Kamakura Corporation
2222 Kalakaua Avenue
Suite 1400
Honolulu, Hawaii 96815, USA
phone 1 808 791 9888
facsimile 1 808 791 9898
www.kamakuraco.com

Advanced Financial Risk Management:

Tools and Techniques for Integrated Credit Risk, Market Risk, Liquidity Risk and Interest Rate Risk Management

April 2012

Donald R. van Deventer
Kenji Imai
Mark Mesler

Part 4: Risk Management Applications, Instrument by Instrument

Chapter 29:
Mortgage-Backed Securities and Asset-Backed Securities

We are now in the 'home stretch' as we prepare for the comprehensive Jarrow-Merton put option as a measure of integrated credit risk, market risk, liquidity risk and interest rate risk. Our task in this chapter is to apply the analysis of option exercise subject to transaction costs from the previous chapter to portfolios or pools of loans. The authors feel strongly that this should be done on a transaction level basis, loan by loan, rather than analyzing the 'pool' or portfolio as if that pool was a single transaction. If there is one lesson from the 2007 to 2010 credit crisis, it is that failure to do the analysis of securitized loans at the transaction level is a recipe for disaster. The complete disappearance of the collateralized debt obligation market, outlined in Chapter 20, came about as more and more investors realized that the complexity of that structure was intentional; it was a convenient way to discourage investors from looking closely at the underlying collateral. It facilitated the un-ending attempts by Wall Street to sell investors lumps of coal packaged as "early stage diamonds." Two key publications are essential reading for investors who own or are considering an investment in securitized transactions:

The Levin Report, featured in Chapter 18 and pictured in Figure 3, describing "Wall Street and the Financial Crisis: Anatomy of a Financial Collapse (United States Senate, Special Committee on Investigations, Majority and Minority Staff Report, April 13, 2011).

Lewis, Michael. "The Big Short: Inside the Doomsday Machine," Thorndike Press, 2010.

Both of these exceptionally detailed books describe exactly how the process of securitization is designed to hide the flaws of the ultimate collateral in a structured product.

The credit crisis came about because the process of securitization helped to obscure the true nature of the underlying mortgages that were the ultimate collateral in mortgage-backed securities that were in turn, after being split into tranches, used as collateral in CDOs. These mortgages presented loans whose principal was greater than the value of the underlying house and whose borrowers where the prototypical NINJAs ("no income, no job, and no assets"). Investors ignored these huge red flags because they didn't want to spend the time or money to investigate what they were buying at the transaction level. If ever the reader feels the same way, go home immediately or call in sick to work. No structured product should ever be purchased without transaction level due diligence. If the reader's institution cannot afford the time or money to do this, it should not be buying structured products.

From a computer science and financial theory perspective, there is no reason for allowing unnecessary data aggregation from the structuring process to obscure our analysis of the risk-adjusted return on the mortgages and consumer loans that make up the ultimate collateral in a structured product. We illustrate the reasons for this view during the course of the chapter and illustrate how the traditional approaches to prepayment is in fact 'off track' because of attempts to value a pool of loans as if they were a single very large loan. Even if one ignores the impact of home prices on defaults and prepayments, the traditional methods for measuring the interest rate risk of pools of loans were misguided attempts to short cut a loan by loan analysis.

We use the concept of "transactions costs" in Chapter 28 to illustrate the pitfalls of common Wall Street short-cuts like prepayment rates, prepayment tables, prepayment functions, and "involuntary prepayment," known to normal human beings as "default." We also introduce a more model approach to joint modeling of default and prepayment on pools of loans, multinomial logit. Finally, we close with some comments on the ultimate securitized product, the valuation of mortgage servicing rights.

Transactions Costs, Prepayments, Default and Multinomial Logit

In Chapter 28, we analyzed "rational prepayment," in which the decision rule was simple: which strategy minimizes the net present value of my borrowing costs, prepaying now or deciding one period from now? In Chapter 28, the probability of prepaying now was assumed to be 100% if that option lowered the present value of costs. In this section, we add another tool to the transactions cost technique for analyze prepayments and defaults.

Traditional asset and liability management (ALM) has ignored mortgage defaults and focused on interest rate-driven and mortgage age-driven prepayment. The events of the 2007 to 2010 credit crisis have made it more obvious that prepayment and default are intimately linked and that home prices are a critical driver of both probabilities. This section explains how mortgage prepayment and default are modeled on an integrated basis using multinomial logit. This technique can be applied both to a single mortgage loan and (more likely) to a pool of mortgage loans on a loan by loan basis.

