NBER WORKING PAPER SERIES

THE HOUSING MARKET(S) OF SAN DIEGO

Tim Landvoigt
Monika Piazzesi
Martin Schneider

Working Paper 17723
http://www.nber.org/papers/w17723

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2012

We are endebted to Trulia for access and help with their data. We thank Markus Baldauf, John Campbell, Boyan Jovanovic, Lars Ljunqvist, Hanno Lustig, Ellan McGrattan, Francois Ortalo-Magne, Sven Rady, Sergio Rebelo, and seminar participants at Boston University, Carnegie-Mellon University, ECB, Federal Reserve Board, Kellogg, Munich, NYU, Princeton, UNC as well as conference participants at CASEE 2010 at Arizona State, ERID Duke Conference 2010, Duke Macoeconomics Conference 2011, EEA Congress 2010 in Glasgow, HULM Conference 2010 in St. Louis, NBER AP Summer Institute 2011, Richmond Fed 2009 PER Conference, SED 2010 in Montreal, SITE 2010, the 2010 USC Macro Conference, and the 2011 Jerusalem Conference. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed a financial relationship of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w17723.ack

The Housing Market(s) of San Diego
Tim Landvoigt, Monika Piazzesi, and Martin Schneider
NBER Working Paper No. 17723
January 2012, Revised January 2012
JEL No. E21,G10,R20

ABSTRACT

This paper uses an assignment model to understand the cross section of house prices within a metro area. Movers' demand for housing is derived from a lifecycle problem with credit market frictions. Equilibrium house prices adjust to assign houses that differ by quality to movers who differ by age, income and wealth. To quantify the model, we measure distributions of house prices, house qualities and mover characteristics from micro data on San Diego County during the 2000s boom. The main result is that cheaper credit for poor households was a major driver of prices, especially at the low end of the market.

Tim Landvoigt
Department of Economics
Stanford University
579 Serra Mall
Stanford CA 94305-6072
timlandvoigt@googlemail.com

Monika Piazzesi
Department of Economics
Stanford University
579 Serra Mall
Stanford, CA 94305-6072
and NBER
piazzesi@stanford.edu

Martin Schneider
Department of Economics
Stanford University
579 Serra Mall
Stanford, CA 94305-6072
and NBER
schneidr@stanford.edu

1 Introduction

During the recent housing boom, there were large differences in capital gains across houses, even within the same metro area. Figure 1 illustrates the basic stylized fact for San Diego County, California. Every dot corresponds to a home that was sold in both year 2000 and year 2005. On the horizontal axis is the 2000 sales price. On the vertical axis is the annualized real capital gain between 2000 and 2005. The solid line is the capital gain predicted by a regression of capital gain on log price. It is clear that capital gains during the boom were much higher on low end homes. For example, the average house worth \$200K in the year 2000 appreciated by 17% (per year) over the subsequent five years. In contrast, the average house worth \$500K in the year 2000 appreciated by only 12% over the subsequent five years.[1]

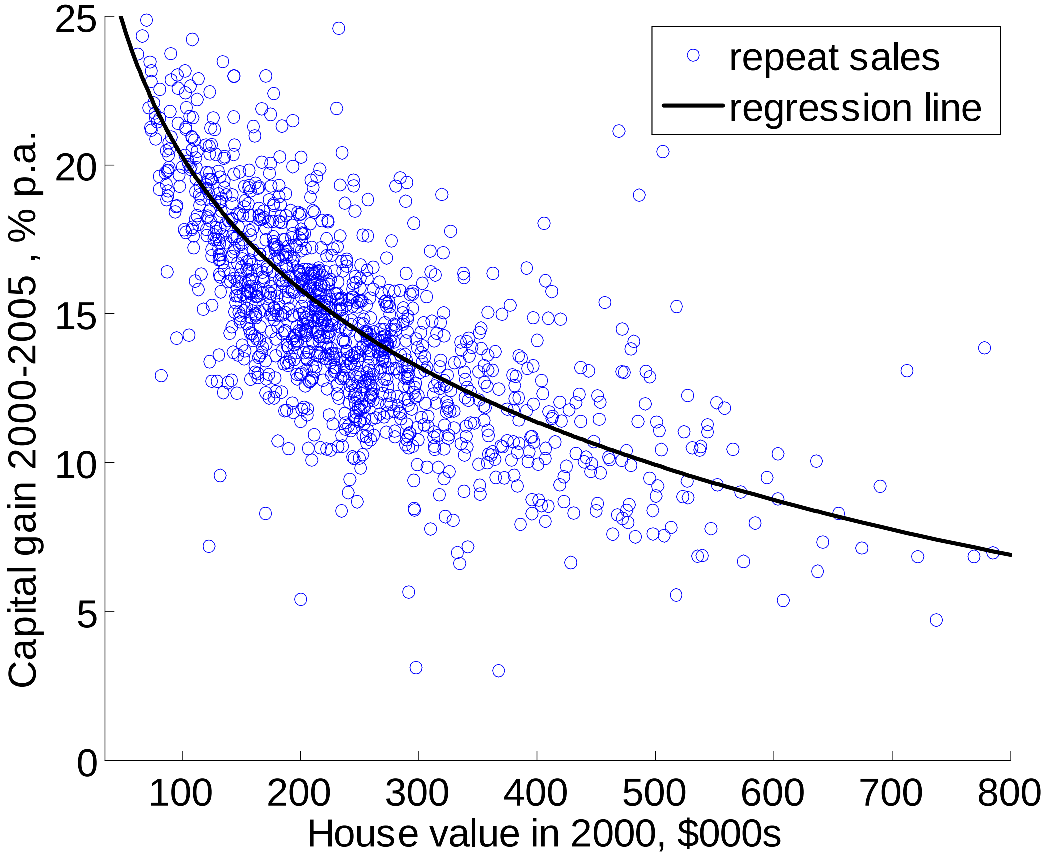

Figure 1: Repeat sales in San Diego County, CA, during the years 2000-2005. Every dot represents a residential property that was sold in 2000 and had its next sale in 2005. The horizontal axis shows the sales price in 2000. The vertical axis shows the real capital gain per year (annualized change in log price less CPI inflation) between 2000 and 2005.

This paper considers a quantitative model of the housing market in the San Diego metro area over the boom period. Its key feature is that houses are indivisible and movers are assigned,

[1]While Figure 1 only has repeat sales from two particular years, Table 1 below documents the basic stylized fact in a joint estimation using *all* repeat sales in San Diego County over the last decade.

in equilibrium, to one of a large number of house types. We use the model to connect various changes in the San Diego housing market (or market*s* – one for each type) with the cross section of capital gains there. In particular, we look at changes in the composition of houses that were transacted, shifts in the distribution of movers' characteristics, and the availability of cheap credit.

The main finding is that the availability of cheap credit has larger effects on housing demand at the low end of the market, with strong implications for relative prices. In addition, shifts in the distributions of houses and movers are important for understanding the cross section of prices. For example, the relatively larger number of low quality houses transacted during the boom led to richer marginal investors at the low end, which also contributed to higher capital gains there. Once we take into account changes in credit conditions, a lifecycle model of housing demand is consistent not only with the large and uneven price changes apparent in Figure 1, but also with the key moments of the joint distribution of house quality, age, wealth and income.

In the model, movers meet houses. Houses differ by quality: there is a continuum of indivisible houses provide different flows of housing services. Movers differ by age, income, and wealth: their demand for housing is derived from an intertemporal savings and portfolio choice problem with transaction costs and collateral constraints. In equilibrium, prices adjust to induce agents with lower demand for housing services to move into lower quality houses. The distribution of equilibrium prices thus depends on the (three-dimensional) distribution of movers' characteristics as well as the distribution of house qualities.

To implement the model quantitatively, we use micro data to measure the distribution of movers' characteristics and the quality distribution of transacted houses. We do this both for the year 2000 and for the year 2005 – the peak of the boom. We then compute model predictions for equilibrium prices in both years and derive the cross section of capital gains by quality. We compare those predictions to a repeat sales model estimated on transaction data. Our numerical solution not only finds equilibrium prices, but also an equilibrium assignment that we compare to the assignment in the data. To understand how the model works, we explore model-implied capital gains for different assumptions on changes in the environment; in particular, we capture cheaper credit by lower interest rates and downpayment constraints.

Our quantitative analysis starts from a benchmark specification that is based on measured changes in the distribution of mover characteristics as well as changes in credit conditions that are perceived to be temporary. This specification generates a substantial housing boom: capital gains between 2000 and 2005 are 13% (7%) per year for houses initially worth $200K ($500K). At the same time, we check that the equilibrium assignment predicted by the model resembles the assignment observed in the data. In particular, the model matches the fact that house quality rises faster with wealth and income for younger cohorts of households. The reason is that in a lifecycle model with nontradable labor income and collateral constraints, younger households choose more levered portfolios and thus invest a larger share of cash on hand in housing.

Our benchmark specification generates capital gains about 5% below those in the data, with a larger discrepancy in the bottom quintile of the quality distribution. We study several extensions that produce capital gains closer to the data. First, we show that if changes in credit conditions are perceived to be long-lived, then demand at the low end is even higher, as in the data. Second, while our baseline specification assumes that price expectations in 2005 exhibit mean reversion, more optimistic expectations help increase prices, especially at the high end. Finally, we consider a specification with systematic changes in quality growth that are chosen to match prices exactly. We show that under this specification changes in credit conditions remain a key driver of capital gains.

Our exercise fits into a tradition that links asset prices to fundamentals through household optimality conditions. For housing, this tradition has given rise to the "user cost equation": the per-unit price of housing is such that all households choose their optimal level of a divisible housing asset. With a *single* per-unit price of housing, capital gains on all houses are the same. In our model, there is no single user cost equation since there are many types of indivisible houses, with marginal investors who differ across house types.[2] Instead, there is a *separate* user cost equation for every house type, each reflecting the borrowing costs, transaction costs, and risk premia of only those movers who buy that house type. Changes in the environment thus typically give rise to a nondegenerate cross section of capital gains.

The fact that there is a *family* of user cost equations is crucial for our results in two ways. First, it implies that changes in the environment that more strongly affect a subset of movers will more strongly affect prices of houses which those movers buy. For example, lower minimum downpayment requirements more strongly affect poor households for whom this constraint is more likely to be binding. As a result, lower downpayment constraints lead to higher capital gains at the low end of the market.

Second, higher moments of the quality and mover distributions matter. For example, we show that the quality distribution of transacted homes in San Diego County at the peak of the boom had fatter tails than at the beginning of the boom. This implied, in particular, that relatively lower quality homes had to be assigned to relatively richer households than before the boom. For richer households to be happy with a low quality home, homes of slightly higher quality had to become relatively more expensive. The price function thus had to become steeper at the low end of the market, which contributed to high capital gains in that segment.

Since we measure the quality distribution of transacted homes directly from the data, we do not take a stand on where the supply of houses comes from. A more elaborate model might add an explicit supply side, thus incorporating sellers' choice of when to put their house on the market, the effects of the availability of land to developers (as in Glaeser, Gyourko, and Saks 2005), or gentrification (as in Guerrieri, Hartley, and Hurst 2010.) At the same time, any model with an explicit supply side also gives rise to an equilibrium distribution of transacted homes that has to be priced and assigned to an equilibrium distribution of movers. The assignment and pricing equations we study thus hold also in equilibrium of *many* larger models with different supply side assumptions. Our results show what it takes to jointly match prices and mover characteristics, independently of the supply side. In this sense, our exercise is similar to consumption-based asset pricing, where the goal is to jointly match consumption and prices, also independently of the supply side.

To our knowledge, this paper is the first to study a quantitative assignment model with a

[2] In contrast, the user cost equation determines a unique price per unit of housing, so every investor is marginal with respect to every house.

continuum of houses and a multidimensional distribution of mover characteristics. Assignment models with indivisible heterogeneous goods and heterogeneous agents have been used in several areas of economics, most prominently to study labor markets where firms with different characteristics hire workers with different skill profiles (for an overview, see Sattinger, 1993.) In the context of housing, an early reference is Kaneko (1982). Caplin and Leahy (2010) characterize comparative statics of competitive equilibria in a general setting with a finite number of agents and goods. Stein (1995), Ortalo-Magne and Rady (2006) and Rios-Rull and Sanchez-Marcos (2008) study models with two types of houses and credit constraints.

We provide new evidence on the cross section of capital gains as well as the composition of trading volume by quality over the recent housing boom. Our results use property level data for San Diego County and several statistical models of price change. Our finding of a nontrivial cross section of capital gains is related to existing empirical studies that compare house price dynamics across price segments within a metro area, for example Poterba (1991), Case and Mayer (1996), Case and Shiller (2005), and Guerrieri, Hartley, and Hurst (2010).[3] Existing studies of volume emphasize the comovement of volume and price changes *over time* as "hot" markets with high prices and high volume turn into "cold" markets with low prices and low volume (for example, Stein 1995.) Our results show that during the recent boom the relationship between prices and volume *in the cross section* was nonmonotonic: volume became relatively higher both for cheap houses and for expensive houses.

Our quantitative model considers jointly the effect of credit constraints and changes in the house quality and mover distributions on prices as well as the cross section of household portfolios. Reduced form evidence has suggested that both credit and changes in distributions can matter for prices. For example, Poterba (1991) points to the role of demographics, whereas Bayer, Ferreira, and McMillan (2007) highlight the importance of amenities (such as schools), and Guerrieri, Hartley and Hurst (2010) relate price changes to gentrification, that is, changes in neighborhood quality. Empirical studies have also shown that credit constraints matter for house prices at the regional level. Lamont and Stein (1999) show that house prices react more strongly to shocks in cities where more households are classified as “borrowing constrained”. Mian and Sufi (2010) show for the recent US boom that house price appreciation and borrowing were correlated across zip codes. Mian and Sufi (2009) show that areas with many subprime borrowers saw a lot of borrowing even though income there declined.

Our exercise infers the role of cheap credit for house prices from the *cross section* of capital gains by quality. This emphasis distinguishes it from existing work with quantitative models of the boom. Many papers have looked at the role of cheap credit or exuberant expectations for prices in a homogeneous market (either a given metro area or the US.) They assume that houses are homogeneous and determine a single equilibrium house price per unit of housing capital. As a result, equilibrium capital gains on all houses are identical, and the models cannot speak to the effect of cheap credit on the cross section of capital gains. Recent papers on the role of credit include Himmelberg, Mayer, and Sinai (2005), Glaeser, Gottlieb, and Gyourko (2010), Kiyotaki, Michaelides, and Nikolov (2010) and Favilukis, Ludvigson, and Van Nieuwerburgh (2010). The latter two papers also consider collateral constraints, following Kiyotaki and Moore (1995) and

[3] Interestingly, these studies do not find a common capital gain pattern across all booms – depending on time and region, low quality houses may appreciate more or less than high quality houses during a boom. This suggests that it is fruitful to study a single episode in detail, as we do in this paper.

Chien and Lustig (2010). Recent papers on the role of expectation formation include Piazzesi and Schneider (2009), Burnside, Eichenbaum, and Rebelo (2010), and Glaeser, Gottlieb, and Gyourko (2010).

The paper proceeds as follows. Section 2 presents evidence on prices and transactions by quality segment in San Diego County. Section 3 presents a simple assignment model to illustrate the main effects and the empirical strategy. Section 4 introduces the full quantitative model.

2 Facts

In this section we present facts on house prices and the distribution of transacted homes during the recent boom. We study the San-Diego-Carlsbad-San-Marcos Metropolitan Statistical Area (MSA) which coincides with San Diego County, California.

2.1 Data

We obtain evidence on house prices and housing market volume from county deeds records. We start from a database of all deeds written in San Diego County between 1997 and 2008. In principle, deeds data are publicly available from the county registrar. To obtain the data in electronic form, however, we take advantage of a proprietary database made available by Trulia.com. We use deeds records to build a data set of households' market purchases of single-family dwellings. This involves screening out deeds that reflect other transactions, such as intrafamily transfers, purchases by corporations, and so on. Our screening procedure together with other steps taken to clean the data is described in Appendix A.

To learn about mover characteristics, we use several data sources provided by the U.S. Census Bureau. The 2000 Census contains a count of all housing units in San Diego County. We also use the 2000 Census 5% survey sample of households that contains detailed information on house and household characteristics for a representative sample of about 25,000 households in San Diego County. We obtain information for 2005 from the American Community Survey (ACS), a representative sample of about 6,500 households in San Diego Country. A unit of observation in the Census surveys is a dwelling, together with the household who lives there. The surveys report household income, the age of the head of the household, housing tenure, as well the age of the dwelling, and a flag on whether the household moved in recently.[4] For owner-occupied dwellings, the census surveys also report the house value and mortgage payments.

2.2 The cross section of house prices and qualities

In this section we describe how we measure price changes over time conditional on quality as well as changes over time in the quality distribution. We first outline our approach. We want to understand systematic patterns in the cross section of capital gains between 2000 and 2005. We establish those patterns using statistical models that relate capital gain to 2000 price.

[4] In the 2005 ACS, the survey asks households whether they moved in the last year. In the 2000 Census, the survey asks whether they moved in the last two years.

