2005 V33 4: pp. 681–710
REAL ESTATE
ECONOMICS

An Empirical Test of a Two-Factor Mortgage Valuation Model: How Much Do House Prices Matter?

Chris Downing,* Richard Stanton** and Nancy Wallace***

This article develops a two-factor structural mortgage pricing model in which rational mortgage-holders choose when to prepay and default in response to changes in both interest rates and house prices. We estimate the model using comprehensive data on the pool-level termination rates for Freddie Mac Participation Certificates issued between 1991 and 2002. The model exhibits a statistically and economically significant improvement over the nested one-factor (interest-rate only) model in its ability to match historical prepayment data. Moreover, the two-factor model produces origination prices that are significantly closer to those quoted in the to-be-announced market than the one-factor model. Our results have important implications for hedging mortgage-backed securities.

The residential mortgage-backed security (MBS) market is one of the largest and fastest-growing bond markets in the United States.[1] Valuing and hedging MBS requires a model for both the prepayment and default behavior of the underlying mortgages. While our understanding of this behavior has improved dramatically over the last two decades, significant challenges still remain. These challenges include, for example, the persistence of model-based MBS pricing errors (quantified in terms of the option-adjusted spread, or OAS), and the need for improved hedging strategies demonstrated by the sizable losses on MBS positions incurred by Askin Capital Management in the early 1990s and by Fannie Mae's recent problems with a large negative duration gap in its

*Jesse H. Jones Graduate School of Management, Rice University, Houston, TX 77005 or cdowning@rice.edu.
**Haas School of Business, University of California at Berkeley, Berkeley, CA 94720 or stanton@haas.berkeley.edu.
***Haas School of Business, University of California at Berkeley, Berkeley, CA 94720 or wallace@haas.berkeley.edu.

[1] According to the Federal Reserve's Flow of Funds accounts, at the end of the second quarter of 2004 the outstanding stock of residential MBS in the United States was $4.2 trillion. Residential MBS outstanding grew at an average rate of approximately 11% per annum from 1998 to 2003, about twice the rate of growth of outstanding corporate bonds issued by U.S. corporations.

portfolio.[2] Further, for firms that participate in the MBS market, recent changes in accounting standards now require firms to account for hedging effectiveness directly on the income statement, raising the visibility of any hedging mistakes.[3]

Some of these problems stem, at least in part, from the widespread use of models for pricing and hedging MBS that focus principally on the effect of interest rates on mortgage prepayment. While interest rates are generally acknowledged to be the most important factor affecting prepayment, there is a substantial literature suggesting that house prices also play an important role. For example, Stein (1995), Archer, Ling and McGill (1996), Mayer and Genesove (1997) and Mattey and Wallace (1998, 2001) emphasize the importance of housing prices as a determinant of regional-level household mobility. To the extent that declines in house prices impinge on borrowers' mobility, housing turnover and hence prepayments would fall. In a similar vein, Longstaff (2004) suggests that declining house prices reduce refinancing activity by impairing borrowers' credit quality, thereby hurting their chances to qualify for new loans. Conversely, several studies have shown that gains in home equity have a significant influence on the propensity to refinance, including, among others, Becketti and Morris (1990), Monsen (1992), Caplin, Freeman and Tracy (1993) and Krishnamurthy, Gabaix and Vigneron (2004). According to survey evidence reported in Canner, Dynan and Passmore (2002), 45% of homeowners who refinanced their mortgages in 2001 and early 2002 used the opportunity to extract equity. In summary, if house price movements affect mortgage prepayment in these and other ways, then any MBS pricing model that omits house prices as a state variable is misspecified.

A related issue is that the effects of mortgage defaults on MBS prices have received much less attention from both practitioners and the academic literature (notable exceptions in the literature include Kau *et al*. (1992, 1995), Schwartz and Torous (1992, 1993), Kau (1995), Deng, Quigley and Van Order (2000)). There are two basic reasons for the focus on prepayment rather than default. First, it is well known that defaults are generally rare events, given that most

[2] Askin is rumored to have lost around $400 million on MBS positions in 1993 when mortgage interest rates rose relative to Treasury rates. The *Wall Street Journal* reported on September 27, 2002, that "... an unprecedented boom in mortgage refinancings due to historically low interest rates—themselves related to weak stock prices—has some investors concerned that Fannie Mae is being tested. The refinancings have created a bigger gap than Fannie Mae would like between the expected lifespan of mortgages on its books and the borrowing it has done to finance them. Fannie Mae's stock is down about 11% because it announced the news." Beckett and Sender (2002).

[3] FAS 122 (May 1995) requires servicers to account for mortgage servicing rights on the balance sheet on a mark-to-market basis, and FAS 133 (June 1998) has introduced new standards for accounting for hedging effectiveness.

MBS are backed by first-lien mortgages protected by homeowner equity equal to 20% of the mortgage principal (80% loan-to-value ratio). Second, the bulk of residential MBS carry a credit guarantee from Ginnie Mae, a federal government agency, or either Freddie Mac or FannieMae, government-sponsored enterprises generally perceived to have the implicit backing of the federal government. Hence, for the bulk of MBS, the effect of mortgage defaults is to shift the return of principal forward in time, much like prepayments, and given the low incidence of defaults over recent history, a reasonable prior would seem to be that mortgage default can be ignored when modeling MBS.

In this article, we attempt to overcome both of these problems by developing and empirically estimating a structural two-factor mortgage valuation model that incorporates both interest rates and house prices as state variables. The structural approach has the potential to deliver a model that can produce informative forecasts in economic environments unlike those seen in the past, because mortgage terminations are the result of optimizing behavior by the agents in the model. Second, to the extent that overall terminations are correlated with house price movements, by incorporating house prices as a factor the model should more accurately describe termination behavior. The model explicitly values a borrower's joint option to prepay or default on his or her mortgage, thus allowing house price movements to affect both prepayment and default, and hence MBS prices. The model also incorporates discrete-time decision making on the part of borrowers and borrower-level heterogeneity in transaction costs, thereby capturing the well-known "seasoning" and "burnout" patterns observed in the termination behavior of home mortgages.

We estimate the parameters of the model using comprehensive data on termination rates for the mortgage pools backing Freddie Mac Participation Certificates issued between 1991 and 2002. The results indicate that house prices play a significant role in determining MBS prices. Specifically, when the two-factor model is compared to the nested one-factor interest-rate model (similar to that in Stanton (1995)), we find that the two-factor model produces a significantly better fit to the observed termination behavior of the pools. In addition, predicted MBS prices are closer to observed prices. We also use the model to examine the sensitivity of predicted MBS prices to movements in interest rates and house prices. This analysis indicates that, while MBS prices are primarily sensitive to interest rate fluctuations, house price movements also have an important effect. Movements in the value of the default option have a significant effect on the value of a mortgage borrower's prepayment option, and hence on the likelihood of prepayment. These results have important implications for hedging. In particular, even a strategy of hedging an MBS position against interest rate risk alone needs to take into account the fact that the optimal hedge ratio will vary substantially with the level of house prices.

The article is organized as follows. The next section reviews the existing valuation literature. The third section introduces the pool-level pricing model and the fourth section discusses how we implement the model. The fifth section describes the Freddie Mac data used in estimating the model, and the sixth discusses our estimation strategy and our results. The seventh section presents our pricing results and sensitivity analysis, and the final section concludes.

Pricing Mortgages

A fixed-rate home mortgage is a callable, defaultable bond whose payments are made by an individual borrower to a bank or other financial institution. Although many different mortgage types exist, we will focus on 30-year, fixed-rate mortgages—the loans backing most MBS. We will use the notation B_t to denote the market value of the remaining scheduled payments in the absence of any options; we refer to this stream of payments as the "underlying bond." Valuing a mortgage amounts to valuing this bond together with its embedded options.

Prepayment

At any time after taking out the mortgage, the borrower may choose to stop making the remaining scheduled payments, and instead pay off the remaining principal amount, F_t.[4] Paying off the loan is equivalent to exercising a call option on the bond B_t, with an exercise price equal to F_t. Under a one-factor interest-rate model, the lower current interest rates are, the higher B_t is, and hence the more in the money the prepayment option is. In a two-factor setting, movements in both interest rates and house prices determine the extent to which the option is in the money. When interest rates and house prices reach a boundary—the exact location of which is determined empirically for a reduced form model, or endogenously as part of the solution to the pricing problem for a structural model—the borrower exercises the option and pays off the loan early.

Default

In addition to choosing whether to make the scheduled monthly payment or to pay off the loan in full, a borrower may choose to default on the loan, handing over the house, the value of which we denote by H_t, and stopping all future mortgage payments. This default option is another call option on B_t, this time

[4] The remaining balance is calculated using the standard annuity formula:

$$F_t = \frac{\text{PMT}}{r}\left[1 - \frac{1}{(1 + c/12)^{12(T-t)}}\right],$$

where PMT is the monthly payment amount, c is the contractual interest rate on the loan and $T - t$ is the remaining time on the loan.

with exercise price H_t. All else equal, the lower is H_t, the more attractive exercise of the default option is. However, it is important to emphasize that the default and prepayment options are not independent of one another. Because exercise of one option precludes exercise of the other, the options are substitutes, and movements in the value of one affect the value of the other. Hence the borrower holds a joint option to terminate the mortgage at any time by either prepaying or defaulting.

