THE GENDER GAP IN HOUSING RETURNS*

PAUL GOLDSMITH-PINKHAM

KELLY SHUE

October 2019

Abstract

Housing wealth represents the dominant form of savings for American households. Using detailed data on housing transactions across the United States since 1991, we find that single men earn one percentage point higher unlevered returns per year on housing investment relative to single women, with couples occupying the intermediate range. The gender gap grows significantly larger after adjusting for mortgage borrowing: men earn 6 percentage points higher levered returns per year relative to women. Data on repeat sales reveal that women buy the same property for approximately 2% more and sell for 2% less. The gender gap in housing returns varies by holding period, and arises because of gender differences in the location and timing of transactions, choice of initial listing price, and negotiated discount relative to the listing price. Gender differences in upgrade rates, preferences for housing characteristics, and listing agents appear to be less important factors. The gender gap varies with market tightness and demographic characteristics, but remains large in regions with high average education, income, and house price levels.

*Respective contact information and affiliations are: paul.goldsmith-pinkham@yale.edu, Yale School of Management; kelly.shue@yale.edu, Yale School of Management and NBER. We thank James Choi, Barry Nalebuff, and Matt Spiegel and seminar speakers at the University of Toronto for helpful comments. We thank Gen Li and Huijun Sun for excellent research assistance. This research was funded in part by the International Center for Finance at Yale SOM.

I. Introduction

Housing wealth accounts for the majority of most American households' wealth, with Americans investing more in the housing market than in the stock market.[1] Housing also differs from other common forms of household savings, such as bank deposits, bonds, and stocks, in that it is an illiquid and heterogeneous asset with prices determined through bilateral negotiation. Motivated by the existing research showing gender differences in financial sophistication, risk aversion, negotiation, and preferences (e.g., Sunden and Surette, 1998; Ayres, 1990; Babcock and Laschever, 2009; Sapienza et al., 2009; Reuben et al., 2015; Exley et al., 2016), we investigate how men and women differ in their financial returns on housing investment.

We use detailed data from CoreLogic covering over 50 million housing transactions and matched property listings across the US from 1991 to 2017. For approximately 9 million transactions for which we can identify homeowner gender, and the initial purchase and eventual sale prices, we estimate the homeowner's annualized realized return. We find that single men earn more than 1 percentage point higher unlevered annualized returns relative to single women. Couples underperform single women on a raw unadjusted basis, but outperform single women and underperform single men after adjusting for location and timing of sales.

Most U.S. home buyers purchase housing using mortgage debt with loan-to-value ratios of 80 percent or higher, and have not paid down a large fraction of the principal at the time of sale. Therefore, the real return earned is typically a levered return. We find that men outperform women by almost 6 percentage points per year after adjusting for leverage. The growth in the gender gap arises because leverage amplifies raw return differences.

The gender gap in housing returns exists in all years within our sample. It remains substantial in regions with high average education, income, and house price levels. However, the magnitude of the gender gap does vary with demographics, location, and time. Zip codes with lower average education, greater average age, and higher fraction single female are associated with larger gender gaps. Controlling for education and age, regions with higher median family income have larger gender gaps. The gender gap is largest in the right tail of the return distribution, although it also present at the median, and does not significantly reverse in the left tail.

Next, we show that approximately half of the gender gap in housing returns can explained by

[1] We plot the breakdown of housing and equity as a share of net worth in Appendix Figure A1 using data from the Survey of Consumer Finances.

gender differences in market timing, i.e., the choice of holding period, where and when to buy, and when to sell. Women earn lower returns on housing partly because they tend to buy when aggregate house prices are high and sell when they are low. The magnitude of the gender gap also varies with the business cycle, consistent with recent findings in Sakong (2019) showing a relation between cyclical housing transactions and wealth inequality.

In addition to market timing, we find that the gender gap arises from several contributing factors. We begin by examining data on repeat sales of the same property. Holding the property fixed, and adjusting for local time trends in prices, we find that women buy the same property for 1-2% more than men and sell for 2-3% less. This difference in transaction prices can be decomposed into gender differences in the choice of listing price and negotiated discounts relative to the listing price. Again using repeat sales data that allows us to hold the property fixed, we find that women purchase properties when they are listed at higher relative prices, and also choose to list for lower relative prices. In addition, women negotiate worse discounts relative to the listing price.

The gender gap varies with the match between the gender of the buyer and seller. Again exploiting repeat sales, we find that the highest transaction prices are associated with male sellers and female buyers, and the lowest transaction prices are associated with female sellers and male buyers. Female sellers and male buyers are associated with the largest negotiated discounts relative to the listing price, while male sellers and female buyers are associated with the smallest discounts.

Together, these findings relating to listing prices and negotiated discounts imply that women experience worse execution prices on real estate transactions at the points of purchase and sale. Differences in execution prices should matter much more for the annualized returns of short-term investors than for the returns of long-term buy-and-hold investors. Consistent with this insight, we find that the magnitude of the gender gap in annualized returns decays sharply with holding length. The gender gap in housing returns is greater for homeowners with shorter tenure in their properties, because they "trade" assets more often, so variation in execution prices matter more for their returns.

We also explore several other potential mechanisms which may drive the gender gap in housing returns. First, men may select properties with characteristics naturally associated with higher returns. In particular, men may purchase riskier properties, such that their higher return represents compensation for the additional risk. Second, men may invest more in housing maintenance and upgrades, such that their real investment return is lower than implied by analysis using only the sale price and purchase price (Harding et al., 2007). Third, women may be older, have more children, or have lower education and income, and these demographic factors may drive the gender gap in housing returns.

We find that differences in preferences for certain types of homes do not substantially affect the gender gap in housing returns. Men and women indeed differ in their choices of home type (e.g., new construction and number of bedrooms) and listing agent, but controlling for these features does not significantly affect the gender gap in housing returns. We also consider whether the higher returns earned by men may represent compensation for holding riskier properties. We do not find evidence of greater downside risk in housing returns for the men in our sample. However, we caution that we do not observe other outcomes such as bankruptcy that may differ by gender.

Then, using data from the American Housing Survey, we replicate our gender gap results in housing returns, and find that the gender gap remains large after controlling for age, education, ethnicity, number of children, and income. We also find no gender differences in reported home maintenance investment. In addition to routine maintenance, men may invest more in non-routine upgrades or renovations. Using CoreLogic listings data, we find that men are more likely to list homes that have been upgraded or renovated, but the difference in upgrade rates across genders is small. The gender gap in housing returns remains large in a restricted sample for which the house listing does not mention upgrades or renovations.

In general, gender differences in maintenance or property risk imply that the the gender gap in annualized housing returns should remain approximately constant with respect to holding period. Instead, we observe a gender gap in housing returns that decays toward zero with holding period. This pattern is more consistent with a gap in returns that arises from gender differences in execution prices.

Finally, we show that the gender gap in returns and prices narrows in markets when sale volume is high relative to the share of outstanding listings. This suggests that liquid housing markets suppress some of the drivers of the gender gap. In liquid markets, buyers and sellers may behave more like price takers, such that differences in negotiation style and preferences matter less. The fact that the gender gap narrows with market liquidity is also inconsistent with mechanisms of gender differences in risk or maintenance investment. Gender differences in home maintenance or risk should not lead to gender gaps in housing returns that vary strongly with market liquidity.

Our findings are related to existing research examining gender differences in stock market participation, portfolio allocation between stocks and bonds, and investment performance (e.g. Sunden and Surette, 1998; Bajtelsmit and VanDerhei, 1997; Hinz et al., 1997; Barber and Odean, 2001). We believe it is equally or more important to study gender differences in housing investment, given that housing represents a much larger proportion of the typical U.S. household's savings portfolio. Housing also differs from other common forms of household savings because prices are determined through bilateral negotiation. In contrast, we would expect men and women to earn the same return on an investment in a S&P500 index fund (holding timing constant), even if one group were more financially sophisticated or derived greater personal utility from owning the asset.

Our paper is also related to a large literature documenting gender differences in the ability, style, and willingness to negotiate. This literature has generally shown that women have more negative outcomes when negotiating in laboratory settings, as well as in labor market and automobile market settings (e.g., Ayres, 1990; Ayres and Siegelman, 1995; Castillo et al., 2013; Exley et al., 2016; Leibbrandt and List, 2014; List, 2004; Morton et al., 2003; Reuben et al., 2015). Our findings that women negotiate worse discounts relative to the listing price suggests that gender differences in negotiation contribute to the gap in housings returns. The fact that women choose to list homes at lower prices also suggests that gender differences in "first offers" within a negotiation framework may play an important role. Our housing setting is also interesting because personal interactions usually end after housing transactions are completed, so the gender gap cannot be explained by women placing greater value on continued relationships, a factor that might impact labor market negotiations (Babcock and Laschever, 2009).

However, we do not claim to rule out other explanations such as gender differences in preferences. For example, women may equal men in negotiation ability, but care more about purchasing a particular home or derive greater utility from a fast, low-risk, or non-confrontational negotiation process. It is not the goal of this paper to disentangle the role of negotiation ability from preferences. Our findings nevertheless show that variation in housing returns are large and may help explain the gender gap in wealth accumulation at retirement (see Appendix Figure A2 and Neelakantan and Chang (2010)).

Finally, our research question is related to contemporaneous work by Andersen et al. (2018) (AMNV), hereafter referred to as AMNV, which examines gender differences in negotiations for real estate transactions in Denmark. Our analysis differs from and complements AMNV in several ways. First, AMNV (along with related earlier research by Harding et al. (2003)[2]) is focused on demographic variation in negotiation and bargaining power, while we are interested in the total gender gap in housing returns from the perspective of wealth accumulation. We are interested in factors contributing to

[2] Harding et al. (2003) uses data from the American Housing Survey and structural modeling methods to estimate how bargaining power varies with demographics.

this gender gap beyond negotiation; for instance, we show that women earn lower returns partly due to market timing. Second, we differ from AMNV in our usage of data on both listing prices and final transaction prices, which allows us to examine negotiated discounts, as well as the matching between the genders of buyers and sellers. Third, we find large gender differences in transaction prices using repeat sales data in the US, while AMNV find smaller and insignificant differences using repeat sales data in Denmark. These results complement each other, and suggest that the gender gap may vary by country and culture.

II. Data and Measurements

In this section, we present a summary of the data sources for our analysis, describe the construction of key measures, including identification of gender, and summarize the overall data set.

A. Corelogic deeds and listings data

Our main housing transaction data comes from data gathered by CoreLogic from county deeds records. This data includes arms-length transactions (sales between two unaffiliated parties), non-arms-length transactions, and mortgage refinancings (non-transaction deed events). Each observation reflects a housing transaction, containing information on the date of the transaction, the sale price (if the property changed hands), the exact address of the property, and the names of both sides of the transaction. This last set of data fields allows us to partially identify the gender of the participant, as well as the number of participants on each side of the transaction (discussed further below).

To supplement the transactions data with time-varying measures of the properties' characteristics, we link the deeds dataset by property location to a dataset of property listings also constructed by CoreLogic.[3] These data come from Multiple Listing Service (MLS) systems operated by regional real estate boards. Each listing includes a large number of fields describing the property and the status of the listing. These include when the property is listed, the list price, and the listed property features such as the number of bedrooms and bathrooms and age of the structure. If the listing sells, we observe the close date and sale price.

B. Identification of gender and relationships

We identify the gender and family structure of the buyers and seller on each transaction using reported names on the deed. For each deed record, CoreLogic reports the full name of the first and second

[3]Properties are uniquely identified via parcel number (assigned by county deeds offices) and county code.

owner on a deed, and in the case of sale, the full name of the first and second seller. We identify two pieces of information from these name fields: first, we parse the fields to identify exactly how many parties exist on each side of the transaction, since in some cases, couples are transcribed as "John and Mary Smith" in one field, rather than being split across fields as "John Smith" and "Mary Smith." Second, we use the first names to probabilistically assign a gender to each party in the transaction. We follow Chari and Goldsmith-Pinkham (2017) and use data from Tang et al. (2011) to measure the probability that a given name is male or female (based on self-reported data to Facebook). Then, for all names with probability greater than 95%, we assign either male or female. For those who do not match, or whose probabilities are less than 95%, we treat as unknown genders.

Identification of the number of parties, and their respective genders, allows us to group parties into four groups (on each side of the transaction): single male, single female, couples (2 individuals with both genders identified), and other, where other is the residual category and will include single individuals without gender identified, couples where only one gender is identified, couples where neither gender is identified, and institutions. For each transaction, this grouping is done both on the buyer and seller side.

These measures of gender and relationships may be subject to measurement error. There are three types of concerns. First, we fail to identify gender for some individuals, and they will be relegated into the "Other" category. This is quite likely with non-Anglo-Saxon names where the gender is less predictable based on name. Second, we may miscategorize some individuals incorrectly by gender. Given our cutoff for gender is above 95% certainty, we are less concerned about this, but it is possible. Finally, some single men and women identified in our data may actually correspond to couples who choose (or follow local convention) in recording only a single name in a real estate transaction.

C. Measuring unlevered returns and levered returns

In our full dataset of transactions, we are able to identify consecutive arms-length market transactions for each property. We focus on realized returns, earned in the period from purchase to sale. Using these consecutive transactions, we can identify the unlevered annualized return for property i in sale year s: $r_{is} = \left(\frac{P_{is}-P_{ib}}{P_{ib}}\right)^{\frac{1}{(s-b)}} - 1$, where P_{ib} is the purchase price, i.e., the previous market transaction price on the property in year s. Because we observe exact dates for transactions, we allow years s and b to be non-integers to better measure the exact holding length of each property. To ensure that we correctly measure r_{is} for single male, single female, and couples, we focus on the subsample of returns that has three restrictions: (1) we have identified the gender and family structure in both periods b and s, (2) the gender and family structure of the buyer in period b corresponds to the gender and family structure in period s and, (3) names of the buyers in period b is sufficiently close to the names of the sellers period s by string matching distance. This final sample is used for our analysis of housing returns. These filters substantially restrict our analysis sample, since we need to observe multiple transactions and correctly identify gender and family structure, but ensures that we are not incorrectly measuring returns. Our final returns sample contains 9.3 million observations.

In reality, the majority of homeowners in the United States buy their homes using debt, with leverage of five-to-one or higher. Moreover, this leverage tends to persist over a long period of time, with long duration mortgages whose fixed amortization schedules pay mainly interest upfront. Therefore, the real return earned is typically a levered return. Ideally, given the mortgage type, term, interest rate and down payment, we could exactly identify the levered return. However, many of these fields are missing from the data. Instead, we provide several simple approximations to provide an estimate of levered returns.

We identify the initial downpayment D_{ib} and initial mortgage amount Mortgage_{ib} used to buy the home. We then calculate the average interest rate in the year-quarter of initial purchase by taking the 30-year fixed rate mortgage rate from Freddie Mac. Using this interest rate ρ_{ib}, we calculate the relative principal pay down at every monthly duration horizon, and use this to identify the share of remaining mortgage principal outstanding at period s (Mortgage_{is}). This allows us to calculate the total cash out payment for the home at the time of sale: $\text{Equity}_{is} = max\{P_{is} - \text{Mortgage}_{is}, 0\}$.[4] We then approximate the time b net present value of equity as the sum of the downpayment plus the discounted value of principal paydown payments: $\text{Equity}_{ib} \approx D_{ib} + \sum_{\tau=b}^{s} W_{i\tau}/(1+\rho_{ib})^{\tau-b}$. As a result, our levered annualized return is $r_{is}^{\text{lev}} = \left(\frac{\text{Equity}_{is} - \text{Equity}_{ib}}{\text{Equity}_{ib}}\right)^{\frac{1}{(s-b)}} - 1$. Note that in the case of a full cash purchase, $r_{it}^{\text{lev}} = r_{it}$.

The key ingredient for our estimation of levered returns is knowing the initial mortgage amount. However, in many cases the mortgage amount is missing. In recent years, we believe missing mortgage amount represents homes purchased with only cash and no mortgage was involved – in 2017, for example, 28.8% of all housing transactions were done with all cash, and in the 2000s this number was around 20%.[5] However, in the 1990s, the share of missing mortgage amounts in our data is substan-

[4]By using the max operator, we are implicitly assuming that homeowners cannot lower more than 100% of their original equity. Removing the max operator yields similar results.

