Limits of Arbitrage: Theory and Evidence from the Mortgage-Backed Securities Market

Xavier Gabaix Arvind Krishnamurthy Olivier Vigneron*

January 22, 2004

Abstract

"Limits of Arbitrage" theories require that the marginal investor in a particular asset market be a specialized arbitrageur. Then the constraints faced by this arbitrageur (i.e. capital constraints) feed through into asset prices. We examine the mortgage-backed securities (MBS) market in this light, as casual empiricism suggests that investors in the MBS market do seem to be very specialized. We show that risks that seem relatively minor for aggregate wealth are priced in the MBS market. A simple pricing kernel based on the aggregate value of MBS securities prices risk in the MBS market. The evidence suggests that limits of arbitrage theories can help explain the behavior of spreads in this market.

JEL Codes: G12, G12, G14, G21

Keywords: Market segmentation, prepayment risk, liquidity, limited capital, hedge funds, limits to arbitrage.

*Respectively: MIT and NBER; Northwestern University; Deutsche Bank. For valuable comments, we are grateful to John Campbell, Kent Daniel, John Geanokoplos, Ananth Krishnamurthy, Andrew Metrick, Jose Scheinkman, Raman Uppal, and Dimitri Vayanos. We also thank Salomon-Smith-Barney for their data. Kripa Freitas provided excellent research assistance. E-mails: xgabaix@mit.edu, a-krishnamurthy@northwestern.edu, olivier.vigneron@db.com. All errors are our own.

1 Introduction

In the Fall of 1998, spreads on mortgage backed securities (MBS) widened substantially over Treasuries. At the same time, spreads in other markets also widened.[1] These events were coincident with turmoil in the hedge fund industry and the well publicized troubles of Long Term Capital Management.

Recent theories ("limits of arbitrage") suggest a causal link running from the losses suffered by hedge funds to the widening of spreads.[2] Arbitrageurs lost money over this period and their ability to take risk-positions decreased as their capital fell. As a result, arbitrageurs sold out of many positions leading to rising spreads.

Figure 1 illustrates the behavior of spreads (over Treasuries) on a sample of mortgage backed securities. In addition to the high spreads in 1998, note the high spreads in late-1993/early-1994. Events in this earlier period also fit the limits of arbitrage theory. A number of mortgage hedge funds lost money in 1993 as interest rates rose. A prominent example was Askins Capital Management which was rumored to have lost around $400 million.

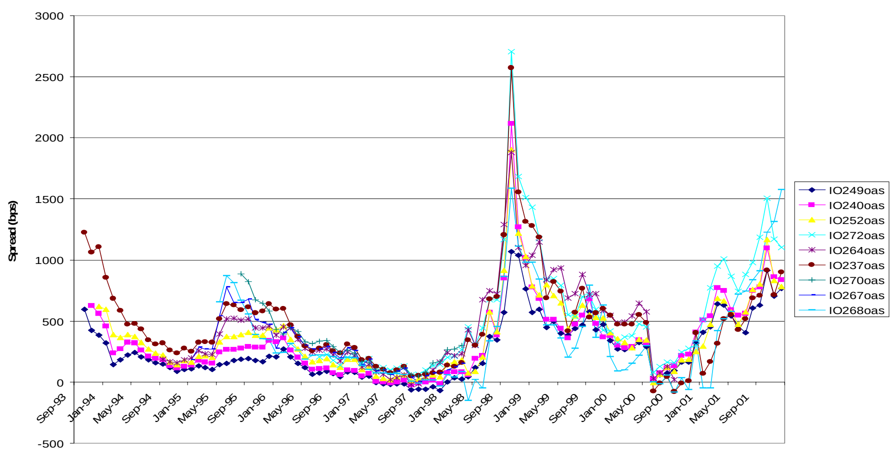

Figure 1: Mortgage Spreads, Sep. 93 - Dec. 01

The option adjusted spreads on a panel of interest-only strips (IO's) is plotted against time. The data from10/93 to 3/98 is from Salomon-Smith-Barney. The data from 4/98 to 12/01 is hand-collected from Bloomberg.

We provide evidence that advances the case for the limits of arbitrage theory. The difficulty in judging the theory based on Figure 1 is that the evidence boils down to just the two data points of 1993 and 1998. Our innovation is to develop and test the implications of the limits of arbitrage theory for the cross-section of mortgage backed securities. The use of cross-sectional data substantially increases the data available to test the theory.

The main idea is as follows. We note that the limits of arbitrage theory requires that

[1] See for example Krishnamurthy (2002) on the 30-year bond spread.

