

NBER WORKING PAPER SERIES

DO HOUSEHOLDS BENEFIT FROM FINANCIAL DEREGULATION AND INNOVATION?
THE CASE OF THE MORTGAGE MARKET

Kristopher Gerardi
Harvey S. Rosen
Paul Willen

Working Paper 12967
<http://www.nber.org/papers/w12967>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2007

The authors thank seminar audiences at Boston University, Boston College, the Federal Reserve Banks of Boston and San Francisco, the 2006 Applied Micro System Meeting, the Fall 2006 Macro System Meeting and Princeton's Public Finance working group for useful comments. We are especially grateful to Mark Aguiar, Alan Blinder, Mark Doms, Karen Dynan, Wendy Edelberg, Jean Grossman, Bob King, Yolanda Kodrzycki, Wayne Passmore, Julio Rotemberg, Scott Schuh, Frank Vannerson and Mark Watson for helpful suggestions. Finally, we thank Jeff Fuhrer for providing us with computer programs. This research has been supported by Princeton's Center for Economic Policy Studies. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

© 2007 by Kristopher Gerardi, Harvey S. Rosen, and Paul Willen. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Do Households Benefit from Financial Deregulation and Innovation? The Case of the Mortgage Market

Kristopher Gerardi, Harvey S. Rosen, and Paul Willen

NBER Working Paper No. 12967

March 2007

JEL No. D14,G21,H89

ABSTRACT

The U.S. mortgage market has experienced phenomenal change over the last 35 years. This paper develops and implements a technique for assessing the impact of changes in the mortgage market on households. Our framework, which is based on the permanent income hypothesis, that allows us to gauge the importance of borrowing constraints by estimating the empirical relationship between the value of a household's home purchase and its future income. We find that over the past several decades, housing markets have become less imperfect in the sense that households are now more able to buy homes whose values are consistent with their long-term income prospects. One issue that has received particular attention is the role that the housing Government Sponsored Enterprises (GSEs), Fannie Mae and Freddie Mac, have played in improving the market for housing finance. We find no evidence that the GSEs' activities have contributed to this phenomenon. This is true whether we look at all homebuyers, or at subsamples of the population whom we might expect to benefit particularly from GSE activity, such as low-income households and first-time homebuyers.

Kristopher Gerardi
Federal Reserve Bank of Boston
Boston, MA 02205
Kristopher.Gerardi@bos.frb.org

Harvey S. Rosen
Department of Economics
Fisher Hall
Princeton University
Princeton, NJ 08544-1021
and NBER
HSR@princeton.edu

Paul Willen
Research Department
Federal Reserve Bank of Boston
P.O. Box 55882
Boston, MA 02205
and NBER
willen968@yahoo.com

1. Introduction

The U.S. mortgage market has experienced phenomenal change over the last 35 years. Gone are the days when most households got a cookie-cutter, 30-year, fixed-rate, level-payment mortgage from a savings and loan that took deposits at 3 percent and lent out at 6 percent. And gone are the days when the typical lender held that mortgage on its books until the maturity of the loan. Today, consumers choose from an extensive menu of mortgages offering flexibility along almost every dimension. Furthermore, most lenders hold the mortgage for a very short time; typically, they sell the mortgage on the secondary market and, more often than not, financial intermediaries then pool that mortgage with many other mortgages and sell the cash flow in a complex security called a collateralized mortgage obligation.

As noted below, many researchers have argued that this transformation has enhanced efficiency by integrating the mortgage market with the broader financial markets. But there has been comparatively little research investigating how or even whether the transformation of mortgage markets has directly benefited the average homeowner. This question is obviously pertinent to both government agencies and consumer advocacy groups that have significant concerns about housing finance at the family level. However, it may also be of interest to a much wider group of researchers and policymakers who are concerned with monetary and fiscal policy. For example, some macroeconomists argue that the well-documented decrease in business-cycle volatility can be partially attributed to the increasing ease with which households can obtain mortgage financing as well as access their existing home equity. Stock and Watson (2003) state, “One explanation for the decreased volatility in residential, but not nonresidential, construction is the increased ability of individuals to obtain nonthrift mortgage financing...” Indeed, Dynan, Elmendorf, and Sichel (2005) have documented that financial innovation more generally played an important role in the reduction of the volatility of economic activity that began in the mid 1980s.

In this paper, we use a novel technique to assess the impact of changes in the mortgage market on individual households. Our analysis starts with an implication of the permanent income hypothesis: that the higher a household's expected future income, the more it desires to spend and consume, *ceteris paribus*. If perfect credit markets exist, desired consumption matches actual consumption and current spending should forecast future income. Since credit market

imperfections mute this effect, we view the strength of the relationship between house spending and future income as a measure of the “imperfectness” of mortgage markets.

We apply this methodology using household-level data from the Panel Study of Income Dynamics (PSID) and consider the forecasting ability of a newly purchased home on future income. We find that households buying bigger houses have higher future incomes, all else being equal. Furthermore, we find that the forecasting relationship of housing purchases and future income has changed over time. The estimated sensitivity of future income to current housing expenditures more than doubled over the length of our sample, from 1969 to 1999. Further, the change was not smooth over time—the application of econometric techniques for locating unknown structural breaks suggests that the relationship changed discretely in the mid 1980s. Following the logic of the previous paragraph, we view the increased sensitivity of future income to house spending as evidence that mortgage markets have become less imperfect over time.

With this finding in hand, we set out to accomplish two goals. First, as success has many fathers, we attempt to establish paternity of the improved system of housing finance. Was it deregulation? Was it the creation of a secondary market? Or was it the activities of the Government Sponsored Enterprises (GSEs), Fannie Mae and Freddie Mac, whose mission is to help families realize the “American dream” of owning a home? Assessment of paternity presents particular difficulties for us, as many of the changes in mortgage markets occurred simultaneously, during the period from 1977 to 1983. To skirt these problems, we focus on the period after 1983 when mortgage market deregulation had mostly run its course and secondary markets were in place, but during which activities by the GSEs fluctuated considerably. We find that the activities of the GSEs had little or no impact on the imperfectness of mortgage markets.

Our second task is to assess the consequences of the new mortgage market for various groups in the population. Advocates for the GSEs, in particular, argue that the transformation of the mortgage market has improved opportunities for the less-well-off and for first-time homebuyers, claiming that the more efficient allocation of risk has enabled lenders to lend to more-marginal borrowers. While we do in fact find evidence of relaxed credit constraints for marginal borrowers, including first-time homebuyers, younger households, and financially impoverished households, we are not able to attribute these improvements directly to GSE activities.

In Section 2, we provide a brief review of the transformation of the mortgage market from 1970 to 1999, the period that we study. Section 3 presents a simple model that illustrates the possible effects of changes in borrowing constraints on the relationship between spending on housing and future income. Section 4 then describes our empirical specification and the data. In Section 5 we present results for our main sample of home purchasers from the PSID as well as results for various subsamples. Section 6 includes an assessment of the sensitivity of our findings to alternative specifications of the model. Section 7 investigates the role of the Government Sponsored Enterprises and is followed by a brief conclusion.

2. The Mortgage Market 1970–1999

This section briefly reviews developments in the U.S. mortgage market from 1970 to 1999. We look first at the deregulation and securitization of the mortgage market, next at the development of new mortgage designs, and then at other developments, including the emergence of the subprime market and the introduction of anti-discriminatory legislation. We conclude the section by reviewing earlier research on innovations in mortgage markets.

2.1 Deregulation and securitization

From the Depression through the late 1970s, deposits in savings accounts provided almost all financing for home loans. Depression-era regulations, updated at various points, channeled low-cost deposits to the thrift industry (Mason 2004). Regulations took many forms and included usury ceilings, interstate banking prohibitions, limits on branching, and perhaps most infamously, Regulation Q, which capped deposit rates and forbade banks from paying interest on checking deposits (England 1992, Gilbert 1986). While these regulations provided some stability, they also, predictably, led to major inefficiencies. Most significantly, by making bank deposits the principal source of funds for mortgages, regulators forced lenders to finance long-term assets with short-term liabilities, a situation referred to in the industry as the mismatch problem (Modigliani and Lessard 1975).

Despite its flaws, the system basically worked until the mid 1960s. Low inflation and stable interest rates meant that the usury ceilings, interest-rate caps, and the mismatch problem did not generate major difficulties. However, starting in the mid 1960s, inflation and interest rates rose, driving up the cost of funds for the savings and loan industry. This forced Congress to

act, beginning a process that would culminate, 20 years later, in the transformation of the U.S. mortgage market from a largely deposit-financed system to a largely capital-markets-financed system.

In 1968, Congress moved aggressively to develop a secondary market for mortgages, that is, a market in which banks could sell mortgages they had originated to other investors. It took an old government agency charged with creating a secondary market, the Federal National Mortgage Association (FNMA, founded in 1938 and now known as Fannie Mae) and divided it into two separate entities. The first was a government agency called GNMA (Government National Mortgage Association, later known as Ginnie Mae), which bought mortgages guaranteed by the Federal Housing Authority and Veterans Administration, and the second was a shareholder-owned but government-sponsored enterprise still called FNMA, which bought other mortgages. The Federal Home Loan Bank Board (FHLBB) created the Federal Home Loan Mortgage Corporation (later known as Freddie Mac) in 1970, with a mandate to buy loans from members of the Federal Home Loan Bank system. This period also saw the emergence of the mortgage-backed security, a bond whose cash flows are backed by homeowners' mortgage payments. Ginnie Mae issued its first security of this kind in 1970, and Fannie Mae and Freddie Mac followed shortly thereafter (Bartlett 1989). Within a couple of decades, this innovation would transform the industry.

Despite these moves by Congress, the problems continued to mount. Boston Fed President Frank Morris described efforts to stabilize mortgage finance over the 1966–1975 period as a “decade of failure” (Morris 1975). Secondary markets were slow to develop; deposits into banks and savings and loans remained the chief source of funds for home mortgage lending. Secondary markets seem an obvious solution to the mismatch problem now, but they did not appear that way to contemporary observers, even to brilliant economists. Franco Modigliani assumed that the solution to the mismatch problem lay in mortgage design, and not one participant at a conference he organized in 1975 to explore the issue made any mention of secondary markets in the accompanying conference volume (Modigliani and Lessard 1975). The first privately issued mortgage-backed security appeared in 1977 and was generally considered a failure (Ranieri 1996).

Continued instability and high interest rates in the late 1970s initiated the final phase of the reinvention of housing finance in America. In 1977, Merrill Lynch invented the Cash

Management Account, in effect allowing non-banks to circumvent Regulation Q (Cocheo 2003). This innovation, combined with many others, severely reduced the availability of funds for the thrift industry, which was still bound by Regulation Q. Even when regulators finally allowed them to pay competitive interest rates, thrifts confronted state usury laws. These laws often meant that the thrifts could not lend profitably, and so they simply stopped lending altogether (Shaman 1979).

The impending collapse of the thrift industry spurred Congress and regulators to action, and over the next six years, legal and regulatory changes transformed mortgage lending (Bartlett 1989). The Depository Institutions Deregulation and Monetary Control Act of 1980 ordered the phase-out of Regulation Q over the next six years and overrode or pre-empted state usury ceilings. In 1982, Congress passed the Garn-St. Germain Depository Institutions Act, which extended the 1980 act, pre-empting state laws that constrained the types of mortgage products originators could offer. In 1984, the Secondary Mortgage Market Enhancement Act solved many of the technical problems facing mortgage-backed securities. Finally, the Tax Reform Act of 1986 created an investment vehicle called a REMIC (Real Estate Mortgage Investment Conduit) that allowed securitizers to divide the flows of principal and interest from mortgage-backed securities into different classes (“tranche securities”) tailored to the needs of different investors.

Two key events dramatically accelerated the development of a secondary market. The first occurred in October 1981, when the FHLBB, the main regulator of thrifts, introduced a change in accounting rules that had the effect of allowing lenders to sell mortgages on the secondary market without booking a large accounting loss (Mason 2004). This change created a liquid secondary market for mortgages virtually overnight (Lewis 1989). Secondary market sales of mortgages increased more than four-fold, from \$12 billion in 1981 to \$52 billion in 1982 (Bartlett 1989).

The second influential event was the realization that issuers could skirt many of the problems that had bedeviled early mortgage-backed securities, by enlisting one of the Government Sponsored Enterprises (Ranieri 1994). The GSEs’ federal charters meant, for example, that their securities were exempt from state investor protection laws. In addition, investors believed, perhaps erroneously, that Fannie Mae and Freddie Mac securities were backed by the full faith and credit of the federal government. Freddie Mac, initially, and later

Fannie Mae, worked closely with Wall Street firms and became the largest issuers of mortgage-backed securities.

Figure 1 displays the evolution of GSE securitization activity from 1970 to 2003. The figure shows a direct measure, the value of mortgage-backed securities (MBS) securitized by Fannie and Freddie. The top panel of Figure 1 (solid line) displays the stock of MBS normalized by total home mortgage debt outstanding, while the bottom panel (solid line) displays the corresponding flow normalized by originations of home mortgages. These ratios have grown substantially over time. The percentage of the stock of MBS securitized by the GSEs increased dramatically from approximately 0 in 1975 to almost 30 percent by 2003. In addition to securitizing mortgages, the GSEs issue debt to purchase mortgages and mortgage-based securities. This “retained portfolio” of mortgages and mortgage-related securities held internally by the GSEs, whose time path is also displayed in Figure 1, has increased enormously over time as well. We discuss issues related to the retained portfolio below.

2.2 Innovation in mortgage design

The menu of available mortgage choices in 1999 vastly exceeded the options that were available in the 1970s. In the 1970s, because of a combination of regulation and inertia, the mortgages available to borrowers consisted almost exclusively of fixed-rate, level-payment instruments. Among other things, consumer groups, as they do to this day, viewed features like variable interest rates as dangerous, and they worked assiduously to prevent adoption (Guttentag 1984). Even when they allowed variable-rate mortgages, regulators established restrictions that severely limited their usefulness. For example, the FHLBB allowed variable-rate mortgages, but the rate could not change by more than 50 basis points every six months, nor could it rise by more than 2.5 percentage points over the life of the loan (Macauley 1980).

However, the 30-year, fixed-rate mortgage was particularly unsuitable for the high-inflation, high-volatility environment of the 1970s. The combination of high nominal interest rates and fixed payments over time meant that households faced very high real payments early in the life of the loan. Some borrowers would have been better served by mortgage designs that allowed relatively lower nominal payments early in the life of the loan, so that the real burden would be distributed more evenly over time. Regulators eventually relented in their opposition to alternative mortgage designs and allowed lenders to offer innovations, including the Graduated-

Payment Mortgage (GPM, first offered in 1977) and the forerunner of today's Option ARM (first offered in 1980), both of which allow borrowers to make a monthly payment that falls short of the interest due on the mortgage (Phalan 1977, Harrigan 1981). These alternative mortgage designs have been popular since the 1980s, and by no means confined to very financially sophisticated households. For example, the Option ARM accounted for a significant share of mortgages in California during the 1980s (Kettell, 2006). In 1996, so-called COFI ARMS (a common term for Option ARMs originated in California) accounted for approximately one half of adjustable-rate mortgages outstanding in California as well as a significant portion of originations (Stahl, 1996).

Initially, new mortgage designs merely inoculated borrowers against high inflation. With high inflation, a graduated-payment mortgage offered a flow of real payments comparable to that of a traditional mortgage with low inflation. However, although the appearance of high inflation enlarged the mortgage menu, its disappearance in the mid 1980s did not shrink the menu. Regulators made no effort to prohibit the alternative mortgages developed in the late 1970s and early 1980s. The use of these new products in a low-inflationary environment allowed lenders to offer borrowers much less rigorous repayment schedules than had prevailed even in the 1950s and 1960s. In other words, the net effect of regulators' responses to high inflation was to liberalize mortgage markets considerably when compared to the traditional system of the 1950s and 1960s.

