the pipeline

ANDREW DAVIDSON & CO

ISSUE NO. 181 | February 2023 | SPECIAL SFVEGAS EDITION

THE MEANDERING S-CURVE: UNRAVELING DYNAMIC PREPAYMENT INCENTIVE BEHAVIOR

By Tom Parrent

Summary

At Andrew Davidson & Co., Inc. (AD&Co), we believe that mortgage borrowers behave consistently across time when faced with similar circumstances and incentives. Relying upon this belief, we build stable, reliable, and accurate prepayment and credit models, so that we do not create models that constantly chase changing behavior. While this understanding of borrower behavior is backed by sound economic and behavioral theory, as well as copious amounts of historical performance data, we also have a responsibility to our clients to continually challenge accepted ideas. We look for patterns that both support and contradict our assumptions; this helps guard against complacency and keeps us on our analytic toes. In this Pipeline article, we explore patterns that could contradict our most fundamental assumption about consumer responses to financial incentives.

Introduction

The refinance incentive function is key to all mortgage prepayment modeling. Theoretically, the curve should look something like Figure 1, thus leading to its common name, the "S-Curve." If empirical results lined up nicely with such theoretical constructs, we could quickly build a stable prepayment model and move on to more interesting activities. However, there are two problems: First, prepayment behavior is far more complex than the idealized case; and second, call us strange, but the analysts at AD&Co cannot imagine much that is more interesting than solving the puzzle of borrower behavior.

Figure 1. Simple Refinance Incentive S-Curve

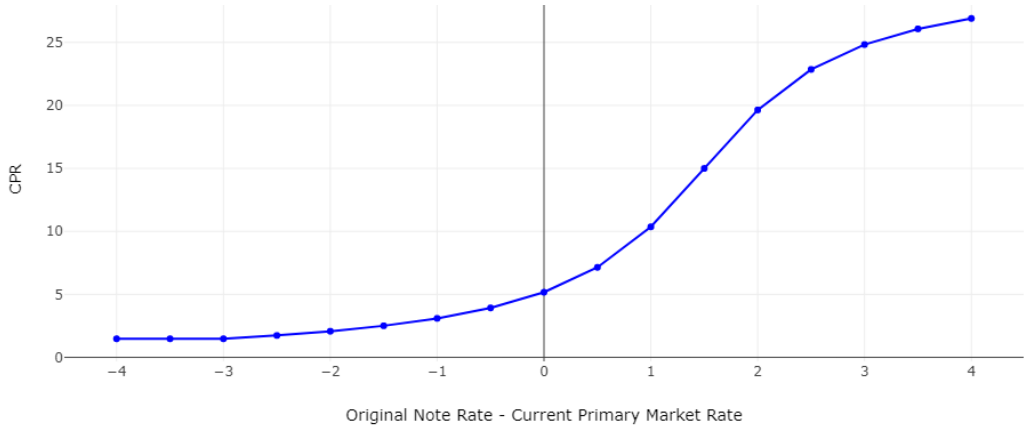

We invite you to grab an analytics raft and join us on a journey of discovery, as we drift along with the "meandering S-Curve," trying to figure out just what makes it move in such odd ways. You will find that the S-Curve is very dynamic, and unraveling its secrets requires meticulous step-by-step analysis. We hope you will find both the journey and the destination enjoyable.

S-Curve Details

We start by looking at the theoretical refinance incentive curve in a bit more detail. In Figure 1 above, the x-axis shows the difference in primary mortgage rates between the time a borrower took out the loan and the observation date. Presented as "Original Note Rate – Current Primary Market Rate," the refinance incentive tells us that borrowers who can lower their payment by refinancing at rates lower than their loan rate will prepay faster than those without such a rate incentive. That behavior is represented by the ascending curve starting around zero and going up and to the right. These loans are considered "in the money."

On the other hand, if current rates are significantly above original note rates, prepayments will be quite slow. Some borrowers will still prepay—due to selling their house, changing loan terms, and a variety of other reasons—but lack of rate incentive will result in generally low and fairly stable prepayment rates. This behavior is referred to as "turnover," representing loans that are prepaid for reasons other than rate incentives.

How do Borrowers Actually Behave?

The S-Curve we presented in Figure 1 was a simple construct. In fact, we know that quite a few factors add complexity. Burnout for deep-in-the-money loans and lock-in effects for out-of-the-money loans are just two such factors. Time now to embrace all the complexity by using real data! For this analysis we use loan-level FNMA 30-year fixed rate, conforming, single family loans originated from January 2010 through May 2022 with observation dates spanning January 2016 through May 2022.

Figure 2 presents actual prepayment rates against the Original Note Rate – Current Primary Market Rate refinance incentive definition. We see that the general shape conforms to the theoretical S-Curve, but there are some significant differences. First, the line drops down on the far right when there is high refinance incentive. It also continues down on the left side of the chart rather than remaining constant. The forces driving these factors, prepayment burnout and turnover, respectively, deserve their own analysis in a future article.

Figure 2. Single Family FNMA 30 Year Loan Observation Dates: 201601 - 202205

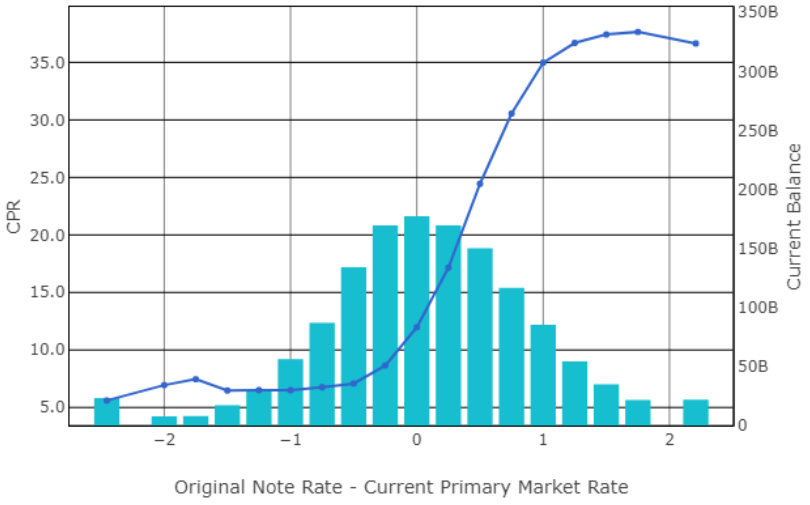

Is Borrower Behavior Consistent Across Time?

When formulating theoretical models, assuming constant behavior of borrowers in response to a given incentive seems reasonable and conveniently simplifying. If borrowers with a two-point rate incentive at a particular time prepaid at an average rate of, say, 35% CPR, we could assume that a different set of borrowers with a two-point rate incentive at another point in time would also prepay at a rate close to 35%.

Figure 3 shows us this assumption would lead us dangerously astray! We see here that the S-Curve meanders all over the chart from one observation year to the next. Loans with a 1.1875% rate incentive in 2016 prepaid at 24% CPR while loans with the same rate incentive paid at 41% CPR in 2020!

