A New Strategy for Dynamically Hedging Mortgage-Backed Securities

Jacob Boudoukh
is associate professor of finance and international business at the Stern School of Business of New York University in New York.

Matthew Richardson
is associate professor of finance at the Stern School.

Richard Stanton
is assistant professor of finance at the Haas School of Business of the University of California at Berkeley.

Robert F. Whitelaw
is assistant professor of finance at the Stern School.

This article develops a new approach for hedging mortgage-backed securities (MBS) that involves estimating the joint distribution of returns on MBS and T-note futures, conditional on economic conditions. The resulting hedge ratio is calculated by differentially weighting past pairs of MBS and T-note futures returns, where the weights depend on how close current economic variables are to their historical values.

In an out-of-sample hedging exercise, using weekly and monthly returns on 8%, 9%, and 10% GNMAs over the 1990-1994 period, the dynamic approach is very successful at hedging out interest rate risk. For example, in hedging weekly returns on 10% GNMAs, the method reduces the volatility of the return from 41 to 24 basis points, while a static method achieves only 29 basis points of residual volatility. Moreover, only 1 basis point of the volatility of the dynamically hedged return can be attributed to risk associated with U.S. Treasuries, compared to 14 basis points of interest rate risk in the statically hedged return.

Institutions hold significant positions in mortgage-backed securities (MBS) for a variety of reasons. Whether these positions reflect trades on relative value or involve inventory holdings due to core businesses, hedging the interest rate risk of these securities is an important concern. This is especially true given the well-documented cases of huge monetary losses incurred by financial institutions and investment groups with respect to their MBS portfolios.

MBS valuation (and, by extension, hedging) is not a straightforward exercise. While fixed-rate MBS issued by government agencies represent default-free claims to the interest and principal of the underlying mortgages, the timing of these cash flows depends on the prepayment behavior of the pool. In particular, as interest rates fall, people have an incentive to refinance existing mortgages at the new lower rates. Thus, fixed-rate MBS investors are implicitly writing a call option on the corresponding fixed-rate bond.

Even though prepayments can occur for reasons not associated with interest rate movements, interest rates are the predominant factor in valuing MBS. For this reason, U.S. Treasury securities, or, more specifically, Treasury note (T-note) futures, are often used to hedge MBS. The reasons are twofold: 1) T-note futures are very liquid derivatives, and 2) the prices of these instruments are determined by the underlying term structure of interest rates and thus relate directly to the value of MBS.

There are two common approaches to hedging MBS using T-note futures. The first is purely empirical and involves the regression of past returns on MBS against past returns on T-note futures. The resulting relation can then be used to hedge the interest rate risk of MBS using the risk in T-notes.

The advantage of this method is that it does not involve strong assumptions regarding the underlying model for the evolution of interest rates or prepayments. The disadvantage, however, is severe. This method is static in nature. It does not explicitly adjust the hedge ratio for changes in interest rates and mortgage prepayments.[1] That is, the observations used in the regression represent an average of the relation between MBS and T-note futures only over the sample period, which may or may not be representative of the current period.

As an alternative, the second approach is model-based. It involves specification of the interest rate process and a prepayment model. These assumptions then help map an MBS pricing functional to interest rates and possibly other factors.[2] The approach represents a dynamic method for determining comovements between MBS prices and T-note futures prices. Conditional on current values of the relevant economic variables and on particular parameter values, these comovements are completely specified.

The problems with this approach are twofold. First, there is no consensus regarding what is a reasonable specification of how the term structure moves through time, and how these movements relate to prepayment behavior. Any model price is going to be tied closely to these possibly ad hoc assumptions.[3]

Second, and more subtle, is the recognition that the parameter values themselves may often be "chosen" or estimated from a static viewpoint. For example, empirical prepayment models often reflect ad hoc projections of prepayment rates on sets of housing and interest rate factors. Do the resulting coefficients, which represent an average of the relation in the past, have the same link to the variables describing the current period? Many of the well-documented MBS hedging fiascoes would imply that this is not the case.

In this article, we provide a first pass at bridging the gap between the regression and the model-based approaches. Our goal is to maintain some of the distribution-free properties of the purely empirical method, while recognizing the important dynamic properties of MBS and other fixed-income instruments.

In particular, we propose a method based on estimating the conditional probability density of MBS returns, T-note futures returns, and relevant current information (such as the level, slope, and curvature of the term structure, interest rate volatility, and prepayment history). This method allows us to hedge MBS with T-note futures, conditional on current information. Thus, the hedge ratio is derived in a similar way as hedge ratios using existing empirical methods, but is dynamically adjusted depending on the current state of the economy.

Using data over the period 1987-1994, and performing an out-of-sample analysis, we find that this dynamic hedging method performs considerably better than the static regression method. For example, in hedging weekly returns on 10% GNMAs, our dynamic method reduces the volatility of the GNMA return from 41 to 24 basis points, while a static method manages only 29 basis points of residual volatility. Moreover, only 1 basis point of the volatility of the dynamically hedged return can be attributed to risk associated with U.S. Treasuries, compared to 14 basis points of interest rate risk in the statically hedged return.

Our analysis provides several additional insights. First, the technique is straightforward to perform for any return horizon; thus, the hedge can be tailored to portfolio rebalancing over any interval, be it a day, a week, or a month. We investigate this feature empirically and comment on the differences between the hedge ratios for weekly and monthly returns.

Second, we perform a detailed comparison between the static and dynamic empirical methods for hedging MBS with T-note futures. We show that the dynamic hedge is much more successful at reducing the risk associated with interest rate movements.

Third, to the extent that some MBS return volatility remains after the T-note futures hedge, an empirical analysis of the possible determinants of the unexplained variation is provided. The results suggest that systematic volatility does exist, but that it is not associated with movements in the term structure.

Fourth, the MBS hedge is created under the assumption that the relation between MBS returns and T-note futures returns is linear. We evaluate this assumption non-parametrically, and comment on the potential need to adjust the hedge for movements in interest rates (and thus T-note returns) of varying magnitudes.

I. DYNAMIC HEDGING OF MBS

The basic approach is to estimate a conditional hedge ratio between returns on an MBS and returns on a T-note future. The hedge ratio is conditional in the sense that we account for relevant current information. This is important for MBS because, as interest rates change, expected future prepayments change, and thus the timing of the future cash flows also changes.

In order to estimate this conditional hedge ratio, a structural model is usually required (as with model-based MBS valuation approaches). Unfortunately, this requirement involves making a number of assumptions on the underlying processes, which may or may not be reasonable.

We take a different approach toward estimating the conditional hedge ratio. Using estimates of the joint and marginal probability densities of the return series and relevant variables, we estimate the conditional hedge ratio directly.[4] The advantage of this approach is that it does not require strong, model-specific assumptions. Of course, the disadvantage is that our approach introduces estimation error.

The Multivariate Density Estimation Method

In general, multivariate density estimation (MDE) is a method for estimating the joint density of a set of variables. Given the joint and marginal densities of these variables, the corresponding conditional distributions and conditional moments, such as the mean, can be calculated. Thus, one can imagine relating the expected return on an MBS to the return on a T-note futures, conditional on relevant information available at any point in time.

Suppose we have T observations, $\mathbf{z}_1, \mathbf{z}_2, ..., \mathbf{z}_T$, where each $\mathbf{z}_t$ is an m-dimensional vector. For our application, the vector $\mathbf{z}$ might include the MBS and T-note futures return, as well as several variables describing the state of the economy.

One popular consistent measure of the joint density is the Parzen (fixed window width) density estimator. The density at any point $\mathbf{z}^*$ is estimated as the average of densities centered at the actual data points $\mathbf{z}_t$. The farther a data point is away from the estimation point, the less it contributes to the estimated density. Consequently, the estimated density is highest near high concentrations of data points and lowest when observations are sparse.

More formally:

$$\hat{f}(\mathbf{z}^*) = \frac{1}{Th^m}\sum_{t=1}^{T} K\left(\frac{\mathbf{z}^* - \mathbf{z}_t}{h}\right)$$

where $\hat{f}(\mathbf{z}^*)$ is the estimate of the probability density at $\mathbf{z}^*$; h is the window width or smoothing parameter (which helps determine how tight the kernel function is); and K (·) is called the kernel function (with the property that it integrates to unity) and is often chosen to be a density function.