2.3 Other changes

Three other changes in mortgage markets in this period are worth noting. First, concerns emerged in the early 1970s of race and gender-based discrimination in mortgage markets, leading to the passage of the Equal Credit Opportunity Act (ECOA). The original ECOA, passed in 1974, prohibited discrimination on the basis of sex and marital status. In 1976, Congress substantially expanded the law, adding age, race, color, religion, and national origin to the list of factors on which lenders could not discriminate (Elliehausen and Durkin 1989). Subsequent research on the impact of ECOA has yielded mixed conclusions with respect to its effects (see, for example, Ladd 1982 and Munnell *et al.* 1996).

Second, in the mid 1990s, lenders adopted automated underwriting procedures, which reduced the cost of originating new mortgages (Straka 2000). This change was driven, at least

partly, by two things. The successful experience of credit-card issuers with numerical credit scores allowed lenders to substitute these scores for loan officer judgment in the analysis of the creditworthiness of the borrower. Also, the anonymity of automated underwriting procedures allowed lenders to refute claims of racial discrimination more easily.

Third, the subprime market, the part of the mortgage business dedicated to borrowers with less-than-perfect credit histories, emerged in the mid 1990s, following the development of credit scoring. Until the mid 1990s, a borrower was either prime and got a loan at the going rate, or was subprime and did not get a loan at all (Munnell et al. 1996). In the mid 1990s, a new generation of lenders began to offer loans to subprime borrowers, but they demanded much higher interest rates as compensation for the added risk. Subprime originations grew from \$65 billion in 1995 to \$332 billion in 2003 (Chomsisengphet and Pennington-Cross 2006).

2.4. Prior research on innovations in the mortgage market

What was the impact of all these institutional changes in the mortgage market? Empirical researchers have approached this question in two ways. The first examines the extent to which mortgage markets and capital markets have become integrated over time. The second focuses on the role that new activities of the GSEs and various institutional changes have played in these developments. We now discuss them in turn.

One way to investigate the extent to which mortgage markets and capital markets are integrated is to look at the time-series relationship between interest rates on mortgages and Treasury yields. The idea is that if capital can flow freely in and out of the mortgage market, then Treasury yields and mortgage-market yields should move together over time. A variety of studies using this general approach have found that, in fact, the correlation between Treasury yields and mortgage yields was greater in the 1980s than in the 1970s. (See, for example, Devaney and Pickerill 1990, Hendershott and Van Order 1989, Goebal and Ma 1993, and Devaney, Pickerill, and Krause 1992).¹

An alternative approach is to examine mortgage markets across regions. The idea here is that in a well-functioning mortgage market, regional conditions should reflect credit availability in the national capital market rather than in a particular region. Of course, other characteristics

¹ Statistical models based upon the Arbitrage Pricing Theory from finance suggest that integration of the mortgage market with traditional capital markets increased during the 1980s. See Bubnys, Khaksari, and Tarimcilar (1993).

that might affect mortgage rates and that vary systematically across regions must also be taken into account. Rudolph and Griffin (1997) found that the coefficient of variation of mortgage rates across Metropolitan Statistical Areas decreased from 1963 to 1993, a finding that is consistent with more integration over time.² In the same spirit, in a well-functioning mortgage market, the terms of a family's mortgage should be independent of the particular institution that originates the mortgage. Loutskina and Strahan (2006) show that, at least in certain segments of the mortgage market, this has in fact become the case.³

In short, then, both strains of the empirical literature point in the same direction—over time, the mortgage market has become more integrated with national capital markets, and from this perspective, the mortgage market has become a better-functioning market.

While the improvement in the operation of housing finance markets has been well-documented, the contribution of the GSEs' securitization activities has received little attention.⁴ Interestingly, several papers run counter to the widely held view that securitization played a major role. As we mentioned above, some securitization occurred in the 1970s, but Rudolph, Zumpano, and Karson (1982) found that regional variation in contract rates did not decrease over this period and concluded that “the secondary market has not significantly reduced differences in local markets.” In this context, a paper by Goebel and Ma (1993) is of more relevance, because their sample period encompasses the 1980s, which, as noted above, witnessed the huge expansion of Fannie's and Freddie's securitization activities. Goebel and Ma's vector autoregression analysis of the relationship between Treasury and mortgage-market yields suggests that “the two markets were already integrated before the full development of the secondary mortgage markets between 1984 and 1987.”

An important aspect of all the studies discussed above is the centrality of the relationship between Treasury rates and mortgage rates. While important and interesting, this tack removes the focus from where we believe it really belongs, which is the impact of securitization on

² Rudolph, Zumpano, and Karson (1982) examine how a variety of attributes of mortgage contracts (contract rate, loan initiation fees, maturity, and loan-to-value ratio) varied by Standard Metropolitan Statistical Area between 1968 and 1978. They find that while the variability of fees and charges and loan maturity declined over this period, the contract rate and loan-to-value ratio did not change. This period predates both the big increase in GSE securitization activity and other important changes in housing finance institutions.

³ They highlight the role of securitization, arguing that credit supply evolves independently of conditions at particular banks in the highly securitized “conforming” loan market (the part of the market for which the GSEs are permitted to buy up mortgages).

⁴ In contrast, there has been a great deal of attention focused on the impact of the GSEs on the level of mortgage rates. See, for example, Ambrose and Thibodeau (2004) and Passmore, Sparks, and Ingpen (2002).

households. What we care about ultimately is not the correlation between Treasury rates and mortgage rates *per se*. Rather, the key question is whether securitization (or any other development in the housing finance market) enhanced the likelihood that households could borrow enough to buy a home that maximized their utility, given their lifetime income.⁵ To address this question, we develop a model that allows us to investigate the impact of mortgage-market innovations at the household level.

3. Model

In this section, we develop a model of how changes in the market for housing finance affect individual households. We first present a simple model with two types of households and two types of houses. A key result is that relaxing constraints on mortgage lending increases the coefficient in a regression of future income on housing expenditures.⁶ This particular implication of the model is significant because it provides us with an interpretation of the results from our empirical analysis below. We then present a necessary condition for obtaining such a result in a more general model with any number of house types and household types. We also show that, contrary to what one might guess, relaxing borrowing constraints also increases the coefficient in a regression of house spending on future income.

Suppose we have two types of families who differ in future income, low (Y_L^f) or high (Y_H^f), and who face a choice between two types of house, small (H_S) or big (H_B).⁷ Assume further that households have different levels of wealth and that they can borrow to finance their home purchase. Otherwise, the families are identical; in particular, they have the same level of *current* income, so that we can use “high future income” and “high future income growth” interchangeably. For purposes of this example, we assume that, if allowed to borrow an

⁵ Linneman and Wachter (1989) is the only study to our knowledge that attempts to assess the impacts of mortgage-market developments at the household level. They find a diminished impact of borrowing constraints on tenure choice over time, and they attribute their findings to the development of adjustable rate mortgages (ARMs) and the increased use of seller financing and other “non-traditional” financing schemes. Campbell and Cocco (2003) examine the effects of different mortgage designs in a dynamic life-cycle model with borrowing constraints.

⁶ See Artle and Varaiya (1978) for the first theoretical analysis of the implications of borrowing constraints on homeownership, Brueckner (1986) for a 2-period version of the same model in discrete time, and Engelhardt (1996) for an empirical implementation. For a discussion of the user cost of housing in a model without borrowing constraints, see Himmelberg, Mayer, and Sinai (2005).

⁷ This discussion implicitly assumes that housing is entirely a consumption decision. In fact, housing is also purchased for investment purposes. However, under the conditions we assume, the basic results are independent of whether the motive is consumption or investment.

unlimited amount, the low-income-growth families would choose small houses, while the high-income-growth families would choose big houses. However, in the presence of credit constraints, this separation of types might not emerge. For example, lenders typically require that monthly housing expenditures fall below a certain percentage of current income. Thus, current income fixes some maximal amount the household can borrow. This need not limit the size of the house the family can buy if the family has access to sufficient other assets such as its own wealth or the wealth of its close relatives. If not, even a high-income-growth household ends up buying a small house.

In this example, credit constraints reduce the difference in average, observed income-growth rates between small- and big-house buyers. To see why, recall that under our assumptions, in the absence of credit-market constraints, every high-income-growth family purchases a big house and every low-income-growth family purchases a small house. In the presence of constraints, some high-income-growth families instead purchase small houses, and thereby drive up the average income growth associated with small-house buyers.⁸ On the other hand, the income growth of large-house buyers stays the same. Hence, the borrowing constraints attenuate the observed relationship between income growth and the size of current home purchases.

This argument is laid out more formally in Figure 2. Suppose initially that $2/3$ of the high-income-growth families buy a small house and that all of the low-income-growth families also buy a small house, so that the average future income of the small-house buyers is $Y(H_S) = \frac{2}{5}Y_H^f + \frac{3}{5}Y_L^f$. On the other hand, only high-income-growth families buy a big house, implying that $Y(H_B) = Y_H^f$. Now, suppose we relax the constraint so that only $1/3$ of the high-income-growth families buy a small house. In this case, the average future income of small-house buyers falls to $Y(H_S) = \frac{1}{4}Y_H^f + \frac{3}{4}Y_L^f$, whereas the future income of big-house buyers stays the same. Figure 2 illustrates how the movement of high-income-growth families from small houses to big houses as the credit market constraint is relaxed raises the sensitivity of average

⁸ Holding current income and wealth constant, households with higher expected future income (Y_H^f) are more likely to be constrained than those with lower expected future income (Y_L^f). This is because high-income-growth households would like to borrow more than low-income-growth households in order to smooth consumption and consume more today, but their borrowing is limited by current income.

future income to house size—that is, line AC is flatter than line BC. In an econometric context, this means that if we estimate a regression of future income on the value of current house purchases, then the less constrained the borrowing environment, the greater the coefficient on the value of current house purchases, *ceteris paribus*. As shown below, this is precisely what we find in the data. We present a formal, algebraic version of this argument in Appendix B.⁹

This simple model also illustrates that if we choose instead to estimate a regression of housing purchase expenditures on expected future income, we should also expect to obtain an increase in the sensitivity of the size of home purchases to future income when borrowing constraints are relaxed, *ceteris paribus*. The argument is laid out in Figure 3, and is virtually identical to the argument described in Figure 2. Assuming initially, as we did in Figure 2, that 2/3 of the high-income-growth families buy a small house, implies that the average house size of high-income-growth households is $H(Y_H^f) = \frac{2}{3}H_S + \frac{1}{3}H_B$. Since we assume that all low-income-growth households purchase small houses, we have $H(Y_L^f) = H_S$. Again, suppose we relax the constraint so that only 1/3 of the high-income-growth families buy a small house. In this case, the house size of high-income-growth families increases to $H(Y_H^f) = \frac{1}{3}H_S + \frac{2}{3}H_B$, whereas the house size of low-income-growth families stays the same. As Figure 3 illustrates, this results in an increased slope of the line connecting the two points corresponding to the average house sizes of low-income-growth and high-income-growth households, respectively. In a regression context, this implies that the less constrained the borrowing environment, the greater the responsiveness of house value to expected future income, *ceteris paribus*. In fact, when we estimate such a regression we obtain results consistent with Figure 3 (see Section 6.4).

The two-type model is intuitively appealing because of its simplicity. However, one may wonder whether the conclusion continues to hold in a model with more types of individuals and houses. In fact, we can provide a very general expression that links changes in borrowing constraints to changes in regression coefficients. This expression provides us with a necessary

⁹ Our simple model assumes that households with a given level of income growth face the same borrowing constraints; they differ only by wealth. However, the analysis extends to any variable that affects a family's ability to buy a house, not just wealth. In particular, one could show that when families have the same wealth but differ in the borrowing limits they face, then if individuals' borrowing limits are all relaxed, the relationship between current house value and expected future income strengthens. This case is particularly important given the evidence in Munnell et al. (1996).

condition for obtaining an increasing regression coefficient in response to the relaxation of borrowing constraints.

Denote the least-squares estimate from a regression of future income on the value of current house purchases as $\hat{\beta}$, and δ as a credit-market constraint (for example, a debt-to-income constraint). We want to sign the impact on $\hat{\beta}$ of relaxing the borrowing constraint, which corresponds to an increase in δ . In Appendix B, we derive the following expression for the partial derivative of $\hat{\beta}$ with respect to δ :

$$\frac{\partial \hat{\beta}}{\partial \delta} = \frac{1}{\text{var}(H)} \left\{ \text{cov} \left(Y^f, \frac{\partial H}{\partial \delta} \right) - 2\hat{\beta} \text{cov} \left(H, \frac{\partial H}{\partial \delta} \right) \right\} \quad (3.1)$$

where Y^f is future income and H is expenditure on a home. Equation (3.1) implies that the relaxation of borrowing constraints generates an increase in $\hat{\beta}$ if

$$\text{cov} \left(Y^f, \frac{\partial H}{\partial \delta} \right) > 2\hat{\beta} \text{cov} \left(H, \frac{\partial H}{\partial \delta} \right). \quad (3.2)$$

In order to interpret this condition, it is useful to begin by seeing how it applies to our simple 2-type model. Consider first the left-hand-side, which measures the relationship between future income and the impact of loosening the borrowing constraint on the size of house purchased. This term is positive, because the high future income households are the ones who increase their house size when liquidity constraints are relaxed (some of the high-future income households move from small houses into big houses when δ increases, while none of the low-future income households switch). Thus, the covariance between future income and house movements is positive. Next consider the right-hand-side, which measures the association between the value of current house purchases and the impact of loosening the borrowing constraint on house size. This term is negative in the 2-type model. Households initially living in big houses are not affected by changes in borrowing constraints, because by assumption they are all unconstrained, high-future income types. However, a portion of the high-future income households initially constrained and living in small houses become unconstrained and move into big houses, in response to lower constraints. This implies that the covariance between the value of current housing purchases and the response of housing to changes in borrowing constraints is negative. The parameter estimate, $\hat{\beta}$, is nonnegative, since we expect a positive correlation

between house purchase expenditures and future income. Given a positive left-hand-side and a negative right-hand-side, condition (3.2) clearly holds.

However, economic theory suggests that we should expect condition (3.2) to hold even in a more general model. The left-hand-side should be positive because the permanent income model of consumption tells us that *ceteris paribus*, households with higher expected future income growth want to do more consumption smoothing than households with lower expected income growth. Thus, households expecting high income growth are more likely to find themselves borrowing constrained than households expecting lower income growth. This implies that households who are expecting higher income growth will be more affected by the relaxation of constraints, that is, the covariance on the left-hand-side is positive. In the same way, theory suggests that the right-hand-side of condition (3.2) should be negative. If housing consumption is a positive function of permanent income, then *ceteris paribus*, families with higher wealth will prefer to live in more expensive houses. Because households with higher levels of wealth are less likely to be borrowing constrained, then we would expect that on average, households living in bigger houses are less likely to buy larger homes in response to a change in constraints.

Figure 4 provides a diagrammatic summary of these arguments. Holding house value fixed, as we move in the vertical direction, toward increasing values of future income, the probability of a household facing borrowing constraints increases. On the other hand, holding future income fixed, as we move left horizontally, toward decreasing house values, the probability of a household facing borrowing constraints increases. In short, theoretical considerations suggest that, in general, the left hand side of (3.2) is positive and the right hand side is negative. Hence, we expect condition (3.2) to hold in a general model with multiple types of households and house values. Thus, such a model predicts an increase in the responsiveness of future income to the value of housing purchased when credit-market constraints are relaxed.¹⁰

¹⁰ In Appendix B, we show that in a regression of the value of housing purchased on expected future income, the condition that ensures an increase in the regression coefficient in response to lower credit constraints is weaker than condition (3.2). Assuming that the least-squares estimate in such a regression is given by $\hat{\gamma}$, the condition ensuring

that $\frac{\partial \hat{\gamma}}{\partial \delta} > 0$ is

$$\text{cov}\left(Y^f, \frac{\partial H}{\partial \delta}\right) > 0.$$

4. Empirical Setup

The previous section provided a simple model to explain why imperfections in housing finance like borrowing constraints can weaken the observed relationship between housing purchases and future income and, conversely, how relaxing credit constraints can strengthen this relationship. Thus, by examining how the relationship between current home purchases and future income has changed over time, we can infer whether the market for housing finance has become less imperfect. This section develops an econometric model to implement this idea. We then discuss the data used to estimate the model. Finally, we provide a straightforward graphical summary of the evolution of the relationship between house values (normalized by current income) and income growth in our data and show that it is consistent with our theoretical framework.

