Macroeconomic Models with Heterogeneous Agents and Housing

KARSTEN JESKE

The author is a research economist and assistant policy adviser in the macropolicy group of the Atlanta Fed's research department. *He thanks William Roberds and Ellis Tallman for helpful comments.*

It is hard to miss the important role the housing sector plays in the macroeconomy. In 2004 the value of real estate held in household portfolios amounted to over $17 trillion, or 143 percent of annual nominal gross domestic product (Board of Governors 2005). This amount is larger than all corporate equity held directly, through mutual funds and pension funds combined. Closely linked to housing is the mortgage market, which now totals over $7.5 trillion. On the average household's balance sheet, a home mortgage is the largest item on the liability side, easily dwarfing consumer credit, which includes credit card debt and car loans, by a factor of almost four to one.

Macroeconomists, therefore, have spent considerable effort on incorporating housing and housing finance into macroeconomic models. In understanding the effects that housing and the various related government policies have on households in particular and the economy in general, recognizing the economic and demographic diversity among households is crucial. Net worth, housing value (gross and net of mortgages), and income vary substantially among households not only by age but within age groups. Macroeconomists thus employ models—called heterogeneous agent models—that can accommodate this diversity.

This article is a progress report on where this line of research currently stands. At the Federal Reserve Bank of Atlanta conference "Housing, Mortgage Finance, and the Macroeconomy" in May 2005, several papers shared a common framework: namely, using macro models with heterogeneous agents.[1] This article discusses four of the conference papers, the issues the researchers explored, the progress they made, and the challenges that lie ahead.

The paper by Fang Yang (2005) deals with the life-cycle behavior of housing versus nonhousing consumption. Her objective is to build a model to account for two peculiar features of consumption over the life cycle.

The paper by Wenli Li and Rui Yao (2005) studies the effect of house price changes on the macroeconomy. They point out that while house price changes may

Table 1
Homeownership Rates by Age Group

Age group	Homeownership rates (percent)
0–29	30.1
30–39	58.4
40–49	74.0
50–59	79.8
60–69	81.7
70+	79.3

Note: The unconditional mean across all ages is 67.7 percent.

Source: Board of Governors of the Federal Reserve System, Survey of Consumer Finances 2001

not have a huge effect on the economy as a whole, different households—in particular those of different age and asset holdings—will be affected very differently by house price changes. An economic model with heterogeneous agents is therefore necessary to measure the effects a price increase might have on agents with different incomes and asset holdings.

The paper by Matthew Chambers, Carlos Garriga, and Donald Schlagenhauf (2005b) asks how different households are affected by the availability of different mortgage contracts: most importantly, which mortgage contract is most successful in allowing younger households (that tend to have less savings and thus might have trouble coming up with the necessary down payment) to purchase a house.

The paper by Karsten Jeske and Dirk Krueger (2005) addresses the important policy question of whether it is desirable to subsidize mortgage interest rates. One important subsidy is the implicit federal guarantee for government-sponsored enterprises (GSEs). Jeske and Krueger examine the distributional effect of such a subsidy. They find that mostly high-income and high-net-worth households benefit while, both in terms of welfare and homeownership, low-income and low-net-worth households will not benefit at all.

Microlevel Data on Housing and Mortgage Finance

Trying to answer the questions mentioned above requires models that are detailed enough to accommodate a housing market and, most importantly, heterogeneity among households. This heterogeneity is the crucial ingredient of any model attempting to determine the distributional effects of housing policy. Before examining the models, I first introduce some empirical facts coming out of microlevel data to demonstrate how the aggregate numbers on income, real estate, and mortgage debt are distributed across the population. Aggregate real estate values are approximately 143 percent of aggregate annual income. But income, real estate wealth, mortgage debt, and net worth are very unevenly distributed both across and within age groups. The data set used in this article's computations is the Survey of Consumer Finances (from the Federal Reserve Board) for the year 2001, which, among other things, collects data on consumer income and balance sheets.

Table 1 shows homeownership rates by age group compared to the unconditional homeownership rate (that is, over all age groups) of 67 percent. The homeownership rate is the lowest, at about 30 percent, for households headed by persons below the age of twenty-nine. The rate increases steeply over the next two age cohorts, peaking at more than 80 percent for the sixty to sixty-nine age group and moderating slightly after that. Not all homeowners are equal, as Table 2 shows. Not only do fewer young households own a home, but their homes tend to be smaller than those of older households. The average homeowner increases real estate holdings from about \$100,000 until reaching his or her fifties, when he or she holds about \$216,000 in real estate, and then gradually downsizes the value of the primary residence. Mortgage debt also displays a hump-shaped profile, though the peak occurs in the thirty to thirtynine age group. This pattern seems to be consistent with the following scenario: Households upgrade the size and obviously the value of the house they are living in. Early in life most of the upgrading is financed with larger mortgages. Homeowners aged forty and older then tend to finance upgrades out of savings rather than larger mortgages.

Table 2
Home Values and Mortgage Debt by Homeowners' Age Group

Age group	Value of primary residence	Mortgage on primary residence
0–29	103.6	65.5
30–39	157.1	91.4
40–49	190.0	82.1
50–59	216.4	68.5
60–69	193.5	35.4
70+	171.7	11.2

Note: Values and debt shown in thousands of dollars.

Source: Board of Governors of the Federal Reserve System, Survey of Consumer Finances 2001

In homeownership rates, even within age groups there is substantial variation. The first panel of Figure 1 plots rates by income quintiles within each age group. The chart shows that homeownership is associated with high income, with households in the bottom quintile having below-average ownership rates across all age groups. The second panel of Figure 1 yields the same qualitative pattern for net worth, but the differences between households with low and high net worth households are even more severe. Households in the bottom quintile have only a 30 percent homeownership rate even in the sixty to sixty-nine age group. In contrast, for all ages after forty, the second through fifth quintiles display homeownership rates above the national average of 67 percent.[2]

Vast differences also exist between homeowners and renters. The first panel of Figure 2 plots median income by age group for the whole population and for homeowners versus renters. For the population as a whole, household income increases from just below $30,000 for the youngest age cohort to almost $60,000 for fifty to fifty-nine age group, after which it declines to about $25,000 for the seventy and over age group. Homeowners tend to have much larger median incomes than renters; the median homeowner makes roughly twice as much as the median renter. Moreover, the hump in the income process is more pronounced for homeowners. The same general pattern holds for average incomes (see the second panel of Figure 2).

