

Learning about Housing Cost: Theory and Evidence from the German House Price Boom

Fabian Kindermann

University of Regensburg & CEPR

Julia Le Blanc

European Commission, Joint Research Center

Monika Piazzesi

Stanford & NBER

Martin Schneider

Stanford & NBER

February 21, 2024

Abstract

This paper studies expectation formation by heterogeneous households during the recent housing boom in Germany. New survey data show that house price forecasts differ prominently by geography and tenure, and much less by other household characteristics. Renters make on average higher and thereby more accurate price forecasts than owners, but their forecasts are more dispersed and their mean squared forecast errors are higher. A quantitative model of learning about housing cost can match these facts. It emphasizes the unique information structure of housing among asset markets. While renters who do not own the asset are less informed about house prices than owners, they are relatively *better* informed about cash flow growth: they pay for housing services that owners simply consume. Renters thus attribute a boom relatively more to rent growth that raises the level of prices rather than to transitory changes in discount rates. Renter forecasts are therefore not only less precise, but also more optimistic on average than those of owners.

Email addresses: fabian.kindermann@ur.de, julia.le-blanc@ec.europa.eu, piazzesi@stanford.edu, schneidr@stanford.edu. We thank Rudi Bachmann, Theresa Kuchler, Ulrike Malmendier, Charles Nathanson, Tom Sargent, Johannes Stroebel, Stijn van Nieuwerburgh, Michael Weber, and seminar participants at Bonn, Chicago, Deutsche Bundesbank, ECB, Federal Reserve Bank at Cleveland, Harvard, HULM, Leuven, Manchester, Mannheim, Munich, "Micro and Macro Implications of Household Finance", NBER Summer Institute, Nuremberg, NYU, Queens, UBC, and UC Santa Cruz for helpful comments.

A large part of this paper was written while Julia Le Blanc was a senior research economist at the Deutsche Bundesbank in Frankfurt. We thank the Deutsche Bundesbank for kindly providing us with their available data on house prices and expectations as well as with an excellent research infrastructure. The views expressed in this paper are solely those of the authors and should not, under any circumstance, be regarded as stating an official position of the European Commission.

1 Introduction

Changes in expectations play a key role in most theories of housing booms and busts. At the same time, direct evidence on how expectations respond to large price swings is scarce. During past booms, such as the US housing boom of the early 2000s, major representative household surveys did not yet contain detailed questions on expectation formation. This is why we know relatively little about the distribution of beliefs or information sets across households or individuals. Existing studies of housing booms have thus typically focused on *average* investor sentiment as a driver of prices. In particular, the goal has been to find a theory of expectation formation such that the average investor's subjective present discounted value of rents is consistent with average survey forecasts and house prices.

This paper uses new household survey data to characterize the *cross-sectional distribution* of house price forecasts during the recent housing boom in Germany. It then proposes a quantitative model of expectation formation such that the cross section of investors' subjective present discounted values is consistent with the cross section of survey forecasts. We show that *learning about housing cost* is a key force shaping heterogeneous beliefs: households form opinions about house values based on noisy signals of prices and rents in their region, for example values of individual properties. Apart from region, the most important dimension of heterogeneity is tenure: owners and renters disagree systematically since they receive different types of signals.

Two novel stylized facts about the German boom guide our analysis. First the average renter is significantly more optimistic than the average owner. Since households overall underestimate actual house price growth, this means that the average renter forecast is more accurate. Second, renter forecasts are much more dispersed than owners' forecasts. As a result, both the best and worst forecasts are made by renters, and the average squared forecast errors for renters is higher than for owners. These facts are difficult to explain with existing theory. In particular, they are not consistent with a common trait shared by all renters, such as superior information on rent or less experience in housing markets. We need to explain not only why renters are on average different from owners, but also why differences among renters are larger than among owners.

The starting point for our theory is the unique information structure of housing markets. To make accurate forecasts about future house prices, households require information about both prices and rents, the cash flow of the housing asset. Housing markets are unique among asset markets in that *non-owners* of the asset have easy access to information about cash flows. They not only pay rent every month, but also tend to have a network of neighbors who share the same experience. In contrast, *owners* of houses directly consume services from their home (as do their owner neighbors) and therefore

need to pay less attention to cash flow. At the same time, they can be expected to be better informed about house prices.

In a housing boom that features both rent and price growth, these different information sets naturally lead to disagreement between renters and owners about its cause: renters attribute the price boom relatively more to rent growth, whereas owners put relatively more weight on factors other than rents, summarized in a present value model as lower discount rates. Rent growth typically has permanent effects on the level of prices whereas variation in discount rates tend to be transitory. It then follows that renters are more optimistic about future house prices than owners: as in the data, optimism during the housing boom is driven by non-owners. At the same time, renters' forecasts are also more dispersed because their signals about prices are less precise.

The paper makes two main contributions to the literature. First, we document facts on expectation formation using new data from the Bundesbank. Most of our evidence comes from the 2014 and 2017 waves of the Panel on Household Finances (PHF). The data set is unique in that (i) it integrates questions on expectation formation with a detailed household finance survey and (ii) it can be matched to location information at a fine geographic level. The latter is crucial given typical regional differences in price growth during booms.¹ We also present evidence from the Bundesbank Online Panel - Households (BOP-HH), in particular results from a question on information acquisition about housing we proposed for the 2019 wave.

Controlling for regional differences in housing market performance, we show that the house price growth forecast of the average renter is 2-3 percentage points (pp) higher than that of the average owner. While the average renter's forecast is typically within 1pp of realized house price growth in a region, owners forecast at most half of realized growth. Tenure and region further subsume the effects of many other observables such as age, income, wealth, risk aversion and financial literacy, rendering each of these variables insignificant in cross-sectional regressions. In addition, forecast dispersion is large: about 5pp for renters and 4pp for owners. Mean squared error can be viewed as squared bias plus squared dispersion. At our measured values, greater dispersion typically implies higher mean squared error for renters.

The survey delivers three pieces of direct evidence on renters' and owners' information about rents and prices. First, we ask households what information they use to evaluate future house prices. The most important source – mentioned by more than 80% of house-

¹The PHF is the German component of the Household Finance and Consumption Survey (HFCS), the European counterpart to the US Survey of Consumer Finances. We draw on waves in 2014 and 2017. During past housing booms, such as the US housing boom of the early 2000s, major representative household surveys did not yet incorporate detailed questions on expectation formation that would have allowed researchers to explore the distribution of forecasts.

holds – is the direct observation of prices or rents, with renters looking more at rents and owners looking more at prices. Renters further indicate that talking to family and friends is important. Other sources that aggregate information, such as classical and social media or financial advisors, play a smaller role. Second, we show that owners who are also real estate investors make more optimistic forecasts close to those of renters, as one would expect if they also receive regular signals on market rent. Third, we show that even though housing is an illiquid asset, owners' perception of their own house value tracks house prices in their local region. Here we exploit a panel component of the PHF to construct property-level price growth that we can match to our local price growth measures.

Our second contribution is to formulate and estimate a model of learning about housing cost for the entire *distribution* of household expectations during the boom. Both owners and renters learn from noisy signals about what drives the housing boom, with renters receiving more precise signals about rents and owners about prices. Learning is difficult because price fluctuations reflect movement in both rent growth and discount rates. While both high rent growth or low discount rates can generate house price booms, the two factors differ in the outlook for house prices: persistent rent growth induces momentum in prices, whereas mean reversion in discount rates generates reversal. In other words, agents who believe the boom is driven by rent growth should expect it to continue, whereas those who attribute it to discount rates should expect it to fade out.

Consider renters and owners who share an understanding of these price dynamics, but receive different signals about current rents and house prices. Renters observe rent increases through quite precise rent signals but receive noisy price signals. As a result, they attribute the boom relatively more to rent growth. Anticipating momentum, they forecast high price growth on average. At the same time their forecasts are very dispersed owing to high noise in their price signals. Owners, in contrast, see price increases through more precise price signals, but are less sure about what is happening to rents. As a result, they attribute the boom relatively more to a drop in discount rates – for example, an easing of credit conditions. Anticipating reversal, their average price growth forecasts are lower than those of renters. Since their signals about prices are more precise, however, their forecasts are also less dispersed.

We quantify our theoretical model by inferring parameters of households' information structure and their subjective beliefs from data on prices, rents and the cross-section of house price forecasts. Formally, our model describes the joint dynamics of prices and rents perceived by households. We restrict beliefs to respect a standard asset pricing equation that expresses prices as values of a discounted stream of future rents, much like the familiar user cost model of housing. Our quantitative exercise shows that signal extraction with differential information can indeed jointly generate first and second moments of the forecast distribution in the 2014 data of the PHF, given the actual path of

prices and rents during the German boom from 2010 to 2014.

According to our estimates, households started out in 2010 with a view that both low rent growth and a high discount rate contributed to low prices. This initial belief is sensible: by 2010, Germany had seen a long period of decline in price-rent ratios, most recently reinforced by the global financial crisis. At the same time, rent growth had been sluggish. As the economy recovered and the housing boom got underway, renters saw rents lead prices, whereas owners interpreted price increases as due to the recovery from the financial crisis. The quantitative success of the model is that the observation of a few years of boom prices can generate sufficient divergence between average owner and renter beliefs, even though signals contain enough noise to generate observed dispersion in forecasts.

We also provide direct evidence to assess a number of alternative explanations for a discrepancy between the house price forecasts of renters and owners. One candidate is that households overweigh negative signals when forming expectations. Renters who fear rent increases (that they perceive as connected to price increases) thus report higher forecasts than owners who worry about price declines. To explore this possibility, we split our sample according to individual risk aversion. We find that the difference between renter and owner forecasts is marginally higher in the sample of households with above median risk aversion. We conclude that worry could play some role in shaping households' house price expectations. However, the majority of the forecast difference between renters and owners we find in our data can not be explained by this hypothesis.

Another interesting possibility is that attention allocation explains our findings. It could be, for example, that owners do not expect to move and hence "sleep through" a housing boom, whereas renters are looking around to perhaps buy. The survey contains a number of questions about households' recent history as well as future plans. We consider subsamples of owner households who recently moved into their house, households who plan to move in the near future, as well as renter households who plan to buy in the near future. We show that forecast differences between owners and renters are not due to the behavior of these subgroups relative to others. These findings lead us to focus on information flow due to tenure *status* as the key distinction between owners and renters, and not information differences that may arise from an intention to trade.

A general takeaway from our results is that housing markets are special in how prices convey information. Our findings confirm the familiar theme that the housing market is illiquid and agents learn by sampling a small number of price signals. In this sense, housing resembles an over-the-counter market for securities. Information about cash flow, however, works very differently for housing than for securities. Indeed, securities deliver pecuniary payoffs – such as dividends or promised fixed income payments – to owners and owners form beliefs about those payoffs in order to value securities. Housing, in contrast, provides nonpecuniary dividends to owners who are not investors. At the

same time, non-owners are especially well informed about the market value of those dividends.²

The lesson for modeling is that inference from both house prices and rents should be incorporated into future studies of expectations, prices and volume. This means that it is not sufficient to simply import mechanisms from security market models to housing without further modification. In particular, housing models should accommodate the informational advantage enjoyed by renters as well as investors when valuing dividends. Interestingly, it is precisely these two sets of market participants who account for most volume during booms and who hence have the strongest impact on average transaction prices.³

Related Literature

We follow a long tradition of studying asset price expectations through the lens of a present-value relationship. Following the seminal paper of Campbell and Shiller (1988) on equity, a large literature has explored time variation in investors' views on cash flows versus discount rates. Most of this literature does not use survey data, but instead imposes rational expectations; the typical conclusion for equity is that views about discount rates are much more volatile than those about prices. Campbell, Davis, Gallin and Martin (2009) present a similar result for aggregate US housing indices. In contrast, De La O and Myers (2021) use analyst forecasts to document large variation in cash flow expectations for equity that can account for much of the volatility in price-dividend ratios.

Like De La O and Myers (2021), our model also characterizes *subjective* beliefs about cash flow and discount rates without imposing rational expectations. In contrast to the literature, we infer an entire *system* of beliefs for a heterogeneous population of agents. Learning from heterogeneous signals leads owners to attribute relatively more price movements to discount rates rather than rent growth, whereas renters take the opposite view. Signal extraction is difficult because prices exhibit familiar statistical properties inherited from the two factors, in particular, short term momentum (Guren, 2018) and long term reversal (Cutler, Poterba and Summers, 1991, Head, Lloyd-Ellis and Sun, 2014).

Our paper also adds to a growing literature on survey expectations of house prices, pioneered by Case and Shiller (1989). During the US housing boom, average price expecta-

²Our data actually allow a comparison of how households think about stocks versus housing: when the survey elicits stock price expectations, about one half of non-owners take the "don't know" option. In contrast, the overwhelming majority of renters – close to 85% – provide a quantitative forecast of house price growth.

³In our German data, 58 percent of homeowners who changed their main residence within the last decade were previously renting. In contrast, only 3 percent of renters who changed their main residence within the last decade were previously homeowners. These numbers are from the 2014 Panel of Household Finances, which is further explained below.

tions in coastal US cities also showed growing optimism surpassed by realizations, as we see in our German data (for example, Case and Shiller, 2003 or Case, Shiller and Thomson, 2012). Our learning model explains such gradual adjustment of expectations by Bayesian updating. On the empirical side, we provide new evidence on the composition of beliefs during the boom.

Our finding that dispersed signals are used to form expectations is consistent with other recent studies of house price expectations. For example, Bailey, Cao, Kuchler and Stroebel (2018) show that individuals' home purchase decisions are related to the experience of their friends on social networks. Kuchler and Zafar (2019), document that households' views about both first and second moments of future house prices are affected by recent local price observations. Armona, Fuster and Zafar (2019) describe an information experiment, in which individuals revise forecasts when they are told the actual recent price experience in their region. In contrast to these studies, our data and model emphasizes differences in information sets between renters and owners.

Our results on the special nature of owners' vs non-owners' information in housing markets stands in contrast to work on experience effects due to ownership of stocks or consumption goods.⁴ For example, D'Acunto et al. (2021) show that agents' inflation expectations reflect price changes in goods they spend more on. Hartzmark, Hirshman and Imas (2021) show that stock market investors and consumers are more optimistic and have more dispersed forecasts about stocks or goods they own themselves; the interpretation is that they pay more attention to those items and believe they are more informed about them. In the housing market, in contrast, *non-owners'* forecasts are more optimistic, more closely reflect house prices, and exhibit higher dispersion.

There is little work on heterogeneity of beliefs between renters and owners. Favara and Song (2014) develop a theoretical model of the housing market with asymmetric information: in equilibrium more optimistic agents sort into ownership, whereas renters are more pessimistic. Our results suggest that asymmetric information is indeed important, but that sorting does not occur right away during a boom, perhaps because of transaction costs. It is plausible, however, that the belief differences we document lead to trade; since renter-owner transitions are a large share of housing market volume, they may thereby have a large effect on transaction prices.⁵ Adelino, Schoar and Severino (2018) measure households' perceptions of house price risk and show that renters view housing as riskier than owners. In our model, renters' subjective variance of house prices is also higher than

⁴The literature on experience effects more generally goes back to Vissing-Jorgensen (2003) and Malmendier and Nagel (2011); see Malmendier (2021); for a survey.

⁵Piazzesi and Schneider (2009) study a search model with heterogeneous beliefs and short sale constraints and show how entry of a small share of optimistic agents can have a large price impact. See Gao, Sockin and Xiong (2020) for a model where information frictions in housing affect migration and local capital accumulation.

that of owners. Indeed, while renters have better signals of rent, they rely on noisier signals of prices. As a result, they perceive high uncertainty about the discount rate, which accounts for the bulk of price volatility.

Finally, our work relates to an emerging literature that incorporates survey evidence on expectations into quantitative models of housing choice and house prices. Landvoigt (2017) as well as Kaplan, Mitman and Violante (2020) have emphasized the role of expectations for the 2000s housing boom in the US. Leombroni, Piazzesi, Schneider and Rogers (2020) study how heterogeneous inflation expectations affected house prices and interest rates in the 1970s. Ludwig, Mankart, Quintana, Vellekoop and Wiederholt (2020) show how heterogeneous beliefs drive price dynamics in the recent Dutch housing boom-bust cycle; they also rely on a dataset with detailed income, balance sheet and expectations information at the household level. In this paper we approach the relationship between survey expectations and prices from the opposite side: rather than establish an effect of the distribution of beliefs on prices, we derive distributions of beliefs from price (and rent) histories.

The paper is structured as follows. Section 2 provides some information about the institutional background and presents historical data on the housing market in Germany. Section 3 introduces the survey data and documents key stylized facts about the cross section of house price growth forecasts. Section 4 provides direct evidence for the information sets of renters and owners. Section 5 presents the model and its quantitative implications. In Section 6 we explore alternative theories of the cross-section of house price forecasts. The appendix contains details about the data, additional empirical results, and model derivations.

2 Housing in Germany

In this section, after describing the institutional features of the German housing and rental market, we provide an overview of recent price and rent movements in Germany. In the aftermath of the financial crisis, Germany emerged from a decades-long housing slump. The boom saw an initial acceleration in rents followed by a boom in prices and price-rent ratios. New regional price data further show substantial regional heterogeneity, with a stronger boom starting earlier in major metropolitan areas.

2.1 Institutional Background

The German housing market is characterized by a low homeownership rate. Between 2010 and 2017 the share of households owning their main residence ranged around 44

percent, compared with about 60 percent in the Euro area as a whole. Germany has high transfer taxes on buying real estate, no mortgage interest tax deductions for owner-occupiers, and strong renter protections, see for example Kaas, Kocharkov, Preugschat and Siassi (2020) and Weber and Lee (2018). Overall, the rental market is very well developed and households are able to find rental housing in essentially all locations and quality segments.

