claims about causation, but it did argue that following residential investment was useful for those making economic forecasts. This was followed by papers by Coulson and Kim (2003), Davis and Heathcote (2005), and Leamer (2007); the last of these was titled "Housing is the business cycle." The papers, all using different methods from mine, confirmed what I found—that residential investment led the broader economy, without being led by it.

Enough time and enough events have passed that it is worth revisiting whether the empirical regularities found in the research cited above remain. In particular, the economy has since 2007—the end of a housing boom—seen a housing bust, a housing price, but not investment, recovery, and, of course, a pandemic. In light of these shocks, it would be reasonable to ask whether the relationship between housing investment and economic activity remains.

As such, this paper will take the techniques used in Green and Coulson and Kim to update the results. It finds that housing continues to lead the business cycle, but that the business cycle, to a limited extent, leads housing. Coulson and Kim recommend a decomposition of GDP to determine the mechanisms through which housing influences and is influenced by the housing market, and this paper will consequently update their results.

The paper is organized as follows. It begins with a brief review of the techniques used in the papers listed above and their findings. It then presents descriptive evidence of how housing has changed relative to the business cycle beginning with and subsequent to the great financial crises. It then presents pretests of the data we use to update the findings of past papers. It then runs a series of bivariate and multivariate tests examining how housing leads and is led by the business cycle and individual components of that cycle. The chapter finishes with some forecasts and implications.

II. Literature Review

A number of papers over the past 20 years have looked at the relationship between housing and the business cycle. One interesting takeaway is that in the US, housing investment is an unusually strong leading indicator (Leung 2003). For this chapter, I will focus on three papers listed in the introduction, as they are the foundation for updating past findings.

Green (1998) uses simple Granger Causality tests using a variety of specifications. The purpose of a Granger test is to take a stationary time series and find whether the lagged values of one variable may improve the ability to forecast another variable, taking into account the lagged value of that other variable. In other words, one begins by estimating the equation:

$$y_t = \alpha + \sum_{i=1}^l \beta_{t-i} y_{t-i} + \sum_{i=1}^l \gamma_{t-i} x_{t-i} + \varepsilon$$

Where y is the variable to be forecast, and x is the variable contributing to the forecast, and ε is an error term. The next step is an F-test of whether the γ_{t-i} coefficients are jointly different from zero at customary levels of significance (Granger 1980). If they are jointly different from zero, we say that x Granger Causes y .

Green (1998) used residential and non-residential investment as x variables in separate regressions and found that residential investment Granger Caused GDP, while non-residential investment did not. At the same time, when the regressions were reversed, GDP Granger caused non-residential investment but did not Granger Cause residential investment, suggesting that changes in residential investment were a function of exogenous shocks rather than embeddedness in the broader economy.

These results were consistent through a variety of specifications, including models with different lag lengths. The results were also subject to the Brown, Durbin and Evans (1975) CUSUM test for coefficient stability. The paper found that there was a structural break in the coefficients in the year 1980, but that running separate, stable regressions before and after that break had no impact on the qualitative results: in both periods, residential investment Granger Caused GDP, non-residential investment did not, and GDP Granger caused non-residential investment, while it did not Granger Cause residential investment. We shall revisit those results with the benefit of an added 28 years of data.

Coulson and Kim (2000) went a step further by investigating the effect of residential investment on various components of GDP: consumption, non-residential investment, and government spending. They then used a technique that allowed them to uncover the channels through which non-residential investment operated. In the end, they determined that residential investment had an impact on consumption, and that it did indeed influence GDP.

Leamer (2007) summed up his views on how housing revealed itself to be a leading indicator:

Of the components of GDP, residential investment offers by far the best early warning sign of an oncoming recession. Since World War II, we have had eight recessions preceded by substantial problems in housing and consumer durables. Housing did not give an early warning of the Department of Defense Downturn after the Korean Armistice in 1953 or the Internet Comeuppance in 2001, nor should it have. By virtue of its prominence in our recessions, it makes sense for housing to play a prominent role in the conduct of monetary policy. A modified Taylor Rule would depend on a long-term measure of inflation having

little to do with the phase in the cycle, and, in place of Taylor's output gap, housing starts and the change in housing starts, which together form the best forward-looking indicator of the cycle of which I am aware. This would create pre-emptive anti-inflation policy in the middle of the expansions when housing is not so sensitive to interest rates, making it less likely that anti-inflation policies would be needed near the ends of expansions when housing is very interest rate sensitive, thus making our recessions less frequent and/or less severe.

Leamer (2007) is thus arguing that housing is an important tool for the transmission of monetary policy.

Leamer's (2007) regressions were similar to Coulson and Kim's (2000), except that he did not use time-series techniques (i.e., he did not use first differences to eliminate non-stationarity). Nevertheless, his result was similar to those papers that preceded his—residential investment with a one-quarter lag was the most reliable predictor of the business cycle. The t-statistic on lagged residential investment was far larger than that of any other explanatory variable for predicting GDP. Davis and Heathcote (2005) and Leung (2003) also found housing to be a leading indicator.