[5]See the discussion in the Wall Street Journal here, for example: https://www.wsj.com/articles/want-that-

tially higher, and suggests that there may also be missing data issues. To address these data issues, we impute the mortgage amount in several ways. In our first case, we assume that mortgage amounts are zero in all transactions with missing mortgage amounts. In the second case, we assume that missing mortgage amounts are 80% of the overall housing value (the most common mortgage loan-to-value ratio). Finally, we provide a simple benchmark for converting all of the unlevered housing returns into levered returns by assuming all purchases were done with an 80% LTV mortgage. This holds fixed any potential leverage differences across transactions, and instead converts the unlevered returns into a measure that captures the modal degree of leverage that homeowners face.[6]

It is important to note that these calculations estimate realized returns conditional on sale. Since the choice to sell a property is an economic decision based on current market conditions, these returns do not represent returns in a counterfactual world in which all homeowners are forced to sell at a fixed horizon. Instead, these measure realized returns for those homeowners who choose to sell.

D. Description of data

In Table 1, we report the averages for a selection of our outcomes, both broken across our gender groups (single male, single female, couple and other) and pooled into overall averages. We are only able to successfully identify the gender and family structure for around 40% of the sample, or around 23 million sales. Of this sample, Couples have the largest share of transactions, followed by single men and then single women. We refrain from interpreting gender differences in raw sale prices in the summary statistics table, because men and women may purchase properties with different average quality. Our subsequent analysis will focus on housing returns or exploit repeat sales, which hold the property constant.

In Panel B, we report the averages for our sample of sales that are linked to listings data. This sample is substantially smaller and covers roughly 20 million sales. Slightly over 50% of our sample's gender group is identified. Unsurprisingly, this sample has higher prices, since the listings data coverage is later in the sample period. We measure purchase and sale discounts as (listing price - transaction price)/listing price $\times$ 100, so a larger purchase discount contributes to a higher return on housing investment and a larger sale discount contributes to a lower return on housing investment.

house-youd-better-pay-in-cash-1512469800

[6]When the mortgage amount implies an LTV of 100% or greater at purchase, the levered return is undefined. In other cases, initial LTV is below but very close to 100%, which can lead to extremely large estimates of levered returns. To reduce the influence of these outliers, we replace all imputed levered returns for observations with initial LTV greater than or equal to 100% or in the top 1% of levered returns with the hypothetical levered return the homeowner would have earned if the initial LTV had been equal to 80%. This procedure reduces the influence of outlier levered return estimates. The gender gap in levered returns is substantially larger and noisier if we do not use this procedure to deal with outliers.

Couples give the lowest sales discount, followed by single men, single women, and then the Other group. On average, all four groups have very similar average days on market.

Finally, in Panel C, we report the averages for our sample where we observe both the purchase and the sale of a property, and we are able to confirm that the seller is the same as the buyer. This restricted sample is smaller at roughly 9.3 million transactions. In this sample, we exclude the other group. We find that, on average, single men have the highest annualized unlevered return, at 8.47 percent, followed by single women at 6.92 percent, and couples at 6.47 percent. Single men hold their properties for half a year less time than single women, and 0.7 of a year less than couples. Single men also receive higher discounts when *purchasing* the home as well.

III. Empirical Results

This section describes our regression methodology and summarizes our main results measuring the difference in returns between single men, single women, and couples. We then assess the various channels that can explain this difference in returns.

A. Estimation approach

Our main analysis takes two forms. Both approaches use a simple linear regression framework to account for potential differences across gender. The first is an analysis of the unlevered and levered annualized returns:

$$r_{is} = \text{Single Female}_{is}\beta_1 + \text{Couple}_{is}\beta_2 + X_{is}\tau + \epsilon_{is}, \quad (1)$$

where Single Female_{is} is an indicator for a single female seller in period s and Couple_{is} is an indicator for a couple seller in period s in our main return sample. As a result, β_1 and β_2 capture the relative effect when compared to Single Male_{is}, the omitted category. X_{is} captures control variables, which most importantly include five-digit zip code interacted with sale-year-month fixed effects. This fixed effect will effectively compare two houses sold within the same year-month and zip code, and capture any unobservable differences based on location or the housing cycle. In some cases, we will report the *residualized* return, which will demean the returns using a regression of r_{it} on zip-by-sale-year-month fixed effects.

Our second set of analyses, focusing on the channels affecting housing return, is similar but uses alternative outcome measures, such as the log(Sale Price$_{it}$):

$$Y_{it} = \text{Single Female}_{it}\beta_1 + \text{Couple}_{it}\beta_2 + X_{it}\tau + \epsilon_{it}. \tag{2}$$

Since these outcomes are not measured in percent differences, we additionally include a property fixed effect in X_{it} to capture unobserved quality in the property that may be correlated with gender or family structure. To better estimate this property fixed effect, we include transactions that are not included in our returns data sample.

B. Baseline Results

We begin by showing how housing returns differ by the gender and relationship status of the homeowner. We use observations at the sale transaction level. The sample is restricted to observations for which the gender of all sellers can be identified, and for which we can match the identity of the seller at the time of sale to the identity of the buyer at the time of initial purchase. In column 1 of Table 2, we find that single women earn 1.6 percentage points lower unlevered annualized returns than single men (the omitted category). Controlling for zip-sale-year-month fixed effects (representing the location and time of sale) in column 2 and property holding length in column 3 shrinks the gender gap slightly. Women underperform men by 1.1 percentage points after adjusting for the location and timing of transactions.

We also find that couples underperform single women on an unadjusted basis. Column 1 shows that couples earn 0.4 percentage points lower unlevered returns than single women. However, the relative returns for couples is very sensitive to the inclusion of controls for zip-sale-year-month fixed effects. After controlling for these fixed effects in column 2 and property holding length in column 3, couples earn higher returns than single women but underperform single men. These results show that couples earn lower returns primarily due to poor market timing, but outperform single women holding the location and transaction period fixed. We examine how market timing impacts returns in more detail in the next subsection. Our findings are consistent with the possibility that couples face time constraints in real estate transactions due to child care and the school calendar system.

Because most home buyers purchase housing using loan-to-value ratios of 80 percent or higher, and have not paid down a large fraction of the principal at the time of sale, the real return earned is typically a levered return. Appendix Figure A6 plots the average loan-to-value (LTV) at the time of initial purchase for each gender group over time. In Panel A, we report the LTV for the sample with mortgage amount data, while in Panel B we report the share of transactions with missing mortgage data information. Single men have higher LTV conditional on non-missing mortgage data, but also have higher rates of missing mortgage data. Because missing mortgage data can represent cash purchases or true missing data, we do not draw strong conclusions about differences in leverage across gender.

In Table 3, we assess variation in housing returns after accounting for mortgage debt. In general, adjusting for leverage leads to a much larger gender gap in housing returns because leverage amplifies differences in raw returns. Column 1 uses a measure of levered returns in which we assume zero mortgage borrowing for observations with missing mortgage data while Column 2 assumes an 80% LTV for all observations with missing mortgage data.[7] Column 3 creates a hypothetical levered return for each observation assuming exactly 80% LTV. We use this latter measure of levered returns as our baseline measure in future tests because it is less sensitive to large outlier levered returns driven by a subset of households with very high LTV or 95% or greater. This measure also has the benefit of representing the expected gender gap for households with the modal LTV of 80% in the data. Using this measure, we find that women underperform men by 5.7 percentage points per year after adjusting for leverage. Couples earn returns in the intermediate range; they outperform women by 1.4 percentage points and underperform men by 4.3 percentage points.

The imputed measure of levered returns in columns 1 yields slightly smaller gender gaps, as expected because leverage amplifies gender differences in returns, and we assume in column 1 that LTVs are equal to zero for all observations with missing mortgage data. The large gap between couples and single men in column 2 results from the fact that couples are more likely to choose initial low initial LTV.

In addition to examining the gender gap in mean returns, we also compare the distribution of returns across gender groups, residualized by zip-sale-year-month fixed effects and with the average level of returns added back in. Figures 1 and 2 show unlevered and levered annualized returns at various percentiles of the return distribution for each gender group. This set of figures reveal that the gender gap exists in all parts of the return distribution except for the left tail where women and men fare equally poorly. However, the gender gap is larger at the 90th percentile of the returns distribution than at the median.

[7] In columns 1 and 2, we assume at 80% LTV for observations with zero downpayment or implied returns in the top percentile of the imputed levered return distribution. This procedure reduces the influence of outlier levered return estimates; the gender gap in levered returns is substantially larger and noisier if we do not apply this procedure.

In Figure 3, we show the density of realized returns. The figures again show that men weakly outperform women at all parts of the return distribution, with the largest differences in the right tail. In Figure 4, we zoom in to the return near zero, where all gender groups have distributions with missing mass just to the left of zero. This dip in the distribution is consistent with the disposition effect (see e.g., Shefrin and Statman, 1985), in which people are reluctant sell at less than their initial purchase price. Finally, Figures 1 through 3 show that men do not have worse left tail outcomes than women in terms of realized returns. Therefore, compensation for greater downside risk in realized returns cannot explain the higher average returns for men in our sample. However, we again caution that we don't observe other adverse outcomes such as personal bankruptcy costs that may differ by gender.

C. Heterogeneity and Timing

In Table 4, we explore how the average gender gap in housing returns within a zip code varies with zip-level demographics from the 2010 American Community Survey. We measure the gender gap in each zip code as the average difference between male and female residual returns (the residual return represents the return after adjusting for the zip-sale-year-month mean). We present the average gender gap across quartiles of various zip-level demographic characteristics. This table reports simple averages within each quartile of each demographic variable, without conditioning for other demographic variables. Observations are at the zip-code level and equally weighted. We find that the magnitude of the gender gap decreases with education, and increases with age, fraction black, and fraction single female, although the relations are not always linear. We also find that the gender gap remains large even in zip codes in the top quartile in terms of education, income, and house prices (measured relative to the state-year-month average).

In Table 5, we regress the zip-level gender gap on zip-level demographics. As before, we find that the gender gap in housing returns is significantly larger in zip codes with a greater fraction of residents with a high school or lower level of education, a greater fraction of elderly residents above the age of 60, and a greater fraction of single female households. Controlling for other demographics, the fraction of black residents within a zip code does not significantly predict the gender gap in returns. Finally, the gender gap increases with median family income controlling for these other zip-level demographic variables.

Figure 6 shows the magnitude of the average difference in unlevered returns between single men and women across all states in our sample.[8] The gender gap is positive in almost all states within our sample with good data coverage. We believe that variation in the gender gap across states could be partly caused by differences in data quality and estimation error across states. As noted earlier in Section B, some single men and women identified in our data may actually correspond to couples who choose (or follow local convention) in recording only a single name in a real estate transaction. To the extent that we estimate gender and couple status with error, we are likely underestimating the single male-female gender gap. If the degree of estimation error also varies across states, that could contribute to variation in the estimated gender gap across states.

Figures 5 and A5 show how the average and median realized returns varied over time and across gender groups. Average and median returns on housing are positive in all years of our sample, but display significant business cycle variation, with the highest returns in the run-up to the housing market crash in 2006. The magnitude of the gender gap in returns appears to increase with average returns, although the gender gap remains large in magnitude in recent years and does not exhibit a strong secular decline over time.

Appendix Figures A3 and A4 show how the composition of transactions by gender group varies over time. Changing composition combined with business cycle variation in average returns implies that gender differences in market timing can play a large role in the overall gender gap. In Table 6, we explore how much of the overall gender gap can be explained by differences in market timing, i.e., the choice of the exact month in which to buy and sell, as well as the holding length. As we move from column 1 to column 5, we introduce more detailed control variables for market timing, including zip-year-month fixed effects for the initial purchase transaction and zip-year-month fixed effects for the sale transaction. We also control for the interaction between year-month of purchase and year-month of sale fixed effects, which subsume the control variable for holding length. We find that market timing has a particularly large impact on the return gap between couples and single men. 80% of the original return gap can be explained by couples being relatively worse at market timing. Market timing also contributes to the male-female gender gap, albeit to a lesser extent. We find that approximately half of the gender gap in returns in Column 2 (our baseline specification, which already includes zip-year-month fixed effects for the sale transaction) can be explained by more detailed control variables for market timing. A potential explanation for these results is that frictions associated with childcare and school year cycles limit the ability of some couples and single mothers to advantageously time real

[8] We exclude states where the number of observations is less than 500.

estate transactions.

D. The Gender Gap in Execution Prices

So far, we have shown that the gender gap in housing returns can be partly explained by gender differences in the market timing of when the home is purchased and sold, as well as the overall holding period. In this section, we explore gender variation in transaction prices, listing prices, transaction discounts, and holding length.

The unlevered annualized return on housing depends mechanically on the ratio of the sale price to the initial purchase price, annualized to account for holding length. To assess gender variation in each transaction price, we exploit repeat sales data and control for zip-year-month fixed effects to account for time trends within a zip code. Each observation in this analysis is a transaction. To better estimate property fixed effects, we do not restrict the sample to buyers or sellers with identified genders and matched names across sales and initial purchase. All observations corresponding to non-identified single women, single men, and couples are included and coded as the "other" category. Thus, our sample size expands to cover over 50 million observations.

The results in Table 7 show that women purchase homes at approximately 1-2% higher prices than men, holding the property fixed and adjusting for local time trends in prices. Women also sell the same property for 2-3% less than men. Couples do worse than single women in terms of purchasing at higher prices, but also outperform women in terms of selling at higher prices.

We can also examine how transaction prices vary with the match between categories of sellers and buyers. In Figure 9, we plot the coefficients from a regression of log transaction price on the interaction of seller gender and buyer gender, controlling for property fixed effects and zipcode by sale-year-month fixed effects. The base category is male buyers and male sellers, and each estimate should be interpreted as the relative effect compared to that group. Among the four possible matches between male and female sellers and buyers, the highest transaction prices occur when there is a male seller and female buyer, and the lowest transaction prices occur when there is a female seller and male buyer.

While these results suggest that there may be strategic reasons to match with female buyers and sellers, we do not find strong evidence of unusual matching patterns. In Appendix Table A6, we examine how sellers and buyers of different genders and family structure match. We find that sellers and buyers of the same "type" (single male, single female, and couple) tend to match with themselves slightly more than we would expect from random matching.

The gender variation in transaction prices can be decomposed into gender variation in list prices at the time of purchase and sale and negotiated discounts relative to the listing price. For this analysis, we restrict the sample in Table 7 to observations that can be matched to MLS data on home listings, leading to approximately 20 million observations, 10 million of which correspond to repeat sales. Holding the property fixed and adjusting for local time trends in listed prices, we find in Table 8 that women choose to purchase the same property when it is listed for approximately 3% higher than when it is purchased by men, and then choose to list the same property for 1-2% less than when it is listed by men. Couples again fall somewhere between men and women after controlling for location and time. The gap between couples and single men in sale listing prices shrinks dramatically with the inclusion of zip-year-month fixed effects, again consistent with market timing being a major factor in the returns earned by couples.

Next, we examine how negotiated discounts vary by gender. We measure purchase and sale discounts as the percentage discount relative to the listing price, (listing price - transaction price)/listing price $\times$ 100, so a larger purchase discount contributes to a higher return on housing investment and a larger sale discount contributes to a lower return on housing investment. In Table 9, we find that women buyers purchase homes at a 0.26 percentage point lower discount relative to men. Women sellers also offer 0.11 percentage points greater discounts at sale. The gender gap in negotiated discounts relative to listing price further lowers women's relative return, because women also choose to purchase homes listed at higher prices and sell homes with lower listing prices.

The discounts negotiated by couple buyers and sellers also display interesting patterns. Couple buyers lie between single men and women in terms of purchase discounts. However, couples negotiate significantly lower discounts than even single men when selling properties. In combination with the earlier results on listing prices, we find that couple sellers list the same property at approximately equal prices set by single men, but are much less likely to agree to a significant discount relative to the chosen list price.

In Appendix Figure A8, we plot the full purchase and sale discount distributions. There is a large mass of discounts at exactly zero, with female buyers bunching the most at zero discount. When we condition on positive discounts, we see that for purchase discounts, female buyers have more mass on the left than male buyers. This behavior flips for sale discounts.

We can also examine how transaction discounts vary with the match between categories of sellers and buyers. In Figure 10, we plot the discount from a regression of discounts on the interaction of seller gender and buyer gender, including zipcode by year-month fixed effects. The base category is male buyers and male sellers, and each estimate should be interpreted as the relative effect compared to that group. Among the four possible matches between male and female sellers and buyers, the highest discount occurs when there is a female seller and male buyer, and the lowest discount occurs when there is a male seller and female buyer.

