

SIEPR

policy brief

Stanford Institute for Economic Policy Research

on the web: <http://siepr.stanford.edu>

When Interest Rates Go Up, What Will This Mean For the Mortgage Market and the Wider Economy?

By Michael Best

These days monetary policy observers are all fixated on the same question: “*How soon will interest rates go up?*” Amid this frenzy of crystal ball staring we might want to stop and ask what we think the inevitable interest rate increases will actually mean for the economy. Will higher interest rates slow the economy down significantly? Or even push it into recession? If so, how will they do so and who will be most affected?

One important channel through which interest rates might affect the economy is by raising the cost of borrowing for households. Facing a higher price of debt, indebted households may want to borrow less and spend less on goods and services today. In aggregate, these cutbacks could

have a substantial effect on total expenditure and hence on GDP.

As we learned painfully during the recent financial crisis, if we want to understand households’ debt dynamics, the first place to start looking is the mortgage market. Mortgages account for 74 percent of total household debt in the U.S. In the U.K. that number is even higher: 89 percent of total U.K. household debt is mortgage debt. If households’ retiming of their expenditures and consumption is an important channel through which monetary policy affects the economy, we would expect to see that households respond significantly to interest rates. In a recent paper, James Cloyne (Bank of England), Ethan Ilzetzki

continued on inside...

About The Author

Michael Best is a postdoctoral fellow at the Stanford Institute for Economic Policy Research and a Research Affiliate at the Centre for Economic Policy Research. His research interests are in public economics and development economics, particularly public policy and the housing market, tax evasion, and public procurement in low and middle-income countries. Michael holds a PhD from the London School of Economics and an M.Phil from the University of Oxford.



SIEPR *policy brief*

(LSE), Henrik Kleven (LSE), and I investigate these issues.¹

We develop a new methodology for estimating how household leverage and the timing of expenditures respond to interest rates. Our methodology overcomes problems that have plagued earlier efforts to estimate these crucial parameters for macroeconomic policy analysis. Applying it to the U.K. mortgage market, we estimate that household leverage is quite responsive to mortgage interest rates. We also find that more leveraged households are much more sensitive to interest rates than those that are less leveraged. However, we estimate that households' expenditures are much less sensitive to interest rates and that this sensitivity doesn't vary significantly across different types of households.

These findings raise a challenge for many of the models that economists traditionally use to study macroeconomic policy, which typically rely on household expenditure responding strongly to interest rates.

The Elusive Elasticity

Perhaps surprisingly, given its central importance in macroeconomics, there is no clear consensus among economists on how much households' leverage and expenditure respond to interest rates. Economists summarize this responsiveness in a single parameter, the Elasticity of Intertemporal Substitution (EIS) — the percentage change in a household's allocation between future and current spending in response to a 1 percent change in the interest rate. Despite many decades of work by economists, accepted estimates of this parameter range from 0 (no response) to around 2 (interest rates going from 0 percent to 1 percent leading to a 2 percent drop in present expenditure relative to future expenditure).

The EIS is notoriously difficult to estimate. The simplest approach would be to look at the response of total consumption as the central bank changes interest rates. As all of our econometrics students would (hopefully) tell you, this approach is problematic. The interest rate is endogenous: The Fed is continuously responding to economic conditions. A change in consumption coinciding with an interest rate hike could reflect the reaction of

the economy to that change — or instead the response of the Fed to the state of the economy.

Correlating consumption with interest rates also faces a measurement problem. Different consumers face different interest rates and it is unclear which aggregate interest rate best captures the rate to which they are responding. For households with significant equity holdings, the expected return on a stock market index might be most relevant for their choices. Other households might be highly indebted so that they are particularly exposed to credit card rates. All in all, these methodological problems have conspired to shroud the elusive EIS from researchers.

In our paper, we take a new and different approach to this question that overcomes these issues; our approach doesn't rely on changes over time in consumption or interest rates, and we are able to observe directly the relevant interest rate.

Leveraging Leverage Notches

While in many ways the mortgage market in the U.S. is very similar to that in the U.K., we take advantage of a couple of key differences. In the U.S., mortgage rates depend on borrowers' characteristics and the types of homes that they wish to finance in very complex ways. By contrast, in the U.K.,

¹ "Interest Rates, Debt and Intertemporal Allocation: Evidence From Notched Mortgage Contracts in the United Kingdom" *Bank of England Working Paper #543*, August 2015.

mortgage rates depend on only a few key characteristics of the borrowers and their homes. The most important of these is the Loan to Value (LTV) ratio of the mortgage — the fraction of the value of the home that the borrower is asking for.