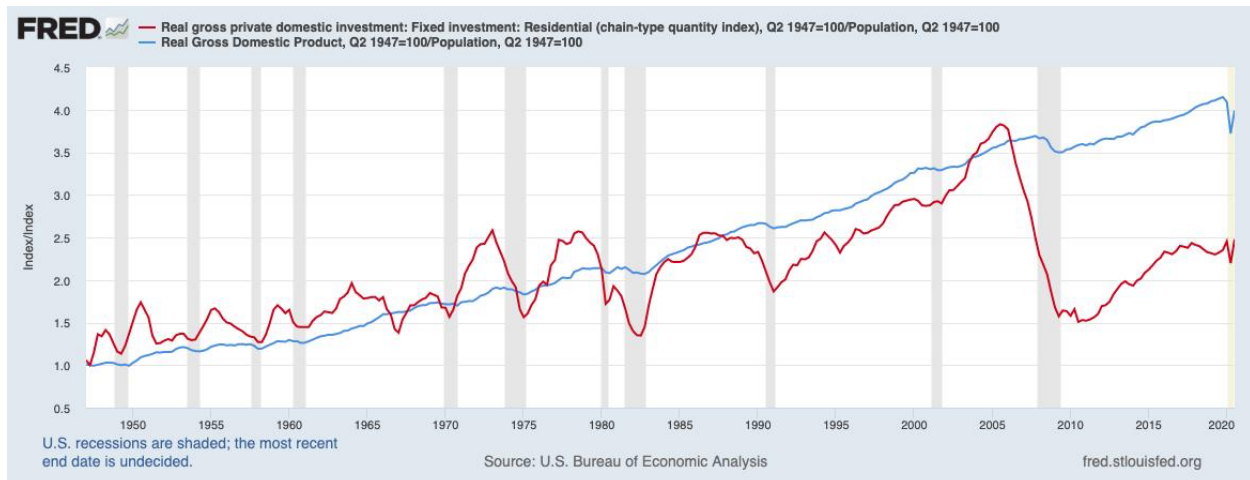
If we look at what has happened since 2007, however, we see a mixed bag. On the one hand, housing absolutely led the economy into a recession that year—residential investment declined by an extraordinary 57 percent between the first quarter of 2006 and the second quarter of 2009, and the contraction that followed that collapse was, by any measure, the worst since the great recession. But then, in the following years, while housing recovered, it did so slowly, and has not recovered its share of GDP, which in the first quarter of 2021 is 9.1 percent lower than it was in the fourth quarter of 2006. The fact that residential investment never recovered, while the broader economy arguably reached full employment before the advent of COVID, suggests that perhaps the link between residential investment and the business cycle has broken. This chapter investigates that possibility.

III. Descriptive data.

We begin with a depiction of United States Gross Domestic Product and Gross Fixed Residential Investment relative to the 1946 beginning of the National Income and Products Account series of the US Commerce Department [Figure 1]. The figure is quite striking in at least three dimensions: (1) it is clear that housing investment is more volatile than the broader economy; (2) up until the Great Financial Crisis, housing investment has a very similar long-

term trend to GDP and; (3) from 2007 on, housing investment dropped relative to GDP, and has remained relatively depressed since.

Figure 1
Residential Investment and GDP relative to 1947.

