QUANTITATIVE PERSPECTIVES

INTEREST RATE MODELING: A CONSCIENTIOUS CHOICE

By ALEX LEVIN

September 2002

ANDREW DAVIDSON & CO., INC.

QP

INTEREST RATE MODELING: A CONSCIENTIOUS CHOICE

By Alex Levin

SEPTEMBER 2002

QUANTITATIVE PERSPECTIVES

Introduction

Intensive developments in the field of interest rate modeling have delivered a bold but confusing model selection choice for financial engineers, risk managers, and investment analysts. Do these modeling issues sound familiar?

- Should a mortgage bank assess the interest rate risk using the lognormal Black-Karasinski model or using the normal Hull-White model?
- Can a portfolio be hedged using different pricing models for assets and derivatives?
- Is there any historical evidence that one model is better than another?
- What does the market "think" about the interest rate distribution? (It must "know" something - otherwise how would the rate options be traded?)

In this issue of *Quantitative Perspectives,* we show that selecting the "best" term structure model is becoming more of a conscientious task than a matter of taste. We show that both recent historical rates and implied volatility skew for swaptions confirm rate "normalization" and reject the conjecture of lognormality. We propose valuing mortgages using the Hull-White model, which can be quickly and accurately calibrated to both the yield curve and the swaption volatility matrix. Being fully aware of these market indications and tendencies, we have expanded and enhanced our Vectors™ suite of analytical models.

Those who read research on how interest rates performed in 1980s and early 1990s are accustomed to the conjecture of lognormality. According to it, interest rates are lognormally distributed, i.e., their logarithm is normally distributed. The rates, therefore, cannot become negative, and their randomness should be naturally and steadily measured by relative volatility. This conjecture had enforced the validity and applicability of the Black-Scholes pricing model to the interest rate option market.

Following Paul Wilmott [1998] we will be measuring volatility by plotting the averaged daily increments versus the rate level. Namely, we can collect all daily rate increments and store them in buckets; each bucket corresponds to some rate level. For example, a 7% bucket includes all daily increments when the rate was between 6.5% and 7.5%. After the data is collected, we average increments using the root-mean-squared formula applied within each bucket and then annualize them. Although the U.S. Treasury rates are currently not the best benchmark for mortgages, they have the longest history (Figure 1).

Figure 1

Daily volatility vs. level for the 10-yr Treasury rate

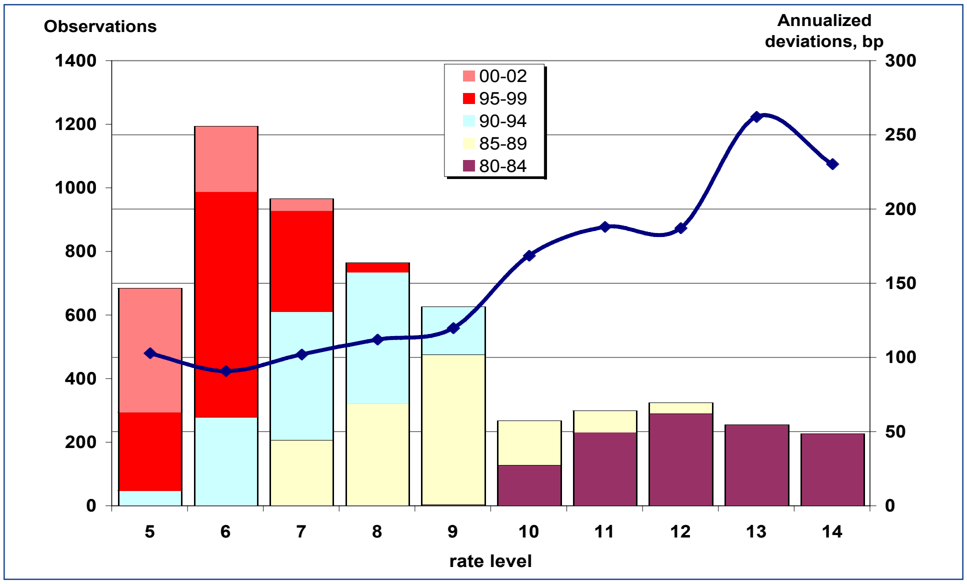

First, let us disregard the bars and look only at the blue line depicting historical volatility measured by annualized deviations (right axis). The absolute historical volatility seems to remain very much independent of rates in the left half of the chart (west of 10%). When the rates are in the double-digits, the same absolute volatility measure grows with the rate level. Now, reading the historical labeled bars, we conclude that the absolute volatility has become rate-independent since late 1980s. This conclusion is confirmed by a similar analysis performed for the 10-yr swap rate history dating back to 1989 (Figure 2).

