

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

Chris Downing, Richard Stanton and Nancy Wallace*

April 25, 2003, 12:01pm

Abstract

Mortgage-backed securities, with their relative structural simplicity and their lack of recovery rate uncertainty if default occurs, are particularly suitable for developing and testing risky debt valuation models. In this paper, we develop a two-factor structural mortgage pricing model in which rational mortgage-holders endogenously choose when to prepay and default subject to i. explicit frictions (transaction costs) payable when terminating their mortgages, ii. exogenous background terminations, and iii. a credit-related impact of the loan-to-value ratio (LTV) on prepayment. We estimate the model using pool-level mortgage termination data for Freddie Mac Participation Certificates, and find that the effect of the house price factor on the results is both statistically and economically significant. Out-of-sample estimates of MBS prices produce option adjusted spreads of between 5 and 25 basis points, well within quoted values for these securities.

*This paper represents the views of the authors and does not necessarily represent the views of the Federal Reserve System or members of its staff. Please address correspondence to (Downing): Federal Reserve Board, Mail Stop 89, Washington, DC 20551. Phone: (202) 452-2378. Fax: (202) 452-5296. E-Mail: cdowning@frb.gov. (Stanton): Haas School of Business, University of California at Berkeley, Berkeley, CA 94720. Phone: (510) 642-7382. Fax: (510) 643-1420. E-Mail: stanton@haas.berkeley.edu. (Wallace): Haas School of Business, University of California at Berkeley, Berkeley, CA 94720. Phone: (510) 642-4732. Fax: (510) 643-1420. E-Mail: wallace@haas.berkeley.edu. The authors 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

1 Introduction

The residential mortgage-backed security (MBS) market is one of the largest and fastest growing bond markets in the United States.¹ Not only are MBS thus worth studying in their own right, but the easy availability of data, their simple structure compared with many other forms of debt, and the lack of uncertainty regarding recovery rates in the event of default make them attractive for studying and testing models of defaultable debt pricing more generally.

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 remain. These challenges include the persistence of model-based MBS pricing errors (option adjusted spread, or OAS), the observed variance in bids for MBS derivative auctions (Bernardo and Cornell, 1997),² and the need for improved hedging strategies demonstrated by Fannie Mae's recent problems with a large negative duration gap in its portfolio.³ Furthermore, for firms that participate in the MBS market, recent changes in accounting standards have significantly increased the need to hedge MBS and to account for hedging effectiveness directly on the income statement.⁴

Many of these problems stem, at least in part, from the widespread use of reduced form models to obtain measures such as mortgage prepayment and default rates. Reduced form models are very convenient to work with, since they are not constrained by anything more than the modeler's choice of functional form. This flexibility makes reduced form models both simple to implement and able to do a good job of describing historical data. Typically, however, prepayment models are used to answer questions about what might happen in the future, rather than just to describe historical data. A functional form that works well when applied to historical data may fail when applied to different data, as many mortgage investors are discovering to their cost today.

¹By the end of the third quarter 2001, the outstanding stock of residential mortgage backed securities in the U.S. was \$3.5 trillion and the year-end growth rate was 12.9 % (Inside ABS and MBS, January, 2002). As of the end of April, 2002, the value of Lehman MBS index exceeded that of the Lehman Treasury index by 15% (Jay and Anderson (2002)).

²Bernardo and Cornell found that the percentage range by which high bids exceeded low bids was 63% in a \$90.6 million auction for agency collateralized mortgage obligations.

³The *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 of the borrowing it has done to finance them. Fannie Mae's stock is down about 11% since it announced the news." (WSJ, 9/27/02, "Stormy Markets Foil Bets by Firms Based on Models", Paul Beckett and Henny Sender.

⁴FAS 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.

A second problem is that most of the current models used for pricing and hedging MBS are based solely on interest rates. While interest rates are generally acknowledged to be the most important factor affecting MBS prices and terminations, there is substantial evidence that other factors may play a significant role. For example, Boudoukh, Richardson, Stanton, and Whitelaw (1995) show that even after doing as good a job as possible hedging MBS using Treasury securities, a sizeable fraction of the original volatility remains. This suggests that there may be one or more omitted variables.⁵ A leading candidate for such an omitted variable is the level of house prices. Not only do housing prices affect the level of default activity, but a low house price also affects the borrower’s ability to qualify for a new loan, and hence may also affect refinancing behavior.⁶

In this paper, we develop and empirically test a mortgage valuation model that avoids both of these problems. First, it is a structural model that explicitly values the prepayment and default options embedded in a mortgage contract as endogenous functions of the underlying sources of uncertainty, rather than using proxies for these options and treating them as exogenous to the valuation (as is typical with reduced form models). The advantage of a structural approach is that it should produce sensible results even when we encounter an economic environment unlike anything we have seen in the past, since the underlying factors are linked to predicted behavior through the optimizing behavior of agents in the model. Second, it is a two-factor model that allows both interest rates and house prices to affect mortgage termination and valuation.

Unlike most prior option-based models (with the exception of Stanton (1995, 1996), Kau and Slawson (2002), and Deng, Quigley, and Van Order (2000)), our model also directly accounts for transaction-cost heterogeneity among mortgage holders. Another important feature of the model is that, in addition to allowing prepayment and default to occur endogenously, we also introduce a channel through which the effects of exogenous background terminations and other transaction-cost related frictions may impede optimal option-exercise strategies. The results of our empirical implementation indicate that the model generates accurate estimates of the pool-level termination rates and reasonable estimates of MBS prices and price sensitivity. Our average pricing errors in a true out-of-sample test against trad-

⁵An alternative potential explanation is that everything *is* driven only by interest rates, but that mortgages cannot be hedged using Treasury securities alone due to the presence of unspanned stochastic volatility (See Collin-Dufresne and Goldstein (2001)). However Boudoukh et al. (1995) show that interest rate volatility does not help to explain much of the residual volatility of MBS prices.

⁶A number of prior empirical studies have documented the importance of home equity on the propensity to re-finance (See, Beckett and Morris (1990), Monsen (1992), Caplin, Freeman, and Tracy (1993)). Stein (1995), Archer, Ling, and McGill (1996), Mayer and Genesove (1997), Matthey and Wallace (1998), and Matthey and Wallace (2001) also emphasize the importance of housing prices as a determinant of regional-level household mobility.

ing prices for MBS forward contracts are about 3.32 percent.⁷ Our model is well suited to applications based on widely available aggregate pool-level performance data and allows us to partition realized aggregate cash flows into the principal dollars that are received as the result of borrower defaults and the principal dollars that are received as the result of prepayment.⁸

The paper is organized as follows. The next section reviews the existing valuation literature. Section 3 introduces the pool-level default/prepayment pricing model. The solution to the pricing model and the discussion of the estimation strategy is presented in Section 4. Section 5 describes the Freddie Mac data used in estimating the model, and Section 6 contains a discussion of our results. Section 7 concludes.

2 Pricing Risky Debt

Two quantitative modeling strategies have emerged from the extensive literature focusing on the valuation and hedging of risky bonds: structural and reduced-form models. Structural models focus on pricing and modeling credit or call events that are specific to a particular class of borrower. These models focus on the underlying dynamics of the asset that is the collateral on the bond. Credit or call events are triggered by random movements in the asset price relative to some threshold (that may be endogenously determined). By modeling credit or call events in terms of the underlying dynamics of asset prices and/or interest rates and the optimizing behavior of agents in the model, the structural methodology functionally links option exercise events to the underlying fundamentals faced by the borrower. Recent examples of structural models of corporate risky debt include Longstaff and Schwartz (1995), Jarrow and Turnbull (1995), Leland (1994), Leland and Toft (1996), Anderson and Sundaresan (2000), Collin-Dufresne and Goldstein (2001), and Huang and Huang (2002). Examples of structural models of MBS risky debt include Dunn and McConnell (1981a,b), Timmis (1985), Dunn and Spatt (1986), Johnston and Van Drunen (1988), and Stanton (1995). Several recent papers consider both default and prepayment. Kau, Keenan, Muller, and Epperson (1992), Kau, Keenan, Muller, and Epperson (1995), and Kau (1995), lay out a structural two-factor option-pricing model for prepayment and default. However, they perform no empirical estimation or testing of their model.

⁷Eom, Helwege, and Huang (2002) find that pricing accuracy is a problem for most current structural bond valuation models. Their in-sample test results for the pricing errors generated by various structural valuation models range between plus 0.50% and minus 10%. Wall Street mortgage valuation models observe the same pricing discrepancies, and practitioners add an additional spread to the interest rate (“Option Adjusted Spread”, or OAS) to bring model prices in line with market prices.

⁸Unfortunately, at this time we have no way to determine the accuracy of these estimated cash flow partitions. The model could also be applied to loan-level data if they were available.

In the reduced-form approach, the value of the borrower's assets and capital structure are not modeled explicitly. Instead, credit and/or call events are modeled as an exogenously specified jump process or hazard rate. These approaches emphasize empirical estimation of the random timing of option exercise events, and then evaluate the conditional expectations under a risk-neutral probability of functions of event timing and realized cash flows. Recent examples of the reduced-form approach to the valuation of corporate risky debt include Jarrow and Turnbull (1995), Duffee (1998), Duffie and Singleton (1999), Collin-Dufresne and Solnik (2001), and Duffie and Lando (2001). Reduced form mortgage valuation models include Schwartz and Torous (1989), Deng, Quigley, and Van Order (2000), and Deng and Quigley (2002).

Despite the fact that corporate bonds and MBS are just different types of risky debt, the two literatures have developed almost independently of each other. In many ways, the characteristics of mortgage-backed securities make these instruments particularly suitable for developing and testing general models of risky debt valuation. First, unlike many other forms of defaultable debt, there are very extensive time series available on the prepayment and default experience of MBS. Second, the structure of mortgage-backed securities is relatively simple; there is only one type of bond in the MBS pass-through structure, so issues related to the seniority of the bonds do not apply. This stands in contrast to many other structured finance vehicles, such as collateralized debt obligations (CDOs), which typically have complicated seniority structures, or corporate bonds, whose values depend upon the entire capital structure of the firm. Finally, when mortgage defaults are triggered, MBS are guaranteed 100% recovery rates on both interest and principal. In contrast, corporate bond recovery rates, for example, vary both in the cross-section and over time. The relative simplicity of residential mortgage backed securities allows us to focus on the performance characteristics of the valuation model, without the typical concerns that an observed empirical rejection is caused by some detail of the specific securities being analyzed.