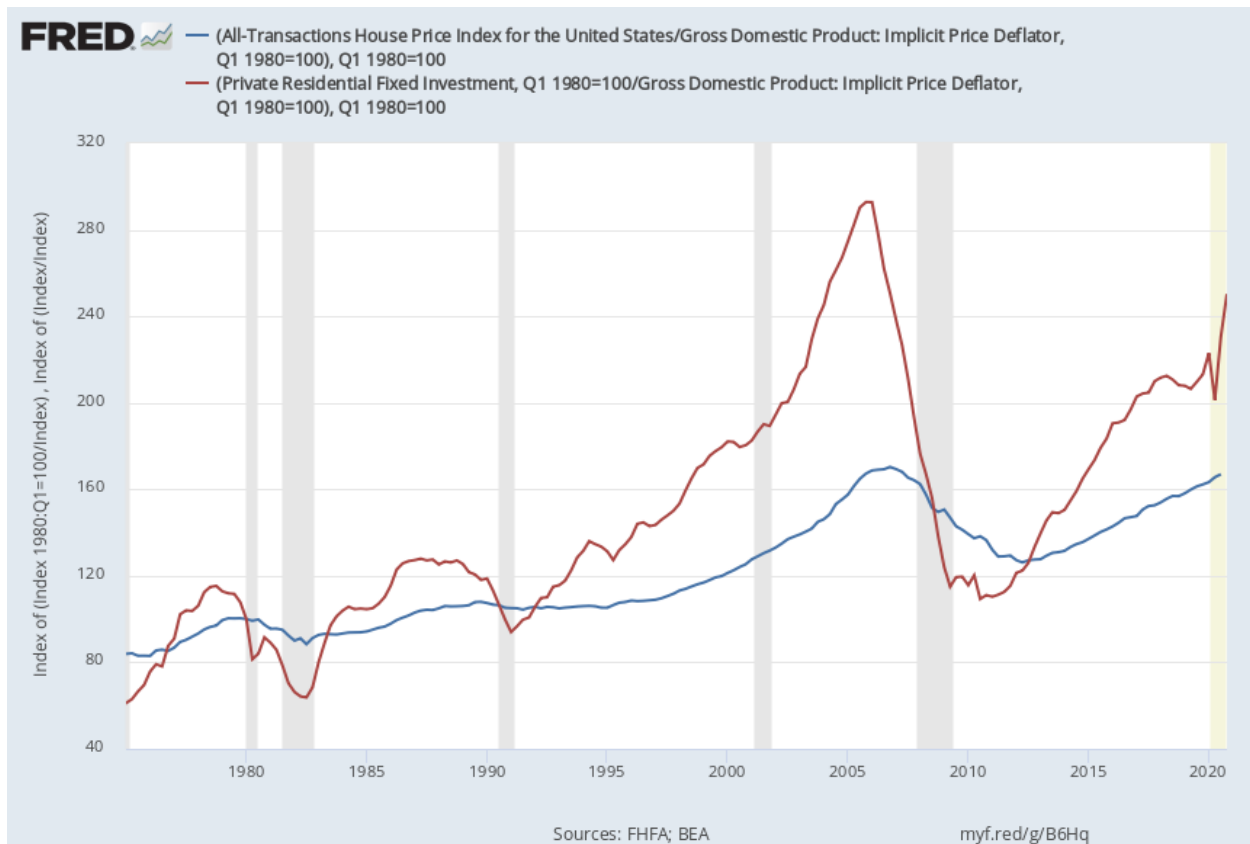


The recovery from the GFC is often characterized as slow: GDP growth in the aftermath of the recession that ended in 2009 was slow relative to growth following previous recessions, although it is worth noting that the expansion of the 2010s was notable for its length and for its ultimate ability to drive unemployment down to sub-four percent levels before the COVID induced recession. That said, the labor force participation rate in the US never recovered to its pre-GFC levels.

One wonders if the slowness of the recovery can be traced to the sluggishness of housing investment. It is also worth noting that the lack of investment in housing comes in the face of two forces that should have encouraged it: rising house prices and low real interest rates.

Figure 2 depicts the real Federal Housing Finance Authority Repeat Sales index, deflected by the GDP deflator, and real gross fixed residential investment from 1975 (the first year of the FHFA series) forward. Note that while in 2020, real house prices had returned to their near highpoint of 2006, residential investment *relative to GDP* remained 9.1 percent below its long-term average—an average that had been pulled down since 2008 (Figure 1).

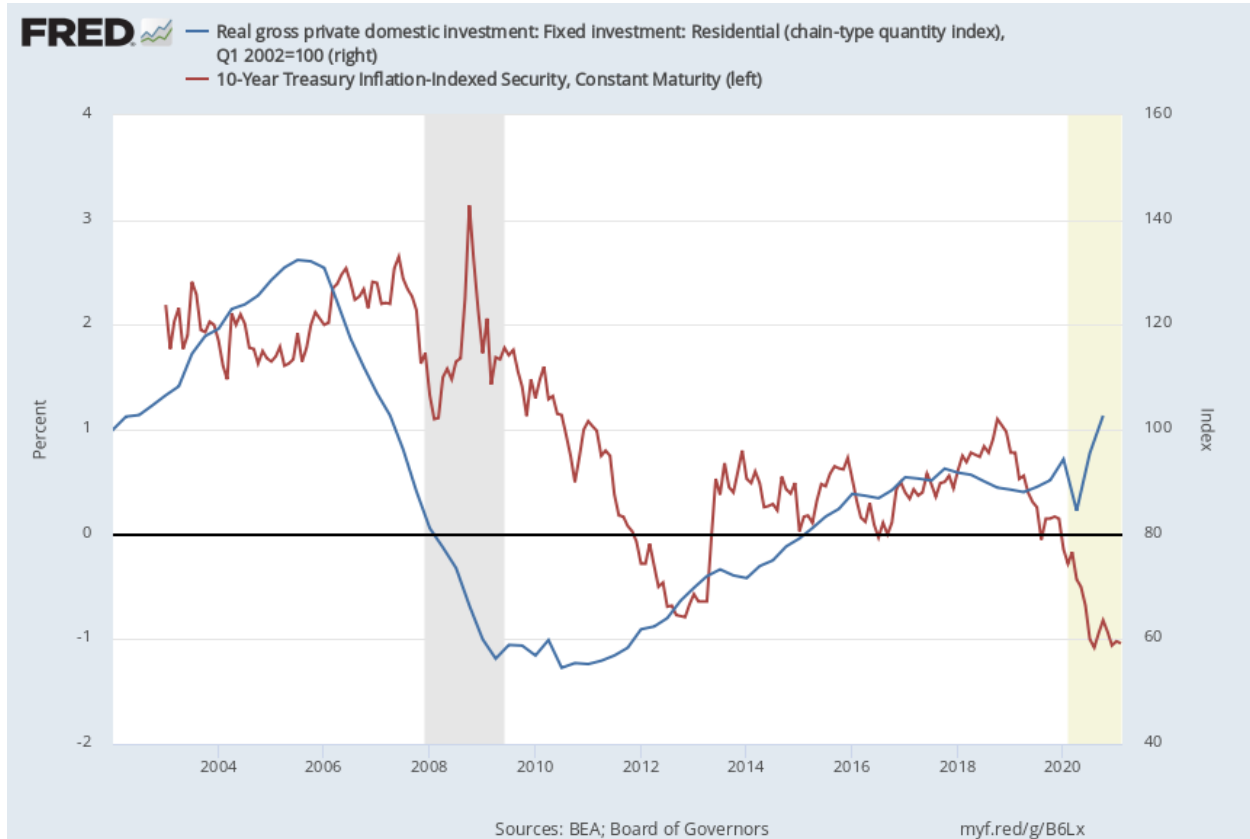
Figure 2
Real House Prices and Residential Investment (1975-2020)