In Tables 7-9, we used the full data sample of transaction prices and listing prices, to better estimate property fixed effects from repeat sales data. The estimated coefficient for the single female indicator represents the gender gap in the dependent variable within a large data sample including home buyers and sellers that are not included in our housing returns sample (inclusion in the returns sample requires that the seller identity match the previous transaction's buyer identity). To isolate the gender gap in transaction prices, list prices, and discounts that correspond to observations in our returns sample, we present supplementary analysis in Appendix Tables A1-A3, in which the indicator variables for male, female, and couples are set equal to one only for observations in our returns sample. In order to preserve our ability to estimate property fixed effects, we categorize all other observations into the "other" category. Hence the male, female and couple dummies should only capture the effects for the subset of transactions we identify in our main returns sample, while the sample size remains the same. A caveat to this analysis is that we do not observe list prices and discounts for all observations in our returns sample, so these tables are not meant to present an exact decomposition of the gender gap in housing returns. Results remain qualitatively similar.

Given the gender differences in listing prices and discounts, one may wonder whether women benefit from less aggressive list pricing and transaction discount negotiation with faster transaction times. For approximately 2 million observations, we also observe the number of days on market between listing and sale resolution. In Table 10, we find that women purchase and sell homes with 3% shorter transaction periods relative to men. Couples have 4% shorter transaction periods than men for home purchases, but have 0.8% longer transaction periods for home sales. This pattern for couples matches the evidence in Table 9 for transaction discounts. Couples appear to value fast resolution when purchasing and are more reluctant to offer large discounts when selling. In column 3, we find that the gender gap in realized returns remains large after controlling for the days on market for purchases and sales. In Appendix Figure A9, we plot the full distribution of days on market for each gender group, and do not see substantial differences in the right tail.

The gender gap in listing prices and transaction discounts together imply that women experience worse execution prices on real estate transactions at the points of purchase and sale. Differences in execution prices should matter much more for the annualized returns of short term investors than for the returns of long term buy-and-hold investors. In a simple model in which women buy properties for δ fraction more and sell δ fraction less then men and hold for t years, we expect that women will earn $2\delta/t \times 100$ percentage points lower annualized unlevered returns than men.[9] In other words, the impact of a gender gap in execution prices on the gender gap in annualized returns should asymptote toward zero with holding length.

In Figures 7 and 8, we plot how the level of annualized housing returns and the gender gap in annualized returns varies with holding length in years. The level of annualized housing returns declines with holding length toward just under 5% per year, consistent with earlier results in Table 2. This pattern may occur due to selection, in which homeowners sell early only if they expect a high sale price that can cover transaction costs. More interestingly, Figure 8 Panel A shows that the gender gap (i.e., the difference in annualized returns for men vs. women) declines sharply with holding length toward zero, consistent with the intuition that differences in execution prices should matter more for the annualized returns of shorter term investors. The gender gap in housing returns is greater for homeowners with shorter tenure in their properties, because they "trade" assets more often, so any advantage or disadvantage in execution prices will matter more for their returns.

We also find that the gender gap in *transaction prices* does not decline toward zero with longer holding lengths. In Figure 8, Panels B and C, we find that women have significantly worse execution prices at purchase and sale for all holding length buckets in our sample. However, worse execution prices at the points of purchase and sale matter less for annualized returns on investment as holding period increases.

In our baseline analysis, we equally weighted each completed housing transaction (each transaction represents the annualized return from initial purchase to eventual sale by the same owner). We estimated the gender gap in housing returns, averaged across all transactions. We can instead weight each transaction by holding length, so that each year in which a property is held receives equal weight. This alternative weighting scheme would place, e.g., six times the weight on a observation for a home

[9] Let P_0 represent the market price of the property at the time of purchase. Suppose that the market value of the property grows by a fraction r each year. Suppose men buy and sell at the market price, so their annualized return equals r regardless of the holding period. Suppose that women buy properties for a fraction δ more and sell for δ less than the market price and hold for a period of t years. We can solve for the annualized return for women r_F such that $[(1+r_F)^t = (1-\delta) * P_0 * (1+r)^t]/[(1+\delta) * P_0]$. Solving for r_F after applying the approximation that $log(1+x) \approx x$ for x close to zero implies that $r_F = r - 2\delta/t$

held for six years relative to an observation for a home held for one year. We present results using this alternative weighting scheme for unlevered and levered returns in Appendix Tables A4 and A5, respectively. We find that the gender gap shrinks significantly: single women earn 0.4 percent lower unlevered returns and 1.5 percent lower levered returns relative to single men. The smaller estimated gender gap is consistent with the fact that gender differences in execution prices should matter less for annualized returns for investments with longer holding periods, and this alternative specification increases the weights on observations with longer holding periods.

E. Other Potential Mechanisms

In this section, we explore several other potential mechanisms which may drive the gender gap in housing returns. First, men may select properties with characteristics associated with higher returns. In particular, men may buy riskier homes, such that their higher return represents compensation for the additional risk. Second, men may invest more in housing maintenance and upgrades, such that their real investment return is lower than implied by analysis using only the sale price and purchase price (Harding et al., 2007). Third, listing agents may have a differential impact on the housing returns earned by women. Fourth, women may be older, have more children, or have lower education and income, and these demographic factors may drive the gender gap in housing returns.

In Tables 11 and 12, we find that gender is predictive of the types of properties held, e.g., square footage, number of bedrooms, whether it was new construction at the time of purchase, popularity of the listing agent, etc. Further, some of these characteristics are predictive of housing returns. However, controlling for detailed home characteristics does not have a large impact on the estimated magnitude of the gender gap in returns (as evidenced by the small difference in coefficients on single female between columns 2 and 3 in Table 11). This analysis shows that women do not sort on average toward a set of housing characteristics that are associated with lower returns.

The results in Tables 11 and 12 also help address another potential explanation. Men may invest more in housing maintenance and upgrades. In particular, men may be more likely than women to purchase fixer uppers, which would explain why men buy at low prices and sell at high prices, holding the property fixed. Given that maintenance and upgrades are costly, men's real investment return may be lower than implied by analysis using only the sale price and purchase price. In Table 12, we find that men are more likely to list homes that have been upgraded or renovated, but the difference in upgrade rates across genders is small. In Table 11, we find that the gender gap remains large after controlling for whether the house has been upgraded or renovated, and in Appendix Table A7, we find that the gender gap in housing returns remains large in a restricted sample for which the house listing does not mention any synonyms for upgrades, renovations, new features, expansions, etc.

Aside from upgrades and renovations that are noted in property listings, men may also invest more in routine maintenance. In Table 13, we use data from the American Housing Survey and find insignificant and close to zero gender differences in home maintenance investment. We are also able to replicate our baseline results of a gender gap in housing returns using this alternate data source. Interestingly, the gender gap in realized returns in Column 3 exceeds the gender gap in estimated returns in Column 2, which uses self-reported estimates of current property values to calculate returns. This comparison suggests that women may underestimate their investment return disadvantage in housing markets.

The empirical patterns in Figure 8 are also inconsistent with a simple model in which the gender gap in housing returns arises *only* from gender differences in home maintenance and property risk characteristics. If men invest more in maintenance each year and purchase riskier properties, then the gender gap in annualized housing returns should not decay toward zero as holding period increases. Empirically, we do observe a gender gap in housing returns that decays toward zero with holding period. This pattern is more consistent with a gap in returns that arises from gender differences in execution prices at the points of purchase and sale (as discussed previously, differences in execution prices of δ predicts that the gender gap equals $2\delta/t$ which approximately matches the shape of the decay in the data).

In Table 13 column 4, we use data from the American Housing Survey to explore whether demographic factors correlated with gender may help explain the gender gap in housing returns. Women outlive men on average, and older individuals may earn worse returns on housing for reasons unrelated to gender.[10] Women may also have more children, lower education, and lower income, factors that may impact housing returns. We find that some of these demographic variables do indeed predict housing returns. However, the gender gap remains approximately constant after controlling for these demographic variables. The results also indicate that the housing disadvantage associated with being female is equivalent to the disadvantage associated with having three additional children.

[10]After the death of a spouse, widows or widowers may sell homes at a discount for a variety of reasons. Such cases are excluded from our returns analysis because we require that the homeowner be single at the time of home purchase as well as the time of home sale in order to be classified as single male or single female. We also restrict our sample to arms-length transactions, which excludes transfers to family members.

The gender gap in housing returns could also arise because of differences in how men and women interact with real estate agents. Housing transactions are typically intermediated by buyers' and sellers' agents who may suffer from agency conflicts. In particular, these agents may prefer higher transaction volume over higher returns for their clients, and men and women may differ in the extent to which they follow the advice of their agents. We unfortunately lack detailed data to explore this potential channel and leave the question of how agents impact the gender gap in housing returns to future research. However, we do find suggestive evidence in column 3 of Table 11 that more popular listing agents are associated with lower returns for home sellers, consistent with the idea that agency conflicts among agents may impact returns.

Finally, we explore how the gender gap varies with market tightness. In general, gender differences in preferences and negotiation style may play a large role in housing markets because these markets are illiquid and decentralized with heterogeneous assets. Thicker markets, where there are many more transactions occurring, may reduce the gender gap by reducing the scope for negotiation and by making buyers and sellers more like price takers. To measure market tightness, we use our listings data to construct a county-by-month measure of the number of sales in a given month scaled by the total number of listings. The larger this number, the more demand for transactions for the given listings. We then replicate our analyses from Table 2, 7 and 9, interacting our measures of gender with market tightness. We find that the gender gap shrinks substantially with market tightness. As market tightness increases, we see that the unlevered return gap between single female sellers and single male buyers shrinks and would completely offset when market tightness is around 1. We find similar offsetting interaction coefficients for transaction prices and discounts. We do not find the offsetting benefit of market tightness for couples in unlevered annual returns, but do find it in the discounts and prices.

Our finding that the gender gap narrows with market liquidity is also inconsistent with alternative explanations relating to gender differences in property risk, characteristics, upgrades, or maintenance. Differences in maintenance or risk should lead to real differences in home value appreciation. They should not lead to gender gaps in housing returns that vary strongly with market liquidity.

IV. Conclusion

We uncover a large gender gap in the returns to housing investment in recent decades in the US. This gender gap is likely to be an important contributor to gender differences in wealth accumulation and welfare, given that housing wealth represents the dominant form of savings for most US households. Using detailed data on housing transactions across the US, we find that single men earn one percentage point higher unlevered returns per year on housing investment relative to single women, with couples occupying the intermediate range. However, the real return on housing earned by most households is a levered return. The gender gap in raw returns grows significantly larger after adjusting for mortgage borrowing. Assuming the modal 80% LTV, men earn approximately 6 percentage points higher levered returns per year relative to women. Using data on repeat sales, we show that women buy the same property for approximately 2% more and sell for 2% less. The gender gap in housing returns arises because of gender differences in the location and timing of transactions, choice of initial listing price, negotiated discount relative to the list price. While the gender gap varies with demographic characteristics, it remains substantial in regions with high average education, income, and house price levels. It also has not displayed a secular decline over time.

In addition to gender gaps resulting from market timing, we find that women experience significantly more negative negotiation-related outcomes in housing markets. Women negotiate smaller discounts relative to the listing price when buying and offer larger discounts when selling. Holding the property fixed, the highest transaction prices occur in cases with a male seller and female buyer, and the lowest transaction prices occur in cases with a female seller and male buyer. However, we caution that these results do not necessarily imply that women make mistakes in housing negotiations. In particular, recent research by Exley et al. (2016) suggests that women can experience even more negative outcomes by "leaning-in" and negotiating more. Given the importance of housing investment for household savings, we believe that further exploration of factors that determine the gender gap in housing is an important direction for future research.

References

Andersen, Steffen, Julie Marx, Kasper Meisner Nielsen, and Lise Vesterlund, "Gender Differences in Negotiation: Evidence from Real Estate Transactions," *Working Paper*, 2018.

Ayres, Ian, "Fair driving: Gender and race discrimination in retail car negotiations," *Harvard Law Review*, 1990, *104*, 817.

_ and Peter Siegelman, "Race and gender discrimination in bargaining for a new car," *American Economic Review*, 1995, pp. 304–321.

Babcock, Linda and Sara Laschever, *Women don't ask: Negotiation and the gender divide*, Princeton University Press, 2009.

Bajtelsmit, Vickie L and Jack L VanDerhei, "Risk aversion and pension investment choices," *Positioning Pensions for the Twenty-First Century*, 1997, *45*, 66.

Barber, Brad M and Terrance Odean, "Boys will be boys: Gender, overconfidence, and common stock investment," *The Quarterly Journal of Economics*, 2001, *116* (1), 261–292.

Castillo, Marco, Ragan Petrie, Maximo Torero, and Lise Vesterlund, "Gender differences in bargaining outcomes: A field experiment on discrimination," *Journal of Public Economics*, 2013, *99*, 35–48.

Chari, Anusha and Paul Goldsmith-Pinkham, "Gender representation in economics across topics and time: Evidence from the NBER summer institute," Technical Report, National Bureau of Economic Research 2017.

Exley, Christine Linman, Muriel Niederle, and Lise Vesterlund, "Knowing when to ask: The cost of leaning-in," *Working Paper*, 2016.

Harding, John P, Stuart S Rosenthal, and CF Sirmans, "Depreciation of housing capital, maintenance, and house price inflation: Estimates from a repeat sales model," *Journal of Urban Economics*, 2007, *61* (2), 193–217.

_ , _ , and Clemon F Sirmans, "Estimating bargaining power in the market for existing homes," *Review of Economics and Statistics*, 2003, *85* (1), 178–188.

Hinz, Richard P, David D McCarthy, and John A Turner, "Are women conservative investors? Gender differences in participant-directed pension investments," *Positioning Pensions for the Twenty-First Century*, 1997, *91*, 103.

Leibbrandt, Andreas and John A List, "Do women avoid salary negotiations? Evidence from a large-scale natural field experiment," *Management Science*, 2014, *61* (9), 2016–2024.

List, John A, "The nature and extent of discrimination in the marketplace: Evidence from the field," *Quarterly Journal of Economics*, 2004, *119* (1), 49–89.

Morton, Fiona Scott, Florian Zettelmeyer, and Jorge Silva-Risso, "Consumer information and discrimination: Does the internet affect the pricing of new cars to women and minorities?," *Quantitative Marketing and Economics*, 2003, *1* (1), 65–92.

Neelakantan, Urvi and Yunhee Chang, "Gender differences in wealth at retirement," *American Economic Review*, 2010, *100* (2), 362–67.

Reuben, Ernesto, Paola Sapienza, and Luigi Zingales, "Taste for competition and the gender gap among young business professionals," Technical Report, National Bureau of Economic Research 2015.

Sakong, Jung, "Cyclical Housing Transactions and Wealth Inequality," *Working Paper*, 2019.

Sapienza, Paola, Luigi Zingales, and Dario Maestripieri, "Gender differences in financial risk aversion and career choices are affected by testosterone," *Proceedings of the National Academy of Sciences*, 2009, *106* (36), 15268–15273.

Shefrin, Hersh and Meir Statman, "The disposition to sell winners too early and ride losers too long: Theory and evidence," *Journal of Finance*, 1985, *40* (3), 777–790.

Sunden, Annika E and Brian J Surette, "Gender differences in the allocation of assets in retirement savings plans," *American Economic Review*, 1998, *88* (2), 207–211.

Tang, Cong, Keith Ross, Nitesh Saxena, and Ruichuan Chen, "What's in a name: A study of names, gender inference, and gender behavior in facebook," in "International Conference on Database Systems for Advanced Applications" Springer 2011, pp. 344–356.