Figure 2

Daily volatility vs. level for the 10-yr swap rate

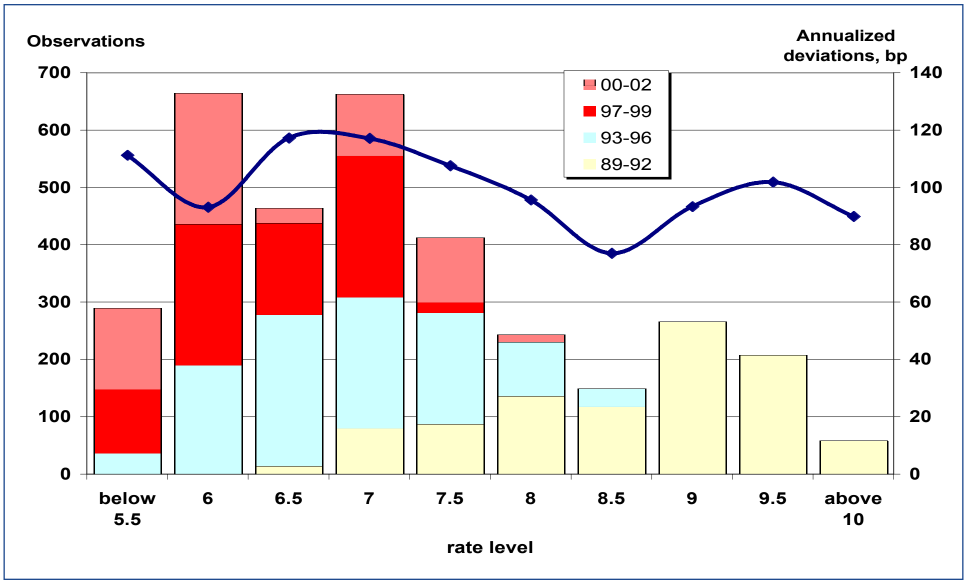

A weak or absent relation between absolute volatility and rate level is a sign of normality rather than lognormality. It also prompts quoting rate uncertainty (and, therefore, option prices) in terms of absolute volatility (such as 110 basis points) rather than relative volatility (say, 20%). Recently, many brokers have begun communicating exactly in that way.

What Does the Swaption Market Think?

Can we recover the rate distribution from the way interest rate options trade? A simple way is to "measure" the implied volatility skew, i.e., dependence of the implied Black volatility on the strike level. If the market participants believed in lognormality, there would exist little reason for the implied volatility to change with the option's strike. Volatility skew testifies against lognormality by the very fact of its existence. In order to discuss a simple "skew measurement" method, let us first introduce a setup that generalizes many known and popular single-factor models, a CEV (constant elasticity of variance) model:

$$dr = (Drift)dt + \sigma r^{\gamma} dz \qquad (1)$$

where r is some modeled rate, σ is the volatility coefficient, γ is the CEV constant. As usual, $z(t)$ is the Brownian motion that disturbs the market, t is time, and the exact specification of the drift term is not very important for our purposes. Mark Rubenstein proposed this model in the 1980s; it has no specific economic meaning but can be viewed as a convenient way to generalize and compare all known popular models.[1] Indeed, for $\gamma = 1$, the absolute volatility is proportional to the rate and we have a lognormal model (with a properly selected drift term), such as Black, Derman, and Toy (BDT) or Black-Karasinski (BK). For $\gamma = 0$, the absolute volatility is rate-independent and can lead to a normal model, such as Hull and White (HW). If $\gamma = 0.5$, we may have a popular family of "square-root" models, such as the Squared Gaussian model (SG), or the model of Cox, Ingersoll, and Ross (CIR). Any "unnamed" values for the CEV are certainly possible, including negative values (hyper-normality) and values exceeding 1 (hyper-lognormality).

Blyth and Uglum [1999] proposed a simple method of recovering the most suitable CEV constant by just looking at the observed swaption volatility skew. They argue that, if a swap forward rate satisfies the random process (1), then the skew should have the following approximate form:

$$\frac{\sigma_K}{\sigma_F} \approx \left(\frac{F}{K}\right)^{\frac{1-\gamma}{2}} \tag{2}$$

In the this formula, F is today's forward rate, K is the swaption strike, σ_F and σ_K are the Black volatilities, for the at-the-money strike (F) and the actual strike (K), correspondingly.

Let us analyze the same set of CEV special values, 0, 1, and 0.5. If $\gamma = 1$, then there will be no skew at all: $\sigma_K \equiv \sigma_F$ for any strike K. This is the Black-Scholes case. For $\gamma = 0$, the skew has a functional form of inverse square root. For $\gamma = 0.5$, it will have the shape of inverse fourth-degree root. It is worth mentioning here that each inverse root function is a convex one (see Figure 3 further in the text), and, therefore, the theoretical skew should not be deemed a straight line (except when $\gamma = 1$). In fact, it should not be confused with a more aggressive convex volatility "smile" that may or may not be present in addition to the skew.[2]

[1]The CEV model is mentioned and analyzed in many published sources, including books by P. Wilmott [1998], and J. James and N. Webber [2000].

[2]This smile can be explained by a "jumpy" nature of the underlying rate. For example, LIBOR caps are traded with a considerable volatility smile because these rates are subject to regulators' interference. Swap rates are much more diffusive than jumpy, therefore, swaptions' smiles are much less pronounced.

The object of this study - the 5-into-10 swaption - was selected with modeling volatility of mortgage rates and valuation of the prepayment option in mind. Figure 3 depicts five skew lines plotted for three "named" CEVs, the actual volatility observations averaged since the beginning of 1998, and the optimal-fit line ($\gamma = 0.23$). The best CEV is therefore found to be generally between the normal case (HW) and the square-root case (SG or CIR). It is also seen that low-struck options are traded with a close-to-normal volatility, whereas high-struck options are traded with a square-root volatility. This phenomenon may be a combination of a slight theoretical smile (see footnote 2) and the broker commission demand.