Background Terminations

Prepayments and defaults can occur for reasons unrelated to interest rate and house price movements. For example, a borrower might prepay a loan after deciding to move to a different house. Default might occur following an uninsured event that damages the house. The possibility that these types of "background" terminations might occur affects the value of the joint option both directly and indirectly. The direct effect is that the joint option might be exercised when it would not be optimal if the decision were based solely on interest rates and house prices. The indirect effect is through a change in the optimal exercise policy for the joint option. For example, if the homeowner knows that he or she is likely to move, and hence might exercise the joint option "suboptimally" at some point in the future, then relative to the case where it is not possible to exercise suboptimally it is more attractive to exercise the option immediately.

Reduced-Form and Structural Models

With the foregoing discussion in mind, we write the value of the mortgage liability as

$$M_t^l = B_t - J_t^l, \tag{1}$$

where J_t^l is the value of the joint termination option to the mortgage holder. Valuing the mortgage requires a model for the exercise of the joint option, that is, a model for prepayments and defaults. Two approaches have emerged in the literature on modeling mortgage termination: reduced-form and structural models.

Reduced-form models. In the reduced-form approach, termination behavior is modeled as a function of a set of exogenous variables supposed to represent factors that influence the likelihood of mortgage termination, such as changes in interest rates, housing turnover and the like. Well-known reduced-form mortgage termination models include Schwartz and Torous (1989), Deng, Quigley and Van Order (2000) and Deng and Quigley (2002), and most of the mortgage valuation models used on Wall Street are based on reduced-form mortgage termination models. The main advantages of this approach are flexibility and the ability to closely mimic the historical data record of mortgage terminations.

However, as is well known, the flexibility of the reduced-form approach comes at the cost of potentially low out-of-sample forecasting power. Moreover, these models are often not well suited for valuation, to the extent that mortgage prices or proxies for the prepayment and default option values are included in the set of exogenous variables used to predict terminations.

Structural models. The structural approach treats mortgage termination as the optimal response of a rational borrower to changes in interest rates and house prices (and potentially other state variables). Referring to Equation (1), borrowers choose when they prepay or default in order to minimize M_t^l. Under this approach, standard contingent-claim techniques can be used to solve for the value of M_t^l, simultaneously calculating the borrower's optimal option exercise policy. This modeling approach was first applied to mortgages by Dunn and McConnell (1981a,b), who modeled the optimal termination behavior of borrowers who could costlessly prepay, but not default. More recent examples of interest-rate-based structural MBS valuation models include Timmis (1985), Dunn and Spatt (1986), Johnston and Van Drunen (1988) and Stanton (1995). Several recent articles, such as Dau *et al.* (1992, 1995) and Kau (1995), consider both default and prepayment but perform no empirical testing of their models.

Well-specified structural models ought to perform well out of sample because termination behavior arises from borrowers' optimizing behavior. However, structural models impose significant constraints on the relation between terminations and the underlying state variables. As a consequence, basic structural models, such as Dunn and McConnell (1981a,b), produce predictions for mortgage prices and termination behavior that diverge in important ways from what we observe in practice. First, these models predict that a mortgage (or MBS) can never trade above par, because borrowers will exercise their prepayment option the instant that the mortgage value exceeds par—what is often referred to as "ruthless option exercise." In practice, and as we will see below, MBS are often observed to trade above par. Second, because all borrowers are assumed to be identical, if any one borrower finds it optimal to refinance, so do all the others. Hence all borrowers should prepay simultaneously. Such behavior is not observed in actual MBS pools.

A Two-Factor Structural Model

Our objective is to develop a structural prepayment and default model that incorporates both interest-rate and house-price dependence and also produces termination behavior and mortgage prices close to those seen in practice. While a simple, frictionless structural model cannot deliver on all of these dimensions, we can overcome these shortcomings by incorporating frictions and transaction costs, along the lines of Stanton (1995).

Borrowers face transaction costs whenever they prepay or default on a mortgage, and these transaction costs probably vary across borrowers. Moreover, borrowers facing identical transaction costs might still take up the question of whether or not to exercise their prepayment or default option at different times. This would be the case if, for example, it is costly for borrowers to make the decision whether or not to prepay or default. Hence, we assume that

- borrowers face transaction costs whenever they refinance or default,
- borrowers are heterogeneous—different borrowers face different transaction costs and
- a borrower takes up the decision of whether or not to exercise his or her prepayment or default option with some probability in any given period of time.

These assumptions are similar to those in the one-factor interest-rate model developed in Stanton (1995), except that here they also have important implications for the default component of the joint termination option.

Transaction Costs and Borrower Heterogeneity

Prepayment involves both direct monetary costs, such as origination fees and mortgage closing costs, as well as implicit costs, such as the time required to complete the process. We model all of these via a proportional transaction cost, $X_p \geq 0$, payable by the borrower at the time of prepayment. Prepayment is optimal for the borrower if

$$M_t^l \geq F(t)(1 + X_p). \tag{2}$$

Different borrowers might face different transaction costs. To account for this possibility, we assume that the costs X_p are distributed according to a beta distribution with parameters β_5 and β_6. This distribution is chosen because it can take many possible shapes, and is bounded by zero and one. Its mean and variance are

$$\mu = \frac{\beta_5}{\beta_5 + \beta_6}$$

$$\sigma^2 = \frac{\beta_5 \beta_6}{(\beta_5 + \beta_6)^2(\beta_5 + \beta_6 + 1)}.$$

In the absence of transaction costs, the borrower will optimally default if the value of the mortgage is greater than or equal to the value of the house. However, like prepayment, defaulting incurs significant direct and indirect costs, such as the value of the lost credit rating. We model these costs via another proportional transaction cost, X_d, payable by the borrower at the time of default. Default is optimal for the borrower if:

$$M_t^l \geq H_t(1 + X_d). \tag{3}$$

As we discuss below, for computational tractability, we assume that $X_d = 0.05$ (5% of house value).

Option Exercise

We describe the probability of a decision on option exercise with hazard functions (Kalbfleisch and Prentice 1980, Cox and Oakes 1984). Informally, if the hazard function governing some event is λ, then the probability that the event occurs in a time interval of length δt, conditional on not having occurred prior to t, is approximately $\lambda \delta t$. As noted earlier, borrowers might also be forced to prepay or default for nonfinancial reasons (such as divorce, job relocation or sale of the house), which we assume is also described by some hazard function. We refer to this as the "background" hazard.

We assume that the probability of prepayment or default in any time interval is governed by the state- and time-dependent hazard function, λ. The value of λ depends on whether it is currently optimal for the borrower to default or prepay, which in turn is determined as part of the valuation of the mortgage. We model the overall hazard rate governing mortgage termination as

$$\begin{aligned} \lambda(t) &= \beta_0 + \beta_1 \text{atan}\left(\frac{t}{\beta_2}\right) P_t + \beta_3 \text{atan}\left(\frac{t}{\beta_4}\right) D_t \qquad (4) \\ &= \lambda_c + \lambda_p + \lambda_d, \qquad (5) \end{aligned}$$

where β_0 denotes the background hazard, the indicator variable P_t is one when prepayment is optimal at time t and zero otherwise, and the indicator D_t is one when default is optimal and zero otherwise. The atan function captures the idea of "seasoning" (see, *e.g.*, Richard and Roll 1989), where *ceteris paribus* new loans terminate more slowly than older loans. In the prepayment region, the termination rate rises over time at a rate governed by β_1 to a maximum rate dictated by the value of β_2. Similarly, in the default region, termination rates rise at a rate governed by β_3 to a maximum given by β_4. For simplicity in what follows, we will use the notation given in Equation (5) to refer to the hazard rates that apply in the various regions of the state space, where $\lambda_c \equiv \beta_0$, $\lambda_p \equiv \beta_1 \text{atan}\,(\frac{t}{\beta_2})P_t$ and $\lambda_d \equiv \beta_3 \text{atan}\,(\frac{t}{\beta_4})D_t$.

The one-factor interest-rate model is nested within the full two-factor model through parameter restrictions. Specifically, we set the transaction cost on default to a prohibitive level ($X_d = 1000$), and we set $\beta_3 = 0$ and $\beta_4 = 1.0e + 9$, which restricts the hazard rate in the default region to zero. Taken together, these parameter restrictions guarantee that cash flows due to default cannot occur for any realization of the state variables.

Implementing the Model

As noted earlier, in our model the state of the world is summarized by two variables, the short-term default-free interest rate, r_t, and the house price, H_t. To implement the model and calculate mortgage values, we need to make assumptions about how these underlying state variables evolve through time.

Interest Rates

We assume interest rates are governed by the Cox, Ingersoll and Ross (1985) model,[5]

$$dr_t = (\kappa(\theta_r - r_t) - \eta r_t)\, dt + \phi_r \sqrt{r_t}\, dW_{r,t}, \tag{6}$$

where κ is the rate of reversion to the long-term mean of θ_r, η is the price of interest rate risk and ϕ_r is the proportional volatility in interest rates. The process $W_{r,t}$ is a standard Wiener process.

We estimated the following parameters for the model using the methodology of Pearson and Sun (1989) and daily data on constant maturity 3-month and 10-year Treasury rates for the period 1968–1998:

$$\begin{aligned} \kappa &= 0.13131 \\ \theta_r &= 0.05740 \\ \phi_r &= 0.06035 \\ \eta &= -0.07577. \end{aligned}$$

It is important to emphasize that these parameter values are held fixed when we estimate the model.