4.1 Econometric specification

Basic setup. We begin by writing down a fairly standard model for forecasting future income in panel data. Specifically, suppose that it is period t and we want to predict household i 's real income in f periods, $Y_{i,t+f}$. The forecasting model is

$$\ln Y_{i,t+f} = \alpha_0 + \alpha_1 \ln Y_{it} + \alpha_2 \ln H_{it} + \beta X_{it} + D_t + \xi_{i,t+f} \quad (4.1)$$

where Y_{it} is current real income,¹¹ H_{it} is the real value of the home that the household purchases in year t , X_{it} is a vector of socio-demographic variables such as age, education, race, and sex, D_t is a set of time effects,¹² and $\xi_{i,t+f}$ is household i 's forecast error of future income at time t .

The only real difference from typical forecasting models is that the model here is augmented with the value of housing purchases on the right-hand-side. According to the permanent income model, consumption and expenditure should reflect households' information about future income. By this logic, if households have better information about their future income than is contained in income's own history, then consumption and expenditure variables should help to forecast income. Housing expenditures are a natural choice in this context,

¹¹ In other specifications we also inserted lagged values of income along with current income as explanatory variables. This did not affect the estimated coefficient on the value of the house, but reduced the estimated coefficient on current income. Because including lagged income values reduces our sample size and because our main concern is the estimated coefficient on house value, we include only current income.

¹² We include time effects to control for aggregate influences such as business-cycle effects.

because for most households the purchase of a home is the largest purchase that it will make over its lifetime.

If we were to confront equation (4.1) with data, we would view a rejection of the hypothesis that $\alpha_2 = 0$ as evidence in favor of the joint hypothesis of forward-looking behavior and the ability of households to make reasonably accurate income forecasts.¹³ Furthermore, a higher sensitivity of future income with respect to housing expenditures implies a greater value for α_2 . As discussed in the previous section, this is critical for our purposes, because it suggests that the impact of changes in the housing finance system on households can be assessed by the estimated elasticity of future income with respect to housing expenditures. This observation allows us to examine from the *household's* standpoint the widespread belief that the housing finance market has become less imperfect over time.

Specifically, in the context of our model, a less imperfect housing finance system suggests an increased observed elasticity of future income with respect to housing purchase expenditures. Hence, if housing market innovations over time have relaxed borrowing constraints, we expect the coefficient on housing expenditures in equation (4.1) to increase over time, *ceteris paribus*. Algebraically, this translates into the value of α_2 growing in magnitude over time. To test this notion, we begin by augmenting equation (4.1) with an interaction term between a linear time-trend and house value:

$$\ln Y_{it+f} = \alpha_0 + \alpha_1 \ln Y_{it} + \alpha_2 \ln H_{it} + \alpha_3 * t \ln H_{it} + \beta X_{it} + D_t + \xi_{i,t+f}. \quad (4.2)$$

If the relationship between housing expenditures and future income is in fact becoming stronger over time, then α_3 should be positive.¹⁴

When it comes to making equations (4.1) and (4.2) operational, the first question one must confront is how to measure the left hand side variable, future income.¹⁵ Including the sum

¹³ If households are forward-looking in their housing decisions but are very bad at predicting their future income, we will not find a statistically significant relationship between the value of a newly purchased house and future income. In effect, then, we are testing the joint hypotheses that households are forward-looking and that they can predict their future incomes with some accuracy.

¹⁴ For purposes of clarity, we interact house value with a time trend rather than interacting each time effect with house value. A linear time trend is a simple way to summarize whether the magnitude of the house value coefficient changes over time. Interacting with a full set of time effects would produce results that would be difficult to interpret.

¹⁵ At least as far back as Friedman (1957, pp. 23–25), empirical investigators have been forced to deal with the ambiguous nature of the time horizon that is appropriate in the permanent income model. Friedman discarded both mean lifetime income and the short-term mean of the anticipated probability distribution of income as proxies for permanent income, instead advocating a measure based on an intermediate time horizon.

of the entire stream of future earnings is clearly infeasible. Instead, we use several alternative measures of future income: realized income two years in the future, realized income 4 years in the future, the average of realized income 2 years and 4 years in the future, and the average of income for each of the 5 years in the future. The advantage of the measures based on averages is that they mitigate concerns that our results will be biased by a change over time in the mix of transitory versus permanent income or in the degree of random measurement error in households' responses. On the other hand, due to the structure of our data set (see below), we have a larger sample size if we use single years rather than averages.

An alternative to equation (4.1) would be to put housing purchase expenditures on the left-hand-side, and expected future income on the right hand side. This is similar to a conventional housing demand equation (see Mayo 1981). The theory laid out in Section 3 suggests that financial innovations should increase the estimated coefficient on future income in such a regression. A serious problem with the housing demand specification is that, as just mentioned, we do not directly observe expected future income and have to rely on various proxies. While virtually every variable is mismeasured to some extent, the measurement of expected future income is particularly problematic, and can lead to an inconsistent estimate of its coefficient. In contrast, when expected future income is put on the left hand side, at least under reasonable assumptions, the effect is to increase the standard error of the regression but not render the parameter estimates inconsistent.¹⁶ In any case, as shown in Section 6.4 below, when we estimate the housing demand specification, our substantive results are essentially unchanged.

4.2 Data

Our primary data come from the Panel Study of Income Dynamics (PSID) for 1968 to 2001. The PSID is an annual panel data survey, which contains, among other things, detailed information on family income and demographic variables. In addition, it includes information on home purchases and their value. Some of the issues involved in using the PSID to estimate our model, including the reliability of the data on homeownership and housing finance, are discussed in Appendix A.

¹⁶ However, in such a regression, measurement error in the housing variable becomes a relevant concern. See Section 6.2 for a test for potential measurement error in the house value variable, and its effects on our results.

Our basic sample includes all households that purchased a home in the period spanning 1969–1999.¹⁷ This includes both first-time homeowners making the transition from renting to owning a home and existing homeowners who are moving into new houses. The PSID’s measure of the purchase price is the value of the home as reported by the household during the first year of occupancy.¹⁸

For some of our specifications, we distinguish between first-time house purchasers and purchasers who were already homeowners. If the household reported in the previous interview that it was renting, and reports in the current interview that it owns a home, then we characterize it as a home-purchaser. Furthermore, if the household has never reported owning a home in previous interviews, then we label it as a first-time homebuyer. There is a subtle issue regarding households for which we do not have a complete tenure history, either because the head was already an adult in the first wave of the PSID in 1968 or because the head entered the PSID in subsequent waves. If such a household owns a house but then buys another house, we can clearly classify it as a repeat buyer. However, if such a household does not own a home and then buys one, we cannot be certain that they are a first-time buyer as the head may have owned and then sold a house prior to accession into the PSID. To address this problem, we adopt the following rule: If the head of the household enters the PSID as a non-home owner and is less than 30 years old, then we assume that the head never owned a home before and thus consider the house purchase part of the first-time homebuyer sample. If the head enters the PSID as a non-homeowner and is over 30, then we assume that the household previously owned a home at some point in the past, and we label it as a repeat-homebuyer. Our results are not sensitive to this cutoff age: We tried decreasing the cutoff to 25, and while it reduced our first-time homebuyer sample by approximately 200 observations, it did not change any of the results.

Altogether, we identify 14,755 house purchases over this period. In several cases, the reported value of the purchased house is extremely low, even after converting from nominal to

¹⁷ We excluded 1968, the first year of the PSID, because of problems with the question pertaining to whether the household had moved within the past year.

¹⁸ The PSID has two variables that allow us to identify house purchases. The first is a question regarding the tenure choice of the household (that is, renting versus owning a home), and the second identifies households that have changed residence between the current and previous interviews. We use information from the house value question to double-check the accuracy of the tenure choice question. The value of the house is one of the few PSID questions that is imputed and thus does not have any missing observations for homeowners. Since this question is missing for renters, we are able to double-check that households that reported a tenure switch from renting to owning did in fact purchase a house in the time between interviews.

real values. In the majority of these cases, the household reports not having taken a mortgage, leading us to believe that these may be inheritances or transactions between family members that are not particularly relevant to our study. Therefore, we eliminate all observations for which the house's real value is less than \$5,000. This reduces our sample of house purchases to 14,495. After deleting observations with missing values for income (current and future), demographic variables, and observations that were part of the Survey of Economic Opportunity (SEO) portion of the PSID,¹⁹ we end up with a baseline sample of 5,277 for our 2-year average future income measure and 4,349 for our 5-year average future income measure. The percentage of households purchasing a home as a share of the total number of households in our baseline PSID sample fluctuates between 5.6 percent (1983) and 9.6 percent (1978) for the 30-year span of our data, and it displays no noticeable patterns.

Our income measure is total income, including both labor and capital income and transfers, received by both the husband and wife and any other individuals in the household. In order to correct for changes in the price level over time, we deflate the income and house value variables by the Personal Consumption Expenditure (PCE) Chain-Type Index, with 2000 as the base year.

Our left-hand-side variable is income f years in the future. As discussed above, we use several alternative measures: realized income two years in the future, realized income four years in the future, the average of realized income 2 years and 4 years in the future, and the average of income for each of the 5 years in the future. These choices are dictated in part by the structure of the data set. Because the PSID switched from an annual to a biennial survey in 1997, it turns out that values for f of 2 and 4 maximize our sample size.²⁰ As far as demographic variables,²¹ the X -vector of equation (4.1), we include: a cubic in age, education (a series of dichotomous

¹⁹ The SEO sample is not a nationally representative sample as it over-samples poor and immigrant families; studies based on the PSID typically exclude it.

²⁰ Because of the 1997 change, we do not have data for the years 1998 and 2000. Therefore, we must discard house purchase data for the observations for which future income, Y_{it+f} , lies in 1998, 2000, or after 2001. For example, if we were looking at income three years into the future, we would need to eliminate house purchases in 1995, 1997, 1999, and 2001; if we were looking five years ahead, we would be forced to throw away 1993, 1995, 1997, 1999, and 2001. Thus, using the average of two and four years as well as each of the individual years themselves, allows us to maximize the sample size. For the case of $f = 2$, we need to discard only 1996 and 2001, while for $f = 4$, we need to discard only 1994, 1996, 1999, and 2001.

²¹ All variables pertain to the head of the household, except the education variables, which correspond to whichever spouse has the highest level of education.

variables), race, whether the head of the household is female, and family size.²² The means and standard deviations of the variables for the different samples are reported in Table 1.

4.3 A Preliminary Look at the Data

Our model posits that to the extent that the market for housing finance has become less imperfect over time, the relationship between income growth and the value of current house purchases should strengthen. To see whether such a relationship is present in the data, we began by computing for each household that purchased a home the ratio of the value of the home to family income. We then computed the mean ratio for each octile of the distribution as well as the mean value of income growth (over a 2-year period) for the families in the respective octiles. Figure 5 contains plots of house value relative to current income (in logs) against income growth for both the first half of our sample (pre-1985) and the second half of our sample (post-1985, inclusive).

Several features of the graph are of interest. First, the relationship between house value relative to current income and income growth during both the beginning and end of our sample period is positive. This is consistent with the notion of forward-looking behavior of households. Second, over time income growth has become more sensitive to normalized house value for households with relatively low levels of housing, but has stayed the same for households with relatively high levels of housing. This makes sense in terms of our theory—the families who bought large homes at the beginning of our sample period presumably were less constrained than those who bought small homes. Therefore, the loosening of constraints that took place over time likely affected their behavior less.

A third and related observation is suggested by a comparison of income growth at the low end of the house value distribution pre- and post-1985. Average income growth for the low-housing group was lower in the post-1985 period. Our framework provides a straightforward explanation for this phenomenon. In the presence of the relatively severe capital market constraints during the pre-1985 period, some households consumed little housing because their expected income growth was low, and some consumed little housing because their expected income growth was high but they couldn't borrow. In effect, there was a mixture of low- and

²² To conserve space, we omit the coefficients on age-squared and age-cubed, both of which are statistically significant at the 1-percent significance level for all of the specifications. The estimated coefficient on age-squared is negative, which is consistent with the hump-shaped income profile commonly estimated in the literature.

high-income growth households at the low end of the housing expenditure distribution. But, with the less severe constraints of the post-1985 period, fewer high-income growth individuals ended up spending relatively little on housing, so there were fewer high-income growth households in this part of the housing distribution. Thus, when we average over income growth rates in this part of the distribution pre- and post-1985, there are more low-growth households post-1985, which accounts for the relative positions of the two graphs.

While it is encouraging that Figure 5 is consistent with our theoretical framework, we must be careful not to place too much stock in it, because it does not take into account effects from other variables that could be contributing to such a pattern in the data. A multivariable analysis is required, to which we now turn.

5. Results

This section presents the results of our basic model, including a specification that allows for the possibility of structural breaks in the relationship between the value of a current house purchase and future income. We then examine how the parameter estimates differ across various subsets of the population and whether the substantive findings are robust to alternative specifications.

5.1 Basic results

The results for estimating equation (4.2) with future income defined as the average of income two and four years ahead are in columns (1) through (5) of Table 2. Consider first the estimated coefficient on the log of house value, α_2 . For our baseline sample (column (1)), the estimate is 0.131. It is significantly different from zero at conventional levels. The implication is that every 10-percent increase in the value of a new home is associated with a roughly 1.3 percent increase in future income.

Our main interest is the coefficient on the interaction of the linear time-trend with house value, α_3 , which tells us whether the relationship between house value and future income has become stronger over time. In column (1) we find that the estimated coefficient on the interaction term for the baseline sample is significant, with a value of 0.004. An estimate of 0.004 implies that in 1969, a 10-percent increase in the value of the house corresponded to a 1.3-percent increase in future income, while in 1999, a 10-percent increase in the value of the house

was associated with a 2.5 percent increase in future income. Thus, the relationship between the value of housing purchases and future income approximately doubled. Within our conceptual framework, this suggests that for the baseline sample, constraints loosened over time, consistent with the notion that developments in housing finance made it easier for households to purchase homes in line with their future income prospects. While this does not constitute a formal test of our model, it is comforting that the result is the same as that suggested by both casual observation and previous econometric work that followed a very different approach.

We now discuss briefly the other coefficients in our basic model in column (1). Most of the demographic variables have statistically significant effects, with signs that are consistent with prior research. The coefficient of 0.358 on current income replicates the usual result that income has an autoregressive component. Future income is increasing in age, *ceteris paribus*, consistent with typical analyses of age-income profiles. The coefficient estimates on the education variables imply that future income is approximately 16 percent higher for high school graduates with some college experience than for high school drop-outs, and about 26 percent higher for college graduates. Households that are neither Caucasian nor Black (*other race*) have future income realizations that are approximately 16 percent lower than those of white households. Female-headed households have future incomes that are almost 34 percent lower than those of male-headed households. Size of household is also statistically significant and positive, although the coefficient's magnitude is small. It suggests that for each one-person increase in the size of a household, its income in the future is approximately 2 percent higher. The coefficient estimates for the time effects are omitted from the table to conserve space; however, they are included in all of the regression models. For the most part, they are statistically significant.

Are these findings sensitive to the way in which future income is measured? In Table 3, we report for each definition of future income the estimated coefficient on the log of house value (α_2) and a time trend interacted with the log of house value (α_3) (To conserve space, we do not report the other coefficients; they are available upon request.). Each panel of Table 3 corresponds to an alternative definition of future income. The first and second rows in each panel show the estimates of α_2 and α_3 , respectively. Comparing the coefficients as we move down the table, we see that the results corresponding to alternative definitions of future income (5-year average, two years ahead, and four years ahead) are virtually identical. Thus, as hypothesized, and consistent

with forward-looking behavior, the value of a newly purchased house has power in predicting future income, even conditional on current income.

5.2 Breakpoint analysis

Our approach so far has been to use the interaction of time and house value to determine whether the observed elasticity of future income with respect to current housing purchases has increased over time. An alternative approach is to let the data determine whether there was a discrete structural change in the relationship between future income and housing purchase expenditures, and if so, to see whether the timing of the change can be related to changes in the operation of the mortgage market. In terms of our basic model, equation (4.4), the question is whether at some point in time there was a discrete change in α_2 , the coefficient on house value.