[2] See, for example, Shleifer and Vishny (1997), Kyle and Xiong (2001) or Gromb and Vayanos (2002).

the marginal investor in the MBS market is a specialized arbitrageur. Thus the constraints faced by these arbitrageurs affects asset prices. In contrast, traditional asset pricing theory assumes that the marginal investor in the MBS market is the same as that in the broad US capital market. If the marginal investor is a specialized mortgage arbitrageur, then the *relative* valuation of assets within the MBS market will also reflect this investor's preferences. We develop a theory and tests around this idea.

An MBS is a security whose cash flows are linked to payments of an underlying pool of (typically consumer) mortgages. Mirroring the underlying consumer mortgage, the MBS is basically a debt security with a declining principal value. However the difficulty in valuing (and hedging) MBS is that the decline in the principal is linked to actual consumer prepayments on the underlying pool of mortgages. These prepayments are not just a function of interest rates, but empirically seem driven by a host of other factors including local macroeconomic variables, demographics, etc.

To price MBS, Wall Street traders have developed sophisticated statistical models of consumer prepayment. These models forecast a pattern of cash flows for a given MBS, which can then be priced using standard techniques. However the traders can directly only hedge the interest rate components of their models. So, to the extent that the actual pattern of prepayments differs from the statistical model (in a way that is not captured by interest rate variation), the traders will be left with residual risk.

We exploit aspects of this residual risk to reach conclusions regarding the marginal investor in the MBS market. In principle, this residual risk should have a price of close to zero. One would be hard-pressed to find a correlation between the mortgage prepayments that are orthogonal to interest rates and aggregate wealth. In our data, the point estimate of the correlation with aggregate stock returns is 0.017. So, the risk should be diversifiable and should not carry a risk premium.

The first piece of evidence is that the residual risk is priced. In order to recover actual security prices in the MBS market, one has to discount with an additional spread - the "option adjusted spread" (OAS) - in valuing the cash flows based on the statistical prepayment model. It is common for traders to quote the OAS when pricing a particular MBS, or to use the OAS to compare different MBS. We have OAS data from a prominent mortgage trading firm for a panel of MBS securities. We measure the residual risk for each of these securities, and find a strong linear relation between a security's residual risk and OAS.

By itself, this result is not conclusive. Note that the OAS is measured relative to a benchmark financial model - it is analogous to an implied volatility for an option priced using the Black-Scholes formula. In the same way, it could be that the benchmark prepayment model is incorrect, and that this bias is most apparent in MBS with high residual risk.

We are more compelled by evidence that the pricing of this residual risk varies in a systematic way depending on the *total value* of the MBS market. This is exactly what one should find if the marginal investor is specialized to the MBS market.

Suppose that the set of buyers and sellers in the MBS market consist of banks and traders whose incentives are such that they are paid more when the MBS market performs well. For example, suppose that their compensation contracts pay them based on their profits on trading in the MBS market. If as a community these traders are the main holders of the MBS market, then their net compensation will in fact be highest when the MBS market rises.

If these banks and traders are the marginal investors in the MBS market, then securities where the residual risk is correlated with declines in the value of the MBS market will have higher risk premia. We find this to be the case.

The effect of prepayments on MBS depends on the difference between mortgage coupon rates and market interest rates. For example, if coupons are much higher than market interest rates, then faster than expected prepayments decrease the value of mortgages. The opposite is true when coupon rates are below market interest rates.

We find that the pricing of residual prepayment risk depends on the difference between the average mortgage coupons and market interest rates. For example, when average coupon rates are close to market interest rates, prepayments have little effect on the total value of mortgages. We find that the price of residual risk is very low in these environments (even for a specific high coupon MBS!)

The evidence forwarded in this paper is consistent with some other recent papers in the literature. Collin-Dufresne, Goldstein, and Martin (2001) study the corporate bond market. They find that a simple Merton (1973) model explains very little of the price changes in a panel of corporate bond prices. Even after including macro factors (stock market, etc.) they are only able to explain about 25% of price changes. The tantalizing evidence they present is that the bulk of the remaining price changes are due to a single factor that is common across all corporate bonds. Unlike us, they are unable to identify either this risk factor or the marginal investor who is pricing this risk[3].

Boudoukh, Richardson, Stanton, and Whitelaw (1997) provide similar evidence for the MBS market. These authors study the pricing of GNMA securities under a benchmark model (a multifactor interest rate model) that they propose. They study the errors of this model in pricing a panel of GNMA securities over a period from 1987 to 1994. Similar to Collin-Dufresne, Goldstein, and Martin, they find that a single (non-interest-rate) factor accounts for 80-90% of the common variation in the pricing errors.