House Prices

The house price, H_t, is assumed to evolve according to a geometric Brownian motion:

$$dH_t = \theta_H H_t\, dt + \phi_H H_t\, dW_{H,t}, \tag{7}$$

where θ_H is the expected appreciation in house prices and ϕ_H is the volatility of house prices. Denoting the flow of rents accruing to the homeowner by q_H, after risk-adjustment house prices evolve according to

$$dH_t = (r_t - q_H)H_t\, dt + \phi_H H_t\, dW_{H,t}. \tag{8}$$

[5] This model is widely used in the mortgage pricing literature. See, for example, Kau *et al.* (1992) and Stanton (1995).

We calibrate Equation (8) as follows:

$$q_H = 0.025$$
$$\phi_H = 0.085.$$

The value of q_H is roughly consistent with estimates of owner-equivalent rents from the BEA, and we estimate the annualized volatility of housing returns from our data on house prices, discussed below. House prices and interest rates are assumed to be uncorrelated.[6]

Given these models for interest rates and house prices, standard arguments show that, in the absence of arbitrage, the value of the borrower's mortgage liability, $M^l(H_t, r_t, t)$, paying coupon c, must satisfy the partial differential equation:

$$\begin{aligned}
&\frac{1}{2}\phi_r^2 r M_{rr}^l + \frac{1}{2}\phi_H^2 H^2 M_{HH}^l + (\kappa(\theta_r - r) - \eta r)\, M_r^l + ((r - q_H)H_t)\, M_H^l \\
&\quad + M_t^l - r M^l + (\lambda_c + \lambda_p)\left(F(t)(1 + X_p) - M^l\right) \\
&\quad + \lambda_d\left(H(1 + X_d) - M^l\right) + c = 0. \qquad (9)
\end{aligned}$$

We also need to impose boundary conditions. The first three of these are

$$M^l(H, r, T) = 0, \qquad (10)$$

$$\lim_{r \to \infty} M^l(H, r, t) = 0, \qquad (11)$$

$$\lim_{H \to \infty} M^l(H, r, t) = C(r, t), \qquad (12)$$

where $C(r,\, t)$ is the value of a callable bond with the same promised cash flows and same prepayment costs as the mortgage, but with no house price dependence.[7] Equation (10) is the terminal condition, reflecting the amortization of the mortgage. Equation (11) arises because all future payments are worthless when interest rates approach infinity, and Equation (12) says that when the house price gets large, default no longer occurs, so we only have to worry about prepayment.

We need additional boundary conditions specifying the free boundary governing optimal default and prepayment. Prepayment is optimal when interest rates

[6] This assumption is made to simplify the interpretation of the results. In terms of solving the pricing problem and carrying out our econometric estimation below, it is straightforward to handle correlated house prices and interest rates.

[7] This value is calculated following the process described in Stanton (1995).

go below some (house price-dependent) critical level, $r^*(H, t)$, and default is optimal when the house price drops below some (interest rate-dependent) critical level, $H^*(r, t)$. At these boundaries, the mortgage value satisfies the conditions

$$M^l(H, r^*(H,t), t) = F(t)(1 + X_p), \tag{13}$$

$$M^l(H^*(r,t), r, t) = H^*(r,t)(1 + X_d). \tag{14}$$

Equation (13) states that, on the optimal prepayment boundary, the mortgage value is just equal to the remaining balance multiplied by 1 + the appropriate transaction cost. Equation (14) states that, on the default boundary, the mortgage is just equal to the value of the house multiplied by 1 + the default transaction cost.[8]

Solving Equation (9) subject to these boundary conditions gives us the value of the borrower's liability, as well as the locations of the optimal default and prepayment boundaries, which in turn determine the values of the prepayment and default hazard rates, λ_p and λ_d.

Given these values, we can now solve for the value of the lender's *asset*, M^a, which is the solution to the following related partial differential equation:

$$\begin{aligned}
&\frac{1}{2}\phi_r^2 r M_{rr}^a + \frac{1}{2}\phi_H^2 H^2 M_{HH}^a + (\kappa(\theta_r - r) - \eta r)\, M_r^a + ((r - q_H)H_t)\, M_H^a \\
&\quad + M_t^a - rM^a + (\lambda_c + \lambda_p)(F(t) - M^a) \\
&\quad + \lambda_d(H - M^a) + c = 0,
\end{aligned} \tag{15}$$

subject to the boundary conditions

[8] There are two additional "smooth-pasting" boundary conditions (see Merton 1973) that ensure the optimality of the boundaries $r^*(H)$ and $H^*(r)$. We use a parallel hopscotch finite difference algorithm to solve Equation (9), in which these conditions are automatically satisfied. The algorithm involves discretizing the PDE in (9) as discussed in Gourlay and McKee (1979). To simplify the numeric solution, before discretizing the problem we first transform to the state space defined by

$$y = \frac{r}{\gamma_y + r},$$

$$z = \frac{H}{\gamma_H + H}.$$

The "packing factors" γ_y and γ_H have the interpretation that they are the centers of their respective spatial grids in the untransformed state space. In practice, we set $\gamma_y = \theta_r$ and $\gamma_H = 1.25$. The packing factors ensure that we make efficient use of the finite set of points on our numeric solution grid.

$$M^a(H, r, T) = 0, \tag{16}$$

$$\lim_{r\to\infty} M^a(H, r, t) = 0, \tag{17}$$

$$\lim_{H\to\infty} M^a(H, r, t) = C(r, t), \tag{18}$$

$$M^a(H, r^*(H, t), t) = F(t), \tag{19}$$

$$M^a(H^*(r, t), r, t) = H^*(r, t). \tag{20}$$

The Data: Freddie Mac Mortgage Pools

Our empirical analysis focuses on the termination characteristics of Freddie Mac pass-through residential MBS. The data for this study consist of Gold Participation Certificate (Gold PC) pools issued by Freddie Mac between January 1991 and December 2002. The underlying mortgages in the Gold program are primarily first-lien residential mortgage loans secured by one- to four-family dwellings. We focus on the pools backed by newly issued, standard 30-year fixed-rate mortgages.[9] We observe pool-level mortgage termination rates from the month of issuance for each pool until the pool is fully paid off or until we reach the end of our observation period.[10] The termination histories of the pools have a median length of 74 months, with a range from 151 months for very low coupon pools to just 5 months for very high coupon pools.

We measure housing values using the Office of Federal Housing Enterprise Oversight (OFHEO) quarterly repeat sales home price indexes, computed at the state level. The state-level series are interpolated to a monthly frequency using standard splining techniques. Freddie Mac reports the initial origination balances by state for each pool. Based on these balances, for each pool we compute a unique composite house price index that is a weighted average of the OFHEO price indexes for the states that are represented in the origination balances. The weight on each state is equal to the state's share of the overall loan balance of the pool.

In Table 1, we summarize the properties of the Freddie Mac pools by year of origination. There are a total of 1,324,898 mortgages in the 19,304 pools that

[9] Specifically, we analyze the subset of pools backing 30-year Gold Participation Certificates with an original weighted average loan age of two months or less and an original weighted average remaining maturity of 350 months or more.

[10] The termination rates are computed as single-month mortality rates following standard formulae such as that on page 205 of Bartlett (1989) for estimating terminations given data on pool factors, time to maturity and coupon interest rates.

Table 1 ■ Summary measures for Freddie Mac pools.

				Housing Returns Average Annual Return Years 1–3 (%)			Interest Rate Average Annual Change Years 1–3 (bps)				SMM End of Year 3		
					Percentiles			Percentiles				Percentiles	
Year	WAC	Mortgages	Pools	Average	25th	75th	Average	25th	75th	No. of Pools	Average	25th	75th
1991	9.500	37,367	1,825	3.2	0.5	6.7	−14	−34	−2	1,753	0.00024	0.00000	0.03782
1992	8.625	79,622	3,819	4.3	2.2	7.1	−18	−38	−3	3,817	0.00028	0.00010	0.00065
1993	7.625	46,989	1,607	5.1	3.3	7.5	−21	−46	−4	1,607	0.00049	0.00021	0.00701
1994	7.875	52,346	1,206	5.7	4.0	7.6	−27	−51	−9	1,205	0.00073	0.00030	0.01098
1995	8.250	31,056	611	6.3	4.9	7.5	−24	−43	−6	611	0.03652	0.00791	0.06431
1996	7.875	44,825	725	7.0	5.8	8.1	−13	−46	−4	723	0.01590	0.00048	0.03756
1997	7.750	65,669	941	7.5	6.6	8.3	−14	−70	−4	941	0.00656	0.00043	0.01629
1998	7.000	150,122	1,940	8.0	7.2	9.0	−43	−75	−6	1,940	0.01030	0.00096	0.01958
1999	7.500	59,556	921	8.6	7.8	9.4	−75	−83	−69	920	0.03410	0.01518	0.05437
2000	7.875	66,186	628	8.5	7.8	9.2	−83	−84	−73	287	0.10586	0.06421	0.13998
2001	6.875	319,446	2,888	NA			NA				NA		
2002	6.375	371,714	2,193	NA			NA				NA		
Total		1,324,898	19,304										