The events of the 2007 to 2010 period have made it clear that the sole focus on a single macro-economic factor, interest rates, has been insufficiently robust to prevent the effective failures of institutions like Lehman Brothers, Washington Mutual, New Century, IndyMac, Countrywide, FNMA, and FHLMC. The more modern approach to risk management is a fully integrated approach to credit risk, market risk, liquidity risk and interest rate risk, the central theme of this book. In the context of traditional interest rate risk management, default was almost always ignored by most practitioners for two reasons:

- Credit risk was in a "separate risk silo" managed by others
- Most silo-focused interest rate risk assessment software was incapable of modeling default even if the analyst had wanted to do so

Both of these constraints, one political and one technical, have now been removed at institutions at the forefront of best practice risk management. One of the tasks triggered by this new approach to integrated risk is that the traditional assumptions of ALM have to be modernized, one by one. In no area is this truer than in mortgage prepayment and default analysis. The techniques traditionally used to mortgage prepayment analysis can be listed briefly as follows:

- Constant prepayment rate, either expressed as a conditional prepayment rate (annualized basis, "CPR") or as a single monthly "mortality" rate. This assumption ignores the impact of interest rates and home prices on prepayment. It is also derived from behavior of mortgage pools rather than individual loan data.
- Prepayment rate as a linear function of mortgage age, interest rates and mortgage attributes. This approach also ignores home prices as a driving factor and is derived from pooled behavior rather than individual loan behavior
- Prepayment tables where the prepayment rate is driven by a table of N explanatory factors, typically the same factors as the linear prepayment function. Again, these tables are typically derived from pools of mortgages and the related prepayment rate on the pool.
- Rational prepayment models, based on the "all or nothing" valuation of the mortgage on an American option basis. Again, default and home prices are often ignored. Chapter 28 explained this approach in detail, with default risk incorporated if the rational prepayment analysis uses a valuation yield curve that reflects the probability of default.
- Rational prepayment subject to transaction costs, which recognizes the true transaction costs of refinancing and other implicit costs that cause the prepayment on mortgages to be less frequent than the fully rational approach would predict. We introduced this approach in the latter half of Chapter 28 as well.

In all of these approaches, the amount of prepayment that occurs, given market conditions, is never random. Given the inputs to each of these techniques, the prepayment amount is known with certainty.

Best practice approaches differ in three key dimensions:

- Prepayment is recognized as being random, just like default, although the probability of prepayment can be estimated just like a default probability
- Prepayment and default are being analyzed at the individual loan level. The fact that those who analyzed mortgage collateral at the loan level fared much better than investors who did not (i.e. asset backed CDO investors and the rating agencies) during the credit crisis has not been lost on senior management.
- Default and prepayment are increasingly being modeled as mutually exclusive events that are modeled together, not separately.

The first step forward in recognizing home prices and other loan-specific inputs in the prepay/don't prepay decision and the default/no default outcome was the use of logistic regression. Logistic regression is commonly used to model events that have only binary 0/1 outcomes, like we saw in Chapter 17 on reduced form credit models. Many analysts have applied this technique to mortgage prepayment and default, estimating separate logistic regressions for these two events. The impact of common risk factors like home prices, interest rates, and mortgage age can be explicitly taken into account on a loan by loan basis.

This logistic regression based approach has many extremely attractive attributes:

- Default and prepayment are recognized as random events, not events that are predicted with certainty
- Previously ignored factors like home prices can be incorporated into both prepayment and default modeling on a consistent basis
- Loss given default can be recognized as related to default probability by modeling home prices forward and using those home prices as inputs to the default and prepayment probability functions. In a default scenario, the simulated value of the home allows one to derive loss given default after applying appropriate liquidation costs and time lags.

Emerging best practice is to use multinomial logit to deal with two practical problems that arise when independent logistic regressions are used:

- With independent logistic regressions, the modeler is uncertain about which probability to apply first in a given period and a given scenario: the simulation of prepay/don't prepay or the simulation of default/don't default
- Similarly, it's possible that the probabilities of prepayment and default could add up to more than 100%

In order to deal with these issues, a number of advanced techniques can be employed. One of the most popular techniques is multinomial logit, which allows N potential outcomes instead of the binary 0/1 outcomes assumed in normal logistic regression. Here are a number of potential specifications of the outcomes from a multinomial logistic regression:

Model A outcomes:

- Default
- Prepayment
- Neither of the above

Model B outcomes:

- Default
- Prepayment for interest rate reasons, remaining in the same house
- Prepayment because of relocation
- None of the above

Model C outcomes

- Default for purely non-medical financial reasons
- Default triggered by medical expenses
- Prepayment because of divorce
- Prepayment because of interest rate reasons
- Prepayment because of job relocation
- Prepayment because of voluntary relocation
- None of the above

These multiple outcomes offer the analyst and risk manager a much more precise assessment of the impact of interest rates and home prices on the value of a pool of mortgage loans. The multinomial logit coefficients, given the explanatory variables chosen by the analyst, are derived either using joint maximum likelihood estimation or by estimating a logistic regression for each outcome and making appropriate adjustments.

For those institutions that place a priority on correctly understanding risk adjusted shareholder value creation and business strategy, the logistic and multinomial logistic approaches are an essential step forward. In the next few sections, we compare this best practice approach with legacy approaches to modeling prepayment. With few exceptions, these legacy approaches ignore default risk.

Legacy Prepayment Analysis of Mortgage-Backed Securities

The valuation and analysis of mortgages, both fixed and floating rate, is one of the most complex analytical problems in U.S. financial markets. As mentioned above, much of this difficulty stems from the confusion between analysis of a single mortgage and analysis of a portfolio or pool of mortgages. Douglas Breeden [1991], one of the leading U.S. researchers in finance, former chief executive officer of mortgage fund manager Smith Breeden, and now Dean of the Duke University Business School, commented about the problem as follows:

> '...mortgages are viewed as far too complicated to value precisely and rigorously, even with the Black-Scholes model and the many improvements developed in the eighteen subsequent years.'