Figure 1: Distribution of unlevered annualized returns by gender group

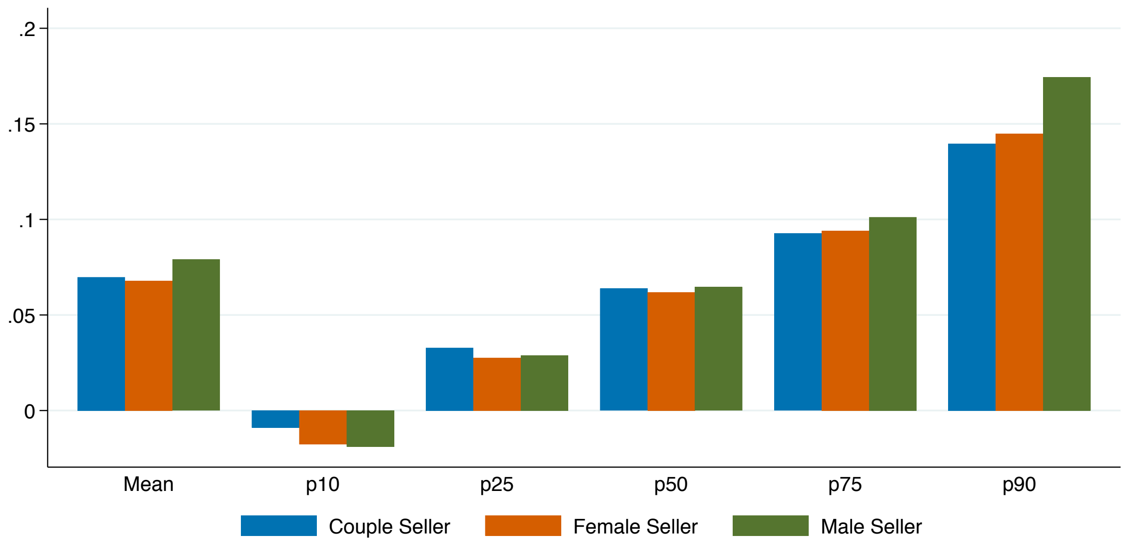

Note: This figure plots summary statistics for the residualized annualized returns for housing transactions by three gender groups: couples, single women, and single men. The residualization is achieved by regressing the returns on zip-by-sale-year-month fixed effects, and taking the residuals. Then, the overall mean is added back to each residual. The first column in each set is for couples, the second is for single women, and the third is for single men. See Section II.B for more details on the definition of gender and family structure.

Figure 2: Distribution of levered annualized returns by gender group

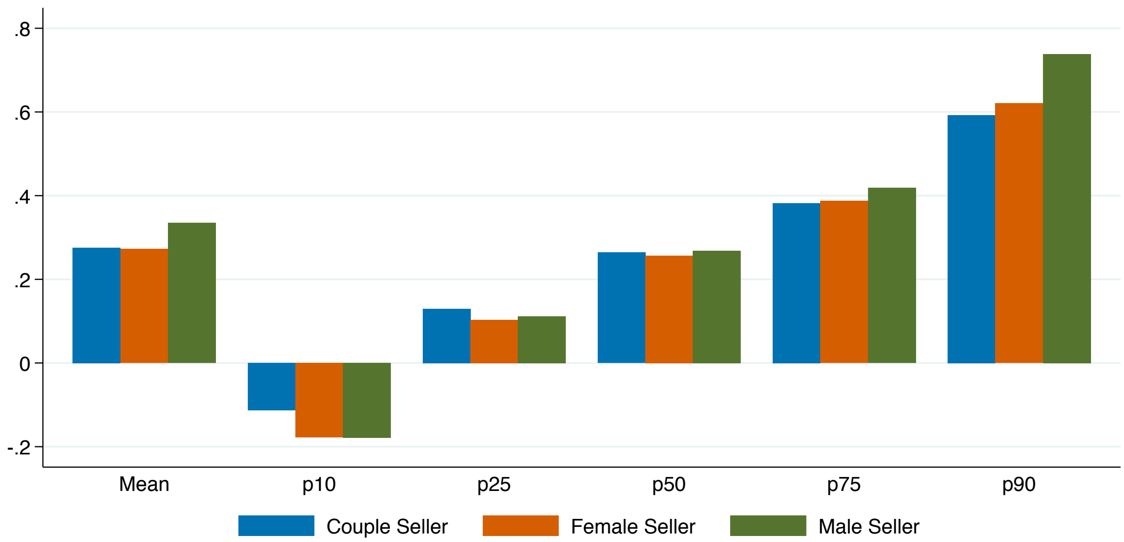

Note: This figure plots summary statistics for the residualized annualized levered returns for housing transactions by three gender groups: couples, single women, and single men. The residualization is achieved by regressing the levered returns, on zip-by-sale-year-month fixed effects, and taking the residuals. Then, the overall mean is added back to each residual. The levered returns are calculated following the formula is Section C and assuming an initial LTV of 80%. The first column in each set is for couples, the second is for single women, and the third is for single men. See Section II.B for more details on the definition of gender and family structure.

Figure 3: Density of unlevered returns by gender group

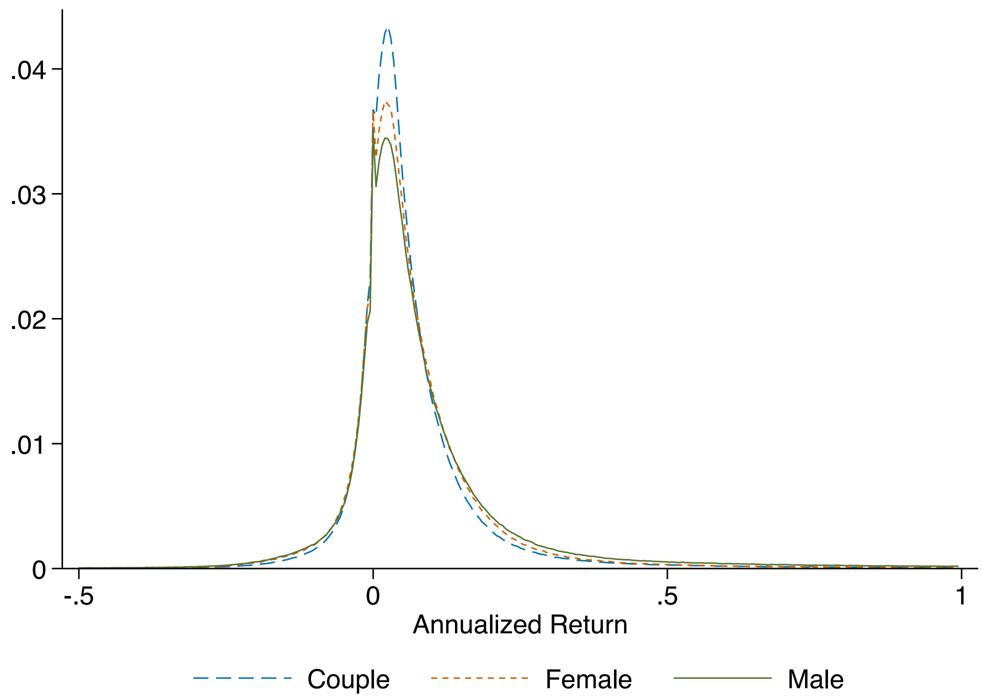

Note: This figure plots the density of the annualized unlevered returns for housing transactions by three gender groups: couples, single women, and single men. Returns are truncated at -50% and +100% for the purposes of this figure. See Section II.B for more details on the definition of gender and family structure.

Figure 4: Loss aversion and unlevered returns by gender group

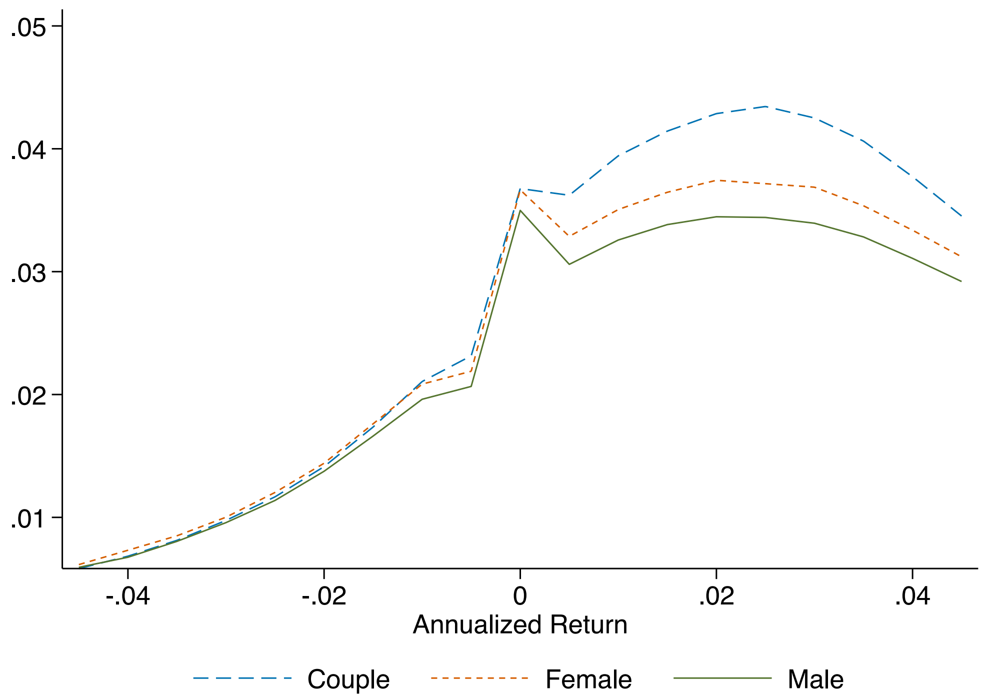

Note: This figure plots the density of the unlevered annualized returns for housing transactions by three gender groups: couples, single women, and single men. Returns are truncated at -4% and +5% for the purposes of this figure. See Section II.B for more details on the definition of gender and family structure.

Figure 5: Average unlevered returns over time by gender group

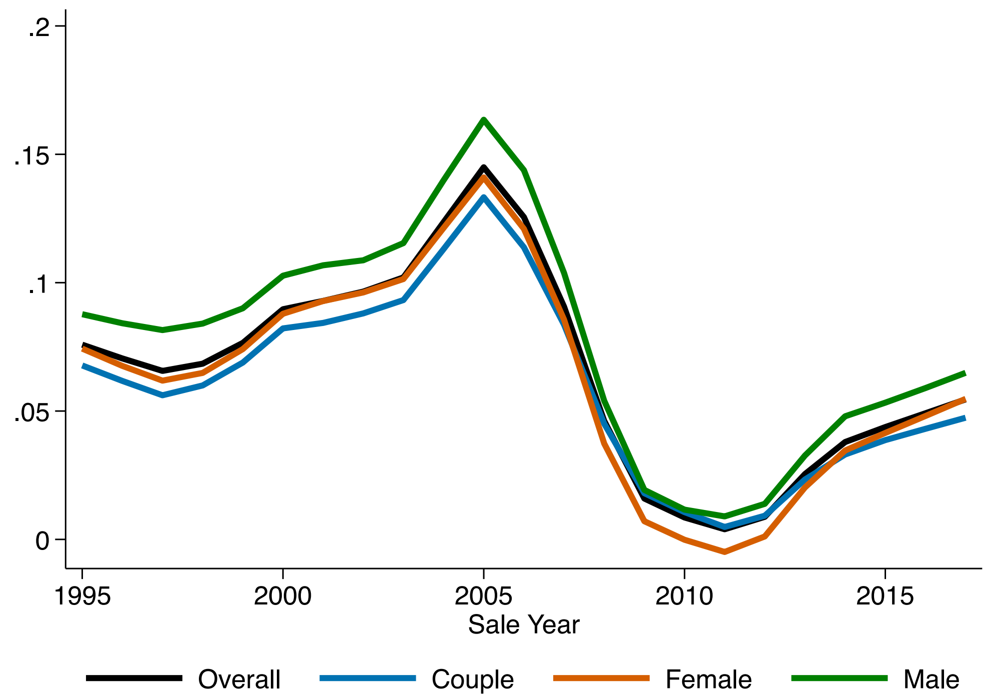

Note: This figure plots the average unlevered annualized return for couples, single women, and single men by sale year. As our sample begins in 1991, we begin this figure in 1995 to allow for sufficient data to avoid truncation. See Section II.B for more details on the definition of gender and family structure.

Figure 6: Difference in unlevered returns between single men and women across states

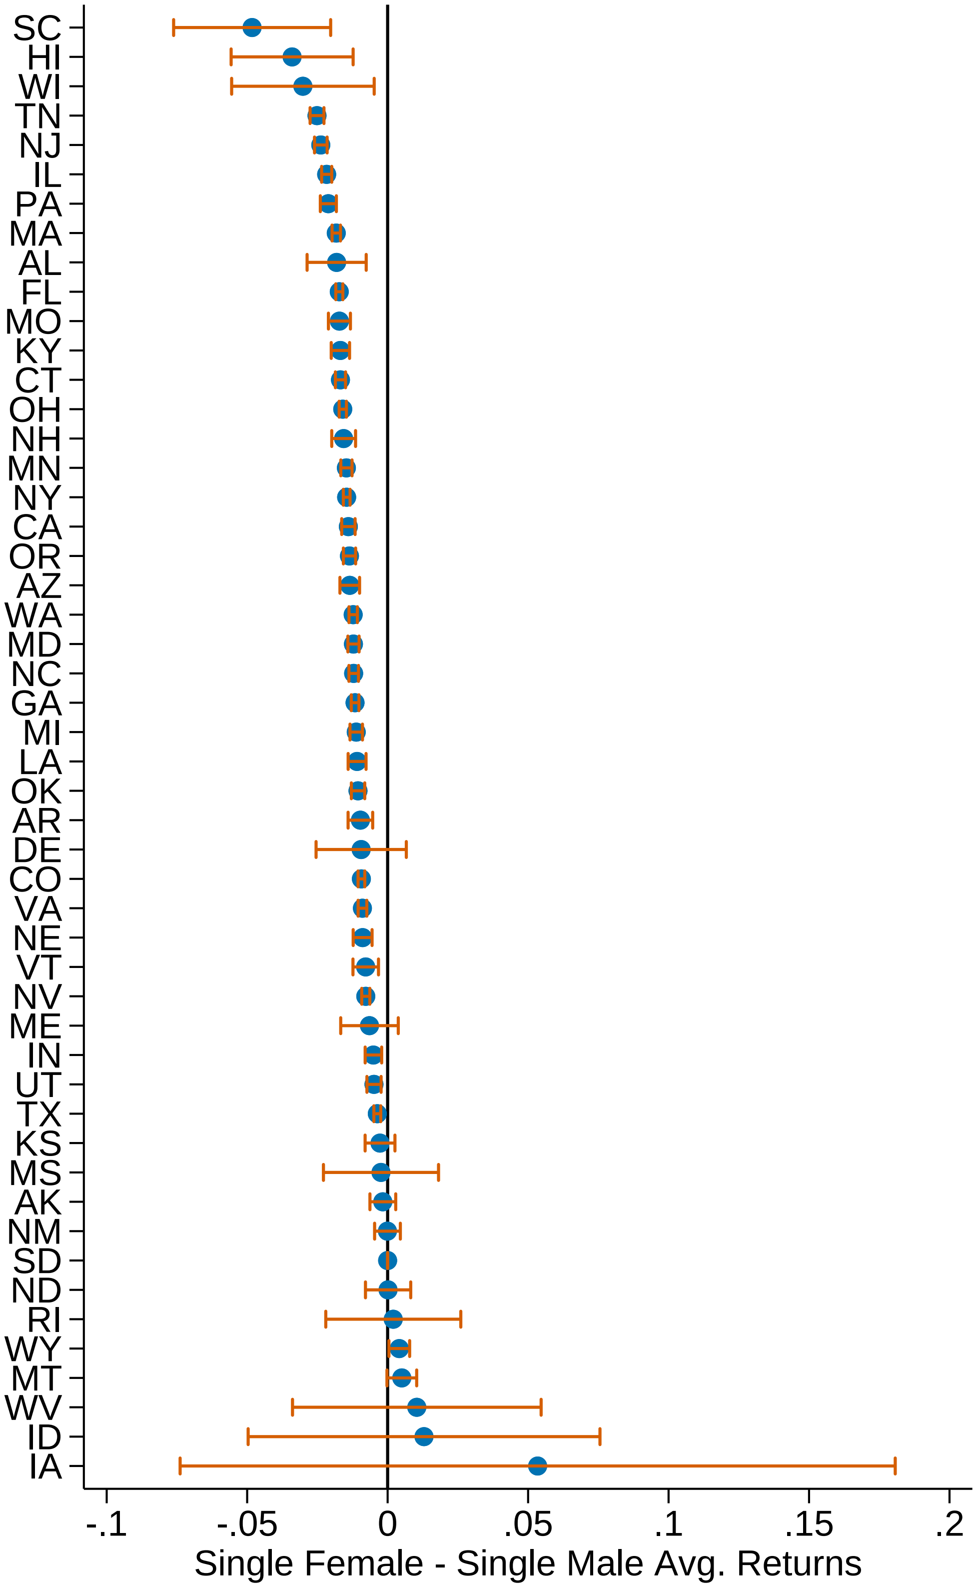

Note: This figure plots the average difference in unlevered annualized returns between single men and women across states, after controlling for zip-by-sale-year-month fixed effects. The points represent the estimated difference in realized returns, while the bars represent the 95% confidence interval for each effect. Standard errors are clustered at the zip-code level. See Section II.B for more details on the definition of gender and family structure.