The simplest such model is the black regression line in Figure 1. Below we describe a more elaborate model of repeat sales as well as a model of price changes in narrow geographic areas – the patterns are similar across all these models.

If there is a one-dimensional quality index that households care about, then house quality at any point in time is reflected one-for-one in the house price. In other words, the horizontal axis in Figure 1 can be viewed as measuring quality in the year 2000. The regression line measures common changes in price experienced by all houses of the same initial quality. More generally, any statistical model of price changes gives rise to an expected price change that picks common changes in price by quality.

There are two potential reasons for common changes in price by quality. On the one hand, there could be common changes in quality itself. For example, quality might increase because the average house in some quality range is remodeled, or the average neighborhood in some quality range obtains better amenities.[5] On the other hand, there could simply be *revaluation* of houses in some quality range while the average quality in that range stays the same. For example, prices may change because more houses of similar quality become available for purchase. In practice, both reasons for common changes in price by quality are likely to matter, and our structural model below thus incorporates both.

Independently of the underlying reason for price changes, we can determine the number of houses in the year 2005 that are "similar" to (and thus compete with) houses in some given quality range in the year 2000. This determination uses a statistical model of price changes together with the cross section of transaction prices. Consider some initial house quality in the year 2000. A statistical model of price changes – such as the regression line in Figure 1 – says at what price the average house of that initial quality trades in 2005. For example, from the regression line we can compute a predicted 2005 price by adding the predicted capital gain to the 2000 price. Once we know the predicted 2005 prices for the initial quality range, we can read the number of similar houses off the cross sectional distribution of 2005 transaction prices.

In our context, we can say more: counting for *every* initial 2000 quality range the "similar" houses in 2005 actually delivers the 2005 quality distribution, up to a monotonic transformation of quality. This is because the predicted 2005 price from a statistical model is *strictly increasing* in the initial 2000 price, as we document below using both parametric and nonparametric specifications. Since price reflects quality in both years, it follows that for a given 2005 quality level, there is a unique initial 2000 quality level such that the average house of that initial quality resembled the given house in 2005.[6]

Of course, we do not know the mapping from initial 2000 quality to average 2005 quality, because a model of price changes does not distinguish between common changes in quality and revaluation. Nevertheless, since we know that the mapping is monotonic, we can represent the

[5] Importantly, changes in quality will be picked up by the expected price change only if they are common to all houses of the same initial quality, that is, they are experienced by the average house in the segment. The figure shows that, in addition, there are also large idiosyncratic shocks to houses or neighborhoods.

[6] In the San Diego housing market, common changes in quality between 2000 and 2005 did thus not upset the relative ranking of house quality segments: the average house from a high quality range in 2000 was worth more—and hence of higher quality—than the average house from a low quality range in 2000. This does not mean, of course, that there were no changes in the ranking of individual houses or neighborhoods—in our statistical model those are captured by idiosyncratic shocks.

2005 quality distribution, up to a monotonic transformation, by the distribution of "similar" 2005 houses by 2000 quality. In other words, the 2000 price can serve as an ordinal index of quality. Quality distributions for both 2000 and 2005 can be measured in terms of this index and then used as an input into the quantitative implementation of our structural model below.

Statistical model of price changes by quality

Consider a loglinear model of price changes at the individual property level. This is the statistical model we use to produce inputs for the structural model. To capture the cross section of capital gains by quality, we allow the expected capital gain to depend on the current price. Formally, let p_t^i denote the log price of a house i at date t. We assume that the capital gain on house i between dates t and $t+1$ is

$$p_{t+1}^i - p_t^i = a_t + b_t p_t^i + \varepsilon_{t+1}^i, \tag{1}$$

where the idiosyncratic shocks ε_{t+1}^i have mean zero and are such that a law of large numbers holds in the cross section of houses. For fixed t and $t+1$, the model looks like the regression displayed in Figure 1.

The model estimated here differs from a simple regression since (1) is assumed to hold for every t in our sample, so that the coefficients can be estimated with data on all repeat sales simultaneously. We find this approach useful since a regression based only on 2000-2005 repeat sales might suffer from selection bias – it would be based only on houses that were bought at the beginning of the boom and sold at the peak.[7] In contrast, under our approach the estimated coefficient a_t reflects *any* repeat sale that brackets the year t. For example, the coefficients for $t = 2004$, say, reflect repeat sales in the hot phase of the boom between 2002 and 2005, but also repeat sales between 2003 and 2008.

Equation (1) differs from a typical time series model for returns in that the coefficients are time dependent. It is possible to identify a separate set of coefficients for every date because we have data on many repeat sales. The coefficients b_t determine whether there is a nontrivial cross section of expected capital gains. If housing were a homogenous capital good, then it should not be possible to forecast the capital gain using the initial price level p_t, that is, $b_t = 0$. The expected capital gain on all houses would be the same (and equal to a_t), much like the expected capital gain is the same for all shares of a given company. More generally, a nonzero coefficient for b_t indicates that quality matters for capital gains. For example $b_t < 0$ means that prices of low quality houses that are initially cheaper will on average have higher capital gains. In contrast, $b_t > 0$ says that expensive houses are expected to appreciate more (or depreciate less).

Suppose $\left(p_t^i, p_{t+k}^i\right)$ is a pair of log prices on transactions of the same house that took place in years t and $t+k$, respectively. Equation (1) implies a conditional distribution for the capital gain over k periods,

$$p_{t+k}^i - p_t^i = a_{t,t+k} + b_{t,t+k}\, p_t^i + \varepsilon_{t,t+k}^i, \tag{2}$$

where the coefficients $a_{t,t+k}$, $b_{t,t+k}$ are derived by iterating on equation (1). We estimate the parameters (a_t, b_t) by GMM; the objective function is the sum of squared prediction errors,

[7] Below we compare GMM estimation results to regression results based on property level, zip code, and census tract prices reported in Appendix B. These results also suggest that selection bias is not a problem.

weighted by the inverse of their variance. The GMM estimation uses data from repeat sales between all pairs of years jointly by imposing the restriction that the multiperiod coefficients $a_{t,t+k}$ and $b_{t,t+k}$ are appropriate weighted sums and products of future a_t and b_t coefficients between t and $t+k$, respectively.

Table 1 reports point estimates based on 70,315 repeat sales in San Diego County that occurred during 1997-2008. Details of how we screen repeat sales are in Appendix A. The first row in the table shows the sequence of estimates for the intercept a_t. For example, the estimated a_t for the year 1999 is the intercept in the expected capital gain from 1999 to 2000. The intercept is positive for expected capital gains during the boom phase 2000-2005, and negative during 2006-2008, reflecting average capital gains during those two phases. The middle row shows the slope coefficients b_t. During the boom phase, the slopes are strongly negative. For example, for the year 2002 we have $b_t = -.09$, that is, a house worth 10% more in 2002 appreciated by .9% less between 2002 and 2003. During the bust phase, the relationship is reversed: positive b_ts imply that more expensive houses depreciated relatively less. The estimated expected capital gains during the boom years are large: a house price between \$100K and \$500K corresponds to $p_t^i \in (11.5, 13.1)$, and the resulting expected capital gain $a_t + b_t p_t^i$ reaches double digits in many years on all houses in this range.

Table 1: Estimated Coefficients from Repeat Sales Model for San Diego

Year	1999	2000	2001	2002	2003	2004	2005	2006	2007
a_t	0.76	1.29	1.41	1.30	0.87	0.60	−0.56	−1.09	−3.18
	(0.04)	(0.04)	(0.04)	(0.04)	(0.05)	(0.06)	(0.07)	(0.10)	(0.12)
b_t	−0.05	−0.093	−0.10	−0.09	−0.05	−0.04	0.04	0.07	0.22
	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)	(0.004)	(0.01)	(0.01)	(0.01)
σ_t^i	8.8	8.3	8.6	8.2	8.0	8.4	9.7	11.4	13.8

Note: This table reports estimates for coefficients a_t, b_t and the volatility σ_t^i in equation (1) for the indicated years. The data for this estimation are the 70,315 repeat sales in San Diego County during the years 1997-2008. The numbers in brackets are standard errors.

In a second step, we estimate the variances of the residuals (σ_t^2) by maximum likelihood, assuming the shocks ε_{t+1}^i are normally distributed and iid over time as well as in the cross section. This is to get an idea of the idiosyncratic volatility of housing returns faced by households who buy a single property. The results are reported in the bottom row of Table 1. Volatility is around 9% on average, slightly higher than the idiosyncratic volatility of 7% reported by Flavin and Yamashita (2002).[8] Another interesting pattern is that idiosyncratic volatility increased by more than half in the bust period.

[8] Table 1A in Flavin and Yamashita (2002) reports a 14% return volatility for individual houses. Their Table 1B reports a 7% volatility for the Case-Shiller city index for San Francisco, which is comparable to San Diego. The difference between these two numbers is a 7% idiosyncratic volatility.

To construct the quality distribution for 2005 below, what matters are the coefficients of the predicted 2005 price given the 2000 price. To ease notation, set t equal to the year minus 2000, so we are interested in $a_{0,5} = 4.75$ and $b_{0,5} = -.322$. The predicted price for 2005

$$\widehat{p}_5 = a_{0,5} + (1 + b_{0,5})p_0, \tag{3}$$

is therefore strictly increasing as a function of the 2000 price. In other words, even though lower quality houses appreciated more during the boom, there were no segments that became systematically more valuable than other segments. In Appendix B, we show that this monotonicity is not an artifact of our loglinear functional form (1). Nonparametric regressions of 2005 log price on 2000 log price reveal only small deviations from linearity, and the predicted price function is also strictly increasing.

Quality distributions

Let Φ_0 denote the cumulative distribution function (cdf) of log transaction prices in the year 2000 (or $t = 0$). With the 2000 price as the quality index, the cdf of house qualities is $G_0(p_0) = \Phi_0(p_0)$. This cdf is constructed from all 2000 transactions, not only repeat sales. The repeat sales model describes price movements of houses that exist both in 2000 and in later years. In particular, between 2000 and 2005 (that is, $t = 0$ and $t = 5$), say, common shocks move the price of the average house that starts at quality p_0 in year 0 to the predicted price (3) in year $t = 5$. Since the mapping from year 0 quality p_0 to year t price is monotonic ($1 + b_{0,5} > 0$), we know that common changes in quality do not upset the relative ranking of house qualities. Of course, the quality ranking of individual houses may change because of idiosyncratic shocks – for example, some houses may depreciate more than others. These shocks average to zero because of the law of large numbers.

We now turn to the quality distribution in 2005. Let Φ_5 denote the cdf of all log transaction prices in $t = 5$. We know that the average house that starts at quality p_0 in year 0 trades at the price $\hat{p}_5$ in year 5. We define the fraction of houses of quality lower than p_0 as

$$G_t(p_0) = \Phi_t(a_{0,5} + (1 + b_{0,5})p_0).$$

By this definition, the index p_0 tracks relative quality across years. If the same set of houses trades in both years 0 and 5, then the quality distributions G_5 and G_0 are identical. More generally, G_5 can be different from G_0 because different sets of houses trade at the two dates. For example, if more higher quality houses are built and sold in $t = 5$, then G_5 will have more mass at the high end.

Figure 2 shows the cumulative distribution functions G for the base year 2000 as well as for $t = 5$. The cross sectional distribution of prices Φ are taken from Census and ACS data, respectively. (We have also constructed distributions directly from our deeds data, with similar results.) The key difference between the two quality distributions is that there was more mass in the tails in the year 2005 (green line) than in 2000 (blue line.) In other words, the year 2005 saw more transactions of low and high quality homes compared to the year 2000.

What does the one-dimensional quality index measure?

Our approach treats San Diego County as a common housing market and assumes a one-dimensional quality index. The index combines all relevant characteristics of the house which

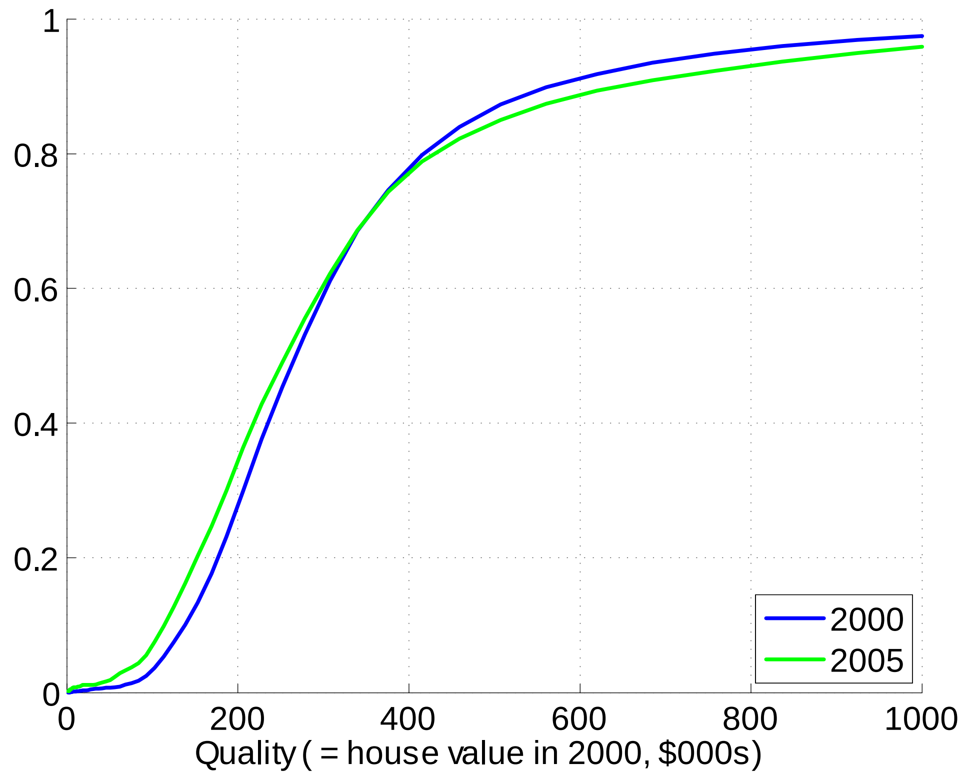

Figure 2: Cumulative distribution function of house qualities in 2000 and 2005

includes features of land, structure, and neighborhood. Above, we have estimated the cross section of capital gains by quality from property-level price data. An alternative approach is to look at median prices in narrow geographic areas such as zipcodes or census tracts. If market prices approximately reflect a one dimensional quality index, then the two approaches should lead to similar predictions for the cross section of capital gains. Moreover, adding geographic information should not markedly improve capital gain forecasts for individual houses.

Appendix B investigates the role of geography by running predictive regressions for annualized capital gains between 2000 and 2005. First, we consider the cross section of capital gains by area, with area equal to either zipcode or census tract, defined as the difference in log median price in the area. We regress the area capital gain from 2000 to 2005 on the 2000 median area price (in logs). We compare the results to a regression of property capital gain on the initial property price as considered above. The coefficients on the initial area price variables are close (between .06 and .07) and the R^2 is similar (around 60%). The GMM estimate for $b_{0,5}$ implied by our repeat sales model above was $b_{0,5} = -.322 = -.064 \times 5$ and is thus also in the same range on an annualized basis. Moreover, the predicted capital gains for the median house ($p^i_{2000} = \log(247,000) = 12.42$) are within one percentage point of each other. We conclude that the price patterns we find are not specific to a repeat sales approach.

Second, we consider predictive regressions for property level capital gains that include not only the initial property price but also the initial area median price. For both zipcode and census tract, the coefficient on the area price is economically small (less than .015), and in the case of census tract it is not significant. The coefficient on the initial property price is almost unchanged. In both cases, the R^2 increases only marginally by 0.01 percentage points. These results are again consistent with our assumption that house prices reflect a one-dimensional index that aggregates house and neighborhood characteristics.

Given these findings, it makes sense to construct quality distributions from property level data. This will allow us to accurately capture shifts in quality that happen within narrow geographic areas. To illustrate this heterogeneity as well as the source of shifts in the quality distribution, Figure 3 shows maps of San Diego County. The left hand panel is a map of 2005 housing transactions in the western half of the county. County area to the east is omitted because it comprises sparsely inhabited rural mountain terrain and the Anza-Borrego desert. Each dot in the map is a transaction, with the colormap reflecting price from light blue (cheap) to pink (expensive). The grey lines delineate zip codes. There is geographic clustering: in the rich suburbs along the Pacific, most traded houses are expensive, whereas in the poor areas around downtown most traded houses are cheaper. However, variation is also apparent within narrow geographic areas, and certainly at the level of delineated zipcodes.

The middle and right panels of Figure 3 illustrate the shift in the quality distribution from 2000 to 2005. In particular, the increase in the share of low quality houses in Figure 2 had two components. First, the share of volume in low quality neighborhoods increased at the expense of volume in high quality neighborhood. The middle panel colors census tracts by the change (between 2000 and 2005) in their share of total countywide volume. Grey areas are census tracts in which the share of total volume changed by less than .05% in absolute value. The warm colors (with a colormap going from red = +.05% to yellow = + 1%) represent census tracts for which the share of volume increased. In contrast, the cold colors (with a colormap going from blue = -1% to green = $-.05\%$) indicates census tracts that lost share of volume. Comparing the left and middle panel, a number of relatively cheaper inland suburbs increased their contribution to overall volume, whereas most expensive coastal areas lost volume.

