Monetary Policy and Rents *

Boaz Abramson

Columbia Business School

Pablo De Llanos

Columbia University

Lu Han

University of Wisconsin-Madison

February 12, 2025

Abstract

We estimate the effects of monetary policy on housing rents. To do so, we construct a new repeat-rent index that has a broader and more granular geographical coverage relative to existing alternatives. Using our rent index, we estimate the impulse responses of rents to monetary policy shocks by employing local projection methods. We find that monetary tightening increases both real and nominal rents. A 25 basis point increase in the 30-year fixed rate mortgage raises real (nominal) rents by 1.7 (1.4) percent 12-24 months following the monetary policy shock. The effect is driven by a shift in household demand from the owner-occupied market to the rental market. Our results highlight the distributional effects of monetary policy and suggest that contractionary monetary policy is limited in its ability to lower inflation.

JEL-Codes: E31, E52, G51, R21, R28, R31.

Keywords: monetary policy, rents, repeat-rent index.

*Abramson: ba2695@columbia.edu. De Llanos: pd2655@columbia.edu. Han: lu.han@wisc.edu. We thank Zhenyang Gong for excellent research assistance and Gadi Barlevy and participants at the Columbia-NYU-Yale Housing Day and Society for Economic Dynamics Winter Meeting for useful comments.

1 Introduction

How does monetary policy affect housing rents? On the one hand, a higher interest rate might lower demand for rental housing and as a result lower rents. Rental housing is a consumption good, and standard intertemporal substitution considerations suggest that a higher interest rate lowers demand for consumption by making savings more attractive (Campbell, 1989; Kaplan, Moll and Violante, 2018). On the other hand, a higher interest rate might shift demand from the owner-occupied market to the rental market, thereby increasing rents. When credit becomes more expensive, financially constrained households are less likely to become homeowners (Ringo, 2024), which crowds in the rental market (Gete and Reher, 2018). Overall, the effect of monetary policy on rents is therefore ex-ante ambiguous.

This paper studies the effects of monetary policy on rents. To do so, we proceed in two steps. First, we construct a new repeat-rent index using a national database of rental listings. Our repeat-rent index provides a quality-constant measure of rent growth that is representative of rent growth in the US. The main advantage of our index is that it has a broader and more granular geographical coverage relative to existing alternative rent indexes. Second, using our rent index, we estimate the impulse responses of rents to monetary policy shocks by employing the standard local projection methods (Jordà, 2005).

Our main finding is that contractionary monetary policy increases both real and nominal rents. A 25 basis point unexpected increase in the 30-year fixed rate mortgage raises real (nominal) rents by 1.7 (1.4) percent 12-24 months following the monetary policy shock. The result is driven by a shift in household demand from the owner-occupied market to the rental market. Using transaction level data, we show that when monetary policy tightens, households are less likely to become homeowners, which crowds in the rental market. We find that the increase in demand for rentals is accommodated by real-estate investors who capitalize on the higher rents by buying houses from owner-occupiers and renting them in out the rental market.

Our results have important policy implications. First, they highlight the distributional effects of monetary policy. Given that renters tend to be lower-income households, our findings suggest that contractionary monetary policy might disproportionately harm the most vulnerable households in the population. Second, our results suggest that contractionary monetary policy is limited in its ability to lower inflation. The finding that monetary tightening increases rent, which is the single largest component of the Consumer Price Index (CPI),[1] limits the extent to which it can lower inflation.

We begin by developing a new repeat-rent index, which we refer to as the ADH-RRI. Our main data source for this apparatus is listing data compiled by Altos Research between 2011 and 2022. Altos com-

[1]See www.bls.gov.

piles a national database of rental listings from online listing platforms and from Multiple Listings Services (MLS). Updated on a weekly basis, the data provides a snapshot of the listings that are observed every week. For each listing, the data records the listed monthly rent, the date in which the listing is observed, the property address, as well as physical characteristics of the listed unit such as the number of beds and baths, the floor size, the year built and the property type. We identify rental units in our data based on their address and physical characteristics.

We show that our listing data is representative of the geographical distribution of rental units in the US. However, it over-represents higher quality rental units within local rental markets. Namely, the median rent in our data is higher than the median rent in the nationally representative American Community Survey (AHS) or American Housing Survey (AHS) and the median rental unit in our data is larger. A potential concern with our repeat-rent index is therefore that it is based on a sample of rental units that is not representative of the US rental market. We address this concern by showing that our repeat-rent index is representative of rent *growth* in the US. In particular, we show that our ADH-RRI aligns remarkably well the CPI-NTRR, an alternative repeat-rent index that is based on nationally representative rental data and is representative of rent growth (Adams et al., 2024). Our index is therefore well-suited for studying the impact of monetary policy on rent inflation.

To construct our repeat-rent index, we follow the repeat-sales methodology of Bailey, Muth and Nourse (1963) and Case and Shiller (1989) and apply it to the rental market. The basic idea is to compare rents listed on the same unit across time to construct a quality-constant measure of rent growth. We construct and analyze various specifications of our index. We develop both a nominal ADH-RRI, which measures growth of nominal rents, and a real ADH-RRI, which measures growth of real rents. We construct our ADH-RRI both at the zip-code level and at the national level, and both at the monthly and quarterly frequency. Our ADH-RRI aligns well with alternative popular rent indexes - Zillow's ZORI index (Clark, 2022), the Marginal Rent Index (Ambrose, Coulson and Yoshida, 2023), and the CPI-NTRR and CPI-ATRR indexes (Adams et al., 2024). Our index, as well as those mentioned above, leads the CPI-Rent index. This is because it tracks rent growth faced by new tenants (as measured by listed rents) rather than rent growth faced by all occupants.

An important advantage of our ADH-RRI relative to these prominent alternatives is that it provides a broader and more granular geographical coverage of rental markets in the US. The ADH-RRI is computed at the zip-code level, while The CPI-NTRR and ACY-MRI are constructed only at the national level. The most granular geographical level for which the CPI-rent index is constructed is the CBSA. ZORI, which to the best of our knowledge is the most granular repeat-rent index to date, is also constructed at the zip-code level. However, its geographical and temporal coverage is limited compared to our ADH-RRI. First, our

ADH-RRI is available starting from 2011, while ZORI is only available starting from 2015. Second, in years where ZORI is available, it covers substantially less zip-codes than our ADH-RRI. The broader and more granular geographical coverage of our index substantially improves the statistical precision of our analysis. It also allows us to study the heterogeneous effects of monetary policy across local housing markets to shed light on the mechanisms that drive the empirical results.

We use the repeat-rent index to study the effects of monetary policy on rents. We estimate the dynamic effects of a monetary policy shock on rents using the standard local projection instrumental variable (LP-IV) framework (Jordà, 2005; Jordà, Schularick and Taylor, 2015; Ramey, 2016; Stock and Watson, 2018). Since our context is the housing markets, we use the 30-year fixed rate mortgage as our (instrumented) monetary policy indicator (Aastveit and Anundsen, 2022; Gorea, Kryvtsov and Kudlyak, 2022). In our baseline specification, we identify exogenous shocks to monetary policy using the Bauer and Swanson (2023*b*) monetary policy surprises series. Bauer and Swanson (2023*b*) construct their monetary policy shocks in two steps. First, they measure high-frequency changes in interest rates around Federal Open Market Committee (FOMC) meetings and around press conferences, speeches, and testimonies made by the Federal Reserve chair. Second, they orthogonalize these monetary policy surprises by regressing them on economic and financial variables that predate the FOMC meetings and Federal Reserve chair announcements, and take the residuals. By doing so, they capture the component of changes to interest rates around monetary policy events that is ex-ante unpredictable. We show that the Bauer and Swanson (2023*b*) monetary policy shock is a valid instrument for the 30-year fixed rate mortgage by establishing that it satisfies the (Stock and Watson, 2018) validity conditions.

Our main result is that contractionary (expansionary) monetary policy shocks increase (decrease) rents. A 25 basis point increase in the interest rate on a 30-year fixed rate mortgage leads to a 1% (1.7%) increase in *real* rents 12 (24) months following the monetary policy shock. This result has important implications in terms of the distributional effects of monetary policy. By making renting more expensive relative to all other prices in the economy, contractionary monetary policy disproportionately harms the most vulnerable households in the population, who tend to be renters. Higher interest rates not only increase real rents, but also *nominal* rents. In particular, a 25 basis point increase in the interest rate on a 30-year fixed rate mortgage leads to a 0.7% (1.4%) increase in nominal rents 12 (24) months following the monetary policy shock. This result has important implications in terms of the effect of monetary policy on inflation. The fact that monetary tightening increases rent, which is the single largest component of the Consumer Price Index (CPI), suggests that it is limited in the extent to which it can lower inflation.

We consider a host of robustness tests for our baseline results. First, The choice of monetary policy shocks can be important for the estimated effects of monetary policy (Ramey, 2016). We therefore replicate our empirical analysis for a host of alternative monetary policy shocks that have been used in the literature. We show that our results are robust to using monetary policy shocks from, for example, Gürkaynak, Sack and Swanson (2005), Nakamura and Steinsson (2018), and Swanson (2021). Second, we verify that our results are not driven by the Covid-19 pandemic. Namely, we show that the effect of monetary policy on rents is similar if we restrict the analysis to the pre-pandemic period.

Next, we turn to the mechanisms underlying our results. Using data on the universe of housing sales in the US from Corelogic, we show that monetary tightening leads to a drop in household demand for owner-occupied housing, which crowds in the rental market. In particular, a 25 basis point increase in the 30-year fixed rate mortgage lowers the number of properties purchased by owner-occupier households by 4% (7%) one (two) years following the shock. We find that the shift in household demand from the owner-occupier market to the rental market is accommodated by real-estate investors who mediate between the two markets. In response to monetary tightening, investors sell less houses to owner-occupiers but continue to buy as many houses from owner-occupiers - effectively increasing the share of housing that they own. Overall, these results suggest that monetary tightening increases household demand for rental housing, which is accommodated by real-estate investors who capitalize on the higher rents by buying houses from owner-occupiers and renting them out.

Related Literature

This paper is among the first to study the effects of monetary policy on rents. A large literature evaluates the impact of monetary policy on house prices (Case and Shiller, 1989; Kuttner, 2014; Williams et al., 2015; Gorea, Kryvtsov and Kudlyak, 2022), generally finding that house prices tend to decrease in response to contractionary shocks (albeit the effect can be small and takes time to materialize). Rents have so far received little attention. In sharp contrast to house prices, we find that rents *increase* in response to contractionary monetary policy. Dias and Duarte (2019), Koeniger, Lennartz and Ramelet (2022) also evaluate the impact of monetary policy shocks on rents and find similar results. An important advantage of our approach is that we construct a repeat-rent index to measure rent growth, while these papers use indexes of average or median rents. A repeat-rent index provides a *quality-constant* measure of rent growth, while measures of average and median rents cannot fully control for the quality of rented units and are biased by changes in the composition of rented units (Bailey, Muth and Nourse, 1963). A second advantage of our paper is that our micro-data allows us to build rent indexes at the zip-code level, while previous papers use measures of rent at the national level. This granularity enables us to study heterogeneous effects across local housing markets and to shed light on the mechanisms that drive the results.

Our work relates to a growing literature on the effects of mortgage lock-in on housing markets. Quigley (1987), Ferreira, Gyourko and Tracy (2010), and more recently Fonseca and Liu (2024), Liebersohn and Rothstein (2025), Batzer et al. (2024) and Aladangady, Krimmel and Scharlemann (2024), show that higher interest rates lower mobility rates of existing homeowners who hold fixed-rate mortgages. This "lock-in" effect occurs because selling a home and buying a new one would require these homeowners to take on mortgages at a prevailing market rate that is higher than the fixed rate on their outstanding mortgages. Mabille, Liu and Fonseca (2024) and Gerardi, Qian and Zhang (2024) develop structural models to study the equilibrium effects of mortgage lock-in. They find that mortgage lock-in decreases the supply of houses for sale and increases house prices. Most related to our paper is De la Roca, Giacoletti and Liu (2024), who study the impact of mortgage lock-in on rents. They find that, by increasing house prices, mortgage lock-in increases demand for rental units, which drives up rents. While these papers focus on mortgage lock-in, we study the effects of monetary policy more broadly. Mortgage lock-in, which limits supply in the owner-occupier market, is one channel through which monetary policy can affect rents. But monetary policy can impact rents even absent mortgage lock-in through various other channels. For example, by increasing borrowing costs, contractionary monetary policy can prevent first-time buyers from becoming homeowners (Ringo, 2024), which might crowd in the rental market (Gete and Reher, 2018).