Closer to our analysis is the work by Froot and O'Connell (1999) on the market for

[3]We conjecture that our theory could shed light on their results. We are developing this idea in ongoing work.

catastrophe insurance. They note that there are times at which the price of catastrophe insurance seems to get unusually high. Froot and O'Connell demonstrate that these are also times in which the capital of all catastrophe insurers is low, and the quantity of insurance transacted is also low. Like us, they argue that the marginal investor in the catastrophe insurance market is an institution (the insurers) rather than the broad capital market. They reconcile the facts as: When the capital in the insurance market is low, the insurers are less willing to write catastrophe insurance and therefore the prices are high and the quantities transacted are low.

Merton (1987) is an early attempt to explore the implications of market segmentation for asset prices. The theory we develop in this paper is closer to the limits of arbitrage literature in which the marginal investor is a specialized institution, and the constraints faced by this institution feed through to asset prices. Allen and Gale (1994) study an environment in which traders must specialize ex-ante in a certain asset market, and this implies that ex-post there is limited market participation and the wealth of the specialized traders is critical in setting prices. Similar ideas are explored by Dow and Gorton (1994), Shleifer and Vishny (1997), Kyle and Xiong (2001) and Gromb and Vayanos (2002). Caballero and Krishnamurthy (2001, 2002) propose a model of emerging market crises in which the crisis is an event in which the marginal investor switches from the world capital market to a trader within the emerging market.

2 Description of the mortgage market

Mortgage backed securities are financial securities that are backed by a pool of underlying mortgages. As of June 2002 there was about $3.9 trillion worth of securitized mortgages.

As mentioned in the introduction, valuing MBS is complex because the cash-flows from an MBS will depend on the prepayment behavior of the consumers in the underlying mortgage pool. Valuing the MBS is typically a two-step process. First, one assumes prepayment behavior as a deterministic function of interest rate paths, housing prices and so on (Richard and Roll 1989, Schwartz and Torous 1990). Then one simulates several interest rate paths, discounting and averaging the cash-flows based on a term-structure model that is calibrated to the current market risk-free rates.

The model-implied prices under this methodology typically differ from quoted market prices. Market participants express this difference by quoting for each security a number called an Option-Adjusted Spread (OAS). More specifically, after the prepayment option has been taken into account through a deterministic prepayment function, the OAS is a spread added to interest rates in the term-structure valuation model in order to recover the quoted market price of the security (equation (7) below). To the extent that the term-structure model is correct, the OAS is a "non-interest rate" risk premium on the security.

We study the OAS on a panel of securities in order to reach conclusions regarding the marginal investor in the MBS market.

Some market participants view the OAS as a risk-premium for bearing model-prepayment risk – this is our assumption as well. But it should be clear that OAS could be capturing many other things. To the extent that the term-structure model is mis-specified, the OAS could reflect an interest rate risk premium. To the extent that the prepayment model does not accurately reflect the prepayment option, the OAS could be capturing the true value of the prepayment option. Finally, although we study fairly liquid securities, the OAS could reflect a liquidity premium. These alternatives will pose a challenge for our empirical methodology.

We have so far used the generic term mortgage-backed securities, which comprises all the securities that are subject both to interest rate risk and prepayment risk. In fact, MBS are very diverse in form. The most typical MBS are the collateral (or passthroughs) which pass both interest and principal payments of the underlying pool of loans to the investor. Collateral is often stripped into two derivatives: the IO (Interest-Only) which passes only the interest payments to the investor, and the PO (Principal-Only) which pays the amount of principal amortized each month (scheduled amortization and prepayments).

3 Model

In this section we describe a very simple environment for studying the pricing of MBS and the OAS. We then develop a general equilibrium model and present our main hypotheses regarding the pricing of MBS.

3.1 Mortgage backed securities with no prepayment risk

Consider a world with a constant interest rate of r and a mortgage-pool with constant prepayment rate of ϕ and coupon of c. At any date t, the amount of outstanding of this mortgage-pool is $a(t)$, where,

$$\frac{da(t)}{dt} = -\phi a(t),$$

given some $a(0)$. We normalize $a(0) = 1$.

The IO is defined as the claim on all of the coupons from this mortgage-pool. Thus, the value of one unit face of the IO is simply,

$$V_{IO} = \int_0^\infty e^{-rt} a(t) c dt = c \int_0^\infty e^{-(r+\phi)t} dt = \frac{c}{r+\phi}. \tag{1}$$

The PO is defined as the claim on the principal repayment on this mortgage-pool:

$$V_{PO} = \int_0^{\infty} (-da(t))e^{-rt}dt = a(0) - r\int_0^{\infty} e^{-(r+\phi)t}dt = 1 - \frac{r}{r+\phi}. \tag{2}$$

Finally the value of the mortgage itself – the collateral – is,

$$V_C = V_{IO} + V_{PO} = 1 + \frac{c-r}{r+\phi} \tag{3}$$

3.2 Prepayment risk

Our aim is to develop an equilibrium model along the lines of a static CAPM to illustrate how prepayment risk is priced.