Figure 7: Unlevered returns by holding period

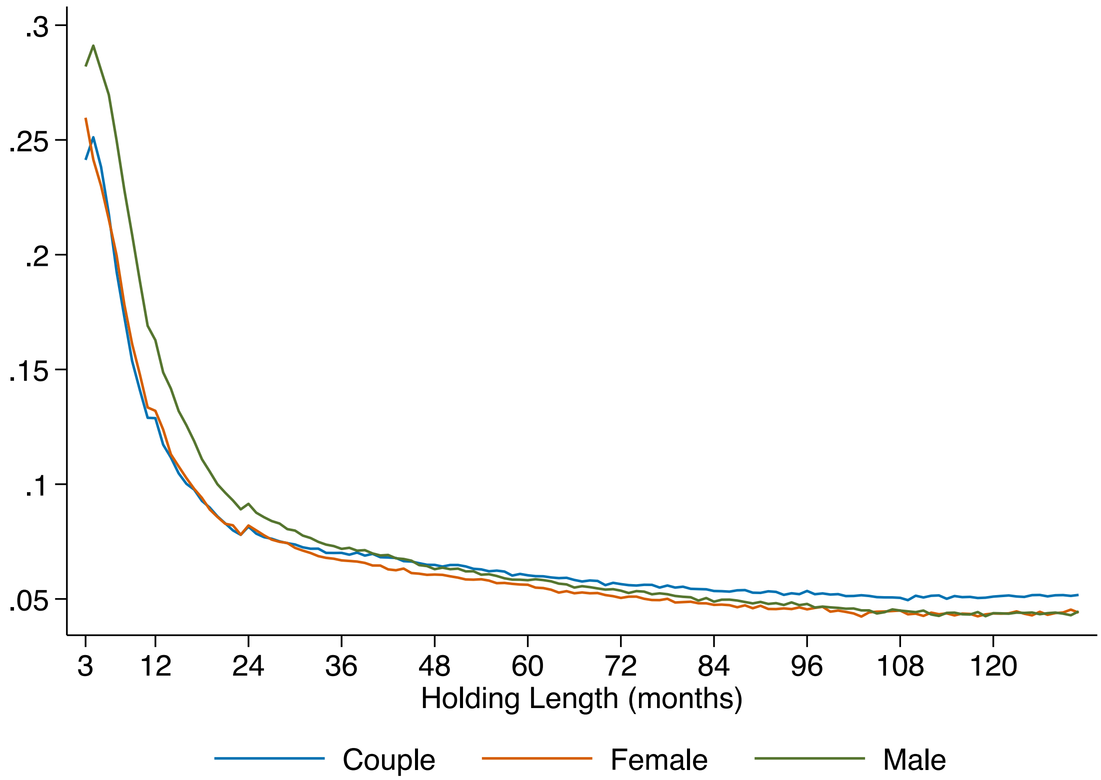

Note: This figure plots the average residualized unlevered annualized returns for couples, single women, and single men by holding period. We exclude holding periods longer than 11 years. See Section II.B for more details on the definition of gender and family structure. See Appendix Figure A7 for the relative distribution of transactions across holding lengths.

Figure 8: Differences between single men and women by holding period

Panel A: Gender gap in unlevered annualized returns

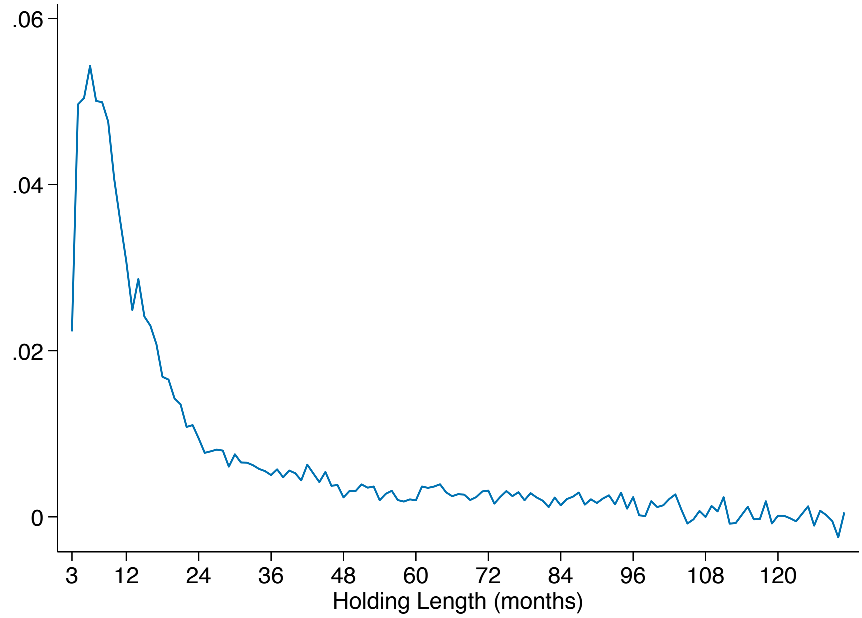

Panel B: Gender gap in purchase price

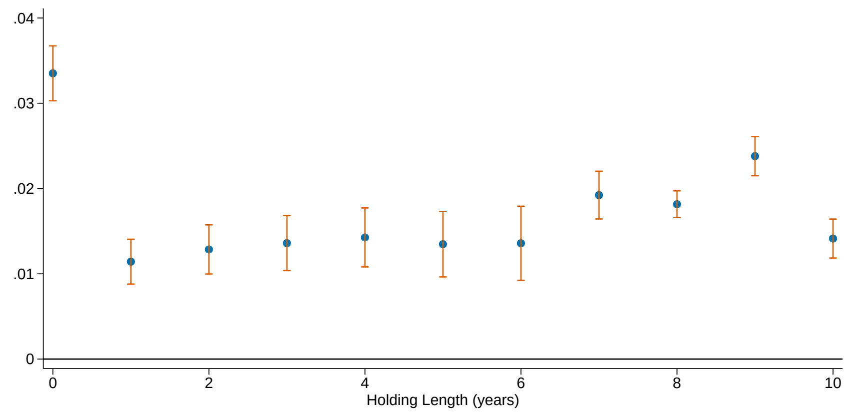

Panel C: Gender gap in sale price

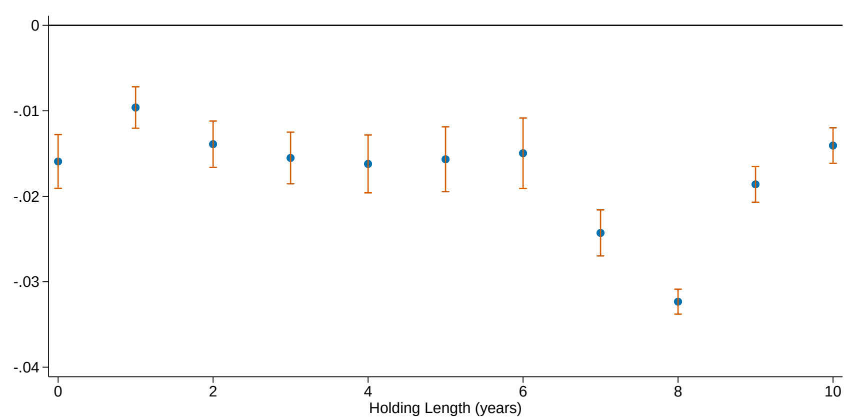

Note: This figure plots the average difference in outcomes between single men and women by holding length for the property. We exclude holding periods longer than 11 years. Panel A plots the male minus female gender gap in returns controlling for zip-sale-year-month fixed effects. Panels B and C plot the female minus male gender gap in log purchase and sale prices, respectively, controlling for property fixed effects and zip-year-month fixed effects.

Figure 9: Differences in sale price between single men and women, split by buyers and seller

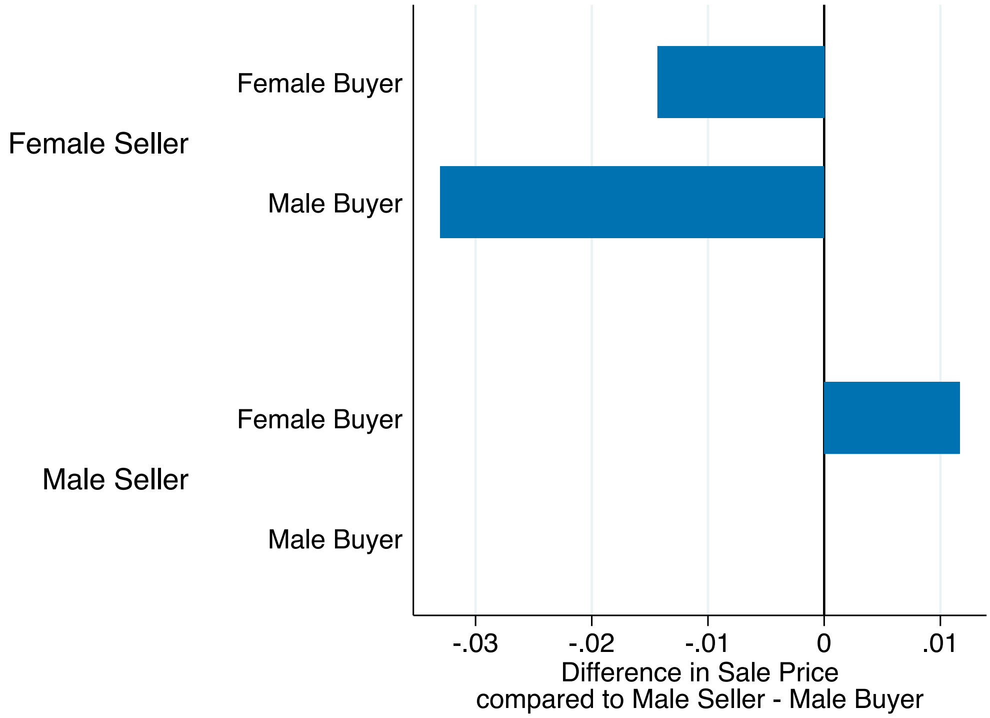

Note: This figure plots the average relative difference in log sale price between single men and women, split by seller and buyer gender. These estimates come from a regression of the form in Table 7 column 4, but allowing for the buyer and seller gender to interact. We plot only the coefficients representing single male or female buyers and sellers, with the male buyer-male seller as the base coefficient. See Section II.B for more details on the definition of gender and family structure.

Figure 10: Differences in discount between single men and women, split by buyers and seller

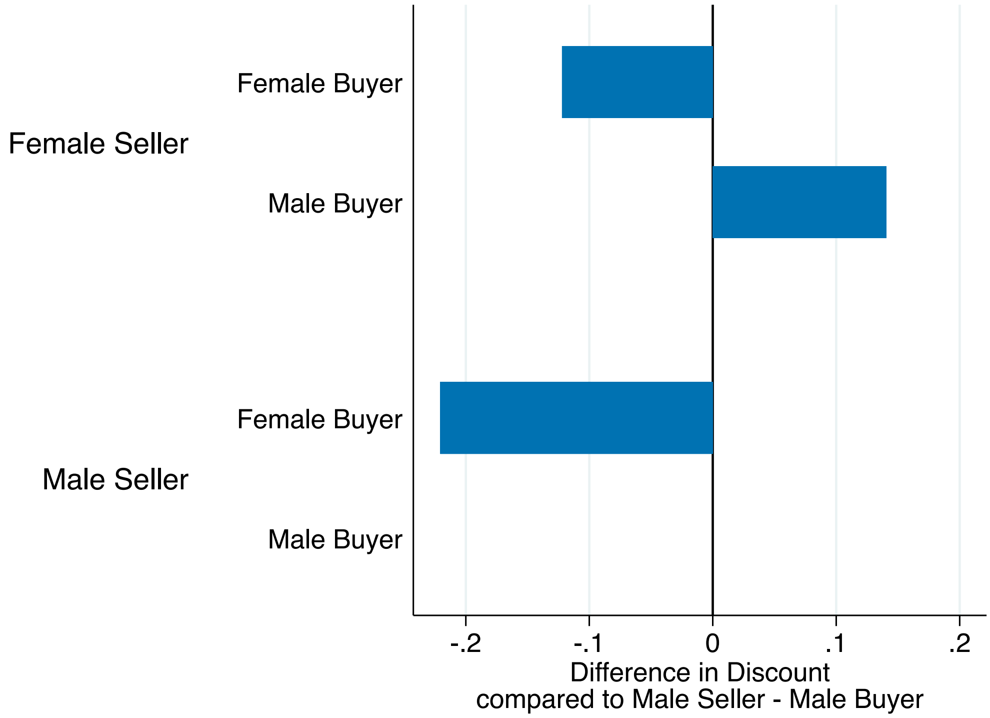

Note: This figure plots the average relative difference in the discount between single men and women, split by seller and buyer gender. We measure discounts as (listing price - transaction price)/listing price $\times$ 100, so a larger discount contributes to a higher return on housing investment for buyers and a lower return for buyers. These estimates come from a regression of the form in Table 9 column 4, but allowing for the buyer and seller gender to interact. We plot only the coefficients representing single male or female buyers and sellers, with the male buyer-male seller as the base coefficient. See Section II.B for more details on the definition of gender and family structure.

Table 1: Summary statistics

	Gender Group				
Panel A: Full Sample	Single Male	Single Female	Couple	Other	Overall
Log(Sale Price)	11.9473	11.9125	12.1383	12.1104	12.0704
Sample Size	7,721,833	5,751,347	10,127,535	29,283,151	52,883,866
Panel B: Listing Sample					
Log(Sale Price)	12.0798	12.0292	12.2709	12.0597	12.1095
Log(List Price)	12.0677	12.0236	12.2689	11.9539	12.0547
Sale Discount (p.p.)	2.8908	3.0368	2.5413	3.0954	2.9261
Log(Days on Market)	3.7339	3.7052	3.7016	3.7851	3.7467
Sample Size	3,100,949	2,728,421	4,689,273	9,524,421	20,043,064
Panel C: Returns Sample					
Log(Sale Price)	12.1429	12.0692	12.3342	-	12.2138
Annualized Unlevered Returns	0.0847	0.0692	0.0647	-	0.0720
Holding Length (Years)	5.2816	5.7174	5.9840	-	5.7029
Log(Purchase Price)	11.8990	11.8313	12.0793	-	11.9663
Purchase Discount (p.p.)	2.8150	2.5388	2.5629	-	2.6379
Sample Size	2,935,077	2,128,157	4,288,185	-	9,351,419

Note: This table reports the summary statistics for the samples used in the analysis. The summary stats are split into the four gender groups (single male, single female, couple and other), and also pooled. Each cell is the overall mean in the sample. There are three samples reported: the first sample in Panel A is the full sample, which uses all sales transactions reported in the data. The second sample in Panel B is the sample of sales successfully matched to listings data. The third sample in Panel C is the sample of transactions where we successfully identify the gender in both the purchase and sale, and also match the names of the buyer and seller across deeds transactions. See Section II.B for more details on how we identify gender and family structure. See Section II.A for more details on how we match the data.