Figure 3. Observation Year Refinance Incentives S-Curves

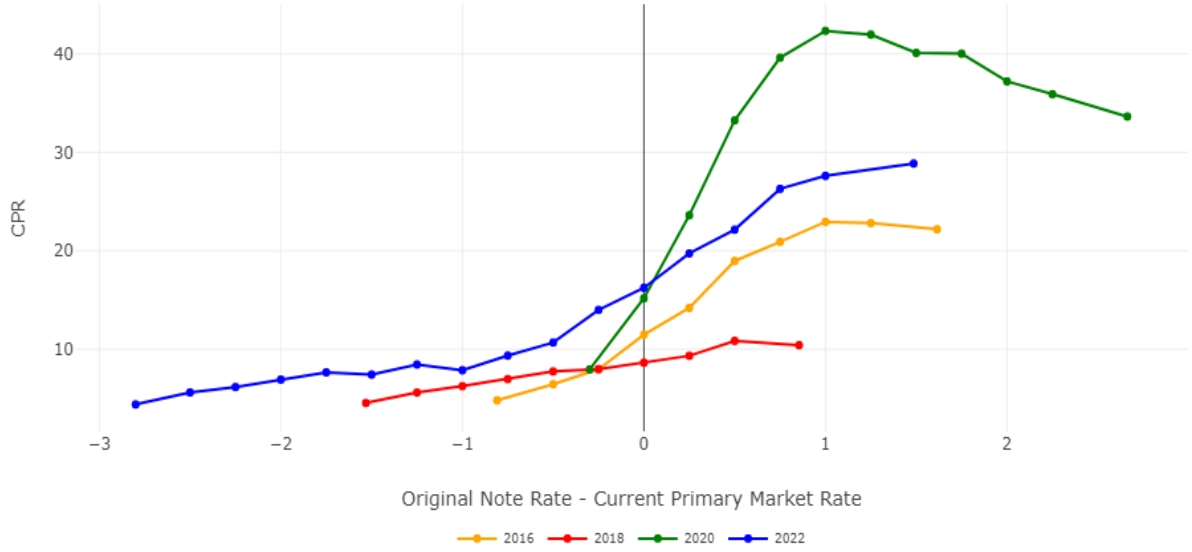

Why would a group of borrowers behave so differently in one year compared to another year? Analytic questions like this set us happily humming away at our keyboards every day! How should we approach this problem? We know that the assumption of constant response to incentive was incorrect. Perhaps we should investigate other assumptions as well.

First, let's examine the implicit assumption that the same borrowers are reacting differently over time. Certainly, borrowers are reacting differently, but the assumption that they are the same borrowers is incorrect! A borrower with a 1.2% refinance incentive in 2016 who actually repaid their loan in 2016 is almost certainly not in the borrower cohort with 1.2% rate incentive in 2017. Even if they are, their loan is less aged than that of those who did not prepay at the same incentive level in 2016.

Looking at the distribution of rate incentives across the observation years in Figure 4, we see that borrower composition must have changed. While movement along the refinance axis does not show that, the volumes do! Some borrowers left the population while many others entered, leading to a net increase in aggregate balance.

Figure 4. Current Loan Balance Distribution by Refinance Incentive and Observation Year

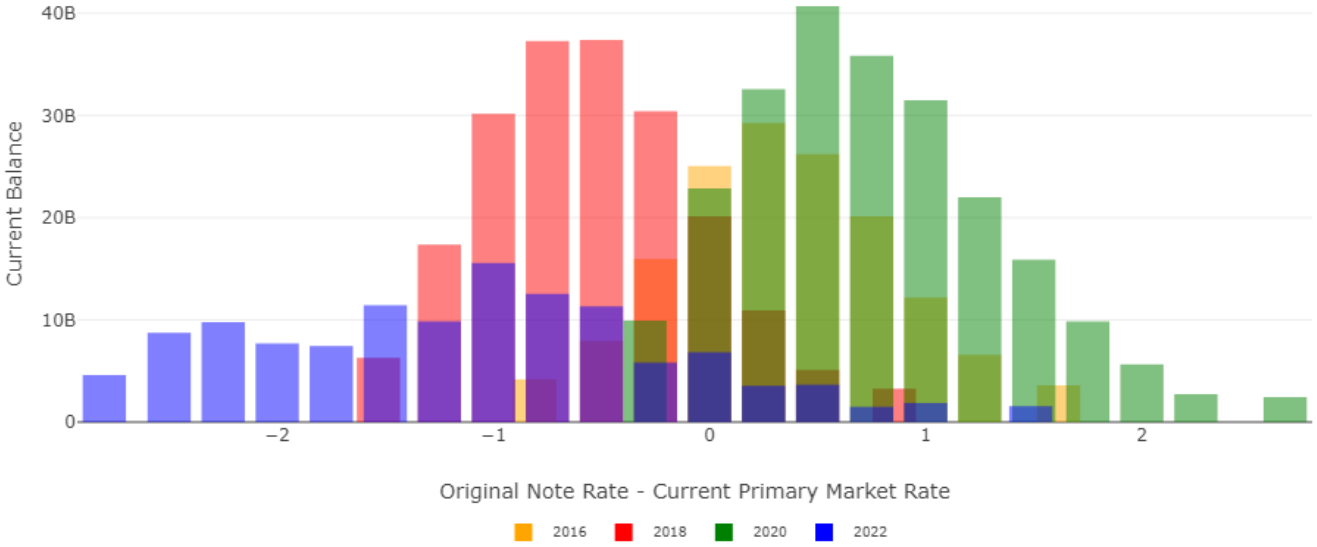

Let's clean things up a bit by grouping the observations from 2016 through 2019 as a single cohort and the observations from 2020 through 2022 as a second cohort. Given those dates, we might be tempted to label them as the "Pre-Pandemic Cohort" and the "Pandemic Cohort," but that might imply potential causality, an assumption without evidence at this point. We will come back to that later.

Furthermore, we limit the refinance incentive to rate differences between -1.6 and +1.6. This ensures that we have sufficient data in the range without the distraction of small volume bins. Figure 5 shows us that these cohorts are now fairly aligned in terms of refinance incentive distribution while still displaying a very significant difference in actual CPR levels for in-the-money loans.

Figure 5. Observation Year Group Refinance Incentive S-Curve

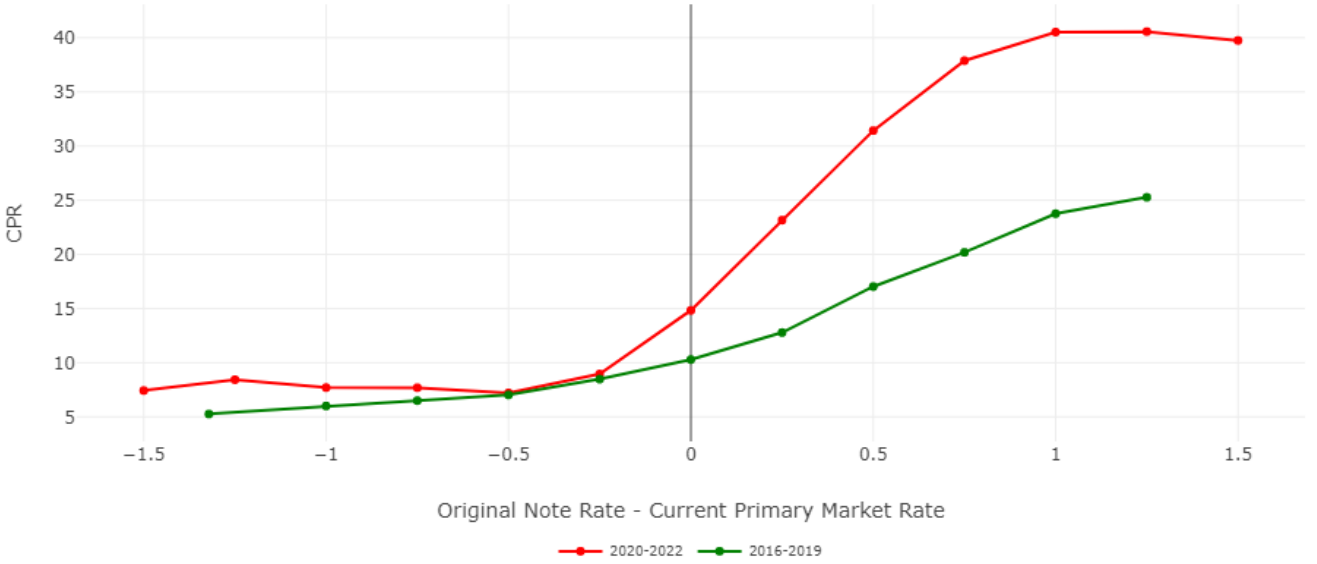

Figure 6. Current Loan Balance Distribution by Refinance Incentive and Observation Year

