

BOARD OF GOVERNORS OF THE FEDERAL RESERVE SYSTEM
DIVISION OF RESEARCH AND STATISTICS

Date: January 8, 2016
To: David Wilcox
From: Raven Molloy
Subject: Measuring Housing Overvaluation using the Zillow Price-to-Rent Ratio

Summary

Zillow computes an aggregate price-to-rent ratio by estimating the rent and market value of each housing unit in the United States, dividing the estimated value by the estimated rent, and then taking the median of these estimates across all housing units. An important advantage of this approach is that it compares rent and value for the same housing unit, so that the ratio is not affected by quality differences between the owner-occupied stock and the rental stock. By contrast, the staff's current method to assess overvaluation compares the price of owner-occupied housing to rents of rental units, and there is a long-run upward trend in this ratio that we partly attribute to differential trends in quality (see Steve Laufer's memo of Dec. 8, 2015).

One disadvantage of the Zillow price-to-rent ratio is that it only begins in October 2010, which makes it difficult to say whether housing is overvalued or undervalued. Thus, until now the staff has not paid much attention to the Zillow price-to-rent ratio. This memo describes a procedure for calculating a historical average for the Zillow price-to-rent ratio based on Census data. The results suggest that the current price-to-rent ratio is closely in line with its long-term average, a conclusion that is similar that drawn from the staff's current overvaluation model. The fact that these two measures give similar signals about overvaluation in the housing market is reassuring because they are derived from very different data sources and methodologies.

Estimating a Historical Average

Zillow estimates the value and rent for each housing unit in the U.S. based on the unit's characteristics and location. To create a historical average for the Zillow data, I follow a similar

modeling strategy using property-level data from the 1960 to 2000 Censuses and the 2011 American Community Survey (ACS). In particular, to impute a rental value and market value for each single-family housing unit in these samples, I regress the contract rent and house value on a set of housing unit characteristics and use the estimated coefficients to predict rent and house value for each single-family unit in the sample (regardless of whether it is owner-occupied or rented).¹ Next I calculate the ratio of estimated price to estimated rent for each housing unit, and take the median across housing units. I use the same model to estimate prices and rents in each year so that changes in the median price-rent ratio do not reflect differences in model specification. Because so few housing unit characteristics available for all Census years between 1960 and 2000, the model is quite simple. Specifically, it includes indicators for: the number of bedrooms, the number of rooms, age of the housing unit, whether the unit is attached to another unit, and the state where the unit is located. This methodology is similar that used in Davis, Lehnert and Martin (“The Rent-Price Ratio for the Aggregate Stock of Owner-Occupied Housing” *Review of Income and Wealth* 2008).²

Although the method described above is designed to mimic Zillow’s methodology, the resulting price-to-rent ratio could be quite different because Zillow’s price and rent models are based on a much richer set of housing unit characteristics and a more complex modeling framework.³ In an attempt to account for these differences, I estimate the median price-to-rent ratio using data from the 2011 ACS, and multiply the price-to-rent ratios estimated in the Census data by the ratio of the Zillow estimate for 2011 and to the 2011 ACS estimate. This normalization is done using 2011 data because 2011 is the first complete year of Zillow data.

Figure 1 shows the median price-to-rent ratios from the Census, the ACS, and Zillow. As mentioned above, the ACS and Zillow averages for 2011 are equal by construction. The fact that the ACS and Zillow estimates for 2012 and 2013 line up fairly well lends credibility to the assumption that changes in the Census price-rent ratio are similar to Zillow, even though the


¹ For the rent regressions, I limit the sample to units where contract rent is less than gross rent (92 percent of rental units from 1960 to 2000 satisfy this condition). Thus, the imputed rental value does not include the cost of utilities. This restriction avoids the need to estimate utility costs for units where this information is not available.

² One important difference is that I limit the sample to single-family homes. In addition, I estimate the value as well as estimating rent, whereas they used house values directly.

³ Moreover, the house values reported in the Census and ACS are owner-occupants’ estimates, whereas Zillow house values are derived from transaction prices.

models used to estimate these ratios are quite different. On the other hand, the Zillow estimate for 2014 was somewhat higher than the 2014 ACS estimate.