An important further advantage of our focus on MBS valuation is that there are pricing data available for mortgage pools with a wide range of coupon rates. These data are available from the forward contract market and provide a good benchmark for the evaluation of model based pricing performance. Recent advances in the development of state-level house price indexes also afford measures of the evolution of loan-to-value ratios for the residential real estate assets that are the collateral for the mortgages in MBS pools. These indexes and conventional methods of modeling the default-free term structure allow us to estimate the underlying exogenous factors driving the exercise of the prepayment and default options embedded in mortgages. Identifying a tractable set of underlying factors in the corporate bond market is a considerably greater challenge and the evaluation of default and prepayment

is further complicated by the presence of various bond covenants, contractual constraints, and the complicated capital structure of most corporations. For these reasons, a structural MBS valuation model may also provide important general insights into the valuation of defaultable callable bonds in a multi-factor setting.

3 The Model

3.1 Valuing the Mortgage and its Embedded Options

A fixed rate home mortgage is a callable, defaultable bond whose payments (usually 360 equal monthly payments) are made by an individual borrower to a bank or other financial institution. Although many different mortgage types exist, we shall focus on 30 year, fixed rate mortgages, the loans backing most MBS. We shall use the notation B_t to denote the market value of the remaining scheduled payments in the absence of any options, and refer to this stream of payments as the “underlying bond”. Valuing a mortgage amounts to valuing this bond together with its embedded options, which we now describe.

Prepayment Option A mortgage borrower with a fixed rate loan possesses a call option on the underlying bond. At any time after taking out the loan, the borrower may choose to stop making the remaining scheduled payments, and instead pay off the remaining principal amount on the loan, F_t , calculated using the standard annuity formula,

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

where C is the monthly payment amount, c is the contractual interest rate on the loan, and $T - t$ is the remaining time on the loan.⁹ Paying off the loan is equivalent to exercising a call option on the underlying bond, B_t , with (time-varying) exercise price F_t . The lower current interest rates are, the higher is B_t , and hence the more in-the-money is the call option (equivalently, the lower are current interest rates, the lower the rate on a new mortgage compared to that on the existing loan). When rates get low enough — exactly how low is determined as part of the valuation — it becomes optimal for the borrower to exercise the call option and pay off the loan early.¹⁰

⁹The monthly payment, C , is calculated using the same formula, setting t to 0, and F_0 to the initial loan amount.

¹⁰We know it is never optimal to exercise an American call option on a non-dividend paying asset before maturity. Early exercise of the mortgage prepayment option *is* optimal at low enough interest rates since the underlying asset makes a payment each month.

Besides paying off the loan when interest rates fall sufficiently, the borrower will also pay off the loan early should he or she move. This needs to be taken into account because it affects the value of the option both directly and indirectly. The direct effect is that the option may now be exercised when a decision based solely on interest rates would say it should not be. The indirect effect is through a change in the optimal exercise policy. For example, if the homeowner knows that he or she is likely to move, and hence may exercise the option “suboptimally” some time in the future, then it is less attractive to keep an option alive and relatively more attractive to exercise it immediately if it is in-the-money.

Default Option 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 an option to exchange one asset (the house) for another (the remaining cash flows on the mortgage), which we can regard either as a put option on the house, or equivalently as another call option on B_t , this time with (stochastic) exercise price H_t . The lower H_t , the more attractive is exercise of the default option. It is convenient to write the value of the mortgage liability as

$$M_t^l = B_t - V_{p,t}^l - V_{d,t}^l, \quad (1)$$

where $V_{p,t}^l$ is the value of the refinancing option to the mortgage holder, and $V_{d,t}^l$ is the value of the default option to the mortgage holder.¹¹ Borrowers choose when they prepay or default in order to minimize the overall value M_t^l . In the absence of frictions or transaction costs, this problem is similar to that studied by Dunn and McConnell (1981a,b), but with the addition of the default option. If we specify some model for interest rates (and house prices), we can value the mortgage and its embedded options in much the same way as we would use a binomial tree to value an American put option: Each period, first discount the value of the security back one period to determine its value if none of the embedded options are exercised, and then decide whether it is preferable to exercise any of the options.

3.2 Frictions and Transaction Costs

The problem with the model so far is that it describes a borrower who does not exist: one who exercises the options in his or her mortgage the moment it becomes optimal to do so, and who faces no transaction costs on doing so. This is the world described by Dunn and

¹¹Note that the values of the two options are not really independent of each other, since exercising one option precludes exercise of the other.

McConnell (1981a,b), but it has some counter-factual implications. First, in such a model, a mortgage (or MBS) can never trade above par, since this would automatically imply that borrowers should already have refinanced. Second, since all borrowers in this world are identical, if any one borrower finds it optimal to refinance, so should all the others, hence we should see all borrowers prepaying simultaneously.

In reality, borrowers do not all act at the same time, and they face costs whenever they decide to refinance or default. We therefore assume (similar to Stanton (1995)) that:

1. Borrowers face transaction costs whenever they refinance or default
2. Borrowers do not necessarily refinance/default immediately once it becomes optimal to do so, but instead act with some probability in any given period of time
3. Borrowers are heterogeneous — different borrowers face different transaction costs.

Transaction Costs There are many costs associated with prepayment, both explicit (in the form of direct transaction costs), and implicit (representing the difficulty of deciding whether or not to prepay, the cost of the time taken to complete the process, etc.). 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

$$M_t^l \geq F(t)(1 + X_p). \quad (2)$$

In the absence of frictions, the borrower will optimally default if the value of the mortgage is greater than or equal to the value of the house. Again, there are many costs associated with default, such as the value of the lost credit rating, which we capture via another proportional transaction cost, X_d , payable by the borrower at the time of default. Default is optimal for the borrower if

$$M_t^l \geq H_t(1 + X_d). \quad (3)$$

Stochastic Option Exercise As in Stanton (1995), we assume that borrowers make exercise decisions only at random discrete time intervals,¹² governed by a hazard function. Informally, if the hazard function governing some event is λ , then the probability of the event's occurring in a time interval of length δt , conditional on not having occurred prior to t , is approximately $\lambda \delta t$.¹³ In addition, borrowers may be forced to prepay or default for nonfinancial reasons (such as divorce, job relocation, or sale of the house), which we assume

¹²This could result, for example, from mortgage holders having to pay some cost to make each decision.

¹³Hazard functions are discussed in detail in Kalbfleisch and Prentice (1980), and Cox and Oakes (1984).

is also described by some hazard function. This results in the probability of prepayment or default in any time interval being governed by some (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

$$\lambda(t) = \lambda_0(t) + \beta_2 D_t + \beta_3 LTV_t P_t, \text{ where} \quad (4)$$

$$\lambda_0(t) = \beta_0 \arctan(t/\beta_1). \quad (5)$$

Here, LTV is the loan-to-value ratio, the dummy variable D is one when default or prepayment is optimal, and zero otherwise, and the dummy variable P is one when prepayment is optimal, and zero otherwise. We refer to $\lambda_0(t)$ as the background hazard.¹⁴ The background termination rates rise over time at a rate governed by β_1 , to a maximum dictated by the value of β_0 . In the prepayment and default regions, the maximum hazard rate shifts up by the amount β_2 , and depending on the sign of β_3 , the maximum hazard rate in the prepayment region will rise or fall with changes in LTV .

Borrower Heterogeneity We have so far described the problem facing an individual mortgage holder. Different borrowers will in general face different frictions, so we assume that the costs X_p and X_d are randomly distributed across borrowers in a pool. For tractability we assume that, for each borrower, $X_p = X_d$, and that, as in Stanton (1995), these costs are distributed according to a beta distribution with parameters β_4 and β_5 .¹⁵

4 Implementing the Model

To implement the model and calculate mortgage values, we need to make assumptions about the underlying state variables. We assume there are two state variables, the short term riskless interest rate, r_t , and the house price, H_t .

¹⁴This functional form captures the idea of “seasoning” (see, for example, Richard and Roll (1989)), where new loans prepay slower than older loans, holding all else constant.

¹⁵This distribution is chosen because it has many possible shapes, and is bounded by zero and 1. Its mean and variance are

$$\begin{aligned} \mu &= \frac{\beta_4}{\beta_4 + \beta_5} \\ \sigma^2 &= \frac{\beta_4 \beta_5}{(\beta_4 + \beta_5)^2 (\beta_4 + \beta_5 + 1)} \end{aligned}$$

Interest Rates We assume interest rates are governed by the Cox, Ingersoll, and Ross (1985) model,¹⁶

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

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:

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

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

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

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:

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

We calibrate equation (8) as follows:

$$\begin{aligned} q_H &= 0.025 \\ \phi_H &= 0.085. \end{aligned}$$

The value of q_H is 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.¹⁷

Given these models for interest rates and house prices, standard arguments show that, in

¹⁶This model is widely used in the mortgage pricing literature. See, for example, Stanton (1995), and Kau, Keenan, Muller, and Epperson (1992).

¹⁷This 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.

the absence of arbitrage, the value of a mortgage $M^l(H_t, r_t, t)$ paying coupon C must satisfy the partial differential equation:

$$\frac{1}{2}\phi_r^2 r M_{rr}^l + \frac{1}{2}\phi_H^2 H^2 M_{HH}^l + (\kappa(\theta_r - r) - \eta r) M_r^l + ((r - q_H)H_t) M_H^l + M_t^l - rM^l + C = 0, \quad (9)$$

subject to the boundary conditions:

$$M^l(0, r, t) = 0, \quad (10)$$

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

$$M^l(H, r, T) = 0, \quad (12)$$

Boundary condition (10) says that, when house prices are zero, the mortgage holder always defaults; because zero is an absorbing barrier for the house price process, the mortgage is therefore worthless. Boundary condition (11) arises because all future payments are worthless when interest rates approach infinity. Equation (12) reflects the amortization of the mortgage.

Valuation and Option Exercise To value a given mortgage and determine its optimal prepayment/default behavior, we need to solve the system of equations (9)-(12). Details of the numerical technique used are provided in appendix A. The basic approach is

1. Pick a set of discrete transaction costs, with associated initial proportions, to approximate the continuous (beta) distribution of costs described above. The exact procedure for doing this is described in Stanton (1995)
2. For each transaction cost, solve for the mortgage value on a grid of r , H and t values, solving in the process for the hazard rate governing prepayment/default for each borrower at each possible interest rate/house price combination
3. Now work forward in time, starting at the issue date of the mortgage
 - (a) Each period, given the observed interest rate and house price, and current proportions of each borrower type in the pool, the current MBS value is a weighted average of the values for each individual type, which can be read off the valuation grids calculated in step 2
 - (b) We can also read off the expected proportion of each borrower type that prepays this period. Use this i. to calculate the expected prepayment rate for the pool as a whole; and ii. to update the proportions of each type remaining in the pool for next period.