As we saw in Chapter 28 and can confirm using multinomial logit, accurate valuation of mortgage-backed securities and other pools of consumer loans is a complex interaction of credit risk, interest rate levels and prepayment behavior. The purpose of this section is to take a second look at mortgage-backed securities in the light of the implications of that chapter: that the transactions cost approach promises dramatic improvements in the speed and accuracy of credit-adjusted valuation and risk analysis for consumer loans on both an individual loan level and a portfolio level. Before looking at the implications of that approach, we look first at traditional industry practice.

Legacy Approaches: Prepayment Speeds and the Valuation of Mortgages

Prior to the 2007 to 2010 credit crisis, Wall Street analysis of mortgage-backed securities relied almost exclusively on 'prepayment models' as a means of analyzing the embedded options in mortgage-backed securities. The prepayment models, in combination with Monte Carlo simulation, produce another number as output—the option-adjusted spread (OAS) on a mortgage-backed security. In the next few sections, we pose a number of questions and attempt to answer them:

- Do prepayment models have an options component?
- Are prepayment speeds predictable enough to use as an input for the option-adjusted valuation of mortgages?
- Is the truth about prepayment speeds obvious enough that sophisticated market participants can reach a consensus about their levels?

Most prepayment models have an objective similar to our objectives throughout this book:

1. Security valuation on a credit-adjusted and option-adjusted basis;
2. Measurement of interest rate sensitivity;
3. Accurate hedges;
4. Guidance on 'rich/cheap' analysis for security selection.

Prepayment speed models come in four basic varieties:

a. Models in which the prepayment speed is a function of time, but not rate levels. Examples are the lifetime prepayment speed models of the Public Securities Association (PSA), Conditional (or 'Constant') Prepayment Rate (CPR), and Single Monthly Mortality (SMM).

b. Tables in which different prepayment speeds are assigned depending on the remaining maturity of the mortgage, the coupon level, and the current level of interest rates. We call this the prepayment table approach.

c. Prepayment functions, which are typically derived from historical data on prepayments and include more inputs than a prepayment table to provide the accurate prepayment speed. Over a long period of time, linear prepayment functions have been the most common functional form used.

d. Logistic regression models, like those we used in Chapter 16, to fit default probability formulas to historical default data-bases. These logistic regression models are relatively new in prepayment analysis, but offer great promise for understanding the probability of prepayment based on historical behavior of individual borrowers.

e. Multinomial logistic regression, which we described above as the best practice approach for joint modeling of default and prepayment.

Analytically, the constant prepayment speed models are the simplest. Prepayment tables, which are one step up in complexity, assume that we can accurately forecast prepayment speeds. Prepayment functions assume the future will be like the historical data set which produced the estimates used in the prepayment function. The same is true for logistic regression models to some extent. How well do these models work?

Constant Prepayment Speeds are a Simply a Principal Amortization Assumption

The most important point to make about the PSA, CPR, and SMM analysis is that they have no options component at all. They represent an assumption about principal amortization and make it quite clear that the rate of principal amortization does not depend on the level of interest rates. As a result, there is no 'rational' consumer response to lower rates at all reflected in these three models. A higher prepayment speed lowers the level of principal outstanding. A higher prepayment speed also has the impact of increasing cash flow on the mortgage in the early years and decreasing it in later years, effectively shortening its duration. These cash flows are not stochastic and don't depend on the level of interest rates.

The single speed approach to prepayment analysis doesn't work very well. One way many market participants have sought to improve the accuracy of prepayment analysis is by creating a prepayment table which specifies what prepayment rates will be at different times in different interest rate environments. Typically, a prepayment table will contain at least two dimensions (and usually more), one for each payment and another for various levels of 'refinancing advantage', the spread between the rate on the mortgage loan being analyzed and the current coupon rate on new mortgages. If there are five interest rate tiers and 360 payments, 1,800 input numbers will be required as inputs to the valuation and hedging analysis.

Douglas Breeden [1994] takes a very interesting look at Wall Street's ability to forecast prepayment speeds, and we use data from his study liberally in this section. First, Breeden shows that Wall Street typically lacks consensus on the best constant lifetime prepayment speed for a given security. Breeden showed that the Wall Street estimates of prepayment speeds on a 10% FNMA mortgage-backed security maturing in 2018 ranged from 20% to 40% as of December 1992.

Breeden also shows that the conditional prepayment rates as a function of the coupon on the mortgage-backed security and the refinancing incentive have been very unstable and have been inconsistent with rationality in many cases. Prepayment speeds for high-coupon mortgages have often been lower than prepayment speeds for low-coupon mortgages (this phenomenon is often called 'burnout', short-hand for the fact that the most interest rate-sensitive borrowers in a pool of loans leave the pool first via refinancing, reducing the interest rate sensitivity of the pool over time). In fact, it is easier to guess the value of an embedded prepayment option directly than it is to guess the prepayment rates needed as inputs to a formula to calculate option values.