Figure 3

Average volatility skew for the 5-yr into 10-yr swaption

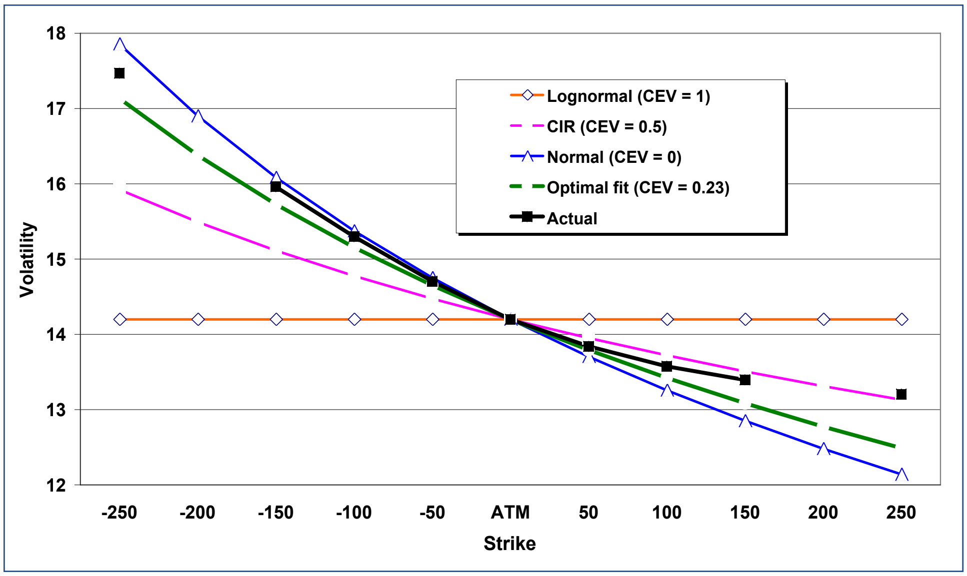

Source: Bank of America; volatility for 200 ITM/OTM was not quoted.

Figure 4 illustrates historical month-by-month skew suggesting that the normalization effect ($\gamma \approx 0$) has been observed lately.

Figure 4

Historical CEV Values

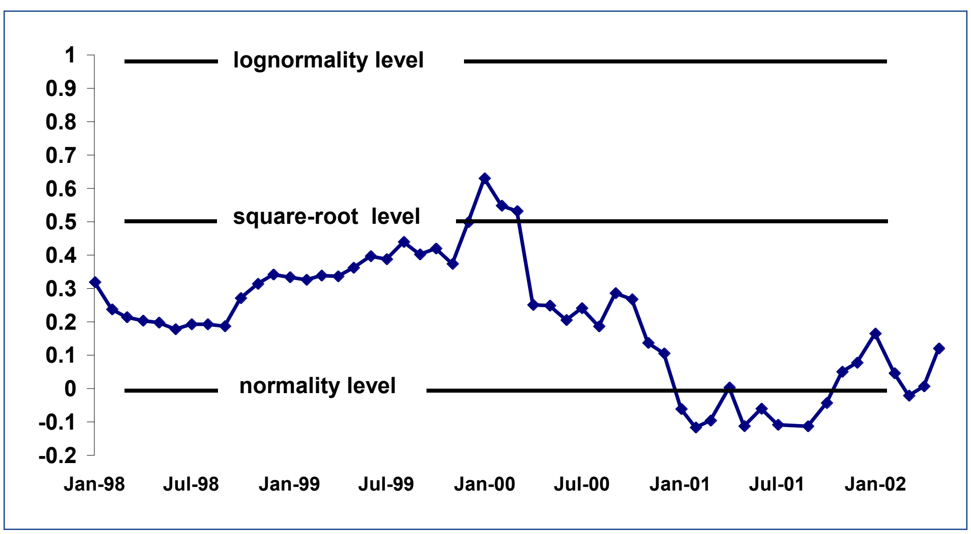

Source: Bank of America

This CEV analysis unambiguously rejects lognormality and reveals a more suitable model. Although the best-fit CEV constant somewhat varies, any volatility model between the normal one and the square root seems to be a decent choice. Due to its analytical tractability and the recent CEV trend (Figure 4), we will focus on the HW model as the chief alternative to the BDT or BK models.

Other Problems with Lognormality

Although the BK model has been the bread and butter of option traders since its inception, its full-scale implementation required for good mortgage analytics is not a simple task.

- Short-rate lognormal models are not analytically tractable. For example, a Monte-Carlo, or any other forward sampling method employed as the primary mortgage pricing tool, will simulate only the short-rate process on its own. Analytics that would map this process into long-rate (a mortgage-rate) dynamics simply do not exist and need time-consuming numerical replacements.