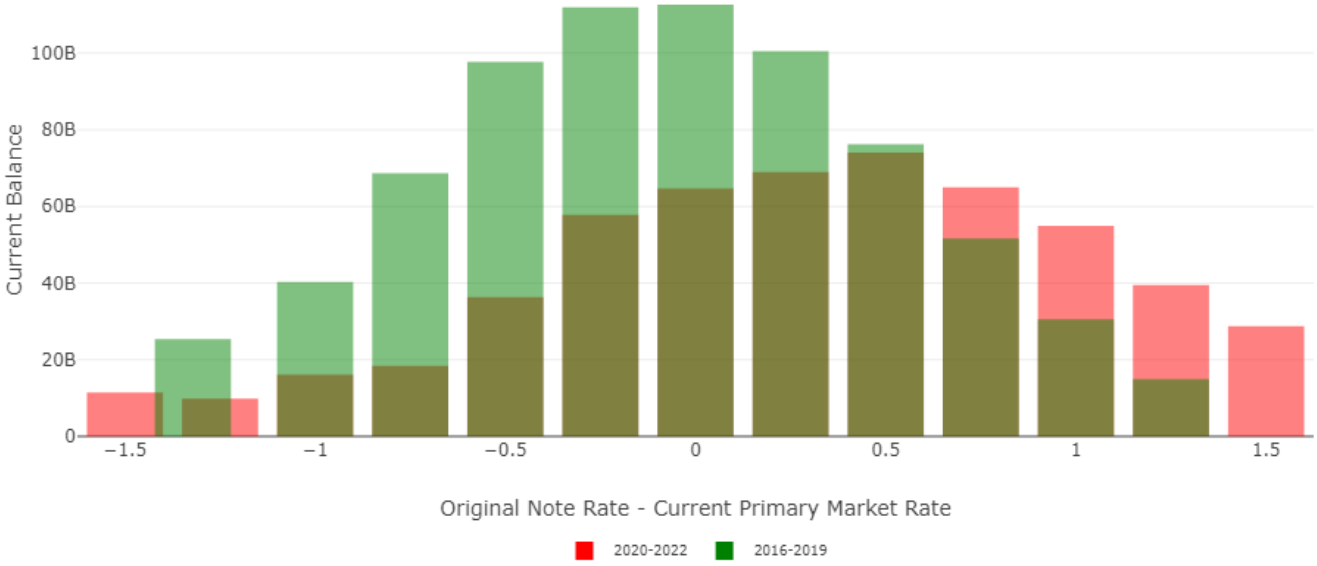

Digging into the Details, One Factor at a Time

While the individual borrowers in the two observation period cohorts are not likely the same, analytically we are interested in the composition of groups of borrowers across time. We will now take you through several possible explanations for the meandering behavior of the refinance incentive, focused on different borrower characteristics that we know drive prepayments.

We will look at each characteristic in the same manner. First, we analyze whether values of a characteristic are associated with different prepayment rates relative to other values of that characteristic. We then check to see if the two characteristic cohorts behaved similarly over the two time periods of interest. If they did, then that particular variable is unlikely to be the cause of the S-Curve shift.

Let's start with a simple example: credit score. From prior research, we know that higher FICO borrowers tend to prepay faster than lower FICO borrowers. Could this cause the behavior we have seen for the two observation period cohorts? Maybe the market portfolio simply had significantly different average credit scores over the two time periods.

Figure 7A show us that FICO does, indeed, affect prepayment speeds along the refinance incentive dimension. Higher FICO loans are more sensitive to rate incentive than lower FICO loans, with faster prepayments in the money and slower prepayments out of the money. Perhaps there were more lower FICO loans in the early period relative to high FICO loans than in the 2020–2022 period. That could explain the shift in the S-Curve!

Figure 7.a. FICO Cohorts

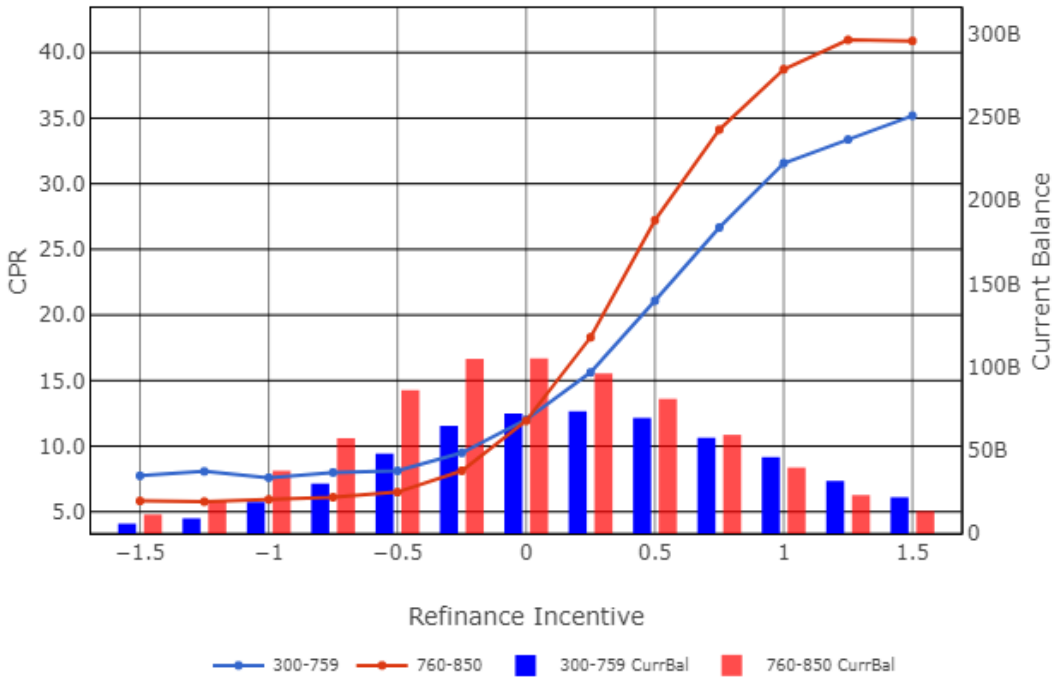

If our hypothesis is that the difference in the two periods is a result of relative FICO volumes, we should see a stable reaction to incentive across the two periods for each FICO cohort. Figures 7B and 7C, alas, show us that even within each FICO cohort, the S-Curve shifted.