The contrasts between homeowners and renters are even more obvious in terms of net worth positions. The third and fourth panels of Figure 2 plot median and average net worth, respectively, by age group. Households accumulate savings during their working years, moreso than does the population as a whole. The median homeowner accumulates more than $200,000 during his or her lifetime. Because the net worth distribution is so skewed, average net worth is considerably larger: The average homeowner accumulates more than $800,000 over a lifetime. The median renter, however, has negligible net worth—consistently under $10,000—for all age groups. In fact, 50 percent of all renters have essentially zero net worth. Even average net worth for renters is below $100,000 in all age groups.

1. The four papers introduced here as well as the program and the remaining papers are available on the Atlanta Fed's Web site at www.frbatlanta.org under "News & Events," "Conferences," "2005."
2. The cutoff for the lowest income quintile is $15,400 in the youngest age group, increases to $26,700 for the fifty to fifty-nine age group, and then declines to $11,300 for households aged seventy and above. The cutoff for net worth is zero for the youngest households and increases to $39,100 for households aged seventy and above.

Figure 1
Homeownership Patterns by Age, Income, and Net Worth

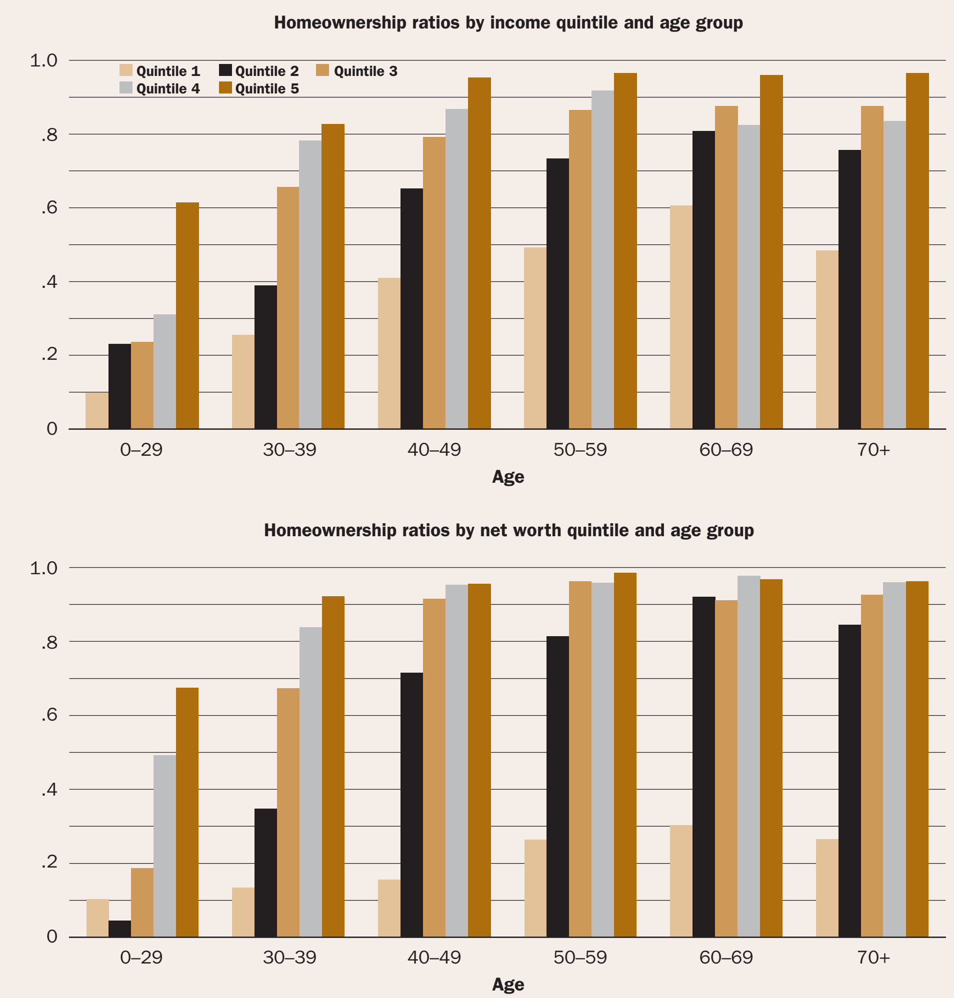

Source: Board of Governors of the Federal Reserve System, Survey of Consumer Finances 2001

In summary, households differ substantially, both between and within age groups, in their net worth positions and asset allocations. As one would expect, when households are young and middle-aged they accumulate savings for retirement, and one form of savings is real estate. Even within age groups, there is considerable heterogeneity of households. Incomes and especially net worth vary substantially across age groups and between renters and homeowners. When trying to answer questions such as "What is the effect of increasing house prices?" or "Should we subsidize mortgage interest rates?" one should take into account that different households will be affected very differently by changes in house prices or government policies. An increase in house prices might be beneficial to existing homeowners, but renters may not be affected at all or, even worse, might suffer if rental rates increase. Likewise, subsidizing mortgage interest—for example, through mortgage interest tax-deductibility or

Figure 2
Income and Net Worth by Age Group

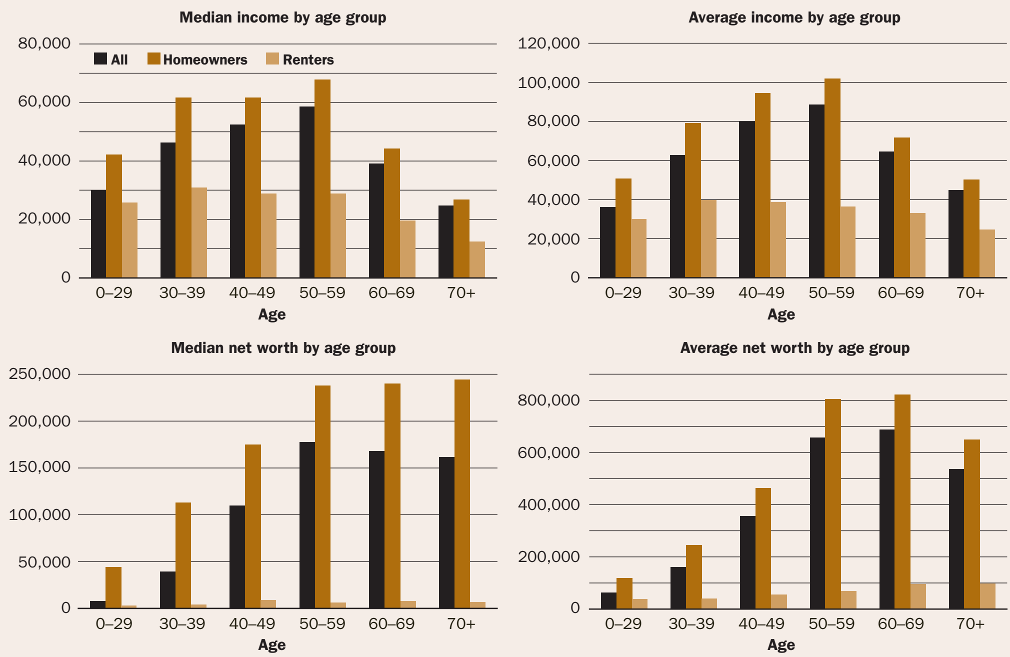

Source: Board of Governors of the Federal Reserve System, Survey of Consumer Finances 2001

government guarantees on mortgage lenders—might mostly benefit existing homeowners, who are well off already because they tend to have above-average income and net worth. This consideration reinforces the importance of studying models with heterogeneous agents to account for the effects that policy has on agents with different demographic and economic backgrounds.