- Relative (Black) volatility is used <u>for quotation only</u>; it is merely a price-volatility conversion tool, not requiring any belief in the Black-Scholes model. As we have seen, volatility changes drastically with the level of rates, and therefore, with expiration of traded options. One constant number cannot describe the entire universe of swaptions that are deemed relevant for mortgage pricing. Using the BK model with constant volatility cannot be recommended in view of a steep yield curve and sharply fallen proportional volatility.

- In contrast to a common belief, in the BDT and the BK models, long rates are not lognormal and have lower volatility than that of the short rate, even in the absence of mean reversion. This may sound like an unpleasant surprise for those who think that a 20% short-rate volatility plugged into the model results in a 20% swaption volatility. Therefore, model calibration to the mortgage-relevant options (not the options on the short rate!) can be complicated.

The short rate in the HW model is driven by a linear stochastic differential equation, which is a special case of the CEV model (1):

$$dr = a(t)(\theta(t) - r)dt + \sigma(t)dz \tag{3}$$

where $a(t)$ denotes mean reversion, $\sigma(t)$ stands for volatility; both can be time-dependent. Function $\theta(t)$ is sometimes referred to as "arbitrage-free" drift. This terminology is due to the fact that, by selecting proper $\theta(t)$, we can match any observed yield curve.

Since (3) is a linear differential equation disturbed by the normally distributed Brownian motion, its output, the short-rate process, will also be normally distributed. In particular, the negative rates are not precluded; we discuss this issue further. Although this fact is well known and never met with enthusiasm among practitioners, there are many advantages in the model that make up for this small drawback.

The model possesses full analytical tractability. For example, the arbitrage-free function $\theta(t)$ is expressed analytically through a given forward curve. The average zero-coupon rates and their standard deviations are also known for any maturity and any forward time. Any long zero-coupon rate r_T of arbitrary maturity T is proven to be also normally distributed and linear in the short rate; volatilities are related as

$$\frac{\text{Long - rate Volatility}}{\text{Short - rate Volatility}} = \frac{1 - e^{-aT}}{aT} \equiv B_T \tag{4}$$

at any moment of time t.

Function B_T of maturity T plays an important role in the HW model. It allows for calibrating the volatility function $\sigma(t)$ to the option market. If the mean reversion is positive, then $B_T < 1$, and the model allows for "quasi-parallel" shocks with rate deviations being gradually depressed along the curve. This feature agrees with the behavior of absolute implied volatility for traded swaptions: it generally falls with the swap maturity. This observation therefore helps to calibrate mean reversion in the model.

If $a = 0$, function B_T becomes identical to 1, regardless of maturity T. This important special case, called the Ho-Lee model, allows for a pure parallel change in the entire curve (every point moves by the same amount). Such an opportunity can be advantageous for standardized risk measurement tests. No other model allows parallel shocks to be mathematically consistent with its internal analytics.

The Hull-White Model: Calibration to ATM Swaptions

Because the standard deviation of any zero-coupon rate can be found explicitly for any bond maturity and any forward time, it can be directly compared with quoted Black volatility. Although market swaps are coupon-bearing instruments, this zero-coupon volatility analysis remains quite accurate within the maturity range deemed relevant for the mortgage market (up to 10 years). Whether the model operates with a time-dependent volatility function $\sigma(t)$ or with a constant volatility parameter, $\sigma(t) \equiv \sigma = const$, it can be optimized to approximate the volatility matrix of traded ATM swaptions. This calibration procedure certainly includes finding the best mean reversion parameter a in the sense discussed above.

Figure 5 presents the calibration results (lines) using a series of ATM options on the 2-year swap and on the 10-year swaps as the input. The bars for six different expirations show known volatility quotes converted into the absolute form (i.e., the relative quote multiplied by the forward rate).[3]

Figure 5

Calibrated volatility term structure - May 2002

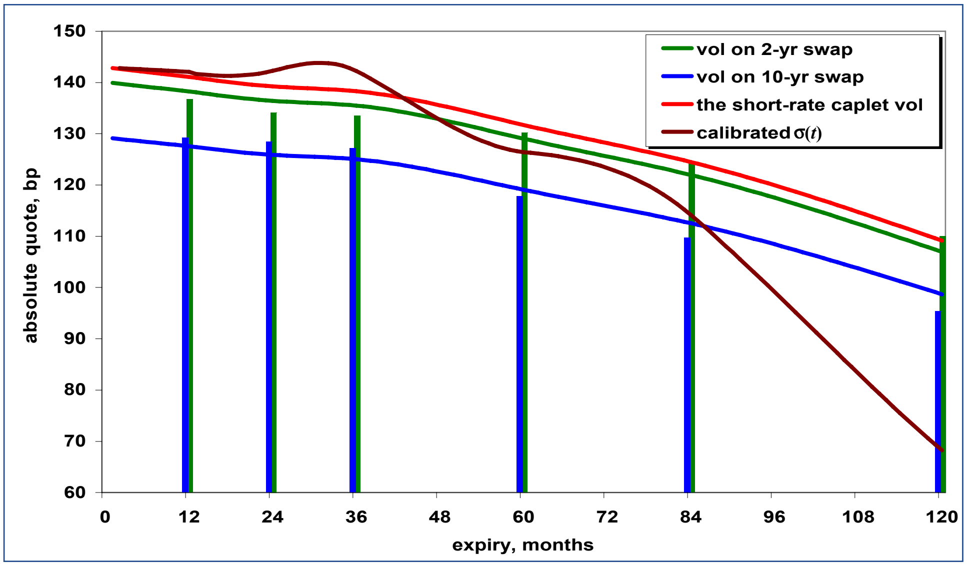

· *Lines: calibrated HW model* · *Bars: ATM market quotes* · *Date: 05/13/2002*

[3]This method of conversion is proven to be more accurate in matching option values under the lognormal and normal versions of Black-Scholes than the formal variance match.