When shopping for a mortgage, borrowers face a menu of available mortgage contracts, each of which has a maximum allowable LTV ratio. All else equal, banks charge higher interest rates at higher LTVs. For example, this means that a borrower faces a higher interest rate on her entire mortgage if she borrows 75.1 percent of the value of her

home than she would if she borrowed only 75 percent of the value of her home. Figure 1 shows the interest rate schedule for different LTV ratios and clearly shows that there are discrete jumps (“notches”) at the common LTV thresholds (60, 70, 75, 80, and 85 percent).²

This step-function schedule of interest rates gives borrowers extremely strong incentives to reduce their LTV to stay below the notches. By moving below a notch, they can save up to 0.5 percent in interest on the entire

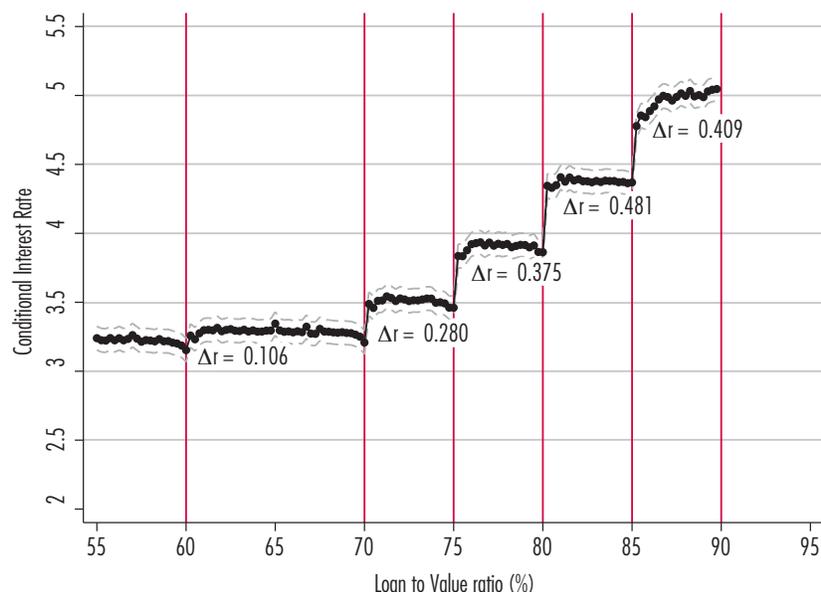
² Contracts with LTV caps at 90, 95, and 100 percent also exist, but there are very few of them during the period our data covers (from September 2008 to December 2014). They were much more prevalent before the housing crisis.

value of their mortgage. So, we should expect disproportionate numbers of mortgages with LTVs just below the notches.

Figure 2 shows that this is exactly what we see. Using data on every mortgage contract origination in the U.K. between September 2008 and the end of 2014, we plot the number of mortgage contracts at each LTV ratio, revealing sharp bunching of contracts exactly below the LTV notches.

If there weren't any notches, we would expect the number of mortgages to evolve smoothly as LTVs increase, without the sharp jumps at the critical LTV ratios. Using information about borrowers' current and previous mortgages, we estimate the LTV they would have chosen without the strong incentives to reduce their LTV to below a notch. The red crossed line in Figure 2 shows our estimate of what the distribution of LTV ratios would have looked like if there weren't any LTV notches in place. As expected, it is smooth and doesn't have a disproportionate number of borrowers at the critical LTV ratios. By comparing the observed distribution of LTVs with this "counterfactual," we are able to infer how much households are willing to reduce their LTV in order to benefit from lower interest rates.

Figure 1
Interest Rates Depend on the LTV with Notches at Contracts' LTV Caps



Now we can see how we have utilized notches to overcome the problems with other methods. By focusing on a decision that a household makes at a single moment in time, we don't need to rely on changes over time in consumption or interest rates. Instead, we rely on interest rate jumps at critical LTV ratios that are exogenous, i.e. not affected by individual choices or the central bank's responses to changing economic conditions. Studying households' responses to these jumps in interest rates overcomes the traditional challenges to estimating the response of debt to interest rates. Furthermore, since we

observe the mortgage interest rates that households could choose, we can be sure we have measured the relevant interest rate for mortgage decisions.³

We estimate that on average households reduce their LTV by $\frac{1}{4}$ percent for every 1 percent change in interest rates. This isn't (yet) an estimate of the Elasticity of Intertemporal Substitution (EIS) described above, but for many policy applications we are interested directly in how households' leverage responds, so this

³ Moreover, relative to other studies of notches (in, for example, tax codes or insurance contracts), we develop a new methodology to create the counterfactual distribution, using the panel nature of our data. The interested reader can refer to our paper for methodological details.