To implement this idea, we use Bai's (1999) likelihood-ratio-type test for multiple structural changes in regression models. The test determines both the number of structural breaks and the location of each break in the data. The particularly novel aspect of the test is that both the null and alternative hypotheses allow for the possibility of breakpoints.²³ This implies that one can use the methodology to test for multiple breakpoints against the null hypothesis of a single (or multiple) breakpoint(s). For example, one can test for three versus two breakpoints. The intuition behind the test is fairly simple. Essentially, it consists of calculating the sum of squared residuals (SSR) for each possible partition of the data corresponding to the number of breakpoints under the null hypothesis, n_0 , and taking the smallest value. Then, one does the same for each possible partition of the data corresponding to the number of breaks under the alternative hypothesis, $n_1 = n_0 + 1$ and compares the minimum SSR under the null to the minimum SSR under the alternative. If the SSRs are not "significantly" different from each other, then the null hypothesis is not rejected and it is assumed that the data contain n_0 breaks. If they are different, then the null hypothesis changes to n_1 breaks, the alternative hypothesis to $n_2 = n_1 + 1$, and the procedure repeats itself until the null is not rejected. For a more detailed explanation of the methodology, the reader is directed to Bai's paper.

The findings from the breakpoint analysis for the baseline sample are reported in Table 3. In each panel, the third row, which is labeled "Breakpoint," shows the estimated breakpoint for

²³ The limiting distribution of the test statistics for tests of only a single structural break are derived assuming the absence of breaks. Thus, when the null hypothesis is rejected for such a test, only a single breakpoint is estimated.

the corresponding definition of future income. For example, when future income is measured as the average of income two and four years ahead, the breakpoint in the baseline sample is in 1985. We find a single breakpoint in 1982 for the 5-year-average specification (panel 2), and single breakpoints in 1985 and in 1984 for the models that predict income two years (panel 3) and four years (panel 4) into the future, respectively. The fact that the locations for the breakpoints are so close across all of the models shows that the algorithm for finding the breakpoints is robust to at least minor changes in specification.

Importantly, when we estimate equation 4.2 separately before and after the breakpoint, we find that the interaction between the time-trend and the value of house purchases is not significantly different from zero in either subsample. This is consistent with the result of only a single breakpoint from Bai's test, which indicates that the ability of house purchases to forecast income was not increasing gradually. This finding sets a high bar for any explanation that depends on a phenomenon that changed gradually over time.

The mid 1980s breakpoint suggests that changes in consumer behavior were caused by structural changes in financial markets as opposed to, say, the anti-discrimination laws (which were passed 10 years earlier) or the development of the sub-prime market (which occurred 10 years later). The 1985 timing might at first glance seem surprising, given that Congress passed the two main acts that deregulated the savings and loan industry in 1980 and 1982. However, the acts explicitly provided for a gradual phase-out of regulations. For example, Regulation Q was not fully eliminated until 1986. In the same way, there were substantial lags in the development of the secondary market. Although nascent secondary markets emerged in the 1970s, they did not really mature until 1986, when the Tax Reform Act of 1986 established the legal framework for mortgage markets that exists today. There was also a lag between the emergence of new mortgage designs (which occurred in the late 1970s and early 1980s) and a serious augmentation of the mortgage choice set. When the new designs originally appeared, they merely maintained reasonable mortgage choices in a high-inflationary environment. Meaningful expansion of the menu of mortgage types required both alternative prototypes and lower nominal interest rates and inflation, and this combination did not occur until 1986, when single-digit mortgage rates appeared for the first time since 1978. In short, given the realities of the evolution of the market for housing finance, a breakpoint in the mid 1980s seems perfectly reasonable.

5.2.A Quantitative magnitude of the break

We turn now from the timing of the structural breaks to an assessment of their quantitative impact. To do this, we augment equation (4.1) with an interaction term between house value and a dichotomous variable that takes a value of 1 for the year of the break and the years thereafter. For example, for the specification using the average of income two and four years ahead on the left-hand-side, in which we find a breakpoint in 1985, the regression model is:

$$\ln Y_{it+f} = \alpha_0 + \alpha_1 \ln Y_{it} + \alpha_2 \ln H_{it} + \alpha_3' (d8501 * \ln H_{it}) + \beta X_{it} + D_t + \xi_{i,t+f}, \quad (5.1)$$

where $d8501$ takes the value of 1 for years between 1985 and 2001, and 0 for years between 1969 and 1984. In the spirit of the breakpoint analysis, equation (5.1) does not include interaction terms between time and house value.

The parameter estimates for equation (5.1) using the baseline sample are shown in column (6) of Table 2. The results reinforce the findings from the breakpoint analysis. The estimate of α_3' is positive (0.064) and statistically significant at the 1-percent level. This means that prior to 1985 the estimated elasticity of future income with respect to house value, α_2 , was 0.161, while after 1985 it increased to 0.225 (= 0.161 + 0.064). This is a substantial increase in the forecasting relationship between house purchases and future income, and it confirms the notion that mortgage markets became much less imperfect sometime in the mid 1980s. As before, we summarize in Table 3 the key results when the model is estimated using alternative definitions of future income. In each panel, the fourth row shows the estimate of α_2 , the estimated responsiveness before the corresponding breakpoint, and the fifth row shows α_3' , the change in the estimated sensitivity after the breakpoint. The results are not very sensitive to the choice of definition of future income.

5.2.B Alternative explanations for the breakpoint

A potential problem with our interpretation of the breakpoint is that the broader macroeconomic environment was changing at around the same time. First, the mid 1980s saw the end of the high-inflation, high-interest-rate environment as well as the end of the great disinflation of the early 1980s.²⁴ Second, as we mentioned in the introduction, macroeconomic

²⁴ For a detailed discussion and documentation of this disinflation and its effects on the macroeconomy, see Goodfriend and King (2005). In addition, recent research suggests that the 1980s saw a secular reduction in macroeconomic activity. See McConnell and Quiros (2002).

volatility appears to have fallen in the 1980s. Third, the skill premium widened considerably. And finally, some economists have argued recently that while macroeconomic volatility fell in the 1980s, household level income volatility actually rose.

Changes in interest rates. The problem arises because our theory suggests that a fall in nominal rates, holding real rates constant, would have the same effect as a relaxation of borrowing constraints. To see why, recall that with traditional mortgages, a high nominal rate drives up the real value of initial monthly payments without raising the real interest cost. High nominal rates increase the monthly payment for a given home, and if the lender's allowable maximum fraction of monthly income that can be spent on housing remains constant, then this effectively makes the borrowing constraint more stringent. However, we doubt that the reduction in nominal interest rates during the mid 1980s is driving our results. If that were the case, we would expect to see two structural breaks, because relatively low nominal rates prevailed both before 1979 and after 1985. The notion that nominal interest rates are driving the process is also implausible given our previous results that found a positive time-trend in the relationship between future income and house value (see Table 3). Given that nominal rates were relatively low at the beginning of our sample period and again at the end, such a time trend would not have emerged if these rates were the driving mechanism.²⁵

Reduction in macroeconomic volatility. Changes in higher moments of the income profile are another possible source of bias in our estimates of α_3 , and α_3' . There is substantial evidence that aggregate income volatility has declined over the past few decades, a phenomenon referred to as the "Great Moderation" in the literature. How might our results be affected if income volatility has also decreased at the household-level, so that households have become more certain about their future incomes or better able to predict them? To explore this issue, suppose that there is considerable uncertainty with respect to future income. Precautionary saving will therefore be strong, and individuals will therefore tend to demand relatively small homes. However, if homes are small (relative to income), then the probability of being credit constrained is relatively low. Now suppose that uncertainty falls. Consequently, precautionary saving decreases, households want to spend more on housing, and credit market constraints are more likely to be binding. In our model, the slope of the relationship between house value and

²⁵ From 1969 to 1977, the average nominal interest rate on a 10-year Treasury bill was 7.09 percent; from 1978 to 1985 it was 11.30 percent; and from 1986 to 2001 it was 6.94 percent.

future income falls, but this is just the opposite of what we find in the data. In short, a reduction in income volatility and/or a better ability to forecast income does not explain our findings.

Changes in the distribution of income: Much research has documented the substantial increase in the dispersion of income over time, particularly between those with and without a college education.²⁶ Although economists differ about the exact timing and extent of the change, most believe that widening of the gap started sometime in the early 1980s. Could this account for the breakpoint we find in the data? It is hard to come up with a convincing story. In the first place, the focus of our model is on the growth of income rather than its level, and it is not clear what has been happening to the dispersion of growth rates over time. More fundamentally, though, there is no reason why a change in the distribution of a right-hand-side variable per se should bias its estimate as long as the underlying relationship with the left-hand-side variable remains the same.

Increase in household level income volatility: Some research suggests that, despite the decrease in aggregate income volatility, there has been an increase in household-level volatility (Dynan *et. al.* 2007). Now, we argued above that a reduction in volatility would result in an increased fraction of the population facing binding constraints and make it less likely that we would find a breakpoint. The same reasoning therefore suggests that an increase in volatility would have the opposite effect and present a plausible alternative explanation for our finding. However, we can distinguish between the hypothesis of financial innovation and the hypothesis of increased household-level volatility because they have opposite implications regarding the magnitude of spending on housing purchases. Specifically, if increased income volatility were driving our results, we would see a reduction in spending on housing because of precautionary saving. On the other hand, as already noted, if financial innovation were at work, we would observe an increase in spending, other things being the same. To determine which view is more correct, we estimate a regression of house purchase values on a time-trend, including current income and all of our demographic variables as well. We find that the time-trend is positive and statistically significant. This indicates that the value of real house purchases has been increasing over time, *ceteris paribus*, which is consistent with the financial innovation hypothesis.

5.3 Potential bias due to omitted variables

²⁶ See Card and DiNardo (2002) for a review of the literature on wage inequality.

In virtually any regression analysis there is a possibility of omitting a key variable that is correlated with both the dependent variable as well as the right-hand-side variables of interest. The omission of such a variable leads to biased estimates of the coefficients on the relevant right-hand-side variables. In our context, we first note that any omitted variables that are correlated with both house purchase expenditures and income growth can result in a biased estimate of α_2 (the coefficient on the value of house purchases), but not necessary a biased estimate of α_3 , or α_3' (which measure how the relationship between the value of house purchases and future income changed over time). This is an important point, because our theoretical framework focuses us primarily on the *change* in the sensitivity of future income to house purchase expenditures (α_3 , and α_3'), and not on the level of the sensitivity (α_2). Put another way, in order to create a bias in the estimates of the parameters that are our main concern, α_3 , or α_3' , an omitted variable would have to generate a change in the correlation between the value of house purchases and income *growth* over time. Moreover, because we find substantial evidence of a single breakpoint in our estimates of α_2 , the change in the omitted variable would have to be discrete. We assess below the potential importance of the omission of several specific variables from our empirical model.

5.4 Differences across sub-groups

Much of the public policy discussion of housing finance has focused on the ability of disadvantaged households to buy homes. In this context, we are particularly interested in the impact of the development of mortgage markets on poor families, female-headed households, black households, and younger households. As well, there have been concerns about first-time homebuyers. In this section, we present the results when the model is estimated separately for each of these sub-groups.

Poverty, race and gender. To investigate the evolution of the housing finance situation facing the poor, we estimate our model using only data from the SEO poverty sample in the PSID, which we refer to as the “poverty sample” hereafter. Summary statistics for the poverty sample are reported in panel 2 of Table 1.

The results when we use only observations from the poverty sample to estimate equation (4.2) are reported in column (2) of Table 2, and are very similar to their counterparts estimated using the baseline sample in column (1). The estimated coefficient on house value in the second column of the table is slightly higher than the corresponding estimate in the first column,

indicating that, for the poverty sample, the relationship between the value of the house purchase and the average of income over two and four years into the future is slightly stronger. The estimated coefficient on the interaction term, $time*loghval_{it}$, is similar in magnitude to its counterpart in the baseline sample, although it is significant at the 10 percent level as opposed to the 1 percent level for the baseline sample. Thus, poorer families exhibit similar forward-looking behavior in their housing purchase decisions, and this behavior has also become stronger over time. In Table 3, the first two rows of each panel in column (2) show how the estimated coefficient on log of house value and its interaction with time vary with the definition of future income. The results are very similar across our different income definitions. For the four-year ahead specification (panel 4), the estimated coefficient on log of house value is 0.142, while for the five-year average (panel 2) and two-year ahead (panel 3) specifications, the estimated coefficient on log of house value is slightly smaller (0.114 and 0.119 respectively). The estimated coefficient on the interaction term is positive and statistically significant at the one percent level for the five-year average and two-year ahead specifications. However, it is not significant for the four-year ahead specification (although the sign is positive).

Breakpoint analysis results for the poverty sample are also reported in column (2) of Table 3. The breakpoints for each definition of future income are in the third row of the respective panels. A single structural break in 1978 (significant at the 10-percent level) is found for the model with future income averaged over two and four years; one break in 1977 for the model using a five-year average of future income; one break in 1980 for the 2-year model, and one break in 1981 for the 4-year model.

The quantitative magnitudes of the various breaks (based on estimation of equation (5.1)) are also displayed in column (2) of Table 3. The shifts are similar in magnitude to those for the baseline sample, with values of 0.09 for the average of two and four years ahead; and 0.10, 0.11 and 0.12, for the 5-year average, 2-year, and 4-year definitions, respectively. These results suggest that the decrease in market imperfectness for poorer households was similar in magnitude to that in our baseline sample.

This begs the question of why the breakpoint for the poverty subsample occurred earlier than for the sample as a whole. Given that the key financial market innovations mostly took effect later, we searched for alternative stories. Because poor families are disproportionately

headed by women and blacks (see Table 1), a natural starting point for such an investigation would be to estimate the model separately by gender and race.

For female-headed households, we find minor evidence of a breakpoint in 1978 for the average of income two and four years ahead, minor evidence of a breakpoint in 1976 for the 5-year-average specification, minor evidence of a breakpoint in 1981 for the 2-year specification and minor evidence of a breakpoint in 1976 for the 4-year specification. We do not find evidence of a breakpoint for black households in any specification.²⁷

A possible explanation for the findings of female-headed households relates to legislation passed in the mid 1970s. In Section 2 we noted that concerns about possible discrimination based on race and gender led to the passage of the Equal Credit Opportunity Act (ECOA) and its amendment in the mid 1970s. Whether or not legislative measures were effective in helping minorities and women was (and still is) a controversial topic. That said, expanded opportunities for women could explain our findings of a breakpoint in the late 1970s and early 1980s, given that some have argued that ECOA was implemented with a lag (see, for example, Munnell et al. 1996). Because households in the poverty sample were twice as likely to be headed by women as in the baseline sample, this could explain our poverty sample results as well. The failure to find a breakpoint in the black subsample could reflect the stubbornness of racial discrimination in lending markets. As we mentioned above, some researchers argue that racial discrimination remained pervasive in mortgage markets as late as 1990.

First-time homebuyers. Another group of people who may be relatively disadvantaged when it comes to housing finance are first-time homebuyers. We expect families who currently own or previously owned a home to have advantages over first-time homebuyers when it comes to financing the purchase of a new house and obtaining a mortgage. For example, existing homeowners have an established credit history and often have a cushion of available equity. Therefore, this group would stand to benefit least from a less imperfect mortgage market, since they are less likely to be borrowing-constrained in the first place. The other side of the same coin is that if mortgage-market improvements really are efficacious, they should have a relatively large impact on first-time homebuyers. To investigate these issues, we construct a first-time homebuyer sample and estimate our models separately for this sample. Because we are interested in households that are most likely to have been borrowing constrained, we restrict our sample of

²⁷ These results are available upon request.

first-time homebuyers to households who financed their purchase with a non-trivial mortgage. Thus, we define a first-time homebuyer as a household that had not previously purchased a home, and had financed their first purchase by obtaining a loan at least half of the value of the purchase price (loan-to-value ratio of greater than or equal to 0.5).²⁸ Summary statistics for these samples are shown in Table 1.