At the same time, as measured by the Ten-Year Treasury Inflation-Protected Rate, real interest rates were negative in 2020, yet, again, housing investment was below trend (Figure 3).

Indeed, the correlation between TIPS and residential investment from 2002-2020 was 0.41. This is perhaps not entirely surprising, as high interest rates are associated with strong economic performance. But suppose we lag TIP rates by two quarters relative to residential investment. In that case, we still get a positive correlation of 0.31, meaning that there is no direct evidence that rising interest rates cool off housing investment, nor that falling interest rates stimulate it. As it turns out, however, the TIP rate is surprisingly not a stationary series, meaning that an appropriate model specification requires differencing TIPS in a vector autoregressive model. We will later show the results of this specification.

Figure 3
Real Interest Rates and Residential Investment



This paper is not about why the expected relationships between prices and investment or real estate rates and investment seem to have broken down—although there is a substantial literature on how availability (as opposed to the cost) of credit affected homebuilding, as well as literature on how land use regulation has prevented homebuilding. But it is striking how housing investment has remained at a very low level by historical standards for more than a decade.

Given past results, it is possible that the sluggish performance of residential investment kept the economy from reaching its potential output until just before the COVID-induced recession. But it is also possible that a once stable relationship between housing investment and GDP has ceased to be stable. The remainder of this paper will test whether the coefficients reported in

Green (1998), Coulson and Kim (2000), and Leamer (2007) have remained undisturbed or whether they have changed.

IV. Pretesting data

Hamilton's (1994) classic text discusses the treatment of time series data necessary to make sure that regression results rooted in such data are not spurious. It involves testing for stationarity of each individual data series and the cointegration of the various dependent variables. If the data series are individually not stationary, and the dependent variables are not cointegrated, the times series regressions must be in first differences not to be spurious.

Following Coulson and Kim (2000), we are going to investigate the relationships among GDP, residential investment, non-residential investment, consumption, and government spending. Table 1 presents the Dickey-Fuller tests for unit root for each of these five series. In all four cases, the series exhibit unit-root, indicating non-stationarity.

Table 1
Dickey-Fuller tests for unit-root with no lags

Variable	Test Statistic	5% Critical Value
GDP	-2.72	-2.88
Residential Investment	-2.07	-2.88
Non-residential Investment	-0.49	-2.88
Consumption	-2.23	-2.88
Government Spending	-5.01	-2.88

N=292. All variables are in natural logs.

We then turn to tests of cointegration. Table 2 presents Engle-Granger tests of cointegration and finds that residential investment and GDP are not cointegrated, but it is not possible to reject the null hypothesis that non-residential investment, consumption, and government spending are not cointegrated with GDP. Nevertheless, the focus of this chapter is on residential investment, and so it will not contain estimates of error correction models.

Table 2
Tests of Cointegration Between GDP and

Variable	Test Statistic	5% Critical Value
Residential Investment	-1.85	-2.88
Non-residential Investment	-3.22	-2.88
Consumption	-3.65	-2.88
Government Spending	-3.20	-2.88

N=292. All variables are in natural logs.

We, therefore, will do all our specifications in the first differences in natural logs of each series—we use natural logs so that we may interpret the first differences as growth rates.

V. Simple Granger Tests

We begin by updating the results from Green (1998), which tests for Granger causality among GDP, residential investment, and non-residential investment. Table 3 column 1 presents Wald Chi-square tests for rejecting the null of Granger Causality in a bivariate setting using two lags (we tried other lags as well, and the results were the same). The period for this column is the entire post-World War II period for the United States. Because the series themselves have unit roots, we estimate the model in natural log differences.

As in the earlier paper, residential investment Granger Causes GDP, and non-residential does not. GDP also continues to Granger Cause non-residential investment, but the reverse is not true.

We test for the stability of these results using the Brown, Durban, and Evans (1975) CUSUM test. This is a recursive regression test that may identify breaks in time series regimes. The idea of a CUSUM test is that in the event that the coefficients in a time series are stable, the sum of residuals as one moves forward through time should have an expected value of zero.