Figure 7.b. FICO = [300 – 759]

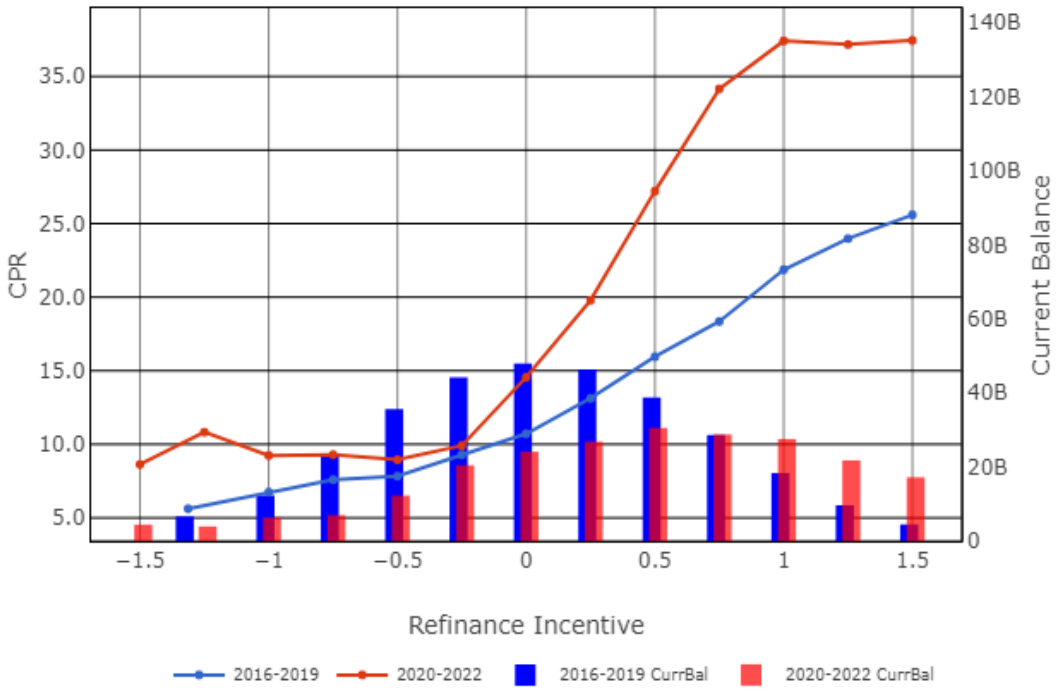

Figure 7.c. FICO = [760 – 850]

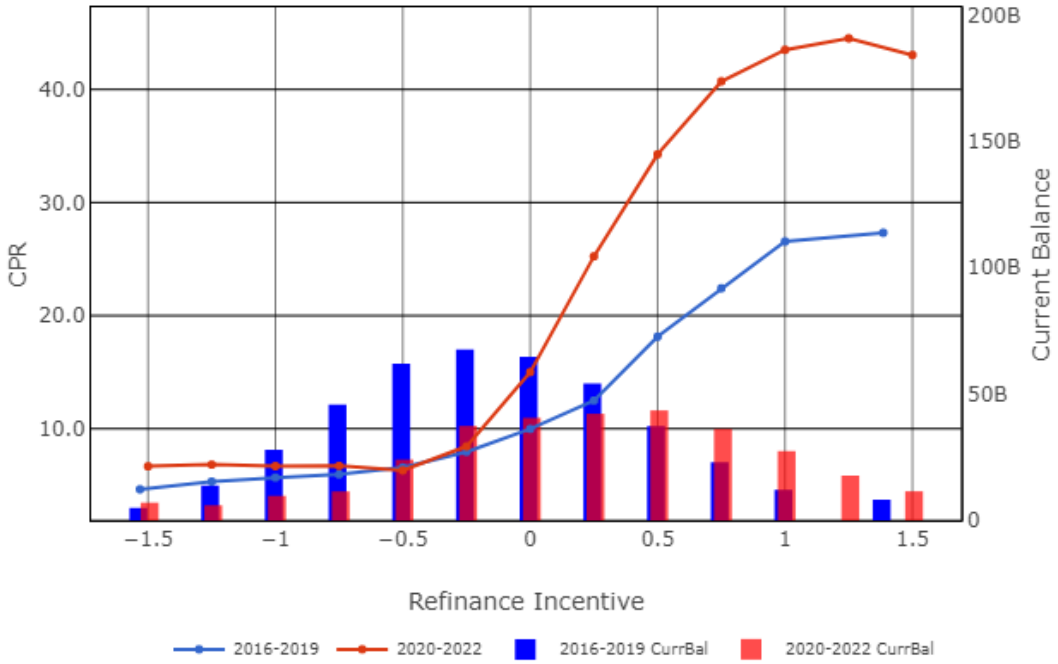

We have plenty more factors to look at in our search for the source of the meandering. Original loan size has potential. A borrower with a large loan has a greater dollar incentive to refinance given the same change in rate than a lower loan borrower does. Maybe we will see that stable behavior within the loan size factor cohorts.

We perform the same analysis we did for FICO, first looking at results this time for loans less than $300,000 compared to those greater than $300,000. Figure 8A shows we are off to a good start with nice separation between the loan size cohorts.

Figure 8.a. Original Loan Size Cohorts

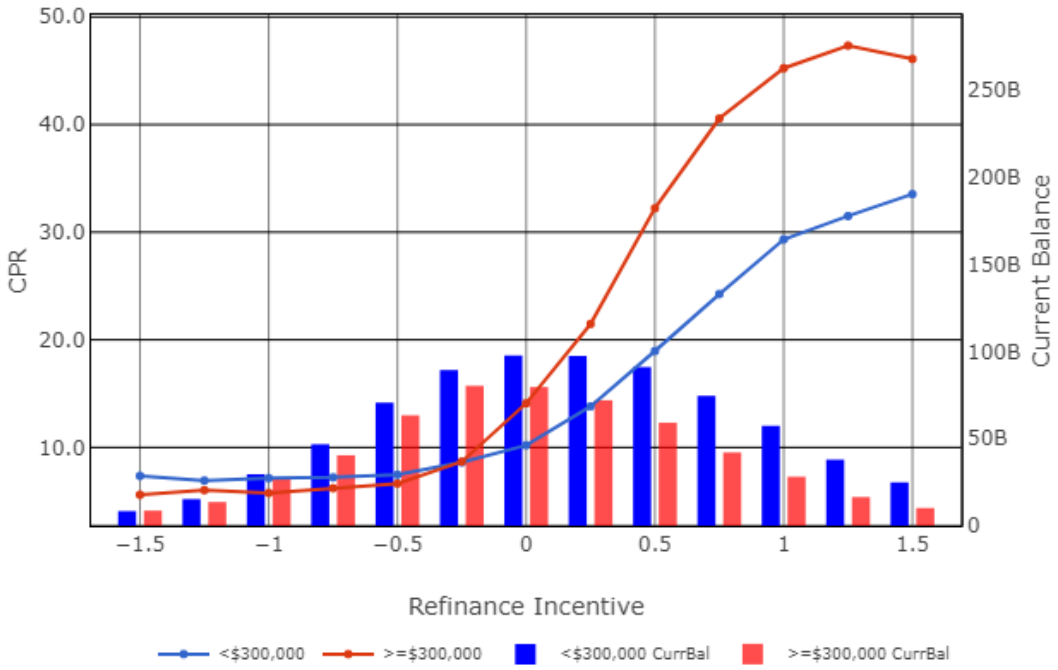

All we need now is some stable profiles across time and… we don't get it. Figures 8B and 8C, once again, show both loan size cohorts meandering. That's disappointing! Such is the nature of research. Formulate a hypothesis, test it, and move on if the evidence does not support it.