The third column in Table 2 shows the parameter estimates from the basic equation (4.2) for first-time homebuyers. The estimated coefficient on house value is small (0.036) and statistically insignificant, however the estimated coefficient on the interaction term, $time*loghval_{it}$ is large (0.006) and statistically significant at the 1 percent level. This implies that by 1999, a 10-percent increase in the value of the house was associated with a 2.2 percent increase in future income for the first-time homebuyer sample. Thus, at the beginning of the sample period, there is no relationship between house expenditures and future income for first-time homebuyers, but by the end of the sample period the relationship is similar in magnitude to that of the baseline and poverty samples. This is strong evidence in support of our above hypothesis that potential disadvantaged families such as first-time homebuyers have benefited the most from improvements in mortgage markets.

The corresponding breakpoint results are consistent with those above. The third column of Table 3 tells us that in all specifications of the model, there is evidence of a single structural change in the coefficient estimate of house value sometime in the early to mid 1980s. While Bai's breakpoint test only finds a statistically significant break for future income defined as a 5-year average (panel 2), the null hypothesis that the estimated coefficient corresponding to house value is the same before and after the break is rejected at the 1 percent significance for almost every specification (5 percent level for income 4 years ahead). This further suggests that the decreased imperfectness of the mortgage market had a significant effect on first-time homebuyers.

Younger versus older households. Households that are just starting families and still in early stages of the life-cycle may not have the same access to credit markets as older households with well-established borrowing histories. We imagine, therefore, that younger families would be

²⁸ The results are robust to slight perturbations of this loan-to-value rule (i.e. $ltv=0.4$ or $ltv=0.6$). However, the results change when we include in the sample households that purchased their first home either without obtaining a mortgage or by obtaining a trivial mortgage (i.e. $ltv<0.2$). We suspect that many such households inherited their homes or were the beneficiaries of family transfers. Unfortunately, the PSID does not contain enough information to test this hypothesis.

more likely to be credit constrained and commensurately more likely to benefit from improvements in the market for housing finance. To investigate this hypothesis, we divide our sample into households under the age of 40 and households over the age of 40 (inclusive), and estimate our model separately for each sample. Summary statistics for these samples are displayed in panels 4 and 5 in Table 1. Column (4) in Table 2 shows the parameter estimates from the basic equation (4.2) for the younger sample. The estimated sensitivity of future income (average of 2 and 4 years ahead) to house purchase expenditures is 0.098, and it has increased significantly over time as evidenced by the coefficient, 0.005, associated with the interaction term between the time-trend and house purchase expenditures. Furthermore, we find evidence of a single break point in α_2 occurring in 1981 (column 4 of the first panel in Table 3). As also indicated in Table 3, for each of the other definitions of future income, there is a break either in 1980 or 1981. Thus, we find substantial evidence of a decrease in the imperfectness of mortgage markets for households with heads under the age of 40.

Conversely, we do not find evidence of improved credit markets for the sample of older households. Column (5) in Table 2 displays the parameter estimates for the over 40 sample. The estimated sensitivity of future income to house value is 0.186, which is larger than all of the other samples. However, the interaction term, $time * \log hval_{it}$ is not statistically different from zero—the sensitivity of future income to house value has not increased over time. This is true regardless of how future income is defined. (See the estimates in the second row, fifth column, of the various panels in Table 3). Furthermore, as indicated in the third rows of column (5) of the various panels, there is no evidence of a breakpoint occurring in any of the specifications for this subsample.

Taken together, these results make perfect sense within our framework. Younger households, which are likely to be constrained in credit markets, *ceteris paribus*, have benefited from innovations in the market for housing finance. Older households, which are less likely to be constrained, have not.

High-housing versus low-housing families. As a final check on our model, we divide our sample according to the ratio of house value to current income at the time of purchase. As already noted, according to our simple theoretical model, the sensitivity of future income to housing purchase expenditures should increase more over time for families who originally purchased relatively small homes, because they are the families who are most likely to have been

constrained. On this basis, we would expect the increase in the sensitivity of future income to housing expenditures to be much more pronounced for households purchasing lower-valued homes relative to their income at the time of purchase.

There is bound to be some arbitrariness in defining a “high” or “low” value of housing relative to current income. We simply set the dividing line where the log of the ratio was equal to 1. For the “low” sample, we find strong evidence of a single breakpoint in 1984 for the 2-year average specification of the model (significant at the 1-percent level) and strong evidence of a structural break in 1980 for the 5-year average specification. There is weaker evidence of a break in 1984 for the 4-year ahead specification (significant at the 10-percent level). For the 4-year ahead specification the breakpoint occurs a year later (although not quite significant at the 10-percent level). In contrast, for the “high” sample of households, we find no evidence in any of our specifications of a breakpoint in the relationship between future income and housing purchases.²⁹ These findings are consistent with the implications of the simple theoretical model presented in Section 3 and with Figure 3, the plot of the distribution of house expenditures versus the distribution of future income from Section 4.3.

5.5 The Role of Labor Market Decisions

Our model posits that people decide on the purchase of their home given their beliefs regarding future earnings, and we find such a relationship in the data. However, an alternative interpretation is that causality runs in the other direction—households decide to purchase a house that is beyond their means, and then work harder in the future to earn enough income to make the mortgage payments.

There are basically two feasible channels through which a household could increase its income over a relatively short time horizon. The first and more likely channel is by increasing labor supply, either by working more hours or by obtaining a second job. To see whether the correlation between the value of a current house purchase and future income is driven by increases in labor supply, we estimate a regression of total annual hours worked on the value of the house purchase, current hours worked, income, and our usual set of demographic variables,

²⁹ We also estimate equation (4.2) for only the families with high ratios of housing to current income; we find that the estimated coefficient of the interaction term between time and house value is not significantly different from zero for any of the model specifications, further reinforcing the findings from the breakpoint analysis. These results are available upon request from the authors.

two years (as well as four) after a house purchase. If the increased labor supply scenario were operative, we would expect to see a statistically significant, positive coefficient estimate associated with house value. However, the estimated coefficient is not statistically different from zero in any of our specifications.

A second channel through which a household might affect its income in the short run is occupational change. The PSID contains information about the occupation of both the head of the household and the spouse.³⁰ We use this information to construct a dichotomous variable that takes a value of 1 if either the head of the household or the spouse switched occupations during the two (four) years after a house purchase, and 0 otherwise. We then estimate a probit model for occupational switches, in which we include the value of the house purchase, current labor income, a set of occupation indicator variables, and a set of demographic variables. If households are buying homes that they cannot afford and then switching jobs in an effort to increase income, we would expect that higher expenditures on housing would increase the probability of switching occupations. However, we find that the value of the house purchase has no statistically discernible effect on the probability of switching careers within two years of a purchase. Exactly the same result is obtained when we look at a 4-year time horizon.³¹

6. Alternative specifications

In this section, we assess the robustness of our results to alternative econometric specifications. We begin by looking at a model that is more in the spirit of the traditional setup in the housing demand literature. Next we analyze some measurement issues. Finally, we consider a series of issues related to the consequences of omitting certain variables from our model.

6.1 Housing demand specification

Consider a conventional housing demand model, in which the value of house purchase is regressed on future income, *inter alia*. As we showed in Section 3, our model predicts that the

³⁰ The coding of occupation in the PSID presents some technical issues. In the early years (up to 1980), occupation was coded at the 1-digit or 2-digit Census level, while in the later years it was coded at the 3-digit level. To construct an occupational code that is consistent for the entire span of the data, we used the 1968–1980 Retrospective Occupation-Industry Files, a PSID supplement that provides 3-digit occupation codes for household heads and spouses pre-1980. We then constructed our own 1-digit code for the entire 32-year sample (12 different classifications).

³¹ These results are available upon request to the authors.

coefficient on future income should increase in response to the relaxation of borrowing constraints in the mortgage market. As we discussed in Section 4.1, we believe that such a model is more susceptible to measurement error than our preferred specification. Nevertheless, as a check on the robustness of our results, we estimate the equation,

$$\ln H_{it} = \gamma_0 + \gamma_1 \ln Y_{it} + \gamma_2 \ln Y_{it+f} + \gamma_3 t^* \ln Y_{it+f} + \beta X_{it} + D_t + e_{it}. \quad (6.1)$$

In results not reported here for brevity, we find that, consistent with the predictions of our theoretical model, the estimates of γ_2 and γ_3 are positive and statistically significant. In addition, we estimate a variant of the model allowing for a discrete change in the relationship in 1985,

$$\ln H_{it} = \gamma_0 + \gamma_1 \ln Y_{it} + \gamma_2 \ln Y_{it+f} + \gamma_3' (d8501 * \ln Y_{it+f}) + \beta X_{it} + D_t + e_{it}. \quad (6.2)$$

Consistent with our previous results, the estimate of γ_3' is positive and statistically significant, suggesting the presence of a break point in the mid-1980s in the sensitivity of house purchase values to future income. The estimates of equations (6.1) and (6.2) are available upon request.

6.2 Measurement Issues

House values in the PSID are self-reported. Measurement error associated with self-reported asset values is a common problem in the empirical literature (Miniaci and Weber 2002), and could lead to inconsistent parameter estimates. To address this problem, we take advantage of the fact that the PSID provides information about whether the value of the house was edited or imputed. When we estimate our models including only those values that were not edited in any way, we find no substantive changes in our results.

Another measurement issue relates to income. In particular, should capital income be included as well as labor income? We believe the answer is yes, as forward-looking households likely consider all sources of income when contemplating a house purchase. However, because capital and business income might be more difficult to forecast than labor income, it is useful to confirm that our results are not sensitive to the way that income is defined. Therefore, we re-estimate all of the forecasting equations using only household labor income.³² We find no substantive differences from the previous results.³³

³² Another issue is whether income should be measured in real or nominal terms. We have used real income, because it is more consistent with the basic permanent income hypothesis. However, to make sure that our results are not sensitive to this distinction or our choice of deflator, we re-estimated the model using nominal magnitudes for house value and income, and we found that the results were essentially unchanged.

³³ The coefficient on current income increases and the cubic term in age is no longer significant, but the estimated coefficient on house value is not noticeably different.

6.3 Forward looking behavior

We noted above that our main interest is the coefficient on the interaction term between the value of house purchase and time, and that it is unlikely that omitted variables are substantially biasing this coefficient. That said, the coefficient on the value of house purchase, α_2 , is of independent interest because a positive sign is consistent with the hypothesis of forward looking behavior. It is therefore useful to consider whether omitted variables might be biasing our estimates of α_2 .

Financial wealth is a possible candidate for an omitted variable. Our forecasting equation follows the conventional tack of not including wealth on the right-hand-side. However, to the extent that wealth is correlated with future income as well as house value, then omitting it could lead to inconsistent estimates of α_2 . Unfortunately, only a limited amount of wealth data is available in the PSID. The PSID provides supplementary information regarding household wealth for just 4 years – 1984, 1989, 1994, and 1999. Using these years, we re-estimate the income forecasting equation with wealth as an explanatory variable. For this sample, which has 1,647 observations, the estimated coefficient on house value is 0.212 with wealth included, while it is 0.204 without wealth. The difference is not statistically significant, leading us to conclude that, on the basis of the available data, the omission of wealth is not biasing our estimates of α_2 .

Another possible omitted variables problem arises from the fact that there might be regional differences in the demand for housing. Perhaps, for example, families in high income-growth regions tend to prefer more expensive houses. The PSID allows us to address this issue straightforwardly. The data include information about the household's state of residence, and we re-estimate our basic models with state and regional effects. We find that their inclusion has no substantive effect on our results.

7. Investigating the Contribution of the GSEs

We now use our econometric model to assess the impact of the securitization activities and portfolio decisions of the housing GSEs, Fannie Mae and Freddie Mac. To begin, we define $GSElev_t$ to be the proportion of the stock of all home mortgages that were securitized by Fannie and Freddie in year t , and $GSEacq_t$ to be the corresponding flow variable. Figure 1 shows the evolution of each of these ratios over time. With our measures of GSE activity in hand, the

question is how to use them to estimate the extent to which the loosening of borrowing constraints documented above can be attributed to the securitization activity of the GSEs. A straightforward approach is to interact one or another measure of their activity in period t , GSE^*_t , with house value in our empirical regression model:

$$\ln Y_{it+f} = \alpha_0 + \alpha_1 \ln Y_{it} + \alpha_2 \ln H_{it} + \alpha_3 GSE^*_t \ln H_{it} + \beta X_{it} + D_t + e_{it}. \quad (7.1)$$

A positive value of α_3 implies that more extensive GSE activity is associated with H_{it} 's being a better predictor of future income, which in turn implies an improved ability for households to purchase homes that are in line with their future income prospects. In short, a positive value of α_3 is consistent with the notion that GSE activity has helped to make the housing finance market less imperfect.

An important question is the time period over which the model should be estimated. As emphasized above, the 1970s and early 1980s were a period of tremendous flux in U.S. financial markets. The technological and regulatory environments changed dramatically. We have no way to parameterize these changes, and hence, no way to discern the independent effect of the GSEs' securitization activities. In mechanical terms, if we were to estimate equation (7.1) with observations going back into the 1970s, the coefficient on the GSE variables would be biased upward because it would be picking up not only the impact of securitization, but also the impact of the omitted changes. In order to deal with this problem, we begin the sample in 1983. By then, most of the deregulatory changes had been enacted, so we can feel fairly confident that the coefficients on the GSE variables are in fact reflecting the impact of their securitization activities, and not other changes in the financial environment. In any case, prior to 1983, GSE securitization activity was almost negligible. Before 1982, outstanding mortgages securitized by the GSEs were never more than 2 percent of all home mortgages outstanding. In 1982 this number increased to 5 percent, while after 1990 the proportion of all home mortgages outstanding securitized by the GSEs never fell below 25 percent and was often close to 30 percent (see Figure 1).

Table 4 shows the estimate of α_3 in equation (7.1), and the associated p-value, for each of our samples, with future income defined as both the average over two and four years as well as the 5-year average.³⁴ (As in Table 3, the coefficients of the other variables are omitted for the

³⁴ Results for the 2- and 4- year ahead definitions are virtually identical and are available upon request.

sake of brevity and available upon request.) Each column exhibits results using data from the indicated sample (baseline, poverty, first-time home buyers, age less than 40, and age greater than or equal to 40). Panel 1 includes $GSElev_t$ in the regression, while panel 2 includes $GSEacq_t$. The results are quite clear for the baseline sample (column (1)): Whether we measure GSE securitization activity in terms of stocks or flows, and whether we are trying to estimate future income as an average over two and four years or as the five-year average, we cannot reject the hypothesis that the coefficient on the interaction term between the log of house value and GSE securitization activity is zero. Thus, we find no evidence that the securitization activities of Fannie and Freddie played any role in making the housing finance market less imperfect.

The GSEs and marginal borrowers. One of the missions of the GSEs is to facilitate the financing of housing for low and middle-income families. Indeed, in 1992 Congress passed legislation requiring the Department of Housing and Urban Development to establish quantitative goals for Fannie and Freddie with respect to their purchases of mortgages from families in these income groups (Ambrose and Thibodeau 2004). At least since that time, the extent to which the GSEs actually help low-income families has been a subject of fierce political debate, with the GSEs claiming that they have helped marginal borrowers the most and their critics arguing the opposite. Taking the GSEs' claims at face value, perhaps even if they have not improved the system of housing finance for the population as a whole, they have improved things for marginal borrowers. In terms of our model, if the GSEs' claims were correct, we would expect to see significant and positive GSE interaction terms in the models estimated with the poverty, the first-time homebuyer, and the under 40 samples.

Columns (2), (3) and (4) display the estimates of α_3 when we estimate equation (7.1) using the poverty, first-time homebuyer, and younger household samples, respectively. The point estimates are statistically insignificant for all three samples in each of the specifications. The results are similar for our sample of older households (column 5). Hence, we cannot reject the hypothesis that GSE securitization activity played no role in reducing the capital market imperfections facing low-income families, first-time homebuyers, and younger households.