To test this proposition, Brown, Durban, and Evans (1975) show that the estimate

$$W_r = \frac{1}{\hat{\sigma}} \sum_{k+1}^r w_j$$

Table 3
Simple Granger Tests of Residential Investment, Nonresidential Investment, and GDP

	(1) Δgdp	(2) Δres	(3) $\Delta nonres$
Δgdp		4.17	2.62
Δres	30.7***		23.4***
$\Delta nonres$	1.13	9.49***	
N	292	292	292

p -values in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

where $w_r = \frac{y_r - x_r' b_{r-1}}{\sqrt{(1 + x_r'(X_{r-1}' X_{r-1})^{-1} x_r)^2}}$, Y_r is a time series of a dependent variable running from time

1 through r , X_r is a time series of a matrix of explanatory dependent variables running from time 1 through r , and B_r is a vector of estimated coefficients derived from the regression running from time t through r , is distributed Chi-squared. The null hypothesis of coefficient stability implies that the sum of the recursive residuals stays within 95 percent confidence bands, which widen as the time series lengthens.

Surprisingly, The CUSUM test does *not* reveal a structural break (Figure 4). Ploberger and Kramer (1992) note that the CUSUM recursive residual test often fails to pick up structural breaks toward the end of a time series and recommend the use of OLS residuals, which should have an expected value of zero throughout the time series period. But when we run CUSUM

tests based on OLS residuals, we still fail to find a structural break, even using 90 percent confidence bands.

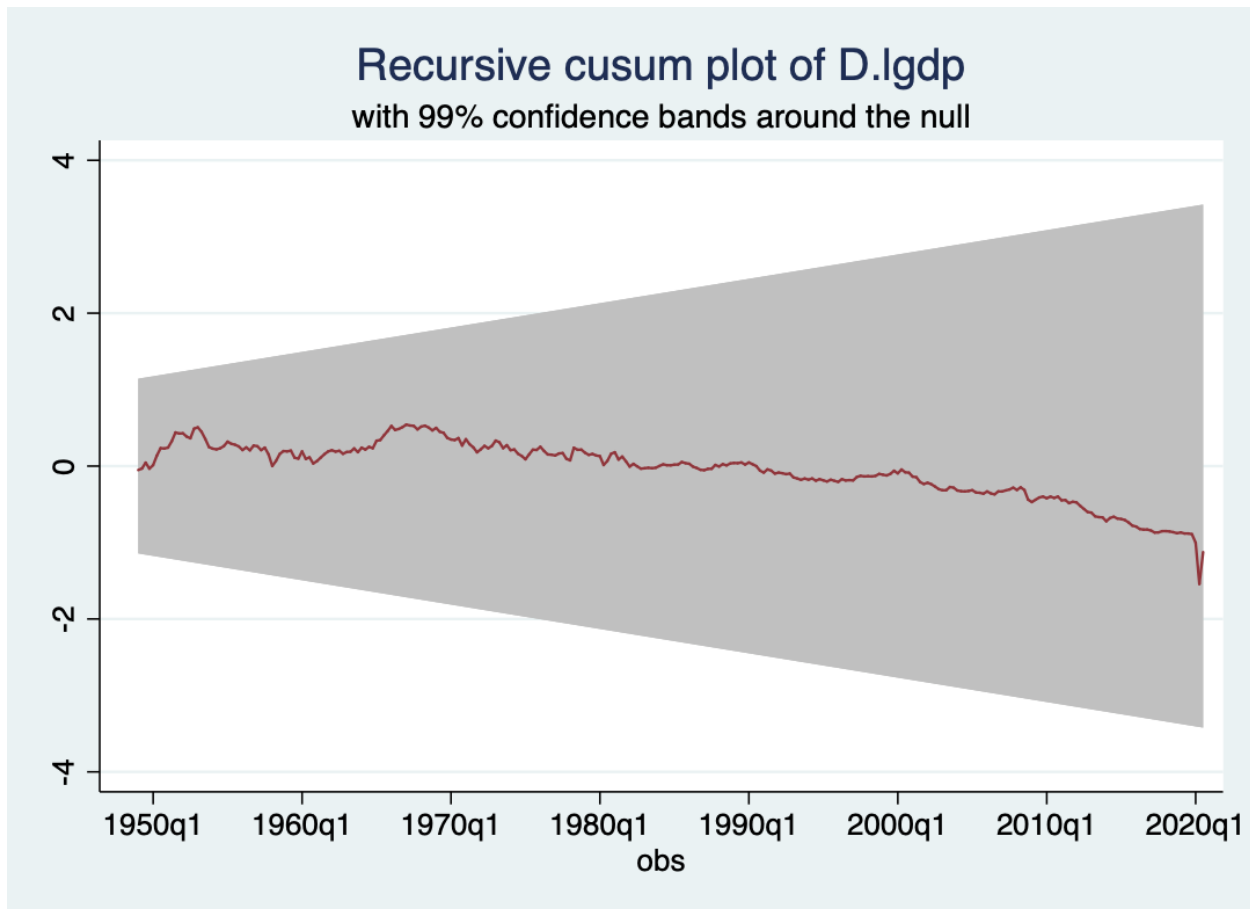
Yet Figure 1 reveals something very strange that happened at the time of the Great Financial Crisis—housing's share of GDP dropped dramatically and never since recovered. We thus turn to another test of structural stability—the Chow Test. This test has us break a combined regression into two pieces, and test whether the coefficients remain stable from one piece to the next. The Chow Test statistic is

$$\frac{\frac{SSR_c - SSR_1 - SSR_2}{k}}{\frac{SSR_1 + SSR_2}{N - 2k}}$$

We thus run regressions on two sets of truncated data. Given the lack of guidance from the CUSUM test, we must choose a year to determine whether there was a structural break. We look at two potential years: 2002 and 2008. In the year 2002, mortgage lending changed dramatically, with the private label market and exotic mortgages swelling (Green and Wachter 2007). And 2008 was the most dramatic year of the Great Financial Crisis.