Figure 1: Median Price-to-Rent Ratios in Census, ACS, and Zillow



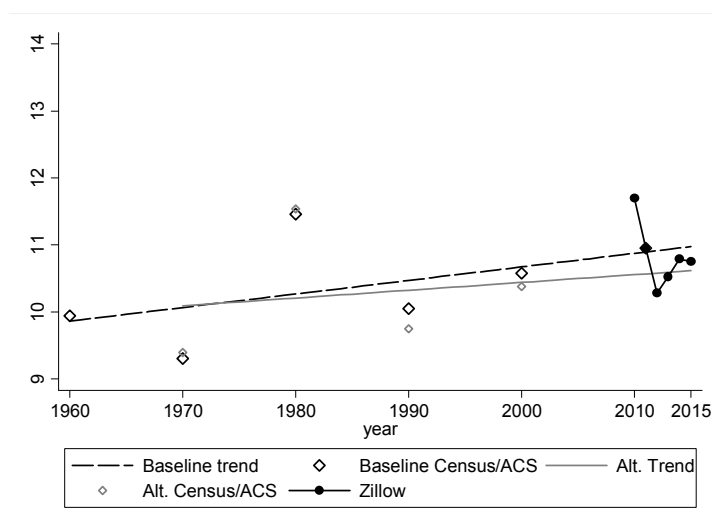
There appears to be a slight trend in the Census estimates from 1960 to 2000, as shown by the dashed line.⁴ Comparing the ACS and Zillow estimates to this trend results in a conclusion about overvaluation that is similar to the staff’s current view: the price-to-rent ratio far exceeded historical norms during the mid-2000s, fell to a level somewhat below historical norms during the financial crisis, and has returned to a level roughly in line with historical norms in recent years. Even if one were to ignore the trend, the current Zillow estimates would still be only about 4 to 5 percent above the simple average of the Census data from 1960 to 2000.

To investigate the robustness of the Census price-to-rent estimates, I calculate a different historical average that is based on price and rent regressions estimated separately by metropolitan area. By allowing all coefficients to vary across locations, this specification is richer than the baseline model. This specification also limits the sample to a balanced panel of 58 metropolitan areas, so changes in the price-to-rent ratio over time are not affected by long-run trends in where housing units tend to be located. However, one drawback of this specification is

⁴ I do not have a good explanation for this trend. We attribute the upward trend in the overvaluation model to trend differences in the quality of owned versus rental housing. The real interest rate was roughly the same in 1960 and 2000, suggesting that the upward trend is not related to borrowing costs. In any case, the trend in the price-to-rent ratio is very slight—it increases by 0.2 percent per year, or 8 percent over a 40-year period.

that metropolitan area indicators are not available in the 1960 Census, so this analysis only uses Census data from 1970 to 2000. Figure 2 shows that the trend estimated from the alternate regression is somewhat less steep than the baseline trend, and therefore the current level of the price-to-rent ratio looks slightly higher than the trend, rather than slightly lower.⁵ Thus, as with the staff’s overvaluation model, the results of this model should be interpreted with caution because they are somewhat sensitive to specification. Nevertheless, the results are consistent with the staff view that the current price-to-rent ratio is fairly close to historical norms.

Figure 2: Alternative Trend in Price-to-Rent using Metropolitan Area-Specific Regressions



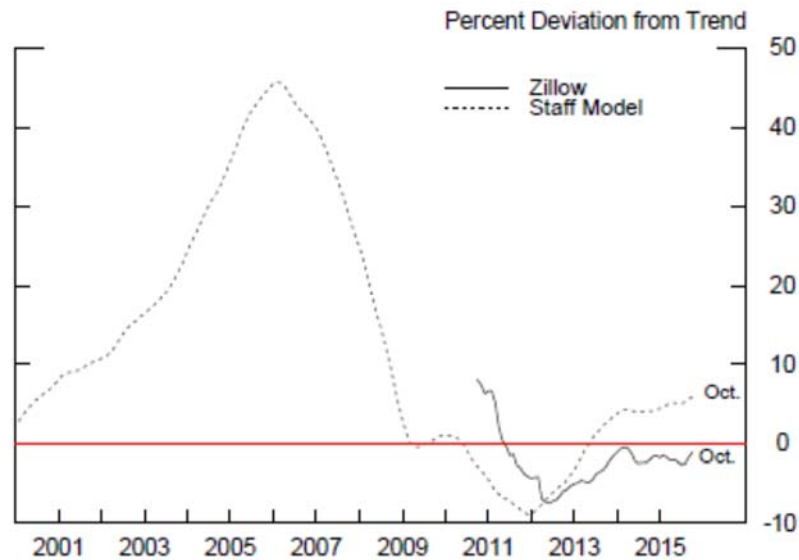
Comparing the Zillow Price-to-Rent Ratio with the Staff’s Overvaluation Model

Figure 3 shows the implication of the Zillow price-to-rent ratio for an assessment of overvaluation in the housing market. Specifically, I divide the Zillow price-to-rent ratio by the trend shown in Figure 1 and take the logarithm of the result so that the deviation from trend can be interpreted in percentage terms. For comparison, the figure shows the analogous metric from the staff’s standard overvaluation model—i.e. the logarithm of the aggregate price-to-rent ratio relative to its estimated trend. The time series patterns of the two measures are similar.

⁵ This result is not due to the shorter sample period, as estimating the trend from 1970 to 2000 in the baseline specification leads to a slightly steeper trend.

Moreover, although the current Zillow estimate is somewhat lower than the staff model estimate, both are consistent with the price-to-rent ratio being near its historical norm.

Figure 3: Overvaluation based on Zillow Benchmark and Staff Model



Conclusion

The Zillow data provide a nice complement to the staff's current method of assessing overvaluation in the housing market, and we plan to follow both measures going forward. In the future we also plan to follow a similar strategy for estimating historical benchmarks for price-to-rent ratios for states and metropolitan areas, which would allow for a better assessment of how overvaluation varies across locations.