Second, the share of low quality volume increased within census tracts, and here the direction is less clearly tied to overall area quality. The right panel colors census tracts by the change (between 2000 and 2005) in the share of census tract volume that was contributed by houses in the lowest quintile of the overall county quality distribution. Here grey areas are census tracts in which the share of total volume changed by less than 5% in absolute value. Warm colors (with a colormap from red = 5% to yellow = 60%) represent census tracts in which the composition of volume changed towards more low quality housing. The cold colors (colormap from blue = -60% to green = -5%) show tracts where the composition changes away from low quality housing. Comparing the left and right panels, many of the inland neighborhoods that saw an overall increase in volume also experienced an increase in the share of low quality housing. At the same time, there is less low quality housing in the downtown area, which did not see unusual volume. Moreover, even some of the pricey oceanfront zipcodes saw an increase in the share of low quality houses.

2.3 Mover characteristics

Below we model the decisions by movers, so we are interested in the characteristics of movers in San Diego in 2000 and 2005. Table 2 shows summary statistics on the three dimensions of household heterogeneity in our model: age, income, and wealth. For comparison, we also

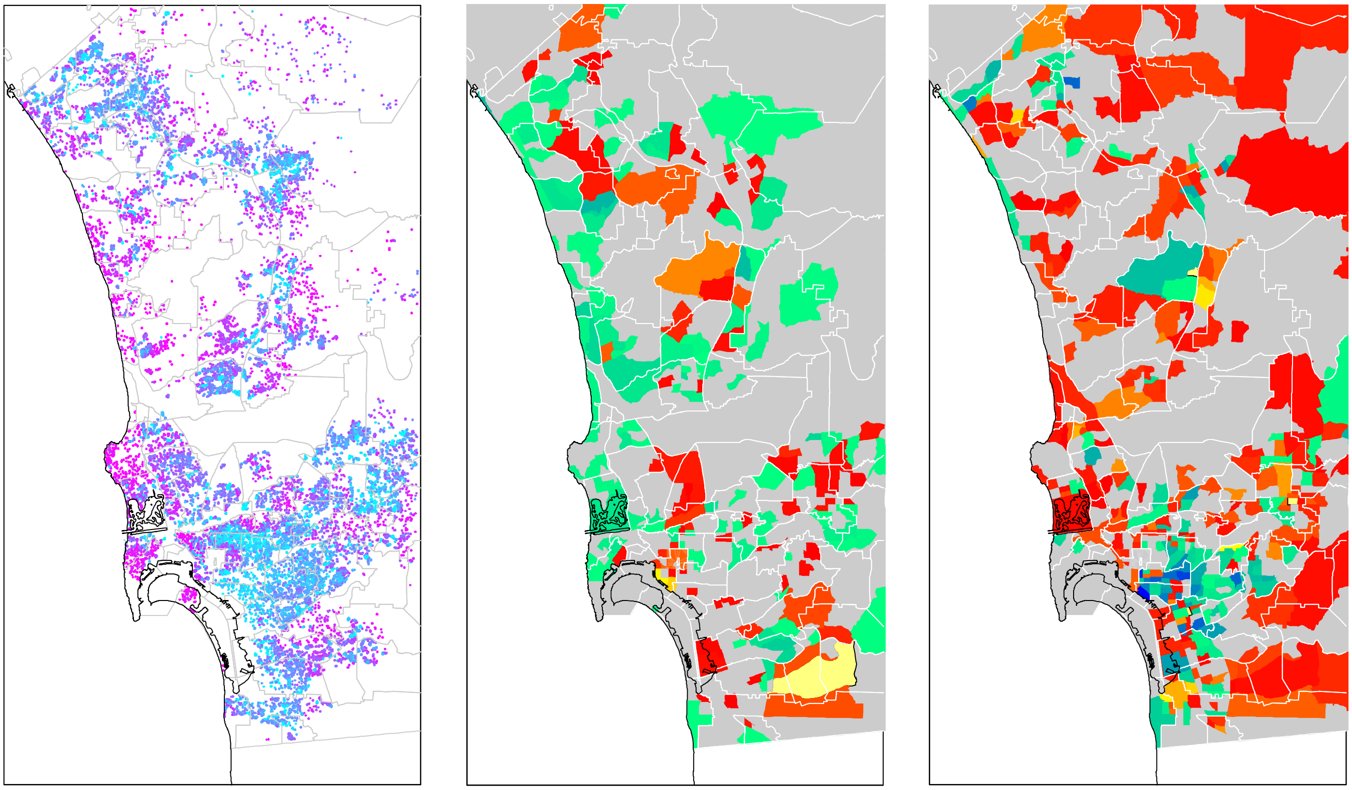

Figure 3: *Left panel*: individual transactions in San Diego County; each dot is a house that was sold in 2005. Color indicates 2005 price ranging from light blue (cheap) to pink (expensive). Grey lines delineate zip codes. *Middle panel*: census tracts colored by change (between 2000 and 2005) in their share of total countywide volume. Warm colors indicate areas where volume increased, with change in share of volume increasing from red to yellow. Cold colors indicate areas where volume decreased, with change in share of volume decreasing from green to blue. Grey areas are census tracts in which the share of total volume changed by less than .05%. *Right panel*: census tracts colored by change in the share of census tract volume contributed by houses in the lowest quintile of the overall county quality distribution. Warm (cold) colors indicate census tracts in which the composition of volume changed towards more (less) low quality housing. Here grey areas are census tracts in which the share of total volume changed by less than 5% in absolute value.

report statistics for stayer households. The difference in mover versus stayer characteristics is particularly pronounced in the year 2005, at the peak of the housing boom. This finding underscores the importance of measuring the characteristics of movers, which are the households whose optimality conditions we want to evaluate.

Table 2 shows that movers tend to be younger than stayers. In San Diego, roughly 13% of stayer households are aged 35 years and younger. Among movers, this fraction is almost three times as large in the year 2000. It further increases to 46% in the year 2005. Table 2 also shows that the median income of younger households is roughly the same as the median income of older households. However, younger households are poorer than older households; older households have about 2.5 times as much wealth as younger households. Finally, movers are somewhat poorer than stayers.

The reported medians mask substantial heterogeneity within each characteristic. For example, the top 10 percent richest households tend to receive roughly 20 percent of the total income earned by their age group, which illustrates the fact that there is income inequality. This inequality is even more pronounced for total wealth, where the top 10 percent households own 50 percent of the total wealth in their age group. The amount of inequality stays roughly the same across the two years, 2000 and 2005.

Table 2: Characteristics of San Diego Movers and Stayers

	Year 2000		Year 2005	
	Movers	Stayers	Movers	Stayers
Fraction of households				
aged ≤ 35 years	0.34	0.13	0.46	0.14
aged > 35 years	0.66	0.87	0.54	0.86
Median Income (in thousands)				
aged ≤ 35 years	74.1	74.8	77.5	86.7
aged > 35 years	82.3	74.4	88.7	78.5
Median Wealth (in thousands)				
aged ≤ 35 years	145.0	161.2	222.3	257.0
aged > 35 years	361.4	402.2	603.3	724.7

Note: This table reports summary statistics for stayer and mover households in San Diego County for the years 2000 and 2005. All dollar numbers are in 2005 dollars, and are thus comparable. The table has two age bins for household heads: aged 35 and younger, older than 35 years. For age and income, we use age of the household head and income reported in the 2000 Census and 2005 ACS. For wealth, we use imputed wealth with data from the Survey of Consumer Finances. The appendix explains the details.

3 Assigning houses to movers

We consider an assignment model of a city. A group of mover households faces an inventory of available houses. Houses are indivisible and come in different qualities indexed by $h \in [0, 1]$. The one-dimensional quality index h summarizes various aspects of housing that households care about (for example square footage, location, views or amenities such as schools.) The inventory of houses is described by a strictly increasing cumulative distribution function $G(h)$. A house of quality h trades in a competitive market at the price $p(h)$.

Every mover household buys exactly one house. Let $h^*(p,i)$ denote the housing demand function of household i. It depends on the house price function p as well as on household i's characteristics. In equilibrium, the markets for all house types clear. For every $h \in [0,1]$, the number of households who demand houses of quality less than h must therefore be equal to the number of such houses in the inventory:

$$\Pr\left(h^*(p,i) \leq h\right) = G(h). \tag{4}$$

The price function $p(h)$ describes a set of house prices at which households are happy to be assigned to the available inventory of houses.

How hard it is to solve a model like this depends on (i) how housing demand h^* is derived and (ii) what the distribution of movers looks like. Our quantitative model derives housing demand from an intertemporal optimization program with uncertainty and frictions (borrowing constraints and transaction costs.) Moreover, households differ by age, wealth, and income so that the distribution of movers is three-dimensional. Both housing demand and equilibrium prices must then be determined numerically.

Before introducing the full model, we illustrate the properties of the setup by studying a simpler version with only one mover characteristic, wealth. While this version is not suitable for quantitative work, it is helpful to get intuition since prices and assignments are available in closed form. Section 3.1 also assumes that demand is derived from a frictionless, deterministic, one-period optimization problem. We use the resulting version to do comparative statics with the house and mover distribution, and to compare our setup to other models of house prices. Section 3.2 provides an extension to intertemporal demand that allows us to discuss some additional issues in a still simple context.

3.1 Static demand, one mover characteristic

Households care about two goods: housing and other (numeraire) consumption. Households start with wealth w and buy a house of quality h at the price $p(h)$. Households also choose their consumption of numeraire c. A household maximizes utility

$$u(c,h) \tag{5}$$

subject to the budget constraint

$$c + p(h) = w.$$

Let $F(w)$ denote the strictly increasing cumulative distribution function of wealth w defined on the nonnegative real line. An equilibrium consists of a consumption and house allocation together with a price function such that households optimize and markets clear.

The first order condition for the household problem is

$$p'(h) = \frac{u_2(c,h)}{u_1(c,h)}. \tag{6}$$

It says that the marginal rate of substitution (MRS) of housing for numeraire consumption equals the *marginal* value of a house $p'(h)$ at the quality level h that the household chooses.

Intratemporal Euler equations that equate MRS and house prices hold in many models of housing. What is special here is that the house price does not need to be linear in quality. The MRS is thus equated to a house price that may differ across quality levels. In this sense, houses of different quality are priced by different marginal investors.

Consider an equilibrium such that optimal house quality is unique and strictly increasing in wealth. The assignment of houses to wealth levels is then given by a strictly increasing function $h^* : \mathbb{R}_0^+ \to [0,1]$. It is convenient to work with its inverse $w^*(h)$, which gives the wealth level of an agent who is assigned a house of quality h. The market clearing condition (4) now says that, for all h,

$$F(w^*(h)) = G(h) \Longrightarrow w^*(h) = F^{-1}(G(h)),$$

which works because F is strictly increasing. The assignment of wealth levels to house qualities depends only on the respective distributions, and is independent of preferences. Of course, prices will depend on preferences through the Euler equations.

The function w^* describes a QQ plot commonly used to compare probability distributions. Its graph is a curve in (h, w)-space that is parametrized by the common cdf value in $[0,1]$. The shape of the graph w^* is determined by the *relative* dispersion of house quality and wealth. If the relative dispersion is similar, the graph of w^* is close to the 45 degree line. In the case where the quality and wealth distributions are identical, the graph of w^* is exactly the 45 degree line. If the distribution of wealth F is more dispersed than the quality distribution G, the graph of w^* is steeper than the 45 degree line. Otherwise, w^* is flatter. For example, if all houses are essentially of the same quality, but there is some dispersion in wealth, then w^* must be close to a vertical line and is thus very steep.

To characterize equilibrium prices in closed form, we specialize further and assume separable log utility, that is, $u(c,h) = \log c + \theta \log h$. From the Euler equations, the marginal house price at quality h must be equal to the MRS between housing and wealth spent on other goods:

$$p'(h) = \theta \frac{w^*(h) - p(h)}{h}. \tag{7}$$

An agent with wealth $w^*(h)$ must be indifferent between buying a house of quality h and spending $w^*(h) - p(h)$ on other goods, or instead buying a slightly larger house and spending slightly less on other goods. An agent who already spends a lot on other goods is willing to pay more for a larger house (because of diminishing marginal utility of nonhousing consumption.) Therefore, if the house of quality h is assigned to an agent who spends more on nonhousing consumption per unit of house quality, then the slope of the house price function must be steeper at the point h.

To obtain closed form solutions for equilibrium prices, we further assume that the distributions G and F are such that the assignment function is a polynomial

$$w^*(h) = \sum_{i=1}^{n} a_i h^i. \tag{8}$$

The lowest-quality house must have a zero price since it is purchased by the buyer who has zero wealth. The unique solution to the ordinary differential equation (7) that satisfies $p(0) = 0$ is given by

$$p(h) = \int_0^h \left(\frac{\tilde{h}}{h}\right)^\theta \frac{\theta w^*\left(\tilde{h}\right)}{\tilde{h}}\, d\tilde{h} = \sum_{i=1}^n a_i \frac{\theta}{\theta + i} h^i. \tag{9}$$

If this solution is strictly increasing, then it is an equilibrium house price function. The price for a house of quality h is the weighted average of MRS for all agents who buy quality *less* than h, with the MRS evaluated at total wealth.

When is the price function linear?

In general, the equilibrium price (9) is a nonlinear function of quality. Higher powers of h matter for prices if they matter for the assignment (8). Linear pricing emerges in equilibrium, however, for particular pairs of distributions. Indeed, suppose the distributions F and G are scaled version of each other, that is, $F(w) = G(w/k)$. The assignment function is then $w^*(h) = kh$. Let $\bar{W} = \int w\, dF(w)$ and $\bar{H} = \int h\, dG(h) = \bar{W}/k$ denote average wealth and house quality, respectively. The price function can be written as

$$\overline{p} = \frac{\theta}{\theta + 1} \frac{\bar{W}}{\bar{H}}. \tag{10}$$

It depends on the distributions F and G only through their respective means. More generally, the segment-specific house price (9) typically depends on details of the distributions through the parameters of the assignment (8).

The emergence of linear pricing as a knife-edge case is not limited to log utility. Indeed, for a given utility function, cdf $F(w)$ for movers and average housing quality $\overline{H} = \int h\, dG(h)$, we can find a cdf $\widetilde{G}(h)$ with mean $\overline{H}$ such that (i) the assignment can be represented by an increasing function $w^*(h) = F^{-1}(\widetilde{G}(h))$ and (ii) the price function is linear $p(h) = \overline{p}h$, where for all h we have

$$\overline{p} = \frac{u_2(c,h)}{u_1(c,h)}. \tag{11}$$

When pricing is linear, the per-unit house price $\overline{p}$ enters the Euler equations of *all* households: The marginal (or, equivalently, the average, per unit) user cost can be read off the Euler equation of any household – in this sense *every* household is a marginal investor for *every* house $h \in [0,1]$. In contrast, in the nonlinear pricing case, the marginal user cost at quality h can be read off only *one* Euler equation (6), that of the marginal investor with wealth $w^*(h)$.

The linear special case is *imposed* in macroeconomic models with divisible housing capital. In these models, there is a production technology that converts different houses into each other which is linear; the marginal rate of transformation between different houses is thus set equal to one. Moreover, the cdf $\widetilde{G}(h)$ is assumed to adjust so as to ensure linear pricing. As a result, the per unit price of housing $\overline{p}$ changes if and only if the marginal rate of substitution of *all* investors changes. In our quantitative approach below, we do not take a stance on the production technology or the house quality distribution $G(h)$. Instead, we measure the cdf $G(h)$ directly from the data – in other words, we let the data tell us whether the pricing is linear or nonlinear.