To aid in understanding how the various components of the model operate, as well as to provide some insights on the identification of the parameters of the model, figure 1 compares the observed and predicted patterns of mortgage terminations for pools with weighted average coupons of 8.25 percent when different components of the model are added sequentially. The vertical axes show the termination rates, expressed in terms of the proportion of the pool remaining that terminates in a given month — the single month mortality rate. The details of the parameterization and the dataset from which the observed terminations are extracted are both discussed at length in the next section. The upper left panel displays the model when the hazard function and the transaction cost distribution are both disabled; we refer to this as the “baseline” model. As is well known in the literature on mortgage backed securities pricing, a structural model without transaction costs produces termination behavior that does not match observed behavior because either an entire mortgage pool is predicted to terminate (when prepayment or default are optimal), or terminations are zero (when continuation is optimal).

The upper right panel of figure 1 displays the predictions of the model when the transaction cost distribution is enabled, but the overall hazard (equation (4)) remains disabled. In this model, when interest rates or house prices fall, terminations rise, but only for the sub-pools of mortgages for which transactions costs are relatively low. This holds down the overall termination rate to more reasonable levels. Because there are no background terminations, the level of terminations is too low when termination is not optimal. The model generates a pattern of terminations that is far less smooth than the observed termination patterns, but this is to be expected since the model is generating expected termination rates; actual termination rates will always fluctuate randomly around the values predicted by the model.

The lower left panel displays the predictions of the model when the background hazard function is enabled only in the continuation region of the state space — the region where neither prepayment nor default is optimal. When prepayment and default are optimal, we allow termination to be certain by using the following specification for the hazard:

$$\lambda(t) = \begin{cases} \beta_0 \arctan(t/\beta_1) & \text{If continuation is optimal} \\ \infty & \text{If either prepayment or default is optimal} \end{cases} \quad (13)$$

The panel thus shows the incremental improvement in fit that is achieved with a background hazard rate. Since the hazard rate is infinite in the prepayment and default regions, the probability that the mortgage holder terminates in these regions is one. The contribution of the background hazard to the fit is most evident in months 40-60, where the average termination rates rise toward the observed rates, and when both the baseline model and the

model with only the embedded transaction cost distribution produce termination rates that are too low.

The lower right panel displays the fit of the full model, using the hazard specification in (4). In general, this model produces the best overall fit because it is for this model that the parameters were optimized. The addition of the hazard rate in the prepayment and default regions serves to smooth the path of terminations. After month 40, terminations have risen slightly compared to the model that uses only the background hazard. This is a reflection of the slowdown in prepayment and default terminations for some pools under the full hazard model — more principal remains in these pools for a longer period of time, holding up termination rates as interest rates and/or house prices fall over time.

Compared to the trajectory of predicted terminations shown at the lower left, terminations have risen early on — a surprising result given that the primary effect of the hazard function in the prepayment and default regions is to slow down terminations. However, this illustrates a secondary effect of the hazard function. Conditional on being in the prepayment or default region, the hazard rate lowers the rate of terminations. However, it is also the case that the imposition of the hazard function shifts the prepayment and default boundaries. This shift occurs because the hazard lowers the probability of terminations in the future. In other words, the hazard lowers the time value of waiting to prepay or default, thus raising current terminations. As shown at the lower right, for this sample the average net effect is to raise terminations early on.

These results highlight an important distinction between our modeling approach and the reduced-form approaches of, for example, Schwartz and Torous (1989) and Deng, Quigley, and Van Order (2000). Reduced-form models of prepayment typically focus on estimating the probability of mortgage termination, and usually include ad-hoc, exogenously determined measures of the extent to which the prepayment and/or default options are in the money. This approach precludes pricing mortgages in an internally consistent way. As the results above indicate, the probability of termination and the values of the prepayment and default options (and hence the price of the mortgage) are endogenously determined. When the probability of termination is modified with the hazard model, the prepayment and default boundaries shift, affecting the timing of default and prepayment, the values of the options, and the value of the mortgage.

4.1 Estimating the Model

As in Stanton (1996), we use a non-linear 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, and we pick parameters for the model based on making the predicted and observed termination rates as close as possible.

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:

$$\chi^2(\Theta) = \sum_{i=1}^N \sum_{t=1}^{T_i} (\omega_{it} - \hat{\omega}_{it}(\Theta))^2 \quad (14)$$

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 $\chi^2(\Theta)$.¹⁸

5 The Freddie Mac Mortgage Pools

Our empirical analysis focuses on the prepayment characteristics of Freddie Mac pass-through residential mortgage-backed securities. The data for this study consists of all Gold Participation Certificate (Gold PC) pools issued by Freddie Mac between January, 1991 and December, 1994. 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.¹⁹ To ease our computational burden, we select all pools with weighted-average coupon exactly fifty basis points above the pass-through coupon rate. Furthermore, our data set contains information on the geographic distribution across the states of the underlying loans in each pool at the time the pool was originated. We use this information to select pools with at least seventy five percent of the

¹⁸Evaluating $\chi^2(\Theta)$ presents a significant computational challenge. A single evaluation of the objective function requires that we solve the mortgage valuation problem for a set of discrete transaction costs at each coupon level. Thus, if there are C distinct coupon levels, N_{trans} discrete transaction costs, and we are solving on a finite difference grid of N_{grid}^2 points, the number of computations will be

$$O(C \times N_{trans} \times N_{grid}^2 \times 360 \times 20),$$

where we have assumed 20 operations at each grid point on the solution grid, a conservative estimate. For our dataset, discussed below, we have $C = 11$, and we set $N_{het} = 50$, $N_{grid} = 60$, producing an operations count that is $O(1.4E + 10)$. To tackle a problem of this magnitude, we implemented our finite difference solution algorithm using Scalapack, a library of scalable parallel linear algebra routines. Optimization was carried out on a network of workstations.

¹⁹Specifically, we subset to pools with 30 year Gold Participation Certificates, an original weighted average loan age of two months or less, and an original weighted average remaining maturity of 350 months or more.

mortgages concentrated in a single state. This leaves us with 1,340 mortgage pools and these pools account for 17,665 mortgages.

We observe the mortgage termination rates for the pools from the month of issuance through June, 2001.²⁰ The factor history of the earliest-issued pools are observed for 120 months, and the latest-issued pools are tracked for 78 months. Table 1 presents the aggregate properties of the state-concentrated pools organized by year of origination and weighted average coupon (WAC). The WACs of the pools drifted down to 8.50 percent during 1992 and continued falling to a low of 7.25 percent at the end of 1993 and beginning of 1994. They then moved back up to 8.75 percent by the end of that year. The highest weighted average coupons, 9.75 percent, were originated in the early part of 1991. By the end of 1992 the WACs had fallen to 9.00 percent. These high coupon pools tend to prepay rapidly and the shortest pool payout period observed was twenty three months.

The full-sample average termination rates from pool inception to June, 1998 was 1.05 percent per month.²¹ On average, the prepayment history of the pools was followed for about 102 months. As shown in table 1, the average termination rates rise to about 3 percent of pool principal per month for the high coupon pools originated in late 1991 through early 1993. In contrast, the low coupon pools that were originated in late 1993 and early 1994 have markedly lower average termination rates of less than 1 percent per month.

We measure housing returns using the Office of Federal Housing Enterprise Oversight (OFHEO) quarterly repeat sales home price indices, computed at the state level. These series are interpolated to a monthly frequency using splines. As shown in table 1, the average change in the house price index was .12 percent per month. There is, however, considerable time series variation in the house price index dynamics over the various vintages and weighted average coupons of the pools. The 1991 vintage pools tend to have lower average percentage changes in the house price indexes than the more recent vintage pools.

We identify weak house price pools using a measure of the probability that the current loan-to-value ratio exceeds one (Deng, Quigley, and Van Order (2000)), denoted by $\text{Prob}(LTV(HP_t) > 1)$.²² For most of the months during which we observe prepayment his-

²⁰The mortgage factor is the percentage of the outstanding balance that remains in the pool at the end of each month.

²¹In computing pool-level termination rates, we follow the formula on page 205 of Bartlett (1989) for estimating terminations given data on pool factors, time to maturity, and coupon interest rates.

²²Our measure is constructed by rolling forward the denominator of the initial loan-to-value ratio of the loans in the pool by the sequence of observed changes in a house price index for their geographic area and extrapolating the numerator by the scheduled decline in the principal balance of the loan. The value of the central tendency of the current loan to value ratio is converted to a probability by assuming that the distribution of logarithmic changes in house prices is normally distributed. The variance of this normal distribution is estimated from the moments of the underlying individual home price data from which the aggregate home price indices are constructed.

tories, the value of $\text{Prob}(LTV(HP_t) > 1)$ is below 1.50 percent. Only about one-quarter of the pool-month observations have a $\text{Prob}(LTV(HP_t) > 1)$ value that exceeds 3.50 percent. However, a small portion (about one-tenth) of the pools experience realizations of our measure above 8 percent. Looking at the distribution of $\text{Prob}(LTV(HP_t) > 1)$ by pool, 162 of the 1,340 pools experience realizations of this variable above 20 percent at some point in their observed prepayment history. Many of these weak house price pools have at least 75 percent of their loans in either California or Hawaii. California experienced a sharp decline in home prices beginning in 1990 and extending through 1996, and the highest proportions of estimated negative equity are concentrated in 1995 and 1996, for pools that were originated in 1991.

In figure 2 we present graphs of the average single month mortality rates for weak²³ and non-weak house price pools using only pools with WACs greater than 9 percent.²⁴ These pools were originated in 1991 and the beginning of 1992 and thus allow a partial control for the effects of interest rate motivated terminations. As shown, the weak house price pools consistently terminate more slowly than the non-weak pools and always lag, by several months, the increases in termination rates experienced by the non-weak pools. This marked dampening of termination rates for high WAC pools, which should have paid off rapidly as interest rates declined, is remarkably coincident with the underlying dynamics of the house price returns for the weak pools as shown in figure 3.

The observed termination rates reported in figure 2 suggest that total terminations for the weak pools declined overall. These observed declines could arise as the result of the pure substitution effect between borrowers' optimal exercise strategies for the embedded prepayment and default options. As house prices decline, the expected value of the competing default option rises and expected value of the prepayment option decreases, thus leading to slower total terminations. Another explanation could be that decreasing house prices reduce mobility and/or increase prepayment transactions costs arising from credit constraints or underwriting requirements that are defined in terms of loan-to-value ratios.

6 Estimation Results

Table 2 reports two sets of parameter estimates. The first column reports parameter estimates for the overall hazard (equation (4)), excluding the effects of the loan-to-value ratio (β_3), and for the location parameters of the transaction cost distribution (β_4 and β_5), assum-

²³These are pools where the $\text{Prob}(LTV(HP_t) > 1)$ is above 20 percent at some point in the observed prepayment history of the pool.