A Generic Model with Housing

A key ingredient in the model is the life-cycle pattern, which implies two important features the model must include. First, it must include the hump-shaped earnings profile that the average household experiences, as shown in Figure 2. Second, the model must generate a realistic life span, which means households of different age groups face different mortality risk. Around the trend path of life-cycle earnings, a household receives an uncertain stream of income over its life cycle and then tries to maximize its lifetime discounted utility subject to a budget constraint.

A household receives utility from consumption of both nonhousing goods c and housing goods h. Income comes both from labor and investment income. The maximization problem of a household of age t in recursive form is then defined in the following way: The value of starting period t with assets a and labor productivity y is the maximum value of current utility plus expected discounted period $t + 1$ value as a function of future assets and future labor productivity.

$$V_t(a,y) = \max_{c,h,b'} u(c,h) + \rho_t \beta E V_{t+1}(a', y'),$$

subject to a number of constraints:

$c + h + b' = wy + a;$
$a' = (1 - \delta^h)h + (1 + r)b';$
$b' \geq -\bar{B};$
$c, h \geq 0.$

In the objective function I use a discount factor β common to all age groups as well as additional discounting in the form of ρ_t, which is the conditional survival probability that depends on age. The first constraint is the budget constraint. Labor income measured as wage rate w times labor productivity y plus assets a can be spent on nonhousing consumption c, housing h, and other savings b'. The second constraint specifies the next period's assets a' as the sum of housing stock net of depreciation and maintenance cost $(1 - \delta^h)h$ and other savings times the gross interest rate $(1 + r)b'$. The third condition is a borrowing constraint that specifies a lower bound on asset holdings—in other words, an upper bound on borrowing. Finally, the fourth condition states that both types of consumption have to be nonnegative. This model closely resembles that of Aiyagari (1994), with some distinguishing features: two types of consumption, a very particular borrowing constraint, and a life-cycle component.[3]

Denote u_c and u_h the derivatives of the utility function with respect to nonhousing consumption c and housing h, respectively. Assuming for now that the borrowing constraint does not bind, we can immediately derive two optimality conditions, also called Euler equations:

$$(1)\quad 1 = \rho_t \beta E \frac{u_c(c', h')}{u_c(c, h)}(1 + r), \text{ and}$$

$$(2)\quad 1 = \rho_t \beta E \frac{u_c(c', h')}{u_c(c, h)}(1 - \delta^h) + \frac{u_h(c, h)}{u_c(c, h)}.$$

The first condition is completely standard: The bond return, weighted by the intertemporal marginal rate of substitution (also called the pricing kernel) and discounted by $\rho_t\beta$, equals one. The second condition has a similar structure but also incorporates today's marginal rate of substitution between housing and nonhousing consumption. Both conditions guarantee optimality by ensuring that no utility-improving substitution exists between either consumption and savings or consumption and housing. Next, combine the two Euler equations into

$$\frac{u_h(c, h)}{u_c(c, h)} = \frac{r + \delta^h}{1 + r}.$$

This condition states that the marginal rate of substitution between consumption and housing equals the cost of one unit of housing computed as the depreciation rate plus the opportunity cost of housing measured as the bond rate r. I also divide by $1 + r$ because the payoff from housing is in the current period while the bond pays off in the next period. The right-hand side is also called the user-cost formula for housing, and if there were a rental market in this model, $(r + \delta^h)/(1 + r)$ would be the rental rate per unit.

To put some structure into the model, assume that the utility function takes the following form:

$$u(c,h) = \frac{\left[\theta c^{\eta} + (1-\theta)h^{\eta}\right]^{\frac{1-\sigma}{\eta}} - 1}{1-\sigma}.$$

This is the form normally used in the literature. The parameter σ is the risk aversion coefficient, η determines the substitutability between housing and nonhousing consumption, and θ is a weighting parameter for how much the household values nonhousing consumption. How good is this model in accounting for the observed data? From the model, one can deduce two implications. First, the user cost formula for housing can be written as

$$\frac{h}{c} = \left(\frac{r+\delta^h}{1+r}\frac{1-\theta}{\theta}\right)^{\frac{1}{\eta-1}},$$

that is, housing over nonhousing consumption stays constant over the life cycle. Next, from the first Euler condition (equation [1]) and the fact that the *h/c* ratio stays constant over time one deduces that

$$E\left(\frac{c'}{c}\right)^{-\sigma} = \frac{1}{\rho_t \beta (1+r)}.$$

Normally it will be the case that $\beta(1 + r) < 1$.[4] Since the survival probability ρ_t is less than one, the right-hand side is greater than one. Since $\sigma > 0$, the second implication of the model has now been shown: Without borrowing constraints, consumption decreases over the life cycle, at least in expected terms. This result makes perfect sense. According to the permanent income hypothesis, the mean path of earnings is irrelevant in the absence of borrowing constraints; all that counts is the discounted expected path of earnings. If the interest rate is lower than the rate of discounting, households will choose a decreasing consumption path though there may be fluctuations around this decrease because of changes in labor productivity.

Life-Cycle Patterns of Housing Consumption

The four conference papers mentioned in the introduction are all extensions of this same basic model. Yang (2005) observes that both implications of the model,

3. Note that the timing convention here is slightly different from that in the Yang (2005) paper. Specifically, I assume here that the housing choice h enters today's utility instead of tomorrow's. This assumption makes the problem slightly easier to handle because agents pick housing goods for the same period when their income uncertainty is revealed. Therefore, housing is a riskless asset, which makes the optimality conditions slightly more tractable. Jeske and Krueger (2005), for instance, use this timing convention. In their paper housing investment will nevertheless be risky because of uncertain depreciation. In the Yang paper, on the other hand, households pick h for the following period when the future productivity shock y' is not known yet, a choice that makes housing a risky asset.
4. Economists usually pick the discount factor β so as to match the model's capital-to-output ratio to that in the data. One can show that for the realistic range of the remaining parameters of the model the equilibrium interest rate r implied by this capital-to-output ratio satisfies $\beta(1 + r) < 1$.

