Proceedings of the 1988 Winter Simulation Conference
M. Abrams, P. Haigh, and J. Comfort (eds.)

Using simulation to determine the costs of offering caps on adjustable rate mortages

Thomas W. McDowell
California Banking Group
Financial Analysis
Bank of America

ABSTRACT

This paper presents a methodology for using simulation to determine the cost of offering caps on adjustable rate mortgages. Caps represent a series of options, too complex to be evaluated with the options pricing model. A simulation model offers the ability to cost a variety of combinations of caps, thus providing flexibility in product offerings. Our results indicated that adjustable rate mortgages tend to be priced quite cheaply in the market. Cap costs appear to be higher than lenders may believe.

INTRODUCTION

Adjustable rate mortgages (ARMs) have become a primary lending instrument for financing single-family residential real estate purchases. Prior to 1981, California national banks could offer only a single adjustable rate product, indexed to the cost of funds for savings and loan associations in the Federal Home Loan Bank Board's 11th District. These loans had very limited periodic and lifetime rate adjustment parameters which made them unattractive to most lenders.

In 1981, the Comptroller of the Currency issued regulations allowing national banks a broader range of products. Since then, additional deregulation has resulted in what is essentially an open market. The only regulatory constraint, based on federal legislation passed in 1987, is that adjustable rate first and second mortgage loans must have some maximum interest rate as part of the contract. For first mortgage loans, this requirement was almost unnecessary. Competition among lenders and the desire of borrowers to have some protection against interest rate increases had already led to the virtual disappearance of uncapped ARMs.

The typical ARM has several product features. It is indexed to an interest rate (such as the yield on one-year United States Government securities or the rate on six-month negotiable certificates of deposit), normally with an additive spread (such as the index plus 2.00%). The loan interest rate is reset periodically, usually at the same frequency as the maturity of the index (hedging interest rate risk becomes very difficult if the reset frequency does not match the index maturity). Most ARMs today have "teaser" rates--initial rates that are lower than the fully indexed rate.

Interest rate caps come in two forms. An overall cap limits the maximum rate on the loan, regardless of how high the index goes. Early ARMs had symmetrical overall caps (a minimum as well as a maximum), but most ARMs today do not have a minimum rate, simply because it is too easy for borrowers to evade the minimum by refinancing. ARMs also have periodic caps which limit the rate change which can occur at any reset. Periodic caps are designed to avoid extreme changes in payment amounts. Usually, periodic caps operate in both directions. A common ARM indexed to one-year Treasury securities might have the following components.

Table 1
Typical ARM Parameters

Initial Index	7.80%
Spread	2.70%
Reset Interval	1 year
Initial Rate	8.50%
Overall Cap	13.50%
Periodic Cap	2.00%
Fully Indexed Rate	10.50%
Origination Fee	2.00%

If the index does not move, the loan rate will reset to the fully indexed rate of 10.50% one year after origination. Should the index rise, the periodic cap will constrain the first reset, allowing the loan rate to increase to only 10.50%. This is a subtle, but very real, cost of offering teaser rates on ARMs.

MODEL DEVELOPMENT

Since ARMs normally have 30-year maturities, there are 29 rate resets which can occur. The combination of overall and periodic caps results in a total of 87 options embedded in the loan contract. Using the options pricing model to determine the cost of the 87 caps would be incredibly difficult, since the caps are not independent. Simulation appeared to be the only way to arrive at an estimate of the cost of offering caps.

Work began on the project in May, 1987. The decision to construct the simulation model in a spreadsheet language (SYMPHONY) on a personal computer rather than on a timesharing system was motivated by a need to reduce development time. Furthermore, it was anticipated that cap costs would be required for other loan products, and the ease of modifying spreadsheets made the personal computer approach even more desirable. Finally, the technical challenge offered another incentive for this solution.

Writing the simulation model involved making several simplifying assumptions about loans. First, payments (and thus amortization) occur semi-annually, rather than monthly. This substantially reduced the run time for the simulation. Second, the minimum rate change feature of ARMs was ignored. At Bank of America, the loan rate is not reset if the change would be less than 0.125%. Incorporating this feature would have made the model more complicated, but would have had little effect upon the operation of caps.

Finally, loan prepayments were assumed to occur at a constant rate over 30 years. A semi-annual prepayment rate of 5.00% results in an average loan life of approximately seven years. This prepayment rate is presumed to be unrelated to interest rate movements. It matches fairly well the actual prepayment experience on Bank of America's ARM portfolio over the past three years, a period of relatively stable interest rates. Unfortunately, ARMs have not existed long enough to establish what happens to prepayment rates over an entire interest rate cycle.