Legacy Approaches: Option-Adjusted Spread

Given the difficulties with prepayment models in general, what are the implications for the option-adjusted spread quotations common on Wall Street? First of all, we should summarize the option-adjusted spread (OAS) procedures, for the term OAS refers more to the procedures used than it does to the true concept of the risky spread on mortgages. Typically, analysts do the following:

- They assume a prepayment function or prepayment table is true.
- They assume that Monte Carlo simulation is a good method for valuing an American option to prepay, ignoring the concerns voiced in Chapter 27.
- They solve for the spread over a risk-free interest rate which provides a calculated value for the mortgage-backed security that matches the observable value.

In addition to the comments in Chapter 27 about the fundamental problems with the use of Monte Carlo simulation in the valuation of American option-related securities, OAS analysis typically incorporates these errors:

a. The refinancing incentive is not measured on a matched maturity basis.

Within the prepayment table, the refinancing incentive in the current market is measured in order to choose the correct prepayment rate from the prepayment table or prepayment function. This is measured as the difference between the rate on the mortgage being analyzed and the rate on a new 30-year fixed-rate mortgage. This sounds good at first glance, but a closer look gives cause for serious concern. If there are five years remaining on the mortgage, we are comparing its five-year rate with the 30-year loan and assuming that refinancing is done on the basis of this ‘apples to oranges’ comparison. In a swap trading unit or an investment management firm, anyone who said ‘30-year bonds are better than five-year bonds because their yield is higher’ would lose their job. This sort of comparison, however, is embedded in almost all OAS analysis. Why? Historical data on mortgage rates includes only rates at new issue maturities (typically 15 and 30 years in the U.S.) and the analysts had no other data to measure refinancing spread. The reason for the choice of measure is clear, but it's an important source of inaccuracy.

b. The refinancing incentive is measured on the basis of a 30-year mortgage rate which does not include a premium for the call provisions of the loan

Within the Monte Carlo simulation routine, there is no capability to calculate how much the premium should be on a new mortgage that's callable compared to a hypothetical ‘non-call’ mortgage. Why? You have a ‘chicken and egg problem’—using Monte Carlo and a prepayment table, we need to know what the option is worth in order to calculate what the option is worth! All of the refinancing incentives measured within the Monte Carlo routine are normally done assuming the mortgage is non-callable, introducing another source of error.

c. The prepayment table should be path-dependent, but it's not

If rates start at 6%, jump to 12% for 10 years, and then fall to 4%, the prepayment function and prepayment tables will produce the same prepayment speed prediction as if rates started at 6%, fell to 3% for 10 years, and then jumped to 4%. Clearly, that should not be the case. Monte Carlo is often relied upon on the rationale that prepayments are ‘path-dependent’. We discuss that comment in more detail below, but it's clear Monte Carlo replaces a path-dependence problem in the prepayment rates themselves with a path-dependence problem in the prepayment table.

d. The OAS will match the market value of the observable security, but not its interest rate risk

Why? For the same reason that the single prepayment speed model didn't work. We have two equations (value and rate sensitivity) we want to solve simultaneously, but only one variable (OAS) to adjust in order to do this. We can match both only by accident.

OAS is the ‘plug’ or balancing number that offsets all of the errors inherent in monte carlo analysis and the four problems above. It's not surprising, given the lack of consensus on Wall Street about prepayment rates, that median broker estimates of OAS show very strange patterns when plotted by coupon level. Breeden [1994] cites data for the third quarter of 1989 and the fourth quarter of 1992 which show investors could double their spread theoretical OAS by simply buying an MBS with a different coupon level, even though this discrepancy is public information widely disseminated by Wall Street firms. In light of the 2007 to 2010 credit crisis and the drop in home prices, traditional prepayment models have gone seriously “off the tracks” and a repeat of the Breeden study would show that even more serious errors exist now.

The fact that this difference in OAS was not arbitraged away is due to the market’s view that the OAS numbers were not accurate and that the spread differential was not real. Wall Street didn't bid up the price of the MBS with the widest reported OAS relative to the MBS with the lowest MBS. Clearly, OAS numbers are taken with a grain of salt by market participants. Gregg Patruno [1994] of Goldman, Sachs & Co. summarizes a common view by saying ’The standard measure of mortgage relative value ... OAS does not account for the complete nature of prepayment risk. Rather, it adjusts for the optionality due to interest rate fluctuations under the presumption that prepayment rates are a known, permanent function of interest rates.’[1]

Patruno goes on to suggest both the reason for the problems with OAS and a potential solution to the problem: ‘Homeowners are willing to refinance if the financial incentive is high enough to meet their requirements. We use the measure of refinancing incentive actually considered by typical homeowners and their bankers—not an abstract interest differential or ad hoc statistical artifact, but the real dollar savings expected on an after tax basis, taking into account an appropriate mix of available mortgage rates and points.’[2] In other words, the only way to effectively model consumer prepayment behavior is to consider the rational value of the option to prepay. Patruno concludes that ‘...the levels of current mortgage rates and transactions costs determine what fraction of the homeowners in the distribution should be considered willing to refinance...’[3] The credit crisis also made it clear that the value of the home, relative to the principal on the loan, had a major impact on the ability of a borrower to prepay.