The researcher has discretion in choice of kernel functions and band widths. For example, one popular class of kernel functions is the symmetric beta family, which includes the normal density, the Epanechnikov [1969] kernel, and the biweight kernel as special cases. For the purposes of this study, we use the multivariate normal density:

$$K(w) = (2\pi)^{-m/2} e^{-1/2w'w}$$

Note that this choice is not equivalent to assuming that the variables of interest are jointly normally distributed. In fact, the estimated joint density can take essentially any form. There are few, if any, relevant results about the effects of the choice of the kernel function in small samples, although Epanechnikov [1969] demonstrates that many reasonable kernel functions generate almost equivalent results in terms of asymptotic efficiency.

With respect to the band width, we use the objective measure suggested by Scott [1992]:

$$h_j = k_j \hat{\sigma}_j T^{-\frac{1}{m+4}}$$

where k_j represents an adjustment factor related to the sparseness of the data around the evaluation points, and $\hat{\sigma}_j$ is the sample standard deviation of the variable z_j.[5]

Note that the choice of band widths can be quite general, e.g., different band widths for each data point (h_t) or each dimension (h_j, j = 1, ..., m) as above. As with the choice of the kernel function, there is little formal guidance as to which band width to use for a particular application. Given this latitude, experimentation with different kernel functions and band widths may improve the performance of the MDE procedure, although we do not attempt to optimize these choices in this article.

Let $\mathbf{z}_t \equiv (R^{mbs}_{t+1}, R^{TN}_{t+1}, x_t)$, where R^{mbs}_{t+1} and R^{TN}_{t+1} are the one-period returns on the MBS and T-note futures from t to t + 1, respectively, and x_t is an (m − 2)-dimensional vector of factors known at time t. We can then obtain the conditional mean, $E[R^{mbs}_{t+1} \mid R^{TN}_{t+1}, x_t]$, i.e., the expected MBS return, given movements in the T-note return, conditional on the current economic state as described by x_t.

Specifically:

$$E[R^{mbs}_{t+1} \mid R^{TN}_{t+1}, x_t] = \int R^{mbs}_{t+1} \times \frac{f(R^{mbs}_{t+1}, R^{TN}_{t+1}, x_t)}{f_1(R^{TN}_{t+1}, x_t)} dR^{mbs}$$

$$= \frac{\sum_{i=1}^{t} R^{mbs}_{t+1-i} K_1^{t-i}(.,.)}{\sum_{i=1}^{t} K_1^{t-i}(.,.)} \qquad (1)$$

where $K_1^{t-i}(.,.) \equiv K_1[(R^{TN}_{t+1-i} - R^{TN}_{t+1})/h^{TN}, (x_{t-i} - x_t)/h]$.

K_1 (.,.) is the marginal density, $\int K(z)\, dR^{mbs}$, which is also a multivariate normal density in our application. The expected return in Equation (1) is simply a weighted average of past returns, where the weights depend on the levels of the conditioning variables relative to their levels in the past. We expand further on this intuition in the context of hedge ratios in the next section.

Given $E[R^{mbs}_{t+1} \mid R^{TN}_{t+1}, x_t]$, a hedge ratio can be formed by estimating how much the return on the MBS changes as a function of changes in the T-note futures return, conditional on currently available information x_t. That is:

$$\frac{\partial E[R^{mbs}_{t+1} \mid R^{TN}_{t+1}, x_t]}{\partial R^{TN}_{t+1}} = \frac{\sum_{i=1}^{t} R^{mbs}_{t+1-i} \frac{\partial K_1^{t-i}(.,.)}{\partial R^{TN}_{t+1}}}{\sum_{i=1}^{t} K_1^{t-i}(.,.)} - \frac{\sum_{i=1}^{t} R^{mbs}_{t+1-i} K_1^{t-i}(.,.) \sum_{i=1}^{t} \frac{\partial K_1^{t-i}(.,.)}{\partial R^{TN}_{t+1}}}{\left[\sum_{i=1}^{t} K_1^{t-i}(.,.)\right]^2} \qquad (2)$$

where

$$\frac{\partial K_1^{t-i}(.,.)}{\partial R^{TN}_{t+1}} = -[(R^{TN}_{t+1-i} - R^{TN}_{t+1})/(h^{TN})^2]\, K_1^{t-i}(.,.).$$

Several comments are in order. First, Equation (2) provides a formula for the hedge ratio between an investor's MBS position and T-note futures. For example, if $\partial E\left[R^{mbs}_{t+1-i} \mid R^{TN}_{t+1}, x_t\right] / \partial R^{TN}_{t+1}$ equals 0.5, then for every \$1 of an MBS held the investor should short \$0.50 worth of T-note futures.

Second, this hedge ratio will change dynamically, depending on the current economic state described by x_t. For example, suppose x_t is an $m - 2$ vector of term structure variables. As these variables change, whether they are the level, slope, or curvature of the term structure, the hedge ratio may change in response. Thus, the appropriate position in T-note futures will vary over time.

Third, the hedge ratio is a function of the unknown return on the T-note futures. If the conditional relation between MBS returns and T-note futures returns is always linear, then the same hedge ratio will be appropriate, irrespective of how T-note futures move. If the relation is not linear, then the investor must decide what type of T-note moves to hedge.

For example, the investor might want to form the MBS hedge in the neighborhood of the conditional mean of the T-note futures return, since many of the potential T-note futures returns will lie in that region. On the other hand, it may be the case that the investor is concerned about the tails of the distribution of T-note futures returns, and thus adjusts the hedge ratio to take account of potential extreme moves in interest rates and T-note futures.

Fourth, the hedge ratio is horizon-specific. In contrast to the instantaneous hedge ratio, our method's implied hedge ratio directly reflects the distribution of MBS returns over the relevant horizon. Thus, different hedge ratios may be appropriate for daily, weekly, or monthly horizons.

Interpretation of the Dynamic Hedge Ratio

At first glance, Equation (2) appears somewhat daunting. Its interpretation, however, is quite simple. To see this, first consider the hedging method of regressing MBS returns on T-note futures. The resulting regression coefficient, or hedge ratio, is given by

$$\beta = \frac{\sum_{i=1}^{t} R_{t+1-i}^{mbs} R_{t+1-i}^{TN} - T \mu_{mbs} \mu_{TN}}{\sum_{i=1}^{t} (R_{t+1-i}^{TN} - \mu_{TN})^2} \quad (3)$$

where

$$\frac{1}{T} \sum_{i=1}^{t} R_{t+1-i}^{mbs} = \mu_{mbs}$$

and

$$\frac{1}{T} \sum_{i=1}^{t} R_{t+1-i}^{TN} = \mu_{TN}$$

There is a clear interpretation of the hedge ratio implied by the β coefficient. That is, the hedge ratio is constructed by taking pairs of past MBS and T-note returns, and then equally weighting these pairs' comovements (in this case, by the variability of the T-note futures return).

The problem with this approach is that all observations get equal weight. Thus, in estimating the hedge ratio today, comovements between MBS and T-note returns in high interest rate environments get the same weight as in low interest rate environments. A static hedge ratio, of course, is not appropriate for hedging MBS.

The dynamic hedging strategy outlined above also has a clear interpretation. In fact, it is possible to rewrite Equation (2) as:

$$\frac{\partial E[R_{t+1}^{mbs} \mid R_{t+1}^{TN}, x_t]}{\partial R_{t+1}^{TN}} =$$

$$\sum_{i=1}^{t} R_{t+1-i}^{mbs} [(R_{t+1-i}^{TN} - R_{t+1}^{TN}) / (h^{TN})^2] \omega_i(t) -$$

$$\left[\sum_{i=1}^{t} R_{t+1-i}^{mbs} \omega_i(t)\right] \times$$

$$\left[\sum_{i=1}^{t} [(R_{t+1-i}^{TN} - R_{t+1}^{TN}) / (h^{TN})^2] \omega_i(t)\right] \quad (4)$$

where $\omega_i(t) = \dfrac{K_1^{t-i}(.,.)}{\sum_{i=1}^{t} K_1^{t-i}(.,.)}$.

Thus, the hedge ratio given in (4) is constructed by taking past pairs of MBS and T-note futures returns, and then differentially weighting these pairs' comovements by determining how "close" $(R_{t+1-i}^{TN}, x_{t-i})$ pairs are to a chosen value of R_{t+1}^{TN} and current information x_t. Equation (4) then is similar in spirit to a regression hedge, except that the weights are no longer constant, but instead depend on current information.