Our paper also relates to the literature on the transmission of monetary policy via the mortgage market (Scharfstein and Sunderam, 2016; Garriga, Kydland and Šustek, 2017; Beraja et al., 2019; DeFusco and Mondragon, 2020; Di Maggio, Kermani and Palmer, 2020; Berger et al., 2021; Fuster et al., 2021; Eichenbaum, Rebelo and Wong, 2022). We show that the rental market also plays a key role in the transmission of monetary policy. Namely, by making mortgages more expensive, monetary tightening increases demand in the rental market and raises rents. From a positive perspective, this limits the ability of monetary tightening to lower inflation. From a normative perspective, it implies that monetary tightening can disproportionately harm the most vulnerable households who tend to be renters. These findings contribute more broadly to the literature on the distributional effects of monetary policy (Doepke, Schneider and Selezneva, 2015; Coibion et al., 2017; Kaplan, Moll and Violante, 2018; Auclert, 2019; Cloyne, Ferreira and Surico, 2020; Luetticke, 2021; Holm, Paul and Tischbirek, 2021; Amberg et al., 2022; Andersen et al., 2023).

Finally, our empirical results complement a literature that develops equilibrium models with both housing markets and rental markets to study the role of rental markets for the transmission of credit shocks in the housing market (Greenwald and Guren, 2021; Rotberg and Steinberg, 2024; Castellanos, Hannon and Paz-Pardo, 2024). In these models, credit-driven demand shocks in the housing market filter to the rental market because of real-estate investors who are active in both markets. By documenting the role of real-estate investors in the transmission of (monetary-policy-driven) credit shocks, our results provide empirical support to these models. In line with these models, we find that when credit becomes more expensive due to monetary tightening, household demand shifts from the owner-occupier market to the rental market, and this shift in demand is accommodated by real-estate investors who capitalize on the higher rents.

2 Data

This section describes our data. We begin by discussing the rental listing data that we use to construct our repeat-rent index. We then describe our measures of monetary policy shocks as well as the instrumental and control variables that we use for estimating the effects of monetary policy on rents.

2.1 Rent Prices

Our main data source is rental listing data compiled by Altos Research between January 2011 and September 2022. Altos compiles a national database of rental listings from online listing platforms and from Multiple Listings Services (MLS) platforms. Updated on a weekly basis, Altos provides a snapshot of listings that are observed during the week. For each listing, the data records the listed monthly rent, the date in which the listing is observed, the street address, zip-code, and geocodes of the unit being listed, as well as physical characteristics of the listed unit: the number of beds and baths, floor size, property type, year built, and whether the property features amenities such as air-conditioning and in-unit washer-dryer.

Sample Selection

We focus on listings of multifamily units and single family homes, and exclude short-term and vacation rentals, commercial properties, mobile homes, and listings of individual rooms. We drop listings where the rent, date, number of beds or number of baths is missing, as well as listing with incomplete information on the unit address. To avoid outliers, we drop listings with listed rents that exceed the 97.5 percentile or are below the 2.5 percentile of contract rents in the AHS. Sample selection is discussed in more detail in Appendix A.

Identifying Rental Units

In Section 3, we use our rental data to construct a repeat-rent index. To facilitate the construction of a repeat-rent index, one must first identify listings of the same unit across time. This is because, as discussed in more detail in 3, a repeat-rent index is constructed by comparing rents on the same unit across time. Since our data does not provide a unit identifier, our strategy is to identify units by their street address, number of beds and number of baths. That is, we assume that listings within the same building that have the same number of beds and the same number of baths correspond to the same unit.

Of course, in reality, multifamily buildings might feature multiple different units that have the same number of beds and baths. However, for the purpose of constructing a repeat-rent index, which is a quality-constant measure of rent growth, this is problematic only if these units differ in their quality. In other words, if units within the same building that have the same number of beds and baths are also of the same quality, then comparing a rent listed on one unit to a rent listed on another unit later in time indeed provides a quality-constant measure of rent growth.

If, in contrast, units within the same building that have the same number of beds and baths do differ in their quality, then comparing rents listed on one unit to a rent listed on another unit later in time does not provide a quality-constant measure of rent growth and would bias the repeat-rent index. Therefore, we drop units that are likely to differ in their quality but that are undistinguishable based on their address, number of beds and number of baths. Namely, units that correspond to tuples of address, number and beds and number of baths for which we observe within a same week multiple listings with different prices, or for which we observe statistically extreme rent fluctuations within a one-month or a one-year period.[2]

Our main analysis is at the monthly frequency. To obtain a monthly panel of listing prices at the unit level, we collapse all listings of the same unit that appear within the same month (e.g. due to the unit being listed on multiple platforms or being listed for several weeks within a month), to one observation. Namely, we keep the last listing observed within the month. Units that are not listed in more than one month are excluded, since they do not inform the repeat-rent index. Finally, to minimize the noise of our repeat-rent index, we drop zip-codes-months in which less than 15 units are listed and zip-codes that are observed in less than 70 months throughout our sample.

Our final panel data contains 30.3 million monthly observations of listed rents. It comprises 6.5 million rental units across 5,092 zip-codes. Each rental units is observed on average for approximately 4.7 months during the sample. The average time on market (i.e. the consecutive number of months a unit is being listed) is 2 months.

Geographical Coverage

Panel (a) of Figure 1 illustrates the geographical coverage of our data. For each county, we compute the percentage of all rental units in our data that are located in that particular county. Counties colored in lighter (darker) shades are counties where we observe a relatively small (large) number of rental units. Not

[2]We drop units that correspond to tuples of address, number and beds and number of baths for which we observe a 4-week (52-week) rent fluctuation (in absolute value) that exceeds the 95th percentile (99th percentile) of the 4-week (52-week) rent fluctuation distribution in the data.

surprisingly, we observe more rentals units in more densely populated areas of the U.S., for example the two coasts.

Panel (b) of Figure 1 illustrates the geographical distribution of rental units in the U.S., as measured from the nationally representative American Community Survey (ACS). Counties colored in lighter (darker) shades are counties with a relatively small (large) number of rental units. Reassuringly, comparing both panels suggests that the geographical coverage of our data aligns well with the geographical distribution of rental units in the U.S. This suggests that our data is representative of the U.S. rental market in terms of its geographical coverage.

Figure 1: Geographical Coverage

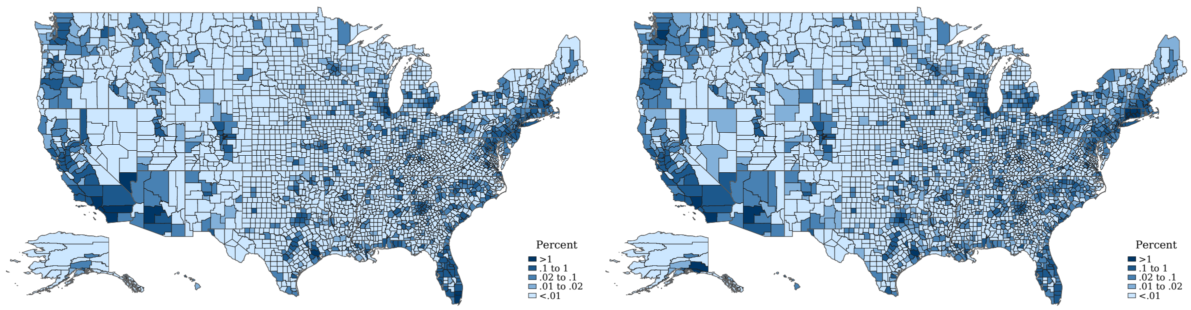

(a) Altos (b) ACS

Note: Panel (a) displays, for each county, the percentage of all rental units observed in our rent data between 2015 and 2019 that are located in that county. Panel (b) displays, for each county, the share of all rental units in the 2015-2019 ACS data that are located in that county. Darker colors correspond to higher shares.

Summary statistics

Table 1 compares summary statistics from our data to summary statistics computed from the American Housing Survey (AHS), from the ACS, and from Zillow. The AHS is a nationally representative survey of the housing stock in the U.S. Since it was redesigned in 2015, we compute the AHS summary statistics based on the 2015, 2017, 2019, and 2021 waves. For consistency, summary statistics from our Altos data, ACS and Zillow are computed based on the same time period. The first column compares the median rent across the different datasets, reported in 2015 U.S. dollars. Columns 2-5 compare physical characteristics of the median unit across datasets, and Column 6 reports the share of rental units that are single-family. Overall, our data over-represents higher tier rental units relative to the nationally representative AHS and ACS. The average rent in our data is higher than the average rent in the U.S, and units in our data are larger and somewhat more likely to be single-family dwellings. Our data aligns with Zillow in terms of the median rent.[3]

[3]Zillow does not provide summary statistics on the characteristics of rental units.

Table 1: Summary Statistics

	Rent ($) (1)	Year Built (2)	Bedrooms (#) (3)	Bathrooms (#) (4)	Sqft (5)	Single-family (%) (6)
ACS	927	1984	.	2	.	30
AHS	967	1974	2	1	875	35
Altos	1400	1981	3	2	1405	36
Zillow	1387	.	.	.	.	.

Note: This table presents summary statistics from ACS, AHS, our Altos data, and Zillow. The first column reports the median rent (in 2015 dollars), columns 2-5 compare physical characteristics of the median unit, and Column 6 reports the share of rental units that are single-family. ACS statistics are computed from the 1-year ACS surveys between 2015 and 2021. AHS statistics are computed from the 2015, 2017, 2019, and 2021 biennial surveys. Altos statistics are computed based on Altos data between 2015 and 2021. The median Zillow rent is computed as the median national ZORI between 2015-2021.

First, I collapse the national ZORI by year, then calculate the mean and median by collapsing the data again. An update: For ZORI, we were still (erroneously) keeping only odd years, as that was how we first mimicked the AHS data. However, we updated this a few weeks ago to correctly include all years. The mean ZORI is now 1,383, *andthemedianis*1, 381.

One plausible explanation for the discrepancy between our data and the AHS and ACS is that our data records listed rents, while the AHS and ACS record contract rents, which might be lower. Note that the Zillow data is also based on online listed rents (see Section 3.1 for a more detailed discussion of the Zillow data). A second explanation is selection. It might be that higher-quality rental units are more likely to be advertised on online listings platforms and therefore disproportionally more likely to be observed in our data.

A potential concern with the repeat-rent index that we construct in Section 3 is that, as illustrated by Table 1, it is based on a sample of rental units that is not representative of the U.S. rental market. We address this concern in Section 3.1 by showing that our repeat-rent index is representative of rent *growth*. In particular, despite differences in underlying samples, our index is consistent with alternative repeat-rent indexes that are based on nationally representative rental data. As discussed in Section 3.1, the advantage of our index is its broader and more granular geographical coverage.

2.2 Monetary Policy Shocks

In our main empirical specification, we measure exogenous shocks to monetary policy using the Bauer and Swanson (2023*b*) monetary policy surprises series. Bauer and Swanson (2023*b*) identify monetary policy shocks in two steps. First, they measure high-frequency changes in interest rates around Federal Open Market Committee (FOMC) meetings and around press conferences, speeches, and testimonies made by the Federal Reserve chair. Second, they orthogonalize these monetary policy surprises by regressing them on economic and financial variables that predate the FOMC meetings and Federal Reserve chair announcements, and take the residuals. We download the monthly Bauer and Swanson (2023*b*) monetary policy surprises from the Federal Reserve Bank of San Francisco data portal.[4] The publicly available series includes only monetary policy surprises measured around FOMC meetings. In our baseline specification, we limit the analysis to shocks occurring prior to the Covid-19 pandemic.[5] We show that the results are robust to including also post-pandemic shocks.

Recent work has shown that high-frequency changes to interest rates around FOMC meetings might not be exogenous. For example, Cieslak (2018), Miranda-Agrippino and Ricco (2021), and Bauer and Swanson (2023*a*) show that these changes are correlated with publicly available macroeconomic and financial indicators that predate these announcements. If high-frequency changes to interest rates around monetary policy events are not exogenous, they are not a valid instrument for estimating the effects of monetary policy (Stock and Watson, 2018). Bauer and Swanson (2023*b*) address this concern by regressing the changes in interest rates around monetary policy announcements on economic and financial data that predate these announcements. Their monetary policy surprises, which are the residuals from this regression, therefore capture the component of changes to interest rates around monetary policy events that is ex-ante unpredictable.