For these reasons, the authors believe that the transactions cost approach of Chapter 28, as analyzed by McConnell and Singh [1994] and Stanton [1995], supplemented by multinomial logistic regression for prepayment and default analysis, represents best practice.

Implications for OAV Spread, CMOs and ARMs

---

1 Patruno [1994], p. 53.

2 Patruno [1994], page 47.

3 Patruno [1994], page 47.

The primary reason for the lack of OAS consensus among Wall Street firms is that all of the assumptions being made are assumptions about the borrower. It is no wonder that consensus is impossible. Instead, we recommend the OAV (Option-adjusted value) spread. Let us consider a corporate treasurer who has to choose between issuing a seven-year non-call bond or a 10-year bond that is non-callable for seven years. We modeled this structure in Chapter 27. He makes the decision in a straightforward way. Firstly, he asks the rate on a comparable Treasury seven-year non-call bond and calculates his spread over Treasuries. Secondly, he calculates (since no observable Treasury is 10 years non-call seven-year bond) what the coupon would be on a 10-year non-call seven-year Treasury bond that trades at par value, using the techniques of Chapter 27. His 'spread' on the 10-year non-call seven-year bond issue is the difference between his coupon and the Treasury's coupon. He issues the bond with the lowest spread.

Collateralized mortgage obligations (CMOs) and adjustable rate mortgages (ARM) can be analyzed using the Heath Jarrow and Morton approach, like the example in Chapter 9, as a variation on our results in Chapter 28.

Logistic Regression, Credit Risk and Prepayment

Logistic regression is another powerful tool that we are familiar with from its use in modeling default probabilities for retail, small business, and large corporations using a historical database of defaults. In a similar way, we can use logistic regression to model the probability of prepayment. With this technique we can combine variables like the following, as explanatory variables for prepayment:

- The net present value benefit of prepaying as of the current observation (which captures the interest rate benefits of prepayment);
- The current value of the house relative to the principal balance of the outstanding loan
- The age of the loan;
- The months remaining before maturity;
- The month (to capture seasonality);
- Direct transactions costs of refinancing (points, fees, etc.);
- The default probability of the borrower, as captured directly by internal models or by some kind of internal credit rating;
- Any other relevant consumer credit bureau approach.

The coefficients of this logistic regression essentially let us derive the transactions costs that we implied previously. They also automate the adjustments for seasonality that we would also have to make in the approach of the previous section. The result is an extremely powerful integration of credit risk, interest rate risk and prepayment risk that significantly advances the accuracy of our ability to model the Jarrow-Merton put option, our integrated measure of risk for the full balance sheet. Multinomial logistic regression is best practice, because it gives integrated probabilities for default and prepayment. For those making the transition from legacy techniques, however, the logistic regression approach is a nice intermediate step on the way to a full multinomial logistic implementation.

Mortgage Servicing Rights: The Ultimate Structured Product

One of the lessons of the credit crisis of 2007 to 2010 was that the conventional wisdom can be a dangerous thing. The conventional wisdom in early 2006 said that "home prices in the United States do not go down" and that the copula method was an accurate method for valuing tranches of collateralized debt obligations. Another complex security valuation issue is becoming increasingly important: the valuation of mortgage servicing rights. This section discusses potential modeling risks one of the most complex "mortgage-backed securities," mortgage servicing rights. The authors would like to thank our colleague Mark Slattery for his helpful comments on this section.

The table below shows the reported valuations for mortgage servicing rights for 8 large U.S. financial services firms between December 31, 2009 and September 30, 2011. The aggregate reported value of mortgage servicing rights has declined from $62.5 billion to $35.6 billion in 21 months. Even at the lower valuation, the very large absolute value attributed to mortgage servicing rights and the tremendous volatility of the valuation in the current market makes an examination of the conventional wisdom essential from both a risk management and corporate governance point of view.

Figure 1 [Chapter29-Figure01.tif]

Kamakura Corporation
Reported Valuations of Mortgage Servicing Rights (USD Millions)
Source: Kamakura Corporation, Company SEC Filings

Mortgage Servicer	December 31, 2009	December 31, 2010	September 30, 2011
Bank of America	19,465	14,900	7,880
Wells Fargo	17,123	15,886	13,769
JP Morgan Chase	15,531	13,649	7,833
Citigroup	6,530	4,554	2,852
US Bancorp	1,749	1,837	1,466
SunTrust	936	1,439	1,033
BB&T	832	830	573
First Horizon	303	207	151
Total	62,469	53,302	35,557

One of the things that these figures make obvious is that the mortgage servicing business is an enormously concentrated business with highly skewed market shares. Another obvious point is that the percentage of total market capitalization for these firms that is attributable to mortgage servicing rights is very large. Mortgage servicing rights are classified as a "Level 3" asset for accounting purposes because there are very few observable traded prices and because many of the inputs to a reasonable valuation are also unobservable and the result, in the end, of educated guesses. If there is model risk in reported mortgage servicing rights valuation, equity market participants need to be informed of these model risks, and those hedging mortgage servicing rights need to take great care because of potential errors in both the amount and *direction* of the hedging needed.