There are two periods, $t = 0, 1$. We assume that the riskless interest rate is constant and normalize it to be one. We assume there are K mortgage pools. In each pool, the mortgage has coupon c^k and quantity θ^k. We assume that mortgages "payoff" at date 1 as a function of c^k, r and ϕ^k. We next describe the payoff function.

We assume that the only uncertainty is in the prepayment rate, ϕ^k, of mortgage-k. The mean forecast of ϕ^k is $\bar{\phi}^k$. Pricing the IO, for example, based on this mean forecast would yield a value of,

$$EV_{IO}^k = \frac{c^k}{r+\bar{\phi}^k}.$$

The problem is that there is model risk as the actual ϕ^k may differ from $\bar{\phi}^k$. Let $\Delta\phi^k = \phi^k - \bar{\phi}^k$ be this variation.[4] We assume that $\Delta\phi^k$ has mean zero and covariance matrix of Ω.

For simplicity, we linearize the above valuation expressions and assume that the date 1 value (terminal payoff in our two-period world) of the k-th IO is,

$$V_{IO}^k = \frac{c^k}{r+\bar{\phi}^k}\left(1 - \eta^k\Delta\phi^k\right). \tag{4}$$

Where $\eta^k = 1/\left(r+\bar{\phi}^k\right)$ and $-\eta^k\frac{c^k}{r+\bar{\phi}^k}$ is the derivative of the IO with respect to the prepayment rate.

Likewise the date 1 value of the k-th PO is,

$$V_{PO}^k = 1 - \frac{r}{r+\bar{\phi}^k}\left(1 - \eta^k\Delta\phi^k\right). \tag{5}$$

Finally, the date 1 value of the k-th collateral is,

$$V_C^k = 1 + \frac{c^k - r}{r+\bar{\phi}^k}\left(1 - \eta^k\Delta\phi^k\right). \tag{6}$$

[4] Unlike our abstraction, in practice interest rates are uncertain. The logical extension of our model to the uncertain interest rate case is to write $\bar{\phi}^k(\tilde{r})$. Then the innovation of $\Delta\phi^k$ is the uncertainty in prepayments that is orthogonal to changes in interest rates. This is the definition we use in the empirical section of this paper.

3.3 OAS

Let P_{IO}^k and P_{PO}^k be the date 0 prices of one dollar face value of the IO and PO. The OAS is defined as the premium to the discount rate of r that is required to recover the market prices of the securities under the mean prepayment forecast. For example, in the case of the IO, the OAS is the solution to,

$$P_{IO}^k = \frac{c^k}{r + \bar{\phi}^k + OAS_{IO}^k} \tag{7}$$

That is the mean prepayment forecast is $\bar{\phi}^k$. Evaluated at this forecast, the value of the IO would be $\frac{c^k}{r+\bar{\phi}^k}$. So, the OAS is the premium to r required to recover the actual market price.[5]

There are two ways to look at the OAS. First, it may simply reflected a mis-specified model of the prepayment option. Perhaps informed market participants have a true model of prepayments which is actually $\hat{\phi}^k$. A naive market participant (and the econometrician) who uses $\bar{\phi}^k$ would have to introduce the additional discount rate of $\hat{\phi}^k - \bar{\phi}^k$ in order to recover the true market prices.

A second way to look at the OAS is that it is a risk premium. Any time that prices differ from expected values, the OAS will be non-zero. However, under this interpretation it may be either an interest rate risk premium or a prepayment risk premium.

In our empirical tests we will try to rule out the alternative hypotheses that the OAS is due to a mis-specified model of the prepayment option or an interest rate risk premium.

Using the same logic as for the IO, the OAS for the collateral is the solution to,

$$P_C^k = 1 + \frac{c^k - r - OAS_C^k}{r + \bar{\phi}^k + OAS_C^k} \tag{8}$$

(i.e. it is the previous valuation expression with an adjustment to r).