The choice of monetary policy shock can be important for the estimated effects of monetary policy (Ramey, 2016). We therefore replicate our empirical analysis for a host of alternative monetary policy shocks that have been used in the literature. In particular, we evaluate the robustness of our results to using monetary policy shocks from Gürkaynak, Sack and Swanson (2005), Gertler and Karadi (2015), Nakamura and Steinsson (2018), and Swanson (2021). In line with Bauer and Swanson (2023*b*), we construct a monthly series for each of these alternative shocks (which are provided at a daily frequency) by summing all daily shocks within each month.

2.3 Instrument and Control Variables

In Section 4, we employ a local projection instrumental variable approach (LP-IV) to evaluate the effects of monetary policy (Ramey, 2016; Stock and Watson, 2018). Here, we briefly describe the instrumented variables and controls that are use in the estimation. As the instrumented monetary policy indicator, we use the Freddie Mac 30-year fixed mortgage rate, downloaded from FRED (series: MORTGAGE30US). Bauer

[4]See https://www.frbsf.org/research-and-insights/data-and-indicators/monetary-policy-surprises/.

[5]Bauer and Swanson (2023*b*) recommend against using the monetary policy shocks that are measured during the Covid-19 pandemic since these shocks tend to be extreme in magnitude.

and Swanson (2023*b*) use the interest rate on two-year US Treasury bonds as their instrumented variable, but since our focus is on the housing market, we use the mortgage rate as our relevant monetary policy variable (as is common in the literature on the effects of monetary policy on housing markets, e.g. Aastveit and Anundsen (2022)). Our results are robust to using the interest rate on two-year US Treasury bonds as the instrumented variable (downloaded from the Federal Reserve Board website, series: SVENY02). For controls, we include lags of Core PCE inflation and lags of unemployment rates at the county level, as well as lags of monetary shocks and the monetary policy indicator. PCE (series: PCEPILFE) is downloaded from FRED. County-level unemployment rates are downloaded from the BLS.[6]

3 Repeat-Rent Index

We construct a repeat-rent index using our rental listings data. Hereafter, we refer to this index as the ADH-RRI. Introduced by Bailey, Muth and Nourse (1963), the repeat-sales method provides a quality-constant measure of price growth. In particular, it uses repeated sales of the same housing unit to control for observed and unobserved time-invariant quality components. Popularized by Case and Shiller (1989), repeat-sales indexes have become the gold standard of house price indexes. Starting with Ambrose, Coulson and Yoshida (2015), repeat-*rent* indexes have also been used by applying the repeat-sales method to the rental market (Clark, 2022; Adams et al., 2024).

Consider a rental unit that is observed in our listing data in both time s and time $t > s$. Assume that the data consists of listings observed in times $\{1, ..., N\}$. The repeat-rent index is constructed by estimating the following regression (Bailey, Muth and Nourse, 1963):

$$\log P_{i,t} - \log P_{i,s} = \gamma_1 D_{i,1} + \gamma_2 D_{i,2} + ... + \gamma_N D_{i,N} + \varepsilon_{i,t}, \tag{1}$$

where $P_{i,s}$ is the listed rent on unit i at time s and $P_{i,t}$ is the listed rent on the same unit i at a later time t. $D_{i,k} = 1$ if the second observation in the pair took place in time k, $D_{i,k} = -1$ if the first observation in the pair took place in time k, and $D_{i,k} = 0$ otherwise. The estimated parameters $\{\gamma_1, ...\gamma_N\}$ represent the percentage change in listed rents relative to the base (omitted) period. The exponent of these estimates constitute the repeat-rent index, where we normalize the value of the index in the base period to 100. That is, the ADH-RRI in time t is given by $ADHRRI_t = 100 \exp(\gamma_t)$.

The error term in Equation 1 is likely heteroskedastic due to variation in the time-gap between pairs of repeated listings (Case and Shiller, 1989). To address this, we follow Calhoun (1996). First, we estimate

[6] See https://download.bls.gov/pub/time.series/la/.

Equation 1 by OLS. Second, we regress the residuals from this regression on a constant, the time-gap between observations and the square of the time-gap between observations, and store the predicted values. Third, we estimate a weighted least squares version of Equation 1 using the inverse of the square roots of these predicted values as weights. We find that our RRI is practically unchanged due to this adjustment (as in Clark (2022) and Adams et al. (2024)).

We construct and analyze various specifications of our ADH-RRI. First, we construct both a nominal index which measures growth of nominal rents, and a real index which measures growth of real rents. The real index is computed by first deflating nominal rents by the non-shelter CPI index. Second, we construct both an "all listings" index and a "new listings" index. The latter is based only on new listings that come on the market (i.e. listings which were not observed in the previous period) while the former is based on all observed listings. Third, we construct an index at both the granular zip-code level and at the national level.[7] Fourth, we consider both monthly and quarterly indexes. Finally, we construct an index for single-family rental units, an index for for multifamily rental units, and an index for for all rental units.

3.1 Comparison to Alternative Rent Indexes

This section compares our ADH-RRI to popular alternative rent indexes - the Zillow Observed Rent Index (ZORI, Clark (2022)), the CPI-rent index, the Marginal Rent Index (ACY-MRI, Ambrose, Coulson and Yoshida (2023)), and the CPI-NTRR Adams et al. (2024). We begin by briefly describing the alternative indexes. We then show that our ADH-RRI aligns well with these indexes and demonstrate that it is representative of rent growth in the U.S. Finally, we discuss the advantages of our index, namely its broader and more granular geographical coverage.

Alternative Indexes

The CPI-rent index and the CPI-NTRR are both constructed from the BLS Housing Survey data. The survey is a nationally representative panel of renter-occupied housing units. For each unit, the survey records the contract rent, the utilities included, unit characteristics and tenants' move-in date. The BLS sample is divided into six panels, and each rental unit is surveyed every six months. For a detailed discussion of the CPI rent index and the CPI-NTRR, see Verbrugge and Poole (2010) and Adams et al. (2024). Below, we provide we brief summary.

[7] We construct the national index as a weighted average of the zip-code level index, with zip-codes weights corresponding to their aggregate rental stock value. That is, the national index is constructed as $ADHRRI_t = \sum_z ADHRRI_{z,t}\omega_z / \sum_z \omega_z$, where $ADHRRI_{z,t}$ is the ADH-RRI for zip-code z at time t and ω_z is the aggregate contract rent in zip-code z, calculated from the 2019 5-year ACS. This approach follows the method used to construct the national Case-Shiller house price index (https://www.spglobal.com/spdji/en/documents/methodologies/methodology-sp-corelogic-cs-home-price-indices.pdf) and ensures that zip-codes with a larger and more expensive rental stock are over-represented in the national index.

The CPI-rent index is constructed at the monthly frequency. Rent growth is measured by first calculating the average six-month rent growth across the units in the panel that is surveyed in that month, and then taking the sixth root of that average. The CPI-rent index measures the rent growth facing all tenants, regardless of their occupancy tenure. It adjusts for aging, structure changes, and changes in utilities included in rent. It is not a repeat rent index. The most granular geographical level for which the CPI-rent index is constructed is the Core Based Statistical Area (CBSA).

The CPI-NTRR (Adams et al., 2024) is a repeat-rent index that measures the rent growth faced by *new* tenants. This is in contrast to the CPI-rent index, which measures rent growth faced by both new and continuing tenants. By limiting the BLS sample only to observations where occupants are new tenants, the CPI-NTRR measures the rent growth that a new renter would face had she signed a new rent contract every period. A main advantage of the CPI-NTRR (and of the CPI-rent) is that it is based on a representative sample of U.S. rental units. The main limitation of the CPI-NTRR is that the sample size of BLS Housing Survey is relatively small. For this reason, the CPI-NTRR is constructed only at the quarterly frequency and only at the national level.

The Zillow Observed Rent Index (ZORI) is a repeat rent index that is constructed from Zillow's proprietary rental listings data and from MLS listing data. As the CPI-NTRR and our ADH-RRI, ZORI measures rent growth faced by new renters. ZORI is based on a sample of rental units that are listed online and is therefore not necessarily representative of the U.S. stock of rental units. For example, as illustrated by Table 1, units listed on Zillow and on MLS are of higher quality relative to the average rental unit in the country. ZORI is constructed at the monthly frequency and at the zip-code level. As we discuss below, the geographical and temporal coverage of ZORI is limited compared to our ADH-RRI. For a detailed discussion of ZORI, see Clark (2022).

The ACY-MRI (Ambrose, Coulson and Yoshida, 2023) is a rent index that measures rent growth faced by tenants in large multifamily buildings. It is constructed in two steps. First, a net rent index (NRI) is computed as the product of the Real Capital Analytics' (RCA) multifamily capitalization rate and the RCA commercial property price index (CPPI), which is a quality-adjusted repeat-sale index of multifamily properties. Second, the ACY-MRI is constructed by rescaling the NRI so that its mean and volatility match the mean and volatility of a previous rent index constructed by (Ambrose, Coulson and Yoshida, 2015). The ACY-MRI is constructed only at the national level.

Index Comparison

This section compares our ADH-RRI to alternative rent indexes. It shows that our index aligns well with the popular alternatives discussed above. Importantly, we show that the ADH-RRI is consistent with the nationally representative CPI-NTRR. This suggests that our ADH-RRI is representative of rent growth in the U.S.

Figure 2 compares the year-over-year rent inflation implied by our ADH-RRI to the rent inflation implied by the CPI-rent index, the CPI-NTRR, ZORI, and the ACY-MRI. Rent inflation is measured at the national level. The figure illustrates that, despite the aforementioned differences in the underlying rental data and index construction methods, our ADH-RRI closely tracks ZORI, the ACY-MRI, and, most importantly, the nationally representative CPI-NTRR. The four indexes clearly capture common rental market dynamics. Our ADH-RRI leads the CPI-rent index, as do the CPI-NTRR, ZORI, and the ACY-MRI. The main reason is that the CPI-rent index measures the rent growth faced by both new and existing tenants, while the other indexes measure the rent growth faced only by *new* tenants. Since rents on existing leases fluctuate less than rents on new leases, the CPI-rent index is less volatile and lags the other indexes.

Table 2 further validates our index by presenting the pairwise correlation coefficients between the year-over-year rent inflation in alternative rent indexes. Correlations between the ADH-RRI, CPI-NTRR, ACY-MRI and CPI-Rent are computed based on the quarters between 2012q3 and 2022q3. The correlations with ZORI are computed based on the quarters between 2016q1 and 2022q3. Our ADH-RRI is highly correlated with the nationally representative CPI-NTRR and with ZORI (the correlation coefficients are 0.95 and 0.96 respectively). The correlation between the ADH-RRI and the ACY-MRI is also high (0.73). The correlation between all indexes and the CPI-Rent index, which measured rent growth for both new and existing tenants, is lower. Overall, Figure 2 and Table 2 suggest that the ADH-RRI is a valid measure for rent inflation. The consistency between the ADH-RRI and alternative rent indexes is in line with previous studies that also document a strong correlation between alternative repeat-rent indexes (Ambrose, Coulson and Yoshida, 2015; Adams et al., 2024).

Advantages of the ADH-RRI

The main advantage of our ADH-RRI relative to alternative prominent rent indexes is its broad and granular geographical coverage. In particular, the ADH-RRI is computed at the zip-code level, while The CPI-NTRR and ACY-MRI are constructed only at the national level, and the most granular geographical level for which the CPI-rent index is constructed is the CBSA. ZORI is also constructed at the zip-code level, but as we discuss below its geographical and temporal coverage is limited compared to our ADH-RRI.