Figure 8.b. Original Loan Size < $300,000

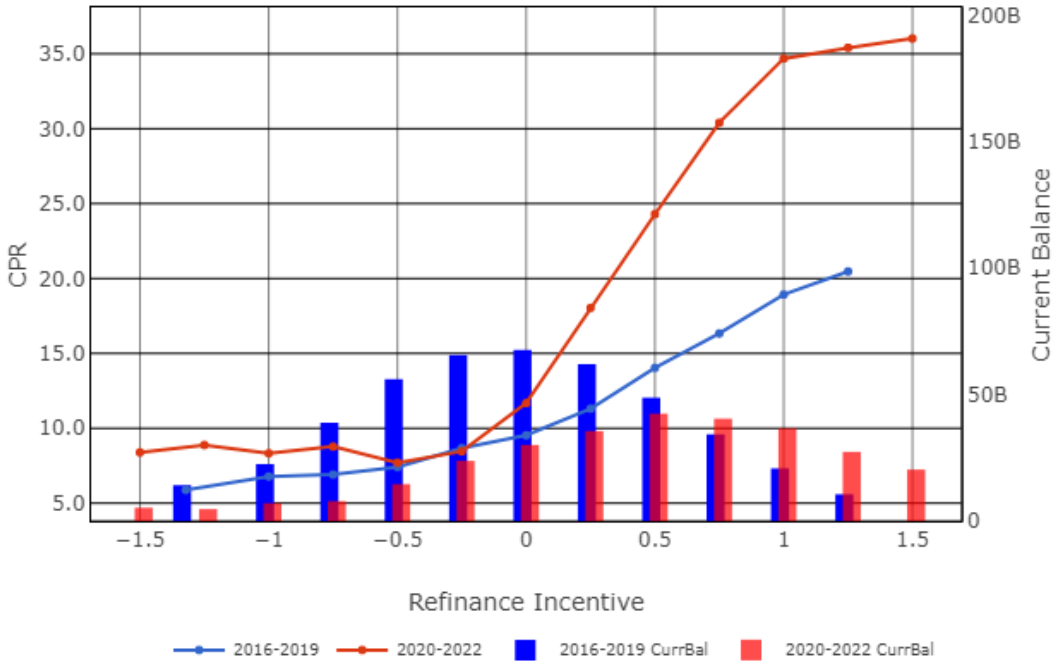

Figure 8.c. Original Loan Size ≥ $300,000

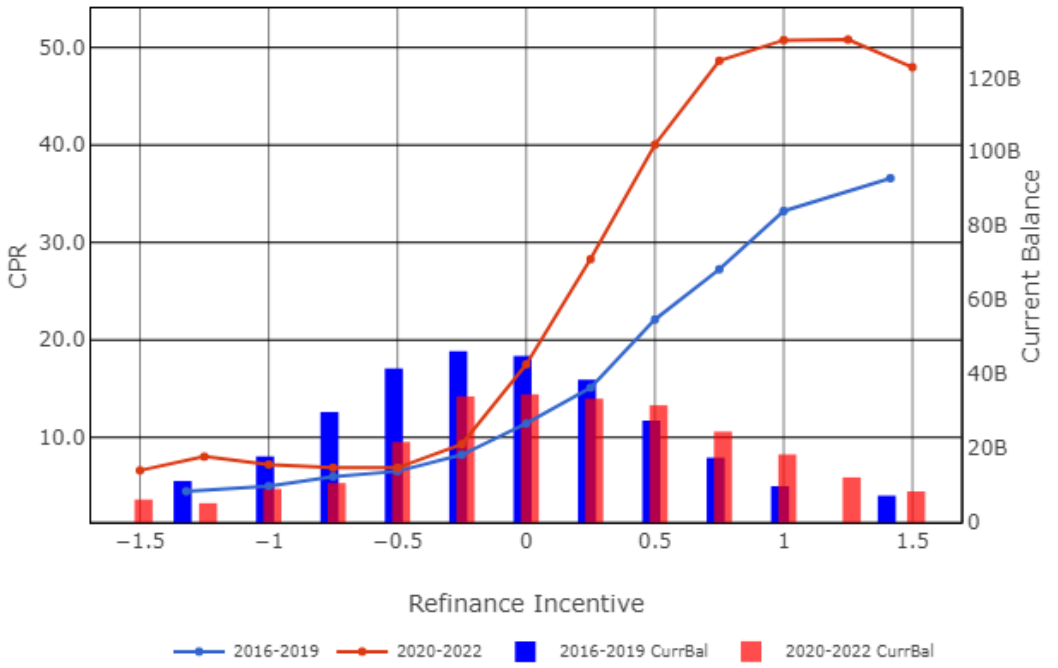

We evaluated many other loan factors, even combining several, such as high FICO and large loan size, all with the same basic results. If borrower and loan factors cannot describe the prepayment sensitivity shift, where else can we look?

That Pesky Pandemic

We mentioned earlier that the curve seems to meander especially strongly between pre-pandemic years and pandemic years. Could there be something in the air pushing the curve around? Instead of looking at loan and borrower factors, perhaps we should evaluate macroeconomic issues!

Several things happened beginning in 2020 that could have altered prior prepayment behavior. First, there was a spike in unemployment, but then forbearance programs muted that effect. Second, interest rates dropped, but that is already covered by the definition of refinance incentive. Finally, house prices rapidly increased.

Perhaps homeowners took advantage of house price appreciation to move or refinance and take out equity. Could the increase in house prices have interacted with the opportunity to refinance at lower rates? We can test that!

We now look at the change from original loan to value (OLTV) ratio to current loan to value (CLTV) ratio, which is affected primarily by house price changes. As the difference [OLTV-CLTV] increases, borrowers can do cash out refinances or perhaps eliminate the need for private mortgage insurance.

We picked out two changes in LTV—no change and a 20 point decrease—to test the hypothesis. Figure 9A shows good separation, with the large decrease in LTV prepaying faster at all rate incentives compared to the no change cohort.

Figure 9.a. LTV Change Cohorts

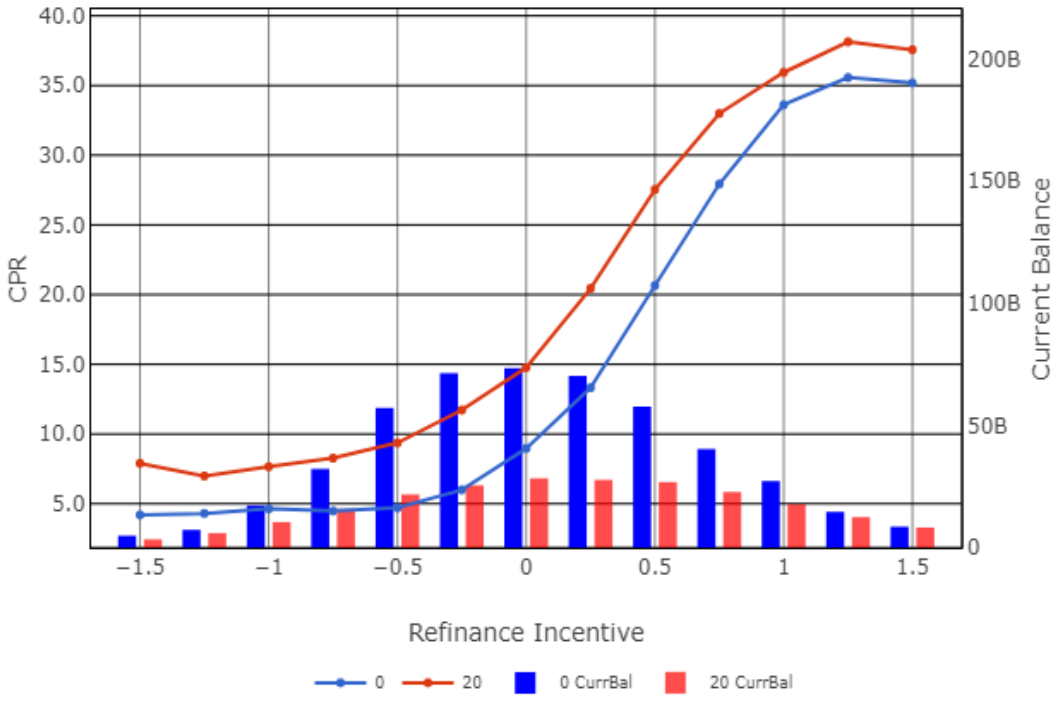

However, Figures 9B and 9C disappoint us yet again, as both cohorts show significant increases in prepayment sensitivity.