Table 2: Gender gap in unlevered housing returns

	Unlevered Ann Return		
	(1)	(2)	(3)
Single Female	-0.016***	-0.013***	-0.011***
	(0.000)	(0.000)	(0.000)
Couple	-0.020***	-0.012***	-0.007***
	(0.000)	(0.000)	(0.000)
Holding Length			-0.006***
			(0.000)
Zip-Year-Month FE	No	Yes	Yes
R-squared	0.005	0.354	0.379
Observations	9,351,419	9,351,419	9,351,419

Note: This table shows how realized unlevered housing returns differ by the gender and relationship status of the homeowner. We use observations at the sale transaction level. The sample is restricted to observations for which the gender of all sellers can be identified, and for which we can match the identity of the seller at the time of sale to the identity of the buyer at the time of initial purchase. The omitted group is single males. Standard errors are clustered by zipcode. *,**, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table 3: Gender gap in levered housing returns

	Lev Ann Ret (missing=0%)	Lev Ann Ret (missing=80%)	Lev Ann Ret (LTV=80%)
	(1)	(2)	(3)
Single Female	-0.033***	-0.056***	-0.057***
	(0.001)	(0.001)	(0.001)
Couple	-0.032***	-0.055***	-0.043***
	(0.001)	(0.001)	(0.001)
Holding Length	-0.035***	-0.047***	-0.037***
	(0.000)	(0.000)	(0.000)
Zip-Year-Month FE	Yes	Yes	Yes
R-squared	0.349	0.346	0.330
Observations	9,351,419	9,351,419	9,351,419

Note: This table shows how levered housing returns differ by the gender and relationship status of the homeowner. Levered returns are calculated as explained in Section C. Column 1 uses the homeowner's actual initial LTV if available, and assumes an LTV of 0% if missing. Column 2 uses the homeowner's actual initial LTV if available, and assumes an LTV of 80% if missing. Column 3 assumes an intial LTV of 80% for all observations. In all other tables and figures, we use levered returns assume an initial LTV of 80% as the our measure of the levered return unless otherwise noted. Standard errors are clustered by zipcode. *,**, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table 4: Heterogeneity by zip-level demographics: quartile averages

Male - Female Unlevered Ann Return	Quartile 1	Quartile 2	Quartile 3	Quartile 4
	(1)	(2)	(3)	(4)
Frac Black	0.0101	0.0102	0.0105	0.0140
Frac HS Education or Less	0.0089	0.0100	0.0118	0.0119
Frac 60+	0.0065	0.0091	0.0130	0.0127
Frac Single Female	0.0089	0.0107	0.0100	0.0151
Median Family Income	0.0114	0.0114	0.0103	0.0103
House Price	0.0146	0.0103	0.0094	0.0063

Note: This table presents the average gender gap in unlevered annualized housing returns across quartiles of various zip-level demographic characteristics from the 2010 American Community Survey. This table reports simple averages within each quartile of each demographic variable, without conditioning for other demographic variables. Observations are at the zip-code level and equally weighted.

Table 5: Heterogeneity by zip-level demographics: regressions

	Male - Female Unlevered Ann Return	Male - Female Levered Ann Return
	(1)	(2)
Frac Black	0.004	0.019
	(0.006)	(0.028)
Frac HS education or less	0.024***	0.089**
	(0.009)	(0.043)
Frac 60+	0.025***	0.133***
	(0.009)	(0.047)
Frac Single Female	0.038***	0.188***
	(0.012)	(0.057)
Log Median Family Income	0.011***	0.052***
	(0.003)	(0.014)
R-squared	0.003	0.003
Observations	14,310	14,310

Note: This table shows how the average gender gap in housing returns within a zip code varies with zip-level demographics from the 2010 American Community Survey. We measure the gender gap in each zip code as the average difference between male and female residual unlevered returns (the residual return represents the return after adjusting for the zip-year-month mean). We regress the zip-level gender gap on zip-level demographics. Standard errors are adjusted for heteroskedasticity. *,**, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table 6: Unlevered returns: market timing

	Unlevered Ann Return				
	(1)	(2)	(3)	(4)	(5)
Single Female	-0.016***	-0.012***	-0.011***	-0.010***	-0.009***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Couple	-0.020***	-0.014***	-0.007***	-0.007***	-0.004***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
Holding Length		-0.008***	-0.006***	-0.001	
		(0.000)	(0.000)	(0.001)	
Zip-SaleYM FE	No	No	Yes	Yes	Yes
Zip-BuyYM FE	No	No	No	Yes	Yes
SaleYM FE x BuyYM FE	No	No	No	No	Yes
R-squared	0.005	0.069	0.379	0.534	0.592
Observations	9,351,419	9,351,419	9,351,419	9,351,419	9,351,419

Note: This table shows how housing returns differ by the gender and relationship status of the homeowner, after introducing control variables for the timing of housing transactions. Standard errors are clustered by zipcode. *,**, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table 7: Transaction price

	Log(Purchase Price)		Log(Sale Price)	
	(1)	(2)	(3)	(4)
Single Female	0.013***	0.018***	-0.032***	-0.027***
	(0.001)	(0.001)	(0.001)	(0.001)
Couple	0.023***	0.029***	0.007***	0.014***
	(0.002)	(0.001)	(0.001)	(0.001)
Other	0.085***	0.015***	-0.064***	-0.054***
	(0.005)	(0.003)	(0.002)	(0.001)
Property FE	Yes	Yes	Yes	Yes
Year-Month FE	Yes	No	Yes	No
Zip-Year-Month FE	No	Yes	No	Yes
R-squared	0.794	0.886	0.793	0.887
Observations	52,883,866	52,883,866	52,883,866	52,883,866

Note: This table examines gender variation in transaction prices. We use repeat sales data that allows us to control for property fixed effects. We also control for zip-year-month fixed effects to account for time trends within a zip code. Each observation is a transaction. To better estimate property fixed effects, we do not restrict the sample to buyers or sellers with identified genders and matched names across sales and initial purchase. All observations corresponding to non-identified single women, single men, and couples are included and coded as the "other" category. Single males are the omitted category. Standard errors are clustered by zipcode. *,**, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table 8: List price

	Log(Purchase List Price)		Log(Sale List Price)	
	(1)	(2)	(3)	(4)
Single Female	0.035***	0.033***	-0.019***	-0.015***
	(0.001)	(0.001)	(0.001)	(0.001)
Couple	0.017***	0.015***	-0.025***	0.002**
	(0.001)	(0.001)	(0.002)	(0.001)
Other	-0.076***	-0.060***	-0.164***	-0.093***
	(0.002)	(0.002)	(0.004)	(0.002)
Property FE	Yes	Yes	Yes	Yes
Year-Month FE	Yes	No	Yes	No
Zip-Year-Month FE	No	Yes	No	Yes
R-squared	0.784	0.842	0.786	0.842
Observations	10,984,588	10,984,588	10,984,588	10,984,588

Note: This table examines gender variation in list prices chosen by home sellers. We use repeat sales data that allows us to control for property fixed effects. We also control for zip-year-month fixed effects to account for time trends within a zip code. Each observation is a listing matched to a sales transaction and the sample is restricted to properties for which we observe at least two transactions. To better estimate property fixed effects, we do not restrict the sample to buyers or sellers with identified genders and matched names across sales and initial purchase. All observations corresponding to non-identified single women, single men, and couples are included and coded as the "other" category. Single males are the omitted category. Standard errors are clustered by zipcode. *,**, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table 9: Discounts relative to the listing price

	Purchase Discount		Sale Discount	
	(1)	(2)	(3)	(4)
Single Female	-0.283***	-0.260***	0.146***	0.107***
	(0.007)	(0.005)	(0.007)	(0.005)
Couple	-0.141***	-0.059***	-0.350***	-0.190***
	(0.013)	(0.005)	(0.013)	(0.005)
Other	0.452***	0.475***	0.205***	0.167***
	(0.012)	(0.007)	(0.018)	(0.011)
Zip-Year-Month FE	No	Yes	No	Yes
R-squared	0.003	0.210	0.002	0.209
Observations	20,043,064	20,043,064	20,043,064	20,043,064

Note: This table examines how negotiated transaction discounts vary by gender. We measure purchase and sale discounts as (listing price - transaction price)/listing price $\times$ 100, so a larger purchase discount contributes to a higher return on housing investment and a larger sale discount contributes to a lower return on housing investment. Standard errors are clustered by zipcode. *,**, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table 10: Days on market

	Sale Log(Days on Mkt)	Purchase Log(Days on Mkt)	Unlevered Ann Return
	(1)	(2)	(3)
Single Female	-0.031***	-0.034***	-0.013***
	(0.003)	(0.003)	(0.000)
Couple	-0.041***	0.008***	-0.016***
	(0.003)	(0.003)	(0.000)
Sale Log(Days on Mkt)			-0.003***
			(0.000)
Purchase Log(Days on Mkt)			-0.003***
			(0.000)
Zip-Year-Month FE	Yes	Yes	Yes
R-squared	0.415	0.309	0.398
Observations	2,024,580	2,024,580	2,024,580

Note: This table examines how days on market, measured as the number of days between listing and sale, varies with gender, and contributes the gender gap in housing returns. Standard errors are clustered by zipcode. *,**, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table 11: Controlling for timing and property characteristics

	Unlevered Ann Return			
	(1)	(2)	(3)	(4)
Single Female	-0.010***	-0.010***	-0.006***	-0.006***
	(0.000)	(0.000)	(0.000)	(0.000)
Couple	-0.018***	-0.017***	-0.005***	-0.004***
	(0.001)	(0.001)	(0.000)	(0.000)
Sale Discount		-0.003***	-0.002***	-0.001***
		(0.000)	(0.000)	(0.000)
Purchase Discount		0.003***	0.003***	0.003***
		(0.000)	(0.000)	(0.000)
Log(Age of Unit)				0.010***
				(0.000)
Foreclosure				-0.002
				(0.003)
Garage				-0.005***
				(0.001)
Pool				-0.003***
				(0.000)
Cooling				-0.003***
				(0.001)
Fireplace				-0.004***
				(0.000)
Basement				-0.001
				(0.001)
Waterfront				-0.001*
				(0.001)
Short Sale				-0.046***
				(0.001)
Bathrooms				0.000
				(0.001)
Log(Sq Ft)				-0.005***
				(0.001)
Bedrooms				0.002***
				(0.000)
Log(List Agent Popularity)				-0.001***
				(0.000)
Upgraded				0.008***
				(0.000)
New Construction				-0.002***
				(0.001)
Property Type FE	No	No	No	Yes
Zip-SaleYM FE	No	No	Yes	Yes
Zip-BuyYM FE	No	No	Yes	Yes
SaleYM FE x BuyYM FE	No	No	Yes	Yes
R-squared	0.005	0.022	0.703	0.709
Observations	1,132,921	1,132,921	1,132,921	1,510,609

Note: This table examines how the gender gap in housing returns varies with additional control variables for market timing and property and listing agent characteristics. The sample is restricted to observations for which we have listings data on property characteristics. Standard errors are clustered by zipcode. *,**, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table 12: Selection of property characteristics

	Upgraded	New Construction	Log(House Age)	Log(Sq Ft)	Log(Agent Popularity)
	(1)	(2)	(3)	(4)	(5)
Single Female	-0.009***	0.002***	-0.020***	-0.066***	-0.019***
	(0.001)	(0.000)	(0.003)	(0.001)	(0.002)
Couple	0.000	0.033***	-0.137***	0.143***	0.148***
	(0.001)	(0.001)	(0.004)	(0.002)	(0.003)
Zip-Year-Month FE	Yes	Yes	Yes	Yes	Yes
R-squared	0.299	0.274	0.515	0.448	0.255
Observations	3,542,111	9,351,419	2,211,953	2,007,061	4,000,582

Note: This table examines gender differences in preferences for property and listing agent characteristics. Standard errors are clustered by zipcode. *,**, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table 13: Maintenance and Housing Returns in the American Housing Survey

	Maint/price	Est Unlevered Return	Real Unlevered Return	
	(1)	(2)	(3)	(4)
Single Female	-0.000213	-0.00602**	-0.0116**	-0.00969*
	(0.000150)	(0.00240)	(0.00506)	(0.00508)
Couple	-0.0000813	-0.00596***	-0.00595	-0.00359
	(0.000126)	(0.00215)	(0.00503)	(0.00476)
Other	-0.000219	0.0185*	-0.00444	-0.000308
	(0.000345)	(0.0105)	(0.00936)	(0.00953)
Age of house	0.0000676***	0.000493***	0.000262***	0.000263***
	(0.00000198)	(0.0000363)	(0.0000504)	(0.0000503)
Holding length	0.000330***	-0.00142***	-0.000576***	-0.000606***
	(0.00000735)	(0.000118)	(0.000145)	(0.000148)
Age of householder	-0.0000766***	-0.0000860	0.0000688	0.0000183
	(0.00000307)	(0.0000647)	(0.0000828)	(0.0000990)
Some college				0.000933
				(0.00368)
College degree				0.00107
				(0.00319)
Graduate degree				-0.00114
				(0.00366)
Number of adults				-0.00208
				(0.00180)
Number of children				-0.00287**
				(0.00123)
Black				-0.00680
				(0.00628)
American Indian				-0.0128
				(0.00936)
Asian				-0.0127*
				(0.00697)
Other race				-0.00816
				(0.0142)
Log family income				0.00238*
				(0.00128)
MSA x Survey Year FE	Yes	Yes	No	No
MSA x Sale Year FE	No	No	Yes	Yes
R-squared	0.114	0.026	0.296	0.301
Observations	124,505	135,669	3,716	3,678

Note: This table uses data from the American Housing Survey. In column 1, the dependent variable is reported annual maintenance scaled by home purchase price. In column 2, estimated unlevered return is the annualized unlevered return, calculated using the homeowner's self reported estimate of current home value relative to purchase price. In column 3, real unlevered return is the annualized unlevered return, calculated using the actual purchase price and sale price. Standard errors are double clustered by household and survey year. *,**, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table 14: Variation by market tightness

	Unlevered Ann Return	Purchase Discount	Sale Discount	Log(Purchase Price)	Log(Sale Price)
	(1)	(2)	(3)	(4)	(5)
Single Female	-0.016***	-0.275***	0.055***	0.023***	-0.029***
	(0.000)	(0.018)	(0.014)	(0.001)	(0.001)
Couple	-0.012***	-0.236***	0.018	0.011***	0.016***
	(0.000)	(0.018)	(0.014)	(0.001)	(0.001)
Other		0.018	0.422***	0.030***	-0.061***
		(0.015)	(0.016)	(0.002)	(0.002)
Singe Female X Tightness	0.019***	0.283***	-0.554***	-0.039***	0.013***
	(0.002)	(0.090)	(0.066)	(0.004)	(0.004)
Couple X Tightness	-0.002	0.415***	-0.141**	-0.007	-0.036***
	(0.002)	(0.087)	(0.067)	(0.006)	(0.005)
Other X Tightness		0.076	-0.000	-0.012	0.037***
		(0.072)	(0.074)	(0.008)	(0.009)
Property FE	No	No	No	Yes	Yes
Zip-Year-Month FE	Yes	Yes	Yes	Yes	Yes
R-squared	0.355	0.207	0.208	0.886	0.886
Observations	8,635,824	19,845,356	19,845,356	46,602,251	46,602,251

Note: This table re-estimates the main regressions from Table 2, Table 7, and Table 9, interacting the gender variables with a measure of market tightness. Market tightness is defined as the number of sales in a given county-month, scaled by the outstanding number of listings currently for sale in that county-month. Standard errors are clustered by zipcode. *,**, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Online Appendix for

The Gender Gap in Housing Returns

Paul Goldsmith-Pinkham Kelly Shue

Figure A1: Share of net worth invested in stocks versus housing, Survey of Consumer Finances

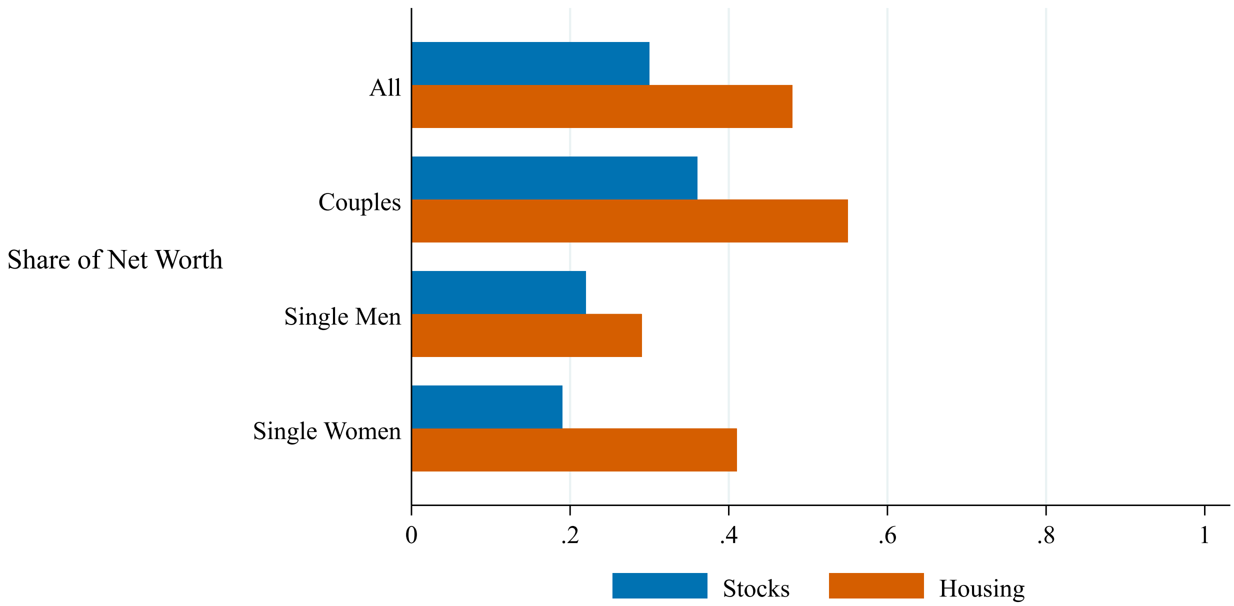

Note: This figure plots the average share of net worth across couples, single males and single females for all individuals with positive net worth. Single women and single men are defined by gender of head of household, and not living with a partner. Couples are defined as those living with partner or married. Net worth share for housing is measured as the equity balance (zero if a non-homeowner) in the primary home, scaled by the overall net worth. The net worth share for stocks is measured as all equity investments (including retirement accounts) scaled by overall net worth: IRA/Keogh accounts, directly held pooled investment funds held by household, directly held stocks held by household, account-type pension plans from the head of household and spouse's current jobs, and trusts investments. We pool across all years in the Survey of Consumer Finance (1989-2016).