Figure 2: Rent Inflation in Alternative Rent Indexes

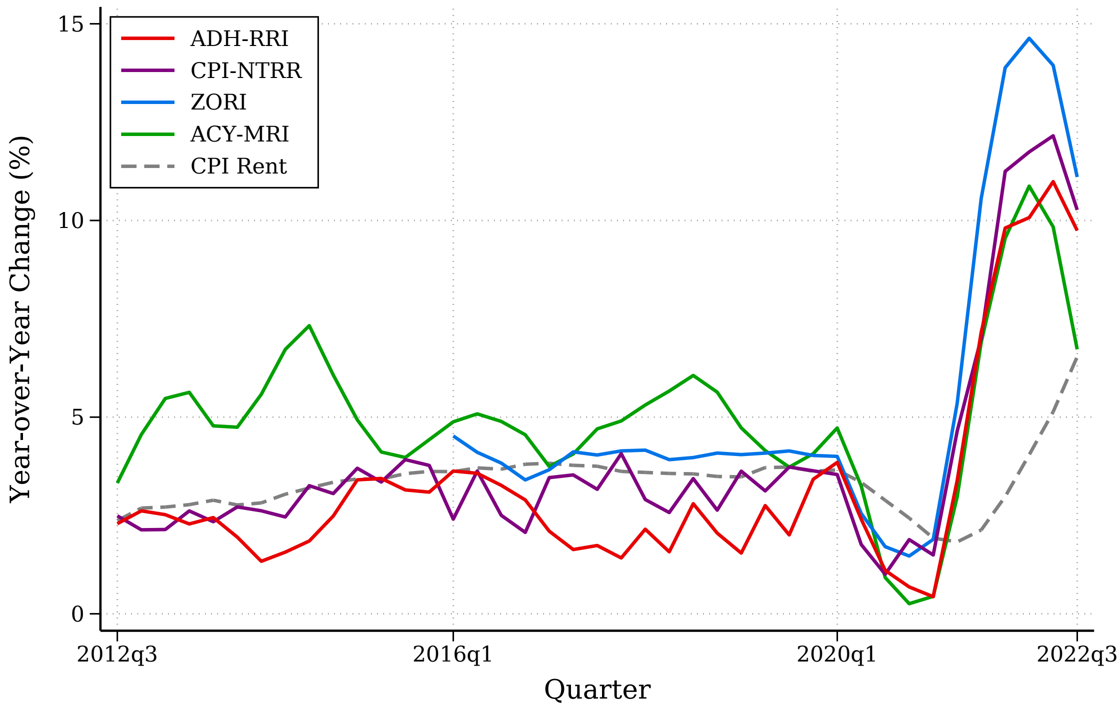

Note: This figure plots the year-over-year rent inflation in alternative rent indexes. Year-over-year inflation is computed for each quarter between 2012q3 and 2022q3. The CPI-NTRR is downloaded from `https://www.bls.gov/pir/new-tenant-rent.htm` at a quarterly frequency. ZORI is downloaded from `https://www.zillow.com/research/data/`, ACY-MRI is downloaded from `https://sites.psu.edu/inflation/`, and CPI-Rent is downloaded from `https://fred.stlouisfed.org/series/CUSR0000SEHA`, all at a monthly frequency. All monthly indexes are converted to quarterly indexes by averaging across months within the quarter.

To illustrate the advantage of the ADH-RRI relative to ZORI, Figure 3 plots the number of zip-codes covered by both indexes across time. First, the ADH-RRI is available starting from 2011, while ZORI is only available starting from 2015. Second, even when ZORI becomes available in 2015, our ADH-RRI covers substantially more zip-codes. In 2015, the ADH-RRI provides measures of rent inflation for 5,278 zip-code, while ZORI covers only 1,004 zip-codes. As time passes, ZORI expands it coverage, but even in 2022 our ADH-RRI covers approximately 1,000 zip-codes more than ZORI. Figure B.1 in the Appendix illustrates which zip-codes are covered by each index between 2011 and 2022. Overall, the ADH-RRI provides broader and more granular measures of rent growth relative to alternative indexes. This is key for studying the effects of monetary policy on rents and for investigating the mechanisms that are in play.

Table 2: Correlation between Alternative Rent Indexes

	ADH-RRI	CPI-NTRR	ZORI	ACY-MRI	CPI-Rent
ADH-RRI	1.00	0.95	0.96	0.73	0.49
CPI-NTRR	0.95	1.00	0.98	0.75	0.49
ZORI	0.96	0.98	1.00	0.88	0.37
ACY-MRI	0.73	0.75	0.88	1.00	0.41
CPI-Rent	0.49	0.49	0.37	0.41	1.00

Note: This table reports the pairwise correlation coefficients between year-over-year rent inflation in alternative rent indexes. The correlations between ADH-RRI, CPI-NTRR, ACY-MRI and CPI-Rent are computed based on the quarters between 2012q3 and 2022q3. The correlations with ZORI are computed based on the quarters between 2016q1 and 2022q3.

Figure 3: Coverage - ADH-RRI and ZORI

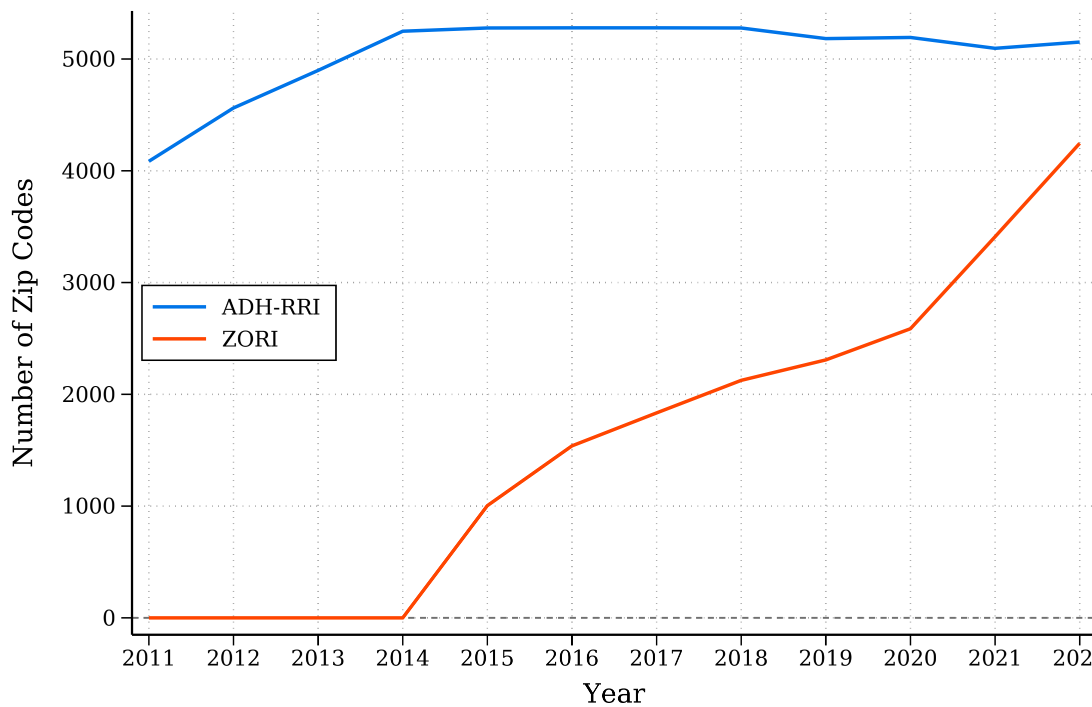

Note: This figure plots number of zip-codes covered by the ADH-RRI and by ZORI for every year between 2011 and 2022.

4 Effect of Monetary Policy on Rent

In this section, we use our repeat-rent index to evaluate the effects of monetary policy on rents. In particular, we estimate the dynamic effects of a monetary policy shock on rents using the standard local projection instrumental variable (LP-IV) framework (Jordà, Schularick and Taylor, 2015; Ramey, 2016). The LP-IV framework combines the Jordà (2005) local projection approach (LP) with instrumental variable (IV) methods and is discussed in detail in Stock and Watson (2018). To perform an LP-IV estimation, we estimate the following LP regression via two-stage least squares:

$$\log ADHRRI_{z,t+h} - \log ADHRRI_{z,t-1} = \beta^{(h)} i_t + \Gamma^{(h)} X_{z,t-1} + \varepsilon_{z,t+h}, \tag{2}$$

for each horizon $h = \{0, 1, ..., 24\}$. The dependent variable is the cumulative growth rate of rent in zip-code z between month $t-1$ and month $t+h$, measured based on our ADH-RRI. i_t is a monetary policy indicator. Since our context is the housing markets, we use the 30-year fixed rate mortgage as our monetary policy indicator (Aastveit and Anundsen, 2022; Gorea, Kryvtsov and Kudlyak, 2022). $\beta^{(h)}$ is the coefficient of interest which captures how a change in the monetary policy indicator impacts rent inflation going forward. $X_{z,t-1}$ is a set of controls which we include to ensure that the LP-IV estimation satisfies the (Stock and Watson, 2018) instrument validity conditions. We discuss the validity conditions and controls in more detail below. Standard errors are clustered by both zip-code (to account for potential serial correlation of errors across time) and by time period (to account for potential heteroskedasticity of errors across zip-codes) and are estimated using the (Cameron, Gelbach and Miller, 2011) multi-way clustering estimation method.

In the first stage, the monetary policy indicator i_t, which is in all likelihood endogenous, is instrumented with a directly observed measure of exogenous monetary policy shock, denoted by s_t. As explained by Stock and Watson (2018), monetary policy surprises measured from high-frequency interest rate changes around FOMC meetings might capture only part of the true underlying (and unobserved) monetary policy shock and might be measured with error. They are therefore instruments for the true monetary policy shock, not the shock itself. Since the true monetary policy shock is unobserved, the instrumented variable in the LP regression is the observed monetary policy indicator i_t.

For the observed monetary policy shock s_t to be a valid instrument for the monetary policy indicator, it must satisfy three conditions (Stock and Watson, 2018). First, the standard relevance condition must hold. That is, conditional on the controls $X_{z,t-1}$, the observed monetary policy shock s_t must be correlated with the monetary policy indicator i_t. Second, the standard exogeneity condition must hold. That is, conditional on controls, s_t must not be correlated with $\varepsilon_{z,t+h}$. Third, the lead-lag exogeneity condition must hold. That is, conditional on controls, s_t must not be correlated with leads and lags of the monetary policy indicator i_t.

The set of controls in Equation 3 are chosen to ensure that the instrument validity conditions hold. They include zip-code level controls, namely zip-code month-of-year fixed effects which control for seasonality at the zip-code level, lags of the growth rate of rent, and lags of the change in the unemployment rate, as well as macro level controls, namely lagged PCE inflation, lagged changes in the mortgage rate i_t, and lagged monetary policy shocks.[8] In Section 4.1, we show that the relevance condition and the lead-lag exogeneity

[8]In particular, we include 12 lags of the monthly zip-code level growth rate of the ADH-RRI, one lag of the year-over-year change in the unemployment rate at the county level, one lag of the year-over-year PCE inflation, one lag of the instrument, and 4 lags of the

condition hold. The exogeneity condition relies on the identification of truly exogenous monetary policy surprises. In our main specification, we use the (Bauer and Swanson, 2023*b*) monetary surprise series. For robustness, we also consider a host of alternative monetary policy shocks that have been used in the literature.

The main appeal of the local projection framework relative to the vector auto-regression (VAR) framework is that it allows for non-linear effects of monetary policy shocks without modeling them as a system (Jordà, Schularick and Taylor, 2015; Ramey, 2016). Local projection is also more robust to misspecification (Jordà, 2005). Since the local projection framework imposes less restrictions, it often results in more erratic and less precisely estimated impulse response functions. However, due to the large size of our dataset, this is not the case in our application.

4.1 Results

Figure 4 illustrates our main result. Panel (a) plots the impulse response function of real rent inflation to an exogenous 25 basis point increase in the interest rate on a 30-year fixed rate mortgage. The dark (light) shaded areas correspond to the 68% (90%) confidence intervals. We find that higher interest rates lead to higher rents. In particular, a 25 basis point increase in the interest rate on a 30-year fixed rate mortgage leads to a 1% (1.7%) increase in real rents 12 (24) months following the monetary policy shock. This result has important implications in terms of the distributional effects of monetary policy. Namely, by making renting more expensive relative to all other prices in the economy, contractionary monetary policy might disproportionately harm the most vulnerable households in the population, who tend to be renters.

Panel (b) plots the impulse response function of *nominal* rent inflation to the same exogenous increase in interest rate. We find that higher interest rates not only increase real rents, but also nominal rents. In particular, a 25 basis point increase in the interest rate on a 30-year fixed rate mortgage leads to a 0.7% (1.4%) increase in nominal rents 12 (24) months following the monetary policy shock. This result has important implications in terms of the ability of monetary policy to effectively control inflation. If contractionary monetary policy increases rents, which account for more than 35 percent of the CPI, this substantially limits the ability of contractionary monetary policy to lower inflation.