Figure 9.b. LTV Change = 0

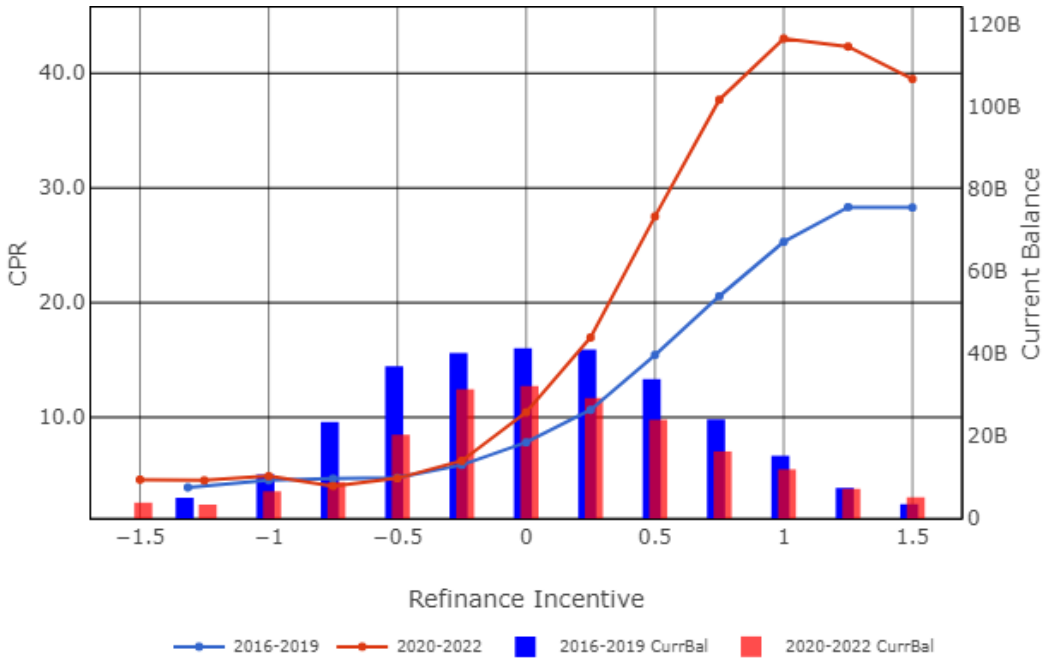

Figure 9.c. LTV Change = 20

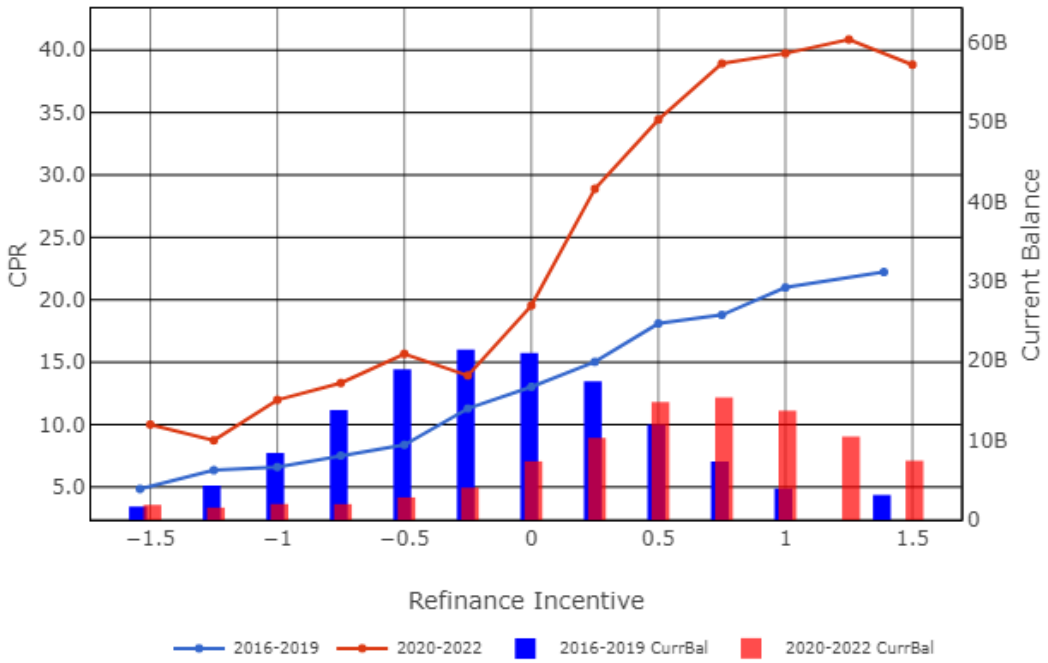

Back to the Drawing Board/Doodling Away

If the change cannot be traced to loan or borrower factors and the largest macroeconomic factor didn't help, what are we left with? Flipping through charts, the pattern was repeated over and over—until something odd caught our eye.

We know that, all other things held constant, prepayments increase as groups of loans age, especially in the first couple of years. That's a fairly fundamental result. Borrowers have just incurred loan origination expenses,

purchasers are less likely to move (thus paying off their loans) soon after closing, and people may simply not want to go through the hassle of applying for another loan. Imagine our surprise when we saw Figure 10A.

Figure 10.a. Loan Age Cohorts

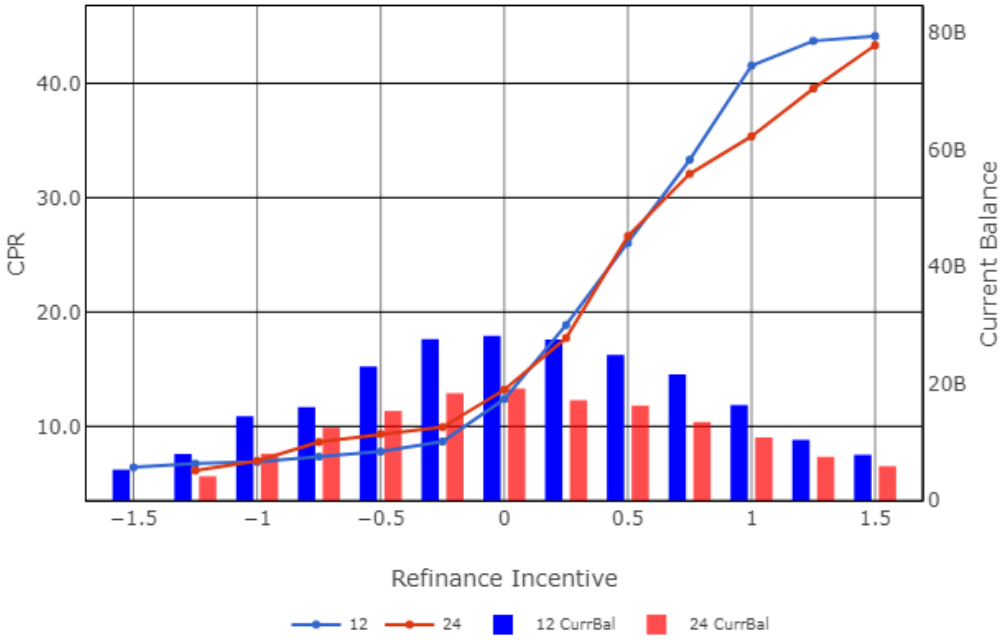

The 0–12 month old loans should definitely be paying slower than the 13–24 month old loans, yet their refinance S-Curves were practically on top of each other for most of the refinance incentive spectrum. Up to this point, we expected to see divergent curves for each of the factors across the whole time period, but had hoped that one factor value would be stable across both periods. We quickly checked the splits across the time periods for loan age.