Figure 3
Consumption over the Life Cycle

Source: Yang (2005) (first panel); author's calculations using data from Yang (2005) (second and third panels) a decreasing consumption profile and a constant housing versus nonhousing consumption ratio, are at odds with the data. The first panel of Figure 3 shows average household consumption on both housing and nonhousing goods by age group (of the head of household). Notice that nonhousing consumption increases from about \$23,000 per year for households aged twenty to an average of \$32,000 per year for age forty-five and then decreases after that.[5]

Contrast this observation to the behavior of housing consumption, which displays an even sharper increase up to about age fifty-five but then does not drop as sharply as nonhousing consumption. Not surprisingly, then, the *h/c* ratio is not constant but rather increases over the life cycle, as shown in the second panel of Figure 3.

Can the increase in the *h/c* ratio be accounted for without radical changes to the model? One could imagine that younger agents face a higher depreciation rate of housing or a higher interest rate.[6] Formally, assume that the consumption ratio at age t now takes the following shape:

$$(3)\quad \frac{h_t}{c_t} = \left[\left(\frac{r+\delta^h}{1+r} + \xi_t\right)\frac{1-\theta}{\theta}\right]^{\frac{1}{\eta-1}},$$

where ξ_t generates differences in households' subjective user cost because of differences in depreciation and interest rates during the life cycle. Is it possible to quantitatively match the observed ratio with realistic values of ξ_t? How positive would ξ_t have to be for young agents and how negative for old agents in order to generate the *h/c* function observed in the data? To compute this cost, I normalize the implicit cost of a forty-five-year-old household to zero and back out the ξ_t from equation (3) using the observed *h/c* ratios. The third panel of Figure 3 shows this implicit cost ξ_t over the life cycle. To account for the observed consumption patterns, older households would have to find housing around 3 percentage points cheaper than those age forty-five do, and, likewise, younger households would have to find housing 6 percentage points more expensive. It is difficult to justify differences of this magnitude—almost 10 percentage points difference between age eighty-five and age thirty—entirely through different borrowing costs or depreciation rates.

Yang's paper can be thought of as finding other explanations to generate this large difference in subjective user costs. Specifically, Yang introduces two crucial ingredients into the model to account for the features observed in the data—binding borrowing constraints and transaction costs—and shows that they can quantitatively account for the increase in the consumption ratio.[7] One can easily add these two features into the model above. First, assume that $\bar{B} = -(1-\gamma)h$, that is, the household can borrow only up to a maximum of a certain fraction of the house value. Parameter γ can be viewed as the minimum down payment on a house. This constraint ensures that with a hump-shaped labor income profile (see the first and second panels of

5. Yang is not the first economist to point out this feature. Earlier work on the consumption profile over the life cycle includes Blundell, Browning, and Meghir (1994); Attanasio and Browning (1995); Attanasio and Weber (1995); and, more recently, Gourinchas and Parker (2002) and Fernandez-Villaverde and Krueger (2002, 2004).
6. For example, one could justify this assumption with the fact that households with a younger head tend to have a larger family size, and thus more persons per square foot cause a higher depreciation rate.
7. Yang introduces even more ingredients, such as a bequest motive and a social security system; however, to illustrate the key intuition, a borrowing constraint and transaction costs are sufficient.

Figure 2) agents will be constrained in their borrowing early in life. In other words, households cannot raise their consumption to the desired level derived from the permanent income model without violating their borrowing constraint. This restriction binds because income is low early in life.

Second, adjusting the level of housing incurs transaction costs. Assume that if a household had a housing level of h_{old} last period and wants to change it to h today, it has to pay a transaction cost of $\psi(h_{old}, h)$. Assume for now that the transaction cost satisfies the following conditions:

$$\psi(h_{old}, h) \geq 0;$$
$$\psi(h_{old}, h_{old}) = 0;$$
$$\frac{\partial^2 \psi(h_{old}, h)}{\partial h^2} \geq 0.$$

That is, the transaction cost function takes on nonnegative values, there is no cost of leaving the size of the house unchanged,[8] and the transaction cost function is convex.[9] In this economy the optimality condition between housing and nonhousing consumption becomes

$$(4)\quad \frac{u_h(c,h)}{u_c(c,h)} = \left(1 - \frac{\mu}{u_c(c,h)}\right)\frac{r+\delta^h}{1+r} + \frac{\mu}{u_c(c,h)}\gamma + \psi_2(h_{old},h) + \rho_t \beta E \frac{u_c(c',h')}{u_c(c,h)}\psi_1(h,h'),$$

where μ is the Lagrange multiplier on the borrowing constraint and ψ_1 and ψ_2 are derivatives of the transaction cost function with respect to the first and second argument, respectively. This formula is the same one as before, linking the marginal rate of substitution between housing and nonhousing consumption with the cost of purchasing housing. Instead of the user-cost formula of housing, one now has a weighted average between the previous user cost $(r + \delta^h)/(1 + r)$ and the down payment ratio γ, where the weight on γ is proportional to the Lagrange multiplier of the borrowing constraint.[10]

One can interpret the first two terms on the right-hand side of equation (4) as the subjective user cost of housing (net of transaction cost). The more the borrowing constraint binds, the higher the housing cost for the household. This result is intuitive: If the constraint binds, the only way to afford more housing is to reduce current consumption, which comes at a cost of γ units of consumption for every unit of housing purchased. The remaining two terms come from the transaction cost function. They determine the marginal cost of housing changes—that is, adjusting housing consumption from h_{old} to h in the current period and adjusting from h to h' in the following period.

Using the same elasticity of substitution between the two consumption types as in the Yang paper—that is, $\eta = 0$—one then obtains the following consumption ratio:

$$(5)\quad \frac{h}{c} = \frac{1-\theta}{\theta}\left[\left(1 - \frac{\mu}{u_c(c,h)}\right)\frac{r+\delta^h}{1+r} + \frac{\mu}{u_c(c,h)}\gamma + \psi_2(h_{old},h) + \rho_t \beta E \frac{u_c(c',h')}{u_c(c,h)}\psi_1(h,h')\right]^{-1}.$$

One way to interpret Yang's paper is that the additional terms in the square brackets tend to decrease in age; thus, inversely, the consumption ratio increases in age. For young agents there are two channels: the borrowing constraint and the transaction cost. A binding borrowing constraint has the same effect as a higher interest rate.