An Introduction to the Valuation of Mortgage Servicing Rights

As is common in the risk management business, the head of a mortgage servicing rights hedging program once said "I don't care about accuracy, I just want to do the same MSR calculation as everybody else." Therein lies the problem. What if "everybody else" in total is only 7 firms? How can they be sure that there are no errors in this complex valuation calculation? One approach is to survey all of the other participants in the market, which is commonly done, to insure consistency. Unfortunately, the assumptions that home prices don't go down and that copulas are appropriate for

CDO valuation were consistent modeling assumptions in US financial markets in 2006. The consistency alone is not sufficient to establish accuracy. The recent controversy about the manipulation of the U.S. dollar Libor market during the credit crisis, when the U.S. dollar Libor panel consisted of 16 firms, illustrates another risk: the risk that the 8 firms who dominate the mortgage servicing business have an economic incentive to agree on a valuation methodology that maximizes the value of this large asset. For all of these reasons, an independent approach of maximum accuracy in valuing mortgage servicing rights is essential.

Starting from the ground up, one can approach the valuation of mortgage servicing rights as the valuation of a fixed income (broadly defined) security subject to default risk and prepayment risk. Mortgage servicing rights, like all of the other securities valued in prior chapters, can be valued in the Heath Jarrow and Morton framework. A recent paper that provides an introduction to this approach is

> Che-chun Lin, Ting-heng Chu and Larry J. Prather, "Valuation of Mortgage Servicing Rights with Foreclosure Delay and Forbearance Allowed," *Review of Quantitative Finance and Accounting*, 26: 41–54, 2006

In the next section, we follow Lin, Chu and Prather and decompose mortgage servicing rights into their component cash flows and compare the "conventional wisdom" and best practice.

Comparing Best Practice and Common Practice in

Valuing and Hedging Mortgage Servicing Rights

There are so many differences between best practice and common practice in valuing and hedging mortgage servicing rights that we will have to be brief to keep this section of reasonable length. We address these differences point by point.

Valuation Yield Curve for Cash Flows

Best practice discounts cash flows using a series of yield curves, depending on the nature of the cash flows, particularly when some cash flows are best valued in nominal dollars and some are best valued in real dollars. Common practice discounts all cash flows at the same yield curve, typically the U.S. dollar Libor-swap curve. The use of this curve is usually justified with the phrase "all mortgage market participants use the swap curve for hedging." Unfortunately, that has nothing to do with whether the curve is appropriate for valuation. As we shall see, the right valuation curve depends on the characteristics of the mortgage servicing right ("MSR") cash flow component's credit risk, prepayment risk, and inflation-sensitivity (real or nominal), and many of the relevant curves are closely tied to home price movements. The Libor-swap curve has a number of serious problems when employed for valuation. We list many of those problems here:

1. As the fines of $450 million levied on Barclays in 2012 for Libor manipulation confirm, the Libor curve has been manipulated and is no longer a legitimate tool for hedging.

2. Market participants commonly assume that fixed rate mortgage yields are a simple spread over an interpolated swap yield at two key maturities, such as the 2 year and 10 year interpolated swap rate at an effective maturity of 6 years. This assumption is false. Over the period from January 8, 2004 to November 17, 2011 using weekly data, the appeal of using swap rates for valuation is deceptively attractive, as shown by the high correlations of swap yields (all over 86%) with the 30 year conventional fixed rate mortgage rate. The true relationship between swap and Treasury yields and fixed rate mortgage yields is more complex, as we discuss next.

3. Using conventional 30 year fixed rate conventional mortgage reported on the Federal Reserve H15 statistical release, the interpolated 2 and 10 year swap rates explain only 82.8% of the weekly variation in fixed rate mortgage yields with a standard error of 30.47 basis points.

By contrast, a simple linear regression using multiple points (13) on both the swap curve and the U.S. Treasury curve explains 98.17% of the weekly variation in fixed rate mortgage yields with a standard error only one-third as big as the conventional wisdom, 9.9 basis points. The conventional wisdom, to benchmark valuation off of two points of the Libor-swap curve or off of the entire swap curve alone, has no basis in the facts. The U.S. Treasury curve plays at least as great a role in the movements of the 30 year fixed rate mortgage yield as does the Libor-swap curve and there is much less risk of manipulation by market participants.

4. The spread between the 30 year swap yield and the matched maturity Treasury yield has been consistently negative since October 24, 2008. The reason for this negative spread is the difference in credit exposures in the two yield curves. A default by the U.S. government, which is no longer an unthinkable possibility, puts potentially the entire principal amount at risk. If there is a default by a swap counterparty, the amount at risk is, at a maximum, the change in the present value of the swap position since its value at origination (which would be zero for a swap done "at market" at the time). With the increase in the perceived risk of the U.S. government, the spread between the swap yield at 30 years and the 30 year U.S. Treasury yield has been negative on 749 business days from February 9, 2006 to November 17, 2011. Over the same period, the 10 year spread between the swap curve and the U.S. Treasury curve has been negative for 67 days, with the first negative spread reported on January 21, 2009.