Now, from (4) and (6) we see that the date 1 payoff on the collateral is equal, state-by-state, to the payoff on a one dollar face of bond plus the payoff on $\frac{c^k-r}{c^k}$ of the IO. Thus, by arbitrage,

$$P_C^k = 1 + \frac{c^k - r}{c^k} P_{IO}^k$$

Using this relation, along with (7) and (8), we arrive at,

$$OAS_C^k = \frac{c^k - r}{c^k + \bar{\phi}^k + OAS_{IO}^k} OAS_{IO}^k \tag{9}$$

[5] We have defined the OAS as the instantaneous risk premium. The "lifetime" risk premium, defined as $\pi_{IO} = E\left[V_{IO}\right]/P_{IO} - 1$ is $\pi_{IO} = OAS_{IO}/\left(r + \phi^k\right)$, where $1/\left(r + \phi^k\right)$ is the duration of the IO. This is analogous to a stock with risky dividends, where the lifetime risk premium is the instantaneous risk premium multiplied by the asset's duration.

The relation between the OAS on the IO and the collateral depends on the coupon on the mortgage relative to market interest rates. In a low interest rate environment ($r < c^k$), the OAS on the IO and the collateral have the same sign. Intuitively this is because shocks lowering the value of the IO – i.e., faster prepayments – also lower the value of the collateral. In the high interest rate environment ($r > c^k$), the converse is true, and the OAS of the collateral has the opposite sign of the IO.

Note that these relations are derived only from arbitrage considerations. We have not made any statements about the equilibrium, or how risks are priced.[6]

3.4 Delegation and the marginal investor

The critical assumption that we make – and for which we provide tests – is that a representative MBS trader is the marginal investor in this market. Formally, we assume that at date 0 there is a set of risk-neutral investors (*"investors"*) with large endowments, as well as a set of MBS traders (*"traders"*) with no endowments. The risk-neutral investors find it unprofitable to invest in the MBS market directly. There is extreme adverse selection: if the investors try to buy mortgage backed securities, then snake oil salesmen will sell them stuff that is worth zero. As a result they give their funds to the specialized MBS trader who invests for them. In order to provide incentives to the MBS trader they give the trader a linear profit share of α. The problem is that the trader is risk averse. He has utility over date 1 wealth of,[7].

$$U(w) = E[w] - \frac{\rho}{2} Var[w] \tag{11}$$

i.e. just a mean-variance maximizer.

[6]The OAS for the PO is defined by

$$P^k_{PO} = 1 - \frac{r + OAS^k_{PO}}{r + \bar{\phi}^k + OAS^k_{PO}}$$

Repeating the arbitrage argument in the text (the payoff on the PO is equal to the payoff on a one dollar face of bond minus the payoff on $\frac{r}{c^k}$ of the IO), we find that,

$$OAS^k_{PO} = -\frac{r}{\bar{\phi}^k + OAS^k_{IO}} OAS^k_{IO}. \tag{10}$$

The OAS on the PO and IO have opposite signs. An increase in prepayment hurts the IO but benefits the PO: Thus the IO and PO have opposite sensitivities to prepayment risk.

[7]See e.g. Holmstrom and Milgrom (1987) for conditions under which a linear incentive is optimal. Given the relatively small risks of the MBS market, it is also plausible that the linear quadratic approximation is a good representation of the trader's risk attitude.

3.5 Equilibrium

At date 0, the investors give w_0 to the traders. From this the traders purchase a portfolio of mortgage backed securities. Let x^k_{IO} and x^k_{PO} be the amount of the k-th IO and PO held in a portfolio. Then,

$$w = w_0 + \sum_k x^k_{IO}(V^k_{IO} - P^k_{IO}) + \sum_k x^k_{PO}(V^k_{PO} - P^k_{PO}) \tag{12}$$

is the date 1 value of the trader's portfolio. Since the trader receives a linear profit share, his problem is to maximize (11) given (12).

This formulation is a variant of the traditional static CAPM. Deriving the first order condition for the trader's portfolio choice problem and then substituting in the market clearing condition of $x^k_{IO} = x^k_{PO} = \theta^k$, yields an expression for the price of the IO,

$$\frac{c^k}{r + \bar{\phi}^k} - P^k_{IO} = -\rho\alpha \; cov\left(\frac{c^k}{r + \bar{\phi}^k}\eta^k \Delta\phi^k, r_M\right) \tag{13}$$

where the market return is defined as:

$$r_M = \sum_j \frac{\theta^j}{\left(r + \bar{\phi}^j\right)^2}\Delta\phi^j (r - c^j) \tag{14}$$

The term on the right hand side of (13) is a risk premium for holding prepayment risk.

3.6 Covariance structure

We make the following simplifying assumption on the covariance structure. We write,

$$\Delta\phi^k = \beta^k \Phi + \epsilon^k \tag{15}$$

where, Φ is a common shock affecting prepayment across all securities, β^k is the loading of security k on the common shock, and ϵ^k is an idiosyncratic prepayment shock. We normalize the variance of Φ to be 1.