²⁴From table 1, we see that there are 411 of these pools and a substantial fraction of these are weak housing price pools.

ing a one factor model in which the riskless interest rate is the sole underlying state variable, and prepayment is the only channel for mortgage termination. The second column of Table 2 reports the estimated parameters of the overall hazard (equation (4)) and transaction cost distribution, for the two-factor model. In the two-factor model the state variables are the riskless interest rate and house prices and there are two channels for terminations: default and prepayment. As shown in the table, the seasoning dimension of the overall hazard (β_0 and β_1) is not precisely estimated in either specification. This result suggests that “ramp-up” in prepayments over the initial life of a Freddie Mac PC pool is not an important feature of termination behavior. The positive and precisely estimated parameter β_2 indicates that the hazard rates of principal return are statistically significantly increased when the prepayment option is in-the-money. In the two-factor model, the effect of LTV, β_3 , on the overall hazard rate is positive and statistically significant in the prepayment region. The positive coefficient estimate indicates that higher LTVs tend to increase the hazard rate if the prepayment option is in-the-money, perhaps because a higher LTV induces borrowers to cash out of the housing market or perhaps to prepay in response to labor market objectives.

An important difference between the parameter estimates for the one- and two-factor specifications is in the parameters β_4 and β_5 determining the distribution of transaction costs in the borrower pool. Both parameters are precisely estimated in the one and two-factor specifications, but the mean of the distribution is somewhat lower in the two-factor model. Figure 4 shows the estimated distribution of transaction costs for the one-factor and two-factor models and contrasts these results with the estimated distribution reported in Stanton (1995). The results suggest that the bulk of borrowers in the pools face relatively small transactions cost - on the order of 11.5 percent of face value or house value - in contrast to the results in Stanton (1995) shown by the light dashed line. This shift from the Stanton (1995) results is consistent with the well documented increase in overall prepayment speeds in the last ten years.²⁵ The inclusion of the house price factor shifts the distribution even further to the left as more of the terminations are accounted for via the drifts and volatilities of the factors, reducing the average transaction cost required to fit the data.

Table 3 reports the results of two-way random effects regressions of the residuals from the one- and two-factor models on the levels and squared levels of house prices, and the level and squared level of the ten-year Treasury rate. The first column of the table shows that the coefficient estimates for the level and squared level of both the ten-year Treasury rate and house prices are statistically significantly different from zero. When house prices have risen (fallen) a lot since mortgage origination, actual termination rates tend to be larger (smaller)

²⁵Market analysts report that recent prepayment speeds are about twice as high as they were in 1992 in interest rate environments where the prepayment options should be exercised (Jay and Anderson (2002)).

than predicted termination rates, although the termination levels increase at a decreasing rate. The effect of the ten-year Treasury rate is similar. As shown in the second column, the residuals from the two-factor model are uncorrelated with the level and squared level of house prices. These results suggest that house prices are an important missing factor in the one-factor mortgage pricing model. The two-factor residuals still retain some dependence on the level and squared level of interest rates.

Table 4 reports the results of our specification tests for the one- and two-factor models. The table displays the results for J-Tests of non-nested model specification (Davidson and MacKinnon (1981)).²⁶ The first line shows the results for the test:

$$\begin{aligned} H_0 : y &= f(r; \Theta_1) + \epsilon_0 \\ H_1 : y &= g(r, H; \Theta_2) + \epsilon_1, \end{aligned}$$

where y is the observed termination rate for pool i at time t , $f(r; \Theta_1)$ is the predicted monthly termination rate under the 1-factor model defined by coefficients Θ_1 , and $g(r, H; \Theta_2)$ gives the predicted monthly termination rate using the 2-factor model defined by coefficients Θ_2 .

Following Davidson and MacKinnon (1981), the test is performed using ordinary least squares to estimate the equation:

$$y - \hat{f} = \alpha(\hat{g} - \hat{f}) + \hat{F}b_1 + \epsilon,$$

where \hat{f} are the predicted values under the 1-factor model at $\hat{\Theta}_1$, \hat{g} are the predicted values under the 2-factor model at $\hat{\Theta}_2$, \hat{F} is the vector of derivatives of f with respect to the elements of Θ_1 , and b_1 is a vector of nuisance parameters (not reported). A statistically significant value of α indicates that the null hypothesis is rejected.

The second line of the table shows the results for the test where the roles of f and g are reversed:

$$\begin{aligned} H_0^* : y &= g(r, H; \Theta_2) + \epsilon_0^* \\ H_1^* : y &= f(r; \Theta_1) + \epsilon_1^*, \end{aligned}$$

This test is performed using ordinary least squares to estimate the equation:

$$y - \hat{g} = \alpha^*(\hat{f} - \hat{g}) + \hat{G}b_2 + \epsilon^*,$$

²⁶Singleton (1985) introduces a test for non-nested alternative specifications with a form that is similar to the J-Test. The Singleton test statistic uses sample cross-products of residuals and instruments in a GMM setting. Given our nonlinear least squares estimation strategy, we follow Davidson and MacKinnon (1981) and use the residuals themselves.

where \hat{G} is the vector of derivatives of g with respect to the elements of Θ_2 , and b_2 is a vector of nuisance parameters (not reported). As before, an estimate of α^* that is statistically significant rejects H_0^* in favor of H_1^* .

The first row of table 4 shows that we can reject the null hypothesis H_0 in favor of the two-factor model, since α is statistically significantly different from zero. However, the second row of the table shows that α^* is also statistically significantly different from zero, indicating that we also reject the null hypothesis H_0^* in favor of the one-factor model. As is often the case with the J-test, the low power of the test does not allow us to definitively reject the one-factor model in favor of the two-factor model. Nevertheless, the combined evidence from the residual analysis and the J-test suggests that the one-factor model is misspecified because it omits controls for the effects of house prices on the termination options.

Figures 5–8 translate the coefficient estimates for the model into average predicted monthly termination rates. The figures show 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.²⁷ The full sample averages reported in figure 5 appear to track the average time series of terminations quite accurately, although prepayment rates do not rise as high as observed levels in the first twenty months. Figure 5 plots the path of predicted terminations for the 9.25 percent mortgage pools, where we observe the most prepayment activity. As shown, the model somewhat underpredicts the first termination wave, but does a better job in later months. The time series of the fit for the 8.25 and 7.25 percent pools are shown in figures 7 and 8. In general, the quality of the fit is in line with that for the 9.5 percent pools. The model does a better job of fitting the initial prepayment wave of the 8.5 percent pools, however, it fits less well in the later months for the 7.5 percent pools.

In part to underscore the hedging implications of our results, in figures 9 and 10 we graphically partition the cumulative principal cash flows for two high coupon pools; a 9.5 percent pool and a 9.25 percent pool. As shown, the primary source of principal paydown for these pools is refinancing. Background terminations and scheduled amortization comprise a very small percentage of the total principal payments over the first eighty months. Our predictions for default occur after month thirty and account for about 6% of the cumulative principal cashflow for the 9.5 percent pool and about 7 % for the 9.25 coupon pools by the end of month eighty. We are unable to evaluate the accuracy of our predictions for the cumulative principal paydowns because pool-level data only provide information on the aggregate levels of terminations arising from scheduled amortization, background termination, prepayments, and defaults which are reported as return of principal due to Freddie Mac’s guarantees.

²⁷The average termination rates are computed across the pools in a given month.

Fortunately, our prior discussion concerning the results shown in figures 5–8 confirms that our predictions for total termination rates are, on the whole, quite accurate.

6.1 Pricing

One important objective of structural modeling strategies is to provide accurate predictions for the valuation of mortgage-backed bonds. We compute the predicted prices at origination for each coupon pool in the full sample and compare these predictions with observed TBA market trading prices for Freddie Mac PC’s at the same date.²⁸ Our use of observed “out of sample” prices for our comparison represents a more rigorous standard than has typically been applied in evaluating the performance of structural bond valuation models (see Eom, Helwege, and Huang (2002)). As shown in Table 5, we consistently underprice the mortgages relative to the observed TBA prices, although our degree of underpricing is well within the ranges previously reported in studies using “in sample” tests (Eom, Helwege, and Huang (2002), Huang and Huang (2002)). Under our more rigorous standard, and without refitting the term structure model for each origination date, we find that the range of mispricing is from a low of about 1.41 percent to a high of 7.12 percent in several of the very small samples. We also compute the estimated option adjusted spread (OAS) for each coupon group again using the TBA prices at the pool origination date. We find OAS levels that vary between 5 and about 25 — 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).

In figure 11, we report a set of duration and convexity calculations for a representative 30-year 8.5 percent coupon mortgage-backed security at origination ($t = 0$) using our structural valuation model. The panels 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), medium (dashed line), and maximum (dotted line); hence we have decomposed the overall MBS duration and convexity into the contributions of the heterogeneous mortgage pools that are distinguished by their transaction costs. The panels on the right display the approximate durations and convexities of the mortgage pools holding interest rates fixed at 5.7 percent, the long-term mean of the interest rate factor. As can be seen in the upper left panel, when 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

²⁸In the to-be-announced (TBA) market, the buyer and seller 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. We gratefully acknowledge Lakhbir Hayre at Salomon Smith Barney for providing the TBA price data.

dramatically. As we move to higher transaction costs, which reduces prepayment, the duration rises for all interest rates. At the maximum level of transaction costs, the likelihood of prepayments or defaults are very low, and the duration-yield curve resembles that of a straight bond.

The upper right panel figure 11 displays the “duration house price” curve, which is computed as $-\frac{1}{M^t} \frac{\partial M^t}{\partial H}$.²⁹ House price duration is negative; in contrast to interest rates, as house prices rise the value of the mortgage pool rises. As was the case with interest rate duration, house price duration rises as transaction costs rise, and approaches zero as transaction costs make default extremely expensive. As indicated in the plot at the upper right, for the zero- to medium-transaction costs, the pool values change by approximately 0.015 percent, and maximum transaction-cost pools change by approximately 0.005 percent, for each penny increase in house value when house prices stand at \$1.10. Given these sensitivities, it is clear that focusing solely on interest-rate duration on the asset-side of a “duration gap” calculation (the difference between asset and liability durations) misses an important risk exposure.³⁰

The lower panels of figure 11 examine the convexity at origination of our representative 8.5 percent weighted average coupon 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, the convexity falls, in line with the decline in duration above. At very high levels of transaction costs, the convexity-yield curve is monotonically downward-sloping, a reflection of the fact that lower interest rates in effect raise the duration of the pool by reducing discounting, thereby shifting the cash flow pattern toward the principal payment. The medium transaction cost calculations are somewhat surprising; at interest rates above about 5.7 percent, the convexity of the representative pool is below that of the zero transaction-cost case. At interest rates below about 5.7 percent, the convexity of the medium transaction-cost pool is well above that of the zero 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

²⁹We have followed the market convention of multiplying duration by minus one, to be consistent with the calculations of interest rate duration.

³⁰FannieMae and Freddie Mac have recently begun to release interest-rate duration-gap calculations for their massive mortgage portfolios. These calculations suggest they also ought to release house-price duration-gap calculations, particularly given the difficulties one faces in hedging house price exposures.

three mortgage pools. The patterns are fairly similar to those of the convexity-yield curves. The zero transaction-cost and medium transaction-cost curves intersect when house prices are approximately \$0.95.