Since the subjective user cost for young agents (before transaction costs) is higher than the cost $(r + \delta^h)/(1 + r)$ for unconstrained agents, the h/c ratio will increase just as in the data because the subjective user cost decreases when the borrowing constraint is relaxed over the lifetime.

So far, the larger implicit housing cost causes young households to substitute away from housing and into nonhousing goods, even without any transaction costs. Transaction costs add to this effect. Since the borrowing constraint induces a hump-shaped profile for consumption (both housing and nonhousing), young agents have the tendency to increase their h. Consequently, the following will be true for any three consecutive housing consumption values: $h_{old} \leq h \leq h'$. In that case $\psi_2(h_{old}, h) \geq 0 \geq \psi_1(h, h')$. Since the second term is discounted by $\rho_t \beta E[u_c(c', h')/u_c(c, h)]$, one would expect the total effect to be positive. Both effects from borrowing constraints and transactions costs explain why for young agents h/c can be lower than the ratio derived in the frictionless version above.

In understanding the effects that housing and related government policies have on households and the economy, recognizing the economic and demographic diversity among households is crucial.

For older agents the borrowing constraint will likely no longer bind. Instead, all of the action comes from the transaction cost terms. The reverse of the argument for young agents applies. Older agents want to reduce consumption because of discounting due to lower survival probabilities. But picking $h_{old} \geq h \geq h'$ implies $\psi_2(h_{old}, h) \leq 0 \leq \psi_1(h, h')$. Again, the last term in equation (5) is subject to discounting, so the $\psi_2(h_{old}, h) \leq 0$ dominates, making the h/c ratio higher than in a frictionless economy. The bottom line of the Yang (2005) paper is that by adding the two modifications the model matches the data qualitatively and quantitatively.

The Effect of House Price Changes

Li and Yao (2005) study the effect of house price changes on the macroeconomy. The first panel of Figure 4 displays the U.S. house price index, adjusted for inflation by deflating the nominal house price index with three different commonly used measures of inflation: namely, the consumer price index (CPI) and the deflators for both gross domestic product (GDP) and personal consumption expenditures (PCE). Between 1975 and 1995 real house prices more or less stagnated if deflated by the CPI but moderately increased by about 20 percent over the twenty years if deflated by GDP or PCE.[11] Of course, the path of house prices was not monotonic. A decline in real house prices between 1979 and 1983 resulted from the high inflation rates in the late 1970s and the two recessions in the early 1980s. Likewise, the 1991 recession

8. Recall that in addition there is maintenance cost, but that enters the budget constraint separately.
9. The shape of transaction costs is slightly different from that used in Yang's paper, or most of the literature for that matter. Using the convex transaction cost, however, it is easier to explain the intuition using the Euler conditions because the necessary optimality conditions will also be sufficient.
10. Notice that in equilibrium it is always true that $\gamma > (r + \delta^h)/(1 + r)$. If this were not the case, one can show that housing investment is strictly more profitable than purchasing bonds, and nobody would ever hold positive amounts of bonds. This outcome cannot be an equilibrium outcome since the bond market has to clear—that is, bond savings have to equal mortgage lending.
11. This result makes perfect sense: GDP and PCE deflators tend to be slightly lower than the CPI because the CPI computes price changes of a fixed basket while in the GDP and PCE consumers obviously substitute into cheaper goods, thus lowering those two indexes.

Figure 4
Consumption over the Life Cycle

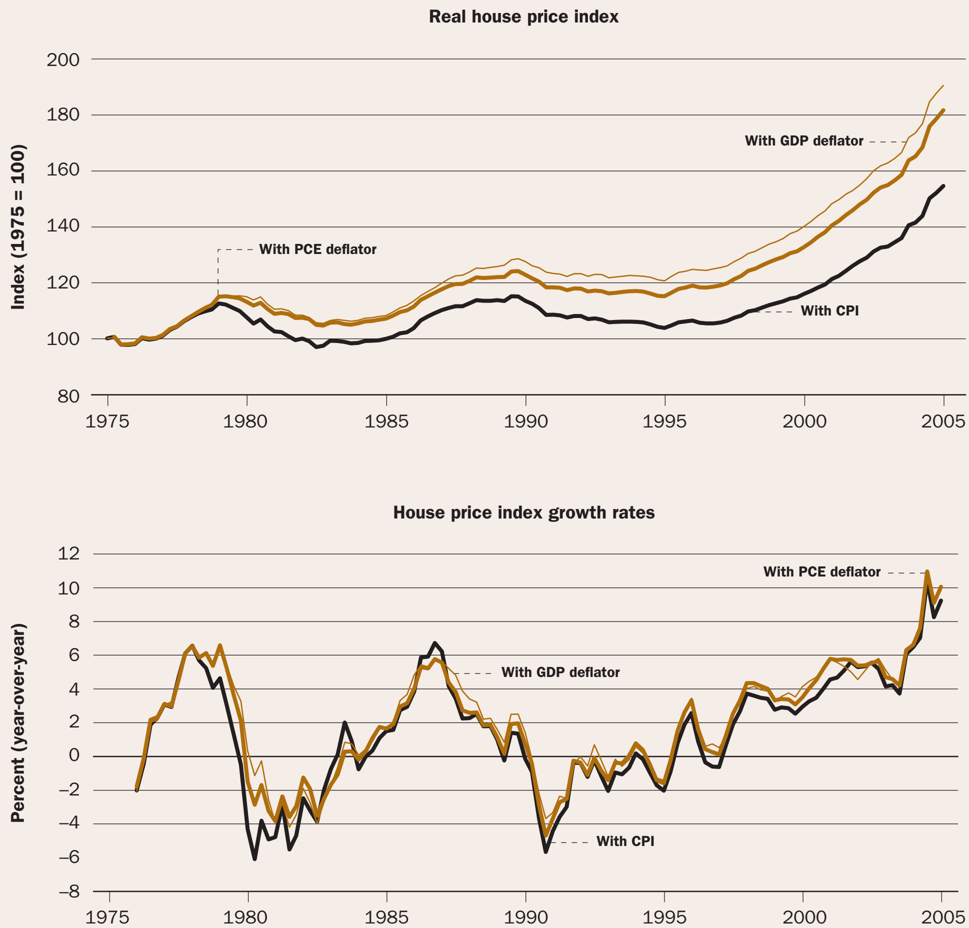

Source: Office of Federal Housing Enterprise Oversight, Bureau of Labor Statistics, and Bureau of Economic Analysis

took a bite out of house prices, and even during the first four years of the recovery, house prices stagnated.