Figure A2: Wealth at retirement, Survey of Consumer Finances

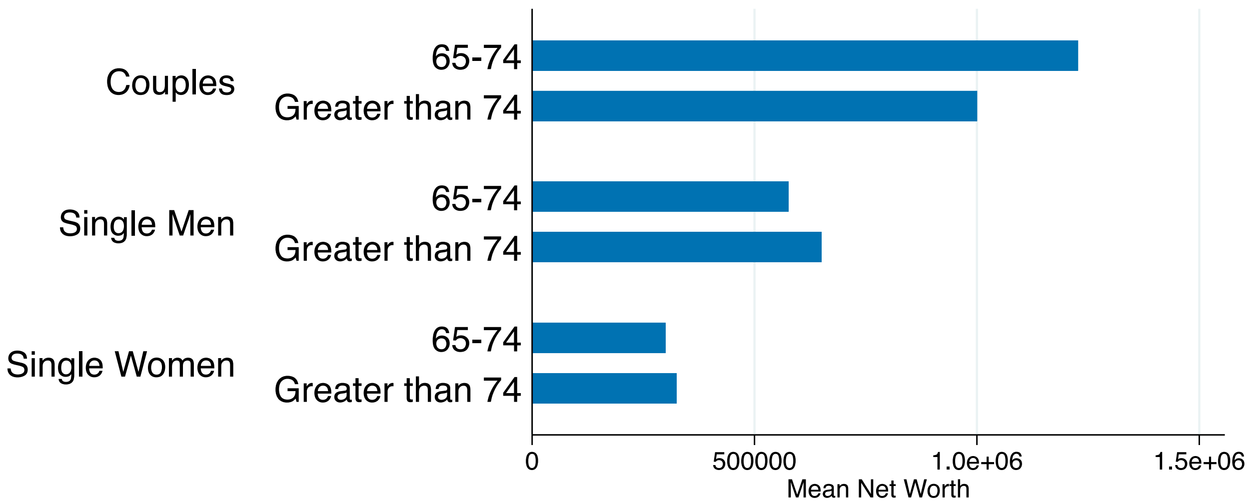

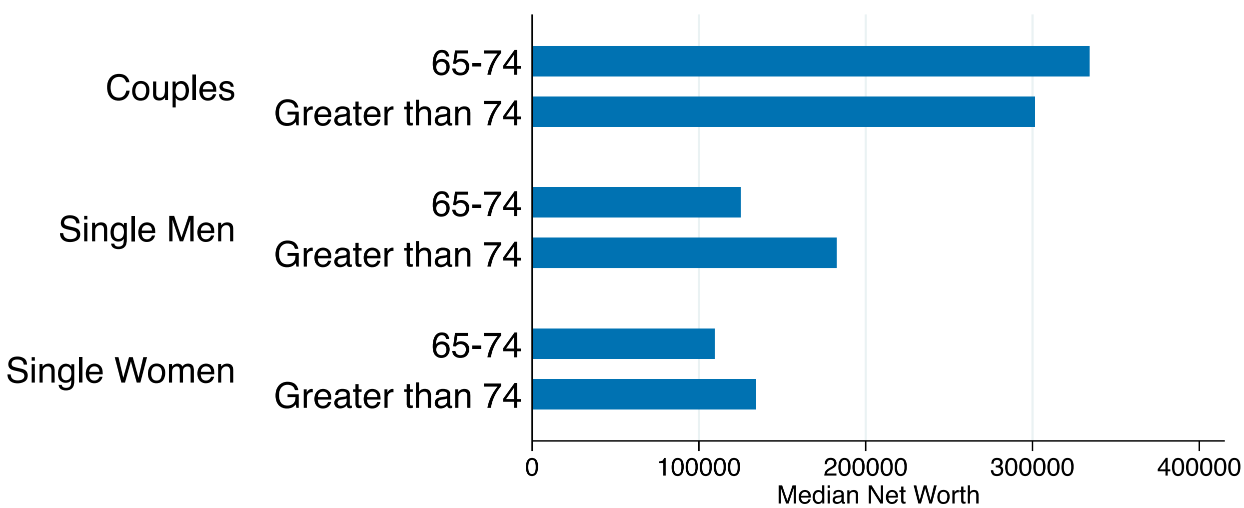

Note: This figure plots the mean and median net worth across couples, single males and single females near retirement age. Age is defined by the head of the household, as reported in the survey. Single women and single men are defined by gender of head of household, and not living with partner. Couples are defined as those living with partner or married. We pool across all years in the Survey of Consumer Finance (1989-2016).

Figure A3: Composition of transactions by gender group over time

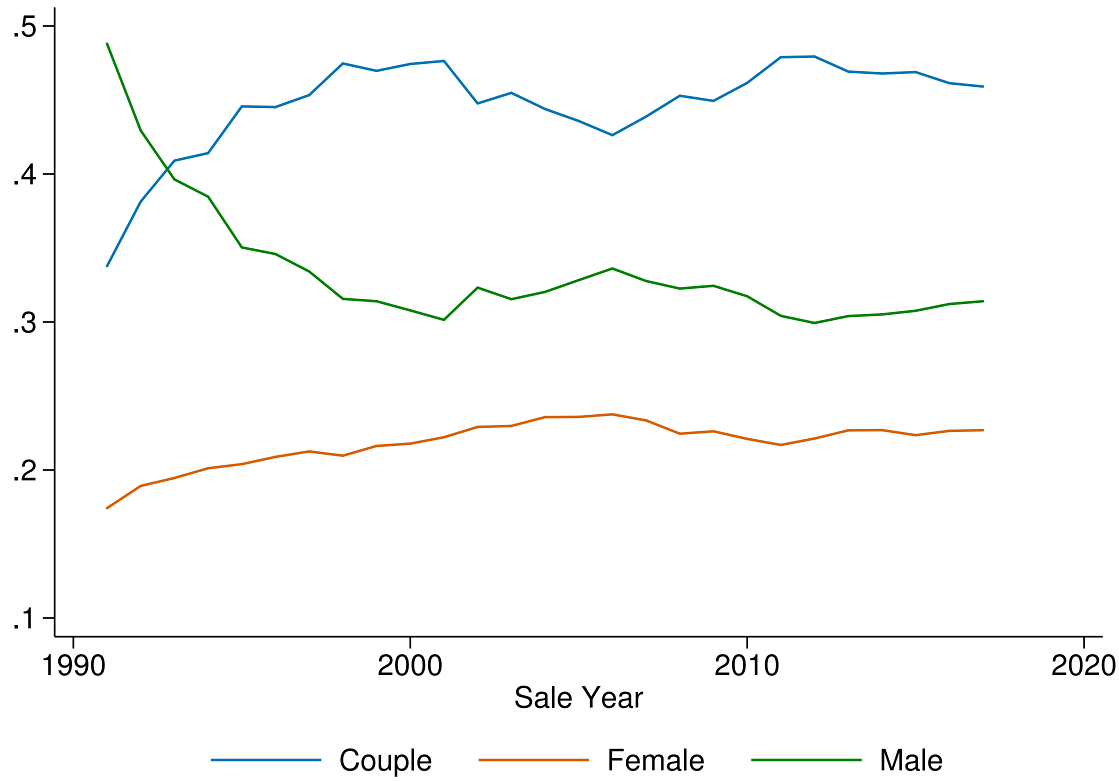

Note: This figure plots the relative composition of sale transactions across couples, single males, and single females within the sample of transactions used for returns estimation. See Section II.B for more details on how we identify gender and family structure.

Figure A4: Transactions over time

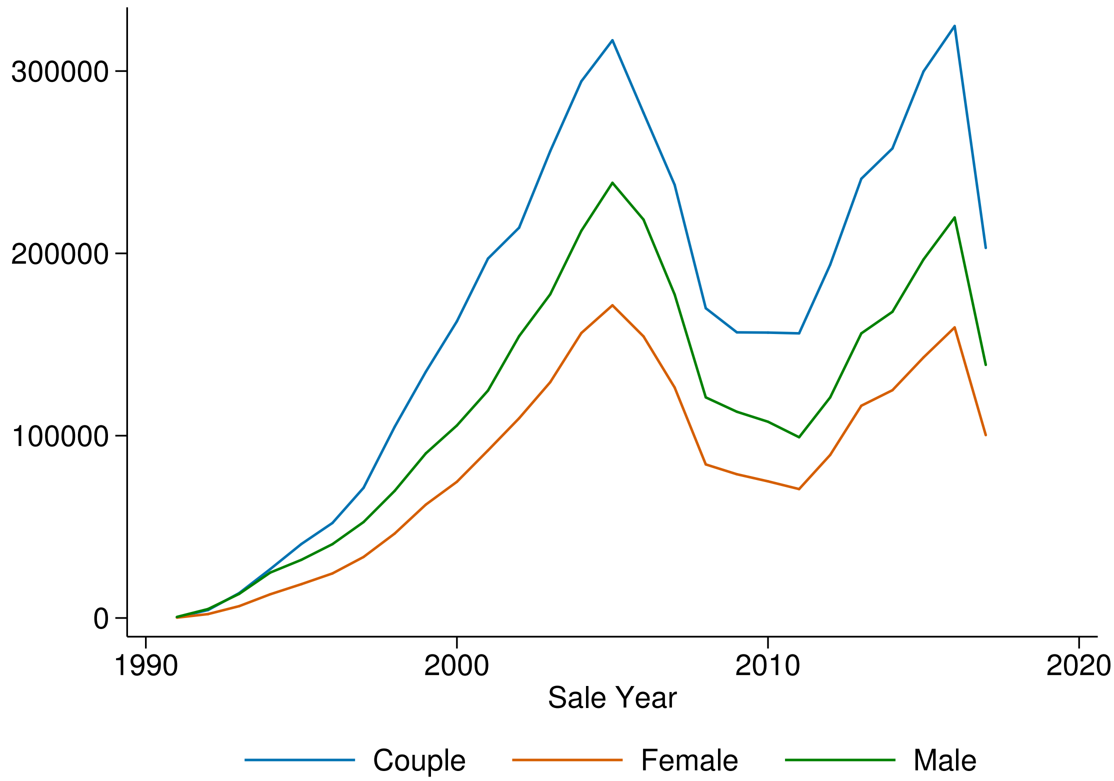

Note: This figure plots the total number of sale transactions across couples, single males and single females within the sample of transactions used for returns estimation. See Section II.B for more details on how we identify gender and family structure.

Figure A5: Median unlevered returns over time by gender group

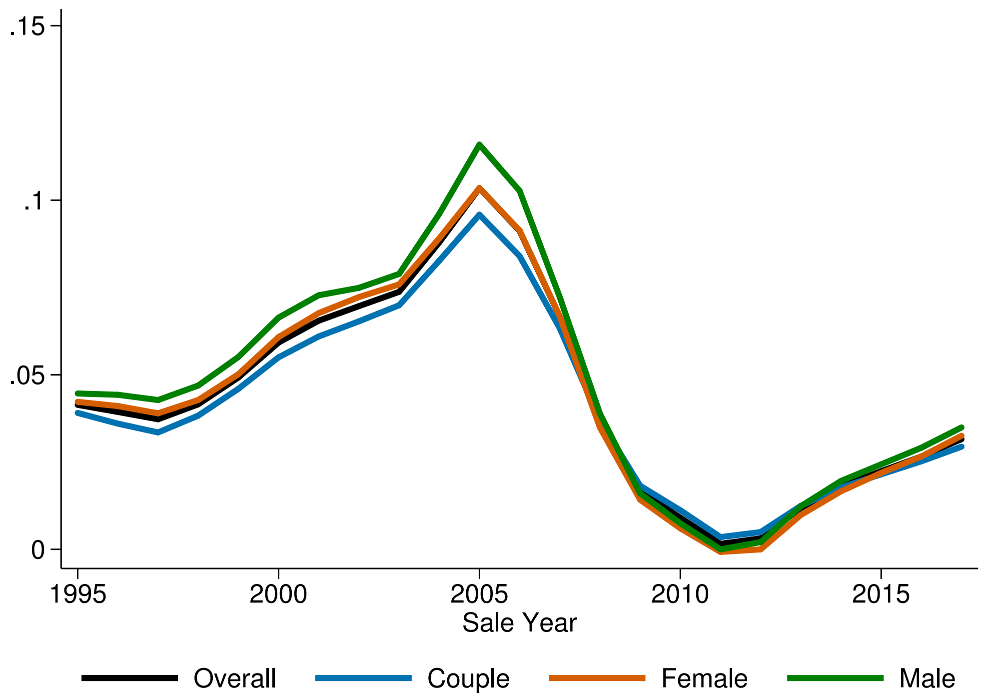

Note: This figure plots the median unlevered annualized return for couples, single women, and single men over the sale year. As our sample begins in 1991, we begin this figure in 1995 to allow for sufficient data to avoid truncation. See Section II.B for more details on the definition of gender and family structure.

Figure A6: Original LTV over time by gender group

Panel A: LTV conditional on Mortgage

Panel B: Missing Mortgage Data

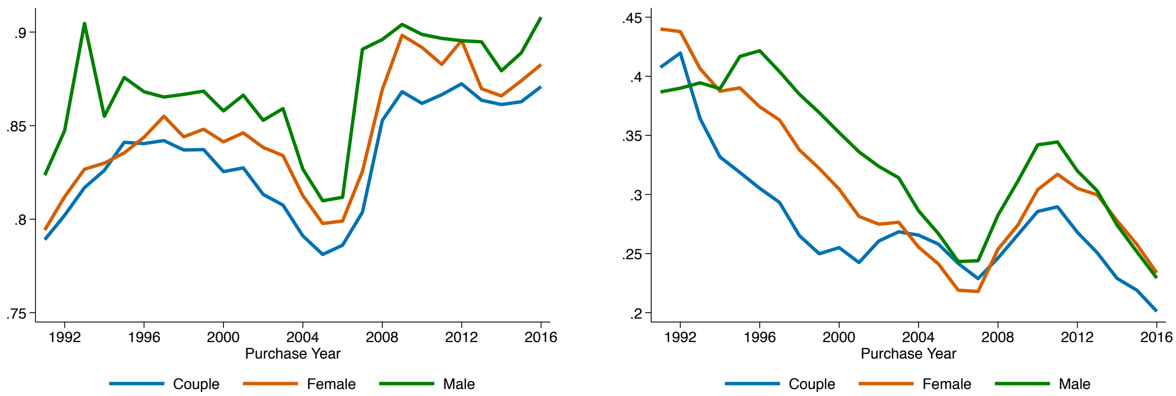

Note: This figure presents information on LTV at time of purchase across couples, single women, and single men. In Panel A, we plot the average LTV at time of purchase, conditional on having information on a mortgage (i.e., conditional on a loan). In Panel B, we plot the share of transactions with missing mortgage data. This combines two forces: full cash transactions and observations with pure missing data.