Of course, it is possible to adapt the regression hedge to incorporate current information by making the regression coefficient a function of variables such as the interest rate. The problem comes in choosing the specific form of the functional relation.

The nice idea behind density estimation is that these weights are not estimated in an ad hoc manner, but instead depend on the true (albeit estimated) underlying distribution of relevant variables. Thus, if the current information x_t is not close to x_{t-i} in a distributional sense, then ω_i puts little weight on the observation pair R^{mbs}_{t+1-i} R^{TN}_{t+1-i}.

Our approach has a clear advantage over the regression hedge given by (3). The hedge ratio in (4) explicitly takes into account the current economic state. For example, if interest rates are currently high, but the term structure is inverted, more weight will be given to past comovements between MBS and T-note futures in that type of interest rate environment. Thus, the hedge ratio adjusts to current economic conditions.

II. DATA DESCRIPTION

Data Sources

We use three main data sources over the period January 1987 to May 1994: 1) mortgage-backed security prices from Bloomberg Financial Markets, 2) ten-year T-note futures prices from Technical Tools Inc., and 3) various term structure information, including the three-month, one-year, five-year, and ten-year yields from the Federal Reserve Board of Governors and Bloomberg Financial Markets.

For MBS data, we collected weekly data on thirty-year fixed-rate GNMA MBS, with coupons of 8%, 9%, and 10%. The prices represent dealer-quoted bid prices on X% coupon-bearing GNMAs traded for forward delivery on a to-be-announced (TBA) basis. The TBA market is most commonly employed by mortgage originators who have a given set of mortgages that have not yet been pooled.

Trades can also involve existing pools, on an unspecified basis. This means that, at the time of the agreed-upon transaction, the characteristics of the mortgage pool to be delivered (e.g., the age of the pool, its prepayment history) are at the discretion of the dealer. Nevertheless, as long as new mortgages with the required coupon are being originated, these pools are likely to be delivered because seasoned pools are more valuable in the interest rate environment that characterizes our sample period. Consequently, GNMA TBAs are best thought of as forward contracts on generic, newly issued, securities. (We address the extent to which the absence of origination affects the results later.)

The definition of a return on a forward or futures contract is somewhat arbitrary, given that no money changes hands upfront. For the purposes of analysis, the returns on both the GNMA TBAs and the T-note futures are defined as the change in the price over the period divided by the price at the beginning of the period.[6]

For the interest rate series, weekly data on the yield on the three-month Treasury bill, the one-year Treasury bill, five-year Treasury note, and ten-year Treasury note were collected for the 1987-1994 sample period. From these series, we construct several potential interest rate factors to be used in the empirical tests. Litterman and Scheinkman [1991] show that almost all term structure movements can be explained by movements in three factors: the level of interest rates, the slope of the term structure, and the curvature. We also include a fourth factor, the volatility of interest rates, because of the option component of MBS.

Specifically, we relate unhedged and hedged GNMA returns to changes in the four factors defined as follows:

- The ten-year yield (level).
- The spread between the ten-year yield and the three-month yield (slope).
- The difference between one-half the sum of the five-year and three-month yields and the one-year yield (curvature).
- Interest rate volatility measured as an exponentially smoothed sum of past daily squared changes in the three-month rate. We consider three different smoothing parameters: 0.99, 0.96, and 0.90. The exponentially smoothed measure is often preferred to historical volatility because it takes into account more recent information, and in fact has a representation much like the popular GARCH (1, 1) process. As an aside, this volatility measure corresponds to J.P. Morgan's RiskMetrics ($\theta = 0.96$).

EXHIBIT 1
YIELD ON THE TEN-YEAR TREASURY NOTE,
WEEKLY FROM JANUARY 1987-MAY 1994

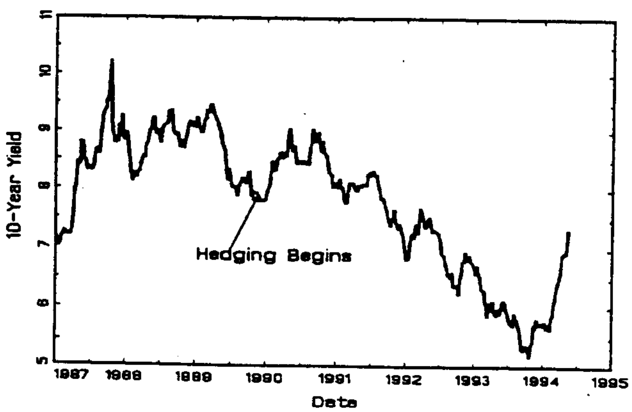

Data Analysis

Exhibit 1 graphs the ten-year yield (the conditioning variable used in our analysis) over the 1987-1994 sample period. We denote this yield as the *level of interest rates*. The figure also shows the point at which the out-of-sample analysis begins.

The figure illustrates the decline in interest rate levels throughout most of the sample, but it is also clear that there is substantial variation in the interest rate level over this period, which suggests the level may have important dynamic information for pricing, and thus hedging, MBS. Since the level of interest rates generally falls over the sample, the moneyness of the prepayment option of the MBS comes into play. This effect is analyzed below.

In our discussion of the MDE method, we describe how MBS returns can be estimated as a function of the unknown return on T-note futures. This is important, because it tells us under what conditions the T-note futures will hedge all movements in MBS returns. As an illustration of the methodology, we focus on the MBS and T-note futures return relation, conditional on the level of interest rates.

Specifically, over the entire sample period, Exhibits 2A and 2B graph the MDE-implied expected weekly return on the 10% and the 8% GNMA against contemporaneous movements in the T-note futures return, conditional on different levels of the ten-year yield. The graph shows the relation over ranges of T-note futures returns and ten-year yields that occurred during the sample period. Within these ranges, the MDE approach interpolates between points. Problems can occur when it is necessary to extrapolate beyond the range of data in the sample.

Several interesting features of the relation are apparent in the figures. First, conditioning on the current economic state does seem to matter — that is, the hedge ratio changes as a function of the level of interest rates.

For example, consider the 10% GNMA in Exhibit 2A. At a ten-year yield of 6%, the line is quite flat, while at 9% it is relatively steep. When long rates are at 6%, the prepayment option of the underlying

EXHIBITS 2A-2B
EXPECTED WEEKLY RETURN ON A 10% (TOP) AND AN 8% (BOTTOM) GNMA AS A FUNCTION OF THE CONTEMPORANEOUS TEN-YEAR T-NOTE FUTURES RETURN, CONDITIONAL ON THREE DIFFERENT LEVELS OF THE TEN-YEAR T-NOTE YIELD

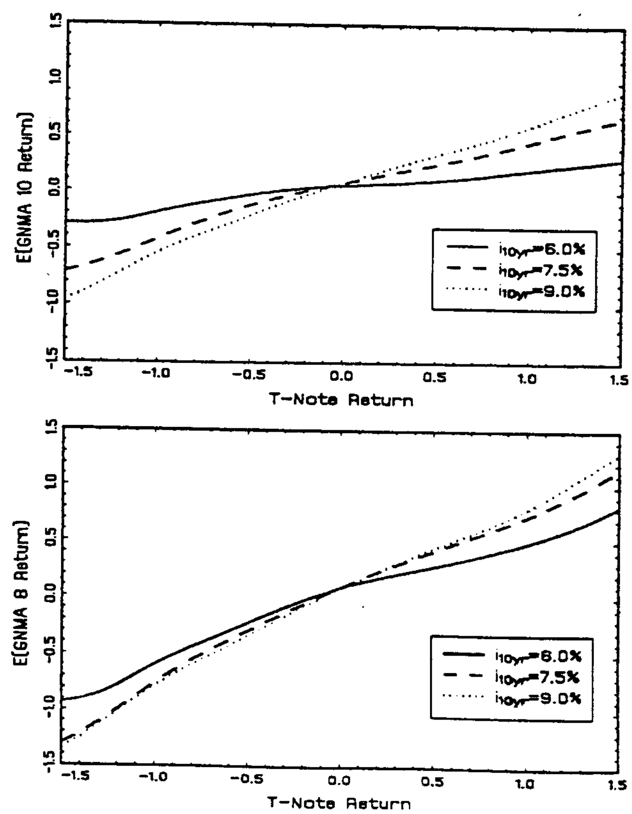

Notes: Relation estimated using MDE January 1987-May 1994. Returns are in percent per week.

mortgages of the 10% GNMA is deep in the money, and therefore MBS returns are relatively insensitive to interest rate changes.