The German rental market offers housing units of different qualities. Compared to the U.S., where the markets for owner-occupied and rental units are highly segmented, the German housing markets are far more integrated. For example, there are rental units among single family homes and townhouses in the top quality segment. As a result, renting is a choice made by a broad part of the population. In fact, in 2014 the homeownership rate in the top income decile was only 71 percent, while the CPS reports a number of almost 84 percent for the US. Even in the top net wealth quintile, there are 12 percent of households who rent their primary residence in Germany.

2.2 Average house price and rent data

Our housing data come from bulwiengesa AG, a leading German real estate data and consulting firm. For the period since 2005, they provide average transaction prices and market rents in each of the 401 German counties, measured in Euros per square meter.⁶ For each county, we observe separate price indices by dwelling type (single-family house, town house as well as existing and newly built apartments), while for rentals we have separate rent indices for existing and newly built apartments. To compare prices to survey forecasts below, we would like to measure the value of the typical house in a region – the subject of the survey question – as opposed to the average transaction price. We would also like to aggregate price data to larger regions for which we have enough data to meaningfully measure average forecasts.

We thus build price indices for Germany as a whole as well as four regions that differ in average growth rates. We aggregate county level indices to larger regions – for example, all of Germany – using survey weights from the 2014 Bundesbank’s Panel on Household Finances (PHF), a representative country-wide survey that we also use to measure forecasts (described in more detail below).⁷ For every household in the survey, we know county and dwelling type. We can thus compute, for any region, the weighted average lo-

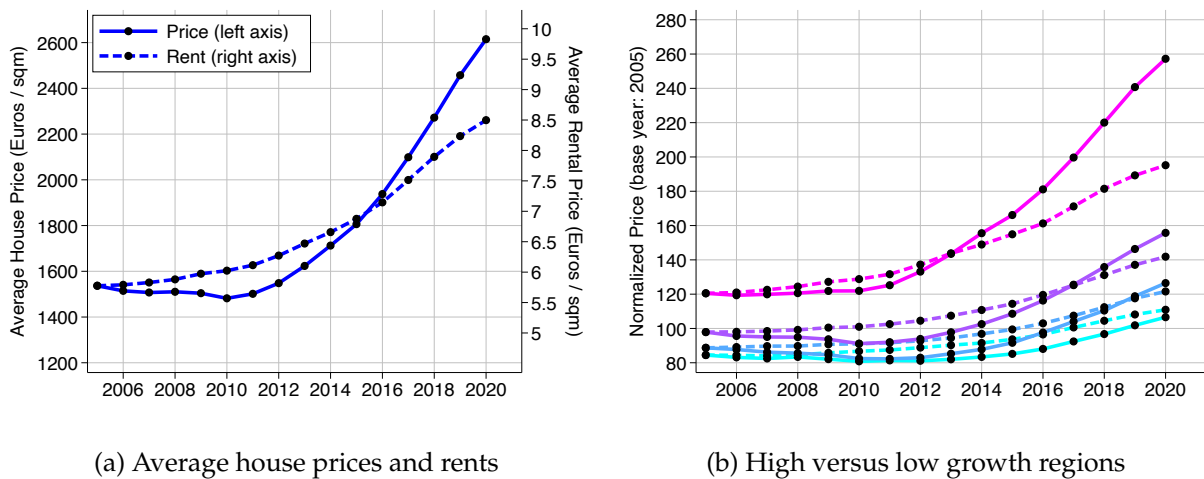
⁶Throughout, we use the term “county” to denote both “Kreise” – the administrative subdivision above a municipality – and “kreisfreie Staedte” – cities that are not part of any “Kreis” but provide the services of a “Kreis” at the municipality level.

⁷Reassuringly, when we weight county indices using Census data on county population sizes in the year 2014 instead of survey weights, we obtain very similar price indices.

cal rent or house price using the information on dwelling types of all survey respondents living in a region. The advantage of this approach – over weighting prices by transactions – is that the composition of our price indices resembles as closely as possible the composition of average forecasts in the survey. If the survey is indeed representative, then so are our actual price indices and average forecasts. In Appendix A we provide more information on the weighting procedures.

Figure 1 summarizes our data on house prices and rents. The left panel shows the behavior of Germany-wide aggregates: average house prices (solid line, left axis) and rental prices (dashed line, right axes) since the year 2005. Both price series are in Euros per square meter. During the years 2005-2020, house prices increased by 68 percent, while rental prices only increased by 46 percent. Consequently, the average price-rent ratio increased by almost 40 percent. We arranged the axes so that the price series start at exactly the same point in 2005 and growth rates are comparable between the two series. The graph shows that house prices stagnated until 2010, while rents were already growing since 2005. This combination led to a fall in the price-rent ratio from 2005 to 2010. After 2010, the growth rate of house prices outpaced the growth rate of rents, so that the price-rent ratio started to increase.

Figure 1: House and Rental Prices Across Regions



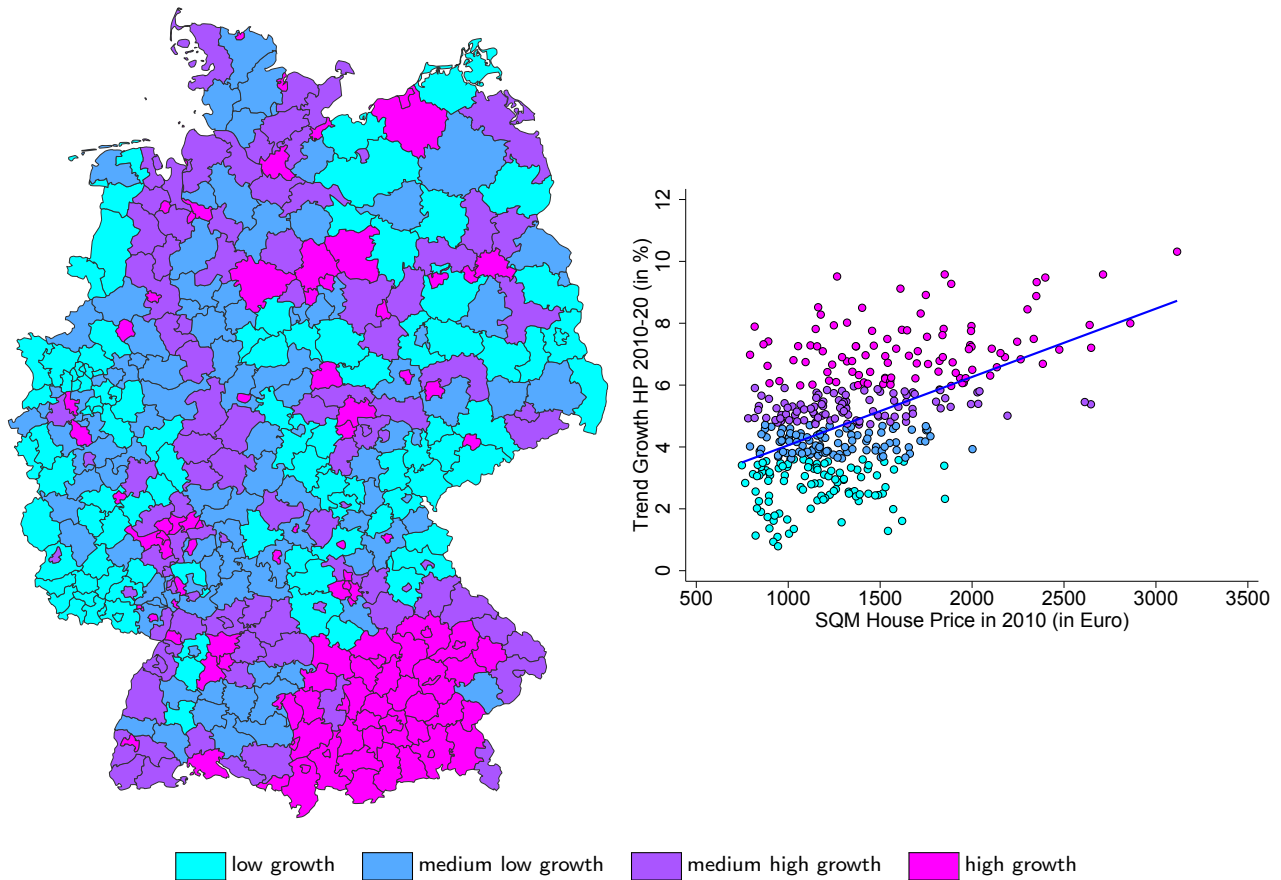
House prices are solid lines, rental prices are dotted lines. Panel (a) measures the national average of house prices on the left axis, rents on the right axis both in Euros per square meter. Panel (b) shows house prices in four growth regions (counties sorted into quartiles based on their trend growth over 2010-2020) normalized so that 100 corresponds to the national average in 2005. The rental series are normalized to the same value as the corresponding price series in 2005. *Source:* Own calculations based on data from bulwiengesa AG and PHF.

2.3 House price and rent data by region

Different regions in Germany have fared very differently during the housing boom. To document this regional heterogeneity, we divide the 401 counties into four “growth regions”, that is, quartiles according to trend nominal house price growth over the decade 2010-2020. Here trend growth is the slope of a linear regression of the log house price on time. The right panel of Figure 1 shows the evolution of house prices (solid lines) across growth regions since 2005, normalized so that a level of 100 corresponds to the national average in 2005. There are large regional differences in the strength of the boom: in the high growth region, house prices more than double during the boom, while in the low growth region house prices increase by merely 25%. Moreover, the boom worked its way slowly through the regions: it started earlier in the highest growth region, with the first strong growth rate in 2011, whereas in the lowest growth region prices began to pick up only in 2014. The figure also shows the evolution of rental prices (dotted lines). Here we have normalized the rent series for each region to start in 2005 at the same values as the house price in the same region, in order to allow for a simple comparison of price and rent within regions. In all regions, rents began to grow before prices, but have now been overtaken in all but the lowest growth region.

To understand which regions belong to the faster growing regions, we produce a map of German counties in the left panel of Figure 2, with different colors indicating differences in the extent of trend house price growth. (The color coding is the same as in the right panel of Figure 1). High house price growth is concentrated in larger cities, and in particular the metropolitan area of München and surrounding Bavarian cities such as Ingolstadt. Low growth is prevalent in more rural areas, but also the Ruhr area, the largest historical coal mining region in Germany now in decline. In the right panel of Figure 2, we correlate trend house price growth (measured on the vertical axis) with the house price at the beginning of the housing boom (in the year 2010, measured on the horizontal axis). Many regions that have been growing the most began with relatively high prices in the years before the boom. Thus, the boom was not a period of convergence for different regions, but instead amplified regional differences. Since the relative ranking in housing affordability of different areas has not changed during this decade, the media often focuses its coverage on the seven most expensive cities, which are referred to as the "Big Seven" (in English language for emphasis) and colored in pink in our map: München, Frankfurt, Hamburg, Berlin, Stuttgart, Düsseldorf, and Köln.

Figure 2: House Prices and Growth Across Germany



Source: Own calculations based on data from bulwiengesa AG and PHF.

2.4 The house price boom in historical perspective


The detailed county-level dataset for Germany covers a relatively short sample. To provide a historical perspective of housing markets in Germany, the green line in the left panel of Figure 3 shows a Germany-wide house price index going back to the 1970s. This "official" index is calculated by the Deutsche Bundesbank and also used, for example, by the OECD and the BIS. The underlying data consists of house prices in a sample of 50 German cities during the 1970s and 1980s. After reunification, the number of sample cities was extended to 100 and successively to 125. Finally, starting from 2005, the Bundesbank uses data on the full set of 401 counties from bulwiengesa AG and combines it with data from other providers. Because of these changes in methodology, there are several structural breaks in the series, which make it rather difficult to interpret its evolution over time.

To obtain a more consistent series, we use house price data for 50 German cities since

the 1970s from bulwiengesa AG, and select those 28 cities that are in the highest growth region during the current house price boom. The house price data only has prices for newly built apartments going back to the 1970s, which are always higher than prices for existing properties. We therefore use the county-level data to compute a short house price series for all properties in the 28 cities for 2005-2020. We then rescale our long series for newly built apartments so that it matches the short series on average (over the years 2005-2020). The left panel of Figure 3 shows the long house price data on newly built apartments (in blue) and the short series based on detailed data (in grey with black dots.) We see that the difference between the two series is only in levels. Rescaling the long-run newly built price series makes the two lines almost indistinguishable. It is reassuring that the long series is actually a good proxy for the overall housing market performance in these 28 cities. Our long series is similar to the official German aggregate house price index during the 1970s and 1980s, but the two series diverge somewhat starting in the 1990s and show a remarkably different behavior starting in 2005.

With our house price index for the same 28 cities, we can now interpret the evolution of house prices over time. The blue line in the left panel of Figure 3 shows that the current house price boom was an unusual development. Nominal house prices approximately doubled over the two decades between 1975 and 1995 and then stagnated for about 15 years. In the recent house price boom starting in 2010, house prices more than doubled over a single decade. This is an exceptional event in the history of German house prices. Note that, when looking at the official price series, one might come to a different conclusion, namely that the current house price growth is merely a return to long-run trend growth and not an exceptional boom. However, the consistent series for 28 cities in the highest growth region over the course of 45 years brings us to a different conclusion.

Figure 3: Long house price and rent data



Source: Own calculations based on data from bulwiengesa AG, PHF and OECD.

The right panel of Figure 3 shows the long-run rent data from bulwiengesa AG (for newly built and existing apartments) which coincides with the detailed data we have available for our 28 sample cities since 2005. The green line in this figure is the OECD's rent price index, where we normalized prices to the same level in 2005. The two rent series historically show quite some similarity, albeit the fact that bulwiengesa uses transaction data while the OECD calculates rents based on the rent component of the CPI. Only in the current boom do we see a strong divergence between the series, which is quite consistent with what we observe for prices. The reason is that our 28 cities are a select sample of what happened in Germany as a whole. However, they are a consistent sample over the course of 45 years, allowing us to analyze the strength of the German house price boom.


Figure 4 shows the house price-rent ratio (left panel) as well as rent growth (right panel) calculated from the price and rent series of the 28 cities presented above. During the inflation of the 1970s and '80s, the price-rent ratio was extraordinarily high (as in many other countries, see Figure 2 in Piazzesi and Schneider (2008)). Starting in the mid 1980s, the price-rent ratio began to fall continuously, owing to the poor performance in house price growth. This long episode of a falling price-rent ratio is only briefly interrupted by some turbulence in the rental market after German reunification in the 1990s. The price-rent ratio reached its trough at around 21.5 in 2010 and then increased to almost 30 in 2020, about a 40% increase. While extraordinary in speed, the 2010s boom in house prices still represents a recovery from a slump in the price-rent ratio since the Great Inflation. Rent growth, in the right panel of Figure 4, spiked once in the 1980s and then again quickly after reunification, as migrants from the East pushed into West German cities. This last episode of high rent growth, however, was accompanied by a "correction" period of negative nominal rent growth in the late 1990s. After the financial crises, rents again picked up growth, much earlier than house prices.

3 Household forecasts of house price growth

Our main source for households' price expectations is the *Panel on Household Finances (PHF)*, a representative survey of German households conducted every three years by the Deutsche Bundesbank. The survey collects data on household portfolios, consumption and income; the extent of detail on financial decisions resembles the US Federal Reserve Board's Survey of Consumer Finances. We mostly focus on the second survey wave which sampled 4,168 households in early 2014. We also perform some robustness checks drawing on 2017 data.

Our analysis makes use of two key strengths of the PHF. First, recent waves of the survey ask not only about actual decisions, but also contain a large number of questions on

Figure 4: House Price-Rent Ratio and Rent Growth in Germany



Source: Own calculations based on data from bulwiengesa AG and PHF.

expectations of asset prices, as well as on households' attitudes and investment plans. Second, we can study expectations by region. Indeed, the survey question at the center of our analysis asks households about *regional* house price growth over the next 12 months. We have access to a restricted version of the dataset that allows us to match households to the counties they live in, and hence the growth regions defined in Section 2.

Beliefs about regional house prices are elicited via a two-part question. First, respondents are asked to give a qualitative view about the direction of the housing market (question dhni0900):

What do you think, how will real estate prices in your area change in the next twelve months?

There are six candidate answers: (i) increase significantly, (ii) increase somewhat, (iii) stay about the same, (iv) decrease somewhat, (v) decrease significantly, and (vi) don't know/no answer. Second, respondents who give a positive or negative direction, that is, they answer (i), (ii), (iv) or (v) are then asked the follow-up quantitative question (dhni0950):

What do you think, by what percentage will real estate prices rise / fall in your area over the next 12 months?

We focus on households who have an opinion about the housing market and hence drop everyone who responds "don't know" to the first question; we come back to these households when discussing information sets in Section 4 below. Our quantitative measure of

price expectations then codes response (iii) as zero and uses the quantitative answer for the second question for all other households. After dropping the bottom and top 1 percent of house price forecasts from the distribution in order to guard against outliers, we are left with a total of 3,647 observations.

3.1 The cross section of household forecasts

In this section, we run regressions of household forecasts on household characteristics. The idea is to find out whether there are systematic differences of opinion between households along observable dimensions. We start with simple linear regressions and then proceed to nonlinear specifications. The dependent variable in all specifications is our quantitative measure of one-year-ahead expected house price growth in the region where the household lives. Average expected price growth among households in early 2014 was 3.1%, substantially below the subsequently realized growth rate of 5%.

Our regressors fall into three broad categories. First, we consider characteristics that typically serve as state variables in life cycle models of household behavior, in particular, age, wealth, income and housing tenure. Second, the PHF asks questions that elicit risk aversion, financial literacy and patience. In economic models, these behavioral traits would typically correspond to features of preferences. Finally, to account for heterogeneity in local housing markets, we include variables that capture geography and house quality. Appendix B describes all variables used in regression tables and figures in detail.