The retained portfolio. In addition to securitizing mortgages, the GSEs issue debt to purchase mortgages and mortgage-based securities. This so-called "retained portfolio" has increased enormously over time. From 1990 to 2003, it increased from 5.2 percent to 22.2 percent of total home mortgages outstanding, as displayed in the top panel of Figure 1 (dotted

line). The retained portfolio has emerged as an extremely contentious issue. Critics of the GSEs, including former Federal Reserve Board Chairman Alan Greenspan, have argued that the hedging activities associated with the retained portfolio are a dangerous source of systematic risk to the U.S. financial system.³⁵ From this perspective, the retained portfolio does nothing to help homeowners; it is merely a way for the GSEs to turn themselves into very profitable but risky hedge funds. In contrast, the GSEs argue that their demand for mortgage-backed securities is an important source of liquidity in the housing market and therefore leads to substantial benefits to homeowners.

Our model provides a natural framework for assessing the impact of the retained portfolio on the financial environment for homeowners. To begin, we construct both stock and flow measures of the retained portfolio, normalized by the total stock of home mortgages outstanding and the total number of new originations of home mortgages, respectively.³⁶ Following the tack described above for estimating the effect of GSE securitization activity, we interact both of these variables with house value and include the interaction terms in our regression models. We then estimate the models using the post-1983 portion of our sample. The results for the flow measure of the retained portfolio are shown in panel 3 of Table 4.³⁷ They provide no evidence that the retained portfolio has made the housing finance market less imperfect for households. More precisely, whether we measure the GSE retained portfolio in terms of stocks or flows, and whether our left-hand-side variable is future income measured as a 5-year average, or the average of two and four years ahead, one cannot reject the hypothesis that the coefficient on the interaction term between the log of house value and the GSE retained portfolio is zero.

7. Conclusions

Taken together, our results suggest the following conclusions: First, the housing finance market has become substantially less imperfect over time. Second, for the population as a whole, there appears to have been a discrete improvement in the housing finance markets in the early to mid 1980s. We conjecture that this was due to a combination of innovative mortgage products, deregulation, and the development of a secondary market in mortgages. Third, one cannot reject

³⁵ See Kopecki (2006). For further discussion of the systemic risk issue, see Jaffee (2003).

³⁶ The series used for constructing these measures were obtained from the Office of Federal Housing Enterprise Oversight, 2005 *Report to Congress*, Appendix Tables 5 and 15.

³⁷ The results for the stock measure are very similar and are omitted for brevity. These results are available upon request.

the hypothesis that the GSEs and their activities in the secondary market have failed to improve the housing finance environment facing low-income and first-time homebuyers.

More broadly, we have argued that a life-cycle approach to thinking about questions regarding housing finance is both theoretically attractive and empirically tractable. This approach takes advantage of the life-cycle prediction that current behavior can predict future income in the presence of well-functioning credit markets. It might be fruitfully applied in other contexts. Possible issues include measuring the “affordability” of housing, assessing the extent of mortgage market discrimination, and other topics as well.

References

- Ambrose, B., and T. Thibodeau. 2004. "Have the GSE Affordable Housing Goals Increased the Supply of Mortgage Credit?" *Regional Science and Urban Economics* 34: 263–273.
- Artle R., and P. Varaiya. 1978. "Life Cycle Consumption and Homeownership." *Journal of Economic Theory* 18: 38–58.
- Bai, J. 1999. "Likelihood Ratio Tests for Multiple Structural Changes" *Journal of Econometrics* 91: 299–323.
- Bartlett, W. 1989. *Mortgage-Backed Securities*. New York: New York Institute of Finance.
- Brueckner, J. 1986. "The Downpayment Constraint and Housing Tenure Choice." *Regional Science and Urban Economics* 16: 519–525.
- Bubnys, E., S. Khaksari, and M. Tarimcilar. 1993. "Intermarket Efficiency: An Application of Interbattery APT to Mortgage-Backed Securities." *Journal of Real Estate Finance and Economics* 7: 99–115.
- Campbell, J., and J. Cocco. 2003. "Household Risk Management and Optimal Mortgage Choice." *The Quarterly Journal of Economics* 118: 1449-1494.
- Card D., and J. Dinardo. 2002. "Skill Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles." *Journal of Labor Economics* 20: 733-783.
- Chomsisengphet, S., and A. Pennington-Cross. 2006. "The Evolution of the Subprime Mortgage Market." *Federal Reserve Bank of St. Louis Review* 88(1): 31–56.
- Cocheo, S. 2003. "25 years later Merrill's CMA Is Still Making Waves." *ABA Banking Journal*. August.
- Devaney, M., and K. Pickerill. 1990. "The Integration of Mortgage and Capital Markets." *Appraisal Journal* 58: 109–113.
- Devaney, M., K. Pickerill, and F. Krause. 1992. "Cointegration and Causal Relations in Mortgage and Capital Markets." *Journal of Financial Services Research* 5: 341–353.
- Dynan, Karen E., Douglas Wl Elmendorf, and Daniel E. Sichel, "Can Financial Innovation Help to Explain the Reduced Volatility of Economic Activity?" Working Paper, Federal Reserve Board, November 2005.
- Dynan, Karen E., Douglas Wl Elmendorf, and Daniel E. Sichel, "Financial Innovation and the Great Moderation: What Do Household Data Say? Federal Reserve Bank of San Francisco Conference: *Financial Innovations and the Real Economy*" November 2006.

- Elliehausen, G., and T. Durkin. 1989. "Theory and Evidence of the Impact of Equal Credit Opportunity: An Agnostic Review of the Literature." *Journal of Financial Services Research* 2: 89–114.
- Engelhardt, G. 1996 "Consumption, Down Payments, and Liquidity Constraints." *Journal of Money, Credit and Banking* 28: 255–271.
- England, C. 1992. "Lessons from the Savings and Loan Debacle: The Case for Further Financial Deregulation." *Regulation: The Cato Review of Business & Government* 15(3).
- Friedman, M. 1957. *A Theory of the Consumption Function*. Princeton: Princeton University Press.
- Goebel, P., and C. Ma. 1993. "The Integration of Mortgage Markets and Capital Markets." *Real Estate Economics* 21: 511–538.
- Gilbert, R., 1986. "Requiem for Regulation Q: What it Did and Why it Passed Away." *Federal Reserve Bank of St. Louis Review* 68(2) 22–37.
- Goodfriend, M., and R.G. King, 2005. "The Incredible Volcker Disinflation." *Journal of Monetary Economics* 52: 981–1015.
- Guttentag, J. 1984. "Recent Changes in the Primary Home Mortgage Market." *Housing Finance Review* 3: 221–254.
- Harrigan, S. 1981. "Wachovia Has Other Bankers Taking Notice." *Wall Street Journal*. April 27.
- Hendershott, P., and R. Van Order. 1989. "Integration of Mortgage and Capital Markets and the Accumulation of Residential Capital." *Regional Science and Urban Economics* 19: 189–210.
- Himmelburg, C., Mayer, C., and T. Sinai. 2005. "Assessing High House Prices: Bubbles, Fundamentals and Misperceptions." *Journal of Economic Perspectives* 19: 67–92.
- Jaffee, D. 2003. "The Interest Rate Risk of Fannie Mae and Freddie Mac." *Journal of Financial Services Research* 24: 5–29.
- Kettel, R., 2006. "Golden West Financial at Credit Suisse First Boston Financial Services Conference – Final." February 8.
- Kopecki, D. 2006. "Greenspan Seeks to Recast Debate on Fannie and Freddie Portfolios." *Wall Street Journal*. January 20.
- Ladd, H. 1982. "Equal Credit Opportunity: Women and Mortgage Credit." *American Economic Review* 72(2): 166–170.

- Lewis, M. 1989. *Liar's Poker: Rising Through the Wreckage on Wall Street* New York: Penguin Putnam.
- Linneman P., and S. Wachter. 1989. "The Impacts of Borrowing Constraints on Homeownership." *Real Estate Economics* 17: 389–402.
- Loutskina, E., and P. Strahan. 2006. "Securitization and the Declining Impact of Bank Finance on Loan Supply: Evidence From Mortgage Acceptance Rates." NBER Working Paper W11983.
- Macauley, I. 1980. "Banks Offer New Choices In Mortgages." *New York Times*. May 25.
- Mason, D. 2004. *From Buildings and Loans To Bail-Outs: A History of the American Savings and Loan Industry, 1831–1995*. Cambridge: Cambridge University Press.
- Mayo, S. 1981. "Theory and Estimation in the Economics of Housing Demand." *Journal of Urban Economics* 10: 95–116.
- McConnell, Margaret and Gabriel Perez Quiros, 2002. "Output Fluctuations in the United States: What Has Changed Since the Early 1980s?" *American Economic Review*, 90 (5): 1464-76.
- Miniaci, Raffaele and Guglielmo Weber, "Econometric Issues in the Estimation of Household Portfolio Models," in Guiso, Luigi, Michael Haliassos, and Tullio Jappelli, *Household Portfolios*, The MIT Press: Cambridge, 2002.
- Modigliani, F., and D. Lessard. 1975. "New Mortgage Designs for Stable Housing in an Inflationary Environment" *Federal Reserve Bank of Boston Conference Series*.
- Morris, F. 1975. "Solving the Long-Range Problems of Housing and Mortgage Finance" in Modigliani and Lessard (1975).
- Munnell, A., Tootell, G., Browne, L., and J. McEneaney. 1996. "Mortgage Lending in Boston: Interpreting HMDA Data." *American Economic Review* 86(1): 25–53.
- Passmore W., R. Sparks, and J. Ingpen. 2002. "GSEs, Mortgage Rates, and the Long-Run Effects of Mortgage Securitization." *Journal of Real Estate Finance and Economics* 25: 215–242.
- Phalon, R. 1977. "A Mortgage Plan To Help the Young." *New York Times*. May 26.
- Polinsky, A., and D. Ellwood. 1979. "An Empirical Reconciliation of Micro and Grouped Estimates of the Demand for Housing." *Review of Economics and Statistics* 61: 199–205.
- Ranieri, L. 1996. "The Origins of Securitization, Sources of Its Growth, and Its Future Potential." In *A Primer on Securitization*, eds. Leon Kendall and Michael Fishman, 31–44. Cambridge: MIT Press.

- Rudolph, P., and J. Griffin. 1997. "Integration of the Mortgage Market into the National Capital Markets: 1963–1993." *Journal of Housing Economics* 6: 164–183.
- Rudolph, P., L. Zumpano, and M. Karson. 1982. "Mortgage Markets and Interregional Differences in Mortgage Terms." *Real Estate Economics* 10: 94–110.
- Shaman, D. 1979. "Mortgage Seekers Find Few Doors Open." *New York Times*. September 30.
- Stahl, D., 1996. "A COFI break." *Mortgage Banking*. November 1.
- Stock, J., and M. Watson. 2003. "Has the Business Cycle Changed? Evidence and Explanations." Prepared for the Federal Reserve Bank of Kansas City symposium "Monetary Policy and Uncertainty."
- Straka, J.W. 2000. "A Shift in the Mortgage Landscape: The 1990s Move to Automated Credit Evaluations." *Journal of Housing Research* 11: 207–232.

Appendix A

Some PSID issues

The PSID tracks members of its first-wave (1968) families, including all those leaving to establish separate family units. Children born to a member of a first-wave family are classified as sample members and are in many cases tracked as separate family units when they set up their own households. Ex-spouses and other adult sample members who move out of PSID family units are also tracked to their new family units. Thus, new PSID families originate from two sources: Children from original PSID families who grow up and establish separate households, and marriage partners who divorce and go their separate ways. This dynamic structure allows two ways to define and identify a household across time. One possibility, which is consistent with the PSID's definition, is to identify a household simply by its head (usually the husband). Such a definition would not factor in a change of spouse. For example, if the husband and wife divorced, and the husband later remarried, the new family created by the new marriage would be identified as the original household. The second possibility, which we adopt, is to identify a household as a unique husband/wife pair. With respect to the above example, if the husband and wife of the original household divorced, then the household would cease to exist in the sample, and a new household would form with the new marriage. We choose to define a household this way because we feel that it is more consistent with our analysis of housing choice and income behavior.

Another important issue is the reliability of the PSID data on homeownership and housing finance. We performed a comparison of homeownership and housing finance data in the PSID with information from other sources, in order to be sure that we are working with a nationally representative sample. Specifically, we compared our annual homeownership rates from the PSID to those produced by the federal government's Office of Federal Housing Enterprise Oversight (OFHEO). During the last three decades, there is never more than a 2-percentage-point difference between OFHEO's annual rates and those from our PSID data. We also compared PSID mortgage data to corresponding data from Chicago Title's Annual Survey of Recent Home Buyers. The PSID has mortgage information for most of our sample period, including the original value of the mortgage and the annual mortgage payment. Average loan-to-value ratios and payment-to-income ratios for first-time homebuyers displayed very similar patterns over time in each data set.

Appendix B

This appendix presents a more formal version of the model discussed in Section 3. We show that in the context of a simple, two-period model, the relaxation of a debt-to-income constraint (DTI), increases the sensitivity of future income to housing expenditures. Our derivation is comprised of two pieces. First, we show that in a world with two possible types of houses (big and small), household wealth determines choice of dwelling, *ceteris paribus*. If household wealth exceeds some threshold level, w^* , then the household buys a big house; if wealth falls short of the threshold, it buys a small house. Further, we show that w^* is weakly decreasing in δ , the maximum allowable ratio of debt to current income. Second we consider a model with households that differ only in the level of their future income, either low (y_1^L) or high (y_1^H), and show that an increase in the maximum allowable debt-to-income ratio, by changing w^* and thus the mix of buyers of houses of different types, increases the sensitivity of future income growth to current house purchase.

B.1: Derivation of the wealth threshold

We consider a deterministic, two-period world in which a representative household faces the following problem. Given current income, y_0 , future income, y_1 , and current wealth, the household must choose between one of two houses ($H^B > H^S$) and consumption today and in the future. Borrowing to finance a house purchase is represented by θ and the amount of the mortgage is constrained to be less than or equal to a constant fraction of current income, δ , the maximum DTI ratio. We also assume a constant interest rate, r , log-utility in which consumption of non-durables and consumption of housing are additively separable, and for simplicity, no discounting over time on the part of households. The household's problem is to maximize

$$\log(c_0^\alpha H^{1-\alpha}) + \log(c_1^\alpha H^{1-\alpha})$$

subject to:

$$\begin{aligned} c_0 &= y_0 + w - H + \theta, \\ c_1 &= y_1 - (1 + r)\theta, \\ \delta y_0 - \theta &\geq 0 \end{aligned}$$

We separate the decision into two parts. First, the household chooses an optimal consumption profile conditional on a particular house. Letting $V(H; \delta, w, y_0, y_1, r)$ measure the indirect utility conditional on house choice H , the household then chooses

the house that maximizes V .

We establish two facts about this problem. First, we show that holding all else equal, there exists a critical or threshold value, w^* , such that if w falls short of w^* , the household chooses to buy a small house, and if w exceeds w^* , the household chooses to buy a large house. Second, the critical value, w^* , is weakly decreasing in δ , the maximum allowable debt-to-income ratio. Taken together, these facts imply that in households that differ only by wealth, an increase in δ raises the fraction of that population that opts for the large house.