Figure 4

Coefficient Stability Test of Granger Equation on Residential Investment "Causing" GDP



Below are the results of simple VARs for the total sample period (1947-2020), the first 56 years of that period (1947-2002), and the 18 years since.

Table 4
Chow Test of Granger Test between Residential Investment and GDP

	(1) Δgdp	(2) Δgdp	(3) Δgdp
Δgdp_{-1}	-0.0822 (0.239)	0.0800 (0.262)	-0.220 (0.208)
Δgdp_{-2}	0.0750 (0.362)	0.113 (0.104)	-0.258 (0.481)
Δres_{-1}	0.0562*** (0.001)	0.0760*** (0.000)	-0.0466 (0.426)
Δres_{-2}	0.0461** (0.009)	0.0182 (0.247)	0.109 (0.071)
_cons	0.00705*** (0.000)	0.00624*** (0.000)	0.00644** (0.007)
N	292	217	75
rss	0.0346	0.0161	0.0148

p -values in parentheses


* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Unlike the CUSUM test, the Chow test reveals a structural break must happen at some point. The test statistic here is $\frac{(0.0346 - 0.0161 - 0.0148)/4}{(0.0161 + 0.0148)} = 8.5$. The test statistic is distributed $F(4, 284)$ and has $p < .00001$. Testing at 2008 also reveals a structural break, but we use a breakpoint of 2002 for further analysis so as to have more observations to test the impact of residential investment on GDP in more recent years.

When we limit the data to 2002 forward, we find very different results—there is no evidence that residential investment Granger Causes GDP. There is more to this finding than the power of the test—we may use the vector autoregressions that were the foundation of the Granger Cause tests to plot impulse response functions. Figure 5 is an impulse response function (IRF) based on estimates from the full sample period, and Figure 6 is an IRF based on estimates from 2002 forward.

Figure 5

Impulse Response Functions based on regressions from 1947-2020 (top) and 2002-2020(bottom)



The contrast between the two is striking. It is clear that a one standard deviation shock in residential investment growth has a meaningful impact on GDO growth—about 30 basis points after two quarters. Given that average GDP growth is about 70 basis points a quarter, this is

material. Note, also, that zero is well outside the confidence interval of the effect of residential investment on GDP.

Turning to the IRF based on the past 18 years, though, gives us a different picture. It is not only the case that zero is well within the confidence band (which might simply indicate the weaker power of an estimate based on 18 years of data, the predicted impact of residential investment on GDP turns negative.

VI. Granger Tests on components of GDP

We now turn to how residential investment influences or is influenced by the three largest non-investment components of GDP: consumption, government spending, and non-residential investment.

It is clear the channel through which residential investment moved GDP—it was through consumption, and through itself in later years (very likely reflecting the long-held view that new construction leads to demand for consumer durables).

Table 5
Granger tests of housing on various components of GDP

	(1) Δ consumption	(2) Δ government spending	(3) Δ residential investment	(4) Δ non-residential investment
Δc_{-1}	-0.124 (-1.61)	-0.0266 (-0.35)	0.306 (1.37)	0.0919 (0.78)
Δc_{-2}	0.183* (2.41)	0.00310 (0.04)	0.364 (1.65)	0.213 (1.84)
Δrinv_{-1}	0.0430 (1.89)	-0.0360 (-1.62)	0.405*** (6.15)	0.0876* (2.53)
Δrinv_{-2}	0.0621** (2.80)	0.0344 (1.59)	0.111 (1.73)	0.0478 (1.41)
Δnrinv_{-1}	-0.0695 (-1.58)	0.0613 (1.43)	-0.508*** (-3.98)	0.262*** (3.91)
Δnrinv_{-2}	-0.0577 (-1.34)	-0.00902 (-0.21)	-0.344** (-2.75)	0.0346 (0.53)

Δg_{-1}	0.0534 (0.88)	0.474*** (8.04)	-0.244 (-1.39)	-0.122 (-1.32)
Δg_{-2}	-0.0795 (-1.30)	0.167** (2.82)	0.0286 (0.16)	0.0494 (0.53)
_cons	0.00869*** (7.35)	0.00190 (1.65)	0.00734* (2.14)	0.00467* (2.59)
N	292	292	292	292

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

As before, we now limit ourselves to observations beginning with the first quarter of 2002. And again, the results change.