Table 1 reports results from linear regressions. Different columns correspond to different sets of regressors. To save space, we list regressors here only if their coefficients are significantly different from zero at the one percent level in at least one specification. Full results, together with clustered standard errors, are reported in Appendix C. Column (1) shows that typical drivers of savings and portfolio choice – age, income and wealth – are weakly correlated with household expectations. The regressors are dummies indicating ten year age bins for the age of the household head, income quartiles and net worth quartiles, as well as the number of household members and whether the household head is college educated. The main finding is that young (below age 40) and poor (in the bottom 25% by net worth) households expect about 1.5-2pp higher price growth compared to older and richer households.

Column (2) introduces three behavioral traits. The PHF measures self-assessed *risk aversion* by asking households to answer “Are you in general a risk-taking person or do you try to avoid risks?” on an eleven point scale from 0 (“not at all ready to take risks”) to 10 (“very willing to take risks”). Similarly, *patience* is the response to “Are you in general a person who is patient or do you tend to be impatient?” again on an eleven point

Table 1: The cross section of house price growth forecasts in 2014

	(1)	(2)	(3)	(4)	(5)
Demographics, Income, Wealth					
Age Group 30–39	1.365** (0.591)	1.306** (0.585)	0.945 (0.576)	0.703 (0.564)	0.367 (0.581)
1st Net Wealth Quartile	1.903*** (0.494)	1.920*** (0.501)	0.139 (0.566)	0.032 (0.551)	–0.269 (0.553)
2nd Net Wealth Quartile	0.939** (0.448)	0.996** (0.441)	–0.186 (0.466)	–0.136 (0.454)	–0.150 (0.459)
Behavioral Traits					
		yes	yes	yes	yes
Tenure					
Renter			2.445*** (0.382)	2.353*** (0.371)	2.100*** (0.379)
Growth Region					
Low				–2.073*** (0.439)	–1.570*** (0.428)
Medium Low				–1.572*** (0.431)	–1.443*** (0.433)
High				1.355*** (0.416)	1.006** (0.400)
Housing/Regional Characteristics					
City Center \geq 500k Inh.					1.808** (0.719)
Sqm size/100					–1.626*** (0.620)
(Sqm size/100) ²					0.429*** (0.128)
Number of Cases	3647	3646	3646	3646	3598
R ²	0.038	0.042	0.064	0.119	0.142

Robust standard errors in parentheses. **: $p < 0.05$, ***: $p < 0.01$

scale from “very patient” to “very impatient”. Finally, financial literacy is assessed by asking households three common questions that test their understanding of compound interest, nominal versus real rates of return, and portfolio diversification. We aggregate answers into a score equal to the number of correct answers, from zero to three. In a linear regression, adding behavioral traits marginally improves explanatory power, but otherwise does not change coefficients.

Column (3) introduces housing tenure, which is a strong predictor of price growth expectations. Indeed, households who rent the property they live in forecast almost 2.5pp higher price growth than homeowners. Moreover, much of the explanatory power from other demographics in columns (1) and (2) was due to the fact that young and poor house-

holds are more likely to rent – coefficients on age and net worth in column (3) are smaller in magnitude and lose significance. In fact, if we regress respondents' price growth forecasts on tenure alone, excluding all other variables, we also find a highly significant difference of about 2.5pp, as well as an R^2 of 0.05, close to the 0.064 in column (3).

A simple candidate explanation for this result is composition. Suppose that renters live predominantly in larger cities that also experience higher price growth. Suppose further that all residents of a city share the same opinion about local markets. Tenure might then appear like a good predictor of forecasts, because it is the best proxy for region, better than say, age or wealth. It is therefore important to control for region in order to understand the effect of tenure. An analogous composition effect might arise *within* regions. We know that housing booms often feature heterogeneity of capital gains by market segment – for example, lower quality houses, might become relatively scarce and increase more in price than the top end of the market. If those properties are in areas with more renters, we might naturally see higher price forecasts by renters. This effect calls for another set of controls that capture house quality.

Column (4) shows that tenure generates large significant differences in forecasts even controlling for region. We include dummies for three of our four growth regions, with the medium high growth region as base level. We do find that price growth forecasts align with local housing market conditions: residents of regions that have seen higher growth also forecast higher growth. The difference in average forecasts between the lowest and highest growth region is 3.4pp, about half the 6.25pp difference in realized growth rates between these regions over the year 2013, right before the survey was taken. Regional variation is important for explaining variation in forecasts: the R^2 now increases to 0.12. At the same time, the coefficient on tenure barely changes: unlike age and wealth, tenure is not a proxy for region. Appendix C shows similar results for a specification with more granular controls for the household's past house price growth experience.

Column (5) shows that the result persists when we add additional regional and housing characteristics. In particular, we control for the time span the household already lived in its current residence, the community size on a 10 point scale, whether the household lives in a city center or in the periphery, whether the building the household lives in needs renovation and a general rating of the dwelling quality on different levels in between "Very Simple" and "Exclusive". The one variable that plays a sizeable role in accounting for forecast variation is whether the household lives at the center of a very large city. Large city dwellers forecast higher price growth, and since they often rent, this reduces the coefficient on tenure slightly, to just above 2pp.

We also have data on the square meter size of the house or apartment, and column (5) includes both size and its squared value to accommodate possible nonlinear effects. While both coefficients are significant, the overall relevance of size is minor. The average size


of a residence in our sample is around 100 square meters with most variation between 50 and 200 square meters. The coefficients thus imply a significant negative relationship between the size of the household’s residence and house price forecasts. At the same time, size adds little explanatory power. If we include only square meter size and its squared value into the regression on top of the regressors in column (4) the R^2 only increases to 0.126, indicating that the predictive power of the size of the residence is limited. We conclude that tenure and local housing market conditions subsume the effects of many other observables such as age, income, wealth, risk aversion and financial literacy. Hence, they act as sufficient statistics for explaining cross-sectional differences in house price growth forecasts.

3.2 Forecast means and dispersions by growth region

In this section, we take a closer look at the distributions – mean and dispersion – of renter and owner forecasts by growth region. The figures here interact the key predictor of forecasts established in the previous sections. They also provide us with targets that our quantitative model below will be required to match.

Figure 5 shows *average forecasts* by region and how they compare to realized price growth. The wide red bars in the left panel are average renter forecasts in the four regions; whiskers indicate 95 percent confidence intervals. Narrow bars represent realized price growth in the respective growth region. Any differences in the height of bars for a region reflect forecast errors made by the average renter in 2014. The right panel repeats the exercise for owners, whose average forecasts are wide yellow bars; the narrow bars represent the same realizations as in the left panel.

Figure 5: House Price Growth Forecasts by Tenure and Growth Region in 2014




Source: Own calculations based on data from bulwiengesa AG and PHF.

The figure summarizes three robust patterns. First, households generally underpredict house price growth. With the exception of renters in the lowest growth region, all average price growth forecasts lie below realized price growth, and most of them are significantly different from realized growth. Second, households' forecasts are consistent with regional differences in the sense that forecasts in low growth regions are lower than those in high growth regions, regardless of tenure status. Third, within each of the different growth regions, we find that renters consistently make higher forecasts than owners.

Figure 6 considers the *cross-sectional mean squared error (MSE)* in households' forecasts. Again the left panel shows renters and the right panel shows owners, each with four bars for the four growth regions. An individual household's forecast error is defined as the squared difference between realized growth – common to all individuals in the region – and the individual forecast. The MSE for a group of households can therefore be decomposed into two parts: the squared average forecast error – indicated by light colors in the figure – and the (cross-sectional) variance of the forecasts. The squared average forecast error reflects the mistake made by the average owner or renter, as shown already in Figure 5. The variances reflect differences of opinions within the groups of renters and owners.

Figure 6: Cross sectional MSE of House Price Growth Forecasts in 2014



Source: Own calculations based on data from bulwiengesa AG and PHF.

In almost all growth regions, renters exhibit a larger MSE than owners. The result is entirely driven by the wide dispersion in renters' individual forecasts. As we have seen above, renters' average forecast is much closer to the actual growth realization everywhere. At first sight, this finding might speak against an informational explanation for the forecast differences observed between renters and owners in the data. Indeed, in a simple model where agents predict an unknown parameter from noisy signals, the unconditional MSE of a better-informed agent (that is, an agent with a more precise signal)

is always below that of a less-informed agent. However, the result found here is about errors conditional on a particular realization. Renters, even though their information is more noisy, can thus be “in the right place at the right time” during a boom that features strong rent growth. Our model below formalizes this point and shows how learning from different information can jointly rationalize the empirical patterns of forecast means and dispersions in the data.

3.3 Expectations of renters and owners over time

So far, we only looked at results from wave 2 of the PHF in 2014. Yet, there are additional data available that allow us to track the differences between households of different housing tenure over time. As the patterns we have shown in this section – as well as others that follow below – are qualitatively very similar across growth regions, it is useful to aggregate findings by region into single numbers for Germany as whole. This allows us to streamline the exposition by presenting summary numbers at the national level, rather than always showing each region separately. However, aggregation must take into account that forecasts reflect regional growth, and renters and owners are not equally distributed across Germany. We therefore apply a reweighting procedure for aggregate statistics that eliminates the effects of an unequal distribution of renters and owners across different growth regions, see Appendix C.5 for details.

Wave 3 of the PHF asked households the same forecasting questions in the year 2017, where the house price boom had already arrived in all German regions. In addition, we can draw on data from the pilot of the *Bundesbank Online Survey - Households* (BOP-HH), a survey that was initiated by the Deutsche Bundesbank in 2019. This representative survey again asks respondents about their forecasts of regional price growth, collects their tenure status and can be matched to our local house price growth database.⁸

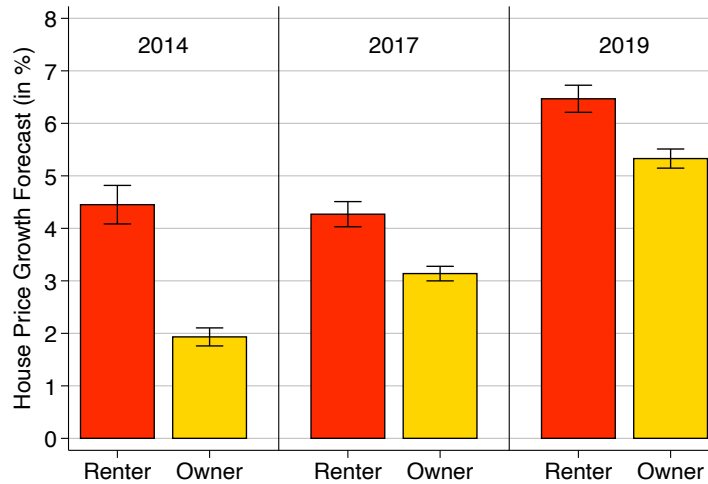
Figure 7 shows the average house price growth forecasts of renters and owners over time. The first thing we see is that price forecasts increase especially in the group of owners, but also in the group of renters. As house prices grow for an extended period of time, households seem to adapt their expectations about the future accordingly. While the gap between forecasts of renters and owners narrows a bit between 2014 and 2017, there is still a sizable difference of more than one percentage point left in 2017 and 2019. Hence, the fact that renters make significantly higher price growth forecasts than owners persists over time and across different surveys.⁹ Summing up, this section leaves us with two

⁸The survey has much less details compared to the PHF when it comes to household characteristics, income and wealth. The questions about house price expectations, however, were framed in exactly the same way. We process the data so that it is comparable to the PHF.

⁹In Appendix C.4 we also provide the results of cross-sectional regressions for the 2017 wave of the PHF.

robust stylized facts: (i) across regions of different house price growth and across time, renters make significantly higher price growth forecasts than owners, and (ii) while being more accurate on average, the forecasts of renters are vastly more dispersed than those of owners.

Figure 7: House Price Growth Forecasts by Tenure Over Time



Source: Own calculations based on data from PHF and BOP-HH.

4 Direct evidence for different information sets

The facts presented in the previous section lead us to explore differences in information sets as a possible explanation for differences in house price forecasts of renters and owners. In particular, our theory postulates that renters have better information about the dividends of housing than owners. We now use additional Bundesbank data to provide more direct evidence on information sets. These facts will then guide specific assumptions we make in our model below. We proceed in three steps. Section 4.1 provides evidence on how households get their information about housing, again emphasizing the role of direct price and rent observations. Section 4.2 argues that owners who are also real estate investors and therefore get better signals about rents make forecasts that are comparable to renter forecasts. Finally, Section 4.3 shows that owners are well aware of current house price growth, but with large confidence bounds.

4.1 Sources of information

To find out more directly how households acquire information about housing markets, we draw on a new question we proposed for the pilot of the Bundesbank's Online Survey – Households (BOP-HH). The BOP-HH was first fielded in spring 2019; it elicits less detail on household income and balance sheet, but instead focuses on expectation formation. Question 306 of the BOP-HH survey reads:

How important are each of the following sources of information for you to evaluate future house prices?

Respondents are presented the following seven potential information sources:

1. Relatives, friends and neighbors
2. Classical media (newspapers, tv, etc.)
3. Social media (like Facebook and Twitter)
4. Online real estate platforms
5. Financial consultants
6. Direct observations of rents in your neighborhood
7. Direct observations of house prices in your neighborhood

For each candidate source, households can check one of four intensities, “Not important at all”, “Somewhat important”, “Quite important”, and “Very important”. To summarize average choices by a single numerical score, we code these intensities as 0, 33, 66, and 100, respectively. Table 2 lists the seven answers in order of importance, measured by the fraction of all respondents who labeled an answer either “Quite important” or “Very important”, reported in the first column. The second and third columns report the average score of the information source among renters and owners, respectively, and the final column measures the difference.

The primary source of information for households to forecast house price growth is the direct observation of prices. In fact, more than 80 percent of households look to rents and almost 80 percent to house prices. Online real estate portals, friends, advisors and especially social media are much less important sources of information. The magnitude of the scores for the seven sources are broadly similar across renters and owners. There are however two significant differences. First, while all households rely on direct observations, renters look more at rents, whereas owners look more at prices. Second, renters rely more on social media, online real estate portals and especially family and friends when gathering information. We view the results as broadly supportive of an approach that emphasizes different information sets that relate to current market experience.

Table 2: Sources of information by tenure

Source	Perceived Important by	Coded average		Difference
		Renters	Owners	
Direct observation of rents	83.34	72.02	69.25	-2.773*** (1.061)
Direct observation of prices	78.24	66.04	69.83	3.785*** (1.148)
Classical Media	73.53	60.32	60.37	0.057 (0.993)
Online Real Estate Portals	66.76	57.72	55.29	-2.432*** (1.108)
Family & Friends	52.50	52.62	48.11	-4.508*** (1.101)
Financial Advisors	46.39	44.68	45.47	0.791 (1.158)
Social Media	12.26	23.81	20.86	-2.963*** (0.943)

4.2 Forecasts by real estate investors

Our model assumes that renters receive better information about rents than owners. Now, households who own their main residence may also own other real estate that they rent out. While these households are owners, they have more information about rents than other owners, similar to renters. From the perspective of our theory, the forecasts for these owner-landlords should be higher than the forecasts or regular owners, and they should be comparable to renter forecasts. In this subsection, we present evidence that supports these predictions.

It is not obvious how to best identify owner-landlords in the survey. Some households who own their main residence also own a vacation home that they rent out occasionally. Other households own property that they leave vacant, or they only rent out to family members at below market rents. We therefore define owner-landlords to be those households who own their primary residence, own at least one additional property (apartment, single-family home, or townhouse), and receive some minimal rental income (more than 650 Euros per month). Below, we use alternative definitions that try to identify more serious real estate investors.

Column (1) in Table 3 repeats our baseline finding from Table 1 about the importance of tenure and region for the cross section of forecasts. To save space, we omit the estimated coefficients on the growth region dummies. Column (2) in Table 3 shows that owner-landlords indeed make higher house price growth forecasts (by 0.42 percentage points)

than other owners. However, the coefficient is not precisely estimated. One issue is that there are not many households who own their main residence and also other real estate. In the survey, only 12 percent of households are owner-landlords. A second issue may be that the additional property owned by an owner-landlord can be located in a different region in Germany.¹⁰ As such, the rent information an owner-landlord receives from her or his investment property may carry only little information about local house price movements.

Table 3: Forecasts by real estate investors

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Tenure</i>							
Renter	2.087*** (0.378)	2.107*** (0.380)	2.125*** (0.381)	2.109*** (0.380)	2.099*** (0.381)	2.097*** (0.380)	2.095*** (0.380)
<i>Is Owner-Landlord</i>							
Yes		0.422 (0.523)	0.021 (0.585)	0.115 (0.533)	0.092 (0.403)	-0.053 (0.386)	-0.256 (0.415)
<i>Is Owner-Landlord × Has above median LTV on other property</i>							
Yes × Yes			1.893 (1.253)				
<i>Is Owner-Landlord × Has more than one rental property</i>							
Yes × Yes				0.958 (1.214)			
<i>Is Owner-Landlord × Rental income share</i>							
Yes × ≥ 0.3					1.096 (1.410)		
Yes × ≥ 0.4						1.963 (1.667)	
Yes × ≥ 0.5							3.586** (1.693)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Number of Cases	3598	3598	3598	3598	3598	3598	3598
R-Square	0.142	0.142	0.142	0.142	0.142	0.143	0.144

Robust standard errors in parentheses. **: $p < 0.05$, ***: $p < 0.01$

To identify serious real estate investors, we consider the following alternatives. Column (3) defines investors as owner-landlords who have an above median loan-to-value ratio on their investment property. The estimated coefficient is comparable to the coefficient on the renter dummy, but again standard errors are high. Column (4) shows that owner-landlords who have more than one rental property also make higher forecasts. Columns


¹⁰While our data allows to identify the location of the primary residence of an owner, we do not have information on the location of investment property.