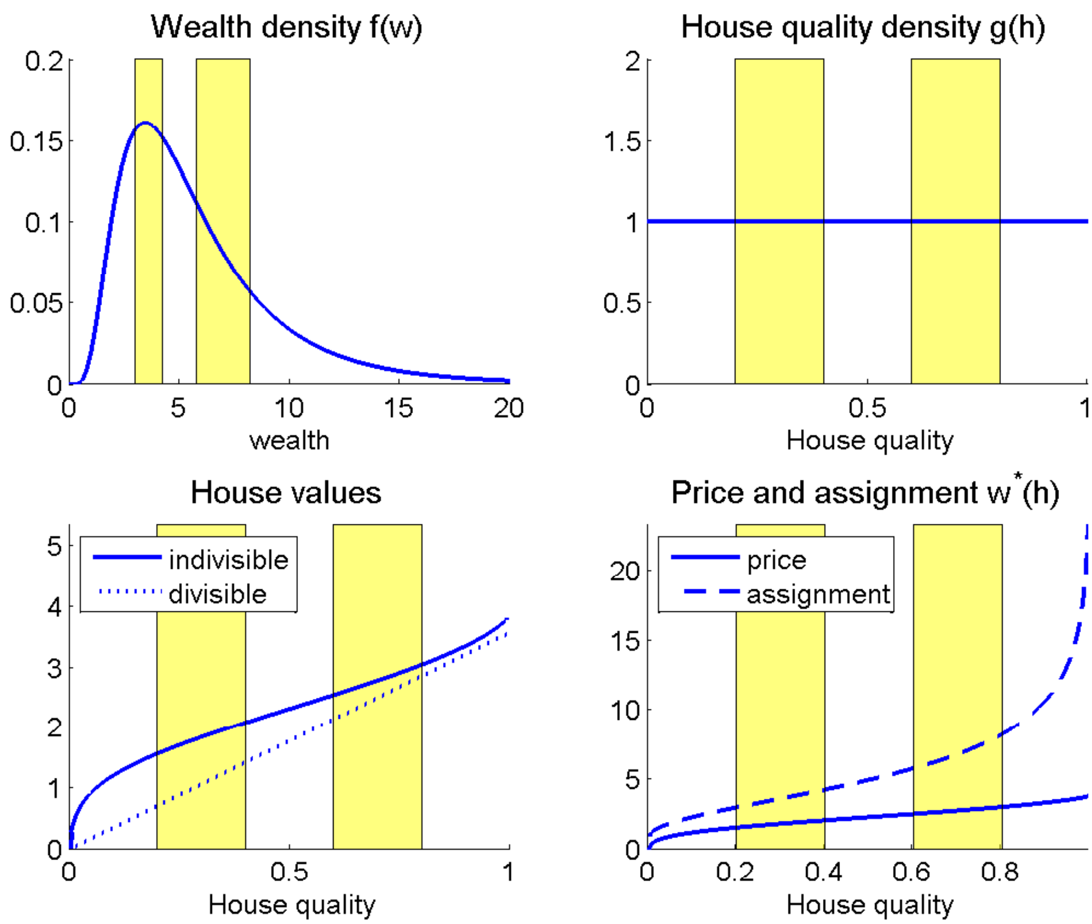

Figure 4: Equilibrium prices and assignment with lognormal wealth density $f(w)$ and uniform house quality density $g(h)$. Top left: wealth density $f(w)$. Top right: house quality density $g(h)$. Bottom left: the equilibrium house price function $p(h)$ in the indivisible model (solid line) and the divisible model (dotted line). Bottom right: assignment $w^*(h)$ and house prices. Shaded areas indicate the quintiles of the distribution.

Simple graphical example

Figure 4 compares equilibria of the indivisible and divisible model with separable log utility. The top left panel shows a lognormal wealth density $F'(w) = f(w)$. The top right panel shows a uniform house density $G'(h) = g(h)$. In both panels, the second and fourth quintile have been shaded for easier comparison. The bottom left panel shows house prices. The solid line is the house price function for the indivisible model. The dotted line is the price function for the divisible model. As shown above, it may equivalently be interpreted as the price function in an indivisible model with a uniform wealth distribution or a lognormal house quality distribution. Finally, the bottom right panel compares the price to the QQ plot $w^*(h)$, which provides the wealth level of an agent who buys a house of quality h.

The house price function of the indivisible model is nonlinear. This reflects differences in the shapes of F and G that lead to a nonlinear assignment w^*. The bottom right panel shows that the function $w^*(h) = F^{-1}(G(h))$ is relatively steep for low and high house qualities, and relatively flat in between. The shape of w^* is determined by the relative dispersion within quintiles in the top two panels. For the uniform distribution of house qualities, dispersion is the same within each quintile. In contrast, the dispersion of wealth is relatively high in the first and fifth quintile, but relatively low in the second and third quintile. To achieve an assignment with house quality increasing in wealth, wealth must thus rise more with house quality in the former quintiles than in the latter.

The house price function is determined by the segment-specific Euler equation (7). For a given quality h, the price function is steeper if the marginal investor spends more on nonhousing consumption per unit of house quality. In particular, for a given $p(h)$, the house price must rise more if the marginal investor is richer. The house price function thus inherits from the assignment w^* the property that it steepens for very high and very low qualities. There is less steepening at high qualities where wealth per unit of quality responds less to w^*. In addition, the house price function must be consistent with pricing by all marginal investors at qualities less than h. More dispersion in wealth relative to house quality in lower quintiles thus leads to higher prices in higher quintiles.

Comparative statics

Figure 5 shows what happens if the house quality distribution becomes fatter tailed; it has more mass at both the low and high end. The new density – the green line in the top right panel – comes from a beta distribution with mean one half. The mean house quality is thus unchanged from the uniform distribution. Therefore, the equilibrium price in the divisible model is the same as before. In contrast, prices change in the indivisible model to reflect the change in distribution. The green line in the bottom left panel is the new price function.

The bottom right panel shows capital gains by house value implied by the change in distributions: it plots the log house value in the "blue economy" (the log of the blue line in the bottom left panel) on the horizontal axis against the capital gain from blue to green (the log difference between the green and blue lines in the bottom left panel) on the vertical axis. The main result here is that capital gains are much higher at the low end than at the high end.

To understand the intuition, it is again helpful to consider the relative dispersion of wealth and quality within quintiles. In the bottom quintile, the dispersion of house qualities has now decreased, making wealth even more dispersed relative to quality. The assignment w^* must become steeper in this region as richer agents must buy lower quality houses. A steeper assignment in turn implies a steeper house price function. Starting from the smallest house, prices rise faster to keep richer marginal investors indifferent. For higher qualities, for example in the third quintile, the effect is reversed: as the house distribution is more dispersed than the wealth distribution, poorer marginal investors imply a flatter price function.

Figure 6 provides an example of a shock to a subpopulation. We assume that all agents with wealth less than 4 develop a higher taste for housing, as measured by the parameter θ. Over that range, we choose the increase in θ to be linearly declining in wealth, with a slope small enough such that the assignment is still monotonic in wealth. The wealth and quality distributions are the same as before. The bottom panels compare the price and capital gain effects. In the divisible model, the price rises to reflect higher demand for housing. In the indivisible model, the Euler equation predicts that the slope of the price function becomes steeper for low house qualities (where the poor households buy.)

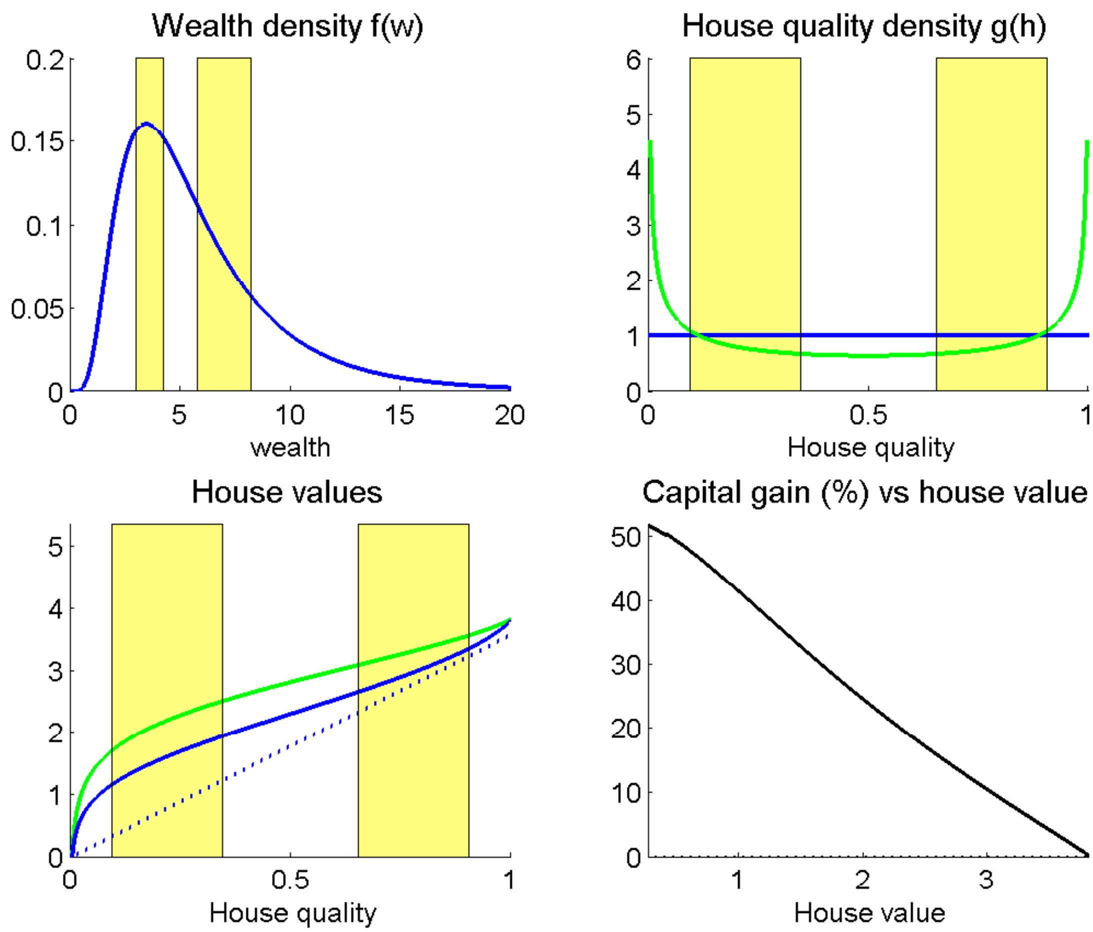

Figure 5: Changing the distribution of house qualities. Top left: wealth density $f(w)$. Top right: quality density $g(h)$ under uniform (blue) and beta (green). Bottom left: blue equilibrium price function $p(h)$ for uniform distribution, green function for beta quality density. Bottom right: capital gain from blue prices to green prices.

Not surprisingly, higher demand leads to higher prices in both the indivisible and divisible models. The divisible model predicts that capital gains are the same for all qualities. Interestingly, the indivisible model implies a higher capital gain at the low end, reflecting the higher demand of households who buy low quality houses. At the high end, the capital gain in the indivisible model is actually lower than under the divisible model.

3.2 Two period demand, one mover characteristic

In the last section, housing demand was derived from a static household problem. It was thus not necessary to distinguish between a house and the service flow from the house, and between house price and user cost (that is, the price of service flow), distinctions that are important when taking the model to the data. Moreover, there was no role for house price expectations. In an intertemporal problem, such as in our quantitative model below, the distinction between price and user cost is present and house price expectations matters. Here we illustrate these

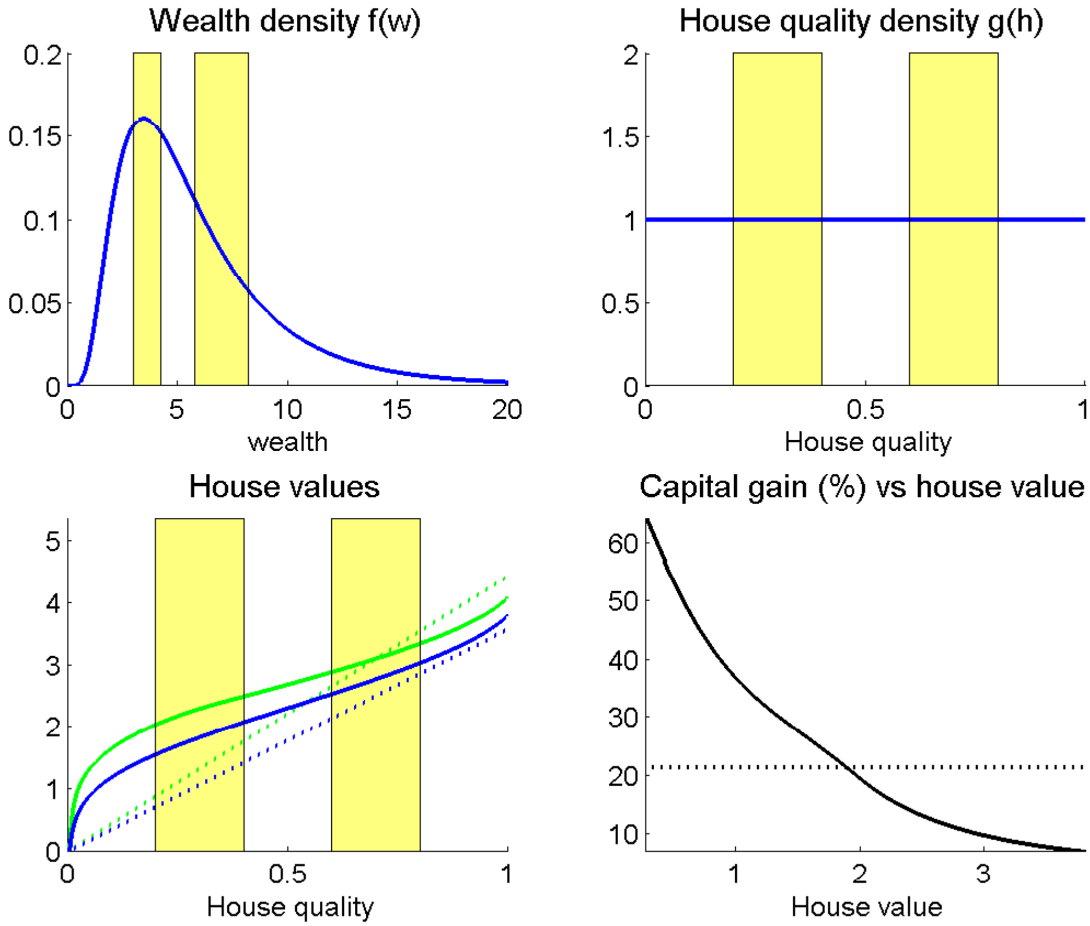

Figure 6: Changing poor households' preferences for housing.

issues using the simplest possible intertemporal household problem – two dates (0 and 1), no uncertainty and frictionless housing and credit markets.

A house of quality h is now an asset that generates service flow $s(h)$ at date 0, where s is strictly increasing with $s(0) = 0$. At date 0, households have income y and can borrow an amount b at the gross riskless interest rate R (where negative b corresponds to the purchase of bonds). Date 1 variables are indicated by a tilde. At date 1, households receive income $\tilde{y}$, sell their house at a price $\tilde{p}(h)$, and pay or receive funds in the credit market so they are left with cash $\tilde{w}$.

A household maximizes utility

$$u(c, s(h)) + v(\tilde{w})$$

subject to the budget constraints for dates 0 and 1

$$c + p(h) = y + b,$$
$$\widetilde{w} = \tilde{y} + \tilde{p}(h) - Rb.$$

Define lifetime wealth by $w = y + \tilde{y}/R$. The household problem can be rewritten as choosing c, h and $\widetilde{w}$ to maximize utility subject to the single lifetime budget constraint

$$c+\widetilde{w}/R+\underbrace{p\left(h\right)-\tilde{p}\left(h\right)/R}_{\text{user cost }\rho(h)}=w,$$

where the expression over the brace is the user cost $\rho\left(h\right)$ of a house of quality h: the cost of purchasing the house less its discounted resale value.

The first order conditions for the household problem include

$$\begin{aligned} u_1\left(c,s\left(h\right)\right) &= Rv'\left(\tilde{w}\right) \\ \rho'\left(h\right) &= \frac{u_2\left(c,s\left(h\right)\right)s'\left(h\right)}{u_1\left(c,s\left(h\right)\right)} \end{aligned}$$

The first condition is standard: it equates the intertemporal marginal rate of substitution to the interest rate. The second condition says that the intratemporal MRS between housing and numeraire consumption equals the *marginal* user cost of housing $\rho'\left(h\right)$ at the quality level h that the household chooses.

From the household problem, the house price only matters to the extent that it affects the user cost $\rho\left(h\right)$. For special cases, we can again obtain closed form solutions.[9] For the two-period model, the lower-left panels in Figures 4-6 represent the user cost function $\rho\left(h\right)$. If house prices $p\left(h\right)$ are proportional to user costs $\rho\left(h\right)$ (which is the case with constant capital gain expectations), the lower-left panels show house prices up to a factor, while the lower- right panels in Figures 5-6 show the actual capital gains.

The relationship between user cost and price depend on how expectations are formed. We consider two scenarios. The first assumes that agents expect all prices to grow at a common gross rate μ, starting from the current equilibrium price function. Agents thus extrapolate forward the behavior of relative prices from what they currently observe. Setting $\tilde{p}\left(h\right)=\mu p\left(h\right)$ in the definition of the user cost, it follows that the equilibrium price is given by

$$p\left(h\right)=\frac{\rho\left(h\right)}{1-\mu/R}$$

In a frictionless model with these *constant capital gain expectations*, house prices are proportional to user costs, and log price changes are equal to log changes in user costs. This scenario is useful for analyzing an economy in normal times when households perceive it to be in steady state. It is also useful when analyzing a boom in which households believe changes in the price pattern to be permanent.