It is unthinkable for market participants to value mortgage servicing rights by discounting at yields below U.S. Treasury yields given the current relative default risk of mortgages and the U.S. Treasury. There is no legitimate argument that justifies such a valuation choice. This graph from December 2, 2011 plots the zero coupon yields, using advanced yield curve smoothing techniques, for the swap curve versus the U.S. Treasury curve. The negative spread at 30 years “contaminates” the spread at much shorter maturities (on this date, maturities 13 years and over), so the choice of the swap curve as a valuation yield curve is a dramatic error. The short end of the curve shows a hump in discount rates for the Libor-swap curve because the credit risk of 1, 3, and 6 month Libor is incompatible with the credit risk of the swap yields used for 1 year and longer maturities, as discussed above:

Figure 2 [Chapter29-Figure02.tif]

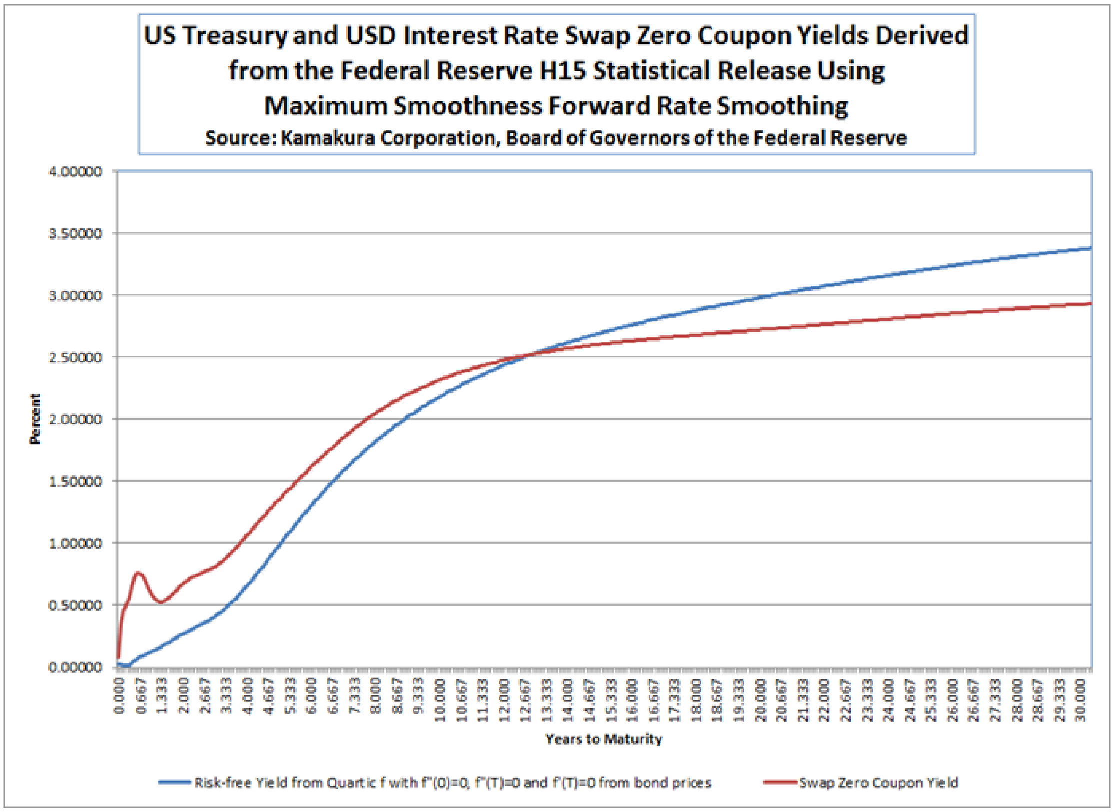

Simulation of Random Movements in Yields

The conventional wisdom for simulating the random interest rate environment when valuing mortgage servicing rights is embedded in legacy asset and liability management systems. For the most part, these systems rely only on one factor models of the term structure. As commonly implemented, these one factor models imply that changes in yields can only (a) all rise together, (b) all fall together, or (c) all remain unchanged. Unfortunately, this implication of one factor term structure models is false as we discussed at length in Chapter 3. With respect to the simulation of random interest rates, the conventional wisdom is dramatically wrong and is in violation of the Basel Committee on Banking Supervision’s December 31, 2010 Revised Basel II Framework for Market Risk, which requires on page 12 that

> "For material exposures to interest rate movements in the major currencies and markets, banks must model the yield curve using a minimum of six risk factors."

Because the conventional wisdom assumes away the yield curve twists that are very heavy influences on mortgage prepayments and defaults, the resulting valuations could potentially be grossly incorrect.

The Role of Home Prices in Defaults and Prepayments

The conventional wisdom in the servicing of a whole loan portfolio is to assume a constant default rate that is not influenced by home prices or interest rates. This assumption is false, and that is well recognized by many. They are constrained by legacy asset and liability systems to use a valuation assumption that they know to be false. Default rates are very sensitive to home prices. A recent study by Kamakura Corporation of the mortgage default rate reported by the U.S. Mortgage Bankers Association found that the change in home prices is statistically significant in predicting defaults, with a t-score equivalent well above the t-score level of 2 that is normally the benchmark for judging statistical significance. Modern enterprise risk systems allow mortgage servicing rights valuation with default rates that are determined by the simulated random movements in home prices, interest rates and other key factors. The erroneous assumption of constant default rates is avoided.