The last two columns of the table 5 display the average predicted interest rate durations and convexities, respectively, for the mortgage pools in our sample. These are the average interest rate durations and convexities of the mortgage backed securities; in other words, for a given mortgage backed bond, the duration and convexity are the weighted average duration and convexity taken across the pools of mortgages that are differentiated by their transaction cost levels. The durations are, on average, very short - none are greater than 5, indicating that expected terminations are high, on average, in part a reflection of the low average transaction cost estimated from the data. The convexities are also relatively low, on average, with none above 150.

7 Conclusions

This paper develops a two-factor structural mortgage pricing model that handles both prepayment and default. Our results indicate that house price fundamentals are important and that a house price factor should be included in valuation models of mortgage contracts with embedded prepayment and default options. Our model also controls for the effects of exogenous background terminations, the effects of a possible credit-constraint channel for the loan-to-value ratio when the prepayment option is in-the-money, and the effects of frictions or transactions costs faced by borrowers. We find that high loan-to-value ratios do not constrain background termination levels for Freddie Mac Participation Certificates. The transaction cost distribution is sensitive to the inclusion of house price fundamentals as a second factor in the model. Including the house price process in the option pricing framework reduces the mean of the transaction cost distribution from 14.5 percent to 11.5 percent. We further verify the importance of the house factor in structural mortgage valuation models by analyzing the residuals of the one-factor model. We find that house prices are a statistically significant predictor of the one-factor residuals, but not the two-factor residuals. Together these results indicate that one-factor interest rate models of mortgage backed security valuation are misspecified and will lead to biased predictions of mortgage termination levels.

Our graphical evaluation of the estimation results suggests that our model succeeds in accurately matching the observed levels of total termination in the Freddie Mac PC pools. We further evaluated our estimation strategy by considering the cumulative breakdown of principal cashflow predicted by the model. We found that the estimated level of principal

returned due to default is a significant source of the cashflow in the high coupon pools. Unfortunately, our data does not allow us to verify these estimates. Finally, we test our results by generating predicted prices for the pools at origination and comparing these estimates with observed “out of sample” TBA prices for Freddie Mac PCs. We found that the model performs reasonably well and pricing errors range from about one to seven percent.

We conclude that our empirical two-factor structural model of risky MBS debt accurately accounts for the timing of termination events and offers a tractable framework to evaluate the effects of interest rate and housing price fundamentals on the valuation of Freddie Mac Participation Certificates. 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 municipal bonds, structured notes, such as collateralized debt obligations, and possibly callable corporate bonds. 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.

A Numerical Implementation of the Model

We need to solve the partial differential equation (Equation 9)

$$\frac{1}{2}\phi_r^2 r M_{rr}^l + \frac{1}{2}\phi_H^2 H^2 M_{HH}^l + (\kappa(\theta_r - r) - \eta r) M_r^l + ((r - q_H)H_t) M_H^l + M_t^l - rM^l + C = 0,$$

subject to appropriate boundary conditions. The interest rate and house price grids are naturally bounded at 0 and ∞ , so we first transform the problem to the unit square by means of the following transforms:

$$\begin{aligned} y &= \frac{r}{\gamma_y + r}, \\ z &= \frac{H}{\gamma_H + H}. \end{aligned}$$

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.

Rewrite equation (9) using the substitutions:

$$M^l(H, r, t) \equiv U(z(H), y(r), t) \tag{15}$$

$$M_H^l = U_z \frac{dz}{dH}, \tag{16}$$

$$M_r^l = U_y \frac{dy}{dr}, \tag{17}$$

$$M_{HH}^l = U_{zz} \left(\frac{dz}{dH} \right)^2 + U_z \frac{d^2z}{dH^2}, \tag{18}$$

$$M_{rr}^l = U_{yy} \left(\frac{dy}{dr} \right)^2 + U_y \frac{d^2y}{dr^2}, \tag{19}$$

to obtain:

$$\begin{aligned} &\frac{1}{2}\beta_r^2(1-y)^4 U_{yy} + \frac{1}{2}\beta_H^2(1-z)^4 U_{zz} + \\ &((1-y)^2 \alpha_r - (1-y)^3 \beta_r^2) U_r + ((1-z)^2 \alpha_H - (1-z)^3 \beta_H^2) U_H - \\ &\lambda_r \frac{y}{1-y} U + U_t + C = 0, \end{aligned} \tag{20}$$

where:

$$\begin{aligned} \beta_r &= \phi_r \sqrt{r(y)}, \\ \beta_H &= \phi_H H, \end{aligned}$$

$$\begin{aligned}\alpha_r &= \kappa(\theta_r - r(y)) - \eta r(y), \text{ and} \\ \alpha_H &= \theta_H H.\end{aligned}$$

We use a parallel hopscotch finite difference algorithm to solve equation (20). This involves discretizing the PDE in (20) as discussed in Gourlay and McKee (1977).

We work backward in time to solve (20). We start at time T , the maturity of the mortgage pool, and step one month backward to $T - 1$. Upon reaching time $T - 1$, we have an approximation to $M_u^l(H, r, t)$, the value of the mortgage when neither the prepayment nor the default option are exercised. We now need to determine whether default, prepayment, or neither, is optimal. If $M_u^l(H, r, t) \geq F(t)(1 + X)$, then it is theoretically optimal to prepay the mortgage. Similarly, if $M_u^l(H, r, t) \geq H_t(1 + X)$, then it is theoretically optimal for the mortgage holder to default. The actual value of the mortgage is a weighted average of M_u^l , $F(t)(1 + X)$, and $H_t(1 + X)$, with the weights equal to the probabilities implied by the model of termination hazard rates.

Dropping function arguments for brevity, given a hazard rate λ , we compute the probability of termination as $P = 1.0 - e^{-\frac{\lambda}{12}}$, as discussed earlier. The value of the mortgage is thus given by

$$M^l = (1 - P)M_u^l + PF(1 + X)I_1 + PH(1 + X)I_2, \quad (21)$$

where I_1 is an indicator function that is one when prepayment is optimal and zero otherwise, and I_2 is a similarly defined indicator function for default.

Table 1: Summary of Freddie Mac Pools with Single State Concentrations

Year	Weighted Average Coupon	Number of Pools	Number of Mortgages in Pool	% Weak Pools	Average SMM	Average % Change in House Prices
1993	7.25	32	520	0.00	0.0041	0.0016
1994	7.25	9	101	0.00	0.0037	0.0024
1993	7.50	62	1203	0.00	0.0056	0.0030
1993	7.50	4	90	0.00	0.0055	0.0027
1993	7.75	38	498	2.63	0.0080	0.0019
1994	7.75	6	120	0.00	0.0061	0.0028
1993	8.00	39	707	0.00	0.0098	0.0025
1994	8.00	8	118	0.00	0.0131	0.0026
1992	8.25	121	1340	0.00	0.0121	0.0011
1993	8.25	11	132	0.00	0.0092	0.0013
1994	8.25	4	47	0.00	0.0091	0.0024
1992	8.50	246	2690	1.22	0.0162	0.0010
1993	8.50	7	122	0.00	0.0153	0.0018
1994	8.50	8	97	0.00	0.0096	0.0028
1992	8.75	216	3036	0.00	0.0166	0.0010
1993	8.75	5	64	0.00	0.0138	0.0004
1994	8.75	1	9	0.00	0.0029	0.0025
1992	9.00	109	1268	1.83	0.0240	0.0022
1993	9.00	3	64	0.00	0.0131	0.0004
1991	9.25	66	932	30.30	0.0270	0.0012
1992	9.25	19	233	0.00	0.0259	0.0026
1991	9.50	299	3994	40.80	0.0308	0.0004
1992	9.50	3	61	0.00	0.0147	0.0020
1991	9.75	24	247	65.20	0.0278	0.0006
Total		1340	17665			

Table 2: Estimation Results

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

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

The loan-to-value ratio at time t is denoted by LTV_t . The dummy variable D_t is one when either the default or prepayment options are exercised, and zero otherwise. The dummy variable P_t is one when just the prepayment option is exercised, and zero otherwise.

The coefficients β_4 and β_5 determine the transaction cost distribution. The mean of the transaction cost distribution is given by:

$$\mu = \frac{\beta_4}{\beta_4 + \beta_5}$$

and its variance is given by:

$$\sigma^2 = \frac{\beta_4 \beta_5}{(\beta_4 + \beta_5)^2 (\beta_4 + \beta_5 + 1)}.$$

Approximate standard errors are shown below each coefficient estimate.

Coefficient	Full Sample	
	1-Factor	2-Factor
β_0	0.00104	0.0000787
	0.02427	0.0243952
β_1	1.54380	1.5179260
	4.26686	3.2689200
β_2	1.49625	0.4417458
	0.31582	0.0009270
β_3		1.1067770
		0.4215456
β_4	0.59317	0.3779162
	0.00013	0.0000489
β_5	3.43872	2.8838000
	0.00014	0.0007916
χ^2	23.72	23.590
N	118,654	118,654

Table 3: Regression of the Monthly Residuals from the One and Two Factor Models on the Ten-year Treasury Rate and the State-level House Price Index: Fixed-Effects Model by Pool

Coefficient	Full Sample	
	1-Factor	2-Factor
Intercept	-0.29967	-0.19600
	0.01610	0.01600
Ten-year Treasury rate	0.07220	0.05222
	0.00437	0.00433
Squared Ten-year Treasury rate	-0.00506	-0.00370
	0.00033	0.00032
House Price Index	0.07399	0.01753
	0.01210	0.01210
Squared House Price Index	-0.02570	-0.00275
	0.00507	0.00506
R^2	0.0037	0.0018
N	118,654	118,654

Note: Standard errors are shown below each estimate.

Table 4: Specification Test Results: J-Test

The table reports the results of our specification tests for the one- and two-factor models. The first row of the table defines the null hypothesis as the one-factor model (the alternative hypothesis is the two-factor model) and the second row of the table defines the null hypothesis as the two-factor model (the alternative hypothesis is the one-factor model). A statistically significant parameter estimate indicates that the null hypothesis should be rejected in favor of the alternative model

	Estimate	Standard	
		Error	t -Ratio
α where $H_0 : y = f(r; \beta_1) + \epsilon_0$	0.654	0.051	12.8
α^* where $H_0^* : y = g(r, H; \beta_2) + \epsilon_0^*$	0.344	0.051	6.8
N = 118,647			

Table 5: Average Predicted and Observed Prices, and Average Predicted Durations, by Year and Coupon

Year	Coupon	N	Average Predicted Price	Average Observed Price	Average Pricing Error	Average Predicted OAS	Average Predicted Duration	Average Predicted Convexity
1991	9.25	74	100.40	103.29	-2.89	15.81	3.3	125.4
	9.50	296	101.31	102.71	-1.41	21.94	3.0	136.8
	9.75	22	101.90	103.72	-1.82	24.77	2.8	134.2
1992	8.25	130	97.73	102.35	-4.62	7.07	4.2	113.3
	8.50	253	98.81	102.98	-4.16	8.87	3.9	123.7
	8.75	221	100.00	103.55	-3.56	11.33	3.6	134.9
	9.00	112	101.10	104.31	-3.21	15.00	3.2	142.9
	9.25	7	101.97	104.71	-2.75	19.29	2.9	139.4
1993	7.25	5	94.67	101.79	-7.12	5.00	4.7	113.6
	7.50	17	95.97	102.85	-6.87	5.00	4.6	116.3
	7.75	20	97.31	103.49	-6.18	5.00	4.3	131.1
	8.00	28	98.04	104.08	-6.03	5.00	4.2	147.1

Note: The 1993, 8.25 coupon group has been omitted due to the small sample size (n=2).