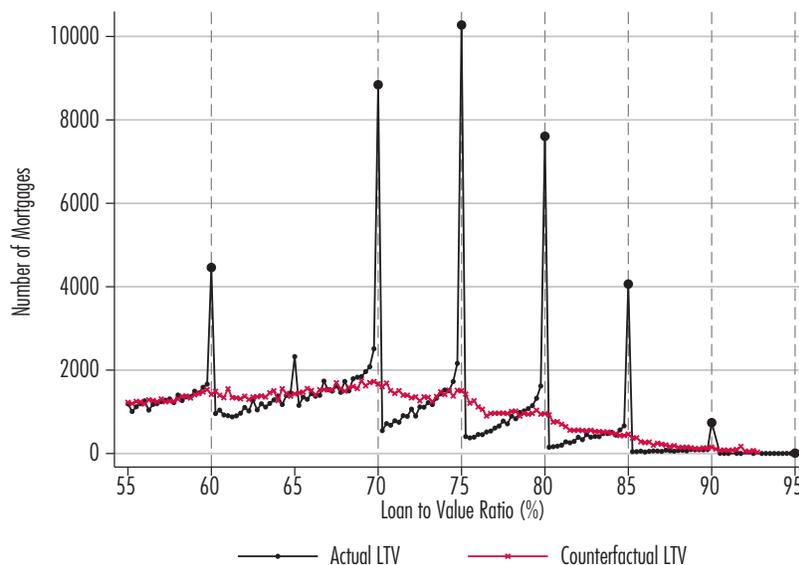
estimate already allows us to make progress in evaluating those policies. Interestingly, the average response we estimate conceals considerable differences among different households. Specifically, in response to a 1 percent jump in interest rates, households with LTVs around 85 percent reduce their borrowing by 1.37 percent, while households with LTVs around 60 percent decrease their debt by only 0.07 percent.

Leverage Responses and the Real Economy

Our methodology for analyzing notches in the mortgage interest rate schedule gives us compelling estimates of how interest rates affect households' leverage decisions. However, households can reduce the LTV on their mortgage in two ways. They can either reduce the amount they borrow (the effect we are interested in) or they can buy a more or less expensive house. To study how households' expenditure and consumption respond to interest rates we want to isolate the effect on the amount they borrow. To achieve this, we focus on households that are refinancing their home. Since refiners stay in the same home while refinancing, their reduction in

continued on flap...

Figure 2
Refinancers Reduce Their LTV to Bunch at Notches



LTV to get below the notches can only be because they are taking out smaller loans, precisely the effect we're looking for. Luckily, in the U.K. refinancing is very common, with households refinancing on average every 3 years, so by studying refincancers we are not focusing on households that are somehow unusual and therefore uninformative for the broader question we're asking.

It is not surprising that different households respond differently to interest rate incentives. Households with high LTVs are making their choices in potentially very different settings than households with low LTVs. Is there a way to contextualize these highly varied responses? This is where the Elasticity of Intertemporal Substitution (EIS) enters. Armed with our estimates of refincancers' leverage and debt responses and data about their specific circumstances, we build a model of household borrowing decisions. Using this model, we are able to transform our estimates of leverage responses into estimates of the EIS — the key parameter we want to estimate.

Using our model, we find

that the EIS is around 0.1, at the lower end of the range of existing estimates. The notches give households enormously strong incentives to reduce their LTV, so even the striking amount of bunching we see in Figure 2 only translates into a modest elasticity when viewed through the lens of our model. Any economic model has numerous assumptions, but our estimates are extremely robust to changing modeling assumptions, which is reassuring. Moreover, the EIS is remarkably stable across different groups of households despite the big differences in leverage responses across different households. This illustrates how households with similar underlying desires to respond to interest rate incentives can display very different behavior when they are put in different situations, an important consideration if we want to extrapolate from our estimates to alternative policy scenarios.

Implications

Our research shows that households' plans about how to allocate their debt, spending, and consumption over time do not respond very strongly

to interest rates. This echoes a growing realization in public economics that individuals respond only slightly to price incentives. The small response to interest rates has important implications both for academic economics and for our analysis of economic policy. Traditional economic models, used by academics and central banks to think about the economy's boom and bust cycles, tend to rely heavily on households reallocating consumption over time in response to shocks. But if the EIS is small, then this effect will be small, creating a problem for these models.

If interest rates have the power to affect the economy, it must be through some other channel. Since the onset of the financial crisis, the economics profession has risen to the challenge and a blossoming of creativity has been considering channels like household leverage, asset prices, and housing values as propagators of business cycles. Given how little interest rates influence households' consumption choices, these new approaches may allow us to provide new answers to age-old questions. Luckily, there are plenty of places for us to start looking.