To begin, let:

$$\Delta V(\delta, w) = V(\delta, H^B, w) - V(\delta, H^S, w),$$

where for notational simplicity we suppress the additional arguments of V . Our first goal is to show that when $w \geq w^*$, $\Delta V(\delta, w) > 0$ and when $w \leq w^*$, $\Delta V(\delta, w) < 0$. We proceed by setting up a Lagrangian for the conditional problem:

$$\begin{aligned} \mathcal{L} = & \log(c_0^\alpha H^{1-\alpha}) + \log(c_1^\alpha H^{1-\alpha}) + \\ & \lambda_0 \cdot (y_0 + w - H + \theta - c_0) + \lambda_1 \cdot (y_1 - (1+r)\theta - c_1) + \lambda_{DTI}(\delta y_0 - \theta) \end{aligned}$$

plus the usual complementary-slackness conditions. By the envelope theorem,

$$\frac{\partial V}{\partial w} = \lambda_0 = \frac{1}{c_0} > 0.$$

If the household is unconstrained, then λ_{DTI} is zero, and the associated initial consumption is

$$c_{0, \lambda_{DTI}=0} = \frac{y_0 + y_1/(1+r) + w - H}{2} \quad (1)$$

If it is constrained then δ_{DTI} is positive, and initial consumption is

$$c_{0, \lambda_{DTI}>0} = (1 + \delta)y_0 + w - H \quad (2)$$

In order to determine the sign of the term $\frac{\partial \Delta V}{\partial w} > 0$, we need to know the relationship between consumption if the family buys a small house, c_0^S , and consumption if the family buys a big house, c_0^B . To think about this, it is useful to consider the various possibilities with respect to the borrowing constraints that the family faces. If the household is constrained whether it buys the big house or the small house, then equation (1) shows that $c_0^S > c_0^B$. If the household is unconstrained whether it buys the big house or the small house, then equation (2) shows that $c_0^S > c_0^B$. In either case, purchase of a house comes at the expense of current consumption. Another regime

is for the borrower to be constrained in buying the small house and unconstrained in buying the big house, but this is impossible. This leaves only the case in which the household is unconstrained in buying the small house, but constrained in buying the big house – the case which we next explore.

We know that for a given house, unconstrained initial consumption is always larger than constrained initial consumption. That is:

$$c_{0,\lambda_{DTI}=0}^B > c_{0,\lambda_{DTI}>0}^B$$

and

$$c_{0,\lambda_{DTI}=0}^S > c_{0,\lambda_{DTI}>0}^S.$$

Furthermore, by equation (1),

$$c_{0,\lambda_{DTI}>0}^S > c_{0,\lambda_{DTI}>0}^B,$$

implying that

$$c_{0,\lambda_{DTI}=0}^S > c_{0,\lambda_{DTI}>0}^S > c_{0,\lambda_{DTI}>0}^B.$$

Thus, $\frac{\partial \Delta V}{\partial w} = \frac{1}{c_0^B} - \frac{1}{c_0^S} > 0$. That is, as wealth increases, so does the gain in utility associated with buying a big house.

Now, $\Delta V(\delta, w^*) = 0$ implicitly defines a function $w^*(\delta)$. Our goal is to show that $w^{*'}(\delta) < 0$. Differentiating both sides of $\Delta V(\delta, w^*) = 0$ with respect to δ yields

$$\frac{\partial \Delta V}{\partial w} w^{*'}(\delta) + \frac{\partial \Delta V}{\partial \delta} = 0 \tag{3}$$

implying that

$$w^{*'}(\delta) = -\frac{\frac{\partial \Delta V}{\partial \delta}}{\frac{\partial \Delta V}{\partial w}}.$$

We have already shown that the denominator is positive. To finish the proof, we simply need to sign the numerator. Using the envelope theorem again,

$$\frac{\partial \Delta V}{\partial \delta} = [\lambda_{DTI}^B - \lambda_{DTI}^S] y_0.$$

If neither constraint is binding, then $w^{*'}(\delta) = 0$ because both of the λ 's are zero. If the constraint binds only for the big house, then $w^{*'}(\delta)$ exceeds zero because $\lambda_{DTI}^B > \lambda_{DTI}^S$. If the constraint binds for both houses, then we use the first-order condition of the optimization problem, which says that

$$\lambda_{DTI} = [\lambda_0 - \lambda_1](1 + r) = \left[\frac{1}{c_0} - \frac{1+r}{c_1} \right].$$

Thus,

$$\frac{\partial \Delta V}{\partial \delta} = \left[\frac{1}{c_0^B} - \frac{1+r}{c_1^B} - \frac{1}{c_0^S} + \frac{1+r}{c_1^S} \right] y_0 > 0,$$

where the last inequality follows from the fact that when the DTI constraint binds, the budget constraint implies that $c_1^B = c_1^S$ and $c_0^B < c_0^S$. Hence, we have shown that w^* is weakly decreasing in δ , the maximum allowable debt-to-income ratio. Algebraically, $w^{*\prime}(\delta) \leq 0$.

B.2: Effect of relaxing the DTI constraint on different types of households

We now assume two types of households that differ only in their future income, y_1^i , where i is either “high” (H) or “low” (L). We assume that there is a continuum of households with measure 1 of each type, indexed by their wealth w . We assume that the wealth distribution is bounded above by \bar{w} and below by \underline{w} , and that it has a cumulative distribution function that is strictly increasing. For each type of household, we can compute a wealth threshold function, $w_i^*(\delta)$. For simplicity, we make the following three assumptions on the wealth distribution.

$$\bar{w} < w_L^*(\infty), \tag{4}$$

$$\underline{w} > w_H^*(\infty), \tag{5}$$

$$\bar{w} > w_H^*(0) > \underline{w}. \tag{6}$$

Conditions (4) and (5) imply complete separation of types when there are no borrowing constraints. When unconstrained, the richest low type of household buys a small house and the poorest high type of household buys a big house. However, condition (6) guarantees that the constraints matter: If we eliminate borrowing altogether, the richest high types still buy a big house, but the poorest opt for a small house.

In Section 3, we made two claims, which we prove in turn. First, we argued that “the borrowing constraints attenuate the observed relationship between income growth and the size of current home purchases.” To see why, start with in a situation with no constraints. In that case, all high types buy big houses and all low types buy small houses. Thus,

$$Y(H_B, \infty) = y_1^H$$

and

$$Y(H_S, \infty) = y_1^L,$$

where $Y(H_j, \delta)$ is the average future income of buyers of house, $j = B, S$, when the debt-to-income constraint equals δ . Now we introduce a finite debt-to-income constraint. By continuity and condition (6), there exists δ sufficiently small, call it $\hat{\delta}$ such that at least some high types buy a small house, meaning that the pool of small-house buyers is now a mix of high and low types, rather than a monoculture of high types, so

$$Y(H_S, \hat{\delta}) > Y(H_S, \infty).$$

But, the income of the big-house buyers does not change, as condition (4) assures us that only high types ever buy the big house, so

$$Y(H_B, \hat{\delta}) = Y(H_B, \infty),$$

implying that the income gap between big-house buyers and small-house buyers shrinks when we introduce a binding constraint.

Our second claim was that relaxing an existing debt-to-income constraint increased the gap between average income growth of big-house buyers and small-house buyers. To see why, suppose we have an economy with debt-to-income constraint level $\hat{\delta}$, as above. If we raise δ , the number of high types that buy the small house falls, since $w_H^*(\delta) < 0$, and the number of low types stays the same, lowering the average income of small-house buyers. The cross-sectional relationship between house size and income growth increases. Because the average future income of big-house buyers again remains the same, it follows that the gap shrinks when we lower δ .

B.2: Discussion of a more general model

The 2-type model derived above, provides a specific example of how relaxing borrowing constraints will increase the estimated coefficient in a regression of future income on current housing expenditures. The model is intuitively appealing because of its simplicity. However, one may be inclined to ask whether generalizing the model to include more types of individuals and houses will change the substantive conclusions?

In this section, we will provide a very general expression that links changes in borrowing constraints to movements in regression coefficients. In the context of a more general model, this expression provides us with a necessary condition for obtaining an increasing regression coefficient in response to the relaxation of borrowing constraints.

We assume that $\hat{\beta}$ is the ordinary least-squares estimate from a regression of future income, y^f , on the value of house purchases h . We are interested in the direction of impact on $\hat{\beta}$ from the relaxation of borrowing constraints, which corresponds to increasing δ in our model. Thus, we must evaluate the quantity, $\frac{\partial \hat{\beta}}{\partial \delta}$. Since

$$\hat{\beta} = \frac{\hat{c}ov(h, y^f)}{\hat{v}ar(h)} = \frac{\frac{1}{T} \sum_i (h_i - \bar{h}) * y_i^f}{\frac{1}{T} \sum_i (h_i - \bar{h})^2}$$

it is easy to derive the following expression:

$$\frac{\partial \hat{\beta}}{\partial \delta} = \frac{1}{\hat{v}ar(h)} \left\{ \frac{1}{T} \sum_i (y_i^f - \bar{y}^f) * \frac{\partial h_i}{\partial \delta} - \frac{2\hat{\beta}}{T} \sum_i (h_i - \bar{h}) * \frac{\partial h_i}{\partial \delta} \right\},$$

which can be rewritten as,¹

$$\frac{\partial \hat{\beta}}{\partial \delta} = \frac{1}{\hat{v}ar(h)} \left\{ \hat{c}ov(y^f, \frac{\partial h}{\partial \delta}) - 2\hat{\beta} \hat{c}ov(h, \frac{\partial h}{\partial \delta}) \right\}. \quad (7)$$

According to equation (7), the relaxation of borrowing constraints will result in an increase in $\hat{\beta}$ if

$$\hat{c}ov(y^f, \frac{\partial h}{\partial \delta}) > 2\hat{\beta} \hat{c}ov(h, \frac{\partial h}{\partial \delta}). \quad (8)$$

Thus, condition (8) tells us that in order to obtain an increase in $\hat{\beta}$, we must place restrictions on the manner in which we shift the distribution of future income types among houses, as well as the manner in which we change the housing distribution itself in response to lower borrowing constraints. It is easy to see that our simple 2-type model above satisfies these conditions. The left-hand side term is positive, since a portion of high-income growth households move from small houses into big houses in response to lower borrowing constraints, while none of the low-income growth households switch. Thus, the covariance between future income and house movements is positive. Meanwhile, the right-hand side term is negative in the model. Households initially living in big houses are not affected by changes in borrowing constraints, since by assumption they are all unconstrained, high-future income types. However, a portion of the high-future income households initially constrained and living in small houses become unconstrained and move into big houses, in response to lower

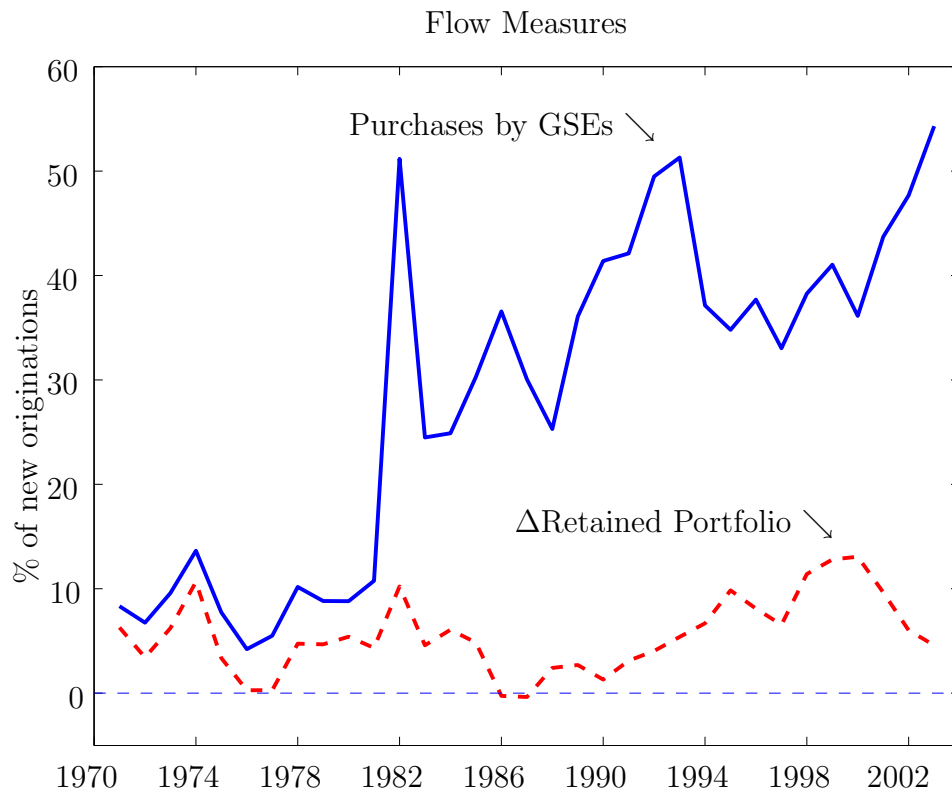
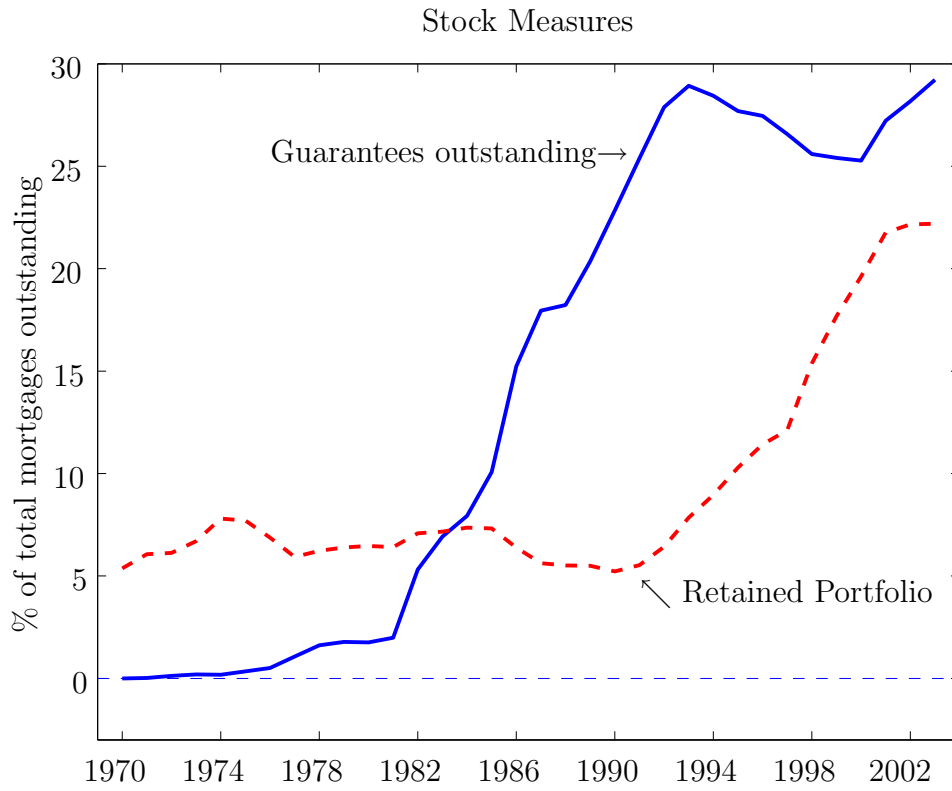
¹Note, that all terms involving partial derivatives of future income drop out, since we are assuming that changes in borrowing constraints do not have a direct effect on income.

constraints. This implies that the covariance between housing expenditures and house movements is negative.

Economic theory suggests that the left-hand side of condition (8) should be positive. The permanent income model of consumption tells us that *ceteris paribus*, households with high, expected future income growth will smooth consumption to a much greater extent than households with low expected income growth. Thus, in the presence of constraints on borrowing, households expecting high income growth are more likely to find themselves borrowing constrained than households expecting lower income growth. This implies that households who are expecting higher income growth will be more affected by the relaxation of constraints.

Theory also suggests that the right-hand side of condition (8) should be negative. If housing consumption is a positive function of permanent income, then *ceteris paribus*, families with higher wealth will prefer to live in more expensive houses. Since households with higher levels of wealth are less likely to be borrowing constrained, then we would expect that on average, households living in bigger houses are less likely to change homes in response to a change in constraints.

Figure 1: GSE Activity Over Time



Sources: *Historical Statistics of the United States* and OFHEO.

Figure 2: Borrowing Constraints, Income Growth and Housing Expenditure




Figure 3: The effect of relaxing constraints in a regression of housing on future income.




Figure 4: The effect of relaxing borrowing constraints on the housing-future income relationship in a more general setting. The key condition for borrowing constraints to raise the sensitivity of house spending to future income is that the smaller the house purchased and the higher the income growth, the more likely it is that a household faces binding constraints. See Section 3 for details.