(5), (6) and (7) isolate owner-landlords for whom the rental income is an increasing share of total household income, ranging from at least 30, 40 to 50 percent. Investors with higher rental income shares make higher forecasts. Moreover, the estimated coefficients for investors with a rental income share above 40 percent are comparable to those of renters, and also significant. We conclude that forecasts for serious real estate investors are quantitatively in the same ballpark as renter forecasts.

4.3 Do owners keep track of house prices?

One central ingredient of our model analysis is that owners track the evolution of local prices. We can shed some light on the validity of this assumption by checking owners' perceptions of the value of their own residence over time. Indeed, a subset of the PHF survey is organized as a panel, following households across waves from 2011 to 2014 and again from 2014 to 2017. In each wave, households are asked to estimate the hypothetical sales price of their current main residence. We match those price estimates for all panel households who did not move between two survey waves. This allows us to calculate a *perceived home price growth rate* for all owners over both the years 2011-2014 and 2014-2017.

Figure 8: Information quality of opinionated owners



Source: Own calculations based on data from PHF.

Figure 8 compares annualized mean perceived home price growth with realized growth rates in the owner's region over the same time span. The left panel shows 2011-14 and the right panel shows 2014-17; in both cases, whiskers on the left hand yellow bars indicate 95% confidence intervals for the mean perceived home price growth rate. Confidence intervals are large, consistent with imperfect information on the part of owners. At the

same time, for both time spans, the realized growth rate lies within or at least very close to the confidence band of the mean perceived growth rate. We thus take away that the *average* owner is aware of local housing market conditions, and that an explanation for differences between owners and renters should not rely on owners simply receiving no market signals at all.

Summing up, the evidence in this section suggests that information sets are important for individual house price growth forecasts. While surveyed households report that observing prices and rents is the most important source of information for forecasting house prices, renters report to look more at rents while owners look more at prices. Owner-landlords who receive both rent and price signals make house price forecasts that are much more aligned with the forecasts of renters. Finally, our data shows that owners are well aware of past house price movements, albeit with significant noise.

5 A model of learning about housing cost

In this section, we develop a simple model that describes the joint distribution of prices, rents, and household forecasts in line with the empirical observations made in the previous sections. Its goal is to show how learning with different information sets can naturally lead to the large differences in the cross-sectional distribution of forecasts that we see in the data. We do not explicitly model household decisions, but only specify the information sets of owners and renters and compute their conditional expectations. We only impose one piece of structure: there is a connection between rents and prices because a developer sector arbitrages between rental and owner occupied houses. Appendix D derives the equations below.

Asset pricing. We consider the valuation of houses in a region. Developers can sell housing units at a price P_t or rent them out at the rental price R_t , both denominated in Euros. Prices and rents should be thought of as regional averages, as in our data presented above. At any date t , developers are indifferent between selling a house at date t today or holding it for one period and receiving rent as “dividend”. The equilibrium house price P_t thus satisfies a standard intertemporal Euler equation

$$P_t = E_t [\tilde{M}_{t+1} (P_{t+1} + R_{t+1})], \quad (1)$$

where E_t is the conditional expectation operator and \tilde{M}_{t+1} is the stochastic discount factor of the developer.

We emphasize that the Euler equation (1) allows for a wide range of frictions and behavior

by developers (or investors in developer firms). The special case where \tilde{M}_{t+1} is perfectly foreseen one period in advance is the familiar “user cost model” for frictionless housing markets under rational expectations, $M_t := E_t [\tilde{M}_{t+1}]$ is the one-period ahead nominal bond price. More generally, the stochastic discount factor may capture (i) risk attitude of developers, (ii) financial frictions that affect the developer’s cost of capital, or (iii) differences in beliefs between developers who value houses and outside observers who know the current price and rent. Large literatures have documented the relevance of time variation in (i) – (iii) on many asset prices. We do not take a stand on what force is most important; all that matters below is that households evaluate the Euler equation with expectations E_t conditional on their time- t information, and they contemplate movements in \tilde{M}_{t+1} as a source of price volatility.

To deal with trends, it is helpful to work with rent growth and price-rent ratios. We denote the gross growth rate of rents by $G_t = R_t/R_{t-1}$ and define the price-rent ratio as $V_t = P_t/R_t$. We can then rewrite (1) as

$$V_t = E_t [\tilde{M}_{t+1} (V_{t+1} + 1) G_{t+1}]. \quad (2)$$

We follow standard practice in assuming that \tilde{M}_t and G_t are jointly stationary, and focus on stationary solutions V_t to the difference equation (2). In other words, rents and prices follow a stochastic trend, but they are cointegrated so the price-rent ratio is stationary. Movements in the stochastic discount factor or the growth rate may therefore lead to divergence of price and rent growth in the short or medium run, but not in the very long run.

To implement our learning model, we use a log-linear approximation of (2) around its deterministic steady state. Suppose that the mean stochastic discount factor is M and the mean growth rate is G , with $MG < 1$. The steady state price-rent ratio is then $V = MG / (1 - MG)$. Denoting logarithms by small letters, the steady state log growth rate is $g = \log G$ and the log price-rent ratio is $v = \log V$. Denoting log deviations from steady state by hats, we obtain the linear difference equation

$$\hat{v}_t = \hat{m}_t + E_t [MG\hat{v}_{t+1} + \hat{g}_{t+1}], \quad (3)$$

where $m_t = E_t [\tilde{m}_{t+1}]$ is the predictable component of the log discount factor and \hat{m}_t is its deviation from steady state. Appendix D.1 contains a detailed derivation of this equation.

The dynamics of prices and rents. We choose functional forms for the stochastic discount factor and growth rate that capture key features of the data presented in Section 2 but at the same time allow for easy application of Bayes’ rule. In particular, deviations from the mean in both the growth rate of rents and the predictable component of the discount

factor are described by Gaussian AR(1) processes

$$\begin{aligned}\hat{g}_t &= \alpha_g \hat{g}_{t-1} + \varepsilon_t^g, \\ \hat{m}_t &= \alpha_m \hat{m}_{t-1} + \varepsilon_t^m.\end{aligned}\tag{4}$$

The innovations ε_t^g and ε_t^m are serially as well as mutually uncorrelated and normally distributed with mean zero. Both processes are persistent but stationary, $\alpha_g, \alpha_m < 1$. The distribution of rent growth can be directly estimated from the data. Persistence in the growth rate allows in particular for an acceleration of rent growth that pushes rents to a permanently higher level.

The role of the persistent discount factor shock ε_t^m is to allow for “excess volatility” in house prices: the price-rent ratio can move even if there is no news about current or future rents (that is, no “cash flow news”). To see this, we solve (3) by the method of undetermined coefficients to obtain the stationary solution:

$$\hat{v}_t = \beta_m \hat{m}_t + \beta_g \hat{g}_t,\tag{5}$$

where $\beta_m = 1 / (1 - \alpha_m MG)$ and $\beta_g = \alpha_g / (1 - \alpha_g MG)$ are positive coefficients. Prices are high relative to rents when developers either (i) discount the future at a lower rate or (ii) expect unusually high growth in rents. Since both types of fluctuations are mean-reverting, their impact on prices depends on their persistence relative to the duration of houses, captured by MG .

Putting together the trend and fluctuations around it, we now summarize the joint dynamics of (log) rents and prices in the region by

$$\begin{aligned}p_t &= v + r_t + \hat{v}_t, \\ r_t &= g + r_{t-1} + \hat{g}_t.\end{aligned}\tag{6}$$

Both prices and rents grow on average at the rate g . Movements in \hat{g}_t induce transitory deviations from trends in both variables. In contrast, movements in \hat{m}_t can drive prices to move above or below trend even if rents simply grow at the trend growth rate. The model thus allows for two types of booms discussed in the literature: swings in \hat{m}_t capture changes in interest rates, credit conditions or investor sentiment, whereas swings in \hat{g}_t capture changes in actual rent growth.

To study household learning, it is helpful to have concise vector notation for the dynamics of prices and rents. We thus define a state vector $x_t = (\hat{m}_t, \hat{g}_t, r_t, 1)^\top$ that contains the predictable component of the discount factor as well as the stochastic components of rents

and the growth rate. We can then represent the distribution of rents and prices as

$$\begin{pmatrix} p_t \\ r_t \end{pmatrix} = Bx_t; \quad x_t = Ax_{t-1} + C\varepsilon_t, \quad (7)$$

for some matrices A , B and C , where ε_t is a 2×1 vector of iid standard normal innovations see Appendix D for a detailed description of the state space system.

Information structure and forecasts. Consider now owner and renter households who answer survey questions. We want to capture the idea that they sample the rents and prices of a few individual dwellings and hear about others from friends or neighbors. We thus assume that an individual household i who is either an owner (type $h = o$) or a renter (type $h = r$) observes a vector of noisy signals of the current average (log) price and rent in the region

$$s_t^{i,h} = \begin{pmatrix} p_t \\ r_t \end{pmatrix} + w_t^{i,h}, \quad h = r, o. \quad (8)$$

Here $w_t^{i,h}$ is a 2×1 vector of idiosyncratic Gaussian shocks with mean zero that are mutually uncorrelated and iid in the cross section of individual households. Their covariance matrix \mathcal{W}^h depends on the household type h : for example, owners may receive less precise (or more noisy) signals about rents relative to prices, and vice versa for renters.

We assume that households know the distribution of our average price and rent data (p_t, r_t) and their own signals $s_t^{i,h}$. We identify an individual's survey forecast with the conditional expectation of average price growth given that individual's history of signals and their initial view of the state. Household i believes that the initial state is $x_0^{i,h}$. This initial nowcast is drawn from a normal distribution that depends on the household type h through its mean \bar{x}_0^h (the average initial nowcast of type h agents) and the covariance matrix Ω_0^h (their cross sectional dispersion). The price growth forecast at date t is then

$$f_t^{i,h} = E \left[\Delta p_{t+1} | s_t^{i,h}, s_{t-1}^{i,h}, \dots, s_1^{i,h}, x_0^{i,h} \right]. \quad (9)$$

Since the data, signals and the initial belief are all jointly normally distributed per (4)-(8), subsequent beliefs are also normally distributed and forecasts can be computed via the Kalman filter.

Given the structure of the system (7), we can choose the variance Σ_0^h of the initial nowcast error $x_0 - x_0^{i,h}$ such that the forecast error variance of type h households is time invariant. Intuitively, this works because households track a persistent hidden state by observing noisy signals. Every signal contains information that lowers uncertainty about the state, but also adds additional noise. When the two forces balance, uncertainty about the state

as well as forecast error variances are constant. Standard results further imply that this choice of initial variance is what one would obtain if agents had seen an infinite sequence of past signals. We make this choice throughout in our quantitative application below – it captures in a parsimonious way the idea that agents are uncertain even at the beginning of our sample.

Characterizing the distribution of nowcasts and forecasts. Given an initial cross sectional distribution of beliefs about the state – that is, a cross section of $x_0^{i,h}$ s – as well as a realization of the data (p_t, r_t) , our model generates a panel of forecasts that we can match to our survey data. The key to understanding the dynamics of forecasts is the evolution of individual “nowcasts” of the current state, denoted $x_t^{i,h} := E \left[x_t | s_t^{i,h}, s_{t-1}^{i,h}, \dots, s_1^{i,h}, x_0^{i,h} \right]$. Indeed, (7) implies that this nowcast is a sufficient statistic for forecasting future prices given an agent’s past information. In particular, up to a constant, the price growth forecast is $f_t^{i,h} = B_{1\bullet}(A - I)x_t^{i,h}$, where $B_{1\bullet}$ is the first row of the matrix B .

The law of motion of the nowcast takes the standard form

$$x_t^{i,h} = Ax_{t-1}^{i,h} + \Gamma^h \left(s_t^{i,h} - BAx_{t-1}^{i,h} \right), \quad (10)$$

where the “gain matrix” Γ^h is constant because of our choice of initial conditions. Nowcasts – and hence forecasts – are updated according to a time invariant rule. To arrive at a nowcast for date t , agents start from the date $t - 1$ forecast of the state $Ax_{t-1}^{i,h}$. Upon receiving signals, they make an adjustment depending on the last forecast error, the term in parentheses. We note that errors occur not only because of the new realization of the data (p_t, r_t) that is common to all agents, but also because of the noise in agents’ signals.

What is the *average* forecast made by agents of type h ? Since forecasts are linear in nowcasts, they depend on the evolution of the average nowcast \bar{x}_t^h of type h agents. By the law of large numbers, the noise in type- h signals washes out in the average, and we obtain a recursion for the average type- h nowcast:

$$\bar{x}_t^h = A\bar{x}_{t-1}^h + \Gamma^h \left(\begin{pmatrix} p_t \\ r_t \end{pmatrix} - BA\bar{x}_{t-1}^h \right). \quad (11)$$

Given an initial average nowcast \bar{x}_0^h for type h , the current average nowcasts – and hence also current rent and price growth forecasts – are deterministic functions of the data (p_t, r_t) . This is the relationship we use below to link observed average forecasts to the observed house price and realizations.

What is the *cross sectional dispersion* of nowcasts? Suppose the date $t - 1$ cross sectional

variance of nowcasts $x_{t-1}^{i,h}$ is Ω_{t-1}^h . From (10), the date t variance is then

$$\Omega_t^h = \left(I - \Gamma^h B \right) A \Omega_{t-1}^h A^\top \left(I - \Gamma^h B \right)^\top + \Gamma^h \mathcal{W}^h \Gamma^{h\top}. \quad (12)$$

The first term reflects the adjustment of nowcasts due to the information conveyed by prices and rents. Since the same data realizations affect all agents, this tends to reduce the dispersion of nowcasts. The second term reflects new date t noise which increases dispersion. We can choose as our initial dispersion Ω_0^h the fixed point of (12) at which the two forces balance. As a result, the dispersion of nowcasts (and hence forecasts) is constant over time.¹¹ The idea is to parsimoniously capture disagreement about nowcasts among individuals at the beginning of our sample, much like our initial variance for individual beliefs captures uncertainty initially perceived by an individual.

Information in prices and rents. Why do renters make higher price growth forecasts in the current German housing boom? To see how learning accounts for this fact, consider price growth under the subjective belief of household i of type h . We can decompose it into a forecast – based on elements of the nowcast vector $x_t^{i,h}$ – as well as an orthogonal forecast error

$$\Delta p_{t+1} = g + \frac{\alpha_m - 1}{1 - \alpha_m MG} \hat{m}_t^{i,h} + \frac{(\alpha_g)^2 (1 - MG)}{1 - \alpha_g MG} \hat{g}_t^{i,h} + u_{t+1}^{i,h}.$$

Here the forecast error reflects both the nowcast error due to imperfect learning up to date t and the new innovations that affect actual prices and rents at date $t + 1$.

The two forces that can generate housing booms in our model thus affect price growth forecasts in opposite directions. Indeed, the coefficient on the nowcast of the discount factor $\hat{m}_t^{i,h}$ is negative: when agents perceive a housing slump due to financial frictions, say, they expect mean reversion and hence forecast high price growth. At the same time, the coefficient on the growth rate $\hat{g}_t^{i,h}$ is positive: agents also predict high price growth when they perceive rents to be rising. Averaging across agents, we have that renters have higher price growth forecasts than owners if they perceive higher rent growth or a lower discount factor.

Now consider an owner at the end of the financial crisis, who initially believes the discount factor and the growth rate of rents are both below average, so both forces contribute to a low price-rent ratio by (5). The owner thus expects prices to rebound from the crisis. For concreteness, suppose further that the owner perfectly observes prices, but not rents. He thus initially sees a recovery with low price growth. As a result, he adjusts his nowcast

¹¹The resulting cross sectional variance of nowcasts is lower than the posterior variance of any individual's nowcast. This is because the latter incorporates uncertainty about the time series evolution of the hidden state, whereas the former reflects only noise in signals.

in the direction dictated by the forecast error – towards lower rent growth and a higher discount factor which both imply a lower price forecast. Contrast this with a renter who also observes rent. As he sees strong rent growth, he attributes the sluggish movement in prices to a low discount factor. Both high rent growth and a low discount factor lead the renter to forecast higher price growth.

What changes with noise? The argument so far has only used average responses to price and rent realizations. A version of it will thus go through when agents observe rents and prices with error; all that changes is that average responses to signals are weaker when signals are less precise. The new feature with noise is that the model can also speak to dispersion of forecasts. There are two opposing forces here: less precise signals make responses weaker, but noise also makes signals more dispersed. Higher noise can therefore make forecasts more or less dispersed. In fact, at the extremes of perfect and completely uninformative signals, forecast dispersion is zero, whereas it is positive at intermediate dispersion.

We also want to understand how renters make more accurate forecasts on average even though as group they make larger mean squared forecasts errors. The forecast error differs from (minus) the forecast only by a constant common to all agents, namely the new price growth realization. The mean squared error can therefore be decomposed into the dispersion of forecasts and the squared bias, that is, the average forecast error:

$$MSE_t^h(\Delta p_{t+1}) = B_{1\bullet} (A - I) \Omega^h (A - I)^\top B_{1\bullet}^\top + \left(\Delta p_{t+1} - B_{1\bullet} (A - I) \bar{x}_t^h \right)^2, \quad (13)$$

where Ω^h is the posterior variance of nowcasts, the fixed point of (12). The first term depends on the information structure only. The second term, in contrast, depends on the specific realization of the two factors driving prices and rents. In particular, if the boom is driven by rent growth that is better observed by renters, then it is possible for the average renter to be closer to the actual realizations, even though he has less precise signals about prices.