The overall effective error of calibration is just 3.6 basis points of absolute volatility, as measured across the swaption matrix. The mean reversion parameter was restricted to be a constant; its best value is found as $a = 2.05\%$. For some particular applications (like standardized risk tests mentioned above), one may prefer having a zero mean reversion, or/and a constant volatility parameter. These restrictions generally reduce the calibration accuracy as shown in Table 1.

Table 1

Comparative calibration for the HW model

Volatility, bp (best fit)	Mean reversion,%	Accuracy*, bp
123.1 (constant)	Set to zero	12.4
145.9 (constant)	3.36 (best fit)	6.8
Time-dependent	Set to zero	7.8
Time-dependent	2.05 (best fit)	3.6

* *Effective RMSE measured across the volatility matrix; Date: 05/13/2002*

Calibrated σ in Figure 5 is rather responsive to the slope and the shape of the input volatility structure. It sharply falls beyond the 6-year horizon, perhaps as a result of market perception about the current versus the long-term volatility. During the stormy market of the first half of 2002, short-dated options were, indeed, traded at unprecedented absolute volatility levels of 125 - 140 basis points, well above their long-term averages. This was not the case in the calm August of 1998, just prior to the Russian crisis (Figure 6).

Figure 6

Calibrated volatility term structure - August 1998

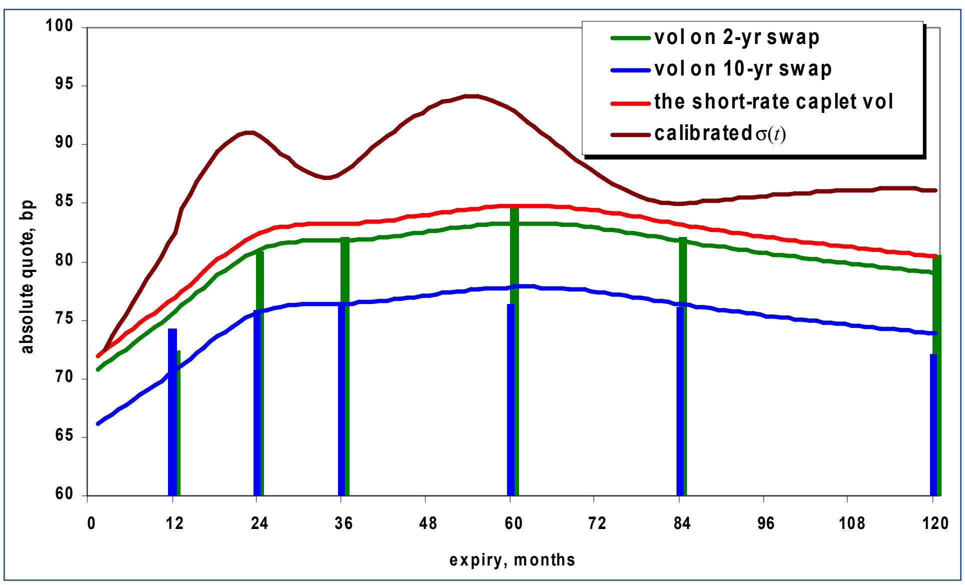

· *Lines: calibrated HW model* · *Bars: ATM market quotes* · *Date: 08/19/1998*

We have demonstrated that the HW model can be calibrated to the traded swaptions. Would it be an even easier task to use caps? After all, the function we seek, $\sigma(t)$, is the "short-rate" volatility function, and the short (LIBOR) rates seem to be good candidates to examine volatility. Although many market participants perceive caps and swaptions traded in unison, they may overlook an important modeling difference--the jumps. Long swaps are chiefly diffusive, and the model disturbed by a Brownian motion [like $z(t)$ in equation (1)] makes sense. Short rates (LIBORs) combine continuous diffusion (small day-after-day changes) with sudden regulatory interferences. It is mathematically not very difficult to add jumps to diffusion (Robert Merton did it in 1974), but the equivalent volatility term structure will become "humpy". Under a jump or a jump-diffusion disturbance, the short-dated Black volatilities come up considerably suppressed. Figure 7 compares volatilities for traded caps (filled marks) with volatilities that came from the swaption-fitted HW model.