[9] Suppose we have a linear services production function $s\left(h\right)=h$ and separable log utility, that is, $u\left(c,h\right)=\log c+\theta\log h$ and $v\left(\tilde{w}\right)=\beta\log\tilde{w}$. The marginal user cost at quality h is:

$$\rho'\left(h\right)=\frac{\theta}{1+\beta}\frac{w^*\left(h\right)-\rho\left(h\right)}{h}. \tag{12}$$

With a polynomial assignment function (8), we get

$$\rho\left(h\right)=\int_0^h\left(\frac{\tilde{h}}{h}\right)^{\theta/(1+\beta)}\frac{\theta}{1+\beta}\frac{w^*\left(\tilde{h}\right)}{\tilde{h}}\,d\tilde{h}=\sum_{i=1}^{n}a_i\frac{\theta}{\theta+\left(1+\beta\right)i}h^i. \tag{13}$$

The second scenario assumes an exogenous price expectation function $\tilde{p}(h)$. It is of interest when thinking about a boom in which households expect prices to mean revert to some earlier price pattern, given by $\tilde{p}$. With a linear services production function and log separable utility, the house price is then

$$p(h) = \tilde{p}(h)/R + \sum_{i=1}^{n} a_i \frac{\theta}{\theta + (1+\beta)i} h^i$$

Under this scenario with *mean-reverting capital gains expectations*, relative house prices still depend on the distributions F and G via relative user costs, but relative price expectations now also play a role. The relative importance of the distribution term for house prices is increasing in the subjective and market discount rates β^{-1} and R. Intuitively, the current assignment matters more for prices if the holding period for houses is longer.[10]

4 A Quantitative Model

For the stylized model in the previous section, housing demand was derived from a frictionless, deterministic, one/two-period optimization problem and households differed only in wealth. In this section, we describe a more general intertemporal problem for household savings and portfolio choice. This problem accommodates many features that have been found important in existing studies with micro data, and thus lends itself better to quantitative analysis. It differs from most existing models because there is a continuum of (indivisible) assets that agents can invest in.

The problem is solved for a distribution of households that differ in age, income, and cash on hand (that is, liquid resources). The distribution is chosen to capture the set of movers in San Diego County in a given year. An equilibrium is defined by equating the distribution of movers' housing demand derived from the dynamic problem to the distribution of transacted houses. Appendix D contains a detailed description of our computations.

4.1 Setup

Households live for at most T periods and die at random. Let D_t denote a death indicator that equals one if the household dies in period t or earlier. This indicator is independent over time but has an age-dependent probability. Preferences are defined over streams of housing services s and other (numeraire) consumption c during lifetime, as well as the amount of numeraire consumption w left as bequest in the period of death. Conditional on period τ, utility for an agent aged a_τ in period τ is

$$E_\tau \left[\sum_{t=\tau}^{\tau+T-a_\tau} \beta^t \left[(1-D_t)\ u(c_t, s_t(h_t)) + (D_t - D_{t-1})\ v(w_t) \right] \right] \tag{14}$$

[10] In the simple model considered here, changing the discount rate corresponds to making the holding period exogenously longer for all agents. In the more general model below, transaction costs induce agents to endogenously choose longer holding periods.

Households have access to two types of assets. First they can buy houses of different qualities $h \in [0, 1]$ that trade at prices $p_t(h)$. Owning a house is the only way to obtain housing services for consumption. A house of size h_t owned at the end of period t produces a period t service flow $s_t(h_t)$ where the function s_t is strictly increasing. It may depend on time to accommodate growth.

Household i borrows b_t at the gross interest rate R_t between period t and $t+1$. The amount b_t measures *net* borrowing.[11] A negative position b_t corresponds to bond purchases. We assume that a household can only borrow up to a fraction $1 - \delta$ of the value of his house. In other words, the amount b_t must satisfy

$$b_t \leq (1 - \delta) p_t(h_t). \tag{15}$$

The fraction δ is the downpayment requirement on a house.

We introduce three further features that distinguish housing from bonds. First, selling houses is costly: the seller pays a transaction cost ν that is proportional to the value the house. Second, every period an owner pays a maintenance cost ψ, also proportional to the value of the house. Finally, a household can be hit by a moving shock $m_t \in \{0, 1\}$, where $m_t = 1$ means that they must sell their current house. Formally, the moving shock may be thought of as a shock to the housing services production function $s_t(\cdot)$ that permanently leads to zero production unless a new house is bought. Of course, households may also choose to move when they do not receive a moving shock, for example because their income has increased sufficiently relative to the size of their current home.

Households receive stochastic income

$$y_t = f(a_t) \, y_t^p y_t^{tr} \tag{16}$$

every period, where $f(a_t)$ is a deterministic age profile, y_t^p is a permanent stochastic component, and y_t^{tr} is a transitory component.

Our approach to incorporating the tax system is simple. We assume that income is taxed at a rate τ. So the aftertax income $(1 - \tau) y_t$ enters cash on hand and the budget constraint. Mortgage interest can be deducted at the same rate τ. Interest on bond holdings is also taxed at rate τ. Therefore, the aftertax interest rate $(1 - \tau) R_t$ enters cash on hand and the budget constraint. We assume that housing capital gains are sheltered from tax.

To write the budget constraint, it is helpful to define cash on hand net of transaction costs. The cash w_t are the resources available if the household sells:

$$w_t = (1 - \tau) y_t + p_t(h_{t-1})(1 - \nu) - (1 - \tau) R_t b_{t-1} \tag{17}$$

The budget constraint is then

$$c_t + (1 + \psi) p_t(h_t) = w_t + 1_{[h_t = h_{t-1} \& m_t = 0]} \nu p_t(h_{t-1}) + b_t \tag{18}$$

[11] In this paper, we focus on the boom period, where the number of mortgage defaults was negligible. For an application to the bust period, it would be important to include these mortgage defaults explicitly in to the model (as in Chatterjee and Eyigungor 2009, Corbae and Quintin 2010, Campbell and Cocco 2011.)

Households can spend resources on numeraire consumption and houses, which also need to be maintained. If a household does not change houses, resources are larger than w_t since the households does not pay a transaction cost. The household can also borrow additional resources.

Consider a population of movers at date t. A mover comes into the period with cash w_t, including perhaps the proceeds from selling a previous home. Given his age a_t, current house prices p_t as well as stochastic processes for future income y_τ, future house prices p_τ, the interest rate R_τ, and the moving shock m_τ, the mover maximizes utility (14) subject to the budget and borrowing constraints. We assume that the only *individual-specific* variables needed to forecast the future are age and the permanent component of income y_t^p. The optimal housing demand at date t can then be written as $h_t^* \left(p_t; a_t, y_t^p, w_t\right)$.

As in the previous section, the distribution of available houses is summarized by a cdf $G_t(h)$. The distribution of movers is described by the joint distribution of the mover characteristics $\left(a_t, y_t^p, w_t\right)$. An equilibrium for date t is a price function p_t and an assignment of movers to houses such that households optimize and market clear, that is, for all h,

$$\Pr\left(h_t^* \left(p_t; a_t, y_t^p, w_t\right) \leq h\right) \leq G_t(h).$$

4.2 Numbers

We now explain how we quantify the model. In this section, we describe our benchmark specification. Section 5 discusses results based on several alternatives. It is helpful to group the model inputs into four categories

1. Preferences and Technology

   (Parameters fixed throughout all experiments.)

   (a) Felicity u, bequest function v, discount factor β
   (b) conditional distributions of death and moving shocks
   (c) conditional distribution of income
   (d) service flow function (relative to trend)
   (e) maintenance costs ψ, transaction costs ν

2. Distributions of house qualities and mover characteristics

3. Credit market conditions

   (a) current and expected future values for the interest rate R
   (b) downpayment constraint δ

4. House price expectations

Our goal is to compare different candidate explanations for house price changes during the boom. We thus implement the model for two different trading periods: once before the boom, in the year 2000, and then again at the peak of the boom, the year 2005.

Preferences and technology are held fixed across trading periods. For each period, we measure the distribution of house qualities and mover characteristics from the 2000 Census cross section (this implementation is labeled $t = 2000$) and then again for the peak of the boom, using the 2005 ACS ($t = 2005$).

We measure credit market conditions in 2000, and assume that households in 2000 were expecting constant capital gains; households were expecting the 2000 relative prices to remain unchanged, and absolute prices to grow at a constant rate with income. The service flow function is chosen to match 2000 house prices at these expectations.

After solving the model for the year 2000, we compute the model for the year 2005 under different scenarios for credit market conditions and house price expectations. We compare the predictions for 2005 equilibrium house prices with 2005 data. Below, we describe all elements in more detail.

Preferences

Felicity is given by power utility over a Cobb-Douglas aggregator of housing services and other consumption:

$$u(c, s) = \frac{\left[c^{\rho}\, s^{1-\rho}\right]^{1-\gamma}}{1-\gamma}, \tag{19}$$

where ρ is the weight on housing services consumption, and γ governs the willingness to substitute consumption *bundles* across both time period and states of the world. We work with a Cobb-Douglas aggregator of the two goods, with $\rho = .2$. If divisible housing services are sold in a perfect rental market, the expenditure share on housing services should be constant at 20%. This magnitude is consistent with evidence on the cross section of renters' expenditure shares (see for example, Piazzesi, Schneider, and Tuzel 2007.) We also assume $\gamma = 5$, which implies an elasticity of substitution for consumption bundles across periods and states of 1/5.

The period length for the household problem is three years. Households enter the economy at age 22 and live at most 23 periods until age 91. Survival probabilities are taken from the 2004 Life Table (U.S. population) published by the National Center of Health Statistics. To define utility from bequests, we compute the utility from receiving an L-year-annuity if a house of average quality can be rented at a rent to price ratio of 7%, the long run average in the US. We select $L = 7.5$ to match to match the age-profile of housing expenditure for households over 65.

The moving shocks are computed based on two sources. First, we compute the fraction of households who move by age, which is about a third per year on average. The fraction is higher for younger households. To obtain the fraction of movers who move for exogenous reasons, we use the 2002 American Housing Survey which asks households in San Diego about their reasons for moving. A third of households provides reasons that are exogenous to our model (e.g., disaster loss (fire, flood etc.), married, widower, divorced or separated.)

We assume that maintenance expenses cover the depreciation of the house. Based on evidence from the 2002 American Housing Survey, maintenance ψ is roughly 1% of the house value per year. The transaction costs ν are 6% of the value of the house, which corresponds to real estate fees in California.

Conditional distribution of income

We estimate the deterministic life-cycle component $f(a_t)$ in equation (16) from the income data by movers. The permanent component of income is a random walk with drift

$$y_t^p = y_{t-1}^p \exp(\mu + \eta_t), \tag{20}$$

where μ is a constant growth factor of 2% and η_t is iid normal with mean $-\sigma_\eta^2/2$. The transitory component y_t^{tr} of income is iid. The standard deviation of permanent shocks η_t is 11% and the standard deviation of the transitory component is 14% per year, consistent with estimates in Cocco, Gomes, and Maenhout (2005).

Distribution of mover characteristics

The problem of an individual household depends on characteristics (a_t, y_t^p, w_t). For age and income, we use age of the household head and income reported in the 2000 Census (for $t = 2000$) and 2005 ACS (for $t = 2005$). The Census data does not contain wealth information. Therefore, we impute wealth using data from the Survey of Consumer Finances. The appendix contains the details of this procedure.

Credit market conditions

The interest rate R_t is set to 3% in 2000, which we measure from three-year interest rate data on TIPS (since our model period is three years.) In 2005, the three-year TIPS rate fell to 1%, so we use that value. Households expect future interest rates to stay at 3%, their 2000 value. For the downpayment constraint δ, we use 20% to describe conditions before the housing boom in 2000, and 10% for the peak of the boom in 2005.

House price expectations

We specify house price expectations for the years 2000 and 2005 separately. For the base year 2000, we assume that households expect constant capital gains across houses. Specifically, households expect all house prices to grow at the same rate μ as labor income:

$$p_{t+1}(h) = p_t(h) \exp(\mu + u_{t+1}(h)). \tag{21}$$

The shock u captures idiosyncratic variation in the house price – it is realized only when the house is sold. We set its volatility equal to 9% per year over the three year model period. We have also investigated adding an aggregate risk component that shifts the price function for all houses. We have found that the results are not very sensitive to adding the modest amounts of aggregate risk that are commonly measured from regional house prices (e.g., Flavin and Yamashita 2002.) We thus omit aggregate regional risk, and model agents' views about the San Diego market as a whole only through different scenarios for the conditional mean (as described below).

For the peak of the housing boom in 2005, we assume that households expect house prices to revert back down to their 2000 levels $p_0(h)$, up to a common growth rate:

$$p_{t+1}(h) = p_0(h) \exp(\mu + u_{t+1}(h)). \tag{22}$$

In this assumption, house price expectations are exogenous. Since equilibrium prices for low quality houses are relatively high in the year 2005, households expect cheap houses to fall more in price. These expectations are thus consistent with the regression evidence in Table 1.

Service flow function

The housing services produced by a house of quality h grow at the same rate μ as income. Starting from an initial service flow function $s_t(h)$, households expect

$$s_{t+1}(h) = \exp(\mu)\, s_t(h). \tag{23}$$

As discussed above, the initial service flow function s_0 is backed out so that the model exactly fits the 2000 price distribution. Constant growth of service flow over time is consistent with evidence on improvements in the cross section of houses discussed in Appendix C.

5 Quantitative Results

In this section we compare pricing by quality segment in our base year 2000 and at the peak of the boom in 2005. We describe a number of different experiments to examine the role of distributional shifts (movers or house qualities) and cheaper credit. In our benchmark scenario for the peak, we use (i) different house quality and mover distributions than in the year 2000, in particular a house quality distribution with fatter tails, (ii) a combination of lower interest rates and lower downpayment constraints, and (iii) mean reversion in price expectations. Other experiments investigate the importance of each feature in isolation, as well as alternatives (e.g., expectations of higher future house prices.)

5.1 Prices and service flow in the base year 2000

The first step in our quantitative analysis considers the base year 2000. Here we take as given (i) the distributions of house quality and mover characteristics, (ii) credit conditions that were prevalent in the year 2000, and (ii) constant capital gains in price expectations. We then determine a service flow function $s_0(h)$ such that the model exactly fits the cross section of observed house prices in the year 2000.

Figure 7 shows the resulting service flow function as a function of house quality. It is strictly increasing, and it is also concave over almost all of the quality range (except for the lowest qualities.) For the base year 2000, the quality index on the x-axis is simply the current house price. It follows that the price of a house is convex in the amount of housing services it provides. Some intuition can be obtained from our analysis of the stylized model in Section 3. In that model, the price function is linear in the service flow if the distribution of service flow across houses is a scaled version of the wealth distribution. In contrast, if the dispersion of the wealth distribution relative to the service flow distribution is larger over a particular quantile range, say the top quintile, than elsewhere, then we would expect the price function to be steeper over that range.

This logic rationalizes the shape of the backed out service flow function. The wealth distribution is indeed more dispersed than the house price distribution. For example, the top

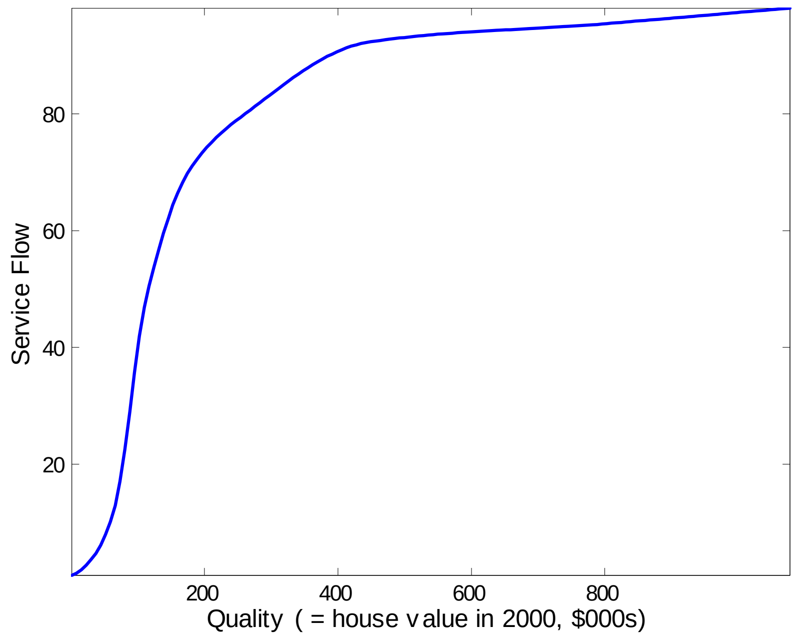

Figure 7: Service flow as a function of house quality (= 2000 house value)

10 percent of households own 50 percent of the total wealth but only 15 percent of the total housing wealth in their age group. If the service flow were, say, linear, then rich households would all try to buy the most expensive houses. For markets to clear, some rich households must be induced to choose cheaper houses. This requires a steep increase in the price per unit of service flow at higher house qualities.

5.2 Changes in house prices from 2000 to 2005

Figure 8 compares prices and capital gains relative to 2000 for our benchmark scenario. In both panels, the horizontal axis measures quality in terms of 2000 price. The left panel shows price as a function of quality for 2005. The green line is the price function in the data, constructed above in equation (3) from the response of prices to common shocks that affect houses of the same quality The blue line is the equilibrium price function from the model under the benchmark scenario. The dashed line indicates 2000 prices, it is linear with slope one since all prices are reported in 2005 Dollars. The right panel shows annualized capital gains between 2000 and 2005 for the data (the regression line from Figure 1) and the model – one fifth the overall log difference between the 2005 price function and the 2000 prices. The shaded areas indicate the quintiles of the 2005 house quality distribution.