Quantitative analysis. We now explore whether the mechanism sketched above can quantitatively account for the mean and variance of the cross section of price growth forecasts in the 2014 survey. We focus on learning in the top growth region, which represents the biggest challenge for our model. We use the data going back to the 1970s on 28 German cities in the top growth region for which we have prices and rents available, see the discussion in Section 2.4. Based on the time series in Figure 4, we estimate the distribution of price-rent ratios and rent growth. For the recent house price boom, we use our high quality price and rent data (p_t, r_t) from bulwiengesa AG to construct prices and rents for these cities. Our focus on the top growth region has the advantage that its boom was leading

the rest of Germany, so the assumption that households learn from local prices and rents is appropriate. In other regions, there may also be regional spillovers in learning.

We pick parameters for the agents' perceived distribution of prices and rents as well as for their information sets in two steps. A first set of parameters can be fixed up front. The model period is a year and we focus on the early boom years starting at the onset in 2010 up to 2014, where the first expectations data is available. We set the average (gross) rent growth rate to $G = 1.033$, the long-run average computed from the data on the top growth region in Figure 4. We then choose the average discount factor $M = 0.932$ to obtain the average price-rent ratio in the data of 25.8. Estimating an AR(1) process for rent growth in the top German cities delivers an AR(1) parameter $\alpha_g = 0.72$ and an innovation volatility of 2.58%. We thus impose in particular that agents correctly perceive the unconditional volatility of rent growth in the data which is 3.72%. Furthermore, we assume that agents have identical beliefs about the persistence of shocks to rents and shocks to the discount factor $\alpha_m = \alpha_g$. Therefore, the process for the price-to-rent ratio is also AR(1) with an autoregression parameter of 0.72. With these parameter choices, the coefficient on rent growth in (5) is pinned down at $\beta_g = 2.36$ and the coefficient on discount factor shocks is $\beta_m = 3.27$. As a result, households' beliefs are consistent with the stylized fact of "excess volatility" of house prices.

We select a second set of parameters by matching a set of target moments. For beliefs, we need to choose noise variances as well as a distribution of priors for each household type. To discipline our model, we make two assumptions. First, we assume priors at the beginning of the year 2010 are such that average households of both types (i) agree on the nowcast of the rent growth rate and (ii) correctly nowcast the beginning of 2010 price-rent ratio. The assumption that the means \bar{x}_0^h are identical across types h implies that any average disagreement between types observed in 2014 must be explained by learning over the boom years 2010 to 2014, and not by disagreement at the beginning. Second, our baseline specification assumes that renters observe rent without noise, and that owners do not observe rents at all. In other words, the noise variances on rent for renters and owners are zero and infinity, respectively. This leaves two noise variances for house price observations of owners and renters, respectively.

Parameters and targets for the moment matching exercise are collected in Table 4. In addition to the parameter values for the prior and noise – the last three lines – we need to find the innovation volatility for the discount factor m . We match the average forecasts and the cross sectional dispersion of forecasts in 2014 of all sample households living in the 28 cities in the top growth region from which we calculate prices and rents. Consequently, renters make higher forecasts than owners while at the same time their forecast volatility is higher. The most subtle point perhaps is how forecast moments help identify the perceived volatility of the discount factor. The argument here is that observed signals

can only lead to sufficient disagreement if they carry enough news about the future. This leads the calibration to a perceived discount factor that has relatively large innovations at about 12%.

Table 4: Model Results

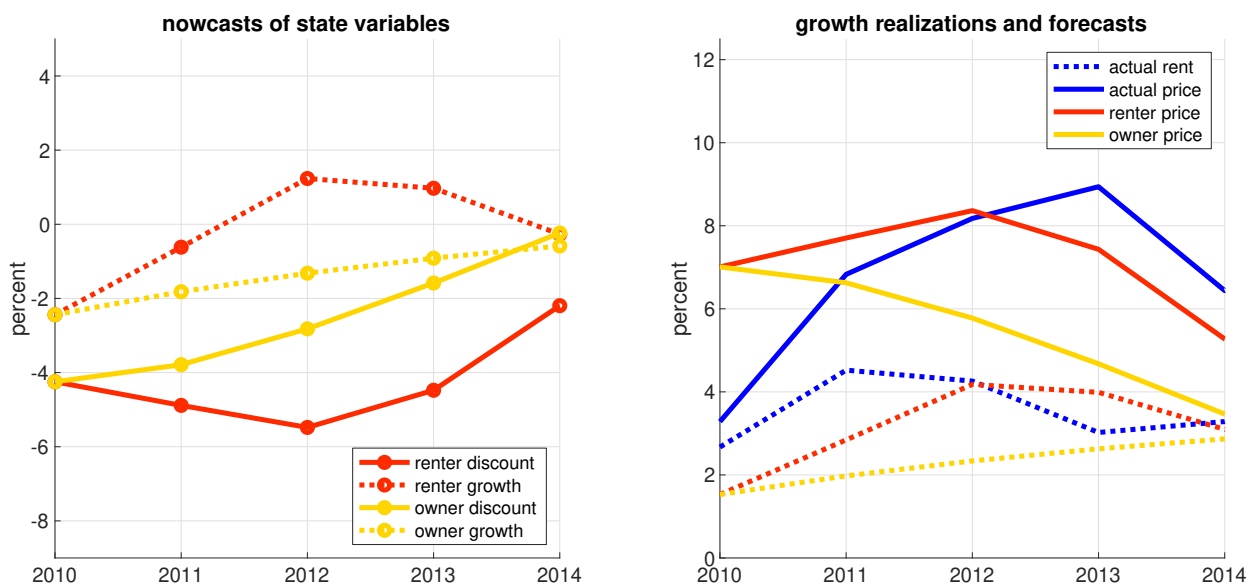
Parameters	Baseline values	Targets	data = model
volatility of discount innovations ε_{t+1}^m	0.121	avg forecast renter	5.27%
initial average rent growth nowcast \bar{g}_0	0.009	avg forecast owner	3.47%
noise volatility for renter	0.217	vol forecast renter	4.83%
noise volatility for owner	0.205	vol forecast owner	3.68%

Consider the calibrated parameter values that speak to the information set. The initial average nowcast of rent growth is low, 0.9%. In order for the average agent to correctly nowcast an initial price-rent ratio that is almost 20% below steady state requires an average nowcast \bar{m}_0 of the discount factor of -4.3% . The coefficient on the discount factor in (5) is 3.27 so the discount factor does most of the work in rationalizing low prices. Renters obtain slightly more noisy signals than owners, although both deal with substantial noise of 21.7% and 20.5%, respectively.

Figure 9 illustrates the belief dynamics that lead the model to account for the data. The color coding is red for renters and yellow for owners. The left panel shows both types' average nowcasts of rent growth (dotted line) and the discount factor (solid line), starting from their common initial condition in 2010. The right panel shows both types' forecasts for rent growth (dotted line) and price growth (solid line) as well as the subsequent actual realizations (in blue), aligned so the realized forecast error is the vertical distance. We thus obtain a quantitative version of the narrative described above. The initial common forecast is based mostly on the common belief that the discount factor coming out of the financial crisis is very low. As the boom starts to develop, opinions diverge: owners perceive a quick rebound of the discount factor, whereas renters take into account rent growth and believe that the rebound of the discount factor is sluggish. Price growth forecasts thus decline, but substantially more so for owners. In 2014, both types make forecasts that are too conservative, with owners about 1.8 percentage points below renters on average, as in the data.

Overall, our calibration strategy is guided by the stated goal of this section: to illustrate that learning about housing cost can lead to large divergence of expectations. The stark difference in information sets in our baseline serves to make our mechanism more powerful. At the same time, assuming initial agreement between average owners and renters also asks it to deliver a lot. An alternative strategy might introduce more or less noise in the rent observations of renters and owners, respectively, and might also allow for some initial disagreement between the average renter and the average owner – after all,

Figure 9: Model Results



it is plausible that the information frictions we emphasize were in place also before the financial crisis. Since we do not have direct evidence on beliefs in this early period, we cannot provide accurate detail here. The quantitative insight we take away is that our mechanism can have first order effects on beliefs.

6 Risk attitude, attention and information structure

In this section we collect evidence on two plausible alternative hypotheses for our main fact that renters are more optimistic than owners. Section 6.1 examines the idea that survey respondents report what they fear, rather than what they expect: renters worry about rent increases whereas owners worry about price declines. Section 6.2 considers attention allocation: renters may be more attuned to an ongoing boom than owners because they are looking to buy. We leverage detailed survey questions on risk attitude as well as household planning, respectively to show that support for these forces is relatively weak. This is what led us to focus our modeling effort on access to information and the unique information structure of housing. As a final piece of supporting evidence, we contrast access to information for housing and equity.

6.1 The role of risk aversion

If renters worry about the size of their rent and therefore about rent and price increases, they might be more willing to forecast higher rent and price growth into the future. Owners, on the other hand, may fear a devaluation of their wealth and hence systematically forecast price growth to be too low. To assess this possibility, we split our sample into two groups according to individual risk aversion. If worry is a driver of responses, we should see more pronounced differences between renters and owners among more risk averse households.

There are two ways to assess risk aversion in the PHF. One is a self-reported general risk aversion measure on an eleven point scale that we already used the cross-sectional regressions in Table 1. A second measure assesses an individual's willingness to take risk when making financial decisions. In particular, households are directly asked about the risk-return tradeoff in an investment context: *If savings or investment decisions are made in your household: Which of the statements on list 5.9 best describes the attitude toward risk? Try to characterize the household as a whole, even if it is not always easy.* Households are asked to select one out of the five statements: "We take significant risks and want to generate high returns.", "We take above-average risks and want to generate above-average returns", "We take average risks and want to generate average returns", "We are not ready to take any financial risks". For both measures, we split our sample at the median and run cross-sectional regressions separately for households with above and below median risk aversion.

Column (1) in Table 5 repeats the cross-sectional regression results from Table 1. The average difference in house price forecasts between renters and owners is 2.1pp. Column (2) reports the results of the same regression for the sample of highly risk averse individuals (above median). We can see that the forecast difference between renters and owners increases somewhat, by about 0.35 percentage points. In the same way, the forecast difference between renters and owners shrinks by about 0.5pp for the below median risk aversion sample in column (3). Hence, we can not fully rule out that fear may play a role in shaping households' house price forecasts. However, even for the sample with low risk aversion, there still is a striking and significant difference of 1.5 percentage points between the house price forecasts of renters and owners. The results are almost identical when using the alternative financial risk-aversion measure. We conclude that fear may have some power in shaping average house price forecasts of renters and owners. However, the majority of the difference between renters and owners reported in Table 1 can not be explained by this hypothesis.

Table 5: Expectations of renters and owners by risk aversion

	<i>General risk aversion</i>			<i>Financial risk aversion</i>	
		<i>Above median</i>	<i>Below median</i>	<i>Above median</i>	<i>Below median</i>
	(1)	(2)	(3)	(4)	(5)
Tenure					
Renter	2.100*** (0.379)	2.469*** (0.535)	1.562*** (0.507)	2.470*** (0.525)	1.519*** (0.487)
Growth Region					
Low	-1.570*** (0.428)	-2.070*** (0.629)	-0.767 (0.523)	-2.165*** (0.567)	-0.409 (0.530)
Medium Low	-1.443*** (0.433)	-1.676** (0.706)	-1.149** (0.471)	-1.575*** (0.584)	-1.101** (0.501)
High	1.006** (0.400)	0.643 (0.575)	1.503*** (0.535)	0.584 (0.538)	1.974*** (0.523)
Controls	Yes	Yes	Yes	Yes	Yes
Number of Cases	3598	1666	1932	2127	1467
R-Square	0.142	0.168	0.190	0.129	0.259

Robust standard errors in parentheses. **: $p < 0.05$, ***: $p < 0.01$


6.2 Forecasts and plans to buy, sell or move

Our theory assumes that renters obtain better information about rents because they observe their own rent as well as that of related properties, and use this information to reason about prices. An interesting alternative hypothesis is that households pay attention to prices in an illiquid market only rarely, namely when they buy or sell. If this were the case, there would be nothing special about tenure per se, but tenure would be an imperfect proxy for incentives to trade.

The PHF includes three questions that speak to this hypothesis. First, it asks renters “Do you intend to buy or build a house or flat for your own accommodation?”. Second, owners are asked about the date they moved into their current residence. Third, in the 2017 wave both renters and owners were asked how long they plan to remain in their current residence. If information improves when households trade, we should expect more accurate forecasts from renters who plan to buy, from owners who have recently bought, as well as from owners or renters who plan to move in the near future.

Figure 10 shows average price forecasts for four types of households: renters with and without plans to buy, as well as owners who moved less than or more than five years ago. The differences among the two renter types are negligible: the high accurate forecast we observe by the average renter is not driven by renters who are planning to buy. For owners, there is a small and borderline significant difference: recent owners make slightly

Figure 10: Price forecasts and incentives



Source: Own calculations based on data from PHF.

higher forecasts. However, the difference between owner types is small relative to the overall difference between owners and renters.

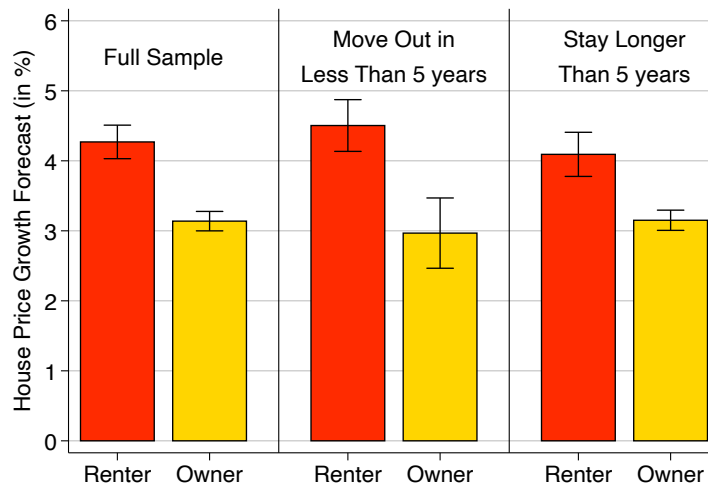
Figure 11 compares average forecasts for owners and renters in the full sample to average forecasts for those households who plan to move within five years as well as those who plan to remain in their current residence longer. Among renters, there is some qualitative support for the idea that agents who plan to shop for a new place soon make higher (and hence more accurate) forecasts. However, for both renters and owners, differences across groups are small and insignificant. Overall, we take away that incentives to trade are not a major factor in driving price forecasts. Our model thus focuses on information advantages that simply reflect tenure, and does not feature endogenous information acquisition.¹²

6.3 Opinions about future prices: real estate vs. equity

The premise of our theory is that the housing market is special among asset markets in that even agents who do not participate – that is, renters – have information about dividends. We now show that this feature indeed differentiates housing from the other major long term assets in modern economies, equity. We make use of the two part structure of the PHF expectations questions: for the question on housing described in Section 3, the

¹²The results of Figure 11 also speak to the hypothesis that households report forecasts of what they fear. If this were the case, we would expect renters who plan to move soon to make higher price growth forecasts, as they fear price or rent increases, whereas owners who plan to move soon should forecast particularly low prices. We do not see relevant differences in the figure.

Figure 11: Forecasts and expected time to move



Source: Own calculations based on data from PHF.


first part gives households the option of responding “don’t know” if they do not want to voice any opinion on the direction of the housing market. An analogous question is available for the stock market.

Housing is a special asset with regard to the share of non-owners who feel confident forming an opinion about price movements. In fact, the vast majority of renters can make a forecast of future house prices. This stands in stark contrast to equity, where we see that a large fraction of non-owners is not capable of making stock price forecasts, which suggests that obtaining signals about house prices is relatively cheap for renters. Figure 12 shows the fraction of opinionated households – who do not answer “don’t know” by participation status for both asset markets. The top panel provides results based on the 2014 survey, while the bottom panel uses the 2017 survey. For equity, “investors” comprise not only households who directly invest in stocks but also those who invest only indirectly via mutual funds or pension funds.

In 2014, the overwhelming majority of the survey population – on average about 90 per cent – is willing to provide a qualitative forecast of future house prices. This number is substantially larger than the 70% who opine on equity. The difference is particularly striking among non-participants: about 85% of renters have developed a view about the housing market, whereas less than 60% of households who do not hold stocks have a view about the stock market. The bottom panel of Figure 12 shows that the same pattern was present in 2017, although the differences between non-owners and owners shrink. Like most countries, Germany has been experiencing an extended period of zero interest rates. Being traditionally mostly active in the risk-free savings market, German households now

have to search for investment alternatives, which raises the incentives for non-owners of real estate or equity to familiarize with those assets.

Figure 12: Opinion formation about prices: real estate vs. equity



Source: Own calculations based on data from PHF.