Figure 10.b. Loan Age ≤ 12 Months

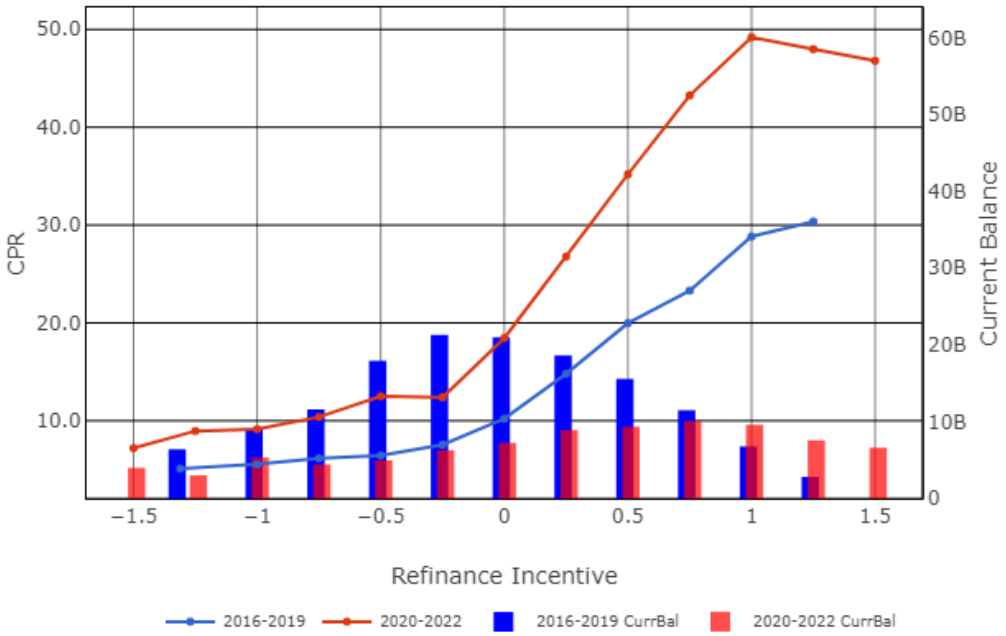

Figure 10.c. Loan Age 13 – 24 Months

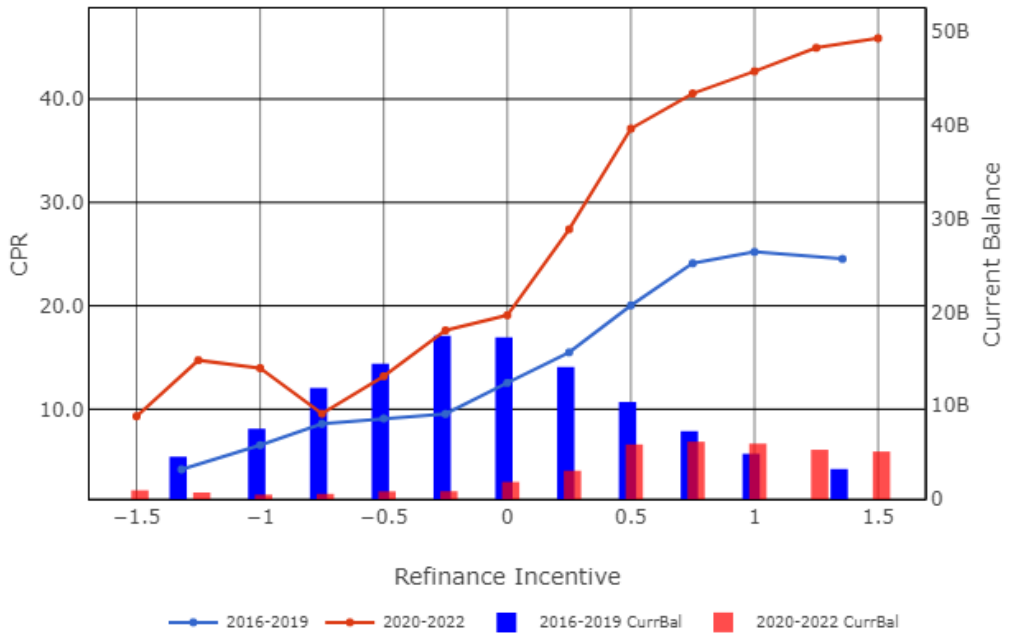

Other than some low volume artifacts in Figure 10C, the splits behave like all the other factors. What could be going on here? All other things equal, older loans should prepay faster than newer loans, given the same prepayment incentive.

All other things equal—but all other things were NOT equal! Let's think back to what happened during the early stage of the pandemic. Unemployment spiked, forbearance took hold, house prices jumped, and mortgage rates quickly fell. Mortgage rates didn't just fall, they quickly fell. The new and more seasoned loans were both exposed to falling rates, but the new loans were exposed very early in their life. There had been no time for baseline prepayments to chip away at loans. The loans closed and then, almost immediately, borrowers could refinance and save.

We pulled out the rate histories and took a close look. Figure 11 looks familiar to everyone in the mortgage industry. Rates have generally fallen since the credit crisis, with occasional blips up, until the Fed started tightening in 2022. The red portion of the line, January 2015–January 2022, highlights falling rates right up until the Fed started tightening.

Figure 11. 30 Year Fixed Primary Mortgage Rates

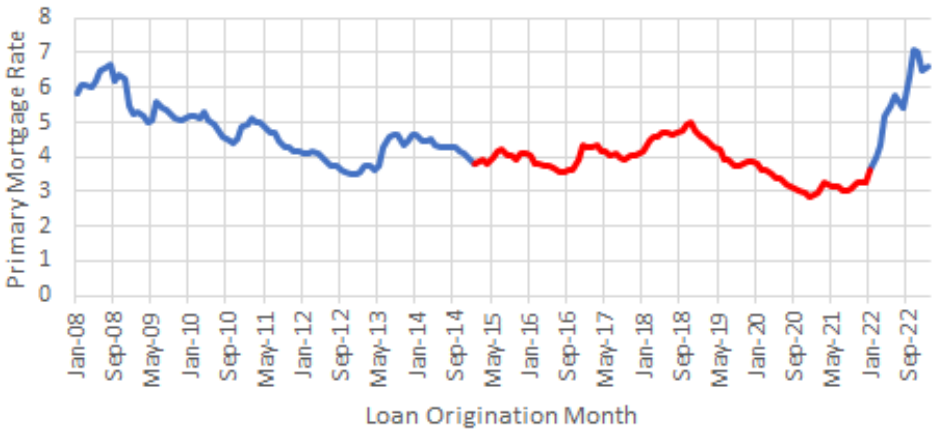

Not only did mortgage rates fall to historic lows, they also fell very quickly. Figure 12 shows us just how unusual the rapid decline was relative to prior periods. We calculated the number of months that passed before a borrower had the opportunity to refinance into a rate at least 100 basis points lower than their existing loan. The blank – spaces in Figure 12 represent loan origination months that have yet to see a full 1% decline in rates.