In contrast, when rates are 9%, the prepayment option is out of the money, and MBS behave much more like the corresponding fixed-rate bonds. When interest rates are high, MBS returns move inversely with interest rate changes, and therefore more closely with interest rate-sensitive securities such as T-note futures.

This observation is even more apparent when we compare Exhibits 2A and 2B. The underlying mortgages of the 8% GNMA become in the money at much lower rates than the 10% GNMA. At each interest rate level, therefore, Exhibit 2B shows that the 8% GNMA return is more sensitive to T-note futures returns, with the slope being steeper at each point.

Of some interest, the hedge ratios for the 8% GNMA are almost identical at interest rate levels of 9.0% and 7.5%. At these levels, small movements in interest rates (as described by the −1.5% to 1.5% weekly return on T-note futures) have similar effects on the MBS return because the prepayment options of the 8.5% mortgages are out of the money. Note that X% GNMAs are backed by X + 1/2% underlying mortgages due to the servicing fee associated with MBS.

Second, for weekly T-note futures returns in the range of −1.5% to 1.5%, the relation between conditional expected returns on GNMA TBAs and T-note futures is, for the most part, linear. Thus, the hedge ratio will not necessarily have to be adjusted for non-linearity using T-note options or a similar interest rate derivative. This said, however, Exhibits 2A and 2B do display a small amount of concavity.

For example, consider the 10% GNMA at an interest rate level of 9%. For a T-note futures return between −1.5% and 0%, the MBS return increases by 1%; in contrast, between 0% and 1.5%, the increase is only 0.8%. From a theoretical viewpoint, a small amount of concavity is to be expected. As interest rates fall, expected returns on MBS become less sensitive to interest rate changes via the prepayment option. Exhibits 2A and 2B show that, for weekly MBS returns, this is not a first-order effect over the range of reasonable values of T-note futures returns.

III. EMPIRICAL ANALYSIS

Because the relation between GNMA returns and T-note futures returns is approximately linear, the hedging analysis focuses on hedge ratios at the conditional mean of T-note futures returns, conditional on the level of interest rates. To the extent there is any non-linearity, the hedge ratio is still appropriate locally for realizations of T-note futures returns in the neighborhood of the mean of its conditional distribution. Since the level of interest rates is the predominant factor underlying both MBS and T-notes, this variable provides a good starting point for a dynamic hedging analysis.

The strategy is to compare the unhedged GNMA returns, the static linear hedged returns [Equation (3)], and the dynamic MDE hedged returns [Equation (4)], where these returns are calculated as follows:

1. The unhedged GNMA returns are calculated for horizons of one and four weeks over the sample period December 1989 through May 1994.
2. The static linear hedged return from t to t + 1 is given by $R^{mbs}_{t+1} - \beta_t R^{TN}_{t+1}$, where β_t is the coefficient from regressing the MBS return on the contemporaneous T-note futures return over the past 150 weeks.[7] β_t is re-estimated every period via a rolling regression, and the monthly horizon hedge ratios are estimated using overlapping observations.
3. The dynamic MDE hedged return takes the same form, except β_t is now estimated using MDE in a rolling estimation over the past 150 weeks of data. In this analysis, we condition on the level of the ten-year yield; thus, β_t takes a measure of the past observations of R^{mbs} and R^{TN}, non-linearly weighted by how close past interest rates are to the current level.

The Hedge Ratio

Exhibits 3A-3D provide the hedge ratio, β_t, for both the linear and MDE dynamic hedges for weekly returns on GNMA 10s, weekly returns on GNMA 8s, monthly returns on GNMA 10s, and monthly returns on GNMA 8s.

Over the 1990 to 1994 sample period, there is

EXHIBITS 3A-3B
HEDGE RATIOS FOR HEDGING WEEKLY 10% (TOP) AND 8% (BOTTOM) GNMA RETURNS USING THE TEN-YEAR T-NOTE FUTURES

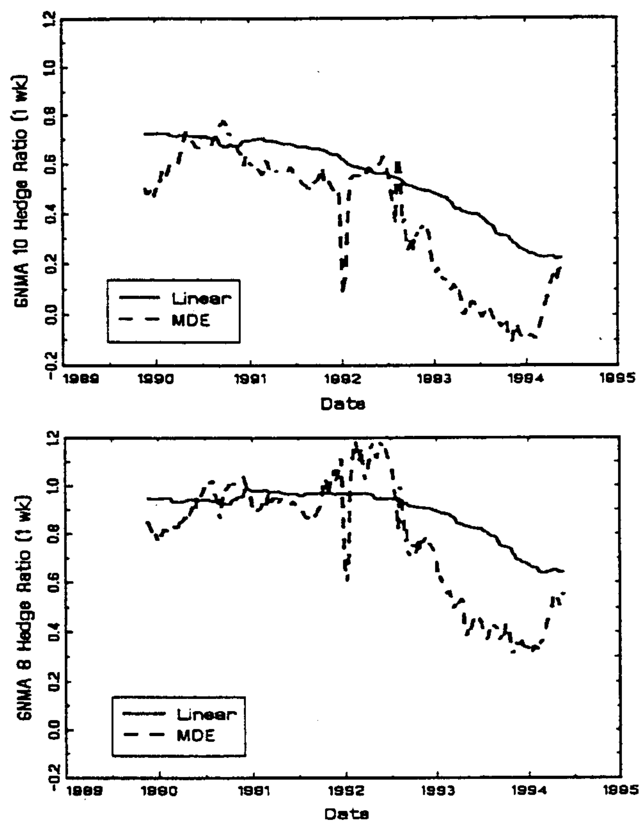

Notes: Hedge ratios estimated on a 150-week rolling basis using both a linear regression and MDE. The MDE hedge ratios condition on the level of the ten-year T-note yield.

a general decline in the hedge ratio, irrespective of the method or horizon length. This is to be expected. Interest rates are declining during this period, and the MBS prepayment option is becoming more in the money. As a result, MBS are less sensitive to interest rate movements — hence, the investor needs to hedge with smaller amounts of T-note futures.

At first glance, it may seem strange that the hedge ratio of the static linear method actually changes through time. Note, however, that the linear regressions are performed on a rolling basis. Thus, as observations enter and leave the estimation window, the estimated hedge ratio may change. Because interest rates are declining for most of the period, the hedge ratio also declines.

A comparison between the hedge ratios for the 10% GNMA (Exhibits 3A and 3C) and that of the 8% GNMA (Exhibits 3B and 3D) is especially revealing. While the overall pattern of the hedge ratios over time is similar, the magnitude of these ratios is quite different. In fact, the GNMA 8's hedge ratio is uniformly higher than the GNMA 10's.

For example, consider the hedge ratio for weekly GNMA returns using the dynamic MDE method. In January 1990, every \$1 at stake in an investment in a GNMA 8 TBA (GNMA 10 TBA) requires an \$0.85 (\$0.50) offsetting position in T-note futures, while in May 1994, a \$0.50 (\$0.20) offsetting