The benchmark describes our preferred case for how house prices respond to changes in the distributions of movers and house qualities, as well as to the credit market conditions in the year 2005. The left panel of Figure 8 shows that equilibrium prices of lower quality houses increased more than the prices of high quality houses from the year 2000 to the year 2005, the peak of the boom. The right panel of Figure 8 shows that capital gains at the low end of the housing market are indeed higher than at the high end. Quantitatively, the benchmark case

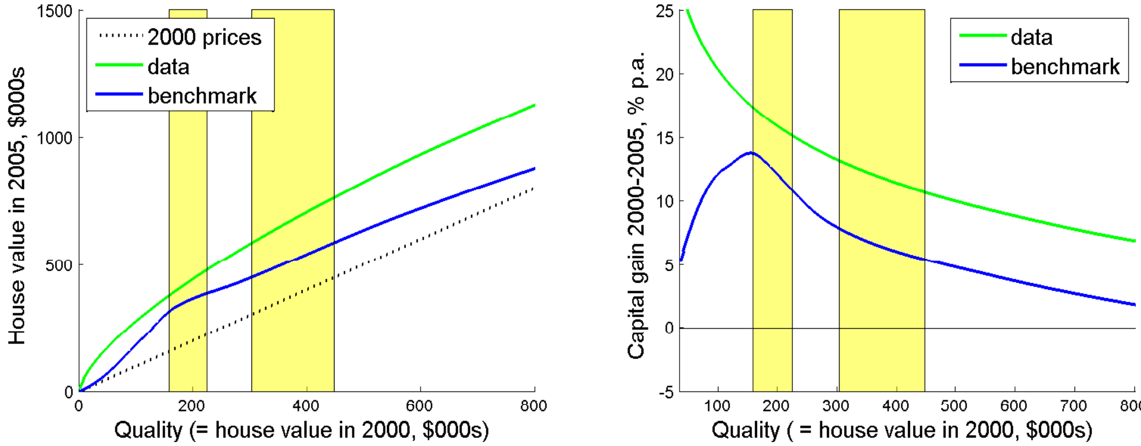

Figure 8: Benchmark results for the year 2005. The left panel shows house prices for 2005. The dashed 45 degree line are 2000 prices. The green line are prices in the data. The blue line are benchmark equlibrium prices. The right panel shows capital gains. The shaded areas indicate quintiles of the house quality distribution. The house quality on the horizontal axis is measured in 2005 Dollars.

matches more than half of the capital gains between 2000 and 2005. At the very low end of the quality spectrum, capital gains are hump-shaped. The peak of the hump is at roughly \$150K, which corresponds to the 20th percentile of the house quality distribution.

5.3 Assignment Properties in Years 2000 and 2005

To understand how the model works, we now consider the assignment of movers to houses. In the stylized model of Section 3, cash on hand was the only dimension of heterogeneity and the assignment could be represented as a line in the plane. Here, movers differ not only in cash (wealth plus income), but also in income and age. Age matters for housing demand if movers' savings and portfolio choice depends on their planning horizon. Income may matter (other than through its effect on cash) because it is used to forecast future income – it can be thought of as a proxy for human capital.

Figure 9 provides a first impression for how the three dimensions of heterogeneity affect the equilibrium assignment. In both panels, the horizontal axis measures house quality and the vertical axis measures cash. Every dot in the picture is a mover household in the 2005 ACS – larger dots indicate larger survey weights. The color of the dots illustrates a second dimension of heterogeneity. In the top panel, this dimension is age, in the bottom panel it is the ratio of income to cash. Both panels show that the model assigns richer households to better houses. However, cash is not the only dimension that matters in the assignment. In particular, for richer movers and higher house qualities, cash can be a bad predictor of house quality, since there are many households with the same cash level but very different houses, and vice versa.

In the top panel of Figure 9, blue dots correspond to younger households. For house quality levels below \$150K, the effect of age appears to be relatively small, and the assignment is driven to a greater extent by cash alone. In contrast, for house qualities greater than \$150K, the

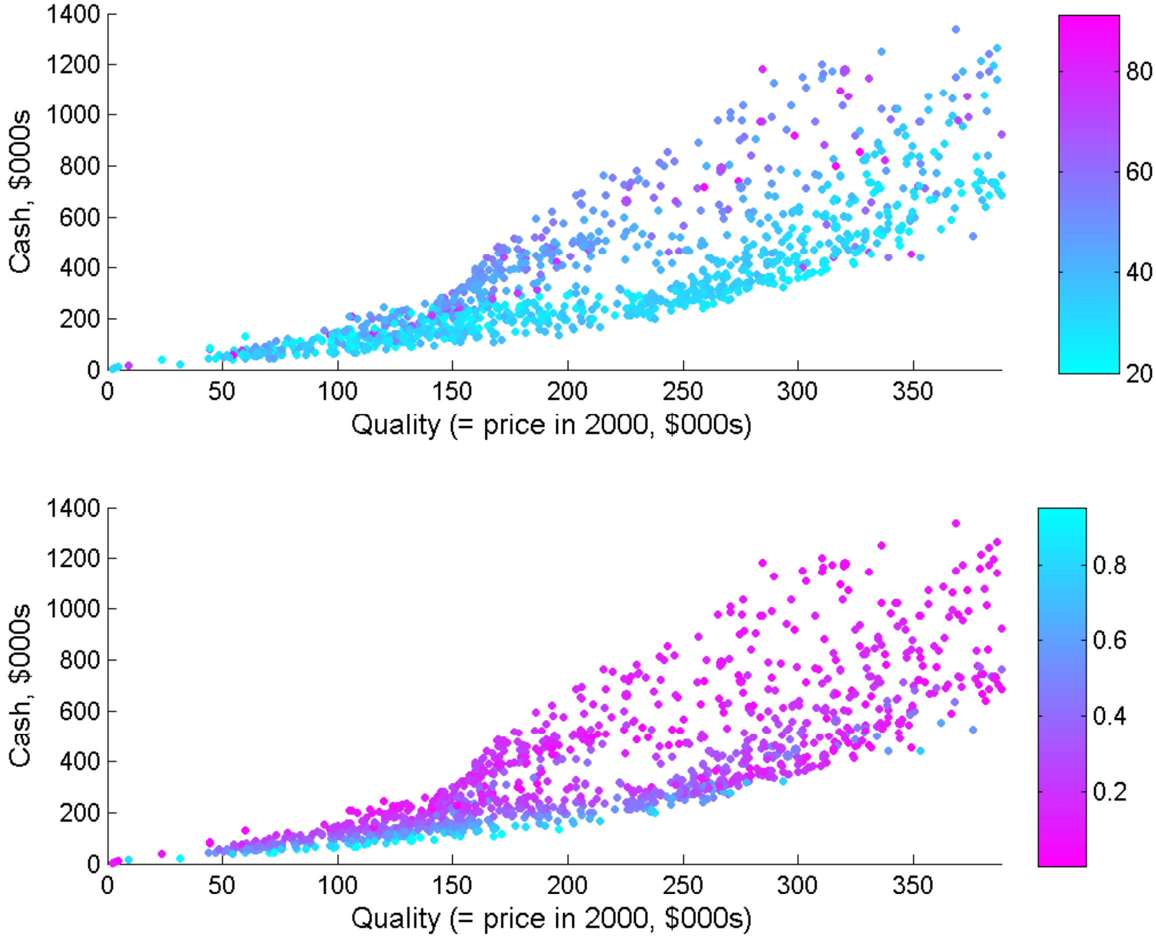

Figure 9: Equilibrium assignment. Both panels are scatter plots of 2005 San Diego ACS observations and their equilibrium assignments. The horizontal axis measures house quality. The vertical axis measures cash (wealth plus income.) The top panel shades dots according to age. The age coloring is indicated on the right bar. The bottom panel shades dots accoring to income-to-cash ratios. The coloring of income/cash is indicated on the right bar. The size of the dots correspond to their sampling weights.

cash/quality relationship is quite different by age. In particular, younger households increase house quality faster with cash than older households.

In the bottom panel of Figure 9, blue dots correspond to a higher ratio of income to cash. The blue dots at the bottom represent households who have virtually no funds other than their current income. Such households also increase quality faster with cash than do households with lower income to cash ratios. One reason is that age and the income to cash ratio are positively correlated; young households have higher income/cash ratios. However, the effect of income/cash ratios on the assignment is visible also at lower house qualities, such as below $150K.

We conclude that, in the model, the assignment at low qualities is driven more by income and cash, whereas at higher qualities age plays an important role.

Assignment by age and wealth – comparing model and data

We now investigate whether the data are consistent with the assignment patterns in the model. Table 3 compares the assignment of house quality to income and cash in the data and the model. Panel A contains results for the year 2000, while Panel B has those for the year 2005. The table has panels with four columns labeled I, II, III and IV that correspond to four house quality bins: houses worth less than $150K, between $150-200K, $200-400K, and above $400K in the base year 2000. The rows of the table report median income and cash together with the top and bottom 10th percentiles of the cash distribution of the indicated movers who buy houses in each of the four bins.

Figure 9 illustrated two predictions the model makes about the assignment. First, higher quality houses are bought by richer households. Table 3 shows that this is also a feature of the data, whether "richer" can be defined by either income or wealth. In the year 2000, the median income increases almost linearly across house bins, both in the data and the model. For households aged 35 years and younger, the average percentage increase from bin to bin is 32% both in the model and the data. For households older than 35 years, the average percentage increase is larger in both model and data (53% and 41%, respectively.) The same pattern emerges in the year 2005, where the average increase for young households is 18% in the model and 26% in the data and again larger for old households (48% and 38%, respectively.)

Across the first three bins, the median cash of households in 2000 also increases almost linearly across bins. The average increase for younger households is 65% in the model and 47% in the data, and for older households 69% in the model and 53% in the data. This increase is stronger for the top house quality which contains more wealth inequality, both in the model and the data. These increases are also larger in the year 2005. Among the 10% poorest households in each bin, the cash increases by 62% in the model and 44% in the data, with a stronger increase at the top quality level. The increase in cash across bins is also stronger among the richest 10% households, again both in the model and the data.

A second prediction of the model is that old households who buy in a certain housing segment are wealthier than the young households who buy in the same segment (as we could see in the top panel of Figure 9.) Table 3 shows that this is also true in the data. In 2000 and 2005, households older than 35 years have on average 68% more cash than younger households who bought a house in the same bin. The model captures most of this difference in 2000 and close to all of it in 2005. Moreover, the cash differences between young and old households are larger at higher house qualities. For example, the data show that old households have 60K more in cash than the young households in the lowest bin and almost 500K more in the highest bin. The model matches 89% of this difference.

Table 3 also reports standard errors for the medians in the data. The model is within two standard error bounds for five medians in Panel A and ten medians in Panel B out of the sixteen total. Hence, the model performs reasonably well in matching the cross-sectional assignment by levels of income and wealth.[12]

[12] This test ignores sampling uncertainty and is therefore tough on the model. A more accurate treatment of uncertainty based on Census replication weight is computationally very costly in our context and therefore omitted. Indeed, suppose we start with the set of replication weights provided by the Census. This would increase the errors around the data medians. Moreover, since the model computation takes the mover characteristics as an input, each set of weights would lead to a different mover distribution and hence a different set of model

Table 3: Assignment Of House Qualities in Data and Model

	Data				Model			
House quality bins	I	II	III	IV	I	II	III	IV
Panel A: Year 2000								
Median Income (in thousands)								
aged $\leq$ 35 years	48.5	68.7	88.1	128.5	43.5	72.7	102.6	115.0
	(1.8)	(2.5)	(4.0)	(7.1)				
aged $>$ 35 years	44.0	63.4	90.3	152.3	30.1	56.4	102.4	146.6
	(2.2)	(2.3)	(4.5)	(8.5)				
Median Cash (wealth plus income, in thousands)								
aged $\leq$ 35 years	112.0	169.7	284.6	646.8	98.8	179.3	361.2	1,455.7
	(5.4)	(7.8)	(15.3)	(39.7)				
aged $>$ 35 years	172.3	284.8	496.4	1,141.8	112.3	322.1	443.1	1,860.5
	(9.2)	(15.4)	(27.6)	(71.5)				
Percentiles of the Cash Distribution (in thousands)								
bottom 10%	60.2	94.9	146.0	310.9	59.7	128.7	204.3	894.7
top 10%	471.1	721.8	1,443.9	3,941.9	191.0	734.0	925.1	4,474.7
Panel B: Year 2005								
Median Income (in thousands)								
aged $\leq$ 35 years	61.3	73.6	101.9	132.5	53.4	86.6	100.9	93.8
	(4.0)	(3.5)	(8.1)	(7.1)				
aged $>$ 35 years	45.9	78.5	91.7	144.7	52.4	86.6	91.7	224.4
	(4.3)	(4.9)	(5.7)	(12.5)				
Median Cash (wealth plus income, in thousands)								
aged $\leq$ 35 years	203.4	258.5	421.3	735.2	125.5	229.7	491.5	1,827.4
	(18.0)	(22.5)	(43.8)	(76.4)				
aged $>$ 35 years	251.6	459.7	712.2	1,645.6	207.0	450.3	775.0	2,427.8
	(29.5)	(45.2)	(70.9)	(187.7)				
Percentiles of the Cash Distribution (in thousands)								
bottom 10%	79.1	125.3	180.4	391.3	68.2	199.0	310.2	1,106.0
top 10%	714.4	1,099.0	1,733.7	5,582.6	324.8	625.7	1,162.6	6,034.5

predictions. We could thus obtain standard errors around the model medians as well.

Note: This table reports moments of the assignment of house quality to income and cash (wealth plus current income) in the data and in the benchmark model. Panel A contains results for the year 2000, while Panel B reports results for the year 2005. Across both panels, house quality is measured in four bins. Bin I contains houses worth less than $150K in the 2000 base year. Bin II contains houses worth between $150-200K. Bin III contains houses worth $200-400K. Bin IV contains houses above $400K. All dollar amounts are reported in 2005 Dollars. The medians are computed for the households who bought a house in the indicated bin. The row "aged $\leq$ 35 years" ("aged $>$ 35 years") reports the medians of households aged 35 years or younger (above 35 years). The row "bottom 10%" ("top 10%") reports the 10th (90th) percentile of the cash distribution. The left columns show the assignment in the data computed from the 2000 Census and the 2005 American Community Survey. The right columns show the assignment in our benchmark model.

Table 4: Housing Wealth relative to Cash (Wealth Plus Income)

	Age			
	below 35	35-50 years	50-65 years	above 65
Panel A: Year 2000				
Data	0.632	0.459	0.369	0.317
Model	0.641	0.474	0.359	0.286
Panel B: Year 2005				
Data	0.968	0.677	0.317	0.387
Model				
benchmark	0.764	0.492	0.312	0.391
2000 credit conditions	0.500	0.322	0.219	0.229
2000 interest rate	0.643	0.426	0.256	0.316
2000 quality distribution	0.705	0.493	0.328	0.355
2005 interest rate permanent	1.044	0.643	0.411	0.458
$\delta = 5\%$	0.917	0.635	0.396	0.530

Policies by age across years

The policy functions in our model explain the assignment patterns in Figure 9. In the model as well as in the data, young households have more human wealth relative to other wealth. Intuitively, human wealth is relatively safe, so that younger households can afford to take more risks in their portfolio than older households. To do that, younger households increase their leverage, which involves buying a house that is large relative to wealth. Younger households thus choose a larger portfolio weight on housing than older households.

Table 4 reports ratios of housing wealth relative to cash (wealth plus income) by age. In the year 2000, younger households had ratios of housing wealth to cash around two thirds, while older household have ratios around one third. In the year 2005, the strong pattern by age is still present. The ratios of young households are now larger by roughly another third, and are close to 100 percent. The ratios of most older cohorts also increase but by a smaller amount. In our benchmark parametrization, the portfolio weights on housing for all cohorts increase. The benchmark ratios for young households are lower than what we see in the data.

The other rows report housing wealth/cash ratios for alternative model specifications. These specifications make benchmark assumptions, except for the feature indicated in the first column of Table 4. For example, the case with "2000 credit conditions" computes the model under benchmark assumptions, except for assuming the credit conditions prevalent in the year 2000.

A similar intuition holds for households with high income to cash ratios. Such households are rich in human capital, which is relatively safe. They increase the risk in their portfolios by taking on leveraged positions. In a world with collateral constraints, these leveraged positions involve house purchases. Households with high income/cash ratios therefore buy houses that are also bought by rich households with low income/cash ratios.

5.4 The role of cheap credit

Figure 10 investigates the role of cheap credit. The top left panel computes the 2005 equilibrium prices under benchmark assumptions, except that we specify 2000 credit conditions. This experiment shows that changes in credit conditions are critical for house prices. It also illustrates that changes in distributions alone induce relatively higher capital gains for lower housing segments.