Initially, the model merely assumed that all loans paid off after seven years. While this type of assumption has often been used in the past for analyzing real estate loans, it is inappropriate for analyzing the risk in offering caps on ARMs. The seven-year payoff model was useful, however, for a validation procedure that will be explained below.

Modelling the changes in the loan interest rate versus the fully indexed rate is not complex, given a set of future values for the index. Coming up with sets of future index rates is the crux of the problem. It was eventually decided to treat the changes in the index rate as a random walk, with a mean change of zero (no secular trend in rates) and a standard deviation derived from historical observations. A reasonable choice would be the most recent 120 monthly observations of the one-year change in the one-year US Treasury yield. Using data from July 1979 through June 1988, this results in a standard deviation of 2.7644%.

It makes a substantial difference which historical period is selected. The same statistic for the last 20 years (July 1969 through June 1988) is 2.2425%, and for the last five years is 1.7865%. As will be shown below, the differences in cap costs are stunning.

Just as important is the question of what distribution to use. The model uses a normal distribution, which has the obvious drawback of occasionally generating negative interest rates. A lognormal distribution eliminates this problem, but has the unfortunate side effect of generating truly astronomical interest rates when a run of positive random numbers appear. An alternative would be a mean reversion process, but this approach demands an assumption about the probability of the mean shifting over a 30-year time period.

The real problem is that the part of the distribution of greatest interest is the tail. How small rate changes are distributed isn't a concern. The crucial question is what happens when rates go to 20.00%, and no one (including the author) has an analytical answer in this area. While we cannot be truly satisfied with the distribution we are currently using, a superior methodology has not emerged to date.

Two more pieces were needed to complete the model. If interest rates fall precipitously, the operation of the periodic cap can result in a customer's loan rate being much higher than the fully indexed rate. At some point, he would simply prepay the loan and refinance, as so many borrowers with fixed-rate real estate loans did in 1986. There is no historical data to indicate what rate differential would trigger this action, but we decided to arbitrarily assume that a differential of 3.00% would automatically cause all loans to prepay. Borrowers would need less than one year to recoup the cost of a new 2.00% origination fee.

The last issue was measuring the cost. The definition of cost used was the difference between the fully indexed rate and the customer loan rate, except during the period up to the first loan reset (any differential during this period is attributed to offering a teaser rate). The model constructs a portfolio, handles prepayments, amortizes loans, resets loan rates, calculates costs every year, and then discounts them to arrive at a present value. Rather than using a constant discount rate, we opted for a variable discount rate which is a function of the index. Since costs are greater in high rate environments, the effect of the variable discount rate is to reduce the present value cost.

The current model is written in LOTUS 123, Version 2.01. A single run of 1000 model iterations takes 19 minutes on an IBM PS-2. This is a major improvement over earlier SYMPHONY versions of the model, which ran on an IBM PC-XT and took over two hours. We discovered that the random number generator in LOTUS 123 does not regenerate the same set of random numbers on a PS-2, even when the system is booted. This problem was solved by creating ten tables of normal random numbers and storing them on the hard disk. The current driver macro is doubly nested, once for reading in the ten random number files, and again for performing 100 iterations. Cap costs from each iteration are stored in a table, then averaged at the end of a run. The mean present value cap cost is converted into a loan spread using a complicated formula involving the initial fully indexed rate, the constant prepayment rate, and a standard discount factor. The spread serves as both an alternate cost measure and a method for analyzing pricing tradeoffs.

RESULTS AND VALIDATION

Results for the loan product described in Table 1 are a present value cap cost of 3.27%, which converts into a spread of 0.84%. The standard deviation around the present value cost is 4.21%. The 95% confidence interval for the cap cost runs from 3.01% to 3.53%. Given the nature of caps, it is not surprising that the confidence interval is rather wide.

Using alternate historical periods with less volatility in the index dramatically reduces cap costs, as shown in Table 2.

Table 2
Present Value Cap Costs
For Differing Volatilities

	10-year History σ = 2.76%	20-year History σ = 2.24%	5-year History σ = 1.79%
Mean Cost	3.27%	2.30%	1.58%
Standard Deviation	4.21%	3.41%	2.57%
Loan Spread	0.84%	0.59%	0.41%

Most loan product managers, when told that the profitability of their loans ought to be assessed 3.27% up front for the cost of caps, first turn white and then question the analysis. Thus the issue of validation became an important one. The most obvious place to look was in the national money markets. Overall caps based on the London Interbank Offered Rate (LIBOR) can be bought in the wholesale markets. Prices for a seven-year cap similar to the overall cap on the ARM were around 3.40% in August, 1987. This is a comparable number, for although LIBOR caps revalue quarterly, and thus the index has more underlying volatility, there is no periodic cap feature. In short, we feel the model generates reasonable results.