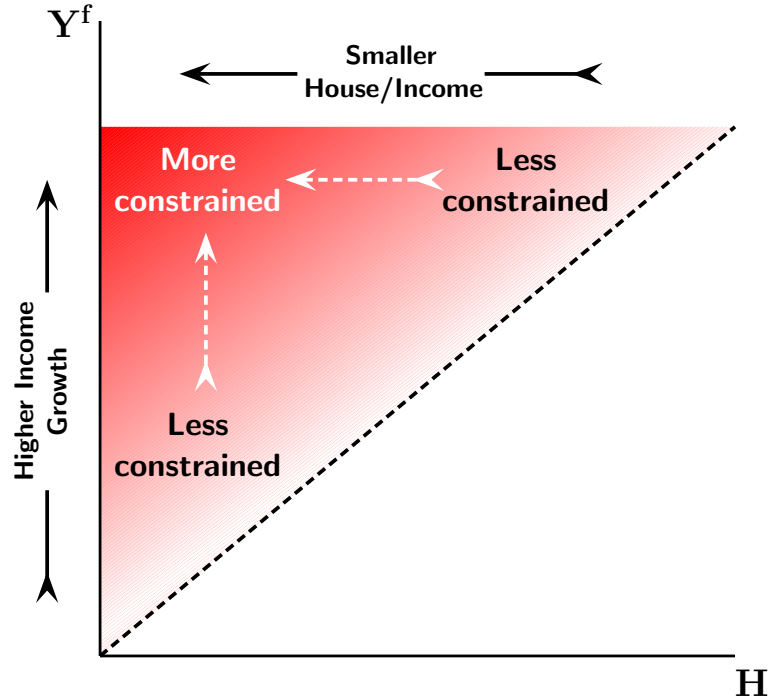
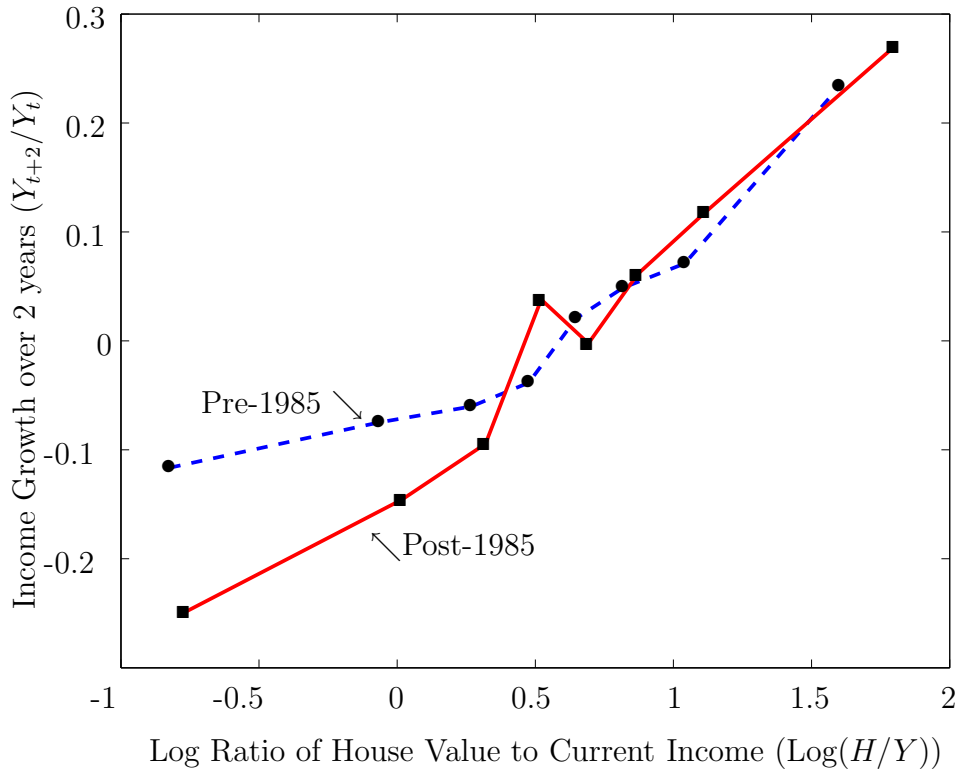


Figure 5: Housing Expenditure and Future Income.



Notes: For each octile of the distribution of the ratio of housing expenditure to income, this figure shows the corresponding average income growth.

Table 1: Summary Statistics.

	(1) Baseline Sample (obs = 5,277)			(2) SEO Poverty Sample (obs = 2,574)			(3) First-Time Buyers (obs = 1,112)		
	Mean	St. dev.	Median	Mean	St. dev.	Median	Mean	St. dev.	Median
Size of household	3.0	1.4	3.0	3.7	1.9	3.0	2.8	1.2	3.0
Age of head	38	14	34	36	12	33	27	4	27
<i>Real Quantities (2000 Dollars)</i>									
Total family income	55,842	43,284	47,296	40,319	27,487	34,347	46,899	23,764	42,918
Labor income	46,439	35,383	41,524	34,084	25,726	30,316	43,482	21,866	40,718
House value	107,247	91,021	88,132	70,284	59,249	56,477	82,903	56,323	74,390
Size of mortgage	78,490	60,274	66,498	58,541	44,384	49,600	67,099	44,357	61,192
<i>Percentages of household heads who are...</i>									
Male		91.0%			80.6%			94.2%	
White		95.3%			38.7%			93.8%	
Black		4.2%			60.1%			5.4%	
Other race		0.5%			1.1%			0.8%	
Educ < High School		6.5%			12.7%			3.6%	
Educ = High School		24.7%			36.7%			25.8%	
High School < Educ < College		37.0%			36.9%			36.4%	
Educ = College		21.0%			8.7%			24.0%	
Educ > College		10.9%			5.0%			10.2%	
First-Time Buyers		21.1%			25.5%			100.0%	

	(4) Age < 40 Sample (obs = 3,453)			(5) Age ≥ 40 Sample (obs = 1,704)		
	Mean	St. dev.	Median	Mean	St. dev.	Median
Size of Household	3.2	1.3	3.0	2.7	1.5	2.0
Age of Head	30	5	30	55	11	52
<i>Real Quantities (2000 Dollars)</i>						
Total Family Income	53,350	34,395	46,232	59,628	55,774	48,892
Labor Income	48,539	31,411	42,804	41,053	41,556	34,866
House Value	103,100	85,765	85,537	113,721	99,570	90,813
Size Of Mortgage	78,095	58,093	67,167	78,793	65,207	63,754
<i>Percentages of household heads who are...</i>						
Male		95.3%			82.2%	
White		95.5%			94.7%	
Black		4.0%			4.8%	
Other Race		0.5%			0.6%	
Educ < High School		3.4%			13.0%	
Educ = High School		24.1%			26.6%	
High School < Educ < College		38.0%			34.2%	
Educ = College		23.8%			15.1%	
Educ > College		10.6%			11.2%	
First-Time Buyers		31.8%			0.7%	

Notes: The baseline sample includes all home buyers in the PSID over the years 1969-2001 excluding those in the SEO Poverty Sample. First-time buyers includes households that bought a house using a mortgage who had never bought a house before. Section 5 provides further details.

Table 2: Estimates of the Basic Model.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Baseline	Poverty	First-Time Buyers	Age < 40	Age ≥ 40	Baseline	Poverty	First-Time Buyers	Age < 40
log(Total Family Income)	0.358 (5.92)	0.32 (5.35)	0.554 (12.93)	0.408 (4.89)	0.308 (3.80)	0.358 (5.92)	0.32 (5.35)	0.553 (13.24)	0.408 (4.89)
log(House Value)	0.131 (5.53)	0.155 (4.59)	0.036 (1.05)	0.098 (3.54)	0.186 (4.63)	0.161 (8.03)	0.141 (4.90)	0.078 (3.39)	0.126 (5.38)
Age	0.047 (3.58)	0.018 (1.08)	0.143 (0.98)	0.2 (1.39)	0.138 (1.81)	0.047 (3.64)	0.019 (1.13)	0.162 (1.15)	0.194 (1.37)
High School	0.085 (2.83)	0.065 (1.95)	-0.008 (-0.13)	0.065 (1.69)	0.078 (1.77)	0.084 (2.83)	0.066 (1.97)	-0.013 (-0.21)	0.064 (1.67)
Some College	0.156 (4.56)	0.187 (4.82)	0.034 (0.51)	0.11 (2.69)	0.196 (3.87)	0.155 (4.57)	0.188 (4.85)	0.030 (0.46)	0.107 (2.68)
College	0.258 (6.45)	0.307 (6.13)	0.194 (3.11)	0.215 (4.59)	0.315 (4.85)	0.258 (6.47)	0.308 (6.15)	0.190 (3.07)	0.211 (4.57)
> College	0.33 (7.49)	0.384 (7.41)	0.219 (3.12)	0.302 (5.93)	0.343 (4.96)	0.329 (7.53)	0.384 (7.46)	0.218 (3.13)	0.298 (5.96)
Black	-0.042 (-1.60)	-0.076 (-3.44)	-0.017 (-0.38)	-0.04 (-1.28)	-0.026 (-0.55)	-0.042 (-1.62)	-0.076 (-3.43)	-0.021 (-0.48)	-0.04 (-1.29)
Other Race	-0.155 (-3.40)	0.084 (1.13)	-0.241 (-3.02)	-0.123 (-1.92)	-0.254 (-4.67)	-0.157 (-3.46)	0.084 (1.16)	-0.240 (-3.09)	-0.124 (-1.98)
Female	-0.337 (7-.71)	-0.379 (-8.69)	-0.245 (-3.37)	-0.334 (-5.63)	-0.341 (-5.99)	-0.337 (-7.71)	-0.38 (-8.70)	-0.247 (-3.41)	-0.337 (-5.66)
Size	0.02 (3.21)	0.021 (3.23)	0.020 (1.08)	0.011 (1.15)	0.033 (3.64)	0.02 (3.16)	0.021 (3.29)	0.020 (1.07)	0.01 (1.08)
Time× log(House Value)	0.004 (3.21)	0.003 (1.84)	0.006 (2.87)	0.005 (2.88)	0.002 (0.98)				
Dummy(Year≥ 1978)× log(House Value)							0.089 (3.46)		
Dummy(Year≥ 1981)× log(House Value)									0.076 (3.24)
Dummy(Year≥ 1985)× log(House Value)						0.064 (3.23)		0.124 (2.51)	
Observations	5,277	2,574	1,112	3,453	1,824	5,277	2,574	1,112	3,453

Notes: Left hand side variable is average of income 2 years and 4 years into the future. Subsamples are described in Section 5 and, briefly, in the notes to Table 1. Columns (1)-(5) provide estimates of equation (4.5) in the text for various subsamples. Columns (6)-(9) provide estimates of equation (5.1) in the text. (There is no second column for the Age ≥ 40 sample because we identified no breakpoints for this sample.) All regressions are estimated using data from 1969 to 2001 and include a constant, a set of year effects, and a cubic in age. Figures in parentheses are t-statistics based on the use of robust standard errors.

Table 3: Selected Coefficients from Alternative Specifications

(1) Left-hand side variable = avg. of 2 and 4 years ahead					
	(1)	(2)	(3)	(4)	(5)
	Baseline	Poverty	First-Time Buyers	Age < 40	Age ≥ 40
(i) log(House Value)	0.131***	0.155***	0.036	0.098***	0.186***
(ii) Time× log(House Value)	0.004***	0.003*	0.006***	0.005***	0.002
(iii) Breakpoint	1985***	1978*	1985	1981**	none
(iv) log(House Value)	0.161***	0.141***	0.078***	0.126***	-
(v) Dummy(Year ≥ Breakpoint)× log(House Value)	0.064***	0.089***	0.124***	0.076***	-
Observations	5,277	2,574	1,112	3,453	1,824
(2) Left-hand side variable = avg. of 5 years ahead					
(i) log(House Value)	0.090***	0.114***	0.006	0.051***	0.164***
(ii) Time× log(House Value)	0.005***	0.007***	0.008***	0.006***	0.001
(iii) Breakpoint	1982***	1977**	1981**	1981***	none
(iv) log(House Value)	0.127***	0.134***	0.067***	0.097***	-
(v) Dummy(Year ≥ Breakpoint)× log(House Value)	0.061***	0.101***	0.087***	0.075***	-
Observations	4,349	2,158	930	2,902	1,447
(3) Left-hand side variable = 2 years ahead					
(i) log(House Value)	0.116***	0.119***	0.068*	0.055	0.204***
(ii) Time× log(House Value)	0.004***	0.005***	0.004*	0.006***	0.001
(iii) Breakpoint	1985***	1980**	1980	1980***	none
(iv) log(House Value)	0.146***	0.134***	0.068***	0.079***	-
(v) Dummy(Year ≥ Breakpoint)× log(House Value)	0.073***	0.099***	0.097***	0.109***	-
Observations	6,717	3,366	1,424	4,335	2,382
(4) Left-hand side variable = 4 years ahead					
(i) log(House Value)	0.101***	0.142***	0.039	0.074**	0.146***
(ii) Time× log(House Value)	0.005***	0.003	0.007***	0.006***	0.005**
(iii) Breakpoint	1984**	1981***	1981	1980	none
(iv) log(House Value)	0.147***	0.103***	0.073	0.107***	-
(v) Dummy(Year ≥ Breakpoint)× log(House Value)	0.089***	0.118***	0.124**	0.091***	-
Observations	5,359	2,638	1,127	3,507	1,852

Notes: Subsamples are defined in Section 5, and, briefly, in the notes to Table 1. Rows (i) and (ii) report results from basic specifications including time trends, corresponding to columns (1) through (5) of Table 2. Row (iii)-(v) exhibit the results of our breakpoint analysis, corresponding to columns (6) through (9) of Table 2. Row (iii) shows the estimated breakpoint. One, two and three asterisks imply that we can reject the hypothesis of no breakpoints at the 10, 5, and 1 percent significance levels respectively. Row (iv) shows the coefficient on the log of house value prior to the breakpoint and row (v) shows the coefficient on the interaction between log of house value and a dichotomous variable which equals one for years starting with the breakpoint. For the coefficient estimates in rows (i), (ii), (iv) and (v), one, two and three asterisks imply that we can reject the hypothesis that the coefficient equals zero at the 10, 5, and 1 percent significance levels, respectively.

Table 4: The Impact of GSE Activity

	(1)	(2)	(3)	(4)	(5)
	Baseline	Poverty	First-Time Buyers	Age < 40	Age \geq 40
(1) GSE secondary market activity (stocks)					
<i>(a) $y_f = \text{avg. of 2 and 4 years ahead}$</i>					
α_3''	0.22	-0.39	0.07	0.06	0.42
p-value	0.23	0.13	0.88	0.78	0.19
observations	2,596	1,225	551	1,652	944
<i>(b) $y_f = \text{avg. of 5 years ahead}$</i>					
α_3''	0.12	-0.05	0.55	0.01	-0.08
p-value	0.51	0.88	0.11	0.94	0.82
observations	1,804	902	405	1,194	610
(2) GSE secondary market activity (flows)					
<i>(a) $y_f = \text{avg. of 2 and 4 years ahead}$</i>					
α_3''	0.21	-0.27	0.32	0.11	0.35
p-value	0.12	0.20	0.16	0.43	0.20
observations	2,596	1,225	551	1,652	944
<i>(b) $y_f = \text{avg. of 5 years ahead}$</i>					
α_3''	0.08	-0.13	0.27	-0.11	0.17
p-value	0.59	0.65	0.32	0.43	0.58
observations	1,804	902	405	1,194	610
(3) GSE retained portfolio activity (flows)					
<i>(a) $y_f = \text{avg. of 2 and 4 years ahead}$</i>					
α_3''	0.00	-1.05	-2.92	-0.65	1.41
p-value	1.00	0.28	0.02	0.38	0.13
Observations	2,596	1,225	551	1,652	944
<i>(b) $y_f = \text{avg. of 5 years ahead}$</i>					
α_3''	-0.88	0.26	-5.19	-1.54	3.79
p-value	0.55	0.91	0.03	0.32	0.19
Observations	1,804	902	405	1,194	610

Notes: For each specification, this table reports the estimate of α_3'' of equation (6.1), which is the coefficient on the interaction of the relevant measure of GSE activity and the log of housing expenditure. The estimates are from models that also include all the covariates in Table 2, which are not reported here.