EXHIBITS 3C-3D
HEDGE RATIOS FOR HEDGING MONTHLY 10% (TOP) AND 8% (BOTTOM) GNMA RETURNS USING THE TEN-YEAR T-NOTE FUTURES

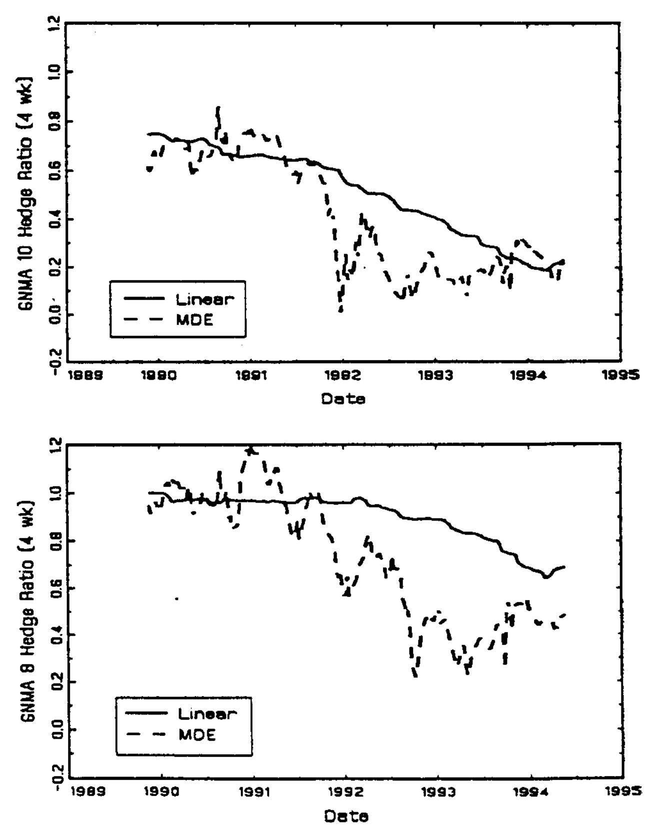

Notes: Hedge ratios estimated on a 150-week rolling basis using both a linear regression and MDE. The MDE hedge ratios condition on the level of the ten-year T-note yield.

position is required. The fact that the GNMA 8s require a larger position in T-note futures reflects their greater sensitivity to interest rates, not just the moneyness of the prepayment option.

For much of the sample, the dynamic MDE hedge requires a smaller position in T-note futures than does the linear hedge. This again should not be surprising. The MDE method conditions on current information, so, as interest rates fall, this fall is impounded immediately into the weights placed on prior pairs of MBS and T-note futures returns [Equation (4)]. The rolling linear regression method, on the other hand, introduces this information slowly, and with equal weight, as the estimation window changes.

The speed at which information gets incorporated also explains why the MDE dynamic hedge ratios are much more variable than the linear static hedge ratio. Since the MDE-based hedge ratio adjusts immediately to information, the variability of the interest rate level (see Exhibit 1) induces variation in the dynamic hedge ratios.

These points are very apparent over the last seventy weeks of the sample (January 1993-May 1994). It is a period in which investors suffered large losses in hedged portfolios due to overestimating the underlying interest rate sensitivity of MBS. At least for the GNMA TBAs studied here, the dynamic method has made an appropriate adjustment during this latter period.

One striking feature of the weekly MDE hedge ratios is the spike at the beginning of 1992. Hedge ratios decrease in magnitude by about 0.5, and then rebound to their original levels within a couple of weeks. This large movement and rapid reversal is a result of the contemporaneous spike in interest rates (see Exhibit 1).

When interest rates fall rapidly to a level that has not been observed in the previous 150 weeks, the MDE procedure has problems extrapolating from the past, and somewhat irrelevant, data. As soon as interest rates rebound, the MDE procedure returns to familiar territory. If interest rates had remained low, the procedure would also have adapted quickly.

Finally, note that the MDE dynamic hedge positions change depending on whether the position is held for one week or one month. For example, consider the GNMA 10 TBA. While the hedge positions for weekly and monthly returns generally decline over the entire sample, the hedge ratio for monthly returns is much more flat than for weekly returns during the 1993-1994 period.

There are two possible explanations for this difference. Even though the full 150 weeks of data are used to form monthly returns, these returns are overlapping and thus include less information than independent observations. Thus, the monthly hedge ratio may be less reliable than the weekly hedge ratio, and the divergence may reflect estimation error.

Alternatively, the hedge ratio may reflect the possibility that the level of interest rates implies different conditional changes in future interest rates over different horizons. Since it is commonly believed that interest rates are mean-reverting, this explanation has some merit. It is an empirical question, however, whether this mean reversion implies enough differences between the conditional distribution of interest changes at weekly and monthly horizons.

Analysis of Hedged MBS Returns

Exhibit 4 provides the mean, volatility, and autocorrelation of the unhedged, linearly hedged, and MDE dynamically hedged 8%, 9%, and 10% GNMA TBA returns for both one- and four-week horizons. Consider the mean returns of these hedged and unhedged positions for one-week horizons. The weekly mean returns for the unhedged GNMAs are 0.078%, 0.077%, and 0.069%, respectively, for the 8s, 9s, and 10s, which translates to between 3.6% and 4.1% on an annual basis. Because TBAs are forward contracts, this mean return represents a risk premium for holding MBS.

Of course, we should note that this period is one in which interest rates generally fell, so these mean estimates may just reflect the particular sample period, rather than an ex ante risk premium per se. Nevertheless, since much of this premium may be compensation for bearing interest rate risk, it should not be surprising that the mean hedged returns drop to 0.006%, 0.020%, and 0.027% for the 8s, 9s, and 10s, respectively (using MDE). This fall represents a decline in overall interest rate exposure due to the hedge itself.

The most interesting results from Exhibit 4 are the standard deviations of the unhedged and hedged returns. First, consider the weekly return horizons for

EXHIBIT 4
COMPARISON OF HEDGING METHODS

	GNMA 8		GNMA 9		GNMA 10	
	1 wk.	4 wks.	1 wk.	4 wks.	1 wk.	4 wks.
Unhedged:						
Mean (%)	0.078	0.326	0.077	0.316	0.069	0.286
Volatility (%)	0.685	1.364	0.531	0.999	0.414	0.746
Autocorrelation	0.011	0.138	−0.019	0.109	−0.045	0.082
Linear Hedge:						
Mean (%)	0.007	0.041	0.017	0.086	0.024	0.126
Volatility (%)	0.318	0.599	0.309	0.573	0.286	0.524
Autocorrelation	0.015	−0.107	0.012	0.021	0.060	0.039
MDE Hedge:						
Mean (%)	0.006	0.123	0.020	0.178	0.027	0.189
Volatility (%)	0.285	0.583	0.245	0.472	0.242	0.440
Autocorrelation	0.011	−0.096	0.016	−0.083	0.036	−0.085

Notes: Mean, volatility, and autocorrelation of unhedged returns on GNMA TBAs and hedged returns compared using two different approaches. The approaches involve hedging GNMAs with T-note futures, resulting in the hedged return, $R_{t+1}^{mbs} - \beta_t R_{t+1}^{TN}$, where R_{t+1}^{mbs} and R_{t+1}^{TN} are the out-of-sample returns on GNMAs and T-note futures, respectively, and β_t, the hedge ratio, is estimated using the prior 150 weeks of data in one of two ways: 1) a linear hedge based on a regression of past R_{t+1}^{mbs} on R_{t+1}^{TN}, and 2) an MDE hedge using the distribution of R_{t+1}^{mbs} and R_{t+1}^{TN}, conditional on the ten-year yield at time t. The estimation is performed on a rolling basis and covers the out-of-sample period, December 1989 to May 1994. Results are reported for both weekly and overlapping monthly returns.

the GNMA TBAs. The unhedged return volatility is 68.5, 53.1, and 41.4 basis points, respectively, for the 8s, 9s, and 10s. The decline in the overall volatility for higher-coupon GNMAs just reflects the effect of the prepayment option vis-à-vis interest rate movements.

Linearly hedging MBS using T-note futures in the static regression substantially reduces this volatility risk to 31.8, 30.9, and 28.6 basis points, respectively. Of particular interest, our dynamic method fares even better — the volatility of the hedged GNMAs is 28.5, 24.5, and 24.2 basis points, a 10%–20% drop relative to the static approach. Under the MDE method, the largest gains are with the GNMA 9s and 10s, since the prepayment option is much more relevant for these securities, given the level of interest rates over the sample period.

This same pattern in results carries through to the four-week horizons. While the unhedged 8%, 9%, and 10% GNMA returns face 136.4, 99.9, and 74.6 basis points of volatility, respectively, the MDE (linear) hedge reduces this risk to 58.3 (59.9), 47.2 (57.3), and 44.0 (52.4) basis points.

To cast this in concrete terms, suppose an institution with a large GNMA portfolio is faced with the possibility of a large negative realization (say, 2 standard deviations) on its GNMA position. This means that, for a $100 million one-month exposure to MBS, the MDE hedge reduces the value at risk from $2.38 million to only $1.04 million for the 8% GNMA.

Exhibit 4 also gives the autocorrelation of the unhedged and hedged GNMA returns. Even if GNMA and T-note futures returns are not serially correlated (which is true for this sample), this does not imply that the hedged return will be serially uncorrelated. The reason is that the estimates of the hedge ratio, β_t, may be correlated through time either due to estimation error (via the rolling regression or MDE) or due to the conditioning variable (i.e., the interest rate level).