Note: The table displays summary measures for the Freddie Mac pools used in our empirical analysis. The pools are broken down by year of origination. The column labeled "WAC" displays the average weighted-average coupon for the pools originated in the given year, expressed in percentage points. The column labeled "Mortgages" displays the number of mortgages at origination in the given year cohort and the column labeled "Pools" shows the number of mortgage pools in the cohort. The columns under the heading "Housing Returns" display univariate statistics for the average annual return to housing measured over the first three years from the origination of a pool. For example, the column labeled "25th Percentile" under this heading indicates that in 1991, 25% of the pools experienced annual average returns on housing of less than 0.5% over the first three years from their date of origination. The columns under the heading "Interest Rate" give univariate statistics for the annual average change in the 10-Year Treasury rate over the first three years of each pool's life. The column labeled "No. of Pools" shows the number of pools outstanding after three years, and the columns under the heading "SMM" display univariate statistics for the single-month mortality rates of the pools at the end of the 36th month from origination.

we track through July 2003. The weighted average coupons (WACs) drifted steadily downward from a high of 9.5% in 1991 to a low of 7% in 1998. WACs rose to 7.875% in 2000 and then fell sharply to 6.375% by 2002. The volume of mortgage originations in part reflects the dynamics of interest rates. Originations rose to a high of 3,819 pools when rates decreased in 1992. The periods of rate increases in 1995 and 2000 led to sizable reductions in pool originations, while volumes increased again with the refinancing boom in 2001 and 2002 when rates fell to 40-year lows.

For each cohort, the annualized average three-year return to housing is summarized in columns 5 through 7 of Table 1. The average three-year return was only 3.2% in 1991 but grew steadily to about 8.6% in 2000 and 2001. The reported percentile ranges indicate that there is dispersion in housing returns across the states in 1991 through 1993. There is also considerable time series variation in the house price indexes over the various vintages of pools. The first few years of the 1991, 1992 and 1993 vintage pools tended to have runs of low and even negative average three-year housing returns, as suggested by the relatively low 25th percentile returns. These pools tended to have large proportions of their initial balances in California mortgages. The observed downward pressure on house prices for the California pools reflects the relatively severe California recession in the early 1990s.

We report univariate statistics for the average annual change in the 10-year Treasury bond rate over the first three years from each pool's origination date in columns 8 through 10 of Table 1. The average changes in the long rate were negative for all the cohorts, although the time series performance of the long rate for the early 1990s vintage pools included periods of both rising and rapidly falling rates. The pools that were originated in 1999 through 2001 experienced continual decreases in the long rate, as is clear from the narrow interquartile ranges for rate changes over these periods.

In column 11 of Table 1, we present the number of pools that remain outstanding 36 months after origination. The median termination rates, or single-month mortality rates, are reported in column 12 and the 25th and 75th percentile ranges for monthly termination rates appear in the last two columns of the table. As shown, the median single-month mortality rates rise steadily over the observation period to a high of about 10% per month. The 25th and 75th percentile ranges reflect the time-series dynamics of pool terminations. Terminations tend to exhibit rather slow levels in the earliest months and then, depending on the level of interest rates and house prices, often experience widely varying levels of termination behavior over the life of the pools. As indicated by the relatively high 75th percentile for the 1991 cohort, some of these relatively high coupon pools exhibited rapid termination rates when interest rates fell in mid-1992 and

1993. In contrast, the low coupon pools that were originated in late 1993 and early 1994 exhibited markedly lower average termination rates after three years.

As an illustration of the important role that house returns play in pool termination behavior, we graphically compare the termination speeds of pools that were originated at the same time but experienced very different patterns of housing returns. We define "negative return pools" to be those pools that experienced negative average returns—as measured by each pool's composite housing price index—over at least one three-year period from the origination date to the pool termination date or the end of the sample period in July 2003. The set of "positive return pools" is the complement of the negative return pools set. According to this definition, 17% of the pools experienced at least one spell of negative returns to housing; these pools were primarily originated in 1991 (1,160 pools), 1992 (1,892 pools) and 1993 (308 pools). The negative-return pools all contain a large share of California mortgages that were exposed to a significant downturn in California house prices in the early 1990s. Most of these pools experienced episodes of negative housing returns early in their payment histories before their principal balances had amortized significantly.

In Figure 1, we present the average monthly termination rates for pools that were originated in 1991. We compare the negative-return pools to the positive-return pools originated in this same year. As can be seen, on average the negative-return pools exhibited markedly lower termination rates in the first 30 months of their payment histories than the positive-return pools. Interestingly, while the mean WACs of these two sets of pools are the same, the borrowers in the negative-return pools terminated their mortgages at much lower rates than the borrowers in the positive-return pools when exposed to sharp decreases in interest rates, consistent with the idea that declines in house prices dampened refinancing behavior through a mobility or credit quality channel. It is striking that once the composite house prices for these pools begin to rise in the later periods, the termination speeds of the negative- and positive-return pools track one another closely. Although not shown, the 1992 and 1993 vintage pools also exhibited important differences in their termination behavior over the first 30 months of their payment histories.

Estimation Strategy and Results

As in Stanton (1995), we use a nonlinear least squares procedure to estimate the coefficients of the model. For any given set of parameters, the valuation procedure described above generates a predicted termination rate for each month. If we have the right parameters, these predicted termination rates ought to be close, on average, to those we actually observe. Our estimation strategy involves searching for the set of parameters for the model that minimizes the sum

Figure 1 ■ Average single month termination rates for pools originated in 1991 that experienced positive and negative housing returns. The figure compares mortgage termination rates for the Freddie Mac PC pools used in our empirical analysis, where the pools are split into those that experienced negative average annual returns to housing for any three-year period over the life of the pool, and those that experienced only positive returns to housing over the life of the pool. All pools were originated in 1991 and are tracked by pool month. The set of "Negative Return" pools contains 17% of the sample (3,360 pools); the set of "Positive Return" pools contains the balance of our sample (15,944 pools).

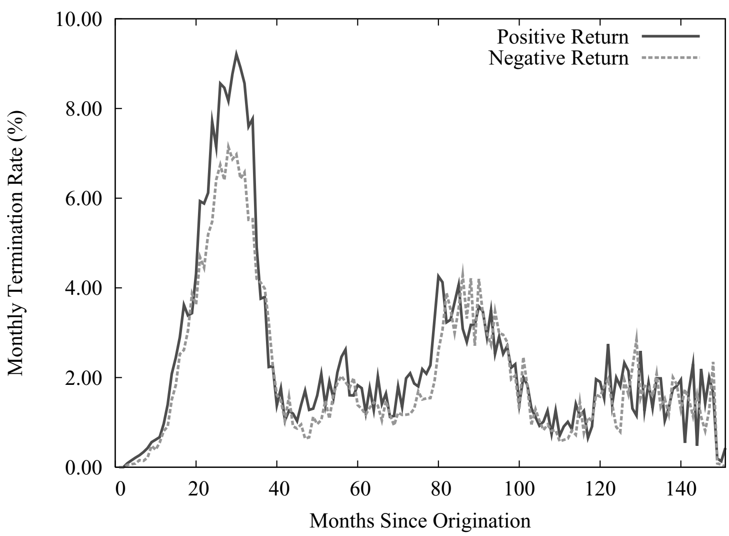

of squared differences between the termination rates predicted by our model and those observed in the data.

Formally, let $\hat{\omega}_{it}(\Theta)$ denote the predicted proportion of the balance of pool i that terminates in month t, as a function of the vector of coefficients to be estimated, Θ. If ω_{it} denotes the observed proportion that terminates (the single-month mortality rate), our objective function is

$$\chi(\Theta) = \sum_{i=1}^{N} \sum_{t=1}^{T_i} (\omega_{it} - \hat{\omega}_{it}(\Theta))^2 , \tag{21}$$

where N is the number of mortgage pools and T_i is the number of observations on pool i. We use the Nelder–Mead downhill simplex algorithm to find the vector of coefficients Θ that minimizes $\chi(\Theta)$.

Table 2 ■ Estimation results.

	One-Factor Model		Two-Factor Model	
Coefficient	Estimate	Std. Err.	Estimate	Std. Err.
β_0	0.03108	0.00015	0.00425	0.00016
β_1	0.85900	0.00032	0.88795	0.00037
β_2	1.22653	0.00087	0.79885	0.00087
β_3			0.82295	0.00276
β_4			0.54126	0.00870
β_5	0.85284	0.00059	0.86703	0.00056
β_6	3.85602	0.00202	5.09893	0.00248
χ	64.8525		64.6690	
N	1,349,180		1,349,180	

Note: The table displays the estimation results for the hazard specification given by

$$\lambda(t) = \beta_0 + \beta_1 \arctan(t/\beta_2)P_t + \beta_3 \arctan(t/\beta_4)D_t.$$

The dummy variable P_t is one when just the prepayment option is exercised, and zero otherwise. The dummy variable D_t is one when either the default or prepayment options are exercised, and zero otherwise.The coefficients β_5 and β_6 determine the transaction cost distribution. The mean of the transaction cost distribution is given by

$$\mu = \frac{\beta_5}{\beta_5 + \beta_6}$$

and its variance is given by

$$\sigma^2 = \frac{\beta_5 \beta_6}{(\beta_5 + \beta_6)^2(\beta_5 + \beta_6 + 1)}.$$

Under the one-factor model, the coefficients β_3 and β_4 are restricted so that the hazard rate is zero. The row labeled χ displays the objective function values under each model, and the sample size is 1,349,180 for each estimation run.