Instrument validity. If the observed monetary policy shock satisfies the three instrument validity conditions specified in (Stock and Watson, 2018), then estimating Equation 3 yields consistent estimates of the effect of the 30-year fixed rate mortgage rate on rent growth. Here, we show that the relevance condition

monthly change in the monetary policy indicator.

Figure 4: Effect of Monetary Policy on Rents

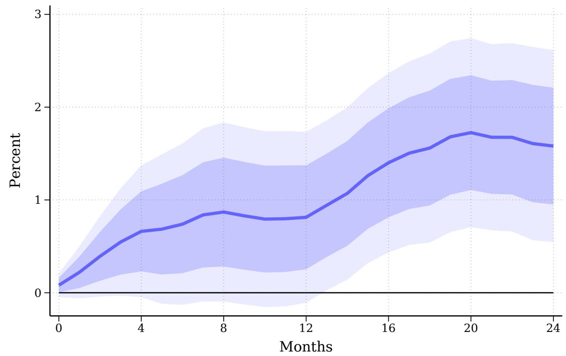

(a) Real Rents

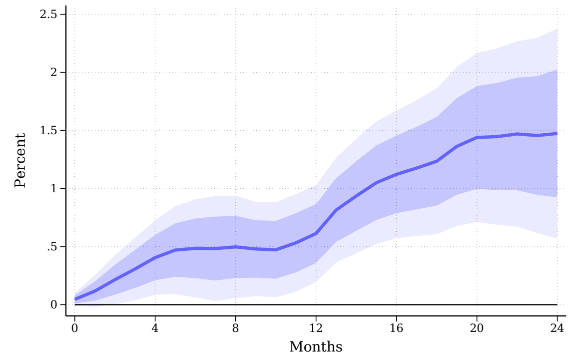

(b) Nominal Rents

Note: Panel (a) (Panel (b)) displays the impulse response function of real (nominal) rent inflation to a 25bps increase in the 30-year fixed rate mortgage. Dark (light) shaded areas represent 68% (90%) confidence intervals. Standard errors are clustered at both the zip-code and month level and are estimated using the (Cameron, Gelbach and Miller, 2011) multi-way clustering estimation method.

and the lead-lag exogeneity condition, which are testable, hold. The exogeneity condition relies on the (Bauer and Swanson, 2023*b*) monetary policy surprises being exogenous and is not directly testable.

To test the relevance condition, we compute the Olea and Pflueger (2013) first-stage effective F-statistic. Figure B.2 in the Appendix plots the effective F-statistic of the first-stage of Equation 3 for the case where the outcome is real rent growth (Panel (a)) and for the case where the outcome is nominal rent growth (Panel (b)), for each horizon $h = \{0, 1, ..., 24\}$. As illustrated by the Figure, the F-statistic is consistently above 10, the rule of thumb cutoff for weak instruments recommended by Staiger and Stock (1997) and Andrews, Stock and Sun (2019). To test the lead-lag exogeneity condition, we follow (Stock and Watson, 2018) and regress the orthogonalized instrument on orthogonalized leads and lags of the dependent variable. Orthogonalizing is done against the set of controls $X_{z,t-1}$. The R-square in these regressions ranges between 0.003 and 0.01. This illustrates that the instrument is unforecastable by leads and lags of the dependent variable, thus satisfying the lead-lag exogeneity condition.

4.2 Robustness

We consider a host of robustness tests for our main result. First, the choice of monetary policy surprises s_t can be important for the estimated effects of monetary policy (Ramey, 2016). We therefore replicate our empirical analysis using a host of alternative monetary policy shocks that have been used in the literature. In particular, we re-estimate Equation 3 using monetary policy shocks from Gürkaynak, Sack and Swanson (2005), Nakamura and Steinsson (2018), Swanson (2021), and the non-orthogonalized shocks from Bauer and Swanson (2023*b*) as alternative instruments.[9] Figure 5 shows the impulse response function of nominal rent inflation for each of these alternative specifications. Figure **??** in the appendix replicates this exercise for real rents. Reassuringly, the results are qualitatively and quantitatively in line with the main specification shown in Figure 4.

Figure 5: Alternative Monetary Policy Shocks

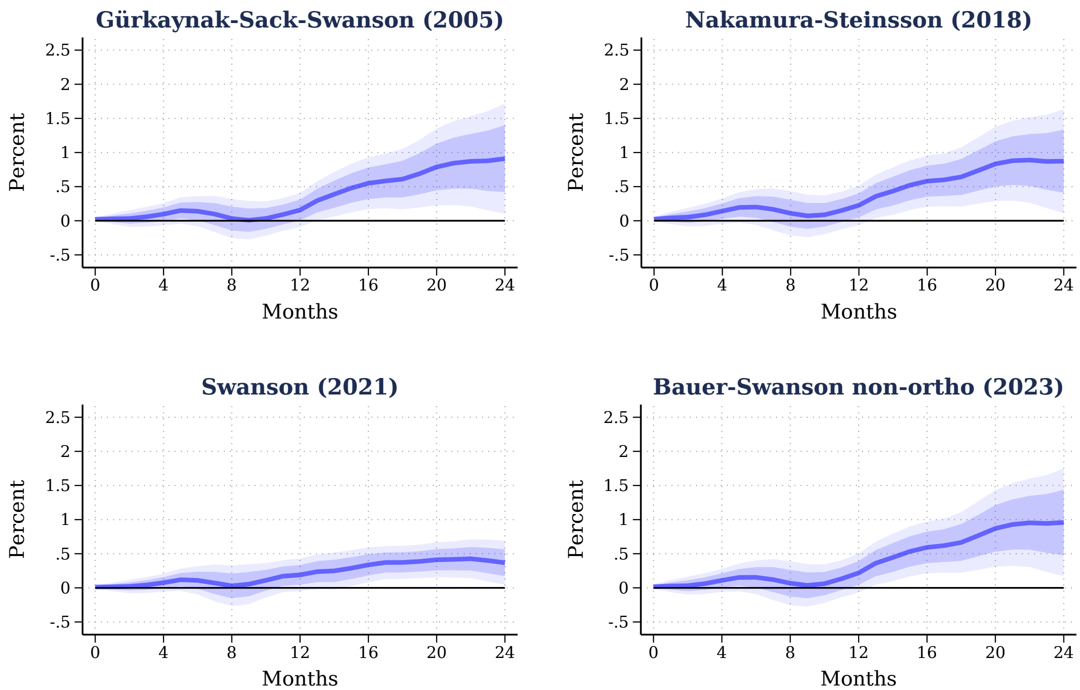

Note: This figure displays the impulse response function of nominal rent inflation to a 25bps increase in the 30-year fixed rate mortgage using alternative monetary policy shocks. The top left panel corresponds to the Gürkaynak, Sack and Swanson (2005) shocks, where we use both the surprise changes in the federal funds rate and the surprise changes in forward guidance as instruments for the monetary policy indicator. The top right panel corresponds to the Nakamura and Steinsson (2018) shocks. The bottom left panel corresponds to the Swanson (2021) shocks, where we use surprise changes in the federal funds rate, in forward guidance, and in large-scale asset purchases (LSAPs) as instruments for the monetary policy indicator. The bottom right panel corresponds to the non-orthogonalized shocks from Bauer and Swanson (2023*b*). Dark (light) shaded areas represent 68% (90%) confidence intervals. Standard errors are clustered at both the zip-code and month level and are estimated using the (Cameron, Gelbach and Miller, 2011) multi-way clustering estimation method.

[9] In these specifications, we also include the controls used by (Bauer and Swanson, 2023*b*) to orthogonalize high-frequency changes in interest rates around FOMC meetings and around Federal Reserve chair announcements. As pointed out by Cieslak (2018) and Miranda-Agrippino and Ricco (2021), high-frequency changes to interest rates around these events might not be exogenous. Bauer and Swanson (2023*b*) address this concern by regressing changes in interest rates around monetary policy announcements on economic and financial data that predate these announcements. These controls include the surprise component of the most recent nonfarm payrolls release from (Bauer and Swanson, 2023*b*), employment growth over the past year as constructed by (Cieslak, 2018), the change in the slope of the yield curve from 3 months before to one day before the FOMC announcement (measured as the second principal component of 1-to-10-year zero-coupon Treasury yields by (Gürkaynak, Sack and Wright, 2007)), the log change in the Bloomberg Commodity Spot Price Index over the same period, the log change in the S&P 500 index over the same period, and the option-implied skewness of the 10-year Treasury yield from (Bauer and Chernov, 2024).

5 Mechanisms

Why does monetary tightening increase rents? In this section we explore the mechanisms driving our main results. We find that a higher interest rate shifts household demand from the owner-occupied market to the rental market. This increase in demand for rentals is accommodated by real-estate investors who capitalize on the higher rents by buying houses from owner-occupiers and renting them in the rental market.

5.1 Sales in the Housing Market

We investigate how monetary policy shocks affect the volume and composition of sales in the housing market. We show that a contractionary shock lowers the volume of houses that are bought by owner-occupier households and increases the share of houses that are owned by real-estate investors. This suggests that household demand shifts from the owner-occupied market to the rental market.

Data. To analyze the impact of monetary policy on housing sales, we use Corelogic data. Corelogic compiles the universe of housing transactions in the US. For each transaction, the data records, for example, the property address, the mailing address of the buyer and the date of transaction. Properties are identified by their Assessor's Parcel Number (APN), which allows tracking the same property across subsequent transactions. We limit our sample to apartments, single family residences, condominiums, and duplexes and to arms-length transactions. For each transaction, we classify the buyer as either an owner-occupier or investor. A buyer is defined as an owner-occupier if her mailing address is the same as the property address and as an investor otherwise. Similarly, a seller is defined as an owner-occupier if her mailing address at the time she bought the property was the same as the property address. We also consider an alternative classification of buyers based on whether or not they are cash-buyers. A buyer is defined as a cash-buyer if she does not take a mortgage to finance the purchase. We interpret cash-buyers as buyers who are likely investors.

Due to the infrequent nature of housing transactions, our analysis of housing transactions is conducted at the quarter frequency and at the county level. For each county and quarter between 2000Q1 and 2024Q2, we compute the total number of transactions and the composition of these transactions based on whether the buyer and seller are owner-occupiers or investors. To mitigate noise, we restrict our sample to counties that historically have a sufficiently high volume of transactions. In particular, for each county we compute the average number of quarterly transactions across the sample and we restrict the analysis to counties that are above the median. We further compute, for each county, the average number of quarterly transactions where the buyer is an investor, and restrict the analysis to counties that are above the median[10].

LP-IV. We estimate the dynamic effects of a monetary policy shock on sales in the housing market using the LP-IV framework described above. In similar fashion to Equation 3, we estimate the following LP regression via two-stage least squares:

$$\log S_{c,t+h} - \log S_{c,t-1} = \beta^{(h)} i_t + \Gamma^{(h)} X_{c,t-1} + \varepsilon_{c,t+h}, \tag{3}$$

for each horizon $h = \{0, 1, ..., 8\}$. The dependent variable is the cumulative change in the number of transactions in county c between quarter $t-1$ and quarter $t+h$. For the number of transactions, we will consider the overall number of transactions, the number of transactions where the buyer is an owner-occupier and the seller is an investor (which we refer to as "owner-to-investor" transactions), the number of transactions where the buyer is an owner-occupier and the seller is an owner-occupier ("owner-to-owner"), the number of transactions where the buyer is an investor and the seller is an investor ("investor-to-investor"), and the number of transactions where the buyer is an investor and the seller is an owner-occupier ("investor-to-owner"). i_t is the monetary policy indicator, namely the average 30-year fixed rate mortgage within the quarter. Following Gerardi, Qian and Zhang (2024), since the transaction analysis is at the quarter level, our instrument is the quarterly average of the (Bauer and Swanson, 2023*b*) monetary surprise series. $X_{c,t-1}$ is a set of controls which include county quarter-of-year fixed effects, lags of the dependent variable, lags of the change in the county unemployment rate, lagged PCE inflation, and lagged changes in the monetary indicator i_t.[11]

Results. Panel (a) of Figure 6 plots the impulse response function of the change in the volume of housing transactions to an exogenous 25 basis point increase in the interest rate on a 30-year fixed rate mortgage. We find that higher interest rates lead to a drop in the number of sales - a hypothetical 25 basis point increase in the 30-year fixed rate mortgage lowers the volume of sales by 4% (7%) one (two) years following the contractionary shock.