Exhibit 4 shows, however, that the serial dependency of β_t does not necessarily lead to serial correlation in the hedged position's returns. For example, for weekly returns, all the autocorrelations are below 4%. With respect to monthly returns, the MDE hedged positions pick up minor autocorrelation — −9.6%, −8.3%, and −8.5% for the 8s, 9s, and 10s, respectively. The magnitudes are consistent with the hypothesis that these autocorrelations are not significantly different from zero.

While reducing the total volatility of the position in GNMA TBAs is an important goal, it is only suggestive of how much interest rate volatility is removed. It is of interest to find out how much of the variation of the returns on the GNMAs is due to variation of T-note futures returns. Exhibit 5 provides such information.

For example, consider the weekly return on the GNMA 8 (GNMA 10). Exhibit 5 shows that 77.0% (50.5%) of the variation of the unhedged return on GNMA 8s (GNMA 10s) can be explained by variation in the contemporaneous T-note futures return. This translates to 60 (30) basis points of volatility due to T-note futures returns. This is in contrast to 10.3% (23.2%) explained variation in the static linear hedged return, which is equivalent to 10 (14) basis points of remaining volatility due to changes in T-note futures prices.

The MDE dynamic hedge fares much better. Only 0.2% (0.3%) of the remaining variation of the hedged return can be related to the T-note futures return — that is, only 1 basis point of T-note risk remains in either case. If the goal of the hedging exercise is to remove risk associated with U.S. Treasuries, the MDE method is a clear winner.[8]

As a final comment on the comparison

EXHIBIT 5
THE SENSITIVITY OF UNHEDGED/HEDGED GNMA RETURNS TO T-NOTE FUTURES RETURNS

	GNMA 8		GNMA 9		GNMA 10	
	1 wk.	4 wks.	1 wk.	4 wks.	1 wk.	4 wks.
Unhedged:						
Correlation	0.878	0.899	0.792	0.809	0.711	0.729
R^2	0.770	0.808	0.627	0.655	0.505	0.532
Volatility (basis pts)	60.2	122.6	42.0	80.8	29.5	54.4
Linear Hedge:						
Correlation	−0.321	−0.367	−0.507	−0.552	−0.482	−0.509
R^2	0.103	0.134	0.257	0.304	0.232	0.259
Volatility (basis pts)	10.2	21.9	15.7	31.6	13.8	26.6
MDE Hedge:						
Correlation	−0.043	0.198	−0.107	−0.046	−0.056	−0.158
R^2	0.002	0.039	0.011	0.002	0.003	0.025
Volatility (basis points)	1.2	11.5	2.6	2.2	1.3	7.0

Notes: Correlation of returns on hedged and unhedged GNMAs with contemporaneous T-note futures returns, the R^2 from the corresponding regression, and the magnitude of the variation of unhedged/hedged GNMA returns explained by T-note futures returns (measured by its volatility in basis points). GNMAs are hedged with T-note futures, resulting in the hedged return, $R_{t+1}^{mbs} - \beta_t R_{t+1}^{TN}$, where R_{t+1}^{mbs} and R_{t+1}^{TN} are the out-of-sample returns on GNMAs and T-note futures, respectively, and β_t, the hedge ratio, is estimated using the prior 150 weeks of data in one of two ways: 1) a linear hedge based on a regression of past R_{t+1}^{mbs} on R_{t+1}^{TN}, and 2) an MDE-based hedge using the distribution of R_{t+1}^{mbs} and R_{t+1}^{TN}, conditional on the ten-year yield at time t. The estimation is performed on a rolling basis and covers the out-of-sample period, December 1989 to May 1994. Results are reported for both weekly and overlapping monthly returns.

EXHIBIT 6

CONTEMPORANEOUS CORRELATIONS OF UNHEDGED/HEDGED GNMA RETURNS — WEEKLY (UPPER TRIANGLE)/MONTHLY (LOWER TRIANGLE)

		Unhedged GNMA			MDE Hedged GNMA		
		8%	9%	10%	8%	9%	10%
Unhedged	GNMA 8%	1.000	0.951	0.876	0.341	0.235	0.149
	GNMA 9%	0.954	1.000	0.952	0.317	0.372	0.241
	GNMA 10%	0.875	0.951	1.000	0.343	0.447	0.438
MDE Hedged	GNMA 8%	0.494	0.423	0.382	1.000	0.824	0.659
	GNMA 9%	0.258	0.330	0.343	0.840	1.000	0.836
	GNMA 10%	0.090	0.177	0.335	0.640	0.844	1.000

Notes: Contemporaneous correlation between MDE dynamically hedged GNMA returns and unhedged GNMA returns. The GNMAs are hedged with T-note futures, resulting in the hedged return, $R_{t+1}^{mbs} - \beta_t R_{t+1}^{TN}$, where R_{t+1}^{mbs} and R_{t+1}^{TN} are the out-of-sample returns on GNMAs and T-note futures, respectively, and β_t, the hedge ratio, is estimated using the prior 150 weeks of data. This MDE-based estimate uses the distribution of R_{t+1}^{mbs} and R_{t+1}^{TN}, conditional on the ten-year yield at time t. The estimation is performed on a rolling basis and covers the out-of-sample period, December 1989 to May 1994. Results are reported for both weekly returns (upper triangular matrix) and overlapping monthly returns (lower triangular matrix).

between the static linear and dynamic MDE hedges, note that the correlations of the linear hedge with the weekly T-note futures return are negative (−32.1%, −50.7%, and −48.2% for GNMA 8s, 9s, and 10s, respectively). This means that the static method continually overhedges the GNMAs with T-note futures.

This is because the linear method is slow to adjust to the decline in interest rates over the sample period. As a result, the hedge is mistakenly set up based on data in relatively high interest rate periods, during which the prepayment option is out of the money and MBS are more sensitive to interest rate movements. Because the dynamic MDE hedge conditions on the current economic state, it does not suffer from this problem (recall also Exhibits 3A-3D).

If the MDE dynamic method for hedging MBS succeeds in removing the primary source of interest rate risk, what is the source of the remaining volatility? As a first pass, Exhibit 6 reports the contemporaneous correlations of the dynamically hedged and unhedged GNMA returns across coupons for both weekly and monthly returns. The contemporaneous correlations across coupons drop as the interest rate risk is hedged away, but substantial cross-correlation still remains. For example, the correlation between unhedged weekly returns on GNMA 8s and 10s is 87.6%, while the correlation between their hedged counterparts is 65.9%.

Is this remaining 66% correlation all due to hedging error, or is there some systematic effect across GNMA TBAs? To check this, we also calculate the correlation between the unhedged and hedged GNMA returns. While the correlations are much smaller, they are not zero, ranging instead from 15% to 50%.

As an example, consider weekly returns on GNMA 8s and 10s. 2.2% (11.8%) of the variation of hedged 8% GNMAs (10% GNMAs) can be explained by unhedged returns on 10% GNMAs (8% GNMAs). This is equivalent to 4.2 (8.3) basis points of volatility risk due to variation of MBS returns, not related to Treasury movements. This result suggests that measurement error cannot account entirely for the remaining volatility.

Sources of Variation of GNMA TBA Returns

While our results so far demonstrate that the MDE dynamic hedge removes the risk associated with T-note returns, the question remains whether this hedge is sufficient to remove all types of fundamental interest rate risk. Exhibits 7A (weekly) and 7B (monthly) provide the R^2 and corresponding basis points risk of unhedged, linearly hedged, and MDE dynamically hedged GNMA returns in relation to changes in the levels of various interest rate factors. These factors broadly span four different types of information: 1) the level of interest rates, 2) the slope of the term structure, 3) the curvature of the term structure, and 4) interest rate volatility.

Several observations are particularly interesting. Both the unhedged and the linearly hedged returns face considerable interest rate risk, which is consistent with the documented ex post correlation with T-note futures returns (see also Exhibit 5). This is in stark contrast to the MDE hedged returns, which show little correlation with changes in any of the interest rate factors. That is, even though only T-note futures are used to hedge MBS returns, the dynamic hedge reduces volatility risk associated with the level, slope, and curvature of the term structure.