Results

Table 2 reports parameter estimates for the one- and two-factor specifications of the model as described in the preceding sections. For the one-factor model, we estimate the background hazard rate, β_0, the coefficients β_1 and β_2 that govern the pace of terminations due to prepayments, and the coefficients that define the prepayment transaction cost distribution, β_5 and β_6. The second and third columns of Table 2 report the estimates and asymptotic standard errors for these coefficients. For the two-factor model, we also estimate the two coefficients that govern the pace of terminations due to default, β_3 and β_4. The results for the two-factor model are displayed in columns 4 and 5.

Given that we are estimating on a sample of 1,349,180 pool-month observations, it is not surprising that all of the coefficient estimates are highly statistically significant. Moreover, with such a large sample, any restriction on the two-factor model will be rejected. Indeed, the standard ratio test statistic for nonlinear least squares indicates that the reduction in the objective function shown in Table 2 is statistically significant at better than the 99% level of significance. Following Amemiya (1985), if we let $\chi^{(1)}$ denote the value of the nonlinear least-squares objective function under the one-factor model (representing two parameter restrictions) and $\chi^{(2)}$ the value under the two-factor model, then the quantity $N[\ln(\frac{\chi^{(1)}}{N}) - \ln(\frac{\chi^{(2)}}{N})]$ is distributed chi-squared with two degrees of freedom. Inserting the relevant quantities produces a test statistic value of about 3,800, while the relevant critical value is approximately nine.

The estimated one-factor background hazard rate, β_0, is considerably smaller than the coefficient obtained under the two-factor model. This result suggests that the background hazard under the one-factor model is in part acting as a proxy for the effects of default on overall terminations. The parameters β_1 and β_2 govern the time-dependent seasoning component of the hazard rate (the so-called "ramp-up") to the maximum hazard rate for prepayment. This initial ramp-up in the termination rate is thought to arise because borrowers may initially face liquidity constraints (having just made outlays for origination costs, for example) and that these constraints dissipate over time. These estimates indicate that the one-factor model predicts slower termination speeds than are typically observed in the first months after origination for the pools.

The estimates for β_1 under the two models are quite similar. Hence the ramp-up rates are similar under the two models. The two-factor specification reaches 95% of its maximum by the tenth month after the origination of the pool. The ramp-up is only slightly slower under the one-factor specification, reaching about 92% of its maximum pace by the tenth month after origination. However, the estimate of β_2 for the two-factor specification is about one-third the magnitude of the estimate for the one-factor model. As a result, the maximum pace of prepayments under the one-factor model is only about 90% of the maximum pace under the two-factor model. Taken together, these results indicate that when interest rates and house prices are configured such that prepayment is optimal under both models, the two-factor model will predict a significantly faster pace of terminations.

The parameters β_3 and β_4 provide for the possibility of a ramp-up in default terminations. As discussed earlier these parameters are restricted to zero under the nested one-factor model. The predicted default ramp-up is relatively slow compared to the refinancing ramp-up, not reaching its maximum pace until about the 81st month after origination. The maximum pace of default terminations is about 91% that of prepayments. These results indicate that, compared to its action in the prepayment region, in the default region the hazard function is

Figure 2 ■ Estimated refinancing cost distribution. The figure shows the estimated cumulative distribution of refinancing costs paid by borrowers within a mortgage pool, expressed as a fraction of the remaining principal balance on the loan. The transaction costs are assumed to be distributed in the population according to a beta distribution with mean and variance given by

$$\mu = \frac{\beta_5}{\beta_5 + \beta_6}$$

$$\sigma^2 = \frac{\beta_5 \beta_6}{(\beta_5 + \beta_6)^2(\beta_5 + \beta_6 + 1)}.$$

Under the one-factor model, the mean refinancing cost is 18.1% of pool balance, with a standard deviation of 16.1%; under the two-factor model, it is 14.5%, with a standard deviation of 13.4%.

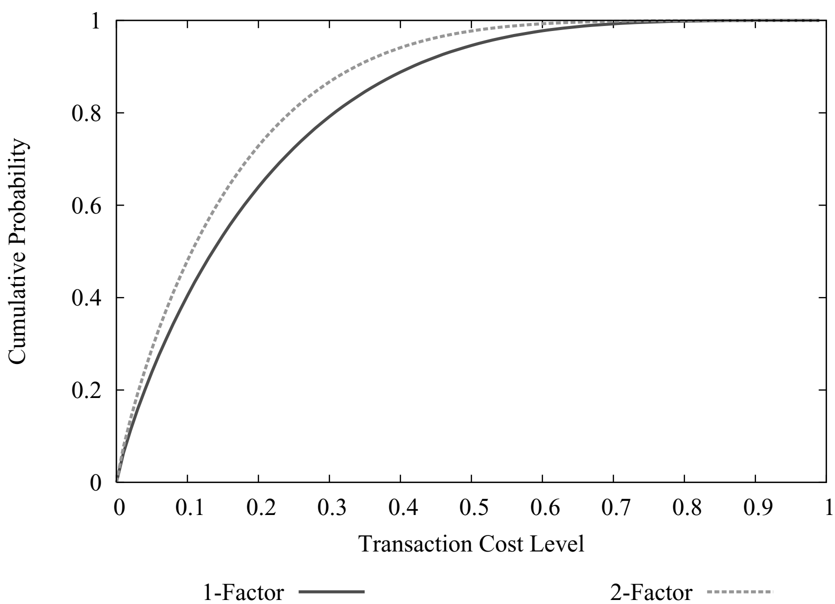

doing more to damp terminations relative to what they would be under a model of purely rational (ruthless) default.

The estimates of β_5 and β_6—the parameters determining the distribution of transaction costs in the borrower pool—reveal an important difference between the one- and two-factor models. The parameter estimates indicate that the mean of the transaction cost distribution is somewhat lower under the two-factor model, as illustrated in Figure 2. The estimates for the two-factor model indicate that the bulk of borrowers in the pools face relatively small transactions cost—averaging 14.5% of face value—in contrast to an average transaction cost of 18.1% under the one-factor model. The estimated standard deviation of transaction costs is smaller under the two-factor model: 13.4% versus 16.1%

Figure 3 ■ Full sample observed and predicted single month termination rates. *Note:* The figure compares actual and model-predicted monthly mortgage termination rates by month since issue, averaged across all pools in the sample.

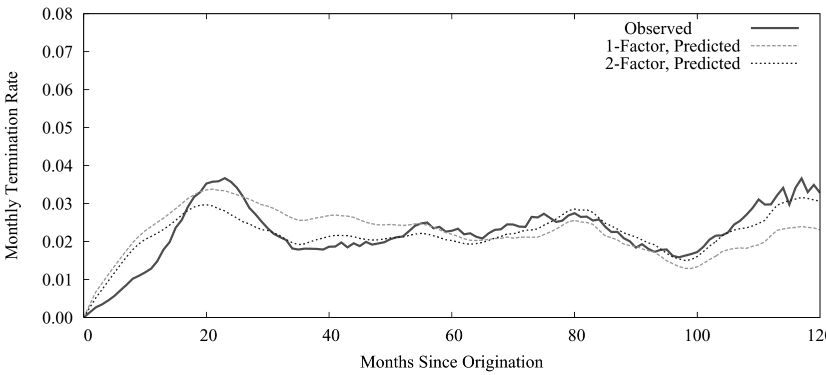

for the one-factor model. These changes in the estimated transaction cost distribution suggest that house prices are a source of heterogeneity across pools that is not controlled under the one-factor model.

Figures 3 and 4 translate the coefficient estimates for the model into average predicted monthly termination rates. The figures display plots of the "in-sample" paths of average predicted terminations against average observed termination rates for pools in the full sample and for pools with different coupon levels.[11] First, it is apparent that the two-factor model more closely tracks the sample averages than the one-factor model. Under both models, predicted terminations do not rise as rapidly as observed terminations in the first 20 months, and they tend to overshoot observed terminations from about 25–50 months. Both models track observed terminations from months 50 onward, though the two-factor model is clearly closer to the observed average pace.

Panel A of Figure 4 plots the path of predicted terminations for the 9.5% mortgage pools, where we observe the most prepayment activity. As shown, the model somewhat underpredicts the first termination wave, and it somewhat overpredicts terminations in the later months. The fits for the 8.5% pools, shown in Panel B, are quite accurate for the two termination booms; however, both models mistakenly predict a termination boom in the 45th month with the one-factor model predicting a pace of monthly terminations about half a percentage point above that of the two-factor model. For the 7.5% pools, shown in Panel C, the models capture the ramp-up effectively, but in general underpredict the spikes and overpredict the quiescent period from the 30th to 40th months from origination.

[11] The average termination rates are computed across the pools in a given month.

Figure 4 ■ Observed and predicted single month termination rates by coupon.
Note: The figure compares actual and predicted monthly mortgage termination rates by month since issue, averaged across all pools in the subsample of pools for the indicated coupon rate.