Under this assumption,[8]

$$OAS^k_{IO} \approx \rho\beta^k\alpha\left(\sum \frac{\beta^j}{(r + \bar{\phi}^j)^2}\theta^j (c^j - r)\right)$$

The sum term poses a problem for us to identify in the data. It is a weighted sum of the coupons of all mortgages in the market, where the weights depend on the amounts

[8]The exact expression is,

$$OAS^k_{IO}\frac{r + \bar{\phi}^k}{r + \bar{\phi}^k + OAS^k_{IO}} = \rho\beta^k\alpha\left(\sum \frac{\beta^j}{(r + \bar{\phi}^j)^2}\theta^j (c^j - r)\right).$$

This expression can be derived from combining (13) with (7), and noting that $\eta^k = 1/(r + \bar{\phi}^k)$.

outstanding of the mortgage and the loading on systematic prepayment risk. To compute the sum requires us to have data on the entire mortgage market – which we do not have. Instead, it is common for mortgage traders to follow whether the market as whole is at a premium or a discount. We compute a weighted average coupon on the liquid benchmarks in the mortgage market (312 securities). The relation we use in our tests is,[9]

$$OAS_{IO}^{k} = \overbrace{\beta^{k}}^{\text{Systematic risk}} \times \underbrace{\rho a\left(\bar{c} - r\right)}_{\text{Market price of risk}} \tag{16}$$

where $\bar{c}$ is the weighted average coupon and ρa is just proportional to the risk tolerance of the MBS trader. The approximation of using the simple weighted average for the coupon is valid when r is in the neighborhood of $\bar{\phi}^{j}$. Alternatively, note that the difference of $c^{j} - r$ is the dominant factor governing changes in the sum for r near c^{j}.

Loosely speaking, the first term in (16) captures the systematic risk of the mortgage, and the term involving the average market coupon captures the market price of risk (recall that ρ is the risk tolerance preference parameter for the MBS trader).

In equilibrium, the market price of risk is proportional to $\bar{c} - r$. Intuitively, when the market as a whole is at a premium – i.e. coupons above r – faster prepayments are costly to the representative trader. Thus securities whose value decrease because of faster prepayments command a positive risk premium. This is the reason that the OAS on the IO is positively related to $\bar{c} - r$. In fact, securities whose values *increase* because of faster prepayments will carry a negative risk premium in this environment. An example of such a security is the PO. A little algebra gives us that the OAS for the PO is equal to,

$$OAS_{PO}^{k} = -\beta^{k} \times \rho a\left(\bar{c} - r\right) \times \frac{r}{\bar{\phi}^{k} + OAS_{IO}^{k}}$$

Another example of a security whose value increases with faster prepayment is a *discount* collateral. A collateral with a coupon below the market interest rate increases in value if the mortgage prepays faster than expected. Given relation (16) and (9), we can write the OAS on the collateral as,

$$OAS_{C}^{k} = \beta^{k} \times \ \rho a\left(\bar{c} - r\right)\left(c^{k} - r\right) \ \times \ \frac{1}{c^{k} + \bar{\phi}^{k} + OAS_{IO}^{k}} \tag{17}$$

Thus the OAS on the collateral depends on both whether the market as a whole is at a premium as well as whether or not a particular security is at a premium. This leads to a quadratic dependence on r. We test this relation in our empirical work.

[9] We have also developed a continuous time model to express the relation between the OAS and prepayment risk. The resulting expressions are very similar to the ones we have derived in the text. See *http://econ-www.mit.edu/faculty/xgabaix/papers.htm* for details.

Finally, all of these relations are reversed when the market as a whole is at a discount. In this case, faster prepayments increase the value of the market. Hence the IO has a negative risk premium, while the PO commands a positive risk premium.

The dependence of the price of prepayment risk on $(\bar{c}-r)$ is really a general equilibrium implication. It seems plausible that the relation between β^k and the OAS could be spurious, or due to model mis-specification. But we think that the fact that it depends on the interaction between β^k and $(\bar{c}-r)$ stems uniquely from equilibrium considerations. Most of our empirical tests are built around this interaction term.

3.7 Testable empirical predictions

The main predictions of the model are contained in equation (16), which we can unpack as:

$$OAS_{IO}^{kt} = \beta^k \lambda_t \tag{18}$$

$$\lambda_t = \rho a \left(\bar{c}_t - r_t\right) \tag{19}$$

where ρa is a constant proportional to the risk aversion of the traders.