Similarly, it is very common in the valuation of mortgage servicing rights to assume a constant, linear, or tabular prepayment function in which home prices play no role. This assumption is also false. Recent publications by Calhoun and Deng (Journal of Real Estate Finance and Economics, 2002) and Campbell and Cocco ("A Model of Mortgage Default," Harvard University research paper, October 2011) are among many research papers which document the very significant role of home prices in prepayment as well.

Other Sources of Cash Flow Related to Mortgage Servicing Rights

The total cash flow from mortgage servicing rights can be broken into components as follows:

- Net servicing fee, equal to the gross servicing fee less guarantee fee, if any
- Net cost to service
- Net cash flow triggered by delinquency and default
- Float on taxes and insurance
- Float on principal and interest
- Float and costs from prepayments

The magnitude and very existence of these cash flows depends on the probabilities of default, partial prepayment and prepayment in full, which have been mis-specified in the conventional approach to valuing mortgage servicing rights as explain above.

There are many other significant errors that can be made in valuing individual components of cash flow. The cost of service, for example, might be assumed to constant, but that assumption is false. The cost of service is affected both by inflation and by the economies of scale in servicing relative to other competitors. The float on taxes and insurance could be calculated on the assumption that taxes and insurance are constant, but that assumption is also false. Both taxes and the cost of assumption are strongly linked to the value of the house. One might assume, as a better approximation, that home prices are correlated with inflation. Unfortunately, looking at correlation between the Case-Shiller home price indices and the consumer price index shows a wild variation in the linkages as documented in the June 1, 2009 Kamakura blog “Risk Management Strategies for Individual Investors, Part 4: Home Prices and Inflation” on www.kamakuraco.com.

The float on principal and interest and the float and costs from prepayments are complex derivatives that depend not just on interest rates but also on home prices. The conventional wisdom ignores the home price linkage. Stated more precisely, analysts recognize the linkages but they are constrained by legacy systems that are unable to take advantage of this economic reality. A modern state of the art enterprise risk management system avoids this valuation error.

Incorrect Hedging of Mortgage Servicing Rights

The conventional wisdom notes that mortgage servicing rights are to be reported at market value under generally accepted accounting principles. The corollary to that fact, however, is wrong—the conventional wisdom concludes that, since these market values are random, mortgage servicing rights need to be hedged. That conclusion is potentially wildly incorrect, for a simple reason. Market participants have simply ignored the “factory” which services mortgages (other than assuming the constant dollar cost above). This factory is not marked to market on the balance sheet under GAAP, but in fact the value of the operations center (buildings, software, hardware, staff, and so on) rises and falls with the value of mortgage servicing rights in a highly correlated way. By ignoring this link and partial or full offset, those who hedge may in fact be INCREASING RISK, not reducing it. Take the example of a scenario like that of 2007 to 2012, where U.S. home prices have fallen sharply. This both reduces the value of mortgage servicing rights (because default rates have increased) but it also has reduced the cost of being a participant in the operational processing of mortgages. If all of the mortgages being serviced, for example, defaulted, what would happen? The present value of labor costs of mortgage servicing would essentially go to zero, because all staff would be laid off. The present value of computer hardware needs would go to zero, because the hardware used would be sold or moved to other businesses. The present value cost of software used would go to zero (in the case of annual rental pricing, since the contract would be terminated) or be sharply reduced (in the case of perpetual license pricing, since the maintenance fee would be terminated). If the mortgage servicing business were terminated, the value of the related buildings and land would go to their value of the next best use.

In short, there is a natural offset of some or all of the variation in the value of mortgage servicing rights. The conventional wisdom, however, has completely forgotten this simply because the accountants treat mortgage servicing rights and the physical and intellectual property of the operations center differently. This is a very common source of risk management disasters, in which risk managers pay attention to the accountants and ignore the true economics of what they are trying to hedge.

Conclusions

Among “mortgage backed securities” broadly defined, the valuation of mortgage servicing rights is both the most complex and the best illustration of why best practices need to be employed in valuing mortgage-backed securities in general. The conventional wisdom in the valuation of mortgage servicing is notoriously “sticky” because there are so few market participants in the business and because reported values are so large. There is no one inside the mortgage servicing industry with a vested interested in pointing out that the Emperor’s wardrobe may be a bit threadbare. Only the accountants and regulators who would have to deal with any mark downs in value have downside risk. The shareholders and taxpayers who would ultimately have to pay were one of the large servicing banks to fail are not given sufficient information to make their own judgments about the value of mortgage servicing rights. As pointed out in the previous paragraphs, tens of billions of dollars in “value” relies on a series of assumptions that are known to be false.

We now turn to another major modeling issue in enterprise risk management and the correct valuation of the Jarrow-Merton put option on the firm: the valuation of “non-maturity” deposits.