Figure 7

Implied cap pricing - May 2002

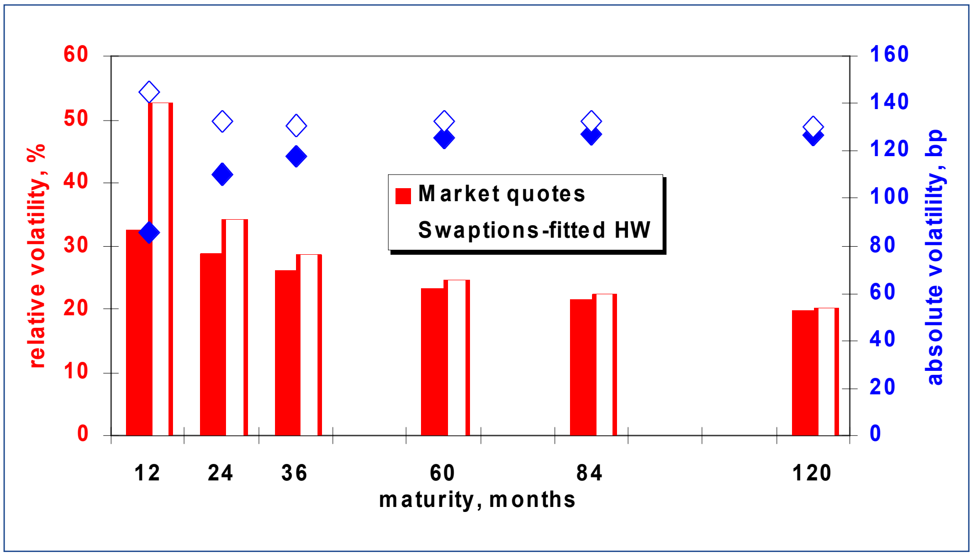

As seen from this chart, the model drastically overstates short-dated cap volatilities, both in absolute and in relative terms. However, as the cap's maturity extends, swaptions and caps seem to converge. Can the cap (rather than the swaption) volatility structure be plugged into a mortgage pricing system? Perhaps, it could, if the system's interest rate model maintained the jump-diffusion setting. Since developers of mortgage analytical systems do not do so, the blind use of caps will understate volatility and, therefore, the prepay option value. In conjunction with traditional diffusion models, we prefer using the swaption market for benchmarking volatility--especially for fixed-rate mortgages. Valuation of ARMs may need additional attention in the view of embedded reset caps.

From BK to HW: Valuation & Risk Management Implications

Let us assume that the models in question, the normal HW model and the lognormal BK model, were independently calibrated to the ATM swaptions. This is to say that these two models should value ATM swaptions identically. However, the volatility skew of the normal model (curve $\gamma = 0$ in Figure 3) is quite unlike the one for lognormal model ($\gamma = 1$). That means that these two models will value any other option differently, except for those employed for calibration, the ATM ones.

Since embedded options found in the mortgage market (prepayment option, ARM caps and floors, clean up calls) are spread over time and instruments, transitioning from the BK model to the HW model will generally result in a change of values. Even more importantly, the interest rate sensitivity measures will change considerably - as a direct result of different volatility specification. Under the BK model, every "up" move in rates proportionally inflates the absolute volatility. Consequently, the embedded prepayment option will additionally reduce the modeled value of the MBS. This can be considered an indirect (via volatility) interest rate effect making the effective duration of mortgages artificially longer. Figure 8 shows a 0.4-year duration reduction when moving from the BK model to the HW model, for the current-coupon agency. This may considerably re-quantify the Delta-hedging needs in secondary marketing and MBS portfolio management.

Figure 8 and Table 2 summarizes comparative valuation results for 30-yr fixed rate agencies obtained under three different term structure models. In each case, the short-rate volatility function was calibrated to the ATM swaptions. As one could expect, "cuspy" mortgages located at the center of ADCo's refinancing curve ("at the money" FNCL7 on Figure 8) are valued in a very close OAS range by all three models. When the prepay option is out of the money (the discount sector), this option will be triggered in the falling rate environment. This sector therefore looks relatively rich under the HW model, whereas the premium sector benefits from using this model. The SG model produces valuation results that are between those of the HW and the BK models.