- A. In the cross-section, the loading of IO-k on the common component of prepayment uncertainty explains the OAS on the IO's.
- B. In the time-series, the difference between the average market coupon, $\bar{c}_t$, and the market interest rate, r_t, explains the evolution of the market price of prepayment risk λ_t.
- C. In the cross-section, the residual prepayment risk of security-k (i.e. $\sigma(\epsilon^k)$) is not priced.
- D. Eq. (17) predicts that the OAS on the collateral is quadratic in the market interest rate, r_t, and is a function of both c^k as well as the average market coupon, $\bar{c}_t$.

The model also points out two alternative explanations for the OAS that our empirical work must contend with: The OAS may reflect an interest rate risk premium; or, the OAS may reflect a mis-specified model of the prepayment option. We discuss both of these possibilities in far greater depth in the next section.

3.8 Discussion

The model we have presented is simplified along many dimensions in order to isolate the relation between the OAS and prepayment risk. We comment on two omissions in this subsection: lack of dynamics, and constant interest rates.

In a dynamic model the current wealth of the arbitrage traders will be an important state variable. For example, the view that high MBS spreads in 1993 and 1998 were due to low arbitrageur capital is an essentially dynamic effect. This is a substantive effect that our model does not capture. Informally, we can think of a changing ρ as proxying for this effect.

To the extent that the aggregate value of the mortgage market is a sufficient statistic for the marginal utility of the representative trader, our cross-sectional pricing equations will be unaffected by the omission of dynamics. Generally, in a dynamic model, the marginal utility will also depend on changes in the investment opportunity set. If preferences are close to unit-elastic, the latter effect will be small and our analysis will remain valid.

In assuming constant interest rates, we have eliminated two potentially important effects. First, there is no place for an interest rate risk premium in the OAS. This is actually easy to fix. We consider a variant of the model with interest rate uncertainty in subsection 5.1 in order to discuss whether or not the OAS includes an interest rates risk premium. The second effect we ignore is the relation between interest rate volatility and the valuation of the prepayment option. It is clear that as interest rate volatility increases, the value of a bond option will rise as well. Since the prepayment option is a bond option, we expect that the value of mortgage backed securities will also be affected. Brown (1999) notes a positive relation between OAS and implied volatilities on Treasury bond options. He argues that this is evidence that market practice is to use a constant volatility in pricing models and this also gives rise to the OAS. We will have to take into account these effects in our empirical work.

Strictly speaking, these latter two observations only apply in the case that our OAS data is calculated under a mis-specified term structure model. Market practice is to use a term-structure model that is calibrated to current market risk-free rates and then discount the cash-flows under the risk-neutral measure implied by the term-structure model. So by construction, the OAS cannot reflect interest rate risk. A small miscalibration of those models is possible, but is unlikely to account for OAS as high as 500bp. In any case, it is easy to reject this alternative based on the evidence we present in the next section.

4 Data and estimation

Our data consist of the OAS for nine IO's and PO's (see Table 1) furnished by Salomon-Smith-Barney. This data is daily and covers a period beginning (for some securities) in August 1993 and ending in March 1998. We also have data on the historical prepayment rates (monthly frequency) of the underlying collateral. The nine strips chosen are very liquid securities and fairly representative in age and coupon of the active secondary market.

The collateral are all FNMA 30-year conventional loans.

Table 1									
IO/PO[a]	249	240	252	272	264	237	270	267	268
Coupon[b] (%)	7.08	7.49	7.95	8.07	8.49	8.48	9.01	8.91	9.64
Age[c]	58	60	63	27	50	70	80	47	110

a: Securities are identified by their pool number.

b: Weighted average coupon on underlying mortgage pool ($\pm 5bp$ over sample.)

c: Age in months as of July 98.

We also have monthly data on the OAS for six generic FNMA 30-year collateral covering a period from October 1987 to July 1994. The coupons on these securities range from 7.5% to 11% and the data was provided by Smith-Breeden. We use this data for some auxiliary tests. We do not have data on prepayment rates for these securities.

The bulk of our analysis is conducted using the IO data. The main reason we focus on the OAS of the IO's (as opposed to the PO's) is that they are large and hence measured with less error.[10] Of course, both the OAS of the IO and the PO contain the same information, as we have shown previously. We return to the OAS on the PO's briefly when we discuss the interest rate risk premium hypothesis. We construct time series of monthly OAS by forming simple averages of the daily figures. This reduces micro-structure effects. The data is an unbalanced panel, with common last observations, but varying initial observations.