Figure 1: Predicted and Observed Termination Behavior for Different Sub-Models: 8.25% Weighted Average Coupon Pools

The figure compares the observed and predicted patterns of mortgage terminations for pools with weighted average coupons of 8.25 percent when different components of the model are added sequentially. The vertical axes show the termination rates, expressed in terms of the proportion of the pool remaining that terminates in a given month — the single month mortality rate. The details of the parameterization and the dataset from which the observed terminations are extracted are both discussed at length in section 4.1. The upper left panel displays the model when the hazard function and the transaction cost distribution are both disabled. The upper right panel displays the predictions of the model when the transaction cost distribution is enabled, but the overall hazard (equation (4)) remains disabled. The lower left panel displays the predictions of the model when the background hazard function is enabled only in the continuation region of the state space — the region where neither prepayment nor default are optimal. The lower right panel displays the fit of the full model, using the hazard specification in (4).

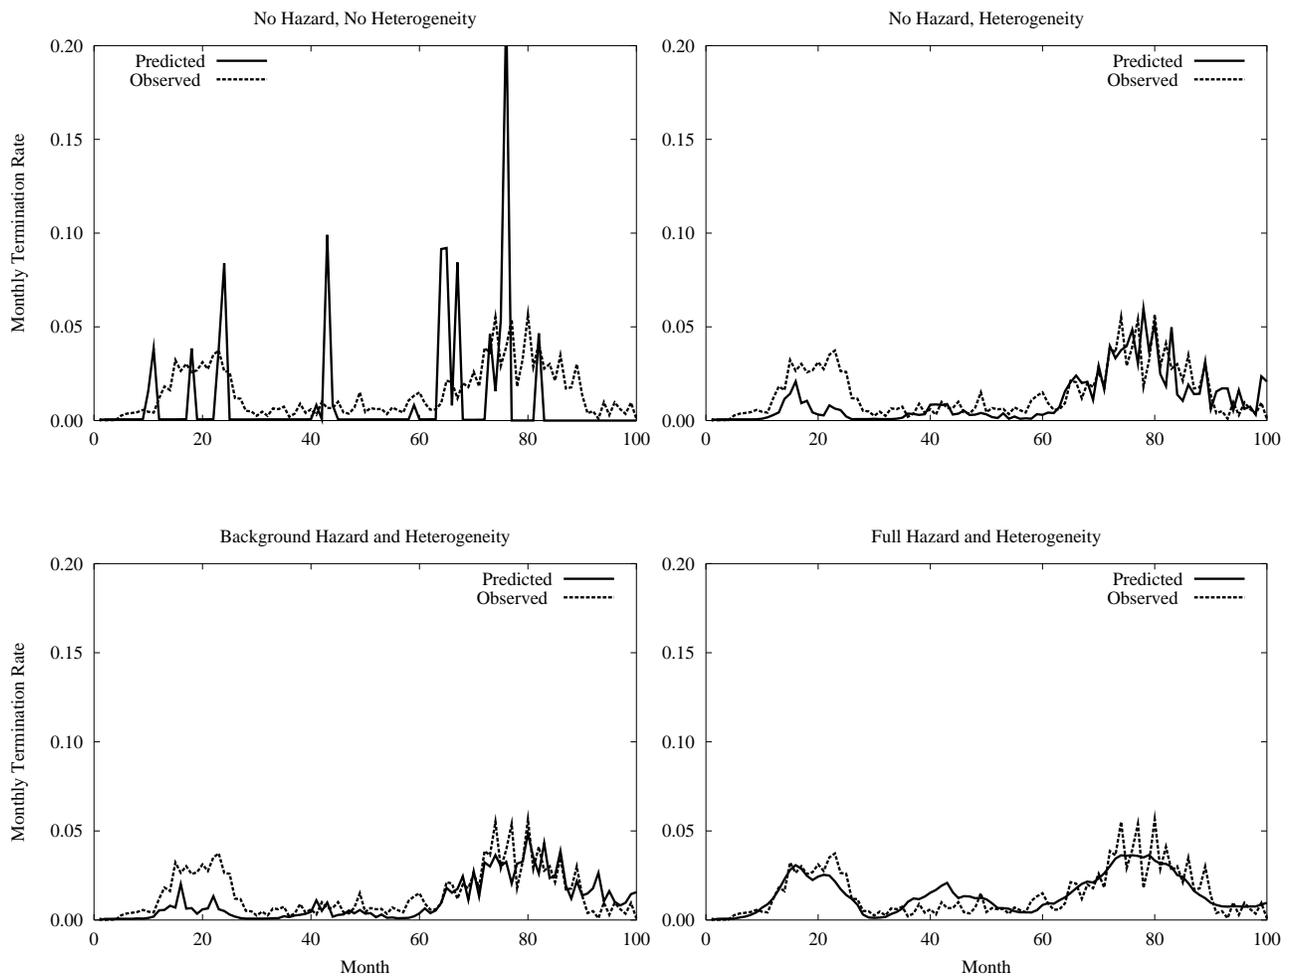


Figure 2: Average Single Month Termination Rates for High Coupon Weak and Non-Weak House Price Pools

The figure compares mortgage termination rates for non-weak house price pools (where the probability that the pool's loan-to-value ratio exceeds one is never greater than 20 percent over the sample period), to the termination rates for weak house price pools (where this probability exceeds 20 percent at some point in the sample period).

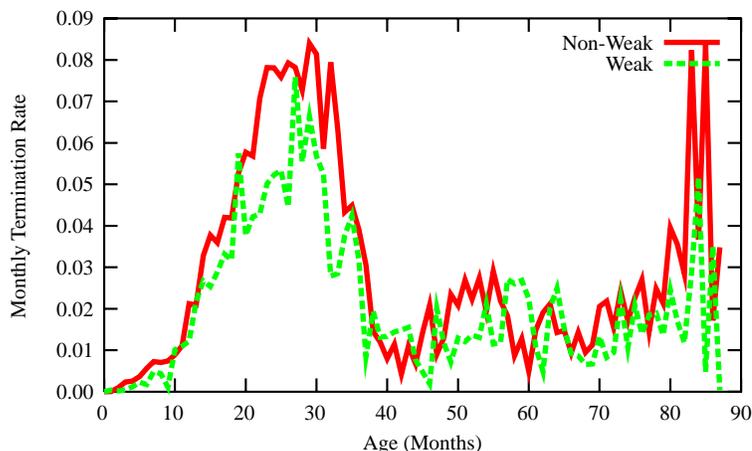


Figure 3: Average Monthly House Price Returns for High Coupon Weak and Non-Weak House Price Pools

The figure compares the average monthly returns to residential single family homes for pools with weighted average coupons equal to or greater than 9% that did not experience weak house prices (where the probability that the pool's loan-to-value ratio exceeds one is never greater than 20 percent over the sample period), to the average monthly returns to homes in weak house price pools (where this probability exceeds 20 percent at some point in the sample period) that have the same weighted average coupons.

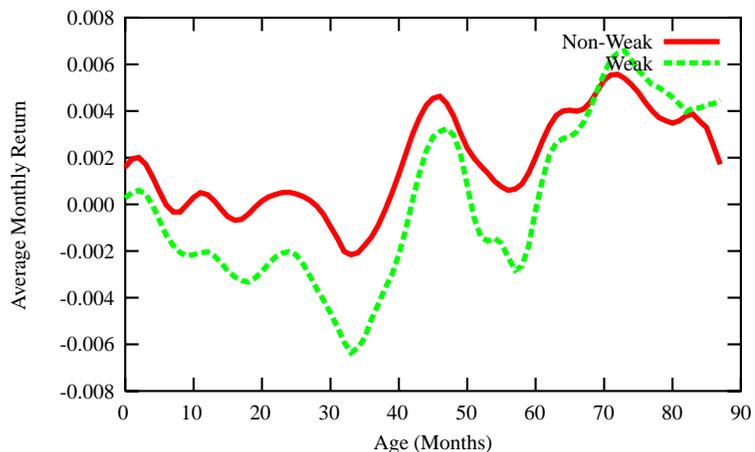


Figure 4: Estimated Transaction Cost Distribution

The figure shows the estimated distribution of transaction costs paid by borrowers within a mortgage pool on either refinancing or defaulting, expressed as a fraction of the remaining principal balance on the loan.

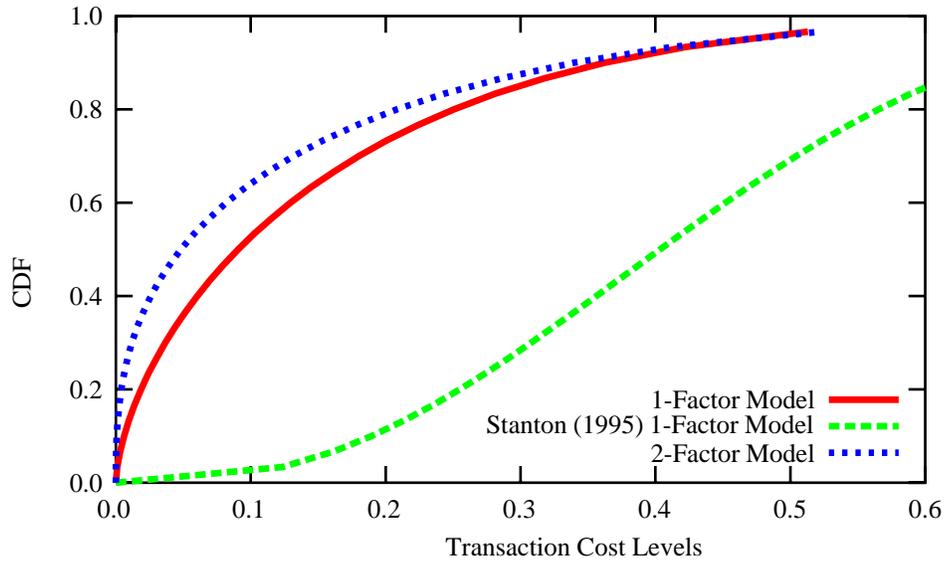


Figure 5: Full Sample Observed and Predicted Single Month Termination Rates

The figure compares actual and model-predicted monthly mortgage termination rates by month since issue, averaged across all pools in the sample.