Importantly, panels (b) and (c) show that the effect is driven entirely by a drop in the number of transactions where the buyer is an owner-occupier. In particular, the volume of housing transactions where the buyer is an owner-occupier drops by 4% (7%) one (two) years following the shock, while the effect on the volume of transactions where the buyer is an investor is statistically undistinguishable from zero. Figure

[10] The final sample is a panel of 918 counties and 89,964 quarterly observations.

[11] In particular, we include 4 lags of the quarterly growth rate of the number of transactions, one lag of the year-over-year change in the quarterly unemployment rate at the county level, one lag of the year-over-year quarterly PCE inflation, and one lag of the quarterly change in the monetary policy indicator.

Figure 6: Effect of Monetary Policy on Housing Sales

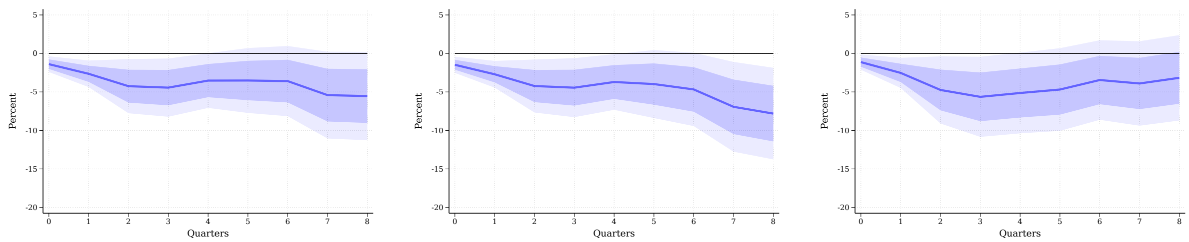

(a) Total Sales (b) Sales to Owner Occupiers (c) Sales to Investors

Note: Panel (a) displays the impulse response function of the change in the volume of housing transactions to an exogenous 25 basis point increase in the 30-year fixed rate mortgage. Panel (b) (Panel (c)) displays the impulse response function of the change in the volume of housing transactions where the buyer is an owner-occupier (investor). Dark (light) shaded areas represent 68% (90%) confidence intervals. Standard errors are clustered at both the zip-code and month level and are estimated using the (Cameron, Gelbach and Miller, 2011) multi-way clustering estimation method.

B.4 shows that this result is robust to an alternative definition of investors and owner-occupiers where buyers are classified as investors (owner-occupiers) if they are a cash-buyer (not a cash-buyer). The result is intuitive. When the interest rate increases, higher borrowing costs deter financially constrained households from buying homes. Deep-pocketed investors, however, are not impacted by borrowing costs. All else equal, we might still expect investors to pull out of the housing market as alternative investments, for example money market funds, become more attractive. However, as we show in Figure 4, rents increase in response to a contractionary monetary shock. This makes investment in housing more attractive, and explains why the volume of investor activity in the housing market is unaffected by the increase in interest rate.

Figure 7 shows that monetary tightening increases the share of houses that are owned by real-estate investors. It plots the impulse responses of the change in the volume of "owner-to-owner", "owner-to-investor", "investor-to-owner", and "investor-to-investor" transactions. While the number of houses that are sold by investors to owner-occupiers drops (Panel (b)), the number of houses that transition from owner-occupiers to investors is unaffected by the monetary policy shocks (Panel (c)). Overall, these results suggest that a higher interest rate shifts household demand from the owner-occupied market to the rental market, and that this shift is accommodated by investors who capitalize on the higher rents.

6 Conclusion

This paper provides new causal estimates of the effects of monetary policy on housing rents across the United States, covering the period from 2011 to 2022. Using rental listings data, we construct a new repeat-rent index that has a broader geographical and temporal coverage relative to alternative popular rent indices. Employing standard local projections methods, we find that contractionary monetary policy

Figure 7: Effect of Monetary Policy on Ownership Composition

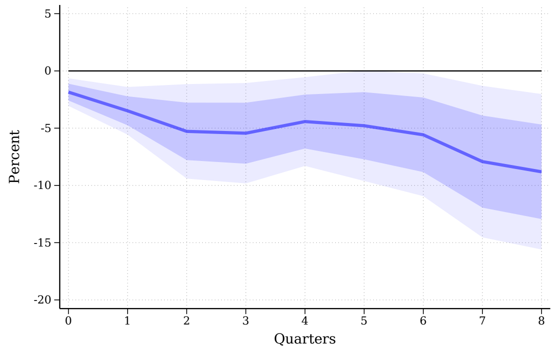

(a) Owner-to-Owner

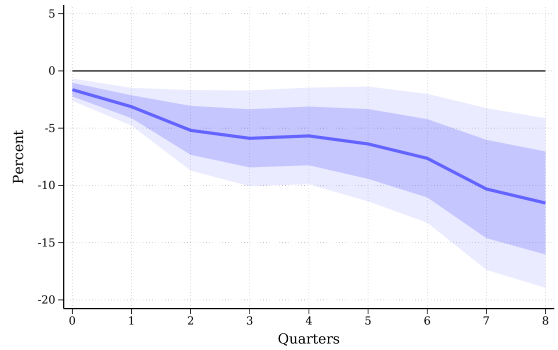

(b) Investor-to-Owner

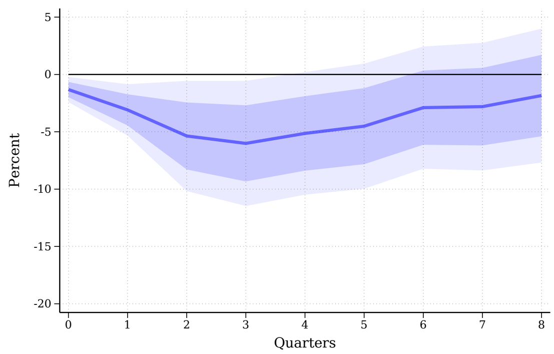

(c) Owner-to-Investor

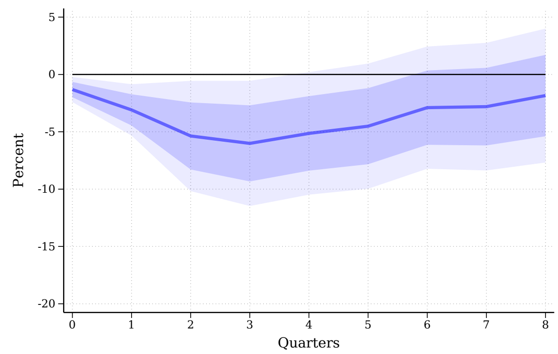

(d) Investor-to-Investor

Note: This figure displays the impulse response function of the change in the volume of "owner-to-owner" transactions (Panel (a)), "investor-to-owner" transactions (Panel (b)), "owner-to-investor" transactions (Panel (c)) and "investor-to-investor" transactions (Panel (d)) to an exogenous 25 basis point increase in the 30-year fixed rate mortgage. Dark (light) shaded areas represent 68% (90%) confidence intervals. Standard errors are clustered at both the zip-code and month level and are estimated using the (Cameron, Gelbach and Miller, 2011) multi-way clustering estimation method.

increases both real and nominal rents. Using housing transactions data, we show that the result is driven by a drop in households demand for owner-occupied housing, which crowds in the rental market. The shift in households demand is accommodated by deep-pocketed real-estate investors who buy houses from owner-occupiers and rent them out to renters.

Our results have important normative and positive implications. On the normative side, the results highlight the unintended distributional effects of monetary policy. By making rents relatively more expensive, monetary tightening makes housing less affordable to the most vulnerable households in the population. On the positive side, our results carry important implications for the transmission of monetary policy to inflation. The finding that monetary tightening increases rent, which is the single largest component of the Consumer Price Index (CPI), suggests that it is limited in the extent to which it can effectively lower inflation.

References

Aastveit, Knut Are, and André K Anundsen. 2022. "Asymmetric effects of monetary policy in regional housing markets." *American Economic Journal: Macroeconomics*, 14(4): 499–529.

Adams, Brian, Lara Loewenstein, Hugh Montag, and Randal Verbrugge. 2024. "Disentangling rent index differences: data, methods, and scope." *American Economic Review: Insights*, 6(2): 230–245.

Aladangady, Aditya, Jacob Krimmel, and Tess Scharlemann. 2024. "Locked In: Rate Hikes, Housing Markets, and Mobility."

Amberg, Niklas, Thomas Jansson, Mathias Klein, and Anna Rogantini Picco. 2022. "Five facts about the distributional income effects of monetary policy shocks." *American Economic Review: Insights*, 4(3): 289–304.

Ambrose, Brent W, N Edward Coulson, and Jiro Yoshida. 2015. "The repeat rent index." *Review of Economics and Statistics*, 97(5): 939–950.

Ambrose, Brent W, N Edward Coulson, and Jiro Yoshida. 2023. "Housing rents and inflation rates." *Journal of Money, Credit and Banking*, 55(4): 975–992.

Andersen, Asger Lau, Niels Johannesen, Mia Jørgensen, and JOSÉ-LUIS PEYDRÓ. 2023. "Monetary policy and inequality." *The Journal of Finance*, 78(5): 2945–2989.

Andrews, Isaiah, James H Stock, and Liyang Sun. 2019. "Weak instruments in instrumental variables regression: Theory and practice." *Annual Review of Economics*, 11(1): 727–753.

Auclert, Adrien. 2019. "Monetary policy and the redistribution channel." *American Economic Review*, 109(6): 2333–2367.

Bailey, Martin J, Richard F Muth, and Hugh O Nourse. 1963. "A regression method for real estate price index construction." *Journal of the American Statistical Association*, 58(304): 933–942.

Batzer, Ross, Jonah Coste, William Doerner, and Michael Seiler. 2024. "The Lock-In Effect of Rising Mortgage Rates." Federal Housing Finance Agency.

Bauer, Michael, and Mikhail Chernov. 2024. "Interest rate skewness and biased beliefs." *The Journal of Finance*, 79(1): 173–217.

Bauer, Michael D, and Eric T Swanson. 2023*a*. "An alternative explanation for the "fed information effect"." *American Economic Review*, 113(3): 664–700.

Bauer, Michael D, and Eric T Swanson. 2023*b*. "A reassessment of monetary policy surprises and high-frequency identification." *NBER Macroeconomics Annual*, 37(1): 87–155.

Beraja, Martin, Andreas Fuster, Erik Hurst, and Joseph Vavra. 2019. "Regional heterogeneity and the refinancing channel of monetary policy." *The Quarterly Journal of Economics*, 134(1): 109–183.

Berger, David, Konstantin Milbradt, Fabrice Tourre, and Joseph Vavra. 2021. "Mortgage prepayment and path-dependent effects of monetary policy." *American Economic Review*, 111(9): 2829–2878.

Calhoun, Charles A. 1996. "OFHEO house price indexes: HPI technical description." *Office of Federal Housing Enterprise Oversight*, 20552: 1–15.

Cameron, A Colin, Jonah B Gelbach, and Douglas L Miller. 2011. "Robust inference with multiway clustering." *Journal of Business & Economic Statistics*, 29(2): 238–249.

Campbell, JY. 1989. "Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence." *Macroeconomics Annual, edited by Olivier Jean Blanchard and Stanley Fischer/MIT Press*.

Case, Karl E, and Robert J Shiller. 1989. "The efficiency of the market for single-family homes." *American Economic Review*, 79(1): 125–137.

Castellanos, Juan, Andrew Hannon, and Gonzalo Paz-Pardo. 2024. "The aggregate and distributional implications of credit shocks on housing and rental markets."

Cieslak, Anna. 2018. "Short-rate expectations and unexpected returns in treasury bonds." *The Review of Financial Studies*, 31(9): 3265–3306.

Clark, Joshua. 2022. "Methodology: Zillow Observed Rent Index (ZORI)." *URL https://www. zillow. com/research/methodology-zori-repeat-rent-27092*.

Cloyne, James, Clodomiro Ferreira, and Paolo Surico. 2020. "Monetary policy when households have debt: new evidence on the transmission mechanism." *The Review of Economic Studies*, 87(1): 102–129.

Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia. 2017. "Innocent Bystanders? Monetary policy and inequality." *Journal of Monetary Economics*, 88: 70–89.

DeFusco, Anthony A, and John Mondragon. 2020. "No job, no money, no refi: Frictions to refinancing in a recession." *The Journal of Finance*, 75(5): 2327–2376.

De la Roca, Jorge, Marco Giacoletti, and Lizhong Liu. 2024. "Mortgage Rates and Rents: Evidence from Local Mortgage Lock-In Effects." *Working paper*.