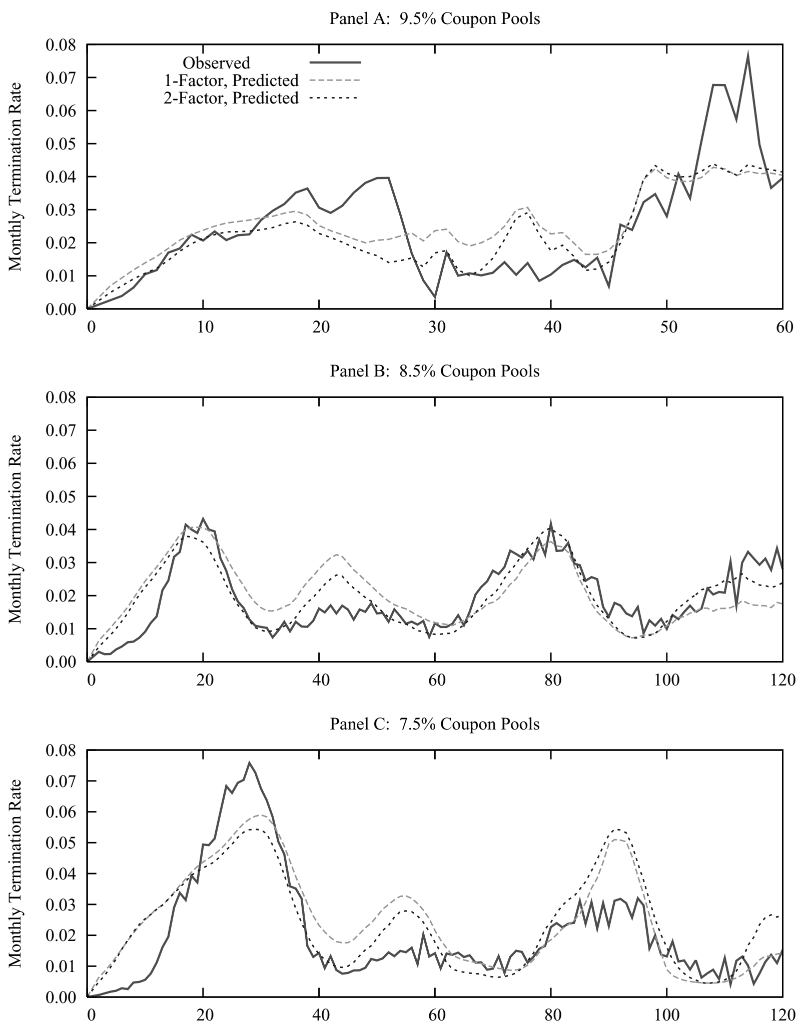

Prices and Sensitivity Analysis

In this section, we compare the pricing accuracy of the one- and two-factor models, and we examine the sensitivity of the predicted prices under the two-factor model to changes in interest rates and house prices. We compute the predicted prices at origination for a subset of the pools in our analysis and compare these predictions with observed prices for Freddie Mac Gold PCs in the to-be-announced (TBA) market. In the TBA market, buyers and sellers decide on general trade parameters, such as coupon settlement date, par amount and price, but the buyer does not know which pools actually will be delivered until two days before settlement.[12] Our TBA data set includes prices for Freddie Mac Gold PCs on a 50 basis point grid of coupons spanning the range from 6.5% to 9.5%, so we subset our data to the pools with coupons that align with the coupons available in the TBA price sample. This subsample amounts to 6,827 pools, or about one-third of the pools used to estimate the models. This out-of-sample test of pricing performance represents a more rigorous standard than has typically been applied in evaluating the performance of structural bond valuation models (see Eom, Helwege and Huang 2004).

As shown in Table 3, pool by pool the two-factor model outperforms, with very few exceptions, the pricing accuracy of the one-factor model.[13] We find that the pricing errors for the two-factor model range from a low of about 5 cents per one hundred dollars of face amount to a high of about \$4.38 per one hundred dollars of face. These errors represent a small fraction of those obtained using the one-factor model and on average the two-factor pricing errors represent a 30% absolute reduction in the one-factor errors.

We also compute the estimated option adjusted spread (OAS) for each coupon group, again using observed TBA prices at the pool origination dates. We find OAS levels that vary between zero and about 25 basis points—well within published OAS levels for these vintages and coupons (Bloomberg reports OAS levels as high as 300 basis points for some of these pools, though their modeling technique is quite different). In summary, these results suggest that the inclusion of house price dynamics leads to a significant improvement in the pricing accuracy of the MBS valuation model.

Price Sensitivity

In Figure 5, we report a set of duration and convexity calculations for a representative 30-year 8.5% coupon mortgage-backed security at origination ($t = 0$) using our two-factor structural valuation model. The panels on the left display

[12] We gratefully acknowledge Lakhbir Hayre at Salomon Smith Barney for providing the TBA price data.

[13] In addition, the degree of mispricing that we find here is well within the ranges reported in studies using "in sample" tests on corporate bond pricing models, such as Eom, Helwege and Huang (2004) and Huang and Huang (2002). We should also note that the coefficients on the underlying SDEs in Equations (6) and (7) are held fixed during this exercise. Presumably, if we were to re-fit the interest rate and house price processes at every date, we could further improve on these results. However, such a refitting procedure is computationally infeasible and is inconsistent with the equilibrium modeling approach we take here.

Table 3 ■ Average pricing errors and OAS for the two-factor model by year and coupon.

				One-Factor Model			Two-Factor Model		
Year	Coupon	*N*	Average TBA Price	Average Predicted Price	Average Pricing Error	Average OAS	Average Predicted Price	Average Pricing Error	Average OAS
1991	9.5	781	103.27	105.10	1.83	11.5	102.74	−0.54	−4.0
1992	8.5	740	103.23	104.71	1.49	7.6	101.89	−1.33	−7.7
1992	9.0	325	104.87	106.18	1.31	7.7	103.58	−1.29	−8.8
1992	9.5	10	106.25	106.28	0.03	0.0	103.88	−2.37	−20.0
1993	7.5	309	103.65	105.51	1.86	9.1	102.10	−1.56	−6.4
1993	8.0	181	104.57	105.83	1.26	7.3	102.72	−1.85	−8.4
1993	8.5	39	105.13	107.46	2.33	14.7	104.59	−0.55	−5.1
1993	9.0	7	105.94	108.78	2.84	20.0	106.06	0.12	0.0
1994	7.5	175	99.50	106.91	7.40	24.8	103.49	3.98	15.6
1994	8.0	69	99.34	105.48	6.14	21.8	102.36	3.02	14.3
1994	8.5	79	100.66	105.75	5.09	23.0	102.91	2.25	12.1
1994	9.0	29	102.53	107.05	4.52	23.5	104.42	1.89	12.0
1995	8.0	53	102.28	106.05	3.77	18.9	102.93	0.66	2.9
1995	8.5	56	102.31	105.45	3.14	15.0	102.61	0.30	1.9
1995	9.0	22	102.10	105.95	3.85	20.2	103.35	1.25	6.8
1996	7.5	69	100.42	106.36	5.94	21.9	102.94	2.52	9.8
1996	8.0	61	102.20	105.56	3.36	15.6	102.45	0.26	0.5
1996	8.5	13	104.12	107.26	3.14	18.9	104.39	0.27	0.3
1996	9.0	28	105.08	108.34	3.26	22.3	105.64	0.56	3.5
1997	7.0	8	100.30	101.99	1.68	5.0	98.18	−2.12	−5.6
1997	7.5	93	101.33	104.15	2.82	11.5	100.74	−0.60	−3.3
1997	8.0	158	102.49	105.55	3.06	15.2	102.44	−0.05	0.5
1997	8.5	9	103.50	106.73	3.22	16.1	103.87	0.37	1.1
1997	9.0	5	105.50	108.00	2.50	20.0	105.33	−0.17	0.0
1998	6.5	70	100.62	103.85	3.24	12.4	99.62	−0.99	−2.2
1998	7.0	168	101.76	105.63	3.88	16.3	101.83	0.08	1.3
1998	7.5	179	102.68	106.94	4.26	21.0	103.52	0.84	4.4
1998	8.0	14	103.65	107.69	4.04	18.6	104.54	0.89	6.0
1998	8.5	2	104.41	109.18	4.77	25.0	106.22	1.82	12.5
1998	9.0	1	106.28	109.43	3.15	25.0	106.67	0.39	5.0
1999	6.5	36	100.72	107.33	6.61	25.0	103.13	2.41	10.0
1999	7.0	27	101.94	109.22	7.28	25.0	105.42	3.48	15.0
1999	7.5	14	102.73	110.61	7.88	25.0	107.11	4.38	25.0
Pooled Averages									
Errors					2.71			−0.26	
Absolute Errors					2.89			1.70	

Note: The table displays the average predicted price, average pricing error and average option-adjusted spread (OAS) for the two-factor model. The pricing errors are calculated against TBA prices for MBS pools and TBA prices for pools with coupons equal to 7.5–9.5% in 50-basis-point increments. We use this subset of pools because TBA price data are not available for the other coupon levels that we use in our analysis. The OAS values are computed to the nearest five basis points for each pool. This subsample includes 6,827 pools, or about 35% of the pools used to estimate the model.