After 1995, however, house prices staged an astonishing increase of almost 50 percent if deflated by CPI and 60 percent if adjusted by either of the deflators. Moreover, no decline occurred in house prices around the 2001 recession, and no stagnation occurred after the recession. On the contrary, as the second panel of Figure 4 shows, year-over-year growth rates in real house prices even accelerated during and after the recession to a record pace of 8 to 10 percent in 2004.

What is the effect of these house price fluctuations, especially the large run-up in prices since the mid 1990s? In a model without heterogeneity, a life-cycle earnings profile, and borrowing constraints, the effect on macroeconomic variables is exactly zero in the following sense: Imagine an economy has only one representative consumer, earning exactly the average income and holding exactly the average house worth 143 percent of average income. An increase in house prices would have a positive wealth effect and a negative effect because housing in the present and future becomes more expensive. One can show, however, that the two effects precisely cancel each other out, so relative price changes between c and h are neutral. Utilizing the earlier model, one can formalize this intuition. Imagine one unit of housing now costs P_h units of the consumption good. Abstracting from transaction costs, the optimality condition for housing versus consumption now becomes

$$(6) \quad \frac{u_h(c,h)}{u_c(c,h)} = \left(1 - \frac{\mu}{u_c(c,h)}\right)\left(1 - \frac{P_h'}{P_h}\frac{1-\delta^h}{1+r}\right) + \frac{\mu}{u_c(c,h)}\gamma.$$

This structure is familiar: The marginal rate of substitution must equal the relative price of the two goods. Relative price is the household's subjective cost, which in this case is a mixture of the user cost formula $1 - (P_h'/P_h)[(1 - \delta^h)/(1 + r)]$ and the cost of the borrowing constraint. Notice that for prices $P_h = P_h'$ this user cost formula reduces to equation (4) without the terms for the transaction cost. Notice also that permanently changing all future real estate prices to the same level will have no effect on the h/c ratio because house prices show up in the optimality conditions only as ratios of two consecutive house prices. One can then show that in the one-person economy the consumption path without the price change is still affordable after the price change. Also, because this path satisfies the new optimality conditions, the optimal behavior of the household is to leave the consumption path unchanged, meaning that a permanent house price change is completely neutral.

Heterogeneity among households is the crucial ingredient of any model attempting to determine the distributional effects of housing policy.

The only way house price changes can have nontrivial effects is if households are heterogeneous. Li and Yao (2005) use a model similar to Yang's, though modified in several dimensions. Most importantly, Li and Yao introduce a rental market in which households can attain housing services via two routes: either purchasing a home, with the purchase subject to a transaction cost, or renting it instead at a cost of αP_h. One can think of the parameter α as the rental return per unit rented.

Li and Yao then simulate the economy after a permanent price change to housing. Not surprisingly, renters suffer from a house price appreciation. As opposed to the representative household mentioned above, renters have no positive wealth effect at all because, by definition, they do not own real estate. Renters are, however, hit with an increase in their future housing costs regardless of their future tenure decision: If they decide to buy a house, they have to pay more, and if they continue to rent, their cost per unit of rental goes up because the rental price is a fixed portion of the house price P_h. Li and Yao can even quantify the loss renters suffer. The loss in utility due to an increase in P_h by two standard deviations (11.5 percent) is equivalent to permanently reducing all current and future consumption by about 4.5 percent.

The effect on renters is obvious and unambiguous, which is not necessarily the case for homeowners. Li and Yao show that house price appreciation can in fact be disadvantageous to some homeowners. The intuition for their result has to do with the pattern of homeownership displayed in Table 2 (on page 17). The average value of the primary residence increases for every age group until around age fifty. As mentioned before, this increase is most likely due to the binding borrowing constraint that keeps young households from attaining consumption levels (both housing and nonhousing) implied by the permanent income hypothesis. Suppose a young household owns $100,000 worth of real estate and experiences a 20 percent increase in house prices. Instead of a $20,000 gain, the household may view this as a $20,000 loss if it was planning to upgrade to a larger house worth $200,000, which now costs $240,000, because the upgrade would now cost $120,000 instead of $100,000.

In conclusion, Li and Yao argue that a substantial portion of the population does not benefit from a house price appreciation. The break-even age for homeowners is at around age fifty, which incidentally is roughly the age at which households finish their upgrading. Older households gain substantially from house price appreciation, partly because many of them start moving into smaller houses or even become renters again and are thus able to cash out some of their capital gains. Even without downgrading, older households experience a welfare gain because the value of the bequest they leave to their heirs increases.

Different Mortgage Contracts and the Tenure Decision of Young Households

One can interpret the paper by Chambers, Garriga, and Schlagenhauf (2005b) as a follow-up to their earlier study (2005a), which uses a life-cycle model much like the ones presented earlier to explore potential reasons for the path of the homeownership rate over the past forty years. Homeownership stayed roughly constant at 64.5 percent between 1965 and 1995 but then increased significantly to now almost 68 percent.

Chambers, Garriga, and Schlagenhauf (2005a) argue that the main reason for this increase is the availability of mortgage contracts requiring lower down payments. They simulate the model economy, once with a tighter and once with a looser down payment constraint, and determine that relaxing the down payment constraint can indeed quantitatively account for the increase in homeownership. Specifically, they study the effect of going from a 20 percent down payment constraint to an 80-15-5 combo loan—that is, an 80 percent first mortgage, a 15 percent second mortgage, and a 5 percent down payment.

How can the arrival of a new mortgage contract increase homeownership, especially among young agents? Recall that younger households, say, between the ages of twenty and forty, are characterized by three features: They are far less likely to be homeowners and have both lower income and lower net worth than the average population. Equation (4) reveals the main culprit for low homeownership among young households: Because of the binding borrowing constraint, the subjective housing cost for young agents is higher than the rental cost, discouraging young agents from buying. Reducing the tightness of the borrowing constraint (lowering parameter γ) is then the most direct way to encourage homeownership among those agents with a positive Lagrange multiplier μ (a binding borrowing constraint). Chambers, Garriga, and Schlagenhauf (2005a) show that lowering γ from 20 percent in the benchmark to 5 percent under the 80-15-5 combo loan will indeed increase homeownership among young households and can quantitatively account for the rise in homeownership.

The great innovation in the two papers by Chambers, Garriga, and Schlagenhauf is that they model mortgages with much more care than in the rest of the literature. The other three papers discussed here—Yang (2005), Li and Yao (2005), and Jeske and Krueger (2005)—take one crucial shortcut when modeling mortgages: Households roll over mortgage debt every period, that is, they choose whatever amortization schedule suits them, constrained only by the borrowing constraint (in Yang and Li and Yao) or the interest rate dependence on the leverage ratio (in Jeske and Krueger). Chambers, Garriga, and Schlagenhauf, however, assume that a mortgage follows a fixed amortization schedule as it would for, say, a thirty-year fixed rate mortgage. One can think of a mortgage contract then as not only an interest rate r and a down payment ratio γ but rather a whole sequence of ratios over the lifetime of the loan. Suppose a thirty-year mortgage with an 80 percent loan-to-value ratio stipulates a sequence of γ starting at 20 percent and increasing very slowly over the first couple of years (because initially most of the mortgage payment goes toward interest rather than principal). Toward the end of the loan more and more of the payment goes toward reducing the principal, and eventually γ reaches 100 percent when the loan is paid off.