References

- Adelino, Manuel, Antoinette Schoar, and Felipe Severino**, "Perception of House Price Risk and Homeownership," NBER Working Paper No. 25090, 2018.
- Armona, Luis, Andreas Fuster, and Basit Zafar**, "Home Price Expectations and Behaviour: Evidence from a Randomized Information Experiment," *Review of Economic Studies*, 2019, 86 (4), 1371–1410.
- Bailey, Michael, Ruiqing Cao, Theresa Kuchler, and Johannes Stroebel**, "The Economic Effects of Social Networks: Evidence from the Housing Market," *Journal of Political Economy*, 2018, 126 (6), 2224–2276.
- Campbell, John Y. and Robert J. Shiller**, "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies*, 1988, 1 (3), 195–228.
- Campbell, Sean, Morris Davis, Joshua Gallin, and Robert F. Martin**, "What Moves Housing Markets: A Variance Decomposition of the Rent-Price Ratio," *Journal of Urban Economics*, 2009, 66 (2), 90–102.
- Case, Karl E. and Robert J. Shiller**, "The Efficiency of the Market for Single-Family Homes.," *American Economic Review*, 1989, 79 (1), 125–137.
- and – , "Is there a bubble in the housing market," *Brookings Papers on Economic Activity*, 2003, 2, 299–362.
- , – , and **Anne K. Thomson**, "What Have They Been Thinking? Homebuyer Behavior in Hot and Cold Markets," *Brookings Papers on Economic Activity*, 2012, 2, 265–298.
- Cutler, David M., James M. Poterba, and Lawrence H. Summers**, "Speculative Dynamics," *Review of Economic Studies*, 1991, 58 (3), 529–546.
- D’Acunto, Francesco, Ulrike Malmendier, Juan Ospina, and Michael Weber**, "Exposure to Grocery Prices and Inflation Expectations," *Journal of Political Economy*, 2021, 129 (5), 1615–1639.
- De La O, Ricardo and Sean Myers**, "Subjective Cash Flow and Discount Rate Expectations," *The Journal of Finance*, 2021, 76 (3), 1339–1387.
- Favara, Giovanni and Zheng Song**, "House price dynamics with dispersed information," *Journal of Economic Theory*, 2014, 149, 350–382.
- Gao, Zhenyu, Michael Sockin, and Wei Xiong**, "Learning about the Neighborhood," *Review of Financial Studies*, 2020, 34 (9), 4323–4372.
- Guren, Adam M.**, "House Price Momentum and Strategic Complementarity," *Journal of Political Economy*, 2018, 126 (3), 1172–1218.
- Hartzmark, Samuel M., Sam Hirshman, and Alex Imas**, "Ownership, Learning, and Beliefs," *Quarterly Journal of Economics*, 2021, 136 (3), 1665–1717.

- Head, Allen, Huw Lloyd-Ellis, and Hongfei Sun**, "Search, Liquidity, and the Dynamics of House Prices and Construction," *American Economic Review*, 2014, 104 (4), 1172–1210.
- Kaas, Leo, Georgi Kocharkov, Edgar Preugschat, and Nawid Siassi**, "Low Homeownership in Germany – a Quantitative Exploration," *Journal of the European Economic Association*, 2020, 19 (1), 128–164.
- Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante**, "The Housing Boom and Bust: Model Meets Evidence," *Journal of Political Economy*, 2020, 128 (9), 3285–3345.
- Kuchler, Theresa and Basit Zafar**, "Personal Experiences and Expectations about Aggregate Outcomes," *Journal of Finance*, 2019, 74 (5), 2491–2542.
- Landvoigt, Tim**, "Housing Demand during the Boom: The Role of Expectations and Credit Constraints," *Review of Financial Studies*, 2017, 30 (6), 1865–1902.
- Leombroni, Matteo, Monika Piazzesi, Martin Schneider, and Ciaran Rogers**, "Inflation and the Price of Real Assets," NBER Working Papers No. 26740, 2020.
- Ludwig, Alexander, Jochen Mankart, Jorge Quintana, Nathan Vellekoop, and Mirko Wiederholt**, "House Price Expectations and Housing Choice," Working Paper, University of Frankfurt, 2020.
- Malmendier, Ulrike**, "Experience Effects in Finance: Foundations, Applications, and Future Directions," *Review of Finance*, 2021, 25 (5), 1339–1363.
- **and Stefan Nagel**, "Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?," *Quarterly Journal of Economics*, 2011, 126 (1), 376–416.
- Piazzesi, Monika and Martin Schneider**, "Inflation Illusion, Credit, and Asset Pricing," in John Y. Campbell, ed., *Asset Pricing and Monetary Policy*, Chicago University Press, 2008, pp. 147–181.
- **and –**, "Momentum Traders in the Housing Market: Survey Evidence and a Search Model," *American Economic Review*, 2009, 99 (2), 406–411.
- Vissing-Jorgensen, Annette**, "Perspectives on Behavioral Finance: Does "Irrationality" Disappear with Wealth? Evidence from Expectations and Actions," *NBER Macroeconomics Annual*, 2003.
- Weber, Jan Philip and Gabe Lee**, "On the Measure of Private Rental Market Regulation Index and Its Effects on Housing Rents: Cross-Country Evidence," IREBS Beiträge zur Immobilienwirtschaft, Heft 21, 2018.

A House Price Data

A.1 Constructing house price and rent indices for Germany

The detailed house price and rent dataset from bulwiengesa AG that we use to construct price indices since the year 2005 contains the following price series for houses in all 401 German counties:

- single-family homes
- apartments (newly built, “Erstbezug”)
- apartments (already existing, “Wiederbezug”)
- town houses (not available for all counties)

As only 7.5 percent of households in the PHF live in town houses and the price series is not available for all German counties, we ignore these building types when constructing our house price index.

We merge the two price series on apartments assuming that about 15 percent of apartments are new and 85 percent are already existing. This is the official number used by the Deutsche Bundesbank, and it is based on estimates of the German Federal Statistical Office. This leaves us with two price series: one for single-family homes and the other one for apartments. We combine the two series to one using data from the PHF. For each growth region, we calculate the share of households living in single-family homes out of the share of households that live in either single-family homes or apartments (more than 90 percent of the population). The respective data is shown in Table A.1. Not surprisingly, single-family homes are less common in high growth regions, as the high growth region contains a lot of cities.

Table A.1: Types of housing in different growth regions

	Growth Region			
	Low	Medium Low	Medium High	High
Single-Family Home	49.50	43.21	41.22	16.12
Apartment	50.50	56.79	58.78	83.88

As for rents, there are two series available:

- apartments (newly built, “Erstbezug”)
- apartments (already existing, “Wiederbezug”)

According to the PHF, more than 80 percent of renters are living in apartments. A rental market, especially for single-family homes, hardly exists. As a result, data on rented

houses is scarce. To construct the series for rent prices, we again combine the two price series assuming that about 15 percent of apartments are new.

A.2 Long-run Data on House Price and Rents


The Bundesbank provided us with the following data for the time span between 1976 and 2020 collected by bulwiengesa AG:

- Purchase prices of apartments (newly built, “Erstbezug”)
- Rents of apartments (newly built, “Erstbezug”)
- Rents of apartments (already existing, “Wiederbezug”)

The data is available for a total of 50 (West German) cities. The series for West-Berlin is discontinued in 1990. We select from the remaining 49 cities those, who are in the highest growth region in the period between 2010 and 2020. This leaves us with a total of 28 cities.

Leveling house price data For the long house price series, we only know the purchase price of newly built apartments, which is consistently higher than the average purchase price of a house in the highest growth region. To make the long price series comparable with the 2005-2020 bulwiengesa data, we re-level the data as shown in the left panel of Figure A.1. The dash-dotted line shows the original long price series for newly built

Figure A.1: Leveling of house price data



Source: Own calculations based on data from bulwiengesa AG and PHF.

apartments. The gray series with black dots is the price series determined for exactly the same 28 cities/counties from the detailed bulwiengesa data we have available since 2005. For the time frame in which both series are available, they move in parallel, but newly

built apartments are just more expensive. To compensate for this effect, we calculate the mean ratio between the two price series at any date in between 2005 and 2020. We then normalize the original series by this ratio. The solid line in the left panel of Figure A.1 shows the resulting adjusted time series which is hardly distinguishable from the detailed data series between 2005 and 2020. This makes us confident that we can use the newly built apartment price series to extrapolate the evolution of prices backward in time until 1976.

In the right panel of Figure A.1, we compare the long-run price series with publicly available time series data for house prices in Germany from the Deutsche Bundesbank, the OECD and the BIS. In the 1970s to early 1990s, all of these data are based on the same common data source that also underlies our price series. The Deutsche Bundesbank uses data on the purchase prices of newly built apartments for all 50 cities, not the 28 we selected, but the price dynamics seem to be fairly similar. With German reunification, the official sample of the Bundesbank was extended to a total of 100 cities and in 1995 to 125 cities, also covering East Germany. This is the first time when our series departs from official sources as we want to provide data for one set of counties over time, while the Bundesbank seeks to approximate house price dynamics all over Germany. The second structural break in the official data series is in 2005, when the detailed *bulwiengesa* data starts. From this point onward, the Bundesbank uses data covering all German counties from different data sources. As we have seen before, the house price boom is much stronger in the highest growth region than in Germany as a whole. Not surprisingly, our price series then departs from official ones at the onset of the German house price boom.

Rents Turning to rent prices, the long-run data contains exactly the same information as the detailed *bulwiengesa* data. Hence, when applying the same weighting as described in the previous section, we obtain the same price series from 2005 onwards, see the right panel of Figure 3. When combining the rent and the price series, we obtain the time series for the price-rent ratio shown in Figure 4.

B Documentation of Variables and Sample Selection

While we already provide detailed information on both the house price and rent data as well as on the assessment of households' price expectations in the main text of the paper, this appendix documents all other survey variables we used in the regressions and figures reported in this paper as well as our sample selection procedure.

B.1 Explanatory Variables Used in Regressions

The following variables are used as explanatory variables in regressions. We use data from multiple waves of the Panel on Household Finances (PHF) as well as the Bundesbank Online Panel – Households (BOP-HH). As variable names might differ across those surveys, but the survey methodology is the same, we only provide the variable name for our baseline survey, the second wave (2014) of the Panel on Household Finances. A full set of variable names and Stata codes are available upon request.

B.1.1 Demographics, Income, Wealth

Age (ra0300): Age refers to the age of the household head. The household head is the “Financially knowledgeable person” of the interview, i.e. the person that knows best about the households’ finances and answers the household questionnaire including all details on wealth holdings for the entire household.

Number of Household Members (anzhmm): Number of all household members currently living in the household, including the household head.

Net Wealth (constructed from several variables): We sum up all safe assets (bonds, pension accounts, life insurances), equity (stocks, mutual funds, businesses), real estate investments and non-interest bearing assets (cash, vehicles, private assets) at the household level and subtract any outstanding debts for the entire household.

Net Household Income (dhi0600): Household heads are asked for the total net disposable household income, consisting of wages, salary, income from self-employment, retirement benefits or pensions, income from public aid, income from renting, income from leasing, housing allowance, child benefits, and other income.

Education (pa0200): This variable describes the highest education level the household head has achieved. Low education refers to having no school degree or at most a lower level secondary school degree (corresponding to the German “Realschule”), Upper Secondary education comprises all upper secondary level degrees (“Abitur”, “Fachhochschulreife”) and tertiary education refers to university or equivalent degrees.

B.1.2 Behavioral Traits

Financial Literacy (dhnm0100, dhnm0200, dhnm0300): Household heads are asked three questions in the areas of interest rate compounding, inflation and diversification, which we use to measure their financial literacy. In particular, those questions are:

Question 1: "Let us assume you have a balance of 100 euros in your savings account. This balance bears interest at an annual rate of 2%, and you leave it there for 5 years. What do you think: How high is your balance after 5 years?"

Possible answers: 1 – Higher than 102 euros / 2 – Exactly 102 euros / 3 – Lower than 102 euros / –1 – Don't know / –2 – No answer

Question 2: "Let us assume that the interest paid on your savings account is 1% per year and the inflation rate is 2% per year. What do you think: After a year, will you be able to buy just as much, more or less than today with the balance in your savings account?"

Possible answers: 1 – More / 2 – Just as much / 3 – Less than today / –1 – Don't know / –2 – No answer

Question 3: "Do you agree with the following statement: "The investment in the stock of a single company is less risky than investing in a fund with stock in similar companies?"

Possible answers: 1 – I agree / 2 – I do not agree / –1 – Don't know / –2 – No answer

We form a literacy index out of the answers to these three questions. Literacy is "Very Low" in case the household head gives a wrong or no answer to all three questions; it is "Low" if the respondent was able to answer exactly one question correctly; it is "Medium" if the respondent was able to answer exactly two questions correctly; it is "High" if the respondent gave correct answers to all three questions.

Patience (zi105): The patience measure comes from a Likert scale where respondents are asked to assess their level of patience. The exact wording of the question is: "How do you view yourself personally: Are you in general a person who is patient or do you tend to be impatient?" Answers can be given on a scale from 0 to 10 where zero refers to being "Very patient" and 10 to "Very impatient". We form three categories for our variable of patience with "High patience" including all respondents who answered at most three on the scale, "Average patience" between 4 and 6 of the scale and "No patience" if respondents rated their patience to be between 7 and 10 on the scale. We chose the division into three groups according to the frequency of answers. Consequently, each of the three groups is populated approximately by an equal share of households.

Risk aversion (zi103): The survey assesses risk aversion in the same way as patience. Households are asked about how risk averse they view themselves. They can again provide a score on a scale from 0 to 10. We split the sample in two groups of about equal size, those "Below Median" and those "Above Median" risk aversion.

Financial risk aversion (hd1800, dhd2800): An alternative measure of risk aversion – financial risk aversion – asks about the household’s investment behavior. The exact wording of the question we use is: “If savings or investment decisions are made in your household: Which of the statements best describes the attitude toward risk?”

Possible answers:

- 1 – We take significant risks and want to generate high returns.
- 2 – We take above-average risks and want to generate above average returns.
- 3 – We take average risks and want to generate average returns.
- 4 – We are not ready to take any financial risks.
- 5 – No uniform classification is possible for the household as a whole.

If the household head chooses answer 5, he or she is asked what investment behaviour he/she has for him-/herself. We again split the sample in two groups of about equal size, those “Below Median” and those “Above Median” risk aversion.

B.1.3 Tenure and Growth Region

Tenure (dhb0200a, dhb0200b, hb0500): We say a household is an “Owner” if the household owns at least 50 percent of the property the household members are living in, the so-called primary residence. All other households are classified as “Renters”.

Growth Region: Our house price data set is based on German counties (Kreise and kreisfreie Staedte). There are 401 such counties in Germany, out of which approximately one half is covered by the Panel on Household Finances. To define growth regions, we take the 401 counties and calculate their “trend growth” over the course of the past 10 years, i.e. starting from the beginning of the house price boom in 2011 to the last year available in the data (2020). To this end, we regress the log of the house price in each county on a year variable. The regression coefficient is our trend growth variable. We divide the 401 counties into four quartiles with respect to trend growth. This leaves us with four growth regions, i.e. those with Low, Medium Low, Medium High and High trend growth in between 2011 and 2020.

B.1.4 Housing/Regional Characteristics

Time Household has been Living in Current Residence (intjahr, dhb0120): Number of years since household has moved into their current primary residence (measured at the time of the interview).

Size of the Local Community (bikgk10): While the county (“Kreis”) a household is living in is the unit on which we measure house prices, this is an even finer descrip-

tion of the size of the local municipality. The official community size measure in Germany is the so-called “BIK Scale” that classifies community sizes along ten categories:

- 1 – < 2,000 inhabitants
- 2 – 2,000 to 4,999 inhabitants
- 3 – 5,000 to 19,999 inhabitants
- 4 – 20,000 to 49,999 inhabitants
- 5 – 50,000 to 99,999 inhabitants and BIK structure type 2 (city region) / 3 (second tier towns) / 4 (third tier towns)
- 6 – 50,000 to 99,999 inhabitants and BIK structure type 1 (metropolitan area)
- 7 – 100,000 to 499,999 inhabitants and BIK structure type 2/3/4
- 8 – 100,000 to 499,999 inhabitants and BIK structure type 1
- 9 – \geq 500,000 inhabitants and BIK structure type 2/3/4
- 10 – \geq 500,000 inhabitants and BIK structure type 1

Building in Which Household Lives Needs Renovation (sc0400): After each interview, the interviewer is asked to fill out a questionnaire on the interview with the household, including questions about the atmosphere, the willingness of the respondent to answer, the living conditions and the state of the house. From these para-data we take the following question to assess the quality of the building that households live in: “Please describe the condition of the building.”

Possible answers:

- 1 – Clean and well maintained
- 2 – A few small cracks in the facade, and isolated cases of peeling paint
- 3 – Badly in need of renovation
- 4 – Dilapidated

Our indicator of renovation needs refers to all answer categories greater than 1.

Quality of the Household's Residence (sc0200): The interviewer is also asked to give a rating of the house. The exact wording of the question is: "Please rate the building."

Possible answers:

- 1 – Exclusive
- 2 – Very good
- 3 – Satisfactory
- 4 – Simple
- 5 – Very simple

Size of Households Residence (hb0100): This is a simple measure of the square meter size of the household's residence. The question refers to the size in sqm of the living space of the house or the apartment, not the size of the full property.

B.2 Variables Used in Figures

The following variables are used as explanatory variables in figures.

Equity Investors (constructed from several variables): We say a household is an equity investor, if the household's portfolio contains any of the following assets: directly held shares, certificates, other securities, mutual funds and private pensions primarily invested in stocks, hedge funds, other risky assets, silent partnerships as well as managed accounts. In our definition, equity also contains the value of self-employment business.

Perceived House Price Growth (hb0900): We can match a subsample of households (the panel households) across the 2011 and 2014 waves as well as the 2014 and 2017 waves of the PHF, respectively. To construct a household's perceived house price growth (see Figure 8), we look at those households who were owners of the same residence in two consecutive waves. For these households, the PHF asks household heads for a price quote of their current primary residence in both waves. Using these price quotes, we calculate each household's perceived annualized house price growth rate and contrast it with the average house price growth rate in the county the household is living in.

Renters with Plans to Buy a House (dnh3000): Renters in the PHF are asked about their intention to buy a property. The exact working of the question is: "Does your household intend to buy or build a house or flat for your own accommodation?". The possible answers are "Yes", "No" and "Don't Know".

Owner’s Tenure Time (hb0700): Owners are asked about the year in which they became owner of the property they currently live in. Note that this date does not have to coincide with the date the household moved into the property. We subtract the date the household became owner of the property from the date of the interview to get to the household’s tenure time.

Intentions to Move (dhh0125): In wave 3 of the PHF, all respondent households are asked about their intentions to stay in their current primary residence. The exact working of the question is: “How long do you expect you/at least one member of the household will continue to live in your primary residence?” The answer is given on a numerical scale indicating “At least another . . . years”. Alternatively the household can choose “Forever” or “Don’t Know”.