Figure 5 ■ Price sensitivity measures for two-factor model: 8.5% pool at time 0. *Note:* The figure reports a set of duration and convexity calculations for a representative 30-year 8.5% coupon mortgage-backed security at origination ($t = 0$). The sub-plots on the left display the approximate duration and convexity of the MBS, holding house prices fixed at \$1.25 (80% LTV). The duration and convexity calculations are made at three levels of transaction costs, zero (solid line), 10% (dashed line) and 40% (dotted line). The sub-plots on the right display the approximate durations and convexities of the mortgage pools holding interest rates fixed at 8.5%. Interest rate duration is calculated according to the standard formula for effective duration, given by $-\frac{1}{M^l}\frac{\partial M^l}{\partial r}$. Note that house price duration, computed as $-\frac{1}{M^l}\frac{\partial M^l}{\partial H}$, is negative; in contrast to interest rates, as house prices rise the value of the mortgage pool rises owing to a decline in the likelihood of default.

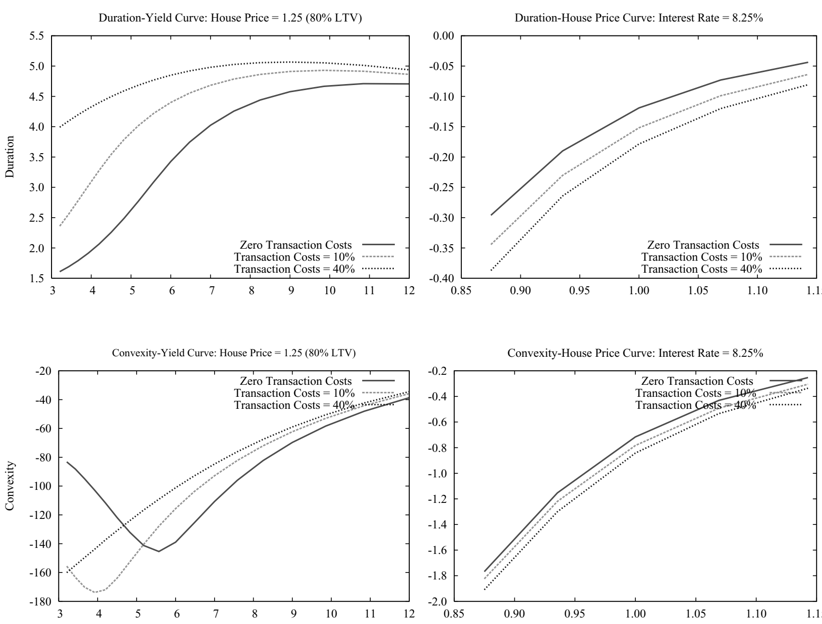

the approximate duration and convexity of the MBS, holding house prices fixed at \$1.25 (80% loan-to-value, or LTV). The duration and convexity calculations are made at three levels of prepayment transaction costs: zero (solid line), 10% (dashed line) and 40% (dotted line). Hence these calculations represent a partial decomposition of the overall MBS duration and convexity into the contributions of sub-pools distinguished by their transaction costs. In this regard, it is important to bear in mind that because the average prepayment transaction cost is low—around 15%—the contributions of the zero and 10% sub-pools to the overall MBS price sensitivity will be substantially greater than that of the 40% sub-pool.

The upper-left panel displays the approximate duration-yield curve holding house prices fixed at \$1.25 (holding the loan-to-value ratio fixed at 80%). As can be seen, when prepayment transaction costs are zero, the duration-yield curve has the characteristics typical of a callable bond: as interest rates fall and prepayment becomes more likely, the duration of the pool shortens dramatically. As we move to higher transaction costs, which all else equal reduce prepayment, the duration rises for all interest rates. At the maximum level of transaction costs, the likelihood of prepayments or defaults is very low, and the duration-yield curve resembles that of a straight bond.

The upper-right panel of Figure 5 displays the "duration house price" curve, which is computed as $-\frac{1}{M^l}\frac{\partial M^l}{\partial H}$.[14] In contrast to interest-rate duration, house price duration is negative: as house prices rise the value of the mortgage pool tends to rise. Here the prepayment transaction costs are only important insofar as shifts in the value of the prepayment option affects the value of the default option. Hence we see that house price duration becomes more negative as prepayment transaction costs rise, reflecting the fact that as prepayment becomes more expensive, the value of the default option rises and hence the sensitivity of MBS prices to house price movements rises. While the house-price durations are about an order of magnitude smaller in absolute value than the interest-rate durations, it is clear that focusing solely on interest-rate duration misses an important risk exposure.

The lower panels of Figure 5 examine the convexity at origination of our representative 8.5% WAC pool. The lower-left panel displays the convexity-yield curve. When transaction costs are zero, the convexity-yield curve of the mortgage pool resembles that of a callable bond: as interest rates fall and prepayment becomes more likely, convexity becomes highly negative, in line with the decline in duration above. At very low interest rates—below about 5%—convexity turns up, reflecting the inflection point in the duration-yield curve induced by its convergence to zero as interest rates fall to very low levels. At very high levels of transaction costs, the convexity-yield curve is monotonically upward-sloping but still negative, a reflection of the fact that higher prepayment transaction costs make the default option more valuable (recalling that we have fixed the default transaction cost to 5% of house value). The medium transaction cost calculations are somewhat surprising; at interest rates above about 5%, the convexity of the representative pool is above that of the zero transaction-cost case. At interest rates below about 5%, the convexity of the medium transaction-cost pool is well below that of the zero

[14] We have followed the market convention of multiplying duration by minus one to be consistent with the calculations of interest rate duration.

Figure 6 ■ Price sensitivity measures for two-factor model by LTV: 8.5% pool at time 0. *Note:* The figure reports a set of duration- and convexity-yield curves for a representative 30-year 8.5% coupon mortgage-backed security at origination ($t = 0$). The upper panel compares the duration–yield curves for mortgages with average LTV ratios of 80% and 100%. The lower panel displays a comparison of the convexity–yield curves for this mortgage.

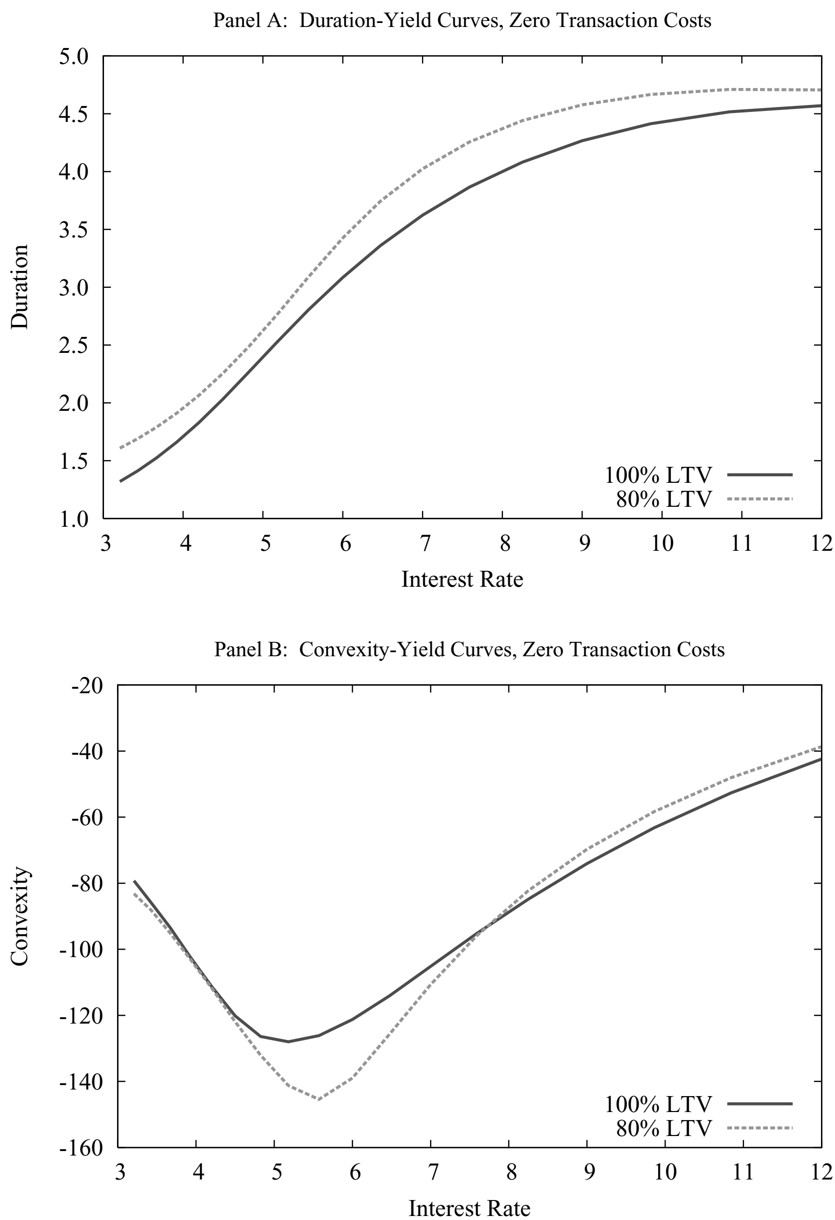

transaction-cost case. Why do the convexity-yield curves cross? Recall from the top left panel that the medium transaction-cost duration-yield curve is above the zero transaction-cost curve over most of its range, but note that for very low or very high interest rates, the curves are nearly identical. Hence, duration must fall more steeply for medium transaction costs as interest rates fall—in other words, the pool has greater convexity. The lower right panel examines the convexity-house price curves for the three mortgage pools. As can be seen, except at very low house prices, the convexity of the MBS in house prices is about two orders of magnitude less negative than its convexity in interest rates.