Relaxing the down payment constraint can greatly increase homeownership among younger households because it reduces the subjective user cost.

Another example would be a balloon loan, in which case γ stays at 20 percent for the duration of the loan and after which the balloon payment comes due, which requires the household to either pay off the loan ($\gamma = 1$) with one large payment or refinance into another mortgage. Naturally, different loan designs imply different payment schedules. Since the model is detailed enough to mimic a whole range of different mortgage contracts, Chambers, Garriga, and Schlagenhauf (2005b) can study how different contract types affect the tenure decision, especially of young agents.

The conclusion of their paper is that mortgage contracts that tend to have low payments early on, such as a balloon mortgage, are most successful in relaxing the borrowing constraint. This conclusion is intuitive: Instead of relaxing just one borrowing constraint, a balloon mortgage has the same effect as relaxing the constraints over the whole life of the mortgage relative to a fixed rate mortgage, with the exception of the last period when the entire principal is due. In general, the authors show that mortgage contracts with an increasing payment schedule and thus lower γ over time tend to encourage more young households to become homeowners but decrease homeownership among older households. This finding makes perfect sense: Young households experience an increasing earnings profile (see the first two panels of Figure 2), which makes a mortgage with an increasing payment schedule a good match for them. The opposite is true for middle-aged and older households whose income, at least on average, is stagnating or even decreasing, making a mortgage with an increasing payment schedule a mismatch to their earnings profile.[12]

Chambers, Garriga, and Schlagenhauf have therefore pointed to the great importance of the borrowing constraint in determining tenure decisions over the life cycle. Relaxing this constraint can greatly increase homeownership among younger households because it reduces the subjective user cost in equation (4) if $\mu > 0$. This effect occurs both for the borrowing constraint in the initial period when a household purchases a house as well as over the duration of the mortgage.

The Effect of Subsidies

The papers discussed so far clearly specify a dual role of housing—namely, a consumption role (h enters the utility function) and an investment role because housing is also a form of saving as long as the depreciation rate is below 100 percent. One

12. Notice that in the model the authors assume that only one single mortgage contract exists because adding multiple contracts into the same model involves too big a computational burden. This result should not interpreted as an argument against, say, balloon mortgages because they reduce ownership among older households. In reality, several types of mortgages are available, so older households can choose mortgages that suit them better than a balloon mortgage.

ingredient missing in the literature so far has been the fact that housing investment is risky. Jeske and Krueger (2005) set up a model in which, in addition to idiosyncratic labor income risk, households face an idiosyncratic house depreciation shock. This setup generates some realistic features: Most importantly, households may end up with negative equity on their home if the depreciation shock is large enough. This feature is one way to generate mortgage foreclosure, which has not been done before in the general equilibrium literature. Jeske and Krueger also assume that there is a deadweight loss from foreclosure, taking the following shape: If a property goes into foreclosure, the bank receives proportionally less than the value of the property. In reality, this foreclosure loss is substantial. Pennington-Cross (2004) estimates the deadweight loss of foreclosure to be 22 percent of the value of the property. This loss becomes crucial in the policy experiment later because a subsidy on mortgage rates will likely increase mortgage default as well as the deadweight loss.

A main finding of Jeske and Krueger is that interest subsidies, if they are indeed passed on to households, can cause overinvestment in housing and larger mortgages.

In addition, Jeske and Krueger allow households to choose any leverage ratio they desire. This feature differs from the papers studied so far, in which households face a sharp borrowing constraint that requires a minimum or even fixed down payment ratio. Jeske and Krueger assume that the interest rate households pay is a function that increases in the leverage ratio they are taking on. Specifically, a bank will price the foreclosure risk into the mortgage, and, if a household chooses higher leverage, a smaller depreciation shock will be required to trigger a default—that is, default is more likely—implying a higher interest rate. Jeske and Krueger must also make some simplifying assumptions. As opposed to the previously discussed papers, they do not consider life-cycle effects—that is, they assume that households are infinitely lived, and they abstract from transaction costs. Both assumptions are necessary to make the computational burden tractable.

Suppose that a government were to provide a subsidy that lowers mortgage interest. Jeske and Krueger try to study the macroeconomic effects of such a subsidy. The subsidy the authors have in mind has a counterpart in reality—government-sponsored enterprises (GSEs) like Fannie Mae and Freddie Mac, which are in the business of borrowing in the bond market to then purchase large portfolios of home mortgages. The GSEs receive a benefit from the federal government in the form of an implicit bailout guarantee.[13] The Congressional Budget Office (CBO) estimates that the subsidy is worth about 42 basis points; that is, GSEs can borrow at rates 0.42 percentage points below what other entities with a similar credit rating would have to pay in the bond market. Passmore, Sherlund, and Burgess (2004) estimate that only a small fraction—7 basis points—actually makes it to the homeowner while the GSEs keep the lion's share of the subsidy to pass on to their shareholders. Blinder (2004), on the other hand, defends the GSEs by estimating that they indeed reduce mortgage rates by almost the entire 42 basis points.

Jeske and Krueger (2005) do not attempt to judge which of the two studies is correct but rather do the following thought experiment: Supposing that the GSEs indeed pass on the entire 42 basis points to homeowners as they claim, then what is the effect on aggregate macroeconomic variables as well as the distribution of wealth and income?

A main finding of Jeske and Krueger is that interest subsidies, if they are indeed passed on to households, can cause overinvestment in housing and larger mortgages.

In fact, mortgage debt increases proportionally more than the housing stock, implying more leverage on housing. This finding, of course, is less than surprising: The subsidy makes housing cheaper and therefore causes more housing investment, and lower mortgage rates encourage more leverage. Thanks to the endogenous borrowing constraint, Jeske and Krueger are able to study one new aspect of housing policy, namely, how mortgage subsidies affect the aggregate level of mortgage default. They find that the subsidy increases the proportion of mortgages in default by one half, which in turn makes the deadweight loss 50 percent larger.

The next provocative question is, What are the distributional consequences of the subsidy? The subsidy has little effect on homeownership rates. In other words, the assistance goes mainly to existing homeowners—those households already well off due to high incomes and high net worth—which are able to afford even larger houses while poor households are almost unaffected in their tenure choice.