Table 6
Granger tests of housing on various components of GDP

	(1) Δ consumption	(2) Δ government spending	(3) Δ residential investment	(4) Δ nonresidential investment
Δc_{-1}	0.191 (0.79)	0.0630 (0.65)	0.985 (1.49)	0.675* (2.01)
Δc_{-2}	0.493 (1.87)	-0.146 (-1.39)	1.366 (1.91)	0.749* (2.06)
$\Delta rinv_{-1}$	-0.145* (-2.40)	-0.00107 (-0.04)	-0.0133 (-0.08)	-0.248** (-2.96)
$\Delta rinv_{-2}$	0.165** (2.75)	-0.0158 (-0.66)	0.409* (2.50)	0.102 (1.22)
$\Delta nrinv_{-1}$	-0.225 (-1.63)	-0.00275 (-0.05)	-0.733 (-1.95)	0.365 (1.91)
$\Delta nrinv_{-2}$	-0.120 (-0.90)	0.0699 (1.31)	-0.483 (-1.33)	-0.356 (-1.93)
Δg_{-1}	0.0433 (0.12)	0.260 (1.86)	-1.166 (-1.22)	-0.513 (-1.05)
Δg_{-2}	-0.645 (-1.84)	0.474** (3.40)	-2.040* (-2.14)	-1.161* (-2.40)

_cons	0.00663*	0.000196	-0.00130	0.00176
	(2.59)	(0.19)	(-0.19)	(0.50)
<hr/> N	<hr/> 75	<hr/> 75	<hr/> 75	<hr/> 75

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Note how the strong relationship between residential investment and consumption becomes attenuated—the sum of the coefficients of the impact of residential investment on consumption drops from 0.105 to 0.02. The impact of residential investment on non-residential investment turns from positive to negative.

VII. Has housing ceased to be an instrument of monetary policy?

We return to Leamer's (2007) discussion of how housing should be an element of the Taylor Rule. Housing investment is easy to measure, and once foretold recessions and recoveries. Under such circumstances, one could imagine a monetary policy rule where interest rates adjusted in response to changes in residential investment.

Yet, we have already seen that the business cycle is not as sensitive (if it is still sensitive at all) to residential investment. This finding undermines the efficacy of housing as an instrument of monetary policy. And under the circumstances, it is worth asking whether the Federal Reserve even still has an influence on residential investment.

We turn to three measures of interest rates and examine their effects on residential investment: the Ten-year Treasury Rate, the Ten Year Treasury Inflation-Protected (TIP) rate, and the Federal Funds rate. We do Granger tests using the full set of data from all three series and then using data from January 2003 on for the Ten-year and Federal Funds rates so that their data length matches those of TIPs, which have only been available since that date. The regressions are in first differences so that the time series are stationary.

We find that long-term interest rates, whether protected for inflation or not, predict residential investment, whether we are looking at long time series or the truncated series (Table 7). This is also true of Freddie Mac 30-year fixed rate mortgage rates. But the Federal Funds Rate has ceased to be a predictor of residential investment activity. Note that the sum of coefficients of the effect of the Federal Funds Rate on residential investment was -0.027 for the entire post-1947 period but was only -0.002 for the Post-2003 period. While the Ten-year

treasury continued to have an impact on housing, it has operated with a longer lag in the post-2003 period than over the entire post-1947 period.

The Federal Reserve's policy of Quantitative Easing (QE) included the purchase of long-Treasury bonds. Because the Federal Reserve acted directly on long-term interest rates, it is possible that short-term interest rates ceased to have an impact on housing. Housing's durability implies long duration, and so its activity should be more influence by long-term rates. Moreover, mortgage rates are generally correlated with 10-year Treasury rates (the correlation is .99 over the longer term, although it is slightly lower, at .98, over the post-2013 period). And indeed, when we look at the impact of mortgage rates themselves on residential investment, they are slightly stronger than the impact of the 10-year Treasury notes. Krishnamurthy and Vissing-Jorgensen (2011) found that QE had a strong impact on both 10-year notes and Agency Mortgage-Backed Securities rates, and so we have some evidence that the channel through which the Federal Reserve stimulated housing shifted from the traditional Federal Funds rate to longer-term securities. Nevertheless, while monetary policy once had a nearly immediate impact on residential investment, in recent periods, it operates with a lag.

Table 7
The Impact of Interest Rates on Residential Investment

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	$\Delta.lrin v$ 54-20	$\Delta.lrin v$ 03-20	$\Delta.lrin v$ 47-20	$\Delta.lrin v$ 03-20	$\Delta.lrin v$ 03-20	$\Delta.lrin v$ 03-20	$\Delta.lrin v$ 71-20
$\Delta lrin v_{-1}$	0.324*** (4.82)	0.0538 (0.46)	0.349*** (5.61)	0.084 (0.67)	0.138 (0.99)	0.038 (0.33)	0.226** (3.09)
$\Delta lrin v_{-2}$	0.191** (2.98)	0.595*** (4.91)	0.253*** (4.09)	0.631*** (4.90)	0.502*** (3.51)	0.696*** (5.69)	0.248*** (3.66)
$\Delta TENYR_{-1}$	-0.029*** (-5.45)			0.013 (0.96)			
$\Delta TENYR_{-2}$	-0.022*** (-3.76)			-0.051*** (-3.83)			
$\Delta TIPS_{-1}$		-0.025 (-1.75)					
$\Delta TIPS_{-2}$		-0.057*** (-3.96)					
$\Delta FedFunds_{-1}$			-0.017*** (-6.29)		-0.013 (-0.75)		
$\Delta FedFunds_{-2}$			-0.010*** (-3.39)		0.010 (0.59)		
$\Delta mortgage\ rates_{-1}$						-0.006 (-0.41)	-0.038*** (-6.96)
$\Delta mortgage\ rates_{-2}$						-0.074*** (-4.96)	-0.018** (-2.95)
_cons	0.00159 (0.66)	-0.00411 (-1.01)	0.00157 (0.70)	-0.00259 (-0.62)	-0.00150 (-0.33)	-0.00401 (-1.02)	0.000302 (0.12)
N	232	68	262	68	68	68	195