B.3 Sample Preparation and Sample Selection in the PHF

In this section, we describe our sample preparation and selection procedure for the PHF, our primary dataset. We proceed in a similar way with the Online Survey – Households. Yet, since the BOP-HH is less rich, the sample preparation and selection is much simpler. A description is available upon request.

In the PHF, we use only the first implicate of the full dataset for simplicity. Since hardly any of the variables we use in this paper are imputed variables, this does not play very much of a role. We then proceed as follows to prepare the sample:

1. **Matching between PHF and house price data:** We first match households to their specific county. This is done via the “Kreiskennziffer”, a unique identifier for each of the 401 German counties. Note that the scientific use files of the PHF do not contain this information, neither do they provide any house price or rent price data on the county level. Users who wish to get access to this data should get in contact with the Research Data and Service Center of the Deutsche Bundesbank.¹³
2. **Classification of Household Portfolios:** We classify household portfolios according to the following schedule:
 - *Safe Assets:* This category contains deposits, directly held bonds as well as mutual funds and private pensions consisting predominantly of money market investments and/or bonds. We subtract from this mortgages as well as any unbacked bank credit.
 - *Real Estate:* This category contains mutual funds and private pensions consisting predominantly of real estate investment, the value of the household’s main

¹³See <https://www.bundesbank.de/en/bundesbank/research/rdsc> for further details.

residence as well as of all secondary residences, all direct real estate investments made for the purpose of renting them out to others or keeping them empty to enjoy capital gains, as well as all business related real estate.

- *Equity*: This category contains mutual funds and private pensions consisting predominantly of stocks, hedge funds or any other risky investments, silent partnerships, managed accounts, directly held shares, certificates and other securities as well as other financial assets. We also include the value of self-employment business in this category.
- *Non Interest-Bearing*: This category contains money other households owe to the observed household, cash, vehicles and valuables. We subtract from this all money the observed household owes to any other household.


The sum of safe assets, real estate, equity and non interest-bearing assets yields total household net wealth. We derive portfolio shares for the four asset categories by dividing their sum on the household balance sheet by net wealth.

3. **Household Expectations**: The PHF asks households about their expectations towards (i) the deposit rate one year ahead of the time of the interview, (ii) house price growth over the next 12 months, (iii) the growth in equity price as measured by the German Stock Market Index (DAX) as well as (iv) growth in the general price level. Except for the question about the expected deposit rate, all expectation questions are asked in the two-step format discussed in section 3. After having coded all households' numerical expectations, we cut off the top and the bottom 1 percent of reported expectations by "missing" them to missing. This is done in order to deal with outliers.

With the full sample at hand, we clean the sample and identify outliers as follows:

- We drop all households who report a non-positive net wealth balance in their portfolio. The majority of those households holds no or only a tiny amount of real estate and equity, but typically has a lot of unbacked consumer credit and credit card debt.
- We then look at portfolio shares. The left panel of Figure B.2 shows a scatter plot of portfolio shares of households. We drop all observations represented by a red dot in the figure, which are rather extreme. This means that we drop households with a more than 300 percent equity share or a 3000 percent real estate share in their portfolio with the understanding that this most likely represents coding or data collection errors. The right panel of Figure B.2 shows the resulting distribution of portfolio shares across households.

Figure B.2: Portfolio Shares and Outlier Identification



C Additional Datawork

In this appendix, we first provide the full regression tables for the regression presented in short in Table 1. Second, we present additional robustness checks using alternative specifications, alternative explanatory variables as well as data from the 2017 version of the PHF. Third, we discuss our reweighting procedure that controls for the fact that renters and owners are unequally distributed across different growth regions.

C.1 The Cross Section of House Price Forecasts in 2014

Table 1 conveyed as simple but important message for our paper: In a linear regression setting, only two variables are relevant for explaining households' forecasts: tenure and location. Table C.2 shows the full regression results for this table.

Table C.2: The Cross Section of House Price Forecasts in 2014 (see Table 1)

	(1)	(2)	(3)	(4)	(5)
DEMOGRAPHICS, INCOME, WEALTH					
Age					
< 30	1.159 (0.707)	1.125 (0.689)	0.696 (0.691)	0.531 (0.660)	0.099 (0.675)
30 - 39	1.365** (0.591)	1.306** (0.585)	0.945 (0.576)	0.703 (0.564)	0.367 (0.581)
40 - 49	-0.227 (0.451)	-0.254 (0.446)	-0.240 (0.447)	-0.229 (0.429)	-0.295 (0.441)
60 - 69	0.337 (0.467)	0.308 (0.467)	0.355 (0.462)	0.402 (0.453)	0.334 (0.474)
≥ 70	-0.171 (0.418)	-0.274 (0.437)	-0.206 (0.435)	-0.251 (0.427)	-0.263 (0.431)
Number of Household Members					
1	0.576 (0.434)	0.555 (0.434)	0.268 (0.436)	-0.054 (0.425)	-0.461 (0.435)
3	0.314 (0.479)	0.298 (0.482)	0.463 (0.472)	0.340 (0.456)	0.320 (0.457)
4	0.371 (0.665)	0.276 (0.654)	0.484 (0.654)	0.398 (0.646)	0.304 (0.611)
≥ 5	0.186 (0.551)	0.218 (0.548)	0.617 (0.547)	0.643 (0.558)	0.744 (0.583)
Net Wealth Quartiles					
1st Quartile	1.903*** (0.494)	1.920*** (0.501)	0.139 (0.566)	0.032 (0.551)	-0.269 (0.553)
2nd Quartile	0.939** (0.448)	0.996** (0.441)	-0.186 (0.466)	-0.136 (0.454)	-0.150 (0.459)
4th Quartile	0.279 (0.292)	0.311 (0.296)	0.562 (0.301)	0.118 (0.296)	0.081 (0.311)

Net Household Income Quartiles					
1st Quartile	0.107 (0.560)	0.108 (0.559)	0.286 (0.558)	0.388 (0.546)	0.607 (0.554)
2nd Quartile	0.030 (0.438)	0.051 (0.436)	0.073 (0.438)	0.169 (0.429)	0.146 (0.422)
4th Quartile	-0.208 (0.320)	-0.216 (0.323)	-0.323 (0.320)	-0.368 (0.310)	-0.476 (0.302)
Education					
Low	-0.249 (0.566)	-0.269 (0.568)	-0.414 (0.563)	-0.397 (0.561)	-0.438 (0.566)
Tertiary	0.394 (0.305)	0.421 (0.317)	0.289 (0.312)	0.154 (0.302)	0.082 (0.304)
BEHAVIORAL TRAITS					
Financial Literacy					
Very Low		1.576 (1.345)	1.616 (1.392)	1.645 (1.330)	2.029 (1.428)
Low		-0.703 (0.737)	-0.777 (0.728)	-0.908 (0.735)	-0.772 (0.771)
Medium		-0.167 (0.346)	-0.133 (0.344)	-0.296 (0.340)	-0.208 (0.352)
Patience					
Not Patient		0.357 (0.386)	0.288 (0.385)	0.293 (0.371)	0.418 (0.367)
Very Patient		0.104 (0.338)	0.041 (0.335)	0.010 (0.329)	-0.006 (0.324)
Risk Aversion					
Below Median		-0.308 (0.306)	-0.352 (0.304)	-0.274 (0.299)	-0.189 (0.302)
TENURE AND LOCATION					
Tenure					
Renter			2.445*** (0.382)	2.353*** (0.371)	2.100*** (0.379)
Growth Region					
Low				-2.073*** (0.439)	-1.570*** (0.428)
Medium Low				-1.572*** (0.431)	-1.443*** (0.433)
High				1.355*** (0.416)	1.006** (0.400)
HOUSING/REGIONAL CHARACTERISTICS					
Time Household has been Living in Current Residence (Years)					
t < 5					0.304 (0.455)
5 ≤ t < 10					0.587 (0.445)
10 ≤ t < 15					0.188 (0.523)

Size of the Local Community (Inhabitants)					
2k < 5k					-1.500 (1.304)
5k < 20k					-0.071 (0.798)
20k < 50k					-0.074 (0.766)
50k < 100k, periphery					0.764 (0.982)
50k < 100k, center					0.915 (0.983)
100k < 500k, periphery					0.756 (0.748)
100k < 500k, center					0.668 (0.769)
≥ 500k, periphery					1.144 (0.831)
≥ 500k, center					1.808** (0.719)
Building in Which Household Lives Needs Renovation					
Yes					-0.413 (0.390)
Quality of the Household's Residence (Interviewer Rating)					
Exclusive					-0.023 (0.628)
Satisfactory					-0.337 (0.386)
Simple					-0.442 (0.565)
Very simple					-0.088 (1.018)
Size of Households Residence					
Sqm size/100					-1.626*** (0.620)
Sqm size/100 squared					0.429*** (0.128)
Constant	1.749*** (0.473)	1.843*** (0.557)	1.502*** (0.556)	2.245*** (0.682)	2.985** (1.165)
Number of Cases	3647	3646	3646	3646	3598
R-Square	0.038	0.042	0.064	0.119	0.142

Robust standard errors in parentheses. **: $p < 0.05$, ***: $p < 0.01$

C.2 Clustered standard errors

The standard errors in Table 1 are not clustered. To check whether clustering affects our results, we divided households into 30 clusters according to local house price growth over the last decade, with a reasonable number of households in each cluster (> 100). Our results do not change much; tenure and region dummies are still highly significant. We also clustered by Bundesland but do not report results here; they are also consistent with those in Table C.3.

Table C.3: Regressions from Table 1 with 30 clusters by house price growth region

	(1)	(2)	(3)	(4)	(5)
Demographics, Income, Wealth					
Age Group 30–39	1.365*** (0.458)	1.306*** (0.459)	0.945** (0.454)	0.703 (0.446)	0.367 (0.449)
1st Net Wealth Quartile	1.903*** (0.531)	1.920*** (0.545)	0.139 (0.587)	0.032 (0.564)	–0.269 (0.559)
2nd Net Wealth Quartile	0.939 (0.474)	0.996** (0.469)	–0.186 (0.476)	–0.136 (0.447)	–0.150 (0.416)
Behavioral Traits		yes	yes	yes	yes
Tenure					
Renter			2.445*** (0.453)	2.353*** (0.413)	2.100*** (0.359)
Growth Region					
Low				–2.073*** (0.400)	–1.570*** (0.250)
Medium Low				–1.572*** (0.439)	–1.443*** (0.396)
High				1.355*** (0.467)	1.006** (0.404)
Housing/Regional Characteristics					
City Center $\geq 500k$ Inh.					1.808 (0.945)
Sqm size/100					–1.626*** (0.511)
(Sqm size/100) ²					0.429*** (0.120)
Number of Cases	3647	3646	3646	3646	3598
R ²	0.038	0.042	0.064	0.119	0.142

Robust standard errors in parentheses. **: $p < 0.05$, ***: $p < 0.01$

C.3 Growth Region Controls vs. Past Local Price Growth Experience

In Table 1 we only control for four different growth regions. In the following table, we show that controlling for past house price growth experience on a finer scale does not change our results. In particular, the difference in house price growth forecasts between renters and owners remains large and significant.

Table C.4: Growth Region Controls vs. Past Local Price Growth Experience

	(1)	(2)	(3)	(4)
Tenure				
Renter	2.100*** (0.379)	1.941*** (0.379)	1.927*** (0.375)	1.997*** (0.376)
Growth Region				
Low	-1.570*** (0.428)			
Medium Low	-1.443*** (0.433)			
High	1.006** (0.400)			
Past local house price growth				
Last year		0.215*** (0.038)		
Last three year average			0.396*** (0.047)	
Last five year average				0.466*** (0.064)
Controls				
Demographics, Income, Wealth	yes	yes	yes	yes
Behavioral Traits	yes	yes	yes	yes
Housing/Regional Characteristics	yes	yes	yes	yes
Number of Cases	3598	3598	3598	3598
R ²	0.142	0.131	0.145	0.140

Robust standard errors in parentheses. **: $p < 0.05$, ***: $p < 0.01$

C.4 The cross section of house price forecasts in 2017

In Table C.5 we test whether we can also replicate the results from Table 1 in the 2017 wave of the PHF. Using the 2017 data, we obtain the same picture as in our baseline regression: tenure and location are the single most important determinants of households' house price forecasts. Yet, there are also some differences. As we already showed in Figure 7, the difference in forecasts of renters and owners shrinks over time. Consequently, the coefficient on the "Renter" dummy variable is smaller in 2017 than in our baseline regression. In addition, it seems that controlling for community size becomes more important, which might reflect the divergence of local housing market performance over the course of the German house price boom. Last but not least, the R-squared value is smaller than in our baseline regression.

Table C.5: The cross section of house price forecasts in 2017

	(1)	(2)	(3)	(4)	(5)
DEMOGRAPHICS, INCOME, WEALTH					
Age					
< 30	-0.737 (0.478)	-0.819 (0.518)	-1.013 (0.508)	-0.970 (0.510)	-1.233 (0.529)
30 - 39	0.174 (0.379)	0.139 (0.406)	-0.043 (0.448)	-0.115 (0.470)	-0.396 (0.474)
40 - 49	-0.258 (0.099)	-0.307 (0.128)	-0.327 (0.128)	-0.324 (0.131)	-0.491*** (0.057)
60 - 69	0.388*** (0.042)	0.376*** (0.049)	0.370*** (0.041)	0.424*** (0.053)	0.475*** (0.112)
≥ 70	0.177 (0.208)	0.218 (0.214)	0.160 (0.187)	0.161 (0.188)	0.363 (0.173)
Number of Household Members					
1	0.785*** (0.108)	0.818*** (0.113)	0.714** (0.149)	0.516** (0.153)	0.273 (0.144)
3	0.418 (0.413)	0.412 (0.365)	0.493 (0.369)	0.505 (0.352)	0.578 (0.413)
4	0.156 (0.476)	0.181 (0.538)	0.320 (0.539)	0.402 (0.555)	0.313 (0.529)
≥ 5	0.587 (0.497)	0.667 (0.492)	0.840 (0.506)	0.901 (0.543)	1.150 (0.491)
Net Wealth Quartiles					
1st Quartile	0.827 (0.643)	0.743 (0.678)	-0.316 (0.516)	-0.397 (0.525)	-0.681 (0.607)
2nd Quartile	0.571 (0.411)	0.546 (0.388)	-0.264 (0.377)	-0.198 (0.400)	-0.249 (0.431)
4th Quartile	-0.085 (0.185)	-0.145 (0.207)	0.105 (0.233)	-0.137 (0.276)	0.085 (0.414)

Net Household Income Quartiles					
1st Quartile	-0.220 (0.204)	-0.213 (0.208)	-0.182 (0.231)	-0.108 (0.244)	-0.079 (0.279)
2nd Quartile	-0.213 (0.470)	-0.220 (0.474)	-0.194 (0.454)	-0.155 (0.455)	-0.116 (0.454)
4th Quartile	-0.297 (0.193)	-0.290 (0.191)	-0.391 (0.164)	-0.424 (0.186)	-0.482** (0.129)
Education					
Low	-0.095 (0.800)	-0.143 (0.810)	-0.262 (0.848)	-0.308 (0.835)	-0.532 (0.617)
Tertiary	0.523 (0.430)	0.544 (0.423)	0.500 (0.424)	0.308 (0.362)	0.283 (0.304)
BEHAVIORAL TRAITS					
Financial Literacy					
Very Low		-0.181 (0.682)	-0.033 (0.681)	0.074 (0.779)	-0.204 (1.113)
Low		0.709 (0.248)	0.710 (0.262)	0.714 (0.240)	0.639 (0.378)
Medium		-0.074 (0.434)	-0.005 (0.432)	0.030 (0.455)	0.039 (0.485)
Patience					
Not Patient		0.604 (0.392)	0.543 (0.414)	0.524 (0.407)	0.529 (0.303)
Very Patient		-0.064 (0.320)	-0.105 (0.317)	-0.120 (0.316)	-0.173 (0.248)
Risk Aversion					
Below Median		0.191 (0.297)	0.134 (0.310)	0.109 (0.325)	0.070 (0.202)
TENURE AND LOCATION					
Tenure					
Renter			1.492*** (0.229)	1.388** (0.241)	1.028** (0.180)
Growth Region					
Low				-0.679*** (0.045)	-0.388 (0.169)
Medium Low				-0.185** (0.051)	-0.214 (0.165)
High				1.103*** (0.061)	0.906*** (0.036)
HOUSING/REGIONAL CHARACTERISTICS					
Time Household has been Living in Current Residence (Years)					
t < 5					0.877 (0.300)
5 ≤ t < 10					0.668 (0.634)
10 ≤ t < 15					0.256 (0.284)

Size of the Local Community (Inhabitants)					
2k < 5k					0.492 (1.021)
5k < 20k					0.941 (0.801)
20k < 50k					1.137** (0.299)
50k < 100k, periphery					1.927 (0.660)
50k < 100k, center					1.045 (0.885)
100k < 500k, periphery					2.410** (0.500)
100k < 500k, center					1.291 (0.603)
≥ 500k, periphery					1.759*** (0.267)
≥ 500k, center					1.934** (0.458)
Building in Which Household Lives Needs Renovation?					
Yes					-0.178 (0.213)
Quality of the Household's Residence (Interviewer Rating)					
Exclusive					-0.383 (0.487)
Satisfactory					-0.046 (0.267)
Simple					0.578 (0.308)
Very simple					0.256 (0.587)
Size of Households Residence)					
Sqm size/100					-1.044 (0.544)
Sqm size/100 squared					0.144 (0.076)
Constant	3.014*** (0.404)	2.793** (0.586)	2.561** (0.523)	2.629*** (0.386)	1.903** (0.502)
Number of Cases	4249	4247	4247	4247	4124
R-Square	0.020	0.027	0.039	0.061	0.087

Robust standard errors in parentheses. **: $p < 0.05$, ***: $p < 0.01$

C.5 Reweighting for aggregate statistics

In this appendix, we describe how we aggregate findings by region into single numbers for Germany as whole. The first two columns of Table C.6 show the distribution of households across growth regions by tenure type using our sample weights: renters are relatively more likely to live in cities where house price growth has been high. Therefore, renters have a higher likelihood of living in a higher growth region. Even if owners and renters in each region made the exact same forecasts, but differed by region, a simple average would therefore show relatively higher forecasts from renters. We do not want this composition effect to inflate national level forecast differences.