This result is consistent with the findings of the generic housing model outlined earlier, especially equation (4): Lowering interest rate r and thus the user cost for housing $(r + \delta^h)/(1 + r)$ may have only a small effect on households that are constrained by the down payment condition.[14] The main part of their subjective user cost, after all, comes from the borrowing constraint (the second term on the right-hand side of equation [4]) lowering current consumption. In the Jeske and Krueger model, another reason why poor households will not necessarily invest in real estate and instead keep renting is that the rental rate is determined endogenously; that is, the rental rate is a market-clearing price ensuring that rental demand equals rental supply. With a mortgage subsidy that causes more investment in housing, the rental price will necessarily drop, making renting more attractive relative to housing. These results may disappoint the proponents of government assistance for mortgage financing. Either the subsidy in the form of implicit guarantees and lower interest rates goes to shareholders of GSEs or the interest rate reduction is passed on to homeowners but ends up in the hands of already well-off households while doing nothing to promote homeownership for poorer households.

The findings, by the way, are consistent with empirical work that suggests that mortgage rate reductions have only a marginal impact on homeownership rates while the borrowing constraint seems to play a much larger role.[15] Jeske and Krueger confirm these findings but are able to do so in a logical and coherent framework using general equilibrium instead of microsimulation techniques.

Conclusion

Households in different age cohorts and different positions in the income and wealth distribution have very different homeownership ratios, real estate wealth, and mortgage debt. Likewise, renters and homeowners look very different in their path of income and asset holdings over the life cycle. This heterogeneity poses a challenge to researchers trying to answer questions regarding housing policy because one cannot study the effect of a policy merely on the average household. Studying housing in the framework of macroeconomic models therefore requires incorporating heterogeneity of households. Likewise, the size of the housing stock is too large to ignore general equilibrium effects.

---

13. See Frame and Wall (2002) for a survey on this topic.
14. One must use caution here: The interest rate itself depends on the leverage ratio, so households will not view it as exogenous as they did in the other three papers. The intuition, though, would be similar in a model with an interest rate as a function of the leverage ratio.
15. See Painter and Redfearn (2002) or Feldman (2001) for a summary of a variety of empirical studies.

Thanks to advances in computational techniques and faster computer hardware, economists are now able to write down and solve models detailed enough to accommodate housing. The conference hosted by the Federal Reserve Bank of Atlanta featured (among others) four papers, all of which are variations of the generic housing model described here even though they touch on a large variety of issues. Yang (2005) points out the importance of transaction costs and borrowing constraints in life-cycle models with housing. Li and Yao (2005) study the effect of house price changes. Chambers, Garriga, and Schlagenhauf (2005b) incorporate financial innovations in the form of different mortgage contracts. And Jeske and Krueger (2005) study the effect of a subsidy on mortgage finance. In each paper, one can deduce the main intuition from the simple equilibrium conditions.

REFERENCES

Aiyagari, S. Rao. 1994. Uninsured idiosyncratic risk and aggregate saving. *Quarterly Journal of Economics* 109, no. 3:659–84.

Attanasio, Orazio, and Martin Browning. 1995. Consumption over the life cycle and over the business cycle. *American Economic Review* 85, no. 5:1118–37.

Attanasio, Orazio, and Guglielmo Weber. 1995. Is consumption growth consistent with intertemporal optimization? Evidence from the Consumer Expenditure Survey. *Journal of Political Economy* 103, no. 6:1121–57.

Blinder, Alan S. 2004. Letter on Federal Reserve working paper to CEO Raines from Dr. Alan S. Blinder, February 23. <www.fanniemae.com/commentary/pdf/022304.pdf> (September 13, 2005).

Blundell, Richard, Martin Browning, and Costas Meghir. 1994. Consumer demand and the life-cycle allocation of household expenditures. *Review of Economic Studies* 61, no. 206:57–80.

Board of Governors of the Federal Reserve System. 2005. Flow of funds statistics. June 9 release. <www.federalreserve.gov/releases/z1/> (September 13, 2005).

Chambers, Matthew, Carlos Garriga, and Donald Schlagenhauf. 2005a. Accounting for changes in the homeownership rate. Florida State University working paper.

———. 2005b. Mortgage contracts and housing tenure-duration decision. Paper presented at the Federal Reserve Bank of Atlanta conference “Housing, Mortgage Finance, and the Macroeconomy,” May 19–20.

Feldman, Ron J. 2001. Mortgage rates, homeownership rates, and government-sponsored enterprises. Federal Reserve Bank of Minneapolis annual report, 3–23.

Fernandez-Villaverde, Jesus, and Dirk Krueger. 2002. Consumption and saving over the life cycle: How important are consumer durables? Proceedings of the 2002 North American summer meetings of the Econometric Society.

———. 2004. Consumption over the life cycle: Some facts from Consumer Expenditure Survey data. University of Pennsylvania working paper.

Frame, W. Scott, and Larry D. Wall. 2002. Financing housing through government-sponsored enterprises. Federal Reserve Bank of Atlanta *Economic Review* 87, no. 1:29–43.

Gourinchas, Pierre-Olivier, and Jonathan Parker. 2002. Consumption over the life cycle. *Econometrica* 70, no. 1:47–89.

Jeske, Karsten, and Dirk Krueger. 2005. Housing and the macroeconomy: The role of implicit guarantees for government sponsored enterprises. Paper presented at the Federal Reserve Bank of Atlanta conference “Housing, Mortgage Finance, and the Macroeconomy,” May 19–20.

Li, Wenli, and Rui Yao. 2005. The life-cycle effects of house price changes. Paper presented at the Federal Reserve Bank of Atlanta conference “Housing, Mortgage Finance, and the Macroeconomy,” May 19–20.

Painter, Gary, and Christian L. Redfearn. 2002. The role of interest rates in influencing long-run homeownership rates. *Journal of Real Estate Finance and Economics* 25, nos. 2–3:243–67.

Passmore, Wayne, Shane Sherlund, and Gillian Burgess. 2004. The effect of housing government-sponsored enterprises on mortgage rates. Federal Reserve Board Finance and Discussion Series No. 2005-6.

Pennington-Cross, Anthony. 2004. The value of foreclosed property: House prices, foreclosure laws, and appraisals. Office of Federal Housing Enterprise Oversight Working Paper 04-1, January.

Yang, Fang. 2005. Consumption along the life-cycle: How different is housing? Paper presented at the Federal Reserve Bank of Atlanta conference “Housing, Mortgage Finance, and the Macroeconomy,” May 19–20.