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

VIII. Why does housing matter less?

A long literature has developed arguing that land use regulation has long had an impact on the production of housing and that regulation has gotten tighter over the years (See Gyourko, Saiz and Summers 2008, Davidoff 2008). Morrow (2013) gives a rather spectacular example of this: the city of Los Angeles, through a series of downzonings, has reduced its zoning capacity from 10 million to slightly more than 4 million people. In city after city, building permits per capita have been falling over time. California's 2020 population loss—the first since it became a state in 1850—was partially the result of the fact that it is not building much housing—barely enough to make up for losses arising from the demolition of old houses. From 2018 to 2019, California only added 10,000 net housing units on a base of 14 million.

The frictions in housing construction are also reflected in the findings here about how interest rates influence residential investment. Rates have become less influential and take more time to exert any influence at all on residential investment. As permitting processes have become more arduous and more stretched out, the ability of monetary policy to influence the business cycle through housing has become attenuated.

At the same time, as the regulatory environment has become more stringent, housing has to some extent become a latent variable. Heckman's series of papers on latent variables may well apply to housing—regulation censors housing much more than it used to, meaning that it may appear to have a smaller influence on GDP than it might if it were permitted to respond more quickly and sensitively to countercyclical policy.

Consider a world where housing limits are so severe that the maximum construction permitted is less than would be constructed at a place in the business cycle where housing demand is at its lowest. The California experience is not far removed from that—note again that it is currently not expanding its housing stock in a period of very low interest rates and very high prices. Under these circumstances, it is impossible to identify the effect of housing on the business cycle because we cannot observe normal fluctuations in the housing market.

Hsieh and Moretti (2019) have shown that zoning as practiced in the US undermines the long-run competitive position of the US economy by creating artificial shortages of housing in places people wish to live. The severing of housing from the business cycle over the past 18 years or so suggests that zoning may also have neutered the ability of the Federal Reserve to do

countercyclical monetary policy.

References

Aura, Saku, and Thomas Davidoff. "Supply constraints and housing prices." *Economics Letters* 99.2 (2008): 275-277.

Brown, Robert L., James Durbin, and James M. Evans. "Techniques for testing the constancy of regression relationships over time." *Journal of the Royal Statistical Society: Series B (Methodological)* 37.2 (1975): 149-163.

Coulson, N. Edward, and Myeong-Soo Kim. "Residential investment, non-residential investment, and GDP." *Real Estate Economics* 28.2 (2000): 233-247.

Davis, Morris A., and Jonathan Heathcote. "Housing and the business cycle." *International Economic Review* 46.3 (2005): 751-784.

Granger, Clive WJ. "Testing for causality: a personal viewpoint." *Journal of Economic Dynamics and control* 2 (1980): 329-352.

Green, Richard K. "Follow the leader: how changes in residential and non-residential investment predict changes in GDP." *Real estate economics* 25.2 (1997): 253-270.

Green, Richard K., and Susan M. Wachter. "The American mortgage in historical and international context." *Journal of Economic Perspectives* 19.4 (2005): 93-114.

Gyourko, Joseph, Albert Saiz, and Anita Summers. "A new measure of the local regulatory environment for housing markets: The Wharton Residential Land Use Regulatory Index." *Urban Studies* 45.3 (2008): 693-729.

Hamilton, James Douglas. *Time series analysis*. Princeton university press, 1994.

Hsieh, Chang-Tai, and Enrico Moretti. "Housing constraints and spatial misallocation." *American Economic Journal: Macroeconomics* 11.2 (2019): 1-39

Krishnamurthy, Arvind, and Annette Vissing-Jorgensen. The effects of quantitative easing on interest rates: channels and implications for policy. No. w17555. National Bureau of Economic Research, 2011.

Leamer, Edward E. Housing is the business cycle. No. w13428. National Bureau of Economic Research, 2007.

Leung, Charles. "Macroeconomics and housing: a review of the literature." *Journal of Housing Economics* 13.4 (2004): 249-267.

Morrow, Greg. The Homeowner Revolution: Democracy, Land Use and the Los Angeles Slow-Growth Movement, 1965–1992. University of California, Los Angeles, 2013.

Ploberger, Werner, and Walter Krämer. "The CUSUM test with OLS residuals." *Econometrica: Journal of the Econometric Society* (1992): 271-285.