Table C.6: Distribution across growth regions by tenure type

Growth Region	<i>PHF Sample Weights</i>		Total	<i>Reweighted for Composition Effect</i>	
	Renter	Owner		Renter	Owner
Low Growth	16.52	21.34	18.84	18.84	18.84
Medium Low Growth	20.42	27.62	23.88	23.88	23.88
Medium High Growth	25.17	29.07	27.04	27.04	27.04
High Growth	37.89	21.97	30.24	30.24	30.24
Sample Share	51.95	48.05	100.00	51.95	48.05

We therefore aggregate forecasts across growth regions by reweighting: we scale the sample weights for households in a tenure cell and growth region so that the distribution of households across growth regions becomes the same for renters and owners. In particular, both distributions become equal to the distribution of all households across regions shown in the third column; since the construction of regions did not weigh counties by populations, the high growth region that contains larger cities is more populated. To get from the original to the reweighted distribution, involves, for example, reducing (increasing) the sample weight of renters (owners) in the top region.

The reweighted distribution does not generate the composition effect discussed above: if all renters and owners in each region made the same forecast but differed by region, then the average forecast for Germany would also be equal. Figure C.3 reports the full sample averages of house price growth forecasts by tenure status using both the original sample weights as well as the weights that control for household composition. We find a relatively small difference between the two weighting schemes; nevertheless we employ our reweighting scheme in what follows to guard against composition effects.

Figure C.3: House Price Growth Forecasts by Tenure



Source: Own calculations based on data from bulwiengesa AG and PHF.

D Derivations of model equations

This appendix provides derivations for all the formulas presented in the paper.

D.1 The log-linearized first order condition

We first log-linearize the developer's first order condition to derive the linear stochastic difference equation (3) for the log price-rent ratio. The first order condition (2) is

$$V_t = E_t [\tilde{M}_{t+1}(V_{t+1} + 1)G_{t+1}].$$

This equation has a positive and finite steady state at

$$V = MG(V + 1) \Leftrightarrow V = \frac{MG}{1 - MG} \quad (\text{D.1})$$

in case $0 < MG < 1$.

An alternative way to write the first order condition is

$$1 = E_t \left[\tilde{M}_{t+1} \frac{(V_{t+1} + 1)G_{t+1}}{V_t} \right] = E_t [\tilde{M}_{t+1} S_{t+1}], \quad (\text{D.2})$$

where S_{t+1} is the return on housing investment for the developer:

$$S_{t+1} = \frac{P_{t+1} + R_{t+1}}{P_t} = \frac{(V_{t+1} + 1)G_{t+1}}{V_t}. \quad (\text{D.3})$$

For any variable X_t , small letters denote logs and hats the deviation from steady state X

$$x_t = \log(X_t) \quad \text{and} \quad \hat{x}_t = \log(X_t) - \log(X) = x_t - x.$$

Taking logs of the return identity (D.3), we obtain

$$s_{t+1} = \log(1 + \exp(v_{t+1})) + g_{t+1} - v_t.$$

In the spirit of a Campbell-Shiller decomposition, we approximate the logarithmic term linearly in v_{t+1} around its steady state value $v = \log(V)$. This yields

$$\log(1 + \exp(v_{t+1})) \approx \log(1 + \exp(v)) + \frac{\exp(v)}{1 + \exp(v)}(v_{t+1} - v) = k + \frac{V}{1 + V}v_{t+1}$$

with $k = \log(1 + V) - \frac{V}{1+V} \log(V)$. Using the identity $MG = \frac{V}{1+V}$ from equation (D.1),

we then obtain

$$s_{t+1} \approx k + MGv_{t+1} + g_{t+1} - v_t.$$

Acknowledging that this equation also needs to hold in the steady state, we can write the return identity in log-deviations from the steady state as

$$\hat{s}_{t+1} \approx MG\hat{v}_{t+1} + \hat{g}_{t+1} - \hat{v}_t. \quad (\text{D.4})$$

Next, we turn to the first order condition (D.2)

$$1 = E_t [\tilde{M}_{t+1} S_{t+1}] = E_t [\exp(\tilde{m}_{t+1} + s_{t+1})].$$

Assuming that \tilde{m}_{t+1} and s_{t+1} are jointly conditionally normally distributed, we obtain

$$\begin{aligned} 1 &= E_t [\exp(\tilde{m}_{t+1} + s_{t+1})] \\ &= \exp \left(E_t[\tilde{m}_{t+1}] + E_t[s_{t+1}] + \frac{1}{2} [\text{Var}_t(\tilde{m}_{t+1}) + \text{Var}_t(s_{t+1}) + 2\text{Cov}_t(\tilde{m}_{t+1}, s_{t+1})] \right). \end{aligned}$$

If \tilde{m}_{t+1} and s_{t+1} are homoskedastic, we can subsume the conditional variances and covariances in a constant term

$$D = \frac{1}{2} \text{Var}_t(\tilde{m}_{t+1}) + \frac{1}{2} \text{Var}_t(s_{t+1}) + \text{Cov}_t(\tilde{m}_{t+1}, s_{t+1}).$$

Taking logs and using our notation $m_t = E_t[\tilde{m}_{t+1}]$ for the predictable component of the log stochastic discount factor, we obtain

$$m_t + E_t[s_{t+1}] + D = 0,$$

or in terms of log-deviations from the steady state

$$\hat{m}_t + E_t[\hat{s}_{t+1}] = 0.$$

Substituting the approximate return identity (D.4), we can write

$$\begin{aligned} \hat{m}_t + E_t [MG\hat{v}_{t+1} + \hat{g}_{t+1} - \hat{v}_t] &\approx 0 \\ \Leftrightarrow \hat{v}_t &\approx \hat{m}_t + E_t [MG\hat{v}_{t+1} + \hat{g}_{t+1}], \end{aligned}$$

which is the linear stochastic difference equation (3).

D.2 The solution for the log price-rent ratio

Suppose that \hat{m}_t and \hat{g}_t follow the Gaussian AR(1) dynamics (4). Our conjecture is that the solution to the linear stochastic difference equation (3) is

$$\hat{v}_t = \beta_m \hat{m}_t + \beta_g \hat{g}_t. \quad (\text{D.5})$$

We plug this conjecture into equation (3)

$$\begin{aligned} \beta_m \hat{m}_t + \beta_g \hat{g}_t &= \hat{m}_t + E_t [MG [\beta_m \hat{m}_{t+1} + \beta_g \hat{g}_{t+1}] + \hat{g}_{t+1}] \\ &= \hat{m}_t + MG [\beta_m \alpha_m \hat{m}_t + \beta_g \alpha_g \hat{g}_t] + \alpha_g \hat{g}_t \\ &= [1 + MG \beta_m \alpha_m] \hat{m}_t + [\alpha_g + MG \beta_g \alpha_g] \hat{g}_t, \end{aligned}$$

where we used $E_t [\varepsilon_{t+1}^m] = E_t [\varepsilon_{t+1}^g] = 0$.

Comparing coefficients on both sides of the equation, we have

$$\begin{aligned} \beta_m = 1 + MG \beta_m \alpha_m &\Leftrightarrow \beta_m = \frac{1}{1 - \alpha_m MG} \\ \beta_g = \alpha_g + MG \beta_g \alpha_g &\Leftrightarrow \beta_g = \frac{\alpha_g}{1 - \alpha_g MG}. \end{aligned}$$

D.3 Aggregate state space system

We start from the definitions of the log price-rent ratio and rent growth, $\log(P_t/R_t) = \log(V_t) = v + \hat{v}_t$ and $\Delta r_t = r_t - r_{t-1} = g + \hat{g}_t$. Combined with the solution (D.5) and the dynamics (4), we can write

$$\begin{aligned} p_t &= r_t + v + \hat{v}_t = r_t + v + \beta_m \hat{m}_t + \beta_g \hat{g}_t, \\ r_t &= r_{t-1} + g + \hat{g}_t = r_{t-1} + g + \alpha_g \hat{g}_{t-1} + \varepsilon_t^g. \end{aligned}$$

To track the evolution of prices and rents within a state space system, the state vector x_t has to contain the current realizations of both AR(1) processes as well as the current rent level r_t (owing to the unit root nature of rents). In addition, we need to add a constant to the state vector. We can write the aggregate state space system as

$$\begin{pmatrix} p_t \\ r_t \end{pmatrix} = Bx_t \quad \text{and} \quad x_t = Ax_{t-1} + C\varepsilon_t \quad (\text{D.6})$$

with

$$B = \begin{pmatrix} \beta_m & \beta_g & 1 & v \\ 0 & 0 & 1 & 0 \end{pmatrix} \quad \text{and} \quad x_t = \begin{pmatrix} \hat{m}_t \\ \hat{g}_t \\ r_t \\ 1 \end{pmatrix} \quad (\text{D.7})$$

as well as

$$A = \begin{pmatrix} \alpha_m & 0 & 0 & 0 \\ 0 & \alpha_g & 0 & 0 \\ 0 & \alpha_g & 1 & g \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad C = \begin{pmatrix} \sigma_m & 0 \\ 0 & \sigma_g \\ 0 & \sigma_g \\ 0 & 0 \end{pmatrix} \quad \text{and} \quad \varepsilon_t = N \left[\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \right]. \quad (\text{D.8})$$

D.4 Characterizing the distribution of nowcasts and forecasts

On the level of the individual, the state space system reads

$$\begin{aligned} s_t^{i,h} &= Bx_t + w_t^{i,h} \quad \text{and} \\ x_t &= Ax_{t-1} + C\varepsilon_t, \end{aligned} \quad (\text{D.9})$$

where $w_t^{i,h}$ is a 2×1 vector of idiosyncratic Gaussian shocks that are uncorrelated across agents and time, $w_t^{i,h} \sim N(0, \mathcal{W}^h)$. The covariance matrix depends on the type h of agent i . This state space system has the same state vector x_t and matrices as in (D.7) and (D.8). If we leave out the individual noise vector $w_t^{i,h}$, we again obtain the aggregate state space system (D.6). Our computations thus work with separate systems for each type. To ease notation, we suppress the dependence of $s_t^{i,h}$ and $x_t^{i,h}$ on type h in what follows.

Forecasts and posterior variance. To compute conditional expectations for the system (D.9), we apply the Kalman filter and thus proceed recursively. Agent i of type h computes his conditional expectation or nowcast x_t^i of x_t given the history of his observables s_t^i . It can be represented as

$$x_t = x_t^i + u_t^i,$$

where the nowcast error u_t^i is orthogonal to the history s_t^i, s_{t-1}^i, \dots and has conditional covariance matrix Σ_t^h . The initial nowcast x_0^i of agent i is drawn from a normal distribution with mean \bar{x}_0^h and a (cross sectional) variance Ω_0^h . The mean and variance depend on the type h of agent i . Standard results on updating with normal distributions imply that x_t^i is also normally distributed.

Note that all agents share the same view on the state space system and use the same now-

and forecasting rules. The only elements that are allowed to depend on type h are the initial mean \bar{x}_0^h and variance Ω_0^h as well as the variance \mathcal{W}^h of the signals. As a result, the covariance matrix Σ_t^h of nowcast errors is identical across individuals of the same type h , but may differ between owners and renters.

For an individual with a signal history $s_{t-1}^i, s_{t-2}^i, \dots$ and a corresponding nowcast x_{t-1}^i , the conditional joint distribution of s_t^i and x_t one period ahead is

$$\begin{aligned} s_t^i | x_{t-1}^i &= B \left(A \left(x_{t-1}^i + u_{t-1}^i \right) + C \varepsilon_t \right) + w_t^i \\ x_t | x_{t-1}^i &= A \left(x_{t-1}^i + u_{t-1}^i \right) + C \varepsilon_t. \end{aligned}$$

The vector $(s_t^{i\top}, x_t^\top)^\top | x_{t-1}^i$ is therefore conditionally normally distributed with known variance covariance matrix

$$\begin{bmatrix} B (A \Sigma_{t-1}^h A^\top + C C^\top) B^\top + \mathcal{W}^h & B (A \Sigma_{t-1}^h A^\top + C C^\top) \\ (A \Sigma_{t-1}^h A^\top + C C^\top) B^\top & A \Sigma_{t-1}^h A^\top + C C^\top \end{bmatrix}.$$

We define the gain matrix

$$\Gamma_t^h = \left(A \Sigma_{t-1}^h A^\top + C C^\top \right) B^\top \left(B \left(A \Sigma_{t-1}^h A^\top + C C^\top \right) B^\top + \mathcal{W}^h \right)^{-1}.$$

The gain again depends on the agent's type h via the noise variance \mathcal{W}^h and its effect on the posterior, because agents only differ in their information set.

We can then write updating of the posterior mean vector and covariance matrix as

$$x_t^i = A x_{t-1}^i + \Gamma_t^h \left(s_t^i - B A x_{t-1}^i \right), \quad (\text{D.10})$$

$$\Sigma_t^h = (I - \Gamma_t^h B) \left(A \Sigma_{t-1}^h A^\top + C C^\top \right). \quad (\text{D.11})$$

Equation (D.10) is the stochastic difference equation (10) for the conditional mean x_t^i . Averaging across agents delivers (D.11) which is the covariance matrix (11). For stationary s_t^i , it has a stationary solution – this is the conditional expectation given the infinite history of signals. The second equation is a matrix Ricatti difference equation. Its positive definite fixed point is the posterior variance given an infinite history of signals. Once we know this fixed point, we also have a time invariant gain matrix Γ^h that summarizes how agents alter their estimates of hidden states in response to observations.

Distribution of forecasts and forecast errors. Prices and rents at date $t + 1$ are $B x_{t+1}$ and their

growth rates are

$$g_t = Bx_{t+1} - Bx_t = B(A - I)x_t + BC\varepsilon_{t+1}.$$

An individual agent's forecast of rent and price based on date t information is therefore Bx_t^i and forecasts of *growth* in price and rent are $f_t^i = B(A - I)x_t^i$. An individual agent's forecast error is

$$g_t - f_t^i = B(A - I)(x_t - x_t^i) + BC\varepsilon_{t+1}.$$

Forecast errors arise when either (i) agents are surprised by the innovation $C\varepsilon_{t+1}$ or (ii) agents have a mistaken nowcast of the hidden state based on their date t information.

To understand the cross sectional distribution of forecasts, we consider the evolution of nowcasts. We can write

$$\begin{aligned} x_t^i &= (I - \Gamma_t^h B) Ax_{t-1}^i + \Gamma_t^h s_t^i, \\ &= (I - \Gamma_t^h B) Ax_{t-1}^i + \Gamma_t^h B (Ax_{t-1} + C\varepsilon_t) + \Gamma_t^h w_t^i, \end{aligned} \quad (\text{D.12})$$

We note that the first and third term vary in the cross section, but are orthogonal, whereas the middle term is aggregate and does not vary across individual agents.

To compute average forecasts, we use the fact that signals in equation (D.12) are unbiased. The cross sectional mean of nowcasts \bar{x}_t^h for agents of type h thus evolves according to

$$\bar{x}_t^h = (I - \Gamma_t^h B) A\bar{x}_{t-1}^h + \Gamma_t^h B (Ax_{t-1} + C\varepsilon_t).$$

Agent types have different nowcast paths if they respond differently to the same news due to their using of different gain matrices.

The cross sectional distribution of forecasts (or forecast errors) depends only on the cross sectional distribution of the nowcasts (or nowcast errors). The cross sections of forecasts and forecasts errors differ only by a constant.

Denote by Ω_t^h the cross sectional variance of the nowcast for type h at date t . From (D.12), it evolves according to

$$\Omega_t^h = (I - \Gamma_t^h B) A\Omega_{t-1}^h A^\top (I - \Gamma_t^h B)^\top + \Gamma_t^h \mathcal{W}^h \Gamma_t^{h\top},$$

a Lyapunov-style difference equation. If we substitute the time invariant gain matrix Γ^h derived above, we can find a fixed point Ω^h that describes the cross sectional variance of nowcasts after an infinite sequence of signals.

Finally, let's turn back to forecasts errors. The average forecast error is

$$g_t - \bar{f}_t^h = B(A - I)(x_t - \bar{x}_t^h) + BC\varepsilon_{t+1}.$$

For the cross sectional variance of forecasts we obtain

$$\begin{aligned} E_t \left[(f_t^i - \bar{f}_t^h)(f_t^i - \bar{f}_t^h)^\top \right] &= E_t \left[B(A - I)(x_t^i - \bar{x}_t^h)(x_t^i - \bar{x}_t^h)^\top (A - I)^\top B^\top \right] \\ &= B(A - I) \Omega^h (A - I)^\top B^\top. \end{aligned}$$

Consequently, we can write the mean squared error of forecasts as

$$\begin{aligned} MSE_t^h &= E_t \left[(g_t - f_t^i)(g_t - f_t^i)^\top \right] \\ &= \underbrace{B(A - I) \Omega^h (A - I)^\top B^\top}_{\text{cross sectional variance of forecasts}} \\ &\quad + \underbrace{\left(B(A - I)(x_t - \bar{x}_t^h) + BC\varepsilon_{t+1} \right) \left(B(A - I)(x_t - \bar{x}_t^h) + BC\varepsilon_{t+1} \right)^\top}_{\text{squared average forecast error}}. \end{aligned}$$