Figure 8

LIBOR OAS and Duration profiles for new FNCLs

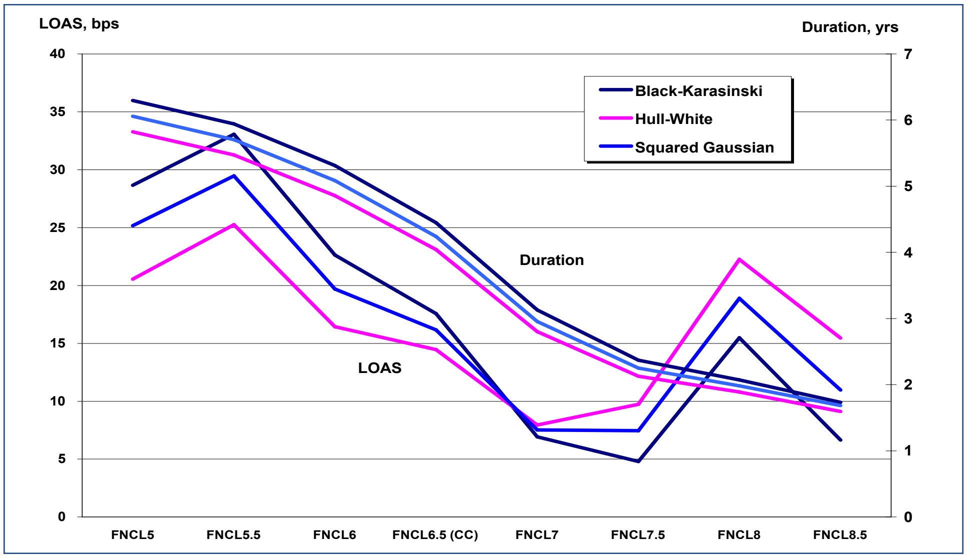

Table 2

Comparative valuation results for 30-yr fixed rate agencies

	Libor OAS			Effective Duration		
	HW	SG	BK	HW	SG	BK
GNSF5	10.5	15.7	19.6	6.13	6.38	6.66
GNSF5.5	15.3	20.1	24.0	5.75	6.00	6.28
GNSF6	4.3	8.0	11.1	5.15	5.41	5.68
GNSF6.5 (CC)	6.4	8.4	10.0	4.28	4.50	4.74
GNSF7	5.9	5.7	5.1	3.20	3.38	3.56
GNSF7.5	7.0	5.1	2.8	2.29	2.43	2.55
GNSF8	21.6	18.2	14.6	2.02	2.14	2.22
GNSF8.5	28.4	24.0	19.6	1.77	1.87	1.92
FNCL5	20.6	25.2	28.7	5.82	6.06	6.30
FNCL5.5	25.3	29.5	33.1	5.47	5.70	5.94
FNCL6	16.4	19.7	22.6	4.86	5.08	5.31
FNCL6.5 (CC)	14.5	16.2	17.6	4.04	4.24	4.45
FNCL7	7.9	7.5	6.9	2.80	2.95	3.13
FNCL7.5	9.7	7.4	4.8	2.13	2.25	2.37
FNCL8	22.2	18.9	15.5	1.89	1.98	2.07
FNCL8.5	15.4	11.0	6.6	1.60	1.68	1.73

Date: April 4, 2002

Although most mortgage instruments will look "shorter" under the HW model, there are some notable exceptions. As we explained above, the primary divergence of HW from BK is found in differing volatility models. Since mortgage IOs and MSRs have drastically changing convexity profiles, they will also have unsteady exposures to volatility, i.e. Vega. For example, Vega is typically positive for an IO taken from a premium pool (case 1), negative for the one stripped off a discount pool (case 2) and about zero for the case when the pool's rate is at the center of refinancing curve (slightly above par, case 3). Therefore, the BK model will generally overstate the rate sensitivity for case 1, understate it for case 2, and will be close to the HW model in case 3 (Figure 9). When arriving at these conclusions, we kept in mind that an IO value, in contrast to regular MBS, grows with the rates. These interesting findings, though effecting Delta-hedging, do not contradict what is well known: IOs, POs, and MSRs are influenced greatly by prepayments and slightly by interest rate models.

Can two different models be used for risk management: one for the assets, another for the hedge? Suppose a mortgage desk blindly uses the BK model, whereas a swap desk trades with a skew. Unless the position is made Vega-neutral, differing volatility specifications in the models may considerably reduce hedge efficiency.

Figure 9

Valuation results for an IO stripped off a new current coupon FNCL pool

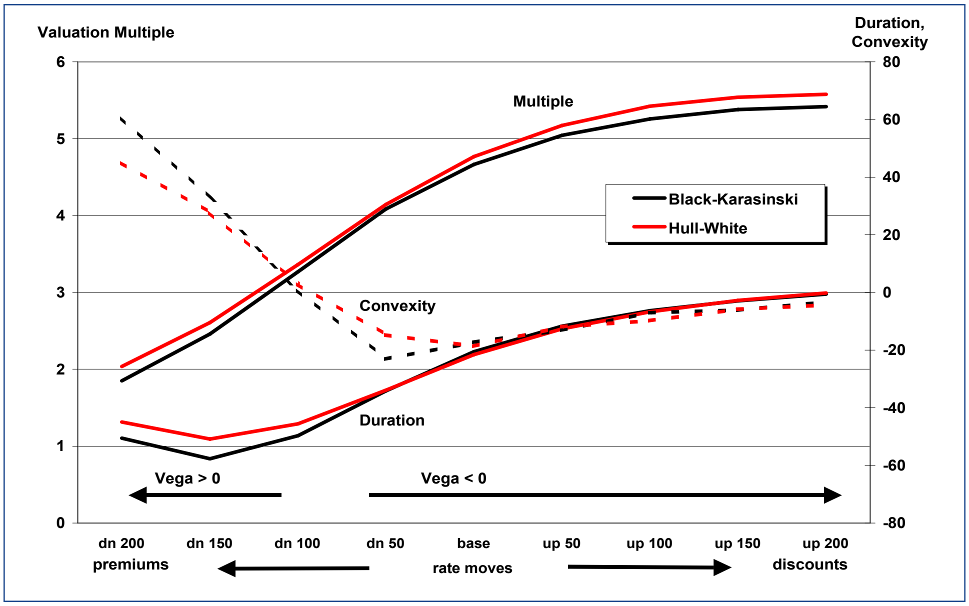

Date: April 4, 2002, OAS = 200 bps for both models.

The Issue of Negative Rates

Knowing that interest rates have never been negative in U.S. history, we should question what detrimental effects might occur when using the HW model. Many often try to estimate the probability of getting negative rates in the model; this naïve approach typically mounts needless scare. Indeed, the odds of such an event is far from infinitesimal, but how badly can it damage the value of an MBS?