For example, consider the weekly return on the 10% GNMA. The MDE dynamic hedge leaves 1.3, 0.3, and 0.4 basis points of volatility due to these three factors, respectively. In contrast, the unhedged (linearly hedged) returns face 28.6 (13.5), 18.5 (9.2), and 3.3 (2.5) basis points of volatility. To the extent that the volatility of the unhedged return is only 41.4 basis points, it is clear how successful the dynamic hedge is.

Similar results carry through to monthly returns (Exhibit 7B). In particular, with respect to interest rate risk, the dynamic hedge is far superior to the static hedging method, although there are two main differences relative to weekly returns.

First, the GNMA returns hedged using the dynamic MDE hedge leave some residual risk that is correlated with interest rate changes. We cannot distinguish whether this is due to estimation error in the hedge ratios (see Exhibits 3B and 3D versus 3A and 3C), or whether it just reflects the difficulty in hedging longer horizon returns using only one instrument (T-note futures).

Second, the curvature of the term structure now plays a larger role. For example, of the MDE hedged 8%, 9%, and 10% GNMA return's volatility of 58.3, 47.2, and 44.0 basis points, respectively, changes in curvature from month to month explain 11.6, 13.0, and 11.1 basis points.

Exhibits 7A and 7B also show that, regardless

EXHIBIT 7A

THE SENSITIVITY OF UNHEDGED/HEDGED GNMA WEEKLY RETURNS TO FACTOR CHANGES

	GNMA 8			GNMA 9			GNMA 10		
	Unhedged	Linear	MDE	Unhedged	Linear	MDE	Unhedged	Linear	MDE
Level:									
Correlation	−0.863	0.301	0.023	−0.775	0.489	0.092	−0.691	0.472	0.053
R^2	0.745	0.090	0.001	0.601	0.240	0.009	0.478	0.223	0.003
Volatility (basis pts)	59.1	9.6	0.6	41.1	15.1	2.3	28.6	13.5	1.3
Slope:									
Correlation	−0.546	0.231	−0.001	−0.487	0.357	0.048	−0.447	0.320	0.012
R^2	0.298	0.053	0.000	0.237	0.127	0.002	0.200	0.103	0.000
Volatility (basis pts)	37.4	7.3	0.0	25.8	11.0	1.2	18.5	9.2	0.3
Curvature:									
Correlation	−0.048	−0.100	−0.023	−0.086	−0.129	−0.080	−0.080	−0.087	−0.016
R^2	0.002	0.010	0.001	0.007	0.017	0.006	0.006	0.008	0.000
Volatility (basis pts)	3.3	3.2	0.7	4.6	4.0	2.0	3.3	2.5	0.4
Volatility:									
Correlation	0.016	−0.028	−0.039	0.022	−0.016	−0.012	0.033	0.007	0.029
R^2	0.000	0.001	0.002	0.001	0.000	0.000	0.001	0.000	0.001
Volatility (basis pts)	1.1	0.9	1.1	1.1	0.5	0.3	1.3	0.2	0.7

EXHIBIT 7B
The Sensitivity of Unhedged/Hedged GNMA Monthly Returns to Factor Changes

	GNMA 8			GNMA 9			GNMA 10		
	Unhedged	Linear	MDE	Unhedged	Linear	MDE	Unhedged	Linear	MDE
Level:									
Correlation	−0.888	0.348	−0.213	−0.799	0.523	0.032	−0.719	0.488	0.150
R^2	0.789	0.121	0.046	0.638	0.280	0.001	0.518	0.238	0.023
Volatility (basis pts)	121.1	20.8	12.4	79.8	30.3	1.5	53.6	25.6	6.6
Slope:									
Correlation	−0.534	0.319	−0.053	−0.437	0.488	0.182	−0.386	0.436	0.252
R^2	0.285	0.102	0.003	0.191	0.238	0.033	0.149	0.190	0.064
Volatility (basis pts)	72.8	19.1	3.1	43.7	28.0	8.6	28.8	22.8	11.1
Curvature:									
Correlation	−0.209	−0.234	−0.199	−0.290	−0.288	−0.275	−0.319	−0.250	−0.253
R^2	0.044	0.055	0.039	0.084	0.083	0.075	0.102	0.063	0.064
Volatility (basis pts)	28.4	14.0	11.6	29.0	16.5	13.0	23.8	13.1	11.1
Volatility:									
Correlation	−0.002	0.055	−0.012	0.005	0.032	0.002	−0.020	−0.024	0.011
R^2	0.000	0.003	0.000	0.000	0.001	0.000	0.000	0.001	0.000
Volatility (basis pts)	0.3	3.3	0.7	0.5	1.8	0.1	1.5	1.3	0.5

Notes: Correlation of returns on hedged and unhedged GNMAs with contemporaneous changes in various factors, the R^2 from the corresponding regression, and the magnitude of the variation of unhedged/hedged GNMA returns explained by these factor changes (measured by volatility in basis points). The factors include contemporaneous changes in the following interest rate variables: 1) ten-year yield (i.e., level); 2) spread between ten-year and three-month yields (i.e., slope); 3) one-half the sum of the five-year and three-month yields minus the one-year yield (i.e., curvature); and 4) interest rate volatility estimate based on an exponentially θ smoothed sum of past squared changes ($\theta = 0.96$).

of the volatility measure used, contemporaneous changes in the volatility of interest rates have very little relation to either the unhedged or hedged GNMA returns series. Because the results are similar for all different values of the exponential smoothing parameter, Exhibits 7A and 7B report the volatility using only one measure, $\theta = 0.96$. While the lack of correlation could be due to estimation error with respect to our estimate of volatility, it should be noted that the measures used are quite standard in the literature [e.g., historical volatility and GARCH (1, 1)].

The conclusions we reach from Exhibits 7A and 7B are twofold. First, the dynamic hedge removes most, if not all, of the different types of interest rate risk inherent in MBS. Of special interest, this hedge is performed using just one instrument (T-note futures) and conditions on just one economic variable (the level of interest rates).

Second, this result means that the volatility of this hedged return is not related to interest rates. It also cannot be completely explained by errors due to either estimation or measurement of GNMA prices. This is because the hedged GNMA returns are correlated with unhedged GNMA returns (which are not estimated) and across different coupons (which makes systematic price errors less likely).

One possibility for explaining these correlations across GNMA returns is that interest rate changes affect these returns non-linearly. Recall from Exhibits 2A and 2B that, due to the prepayment option, the relation between expected GNMA returns is concave (albeit slightly) in T-note returns. Hence, the T-note futures return may not be able to hedge away all MBS risk associated with large movements in interest rates.

Since the correlations documented in Exhibits 5 and 7 may not detect this, it seems worthwhile exploring whether the dynamically hedged returns are related non-linearly to interest rate movements. Given the concavity evident in Exhibits 2A and 2B, squared interest rate movements are a potentially good instrument for detecting this non-linear relation, if it exists.

The correlations of squared interest rate changes with the various GNMAs range from −10% to −20%, which is consistent with a concave relation. This amounts to approximately 4 to 5 basis points of non-linear interest rate risk, not accounted for by the T-note futures hedge. One possible way to avoid this risk may be to add another instrument that will allow us to capture these non-linear movements, e.g., options on T-note futures.

Even if this non-linear interest rate risk is accounted for, the question remains: Is the remaining volatility specific to particular GNMA securities (and thus more in the realm of estimation and GNMA price errors), or is it still systematic across all GNMAs?

If the remaining volatility is systematic, one possible explanation is that GNMA returns are nonstationary due to pool seasoning, and that this nonstationarity induces correlated risk across coupons. To check this, we investigate the periods in which 8%, 9%, and 10% GNMAs are originated. In these periods, given the declining interest rate environment, the newly originated GNMA pools are the most likely pools to be delivered in the TBA market. In contrast, when there are no originations, seasoned pools must be delivered.

A cursory look at the behavior of the hedged GNMA returns vis-à-vis origination does not explain the remaining volatility. For example, from January 1990 to November 1992, a substantial number of 9% GNMAs were originated, while very few originations of GNMA 8s took place. The opposite situation occurred in the latter part of our sample (December 1992-May 1994).

If the remaining volatility is due to the seasoning of the mortgage pools, we would expect the correlation patterns of the 8% and 9% GNMAs to be different. Although there is a general decline in the magnitude of the correlations with unhedged GNMA returns over the sample, the correlation patterns of the 8% and 9% GNMAs are very similar.