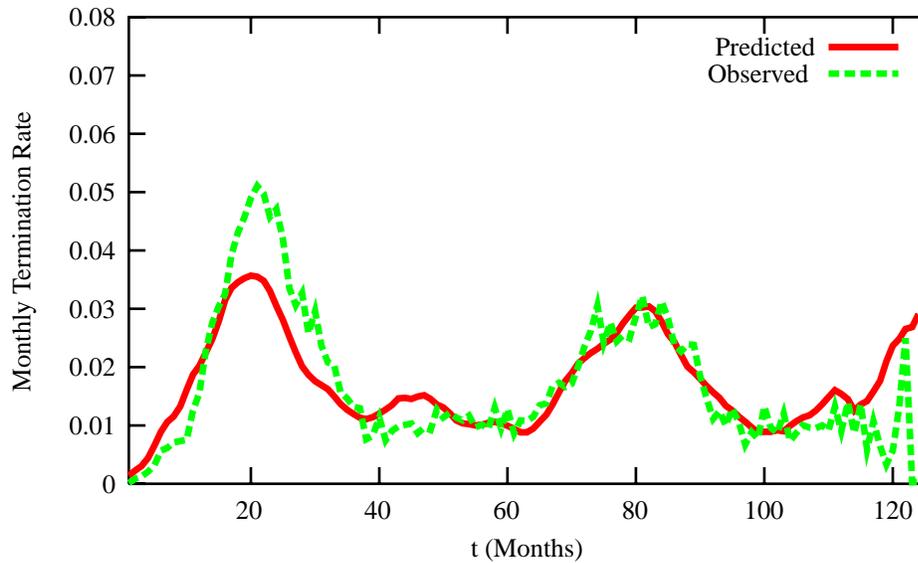


Figure 6: Observed and Predicted Single Month Termination Rates, 9.25% Coupon Pools

The figure compares actual and model-predicted monthly mortgage termination rates by month since issue, averaged across all pools in the sample with a coupon rate of 9.25%.

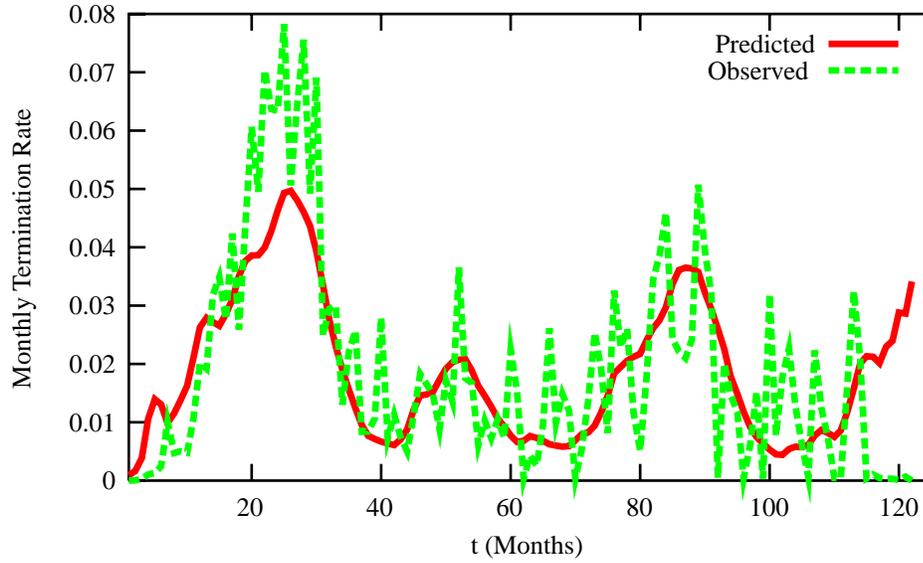


Figure 7: Observed and Predicted Single Month Termination Rates, 8.25% Coupon Pools

The figure compares actual and model-predicted monthly mortgage termination rates by month since issue, averaged across all pools in the sample with a coupon rate of 8.25%.

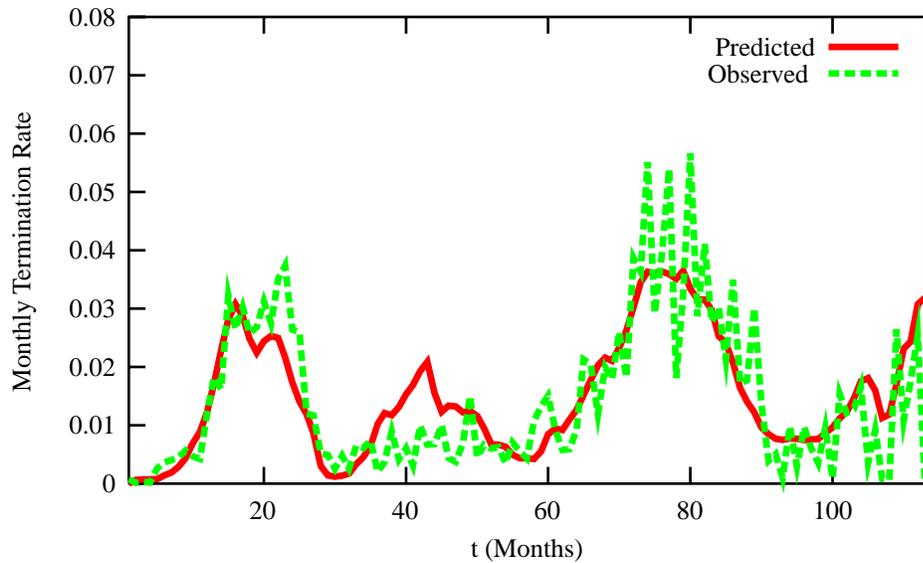


Figure 8: Observed and Predicted Single Month Termination Rates, 7.25% Coupon Pools

The figure compares actual and model-predicted monthly mortgage termination rates by month since issue, averaged across all pools in the sample with a coupon rate of 7.25%.

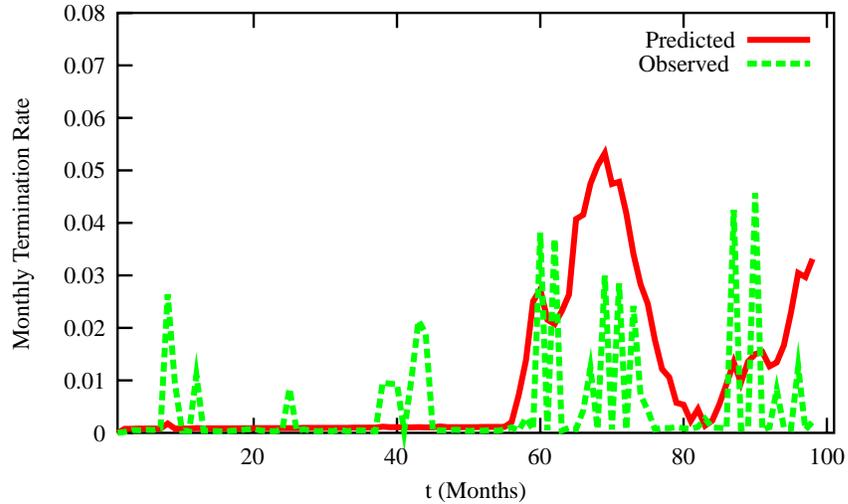


Figure 9: Cumulative Payments of Pool Principal by Source, 9.5% Coupon Pools

The figure depicts a partitioning of the cumulative principal cash flows for a 9.5 percent pool. The solid line depicts the cash flows associated with scheduled amortization. The dashed line depicts the sum of scheduled amortizations and the cash flows associated with background terminations. The short dashed line adds the cumulative cash flows associated with prepayments, and the dotted line gives the total cash flow. The difference between the dotted line and the short dashed line are the cash flows (return of principal by the guarantor) associated with defaults.

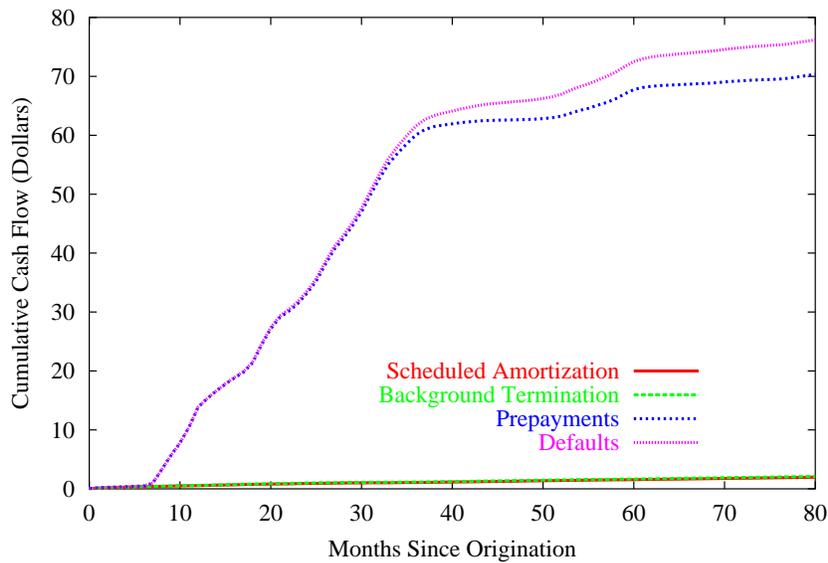


Figure 10: Cumulative Payments of Pool Principal by Source, 9.25% Coupon Pools

The figure depicts a partitioning of the cumulative principal cash flows for a 9.25 percent pool. The solid line depicts the cash flows associated with scheduled amortization. The dashed line depicts the sum of scheduled amortizations and the cash flows associated with background termininations. The short dashed line adds the cumulative cash flows associated with prepayments, and the dotted line gives the total cash flow. The difference between the dotted line and the short dashed line are the cash flows (return of principal by the guarantor) associated with defaults.