Dias, Daniel A, and João B Duarte. 2019. "Monetary policy, housing rents, and inflation dynamics." *Journal of Applied Econometrics*, 34(5): 673–687.

Di Maggio, Marco, Amir Kermani, and Christopher J Palmer. 2020. "How quantitative easing works: Evidence on the refinancing channel." *The Review of Economic Studies*, 87(3): 1498–1528.

Doepke, Matthias, Martin Schneider, and Veronika Selezneva. 2015. "Distributional effects of monetary policy." *Unpublished manuscript*.

Eichenbaum, Martin, Sergio Rebelo, and Arlene Wong. 2022. "State-dependent effects of monetary policy: The refinancing channel." *American Economic Review*, 112(3): 721–761.

Ferreira, Fernando, Joseph Gyourko, and Joseph Tracy. 2010. "Housing busts and household mobility." *Journal of urban Economics*, 68(1): 34–45.

Fonseca, Julia, and Lu Liu. 2024. "Mortgage Lock-In, Mobility, and Labor Reallocation." *The Journal of Finance*.

Fuster, Andreas, Aurel Hizmo, Lauren Lambie-Hanson, James Vickery, and Paul S Willen. 2021. "How resilient is mortgage credit supply? Evidence from the COVID-19 pandemic." National Bureau of Economic Research.

Garriga, Carlos, Finn E Kydland, and Roman Šustek. 2017. "Mortgages and monetary policy." *The Review of Financial Studies*, 30(10): 3337–3375.

Gerardi, Kristopher, Franklin Qian, and David Zhang. 2024. "Mortgage Lock-in, Lifecycle Migration, and the Welfare Effects of Housing Market Liquidity." *Lifecycle Migration, and the Welfare Effects of Housing Market Liquidity (July 28, 2024)*.

Gertler, Mark, and Peter Karadi. 2015. "Monetary policy surprises, credit costs, and economic activity." *American Economic Journal: Macroeconomics*, 7(1): 44–76.

Gete, Pedro, and Michael Reher. 2018. "Mortgage supply and housing rents." *The Review of Financial Studies*, 31(12): 4884–4911.

Gorea, Denis, Oleksiy Kryvtsov, and Marianna Kudlyak. 2022. "House price responses to monetary policy surprises: Evidence from the US listings data."

Greenwald, Daniel L, and Adam Guren. 2021. "Do credit conditions move house prices?" National Bureau of Economic Research.

Gürkaynak, Refet S, Brian Sack, and Eric T Swanson. 2005. "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements." *International Journal of Central Banking*.

Gürkaynak, Refet S, Brian Sack, and Jonathan H Wright. 2007. "The US Treasury yield curve: 1961 to the present." *Journal of monetary Economics*, 54(8): 2291–2304.

Holm, Martin Blomhoff, Pascal Paul, and Andreas Tischbirek. 2021. "The transmission of monetary policy under the microscope." *Journal of Political Economy*, 129(10): 2861–2904.

Jordà, Òscar. 2005. "Estimation and inference of impulse responses by local projections." *American economic review*, 95(1): 161–182.

Jordà, Òscar, Moritz Schularick, and Alan M Taylor. 2015. "Betting the house." *Journal of international economics*, 96: S2–S18.

Kaplan, Greg, Benjamin Moll, and Giovanni L Violante. 2018. "Monetary policy according to HANK." *American Economic Review*, 108(3): 697–743.

Koeniger, Winfried, Benedikt Lennartz, and Marc-Antoine Ramelet. 2022. "On the transmission of monetary policy to the housing market." *European Economic Review*, 145: 104107.

Kuttner, Kenneth N. 2014. "Low interest rates and housing bubbles: still no smoking gun." *The role of central banks in financial stability: How has it changed*, 30: 159–185.

Liebersohn, Jack, and Jesse Rothstein. 2025. "Household mobility and mortgage rate lock." *Journal of Financial Economics*, 164: 103973.

Luetticke, Ralph. 2021. "Transmission of monetary policy with heterogeneity in household portfolios." *American Economic Journal: Macroeconomics*, 13(2): 1–25.

Mabille, Pierre, Lu Liu, and Julia Fonseca. 2024. "Unlocking Mortgage Lock-In: Evidence From a Spatial Housing Ladder Model." *Available at SSRN 4874654*.

Miranda-Agrippino, Silvia, and Giovanni Ricco. 2021. "The transmission of monetary policy shocks." *American Economic Journal: Macroeconomics*, 13(3): 74–107.

Nakamura, Emi, and Jón Steinsson. 2018. "High-frequency identification of monetary non-neutrality: the information effect." *The Quarterly Journal of Economics*, 133(3): 1283–1330.

Olea, José Luis Montiel, and Carolin Pflueger. 2013. "A robust test for weak instruments." *Journal of Business & Economic Statistics*, 31(3): 358–369.

Quigley, John M. 1987. "Interest rate variations, mortgage prepayments and household mobility." *The Review of Economics and Statistics*, 636–643.

Ramey, Valerie A. 2016. "Macroeconomic shocks and their propagation." *Handbook of macroeconomics*, 2: 71–162.

Ringo, Daniel. 2024. "Monetary policy and home buying inequality." *Review of Economics and Statistics*, 1–46.

Rotberg, Shahar, and Joseph B Steinberg. 2024. "Mortgage interest deductions? Not a bad idea after all." *Journal of Monetary Economics*, 144: 103551.

Scharfstein, David, and Adi Sunderam. 2016. "Market power in mortgage lending and the transmission of monetary policy." *Unpublished working paper. Harvard University*, 2.

Staiger, Douglas, and James H Stock. 1997. "Instrumental Variables Regression with Weak Instruments." *Econometrica*, 65(3): 557–586.

Stock, James H, and Mark W Watson. 2018. "Identification and estimation of dynamic causal effects in macroeconomics using external instruments." *The Economic Journal*, 128(610): 917–948.

Swanson, Eric T. 2021. "Measuring the effects of federal reserve forward guidance and asset purchases on financial markets." *Journal of Monetary Economics*, 118: 32–53.

Verbrugge, Randal, and Robert Poole. 2010. "Explaining the rent–OER Inflation divergence, 1999–2007." *Real Estate Economics*, 38(4): 633–657.

Williams, John C, et al. 2015. "Measuring monetary policy's effect on house prices." *FRBSF Economic Letter*, 28: 1–6.

Appendix

A Data Sample Construction

We apply several filters to the raw Altos data to arrive at our final sample. Appendix Table A.1 summarizes the number of observations remaining after each of these filters. First, we drop observations where the address, the date, or the listed rent are missing. Second, we standardize street addresses and drop observations for which the standardized address is incomplete. The address standardization process involves standardizing street suffix abbreviations based on the United States Postal Service (USPS) abbreviations dictionary, standardizing house number suffixes, and truncating unit numbers.[12] Once street addresses have been standardized, we keep only addresses that have complete information on the street name and house number. Third, we drop addresses for which the geocodes are missing or for which the geocodes do not uniquely identify a street address in the data. Fourth, we exclude short-term and vacation rentals, commercial properties, mobile homes, listings of individual rooms, as well as listings for which the unit type is missing.

Fifth, we standardize the number of beds and number of baths. The standardization process involves assigning missing values to the number of beds (baths) in cases where the number of beds (baths) is larger than 10 or is not a multiple of 0.5 (0.25). We then collapse all listings that appear within the same week and have the same address, same number of beds and baths, and same listed rent to one observation. These cases likely reflect duplicate listings for the same unit across different listing platforms or within the same listing platform. While it might be that multiple different units within the same building that feature the same number of beds and baths are listed for rent during the same week, the fact that these units are listed for the same price suggests they are of the same quality. As such, they do not contribute differentially to the repeat-rent index.

Sixth, for each street address, we infer whether the structure is a single-family house or a a multi-family building. Namely, we categorize addresses as corresponding to single-family houses if (1) more than half of the listings associated with the address specify that the unit type is "house" and (2) we never observe multiple listings associated with the address within the same week that list different rents. These conditions suggest that there is only one housing unit in that address. Remaining addresses are classified as corresponding to multi-family units. We drop observations of multi-family buildings that have missing number of beds or missing number of baths.

Seventh, as discussed in Section 2, for the purpose of constructing a repeat-rent index, we identify units

[12] For some listings, the data does record the unit number, but this field is typically missing.

by their street address, number of beds and number of baths. The validity of our repeat-rent index as a quality-constant measure of rent growth therefore relies on units within the same building that have the same number of beds and baths also being of the same quality. we therefore drop units that likely differ in their quality but that are undistinguishable based on their address, number of beds and number of baths. We begin by identifying all tuples of address, number and beds and number of baths for which we observe, within a same week, multiple listings with different prices. We drop observations associated with these tuples. We then identify tuples of address, number and beds and number of baths for which the (unique) weekly rent often fluctuates between consecutive weeks, and drop all observations associated with tuples.[13]

Eight, we drop outliers. Namely, we drop listings with rents that are above the 97.5 percentile or below the 2.5 percentile of contract rents in the AHS.[14] We also drop tuples of address, number and beds and number of baths for which we observe a 4-week (52-week) rent fluctuation, in absolute value, that exceeds the 95th percentile (99th percentile) of the 4-week (52-week) rent fluctuation distribution in the data.

Ninth, we identify and drop listings that likely do not correspond to vacant units. In particular, we define a listing spell as the number of weeks a listing consecutively appears on the market without a break that is longer than 4 weeks. We then truncate the duration of spells where the listed rent is constant throughout the spell to the 90th percentile of spell durations throughout the sample. These cases likely reflect listings that were not taken off the market despite the underlying unit being rented.

Table A.1: Sample Selection

#	Filter	# Listings in Millions (% Raw)
0	-	443.8 (100.0)
1	Exclude Missing Address/Rent/Date	443.2 (99.9)
2	#1 + Exclude Incomplete Standardized Address	422.7 (95.4)
3	#2 + Exclude Missing or Non-Unique Geocodes	417.9 (94.3)
4	#3 + Exclude Short-Term/Mobile/Commercial/Rooms/Missing Type	417.2 (94.1)
5	#4 + Exclude Duplicates of {Week,Address,Beds,Baths,Price}	365.5 (82.5)
6	#5 + Exclude Multi-Family with Missing Beds/Baths	318.4 (71.8)
7	#6 + Exclude Units Non-Distinguishable Based on {Address,Beds,Baths}	140.2 (31.6)
8	#7 + Exclude Outliers (Extreme Rent/Rent Fluctuations)	120.4 (27.2)
9	#8 + Truncate Constant Rent Listings $>$ 90th percentile of Spell Duration	103.4 (23.3)

Note: This table reports the number and share of listings remaining after each data filter.

[13] Namely, we drop all observations associated with tuples of address, number of beds, and number of baths for which we observe, at least in two occasions throughout our sample period, the weekly rent fluctuating between two rents within a 5-week period.

[14] For each of the years 2015, 2017, 2019, and 2021, we compute the 97.5 and 2.5 percentile in the corresponding AHS survey, and drop all listings with rents that are above the 97.5 percentile or below the 2.5 percentile in that year. For all remaining years (in which AHS data is unavailable), we use the (deflated) percentiles computed from the most proximate AHS survey.

In addition to the filters we apply to the raw data, we construct a monthly panel of listed rents, identified at the unit level by collapsing all listings of the same unit that appear within the same month to one observation. Namely, we keep the latest observation within the month. Units that are not listed in more than one month are excluded, since they do not inform the repeat-rent index. We drop zip-codes-months in which less than 15 units are listed and zip-codes that are observed in less than 70 months throughout our sample. Our final panel data contains 30.3 million monthly observations of listed rents. It comprises 6.5 million rental units across 5,092 zip-codes. Each rental units is observed on average for approximately 4.7 months during the sample. The average time on market (i.e. the consecutive number of months a unit is being listed) is 2 months.