Figure 12. Time to First Refinance Incentive of At Least 100 Basis Points

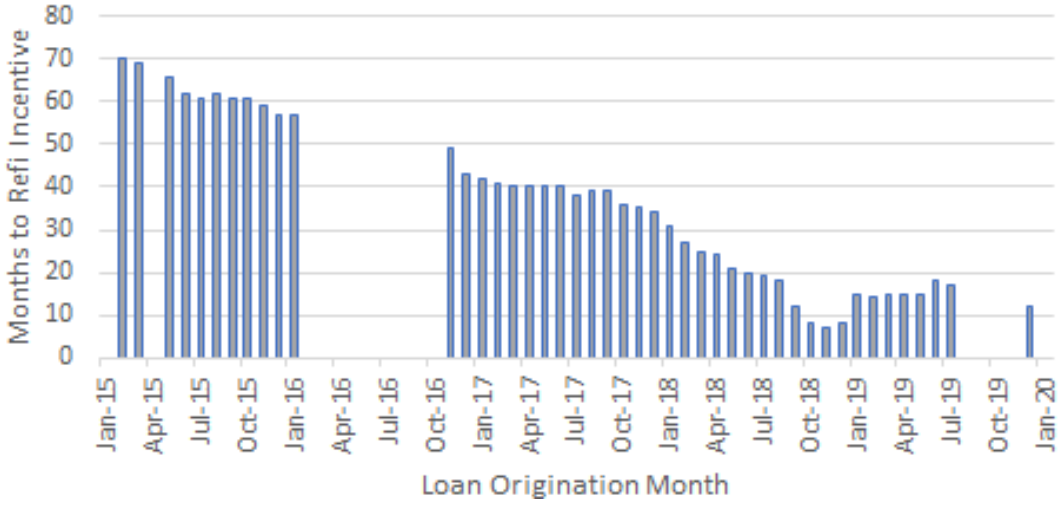

The speed with which mortgage rates fell for loans originated in late 2018 was unprecedented in the post–credit crisis period. We looked specifically at the 2017 and 2018 origination vintages to gauge the effect of the rapid decline. Loans from 2018, of course, were younger, yet, as Figure 13 shows, they paid faster than the 2017 originations when exposed to the same rate incentive.

Figure 13. Origination Year Cohorts

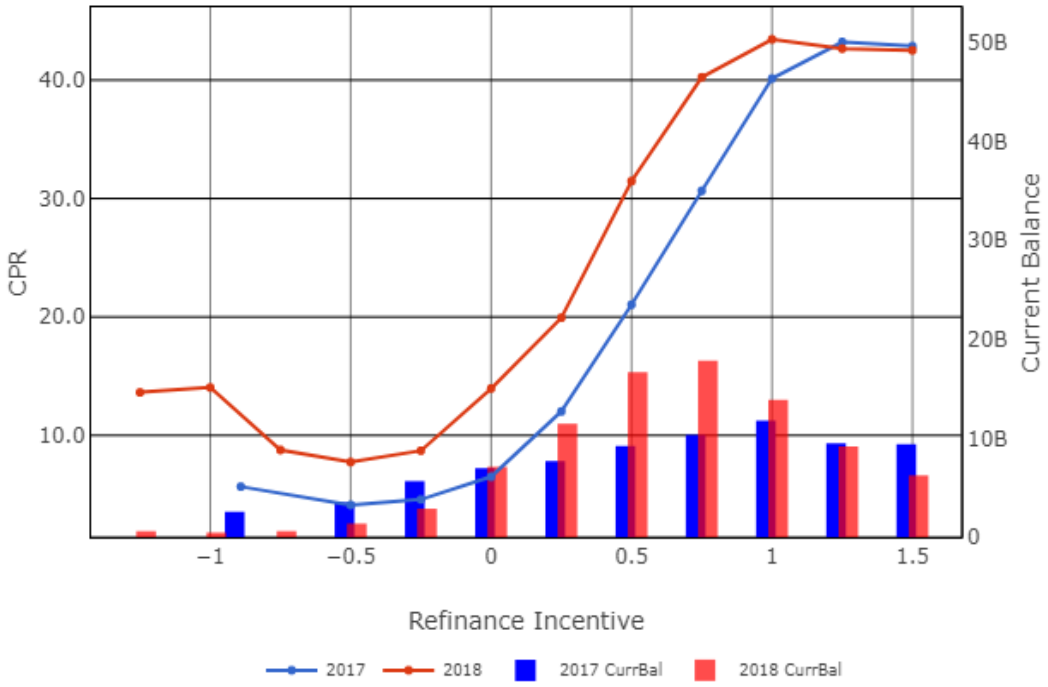

There's a lot going on in Figure 13. First, the speeds at each incentive level are averages from origination through May 2022 for both vintages. That means that total prepayments on the 2017 vintage could be higher than for 2018 because the averages occurred over an additional year on average. However, the response to rate incentive is clearly higher for the 2018 vintage, which was exposed to an immediate, rapid rate decline.

Second, the 2018 vintage speeds are faster across the spectrum, even where there is low incentive. The gap is much smaller for the out-of-the-money region and the 2018 volumes are very low (because rates fell so quickly), so the difference on the left side may be more of an artifact than a real behavioral shift. We are still pulling apart the details.

What's clear is that the speed of rate declines is least correlated with higher prepayment speeds at given incentive levels, and it may be causal as well. Additional research—including performance over other rapid rate decline periods 20 years ago—will help determine how important speed of rate changes is to prepayment sensitivity.

Time To Change the Model? Not So Quick!

Given what we have observed about rate movements and prepayment behavior, do we add this factor to the model? Not at this point. We saw one observation of a macro variable that correlated with one instance of prepayment behavior. That is never enough to justify changing a model. Adding new variables to an already complex model risks overfitting the model to a specific data set or a specific time period. Overfit models may produce nonsensical results outside of the specific circumstances that uncovered the correlation.

Once we identify a candidate driver variable or new interaction between existing variables, we then perform a series of additional tests. First, we back test across a much longer time period to make sure that the effect is persistent and consistent. Second, we project results from the adjusted model through our forward-looking simulations to see if the overall projections still make sense. We do not want to create many new problems by solving one.

For the illustrative example presented above, we have some data from the prepayment surge of the early 2000s that should help test the speed of rate decline hypothesis. Regardless of whether the hypothesis is supported by the additional data, we will have learned something new and useful.

We must also not forget a few other things going on during the time of these observations. A global pandemic, historically low interest rates near zero, an unemployment spike followed by robust economic growth, forbearance, emerging inflation, aggressive Fed policy, housing supply shortage—it was a busy couple of years! The analysis presented here does not purport to identify the true cause of S-Curve variation. Rather, we illustrate some of the techniques we use when analyzing unusual patterns.

Like many forms of exploration, data analysis often leads to unexpected discoveries. We must keep our minds open to observing new patterns while rigorously testing our existing beliefs. That's what makes this work so interesting!

What About the AD&Co Models?

You will have noticed that we have, so far, only looked at actual data. That's always the first step in modeling. We must know what already happened before we can build models predicting what will happen in the future.

We constantly monitor model performance across many dimensions, including refinance incentive. We then evaluate whether any drift in performance justifies tuning the models to reflect changing behaviors. We must balance the desire to be accurate in the short term against the goal of a stable long-term model. In our next Pipeline article on the S-Curve, we will show how well our LoanDynamics Model performs in a dynamic environment.

Finding new factors, such as speed of rate movement, helps us build models that better explain mortgage borrower behavior, while requiring fewer adjustments. These findings also support our core belief that, while individual circumstances and external incentives change, borrowers do, in fact, behave consistently across time when faced with similar situations.

In this article, we have given you a sense of just a small part of the work we do every day in searching for insight. In practice, we use far more sophisticated methods to estimate and test our models but, on occasion, we find it useful to strip away the complicated statistics and fancy optimization techniques and simply look at the raw patterns, and then go on an ad hoc exploration of behavior to challenge our fundamental beliefs.

We hope you enjoyed drifting along with us as we chased the meandering S-Curve!