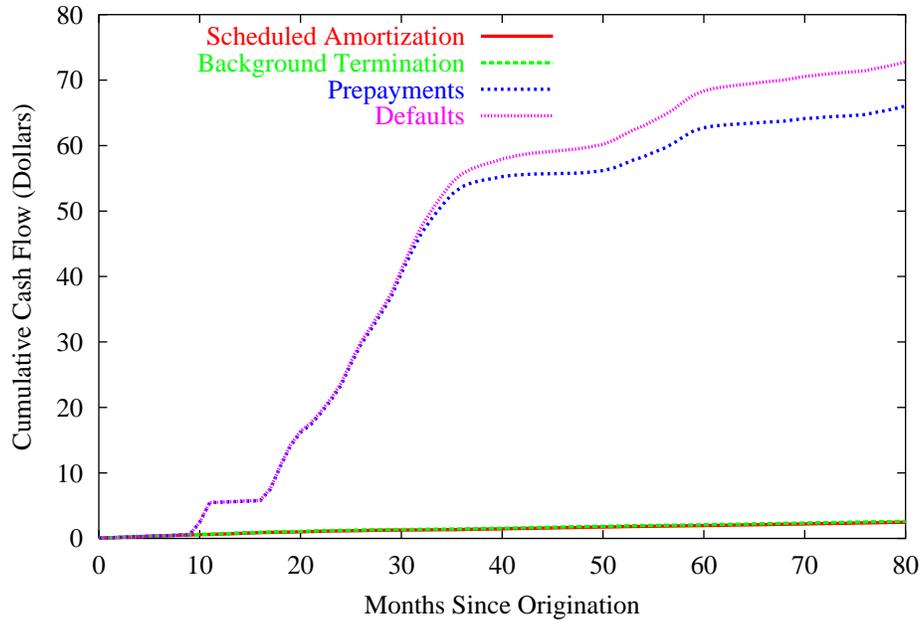
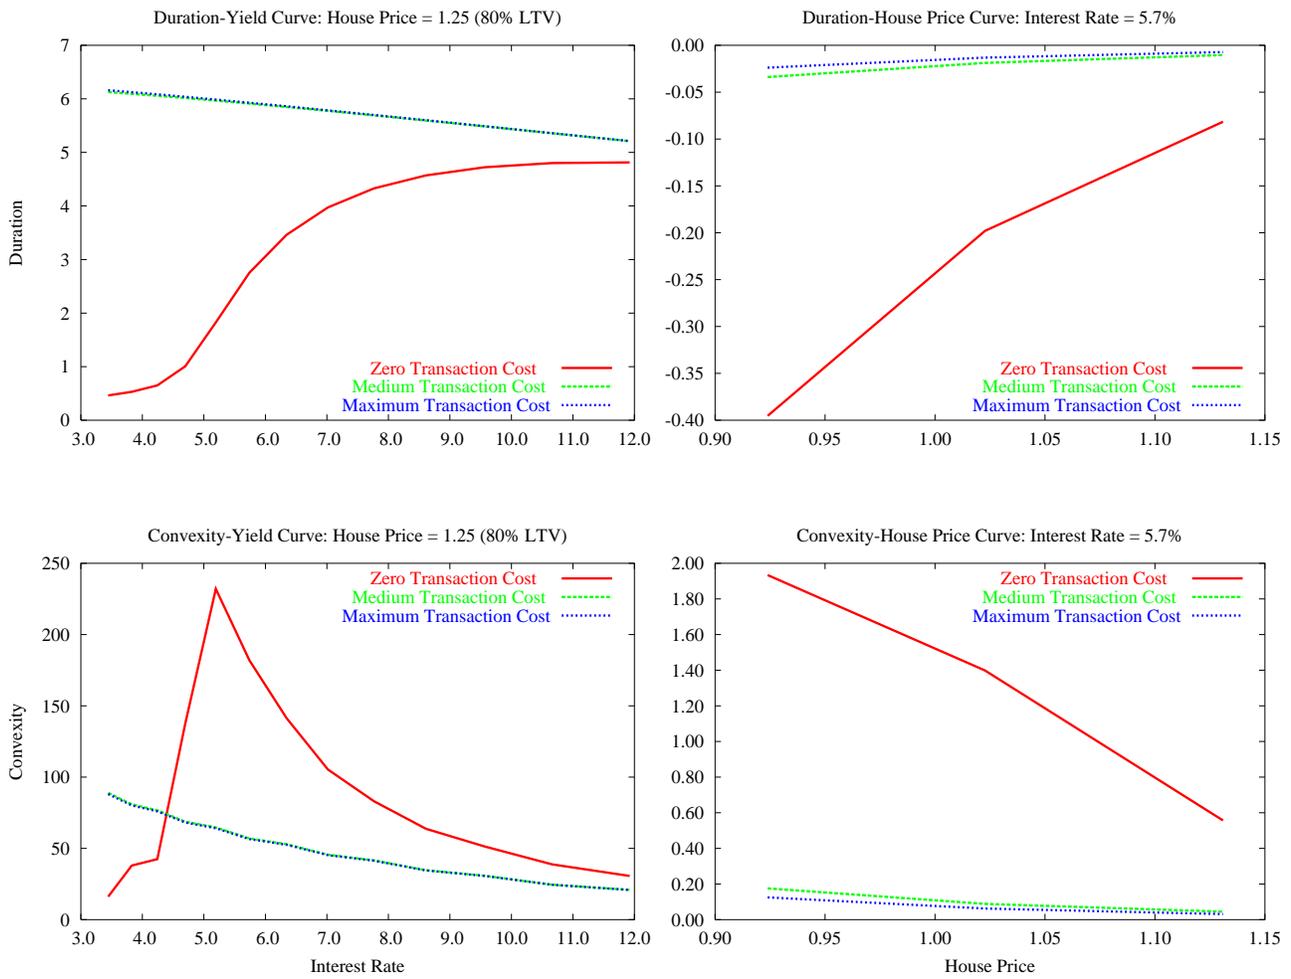


Figure 11: Price Sensitivity Measures: 8.5% Pool at Time 0

The figure reports a set of duration and convexity calculations for a representative 30-year 8.5 percent 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), medium (dashed line), and maximum (dotted line). The sub-plots on the right display the approximate durations and convexities of the mortgage pools holding interest rates fixed at 5.7 percent, the long-term mean of the interest rate factor. Note that house price duration is negative; in contrast to interest rates, as house prices rise the value of the mortgage pool rises.



References

- Anderson, R., and S. Sundaresan, 2000, A comparative study of structural models of corporate bond yields: An exploratory investigation, *Journal of Banking and Finance* 24, 255–269.
- 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, N.J.).
- Beckett, S., and C. S. Morris, 1990, The prepayment experience of FNMA mortgage-backed securities, Working paper, New York University Salomon Center.
- Boudoukh, J., M. Richardson, R. Stanton, and R. F. Whitelaw, 1995, A new strategy for dynamically hedging mortgage-backed securities, *Journal of Derivatives* 2, 60–77.
- Caplin, A., C. Freeman, and J. Tracy, 1993, Collateral damage: How refinancing constraints exacerbate regional recessions, Working Paper 4531, NBER.
- Collin-Dufresne, P., and R. S. Goldstein, 2001, Do credit spreads reflect stationary leverage ratios?, *Journal of Finance* 56, 1929–1957.
- Collin-Dufresne, P., and B. Solnik, 2001, On the term structure of default premia in the swap and LIBOR markets, *Journal of Finance* 56, 1095–1115.
- Cox, D. R., and D. Oakes, 1984, *Analysis of Survival Data* (Chapman and Hall, New York).
- Cox, J. C., J. E. Ingersoll, and S. A. Ross, 1985, A theory of the term structure of interest rates, *Econometrica* 53, 386–407.
- Davidson, R., and J. G. MacKinnon, 1981, Several tests for model specification in the presence of alternative hypotheses, *Econometrica* 49, 781–793.
- Deng, Y., and J. M. Quigley, 2002, Woodhead behavior and the pricing of residential mortgages, Working paper, University of Southern California.
- Deng, Y., J. M. Quigley, and R. Van Order, 2000, Mortgage terminations, heterogeneity and the exercise of mortgage options, *Econometrica* 68, 275–307.

- Duffee, G. R., 1998, The relation between Treasury yields and corporate bond yield spreads, *Journal of Finance* 53, 2225–2242.
- Duffie, D., and D. Lando, 2001, The term structure of credit spreads with incomplete accounting information, *Econometrica* 69, 633–664.
- Duffie, D., and K. J. Singleton, 1999, Modeling term structures of defaultable bonds, *Review of Financial Studies* 12, 687–720.
- Dunn, K. B., and J. J. McConnell, 1981a, A comparison of alternative models for pricing GNMA mortgage-backed securities, *Journal of Finance* 36, 471–483.
- Dunn, K. B., and J. J. McConnell, 1981b, Valuation of mortgage-backed securities, *Journal of Finance* 36, 599–617.
- 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.
- Eom, Y., J. Helwege, and J. Huang, 2002, Structural models of corporate bond pricing: An empirical analysis, Working paper, Ohio State University.
- 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.
- Jarrow, R. A., and S. M. Turnbull, 1995, Pricing options on financial securities subject to default risk, *Journal of Finance* 50, 53–86.
- Jay, W., and K. Anderson, 2002, The role of mortgage hedging in U.S. fixed income markets, Mimeograph, Lehman Brothers and Blackrock.
- Johnston, E., and L. Van Drunen, 1988, Pricing mortgage pools with heterogeneous mortgagors: Empirical evidence, Working paper, University of Utah.
- Kalbfleisch, J. D., and R. L. Prentice, 1980, *The Statistical Analysis of Failure Time Data* (John 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, and J. F. Epperson, 1992, A generalized valuation model for fixed-rate residential mortgages, *Journal of Money, Credit and Banking* 24, 279–299.
- Kau, J. B., D. C. Keenan, W. J. Muller, III, and J. F. Epperson, 1995, The valuation at origination of fixed rate mortgages with default and prepayment, *Journal of Real Estate Finance and Economics* 11, 5–39.
- Kau, J. B., and V. C. Slawson, 2002, Frictions and mortgage options, *Journal of Real Estate Finance and Economics* 11, 5–39.
- Leland, H., 1994, Corporate debt value, bond covenants, and optimal capital structure, *Journal of Finance* 49, 1213–1252.
- Leland, H., and K. Toft, 1996, Optimal capital structure, endogenous bankruptcy, and the term structure of credit spreads, *Journal of Finance* 51, 987–1019.
- Longstaff, F. A., and E. S. Schwartz, 1995, A simple approach to valuing risky fixed and floating rate debt, *Journal of Finance* 50, 789–819.
- 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.
- Mattey, J., and N. Wallace, 2001, Housing price cycles and prepayment rates of U.S. mortgage pools, *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* June, 255–69.
- 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, MIT.
- 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, *Journal of Finance* 44, 375–392.
- Singleton, K. J., 1985, Testing specifications of economic agents' intertemporal optimum problems in the presence of alternative models, *Journal of Econometrics* 30, 391–413.

- Stanton, R., 1995, Rational prepayment and the value of mortgage-backed securities, *The Review of Financial Studies* 8, 677–708.
- Stanton, R. H., 1996, Unobservable heterogeneity and rational learning: Pool-specific versus generic mortgage-backed security prices, *Journal of Real Estate Finance and Economics* 12, 243–263.
- Stein, J. C., 1995, Prices and trading volume in the housing market: A model with down-payment effects, *Quarterly Journal of Economics* May, 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.