The top right panel in Figure 10 computes 2005 equilibrium prices except that we assume a high 2000 value for the interest rate. This experiment isolates the importance of lower downpayment constraints. The dashed line shows that downpayment constraints are responsible for most of the quantitative effect of a change in credit market conditions on house prices between 2000 and 2005. Lower downpayment constraints increase the housing demand of younger households. In Table 4, moving from 2000 credit conditions to a lower downpayment constraint (the row labeled "2000 interest rate") increases the housing weight for the youngest cohort from 50% to 64%. It has smaller effects for the other cohorts. As a consequence, the house prices of low quality houses increase. Intuitively, the relaxation of borrowing constraints acts like the increase in housing demand by poor households in the simple model illustrated in Figure 6.

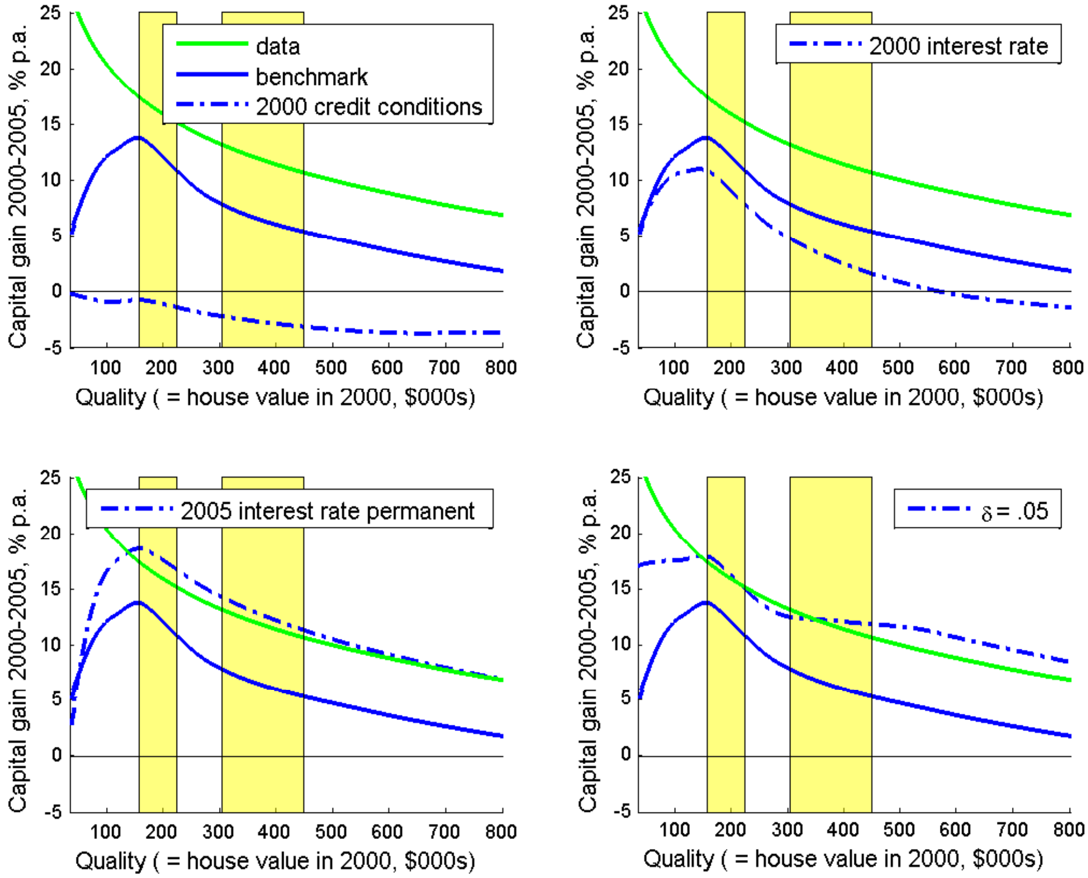

Figure 10: Equilibrium capital gains 2000-2005

In the top right panel, moving from the dashed line to the benchmark line adds the effect of low interest rates. This effect is quantitatively smaller than that of lower downpayment constraints. Moreover, it does not introduce additional tilt towards higher capital gains at the low end of the quality spectrum. If anything, capital gains at the high end rise more. Intuitively, lower interest rates thus affect housing demand in the cross section more evenly than a change in downpayment constraints.

The bottom left panel in Figure 10 computes 2005 equilibrium prices under benchmark assumptions, except that households expect interest rates to stay low (at the 1 percent 2005 level) permanently. In contrast, the benchmark assumes that households expect interest rates to stay at 2005 levels in the short run (for one model period, which represents three years), and then to revert back to 3 percent, their 2000 level. The dashed line in the bottom left panel shows that when households expect low interest rates permanently, the value of future service flows from a house are discounted less, and so current house prices are higher than in the benchmark case. The dashed capital gains from 2000 to 2005 are close to matching the data. Again, a shift towards lower interest rates affects houses of different qualities in a very similar fashion.

The bottom right panel in Figure 10 computes 2005 equilibrium prices again under benchmark assumptions, except with lower downpayment constraints. The benchmark case assumes 10 percent downpayment constraints for the peak of the boom, while the dashed line assumes 5 percent. The dotted line is close to matching the data. The dashed line also improves the fit of the model at the very low end of the market: capital gains in the low quality range are less hump-shaped than in the benchmark case. Once more, a change in downpayment constraints has a disproportionate effect on the low end of the house quality spectrum. Table 4 shows that with the lower downpayment constraints, the 2005 portfolio weights are also closer to matching the data.

5.5 The role of the house quality distribution

The left panel in Figure 11 computes 2005 equilibrium prices under benchmark assumptions, except that we use the 2000 house quality distribution. This experiment illustrates that changes in the quality distribution matter for the relative pricing of houses across qualities. In particular, the dashed line in the left panel of Figure 11 has higher capital gains at the high end of the quality spectrum than the benchmark line. This difference is due to the fatter tails of the 2005 house quality distribution in Figure 2. Richer households move into low quality homes and push up their prices relative to high quality homes. The intuition for this change of the house quality distribution is thus the same as in Figure 5 of the simple model.

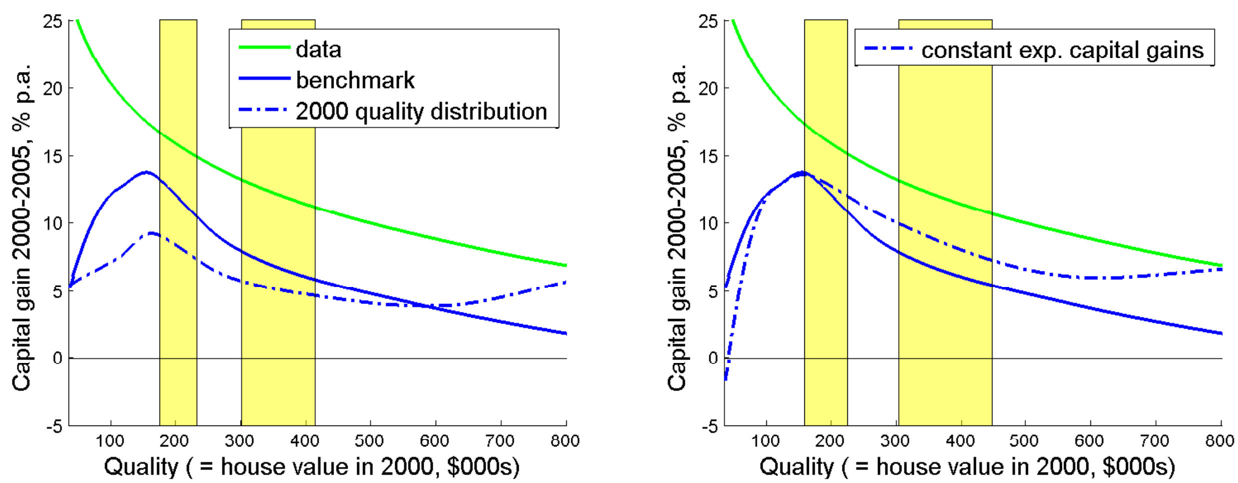

Figure 11: The green/light gray lines represent data on capital gains. The blue/dark gray lines represent equilibrium capital gains under benchmark assumptions.

5.6 Other experiments: expectations and service flow changes

We now check the sensitivity of our results with respect to our assumption on house price expectations and the service flow function. First, the benchmark assumes that households expect house prices to mean revert to their 2000 levels, consistent with the empirical evidence in Table 1. This assumption seems also consistent with evidence from the Michigan survey which indicates that a large majority of households thought that buying a house was not a good deal in the year 2005 (Piazzesi and Schneider 2009.)

To study the importance of this assumption, we recompute the model under benchmark assumptions, except that households expect house prices to stay high at their 2005 level (up to a growth rate.) This expectations scenario corresponds to equation (21), where again expectations are endogenous; they depend on equilibrium 2005 levels. The dashed line in the right-hand panel of Figure 11 represents equilibrium prices with high house price expectations. These high expectations lead to similar equilibrium capital gains for low end houses. But the more optimistic households drive up capital gains for high end houses relative to the benchmark case, which assumes mean-reverting house price expectations.

Second, the benchmark assumes that the service flow function remains the same over time, up to a growth factor. To check the importance of credit conditions under alternative assumptions on this function, we compute a new service flow function that exactly matches 2005 observed house prices. This computation also uses (*i*) 2005 distributions for house qualities and mover characteristics, (*ii*) 2005 credit conditions, and (*iii*) constant capital gain expectations. The left panel of Figure 12 compares the benchmark (blue) and the service flow function that exactly matches 2005 house prices (green.) The new service flow function grows faster for low quality houses than for high quality houses, which helps match the 2005 prices. The right panel of Figure 12 shows the resulting capital gains. By definition, the model-implied capital gains are identical to those in the data.

To again isolate the importance of credit conditions, we also recompute the model with the new service flow function but under 2000 credit conditions. The result is the dashed line in the right panel of Figure 12. It can be compared to the top left panel of Figure 10 which considers 2000 credit conditions when service flow grows at a constant rate for all houses. Figure 12 shows how differential growth in service flow contributes to differences in capital gains at low and medium qualities. However, the overall effect in the absence of cheap credit remains small.

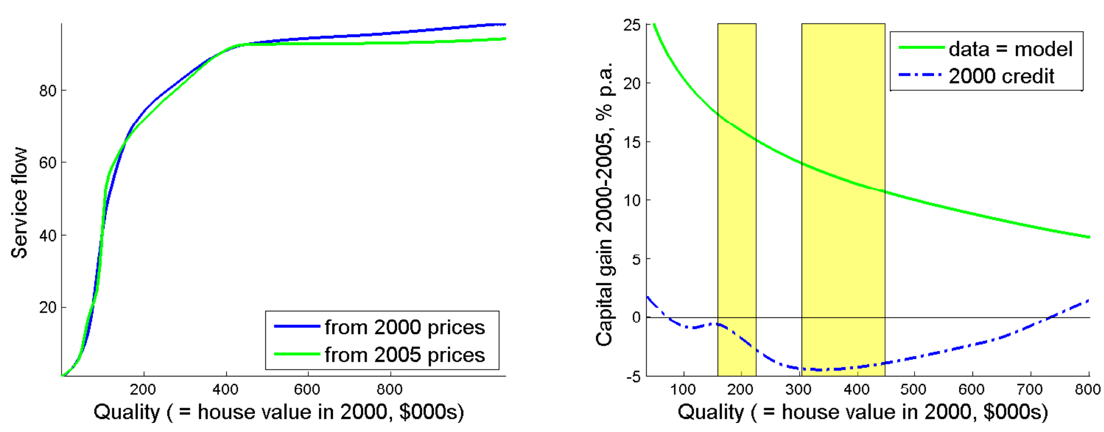

Figure 12: The left panel plots two service flow functions. The blue line matches 2000 house prices as in the benchmark. The green line is a new service flow function that matches 2005 house prices. The right panel shows capital gains under the new service flow function, which are identical to the data. The dashed line computes equilibrium prices with the new service flow function and 2000 credit conditions.

References

Bayer, Patrick, Fernando Ferreira, and Robert McMillan 2007. "A Unified Framework for Measuring Preferences for Schools and Neighborhoods." *Journal of Political Economy* 115, 588-638.

Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo 2011. "Understanding Booms and Busts in Housing Markets." *NBER Working Paper* 16734.

Campbell, John and Joao Cocco 2010. "A Model of Mortgage Default." Working Paper, Harvard.

Case, Karl E. and Christopher J. Mayer 1996. "House Price Dynamics within a Metropolitan Area." *Regional Science and Urban Economics* 26, 387-407.

Case, Karl E. and Robert J. Shiller 2005. "Is there a Bubble in the Housing Market?" *Brookings Papers on Economic Activity* 2.

Caplin, Andrew and John Leahy 2010. "A Graph Theoretic Approach to Markets for Indivisible Goods." Working Paper, NYU.

Chatterjee, Satyajit and Burcu Eyigungor 2009. "Foreclosures and House Price Dynamics: A Quantitative Analysis of the Mortgage Crisis and the Foreclosure Prevention Policy." Working Paper, Philadelphia Fed.

Chien, YiLi and Hanno Lustig 2010. "The Market Price of Aggregate Risk and the Wealth Distribution." *Review of Financial Studies* 23, 1596-1650.

Cocco, Joao F., Francisco J. Gomes, and Pascal J. Maenhout 2005. "Consumption and Portfolio Choice over the Life Cycle." *Review of Financial Studies* 18, 491-533.

Corbae, Dean and Erwan Quintin 2010. "Mortgage Innovation and the Foreclosure Boom." Working Paper, University of Wisconsin.

Favilukis, Jack, Sydney Ludvigson, and Stijn van Nieuwerburgh 2010. "The Macroeconomic Effects of Housing Wealth, Housing Finance, and Limited Risk-Sharing in General Equilibrium." Working Paper, NYU.

Flavin, Marjorie and Takashi Yamashita 2002. "Owner-Occupied Housing and the Composition of the Household Portfolio." *American Economic Review* 92(1), 345-362.

Glaeser, Edward L., Joshua Gottleb, and Joseph Gyourko 2010. "Can Cheap Credit Explain the Housing Boom?" Working Paper, Harvard.

Glaeser, Edward L., Joseph Gyourko, and Raven E. Saks 2005. "Why Have Housing Prices Gone Up?" *American Economic Review* 95 (2), 329-333.

Grandmont, Jean-Michel 1977. "Temporary General Equilibrium." *Econometrica* 45, 535-572.

Guerrieri, Veronica, Daniel Hartley, and Erik Hurst 2010. "Endogenous Gentrification and Housing Price Dynamics." Working Paper, Chicago Booth.

Himmelberg, Charles, Christopher Mayer, and Todd Sinai 2005. "Assessing High House Prices: Bubbles, Fundamentals and Misperceptions." *Journal of Economic Perspectives* 19, 67-92.

Jovanovic, Boyan 2009. "The Technology Cycle and Inequality." *Review of Economic Studies* 76, 707-729.

Judd, Kenneth L. 1998. *Numerical Methods in Economics.* MIT Press, Cambridge.

Kiyotaki, Nobuhiro, Alexander Michaelides, and Kalin Nikolov 2010. "Winners and Losers in Housing Markets." Working Paper, Princeton.

Kiyotaki, Nobuhiro and John Moore 1995. "Credit Cycles." *Journal of Political Economy* 105, 211-248.

Kaneko, Mamoru 1982. "The Central Assignment Game and the Assignment of Markets". *Journal of Mathematical Economics* 10, 205-232.

Lamont, Owen and Jeremy Stein 1999. "Leverage and House-Price Dynamics in U.S. Cities." *Rand Journal of Economics*, 30(3), 498-514.

Mian, Atif and Amir Sufi 2009. "The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis." *Quarterly Journal of Economics* 124(4), 1449-1496.

Mian, Atif and Amir Sufi 2010. "House Prices, Home Equity-Based Borrowing, and the U.S. Household Leverage Crisis." Forthcoming, *American Economic Review.*

Ortalo-Magne, Francois and Sven Rady 2006. "Housing Market Dynamics: On the Contribution of Income Shocks and Credit Constraints." *Review of Economic Studies* 73, 459–485.

Piazzesi, Monika, Martin Schneider, and Selale Tuzel 2007. "Housing, Consumption and Asset Pricing." *Journal of Financial Economics* 83, 531-569.

Piazzesi, Monika and Martin Schneider 2009. "Momentum Traders in the Housing Market: Survey Evidence and a Search Model." *American Economic Review* 99(2), 406-411.

Poterba, Jim 1991. "House Price Dynamics: The Role of Tax Policy and Demography." *Brookings Papers on Economic Activity* 2, 143-202

Rios-Rull, Jose-Victor and Virginia Sanchez-Marcos 2008. "An Aggregate Economy with Different House Sizes." *Journal of The European Economic Association* 6, 705-14.

Sattinger, Michael 1993. "Assignment Models of the Distribution of Earnings." *Journal of Economic Literature* 31(2), 831-880.

Stein, Jeremy 1995. "Prices and Trading Volume in the Housing Market: A Model with Down-Payment Effects." *Quarterly Journal of Economics* 110(2), 379-406.

Appendix

A San Diego County Transactions Data

In this appendix we describe our selection of sales and repeat sales. We begin by describing our sample of sales which not only forms the basis for selecting repeat sales but is also used to illustrate the shift in distributions in Section 2. Our goal is to compile a dataset of households' market purchases of single-family dwellings. We start from a record of all deeds in San Diego County, 1999-2008 and then screen out deeds according to three criteria.