Of course, the residual risk of a given GNMA could be hedged using other GNMAs. We investigate this issue by hedging GNMAs using GNMAs of nearby coupons (i.e., 10s with 9 1/2s, 9s with 8 1/2s, and 8s with 7 1/2s). As expected, the volatility of the hedged returns is considerably lower than that using T-note futures.

For example, for weekly returns on the 10% GNMA, the volatility is 11.2 basis points. Of more interest, however, is the fact that this hedged return has a correlation of −0.327 with interest rate changes, leaving 3.7 basis points of residual interest rate risk (compared with only 1.3 for the T-note futures hedge).

Moreover, there are a number of reasons why hedging GNMAs with other GNMAs may not be practical. First, the transaction costs are higher. While the TBA market is quite liquid (e.g., spreads of 1/32nd to 4/32nd), T-note futures often exhibit lower spreads (e.g., between 1/64th and 1/32nd). Second, this type of hedge may not be sensible if the institution's core business leads to positive holdings of GNMAs of all coupons, or if an institution is trading on the price of MBS relative to U.S. Treasuries.

If the institution/investor does not use other MBS to hedge GNMAs, are there other securities that can serve this function? Without addressing this question specifically, it should be clear that the answer depends on the nature of the remaining volatility risk of GNMAs. If that risk is systematic across the economy (although not related to interest rates), then it is probably the case that some other security can hedge away this risk.

Yet it remains an open question as to 1) what this risk is, and 2) what securities can effectively hedge this risk. If the risk is systematic across only MBS and thus idiosyncratic with respect to other securities, then the risk may not be relevant. That is, it can be diversified away within a portfolio of fixed-income instruments and other assets.

IV. CONCLUDING REMARKS

Our new strategy for dynamically hedging mortgage-backed securities involves estimating the joint distribution between returns on MBS and T-note futures, conditional on current economic conditions. We show that our approach has a simple intuitive interpretation of forming a hedge ratio by differentially weighting past pairs of MBS and T-note futures returns. The weights are determined by how close, in a distributional sense, current variables (like the level of interest rates) are to past values of these variables.

As an application, an out-of-sample hedging exercise is performed for 8%, 9%, and 10% GNMAs over the 1990-1994 period for weekly and monthly return horizons. The dynamic approach is very successful at hedging the interest rate risk inherent in all the GNMAs.

For example, in hedging weekly returns on 8%, 9%, and 10% GNMAs, our dynamic method reduces the risk of the GNMA returns associated with interest rate movements to only 1.2, 2.6, and 1.3 basis points, respectively. This is in contrast to 60.2, 42.0, and 29.5 basis points of interest rate risk if no hedge had been performed.

Clearly, our approach is general enough to be applied to a number of traded derivatives, not only MBS. In fact, the only obstacle in practical application is the need for past data on the derivative. Furthermore, the uses of MDE extend to the estimation of inputs for theoretical models. Therefore, when illiquid derivatives need to be priced and/or hedged, MDE can be used in conjunction with a theoretical model. This point is discussed further in Boudoukh, Richardson, and Whitelaw [1995] in the context of term structure-related inputs.

ENDNOTES

The authors thank The Q Group for financial support and participants at the 1995 Utah Winter Finance Conference and the NBER Universities Research Conference for helpful comments.

[1]For a discussion of some of the problems associated with static hedges, see, for example, Breeden [1991] and Breeden and Giarla [1992]. With respect to linear regression hedges in particular, Batlin [1987] discusses the effect of the prepayment option on the hedge ratio between MBS and T-note futures.

[2]Davidson and Herkowitz [1992] provide an analysis of the various theoretical methodologies for valuing MBS in practice. The advantages and disadvantages of each approach are discussed in detail. With respect to the particular issue of hedging the interest rate risk of MBS, Roberts [1987] focuses primarily on model-based approaches.

[3]Chapter 15 in Hull [1993] provides a comparison and discussion of the major term structure models, with an emphasis on their underlying assumptions.

[4]Elsewhere we develop a non-linear, multifactor pricing model for MBS using density estimation techniques (Boudoukh et al. [1995]). That model can also be used to hedge MBS. In particular, given the MBS price as a function of K factors, K + 1 securities (including the MBS) can be used to hedge away all the relevant risks. Only an instantaneously riskless position, however, can be provided within that framework. The research ignores some of the horizon-specific issues that arise in this article.

[5]One problem with multivariate density estimation is that the fixed window width density estimator produces 1) spurious peaks at points where the observations are sparse, and 2) not enough resolution where the observations are dense. As we discuss later, much of the analysis of the relation between MBS and T-note futures returns is conducted at the mean T-note futures return and the level of the ten-year yield. Since many of the past return observations are centered around the mean, we let $k_{TN} = 0.5$ to produce greater resolution. In contrast, as the ten-year yield tends to lie in sparse regions relative to historical yields (due to the fall in interest rates over the sample), we let $k_{y10} = 2$ in order to avoid spurious peaks. Note that these values are chosen ex ante, and are not necessarily the optimal ones, given the data. Nevertheless, the choice should help alleviate some, if not most, of the spurious estimation problems.

[6]Return series are calculated using the GNMA TBA price series, with adjustments made once a month when the prices of the different contracts are spliced together. Note that the adjustment is made during the splice week using a version of the cost of carry model, modified for prepayments, known as the dollar roll break-even method (see Askin and Meyer [1986]).

[7]The length of the sample period for the rolling linear regression is set at 150 weeks to coincide with the MDE sample length. Reducing the sample period increases the responsiveness of the linear hedge ratio to recent information at the cost of increased estimation error. No effort is made to find the optimal sample size for either the linear hedge or the MDE hedge.

[8]Exhibit 5 shows that similar conclusions are reached for hedges using monthly returns. For the MDE dynamic hedge, there is a little residual risk associated with T-note futures returns, but it is substantially less than both the unhedged and static linearly hedged GNMA returns.

REFERENCES

Askin, D.J., and S.D. Meyer. "Dollar Rolls: A Low-Cost Financing Technique." Mortgage-Backed Securities Research, Drexel Burnham Lambert, 1986.

Batlin, C.A. "Hedging Mortgage-Backed Securities with Treasury Bond Futures." *Journal of Futures Markets*, 7 (1987), pp. 675-693.

Boudoukh, J., M. Richardson, R. Stanton, and R.F. Whitelaw. "Pricing Mortgage-Backed Securities in a Multifactor Interest Rate Environment: A Multivariate Density Estimation Approach." Working Paper, New York University, 1995.

Boudoukh, J., M. Richardson, and R.F. Whitelaw. "The Stochastic Behavior of Interest Rates." Working Paper, New York University, 1995.

Breeden, D.T. "Risk, Return, and Hedging of Fixed Rate Mortgages." *Journal of Fixed Income*, 1 (1991), pp. 85-107.

Breeden, D.T., and M.J. Giarla. "Hedging Interest Rate Risk with Futures, Swaps, and Options." In F. Fabozzi, ed., *The Handbook of Mortgage-Backed Securities*. Chicago: Probus Publishing, 1992, 3rd edition, pp. 847-960.

Davidson, A.S., and M.D. Herkowitz. "A Comparison of Methods for Analyzing Mortgage-Backed Securities." In F. Fabozzi, ed., *The Handbook of Mortgage-Backed Securities*. Chicago: Probus Publishing, 1992, 3rd edition, pp. 627-718.

Epanechnikov, V. "Nonparametric Estimates of Multivariate Probability Density." *Theory of Probability and Applications*, 14 (1969), pp. 153-158.

Hull, J. *Options, Futures and Other Derivative Securities*. Englewood Cliffs, NJ: Prentice Hall, 1993, 2nd edition.

Litterman, R., and J. Scheinkman. "Common Factors Affecting Bond Returns." *Journal of Fixed Income*, 1 (1991), pp. 54-61.

Roberts, R.B. "Hedging Mortgage-Backed Securities: A Review of the Approaches." In F. Fabozzi, ed., *The Handbook of Mortgage-Backed Securities*. Chicago: Probus Publishing, 1987, 2nd edition. pp. 855-888.

Scott, D.W. *Multivariate Density Estimation: Theory, Practice and Visualization*. New York: John Wiley & Sons, Inc., 1992.