Figure A7: Transaction share by holding length

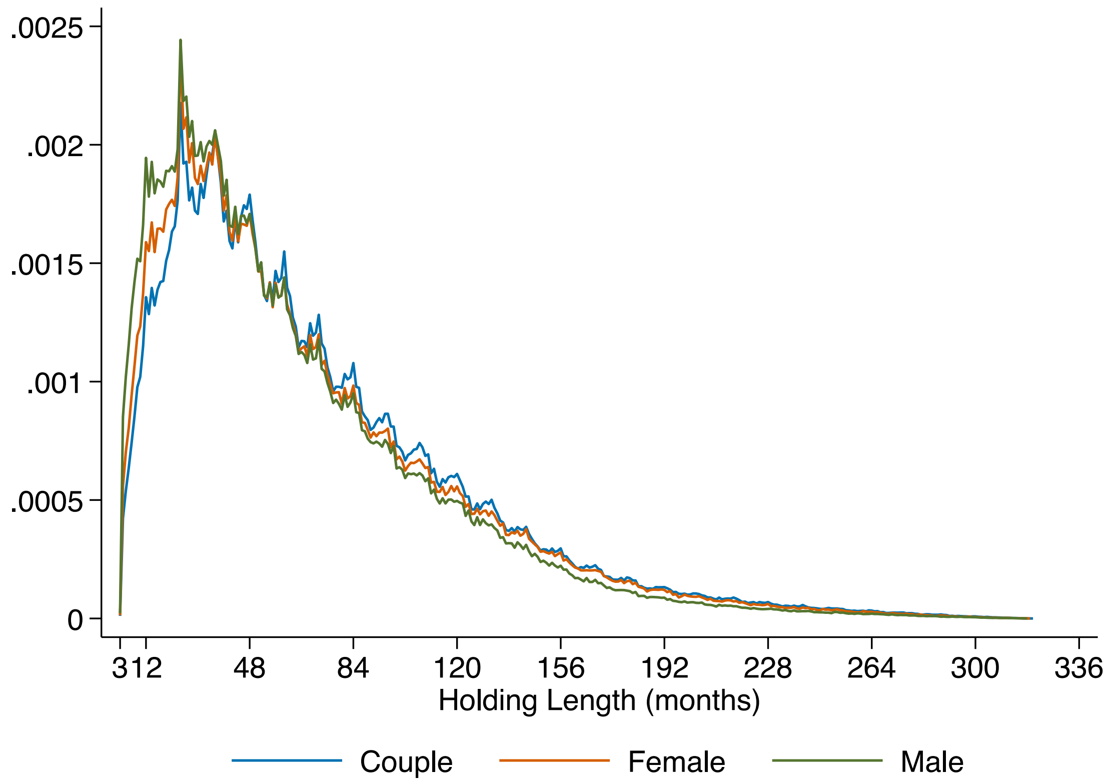

Note: This figure plots the distribution of transactions for couples, single women, and single men across holding lengths for our analysis sample with returns. We restrict our sample to have a minimum of 3 months holding period.

Figure A8: Sale and purchase discount distributions

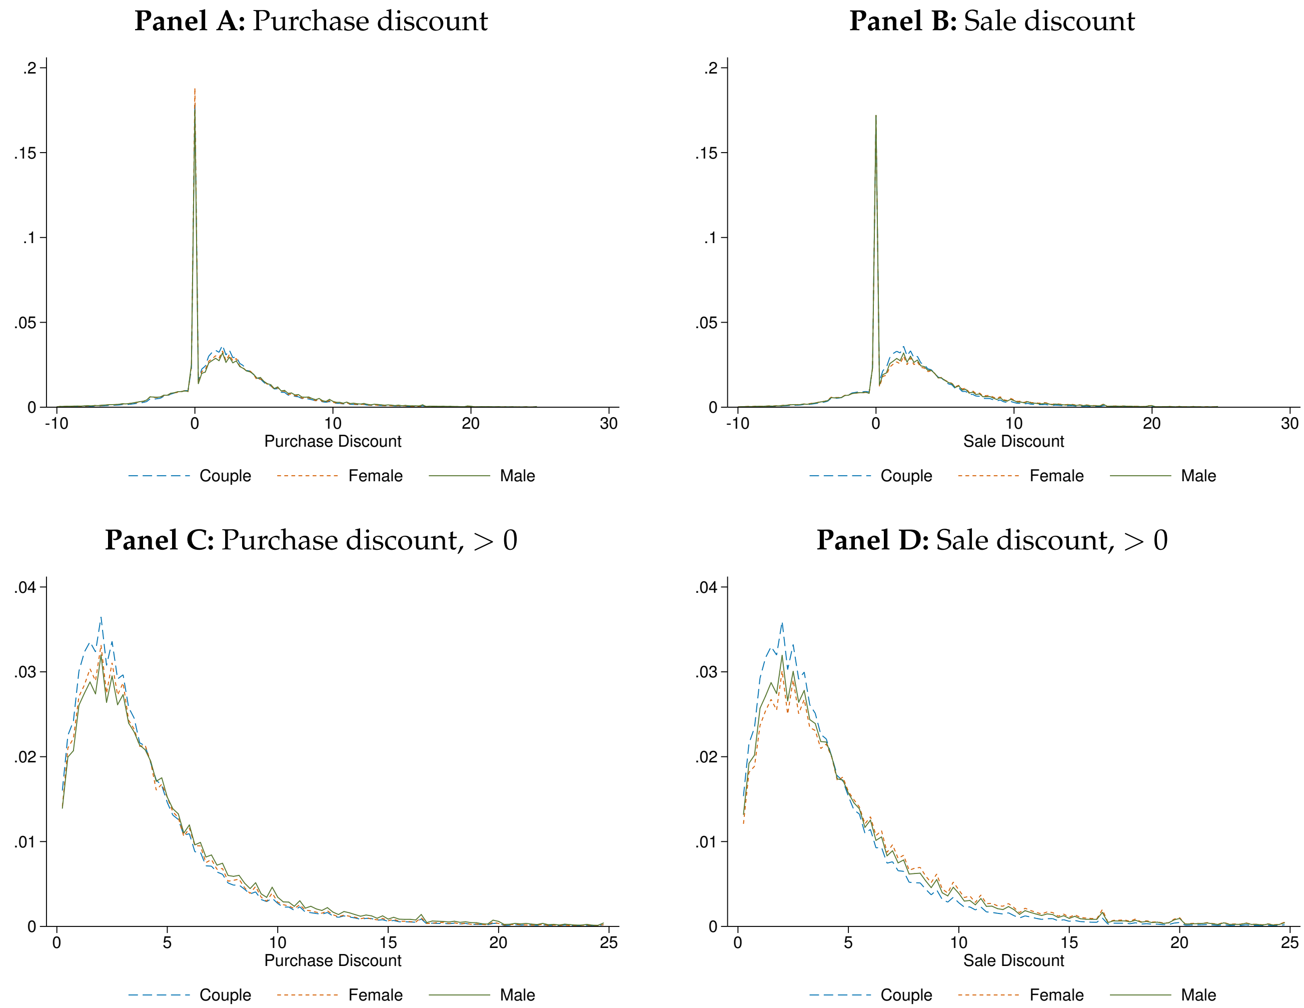

Note: This figure plots the distribution of purchase and sale discounts for couples, single women and single men. We measure discounts as (listing price - transaction price)/listing price $\times$ 100, so a larger purchase discount contributes to a higher return on housing investment and a larger sale discount contributes to a lower return on housing investment. In Panel A and B, we plot the full distributions for all three groups. In Panel C and D, we restrict the distribution to values greater than zero to exclude the spike at 0. See Section II.B for more details on the definition of gender and family structure.

Figure A9: Days on market distribution

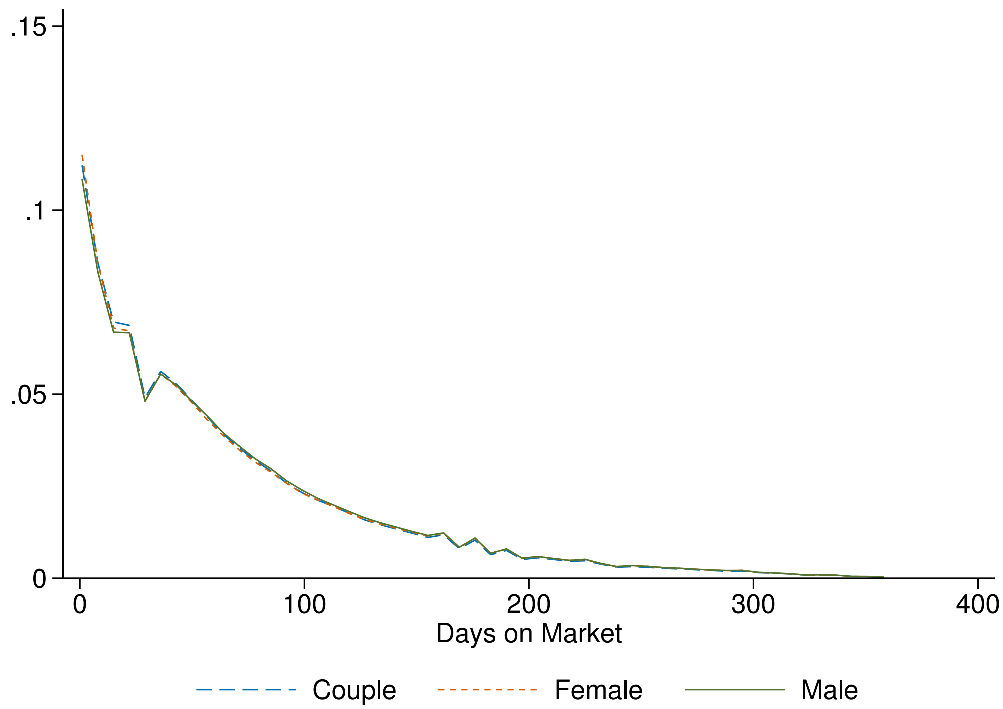

Note: This figure plots the distribution of days on market for listings sold by couples, single women, and single men. See Section II.B for more details on the definition of gender and family structure.

Table A1: Transaction price, focus on subsample with returns data

	Log(Purchase Price)		Log(Sale Price)	
	(1)	(2)	(3)	(4)
Single Female	0.017***	0.018***	-0.020***	-0.018***
	(0.001)	(0.001)	(0.001)	(0.001)
Couple	0.013***	0.010***	-0.002	0.008***
	(0.002)	(0.001)	(0.002)	(0.001)
Other	0.075***	0.027***	-0.080***	-0.044***
	(0.002)	(0.001)	(0.002)	(0.001)
Property FE	Yes	Yes	Yes	Yes
Year-Month FE	Yes	No	Yes	No
Zip-Year-Month FE	No	Yes	No	Yes
R-squared	0.793	0.886	0.793	0.886
Observations	52,883,866	52,883,866	52,883,866	52,883,866

Note: This table is similar to Table 7, except that all indicator variables for male, female, and couples represent observations in our housing returns sample. To preserve the ability to estimate property fixed effects, we categorize all other observations into the "other" category. Standard errors are clustered by zipcode. *,**, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table A2: List price, focus on subsample with returns data

	Log(Purchase List Price)		Log(Sale List Price)	
	(1)	(2)	(3)	(4)
Single Female	0.021***	0.024***	-0.015***	-0.011***
	(0.001)	(0.001)	(0.001)	(0.001)
Couple	0.002	0.005***	-0.014***	-0.001
	(0.002)	(0.001)	(0.001)	(0.001)
Other	0.057***	0.019***	-0.106***	-0.053***
	(0.002)	(0.001)	(0.002)	(0.001)
Property FE	Yes	Yes	Yes	Yes
Year-Month FE	Yes	No	Yes	No
Zip-Year-Month FE	No	Yes	No	Yes
R-squared	0.783	0.841	0.783	0.842
Observations	10,984,588	10,984,588	10,984,588	10,984,588

Note: This table is similar to Table 8, except that all indicator variables for male, female, and couples represent observations in our housing returns sample. To preserve the ability to estimate property fixed effects, we categorize all other observations into the "other" category. Standard errors are clustered by zipcode. *,**, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table A3: Discounts relative to the listing price, focus on subsample with returns data

	Purchase Discount		Sale Discount	
	(1)	(2)	(3)	(4)
Single Female	-0.275***	-0.227***	-0.062***	-0.040***
	(0.011)	(0.008)	(0.009)	(0.006)
Couple	-0.250***	-0.168***	-0.017	-0.007
	(0.017)	(0.008)	(0.015)	(0.007)
Other	0.158***	0.032***	0.764***	0.420***
	(0.015)	(0.007)	(0.017)	(0.008)
Zip-Year-Month FE	No	Yes	No	Yes
R-squared	0.001	0.208	0.004	0.209
Observations	20,043,064	20,043,064	20,043,064	20,043,064

Note: This table is similar to Table 9, except that all indicator variables for male, female, and couples represent observations in our housing returns sample. We categorize all other observations into the "other" category. Standard errors are clustered by zipcode. *,**, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table A4: Unlevered housing returns: weighted by holding length

	Unlevered Ann Return		
	(1)	(2)	(3)
Single Female	-0.006***	-0.005***	-0.004***
	(0.000)	(0.000)	(0.000)
Couple	-0.005***	-0.001***	0.000***
	(0.000)	(0.000)	(0.000)
Holding Length			-0.002***
			(0.000)
Zip-Year-Month FE	No	Yes	Yes
R-squared	0.001	0.384	0.389
Observations	9,351,419	9,351,419	9,351,419

Note: This table re-estimates Table 2, weighting each observation by holding length. Standard errors are clustered by zipcode. *,**, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table A5: Levered housing returns: weighted by holding length

	Lev Ann Ret (missing=0%)	Lev Ann Ret (missing=80%)	Lev Ann Ret (LTV=80%)
	(1)	(2)	(3)
Single Female	-0.009***	-0.014***	-0.014***
	(0.000)	(0.000)	(0.000)
Couple	0.002***	-0.002***	0.004***
	(0.000)	(0.000)	(0.000)
Holding Length	-0.011***	-0.013***	-0.007***
	(0.000)	(0.000)	(0.000)
Zip-Year-Month FE	Yes	Yes	Yes
R-squared	0.337	0.329	0.328
Observations	9,351,419	9,351,419	9,351,419

Note: This table re-estimates Table 3, weighting each observation by holding length. Standard errors are clustered by zipcode. *,**, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

Table A6: Match rates

Panel A: Overall Match Rates

Seller Gender	Buyer Gender Single Male	Single Female	Couple	Overall
Single Male	0.1385	0.0868	0.1010	0.3262
	[0.1207]	[0.0830]	[0.1225]	
Single Female	0.0936	0.0748	0.0752	0.2437
	[0.0901]	[0.0620]	[0.0915]	
Couple	0.1378	0.0930	0.1993	0.4301
	[0.1591]	[0.1095]	[0.1615]	
Overall	0.3700	0.2546	0.3755	1

Panel B: Zip-Year-Quarter Match Rates

Seller Gender	Buyer Gender Single Male	Single Female	Couple	Overall
Single Male	0.1503	0.0997	0.1141	0.3372
	[0.1400]	[0.0936]	[0.1204]	
Single Female	0.1032	0.0869	0.0866	0.2535
	[0.0990]	[0.0736]	[0.0940]	
Couple	0.1533	0.1088	0.2225	0.4642
	[0.1596]	[0.1171]	[0.2000]	
Overall	0.3799	0.2696	0.3958	1

Note: This table presents the joint probability of a seller of a given gender and family type matching with a seller of a given gender and family type. The first number in each cell is the empirical match rate as seen in the data. The number in brackets is the theoretical number if match rates were random (using the product of the two marginal empirical rates). Non-categorized genders and family types are excluded from the matching exercise. Panel A pools the full sample, while Panel B calculates the actual and random match rates at the zip-year-quarter, and then takes the unweighted average across zip-year-quarters.

Table A7: Upgrades

Unlevered Ann Return	No Upgrades	Upgrades
	(1)	(2)
Single Female	-0.010***	-0.017***
	(0.000)	(0.000)
Couple	-0.009***	-0.020***
	(0.000)	(0.001)
Zip-Year-Month FE	Yes	Yes
R-squared	0.416	0.388
Observations	2,406,965	1,135,146

Note: This table estimates the gender gap in housing returns for subsamples of the data with and without any mention of synonyms associated with renovations and upgrades in the listings data. The sample is restricted to observations that are matched to listings with non-missing property descriptions. Standard errors are clustered by zipcode. *,**, and *** denote statistical significance at the 10%, 5%, and 1% level, respectively.