First, we look at qualitative information in the deed record on what the deed is used for. We drop deed types that are not typically used in arms length transfers of homes to households in California. In particular, we keep only grant deeds, condo deeds, corporate deeds and individual deeds. The most important types eliminated are intrafamily deeds and deeds used in foreclosures. Even for the types of deeds we keep, the deed record sometimes indicates that the transaction is not "arms length" or that the sale is only for a share of a house – we drop those cases as well.

Second, we drop some deeds based on characteristics of the house or the buyer. We use only deeds for which a geocode allows us to precisely identify latitude and longitude. We eliminate deeds that transfer multiple parcels (as identified by APN number.) Information about property use allow us to eliminate second homes and trailers. To further zero in on household buyers, we eliminate deeds where the buyer is not a couple or a single person (thus dropping transaction where the buyer is a corporation, a trust or the beneficiary of a trust.)

Third, we drop some deeds based on the recorded price or transaction dates. We drop deeds with prices below $15,000 or with loan-to-value ratios (first plus second mortgage) above 120%. We also consolidate deeds that have the same sales price for the same contract date. We drop deeds that have the same contract date but different prices.

Our repeat sales sample is used to estimate our statistical model of price changes in Section 2.2. A repeat sale is a pair of consecutive sales of the same property within the above sales sample. Since we are interested in long term price changes, we want to avoid undue influence of house flipping on our estimates. We thus drop all pairs of sales that are less than 180 days apart. To guard against outliers, we drop repeat sales with annualized capital gains or losses above 50%.

B Robustness checks

Table B.1 reports additional results for the repeat sales model that incorporate zip code and census tract level information. Regression (i) reports the basic regression of capital gains on the own initial 2000 price (in logs) p_t^i from Figure 1. The regression has a slope coefficient of -0.060 with a standard error 0.0014 and an R^2 of 57.1%. Regression (ii) adds the initial zip code median (again, in logs) p_t^{zip} as regressor. The point estimate of the coefficient on the initial own price is basically unchanged (-0.057 versus -0.060); the difference is not statistically significant. The estimated coefficient on the zip-code median is statistically significant, but

−0.011 is economically small. The added explanatory power of the zip code median is tiny, the R^2 goes from 57.1% to 57.5%. The regression (*iii*) on the zip-code median alone (*iii*) gives an R^2 of 20.6%. Regressions (*iv*) and (*v*) are analogous to (*ii*) and (*iii*), but they use census tract rather than zipcode as the geographical area. The results are quite similar. Regression (*v*) uses only the census-tract medians with an R^2 of 28.5%.

Table B.1 Geographic Patterns in Repeat Sales Model

		const.	p^i_{2000}	p^{zip}_{2000}	p^{census}_{2000}	R^2
(*i*)	$p^i_{2005} - p^i_{2000}$	0.899	−0.060			0.571
		(0.018)	(0.002)			
(*ii*)	$p^i_{2005} - p^i_{2000}$	1.000	−0.057	−0.011		0.575
		(0.035)	(0.002)	(0.003)		
(*iii*)	$p^i_{2005} - p^i_{2000}$	1.016		−0.069		0.206
		(0.048)		(0.004)		
(*iv*)	$p^i_{2005} - p^i_{2000}$	0.879	−0.062		0.004	0.572
		(0.024)	(0.002)		(0.003)	
(*v*)	$p^i_{2005} - p^i_{2000}$	0.837			−0.056	0.285
		(0.031)			(0.003)	
(*vi*)	$p^{zip}_{2005} - p^{zip}_{2000}$	1.014		−0.070		0.672
		(0.068)		(0.006)		
(*vii*)	$p^{census}_{2005} - p^{census}_{2000}$	1.038			−0.071	0.606
		(0.034)			(0.003)	

Note: This table reports results from regressions of the capital gain from 2000 to 2005 in the price series indicated on the left-hand side on the regressors indicated on the headers of the columns. These cross sectional regressions involve individual house prices p^i_t from houses that were repeat sales in the two years 2000 and 2005, zip code medians p^{zip}_t, and census tract medians p^{census}_t in San Diego County.

Regression (*vi*) runs the capital gains in the zip-code medians on the initial zip-code medians. The results are comparable to row (*i*) with a slope coefficient of −0.070, a somewhat higher standard error of 0.0055, and a higher R^2 of 67.2%. Regression (*vii*) does the same exercise for census tracts, again with similar slope as row (*i*) .

Figure 13 plots the repeat sales observations from Figure 1 with the 2005 log price on the y-axis. The black line is the predicted value from a linear regression of 2005 log prices on 2000 log prices. The green line is the predicted value from a nonparametric regression, using a Nadaraya-Watson estimator with a Gaussian kernel and a bandwidth of 0.15. The nonparametric regression line is strictly increasing in the initial 2000 price. This monotonicity

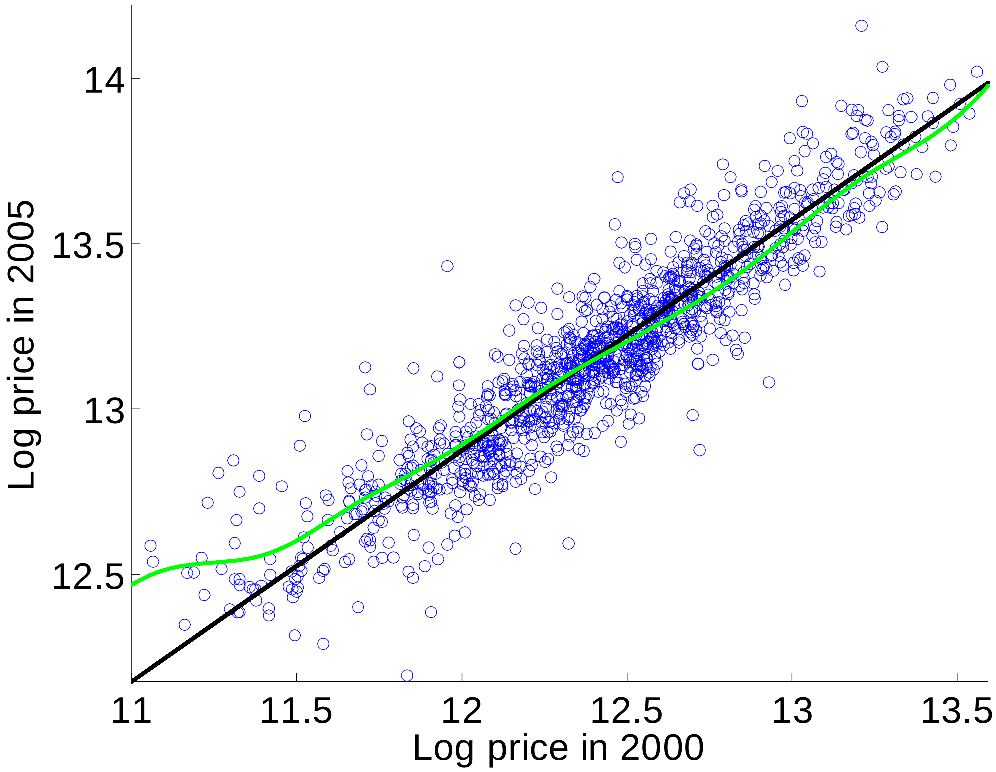

Figure 13: Repeat sales in San Diego County, CA, during the years 2000-2005.

property implies that the relative ranking of houses by quality according to the nonparametric regression is the same as the relative ranking according to the linear regression.

The nonparametric regression line is close to linear for a large range of house values, with the largest deviation at the low end. This deviation does not matter for our approach, because we use the pricing model only to derive an ordinal index. The absolute amount of service flow due to a house of a certain ordinal quality is backed out using the structural model. Section 5.1 uses 2000 house prices to back out a service flow function for that year and assumes a constant rate between 2000 and 2005 to derive the 2005 service flow function. Section 5.6 uses 2000 house prices and 2005 house prices to back out service flow functions for the two years, respectively.

C Data details

This appendix provides details on the calculations of home improvements, house quality and wealth reported in the text.

Improvements

The 2002 American Housing Survey contains data on home improvements in San Diego County. Table C.1 shows the means and medians of annual improvement expenses in San Diego as a percent of house values. We find that San Diego homeowners spend an amount equal to roughly 1% of their house value on improvements each year. The mean percentage spent on improvements is 2% for homes in the lowest bin, worth less than $50,000. However, this higher mean is estimated imprecisely. A test that the mean improvement percentage in the lowest bin and homes in the next bin (worth between $50,000 and $100,000) are identical cannot be rejected at the 10% level. A joint test whether mean improvements across all bins are equal can also not be rejected. We also test whether the data are drawn from populations that have the same median and cannot reject.

Table C.1: Home Improvements in San Diego

	House Value (in thousands)							
	<50	50-100	100-150	150-200	200-300	300-500	500-1,000	>1,000
Improvements (in percent)								
mean	2.11	1.05	0.73	0.75	0.95	0.96	0.8	1.1
	(0.54)	(0.32)	(0.11)	(0.09)	(0.13)	(0.12)	(0.13)	(0.30)
median	0	0.07	0.08	0.10	0.14	0.19	0.13	0.13

Note: This table contains the estimated means of home improvement measured as percent of house value. These statistics are computed for observations within the house price bins indicated on the top of the table. The data are the San Diego County observations of the 2002 American Housing Survey on 'rac' which measures the cost of replacements/additions to the unit. The 'rac' amount is divided by two, because the survey asks about expenses within the last two years. Standard errors (in brackets) are computed using Jacknife replications.

Census house values

The Census data does not contain actual prices but rather price ranges, including a top range for houses worth more than one million dollars. The topcoded range contains 9.6% of houses in 2005 and 1.8% of houses in 2000.To obtain cross sectional distributions of houses sold in a given year, we fit splines through the bounds of the Census house binds. Let p^c a vector that contains those bounds, as well as a lower bound of zero and an upper bound. We can obtain a continuous distribution for every upper bound by fitting a shape-preserving cubic spline through $(p^c, G_0(p^c))$. We choose the upper bound such that the median house value in the topcoded range equals the median in that range in our transaction data. To prepare the imputation of wealth (described below) we set a household's housing wealth to the midpoint of its bin, and we use the median of the topcoded range for the top bin.

Imputation of wealth

For age and income, we use age of the household head and income reported in the 2000 Census (for $t = 2000$) and 2005 ACS (for $t = 2005$). We are thus given age and income, as well as a survey weight, for every survey household. However, Census data do not contain wealth. We construct a conditional distribution of wealth using data from the Survey of Consumer

Finances (SCF). We use the 1998 and 2004 SCF to build the conditional distributions for 2000 and 2005, respectively.

We use a chained equations approach to perform imputations. The estimation is in two steps. In the first step, we use SCF data to run regressions of log net worth on log housing wealth, a dummy for whether the household has a mortgage and if yes, the log mortgage value, and log income for each age decade separately. In the second step, we use a regression switching approach that draws regression coefficients to generate a distribution of wealth for a given set of regressors. For each original household in the census sample, we then create three households with the same income and age, but with different wealth levels given by each of the three possible realizations for wealth using our imputation method. A survey weight for each new household is obtained by dividing the original survey weight by one third.

D Computations

This section describes the computational methods used to solve the quantitative model in Section (4). We need to (i) solve a household problem with a continuum of housing assets with different service flows and prices and (ii) solve for the equilibrium objects (service flow for 2000, and price for 2005) given the three-dimensional distribution of household characteristics and the one-dimensional distribution of house qualities.

Both the price and service flow functions are defined on the interval of available house qualities $[\underline{h}, \overline{h}]$. Both are parametrized as shape-preserving cubic splines, defined by a set $\{h_i, s_i, p_i\}_{i=1}^{I}$, where $h_i \in [\underline{h}, \overline{h}]$ are the break points, $p_i \in [0, \infty)$ is the price p_i at h_i and s_i is the service flow at h_i. We impose strict monotonicity on both functions, that is, $h_j > h_i$ implies $p_j > p_i$ and $s_j > s_i$. Denote the approximating price and service flow functions by $\widehat{p}(h)$ and $\widehat{s}(h)$, respectively. The intertemporal household problem is tractable even with a continuum of assets because agents expect permanent shocks to not alter relative prices across houses. The price function expected in the future equals the cumulative permanent innovation to house prices k plus the price function $\widehat{p}(h)$.

To accurately capture the covariation in the three mover characteristics age, income and wealth, we use the distribution derived from the Census and SCF using the imputation procedure in Appendix C. For every survey household i at date t, we have a tuple $(a_{it}, y_{it,} w_{it})$ as well as a survey weight. We solve the household problem for every survey household i and obtain his preferred house quality. We then use the survey weights to construct a cumulative distribution function for house quality. In equilibrium, this cdf must be equal to the house quality cdf from the data, shown in Figure 2. The equilibrium object (price or service flow) is found by minimizing a distance between those cdfs.

Household problem

The solution to the household problem is calculated using finite-horizon dynamic programming. Value and policy functions are approximated by orthogonal polynomials. Consider the optimization problem faced by a household of age a, with cash w as defined in equation (17), income y, and a house of quality h. Let k denote the cumulative house price innovation. Each period, the household receives an exogenous mobility shock: $m = 1$ indicates that the household must move and $m = 0$ otherwise. The vector of state variables at time t is

$$x_t = [a_t, m_t, h_t, k_t, y_t, w_t].$$

The value function at time t is denoted $v(x_t)$. Income is a separate state variable even though the only shocks to income are permanent. This is because house prices are hit by shocks other than income shocks – the common approach of working with the wealth/income ratio and house/income ratio as state variables does not apply.

It is helpful to separate the household's moving decision from the other choices he makes conditional on moving or staying. Consider first a household who is moving within the period. He decides how to allocate cash on hand (which could come from a prior sale of a house) to consumption, housing or bonds, subject to the budget and collateral constraints. Denote the "mover value function" for this problem by v^m; it depends on the state as well as the approximating price and service flow functions. Consider next a household who is staying in a house of quality h. He decides how to allocate cash on hand to consumption or bonds, again subject to budget and collateral constraints. A stayer household is thus allowed to change his mortgage – this assumption is appropriate for the boom period where refinancing and home equity loans were common.

The beginning-of-period value function v takes into account both forced moves ($m = 1$) and endogenous moves:

$$v(x) = m\, v^m(a, k, y, w; \hat{p}, \hat{s}) + (1 - m) \max \{v^s (a, h, k, y, w; \hat{p}, \hat{s}), v^m(a, k, y, w; \hat{p}, \hat{s})\}$$

The discrete choice of moving can induce kinks in $v(x)$, thus making global polynomial approximation undesirable. Instead, we specify separate approximating functions for $v^m(\cdot)$, $v^s(\cdot)$, as well as the housing policy function $h^m(a, k, y, w; \hat{p}, \hat{s})$ associated with $v^m(\cdot)$. The functions v^m and h^m are smooth in the dimensions of the cumulative price shock k and income y. The down-payment constraint may however induce kinks in the cash dimension. We address this issue by employing a global least-squares approximation using two-dimensional Chebychev polynomials in the (k, y) plane. We then interpolate the coefficients of the approximating polynomials in (k, y)-space at different values of w by shape-preserving cubic splines.

Market clearing

Given a sample of movers with characteristics $\{a_{it}, y_{it,} w_{it}\}_i$ as well as approximating price and service flow function $\hat{p}$ and $\hat{s}$, we calculate the model-implied optimal house qualities as[13]

$$\widehat{h}_{it} = h^m (a_{it}, y_{it}, w_{it}, 0; \widehat{p}, \widehat{s}).$$

We thus obtain a sample of optimal house quality choices $\{\widehat{h}_{it}\}_i$. We then use the survey weights for the movers to compute an empirical cdf of house quality. We smooth this cdf using a cubic spline. We call the resulting cdf $\hat{G}^{dem}(h; \widehat{p}, \widehat{s})$ the demand cdf as it represents optimal housing demands at the given price and service flow functions.

In equilibrium, the demand cdf must equal the quality cdf from the data. The latter is also given as a cubic spline, $\hat{G}$ say, as explained in Appendix C. To get a measure of distance

[13] Note that the policy function is evaluated at $k = 0$ for all households. The cumulative price shock k is only relevant to the extent that it induces uncertainty about the future price level relative to the current price level.

between the demand cdf and the data quality cdf, we define a set of test quantiles $\{g_j\}_{j=1}^{N_G}$, $g_j \in (0,1)$ and compute

$$\sum_{j=1}^{N_G} \left\{ \hat{G}(\left[\hat{G}^{dem} \right]^{-1} (g_j; \widehat{p}, \widehat{s})) - g_j \right\}^2 . \tag{D-1}$$

For our exercises we need to find the equilibrium object (price or service flow), taking as given the respective other function (service or price). In each case, our algorithm chooses the spline coefficients of the equilibrium function to minimize the distance (D-1). For the reported results we use 7 break points and the test quantiles are the nine deciles between 10% and 90% as well as the 1st, 5-th and 95-th percentiles. The error is within one percentage point at every test quantile.