B Additional Figures and Tables

Figure B.1: ADH-RRI vs. ZORI: 2011-2015

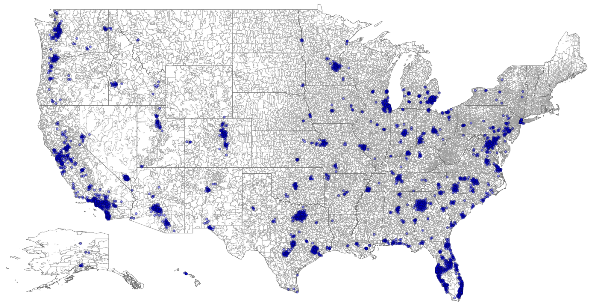

(a) ADH-RRI, 2011 (4,085 zip-codes)

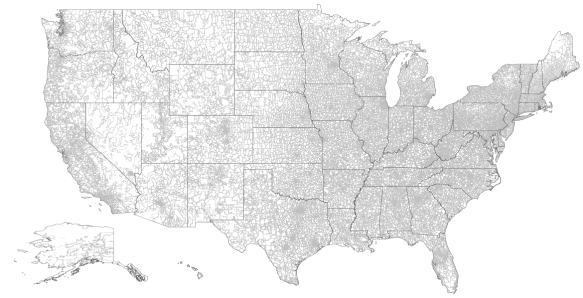

(b) ZORI, 2011 (0 zip-codes)

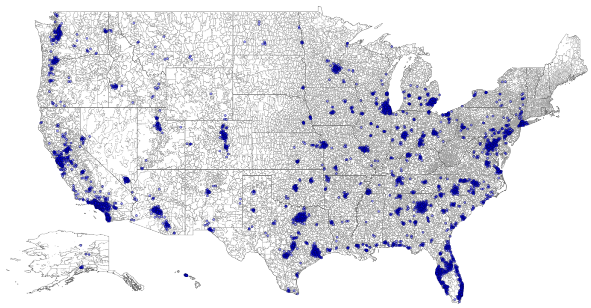

(c) ADH-RRI, 2015 (5,278 zip-codes)

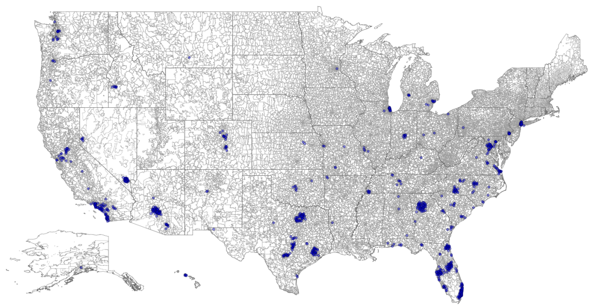

(d) ZORI, 2015 (1,004 zip-codes)

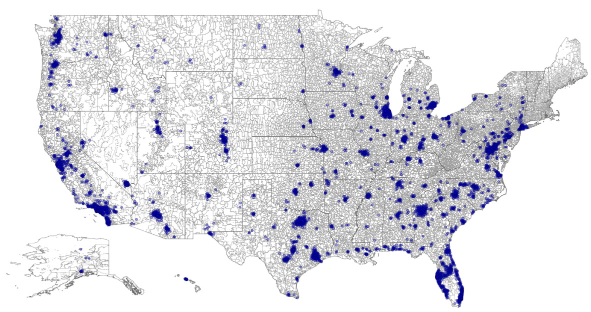

(e) ADH-RRI, 2019 (5,183 zip-codes)

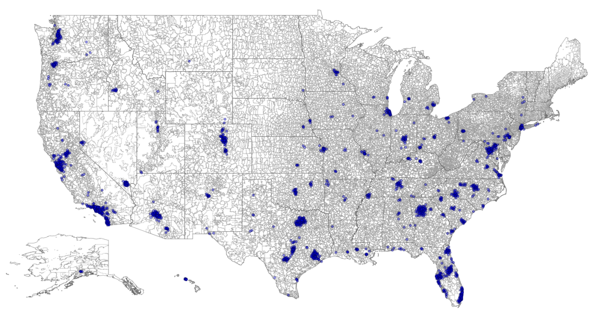

(f) ZORI, 2019 (2,308 zip-codes)

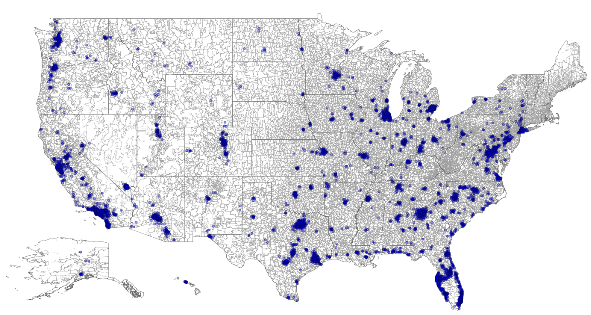

(g) ADH-RRI, 2022 (5,152 zip-codes)

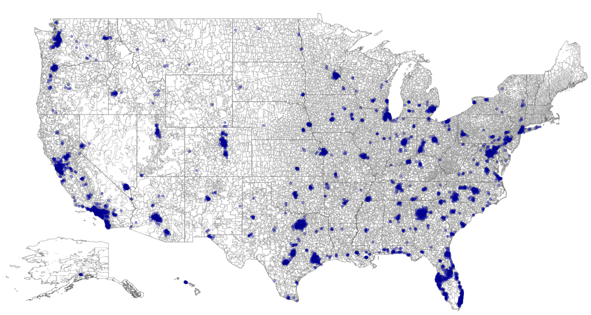

(h) ZORI, 2022 (4,246 zip-codes)

Note: This Figure displays, in blue, the zipcodes that are covered by the ADH-RRI and by ZORI in 2011, 2015, 2019 and 2022. The number of covered zipcodes is in parenthesis.

Figure B.2: First-Stage F-Statistic

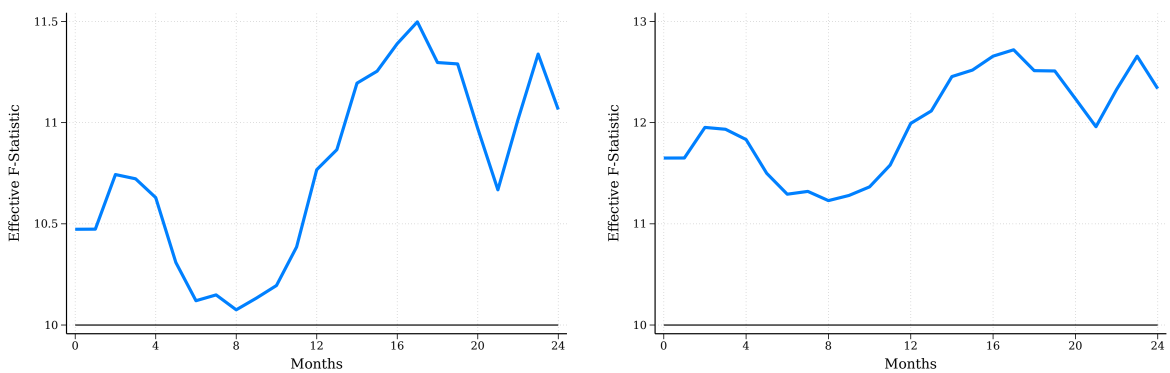

(a) Real Rent

(b) Nominal Rent

Note: This figure displays the Olea and Pflueger (2013) effective F-statistic of the first-stage of Equation 3 for the case where the outcome is real rent growth (Panel (a)) and for the case where the outcome is nominal rent growth (Panel (b)), for each horizon $h = \{0, 1, ..., 24\}$. The F-statistic may differ between Panel (a) and Panel (b) since the controls in Equation 3 include lags of the outcome variable.

Figure B.3: Alternative Monetary Policy Shocks

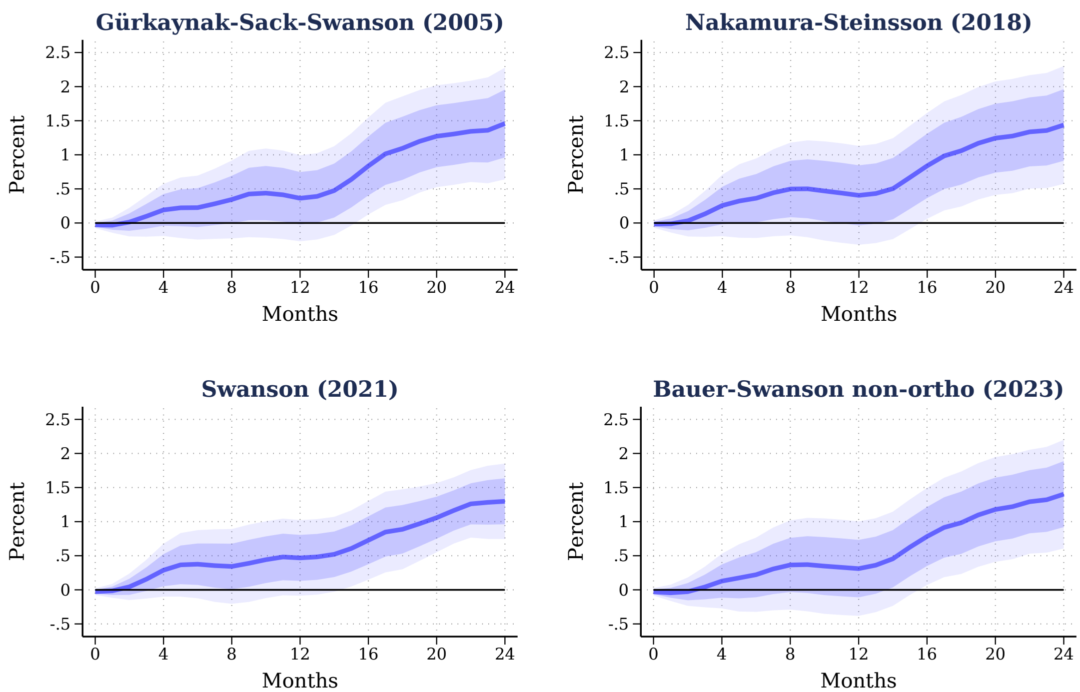

Note: This figure displays the impulse response function of real rent inflation to a 25bps increase in the 30-year fixed rate mortgage using alternative monetary policy shocks. The top left panel corresponds to the Gürkaynak, Sack and Swanson (2005) shocks, where we use both the surprise changes in the federal funds rate and the surprise changes in forward guidance as instruments for the monetary policy indicator. The top right panel corresponds to the Nakamura and Steinsson (2018) shocks. The bottom left panel corresponds to the Swanson (2021) shocks, where we use surprise changes in the federal funds rate, in forward guidance, and in large-scale asset purchases (LSAPs) as instruments for the monetary policy indicator. The bottom right panel corresponds to the non-orthogonalized shocks from Bauer and Swanson (2023*b*). Dark (light) shaded areas represent 68% (90%) confidence intervals. Standard errors are clustered at both the zip-code and month level and are estimated using the (Cameron, Gelbach and Miller, 2011) multi-way clustering estimation method.

Figure B.4: Effect of Monetary Policy on Cash and Non-Cash Sales

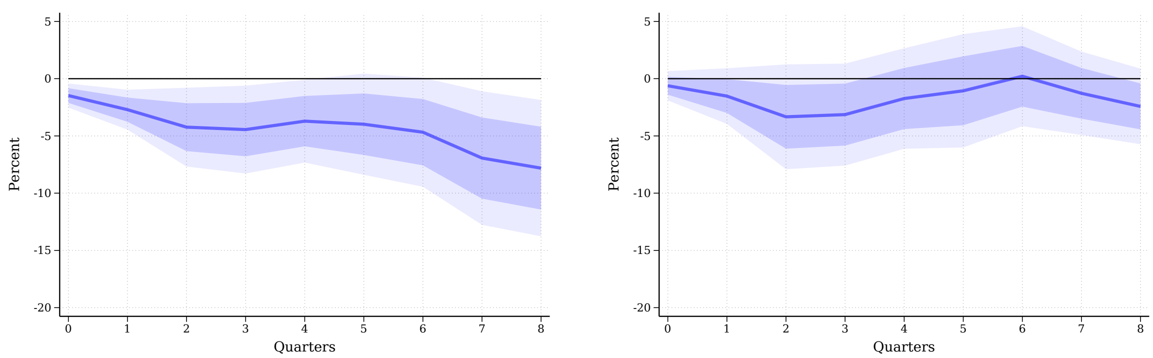

(a) Sales to Non-Cash Buyers

(b) Sales to Cash Buyers

Note: Panel (a) (Panel (b)) displays the impulse response function of the change in the volume of housing transactions where the buyer is an non-cash (cash) buyer to an exogenous 25 basis point increase in the 30-year fixed rate mortgage. Dark (light) shaded areas represent 68% (90%) confidence intervals. Standard errors are clustered at both the zip-code and month level and are estimated using the (Cameron, Gelbach and Miller, 2011) multi-way clustering estimation method.