Nevertheless, movements in house prices and thus the value of the default options have important hedging implications. Figure 6 compares the duration- and convexity-yield curves for high- (100%) and low-LTV (80%) MBS. As before, we set the coupon to 8.5%, maturity is 30 years, and we consider the zero transaction cost sub-pool in each case. The duration-yield curve for the high-LTV pool lies everywhere below that for the low-LTV pool; as house prices fall and the default option moves closer to being in-the-money, interest-rate sensitivity falls. Put another way, because prepayment and default are substitutes, a rise in the value of the default option must be accompanied by a fall in the value of the prepayment option. The effect is greatest when interest rates are between about six and eight percent. In this range, interest-rate duration falls by about 0.5 years, with somewhat smaller declines for rates below 6% and 8%.

The duration-yield curves in Panel A of Figure 6 have different slopes, as reflected in the convexity-yield plots shown in Panel B. Because the low-LTV duration-yield curve rises more steeply for interest rates from about 4.5% to about 7%, the low-LTV convexity-yield curve lies below the high-LTV curve over this range. At very high interest rates, the prepayment option is nearly worthless, and so the duration-yield curves at both LTV levels approach one another. Hence, the convexity-yield curves are also nearly identical at high rates. Nevertheless, for a wide range of interest rates movements, house prices have important effects on *interest-rate* duration and convexity. Hence strategies of hedging against interest-rate risk that do not consider how the hedge ratios move with house prices remain exposed to an important source of risk.

Conclusions

This article develops a two-factor structural mortgage pricing model that treats both prepayment and default as the optimal response of mortgage borrowers to changes in interest rates and house prices. The model is estimated on a comprehensive data set of the termination behavior of the Freddie Mac mortgage pools backing Gold PCs issued from 1991 to 2002. Compared to a single-factor (interest-rate only) version of the model, the two-factor model produces significantly better fits of observed prepayment behavior and market prices. We conclude that our model offers a tractable framework for evaluating the effects of interest-rate and housing-price fundamentals on the valuation of mortgage-backed securities, and that modeling mortgage termination without including the effects of house-price changes will result in significant biases to both prices and hedge ratios.

The model is sufficiently flexible that it is suitable for both pool-level and loan-level mortgage applications. Furthermore, our analytical framework could also be extended to the valuation of commercial mortgage-backed securities, collateralized debt obligations and other securitization structures. All of these bonds have similar embedded prepayment and default options, and their performance is monitored using termination rates that are similar to those used in the mortgage-backed security market.

We would like to thank seminar participants at The University of North Carolina, Chapel Hill, Rice University, Columbia University, NBER Summer Institute, The Federal Reserve Board, Freddie Mac, Fannie Mae Foundation, Stockholm School of Economics, Université de Cergy Pointoise, Centre de Recherche THEMA, Institut National de la Statistique et des Études Économiques, Keio University, The University of British Columbia and the University of Southern California.

References

Amemiya, T. 1985. *Advanced Econometrics*. Harvard University Press: Cambridge, MA.

Archer, W.R., D.C. Ling and G.A. McGill. 1996. The Effect of Income and Collateral Constraints on Residential Mortgage Terminations. *Regional Science and Urban Economics* 26: 235–261.

Bartlett, W. 1989. *Mortgage-Backed Securities: Products, Analysis, Trading*. Prentice Hall: Englewood Cliffs, NJ.

Beckett, P. and H. Sender. 2002. Stormy Markets Foil Bets by Firms Based on Models. *Wall Street Journal*, September 27, 2002.

Becketti, S. and C.S. Morris. 1990. The Prepayment Experience of FNMA Mortgage-Backed Securities. Working Paper. New York University Salomon Center, New York.

Canner, G., K. Dynan and W. Passmore. 2002. Mortgage Refinancing in 2001 and Early 2002. *Federal Reserve Bulletin*, December: 469–481.

Caplin, A., C. Freeman and J. Tracy. 1993. Collateral Damage: How Refinancing Constraints Exacerbate Regional Recessions. Working Paper 4531. National Bureau of Economic Research, Cambridge, MA.

Cox, D.R. and D. Oakes. 1984. *Analysis of Survival Data*. Chapman and Hall: New York.

Cox, J.C., J.E. Ingersoll, Jr. and S.A. Ross. 1985. A Theory of the Term Structure of Interest Rates. *Econometrica* 53: 385–407.

Deng, Y. and J.M. Quigley. 2002. Woodhead Behavior and the Pricing of Residential Mortgages. Working Paper. University of Southern California, Los Angeles.

Deng, Y., J.M. Quigley and R. Van Order. 2000. Mortgage Terminations, Heterogeneity and the Exercise of Mortgage Options. *Econometrica* 68: 275–307.

Dunn, K.B. and J.J. McConnell. 1981a. A Comparison of Alternative Models for Pricing GNMA Mortgage-Backed Securities. *The Journal of Finance* 36: 471–483.

——. 1981b. Valuation of GNMA Mortgage-Backed Securities. *The Journal of Finance* 36: 599–616.

Dunn, K.B. and C.S. Spatt. 1986. The Effect of Refinancing Costs and Market Imperfections on the Optimal Call Strategy and the Pricing of Debt Contracts. Working Paper. Carnegie-Mellon University, Pittsburgh, PA.

Eom, Y.H., J. Helwege and J.-Z. Huang. 2004. Structural Models of Corporate Bond Pricing: An Empirical Analysis. *Review of Financial Studies* 17: 499–544.

Gourlay, A.R. and S. McKee. 1977. The Construction of Hopscotch Methods for Parabolic and Elliptic Equations in Two Space Dimensions with a Mixed Derivative. *Journal of Computational and Applied Mathematics* 3: 201–206.

Huang, J. and M. Huang. 2002. How Much of the Corporate-Treasury Yield Spread is Due to Credit Risk? A New Calibration Approach. Working Paper. Stanford University, Stanford, CA.

Johnston, E. and L. Van Drunen. 1988. Pricing Mortgage Pools with Heterogeneous Mortgagors: Empirical Evidence. Working Paper. University of Utah, Salt Lake City.

Kalbfleisch, J.D. and R.L. Prentice. 1980. *The Statistical Analysis of Failure Time Data*. Wiley: New York.

Kau, J.B. 1995. An Overview of the Option-Theoretic Pricing of Mortgages. *Journal of Housing Research* 6: 217–244.

Kau, J.B., D.C. Keenan, W.J. Muller, III and J.F. Epperson. 1992. A Generalized Valuation Model for Fixed-Rate Residential Mortgages. *Journal of Money, Credit and Banking* 24: 279–299.

——. 1995. The Valuation at Origination of Fixed Rate Mortgages with Default and Prepayment. *The Journal of Real Estate Finance and Economics* 11: 5–39.

Krishnamurthy, A., X. Gabaix and O. Vigneron. 2004. Limits of Arbitrage: Theory and Evidence from the Mortgage-Backed Securities Market. Working Paper. Northwestern University, Evanston, IL.

Longstaff, F.A. 2004. Borrower Credit and the Valuation of Mortgage-Backed Securities. Working Paper. University of California, Los Angeles.

Mattey, J. and N. Wallace. 1998. Housing Prices and the (In)stability of Mortgage Prepayment Models: Evidence from California. Working Papers in Applied Economic Theory 90-05. Federal Reserve Bank of San Francisco.

——. 2001. Housing Price Cycles and Prepayment Rates of U.S. Mortgage Pools. *The Journal of Real Estate Finance and Economics* 23: 161–184.

Mayer, C. and D. Genesove. 1997. Equity and Time to Sale in the Real Estate Market. *American Economic Review* 87: 255–269.

Merton, R.C. 1973. Theory of Rational Option Pricing. *Bell Journal of Economics and Management Science* 4: 141–183.

Monsen, G. 1992. The New Thinking on Prepayments. *Mortgage Banking* October: 48–56.

Pearson, N.D. and T.-S. Sun. 1989. A Test of the Cox, Ingersoll, Ross Model of the Term Structure of Interest Rates Using the Method of Maximum Likelihood. Working Paper. Massachusetts Institute of Technology, Cambridge, MA.

Richard, S.F. and R. Roll. 1989. Prepayments on Fixed Rate Mortgage-Backed Securities. *Journal of Portfolio Management* 15: 73–82.

Schwartz, E.S. and W.N. Torous. 1989. Prepayment and the Valuation of Mortgage-Backed Securities. *The Journal of Finance* 44: 375–392.

——. 1992. Prepayment, Default, and the Valuation of Mortgage Pass-Through Securities. *Journal of Business* 65: 221–239.

——. 1993. Mortgage Prepayment and Default Decisions: A Poisson Regression Approach. *AREUEA Journal* 21: 431–449.

Stanton, R. 1995. Rational Prepayment and the Value of Mortgage–Backed Securities. *Review of Financial Studies* 8: 677–708.

Stein, J.C. 1995. Prices and Trading Volume in the Housing Market: A Model with Downpayment Effects. *Quarterly Journal of Economics* 110: 379–406.

Timmis, G.C. 1985. Valuation of GNMA Mortgage-Backed Securities with Transaction Costs, Heterogeneous Households and Endogenously Generated Prepayment Rates. Working Paper. Carnegie-Mellon University, Pittsburgh, PA.