A more intriguing validation was a survey of 16 financial analysts and economists in Bank of America, conducted in September 1987. They were given four interest rate scenarios over the next seven years: a flat rate environment, a rising environment, a falling environment, and an environment with rates moving up and down. Each scenario had specific values for the index. Our experts were asked to assign a probability to each scenario, summing to one--they were not allowed to generate their own forecasts. The seven-year horizon was selected because it matched the average life of ARMs and was an acceptable time period for forecasting.

It was quite easy to run the four scenarios against a deterministic version of the simulation model in which all loans pay off at the end of seven years, calculate the cap costs for each scenario, and then weight the costs by the probabilities our 16 experts had assigned. The difference in the present value cap cost was a miniscule 0.06%.

These two validations seem to confirm that what we are doing makes sense, but the real test would be in the secondary market. Unfortunately, the major secondary market for ARMs is in loans indexed to the Federal Home Loan Bank Board's 11th District Cost of Funds. This index is only a market interest rate in the broadest sense, and its history can be used for determining volatility only by assuming that the structure of the savings and loan industry has remained unchanged through the interest rate upheavals of 1979-1980, the subsequent deregulation, and the current difficulties certain institutions are experiencing.

MODEL USAGE AND IMPLICATIONS

Theoretically, the model can be used to price ARMs. In fact, the number of institutions offering standard ARM products is large and consumers are very price sensitive. Pricing with the market is virtually mandatory. A more practical use, should broad secondary markets develop, would be to decide whether to hold or sell ARMs. Another use of the model is to test product options. Table 3 below shows present value cap costs for differing versions of the product described in Table 1.

Table 3
Present Value Cap Costs
For Differing Product Features

Periodic Cap	Overall Cap 13.50%	14.50%	15.50%
None	2.46%	1.81%	1.32%
3.00%	2.71%	2.17%	1.82%
2.00%	3.27%	2.87%	2.62%
1.00%	4.91%	4.77%	4.68%

Thus an alternative to the standard product might be a 14.50% overall cap, but an origination fee of only 1.75%. All of these results were generated from the same set of random numbers, in order to obtain the best possible estimate of the relative cost of pricing options.

The cap costs computed above have been used in conjunction with other financial data to determine the overall profitability of making adjustable rate mortgages. The return on assets, given any set of reasonable assumptions about origination and servicing costs, is invariably substandard once the full cost of caps is considered. ARMs appear to be a relative bargain for the consumer. Why has the market priced them so cheaply?

Several potential answers come to mind. One is that lenders just don't realize that the cost of caps is so large. Unlike the cost of funds or operating costs, determining cap costs is a venture into uncharted territory for most institutions.

Another explanation is the one currently in vogue for all problems of American business: too much of a focus on the short term. After all, the biggest cost of offering caps on 30-year loans is most likely to occur far in the future, where uncertainty is greatest. By that time, senior management will have retired. Yet the wide prevalence of teaser rates on ARMs would seem to partially negate this argument.

A third possible answer is that the combined wisdom of the majority of lenders is that the bad days of the early 1980's will never return. In effect, lenders are betting on interest rates. If so, they are going against the combined wisdom of the money market traders.

Lastly, it may be that real estate lenders, and regulators of real estate lenders, want to avoid a repetititon of the early 1980's at all costs. In their effort to avoid interest rate risk by matching the apparent rate maturities of assets and liabilities, lenders have been willing to give away caps. The problem is not as serious as the huge mismatches of the past, but it is still one that is only likely to be corrected by the next round of major interest rate increases.

ACKNOWLEDGEMENTS

The author wishes to extend his special thanks to two major contributors to this project. Mark Williams in Bank of America's Corporate Treasury Department provided much insight into the distribution of interest rates, as well as backstopping some of the important assumptions in the model. Greg Mischel in Financial Analysis and Planning reviewed the overall structure of the model and made the suggestion to switch to the use of a constant prepayment rate with a full 30-year analysis. Key support was also received from Joe Barnard, the author's department head, and Dave Hanna, finance manager for the California Banking Group.

BIOGRAPHY

THOMAS W. McDOWELL is a manager of finance at Bank of America, working in the California Banking Group's Financial Analysis Department. He received a Bachelor of Science degree from Stanford in 1966 and obtained his MBA from Northwestern in 1970. Since then, he has held a variety of positions, all involving financial analysis, at Bank of America. He is a member of The Institute of Management Sciences and a past secretary-treasurer of the Northern California Chapter.

Bank of America
CBG Financial Analysis #3582
Post Office Box 37000
San Francisco, CA 94137
(415) 622-3125