Answering this question may be simpler that it sounds. Here is a hint: consider a LIBOR floor struck at zero. This non-existing derivative will have a sure zero practical value but not under the HW model. We have priced this hypothetical instrument using the worst-case scenario: market as of May 2002 (low rates, high volatility) with model's mean reversion set to zero. The volatility function $\sigma(t)$ was calibrated to the swaption market and extrapolated beyond the 10-year expiration. The value of a zero-struck floor was found insignificant for the average life range relevant in mortgage pricing (up to 10 years). Thus, for the 10-year non-amortizing floor, the value was 7 basis points, which is equivalent to a 0.8 basis point error in the OAS. The issue grows with the horizon: the 30-year floor was priced at 35 basis points, which would lead to 2.5 basis points of spurious OAS. We can conclude that the HW model is rather harmless; it will not lead to a sizable mispricing even in the worst, mortgage-irrelevant case. This conclusion certainly needs periodic reviews as interest rates continue to fall.

New Vectors™ Suite of Term Structure Models

Andrew Davidson & Co., Inc.'s interest rate library includes all three models considered above, the Black-Karasinski model, the Squared Gaussian model, and the Hull-White model. Each of these models, implemented on a lattice, can operate with time-dependent volatility as well as with constant volatility. One can fit them (with mean reversion) to the observed volatility term structure for swaptions rather accurately and quickly.

Special analytics for convexity adjustment, exact for the HW model, very accurate for the SG model, and approximate for the BK model, has been implemented to facilitate the no-arbitrage condition between the user-defined benchmark nodes. The relevant set of long rates requested by the user is produced on the lattice and agrees with the underlying short rate process.

The functionality of this library is integrated with the OAS system or can be used on its own. We recommend using the HW model with time-dependent short-rate volatility function and constant mean reversion calibrated to a set of swaptions on 2- and 10-year swaps with all available expiries ranging between 1 year to 10 years.

■

Acknowledgements This issue of *Quantitative Perspectives* has greatly benefited from joint research with Andrew Davidson (many views are as much his as mine) and insightful comments from Yung Lim, Rob Landauer, and Jim Barrett. The volatility skew data used in the analysis were obtained courtesy of Craig Lindemann and Krystn Paternostro. I would like to thank Steve Heller and Jay Delong who incorporated the new library of interest rate models into the OAS and other ADCo products. This publication would not be possible without the diligent production efforts of Ilda Pozhegu.

References

Black, F and Karasinski, P, Bond and option pricing when short rates are lognormal, *Financial Analysts Journal*, July-August 1991, pp.52-59.

Blyth, S., Uglum, J., Rates of skew, *Risk*, July 1999, pp. 61-63.

Hull, J., *Options, Futures, and Other Derivative Securities*, (4th edition), Prentice Hall, Inglewood Cliffs, NJ 2000.

James, J. and N. Webber, *Interest Rate Modeling: Financial Engineering*, John Wiley and Sons, 2000.

Levin, A., Deriving closed-form solutions for Gaussian pricing models: a systematic time-domain approach, *International Journal of Theoretical and Applied Finance*, vol.1, No.3 1998, pp. 349-376.

Wilmott, P., *Derivatives*, John Wiley & Sons, 1998.

Further Related Reading

A Lattice Implementation of the Black-Karasinski Rate Process June ‘00

Got to www.ad-co.com for our *Quantitative Perspectives* Library

About the Author

Alex Levin is a Senior Quantitative Developer/Consultant responsible for research and development of analytical models for mortgages and other financial instruments and related consulting work at Andrew Davidson & Co., Inc.

Until March of 2002, Alex was a Senior Vice President and Director of Treasury Research and Analytics at The Dime Bancorp (the Dime) in New York. At the Dime, he was in charge of developing efficient numerical and analytical tools for pricing and modeling mortgages, options, deposits, and other complex term-structure-contingent derivatives, risk measurement and management. He has authored Mortgage Solutions, Deposit Solutions and Option Solutions, the Dime's proprietary pricing systems that were intensively used for both security trading and risk assessment. In addition, he was regularly involved in measuring market risk (including a counter-party value-based risk) and hedging positions in various lines of the banking business as well as helping traders and analysts understand quantitative aspects of pricing and risk.

Prior to his employ at the Dime, Alex taught at The City College of NY and worked at Ryan Labs, a fixed income research and money management company.

Alex is a regular speaker at the Mathematical Finance Seminar (NYU, Courant Institute) and has published a number of papers. He holds an M.S. in Applied Mathematics from Naval Engineering Institute, Leningrad and a Ph.D. in Control and Dynamic Systems from Leningrad State University.

ANDREW DAVIDSON & CO., INC.
520 BROADWAY, EIGHTH FLOOR
NEW YORK, NY 10012
212•274•9075
FAX 212•274•0545
E-MAIL mail@ad-co.com

QUANTITATIVE PERSPECTIVES is available via Bloomberg and the Internet.
BLOOMBERG ADQP<GO>
WEBSITE www.ad-co.com

We welcome your comments and suggestions.

ANDREW DAVIDSON & CO., INC.