There are two steps in testing (18)–(19). We need an estimate of β_k, and we need an estimate of $\bar{c} - r$. To form $\bar{c}$, the average coupon oustanding, we take the "market" to be represented by 312 liquid securities over the period 1986 to 1998, and then compute the weighted average of the coupons at each date (weights are amounts outstanding). This gives us a monthly series of $\bar{c}_t$. We use the 10-year constant maturity Treasury yield as r_t. Most of the underlying mortgages have durations around 10 years.

The estimate of β_k is a bit more involved. We first develop a bare-bones statistical prepayment model. For each IO, we have the historical paydown of its collateral month by month, expressed as a series s_{kt} (single monthly mortality, or monthly prepayment rate). The prepayment model is,

$$s_{kt} = \alpha_{0k} + \alpha_{1k} r_t + \alpha_{2k}(r_t - r_{t-1}) + \alpha_{3k} age_t + \epsilon_{kt}$$

where, age_t is the age of the mortgage (note that coupon is absorbed into α_{0k}). We assume that the error follows an AR(1) process,

$$\epsilon_{kt} = \rho \epsilon_{kt-1} + u_{kt}$$

[10]This is consistent with the evidence in Breeden (1994), who shows that IO's have higher signal-to-noise ratio than PO's and pass-throughs.

This procedure results in a time-series of $\hat{u}_{kt}$'s for each security. Note that by construction the $\hat{u}_{kt}$'s are orthogonal to interest rates.

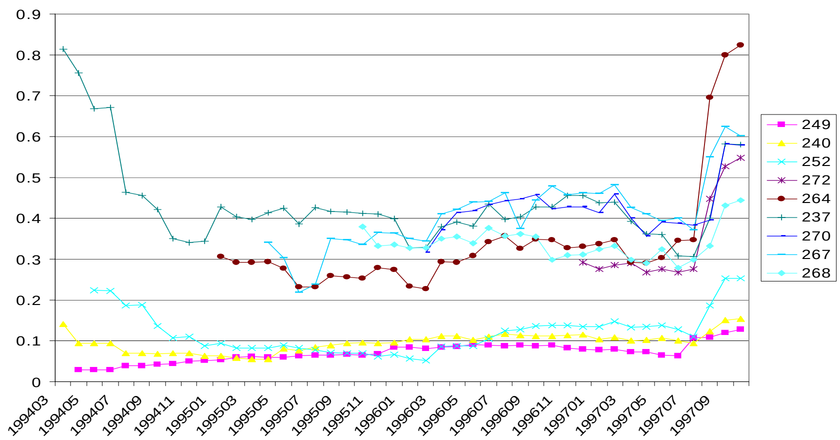

Figure 2: Standard deviation of prepayment errors

The one-year rolling standard deviation of the errors from the empirical prepayment model are plotted over time. There are nine mortgage pools that we study. As a result, there are nine standard-deviation series.

Figure 2 shows the time series of

$$\left(\sum_{s=t-6}^{t+6} \frac{\hat{u}_{ks}^2}{13} \right)^{1/2}.$$

The figure present a rolling one-year standard deviation of the errors. The standard deviations are higher in the beginning of the sample and at the tail end of the sample, but more or less constant at other times. For this reason, we take the prepayment risk β^k to be constant throughout our sample.

The more striking pattern in the figure is that the rankings by standard deviation are fairly well preserved over time. Our estimate of β_k is based on this relation.

Table 2

IO/PO	249	240	252	272	264	237	270	267	268
β^k (st. dev.)	0.083	0.122	0.181	0.465	0.504	0.561	0.479	0.483	0.384
β^k (PCA)	0.058	0.120	0.197	0.431	0.612	0.513	0.475	0.538	0.338
idiosync^k (PCA)	0.091	0.078	0.086	0.174	0.270	0.181	0.259	0.190	0.213
β^k (mtge. model)	1.06	2.10	3.29	4.80	6.67	5.63	6.58	6.63	5.82

We use two proxies for β^k. In the first line of Table 2 we present the sample standard deviations for the errors. In some of our tests we use these standard deviations as β^k.

However, as we have noted before, the idiosyncratic component of the prepayment risk should not be priced. We do not have the prepayment rates for the entire mortgage market. However, on the assumption that our sample is representative, we use a principal-component's analysis to extract the common component and the idiosyncratic component of the prepayment risk.

We focus only on the overlapping observations (22 months) for this analysis.[11] The first eigenvector accounts for 83% of the variance which suggests that equation (15) is a good representation of the data. The second and third components account for 9% and 3.2% respectively. Table 2 (second line) presents the loading on the first eigenvector for each security as well as the standard deviation of the residual (third line). We use the loading on the common factor as β_k, and the residual standard deviation as our measure of idiosyncratic risk. Unfortunately the two vectors are very similar, and as we will see, the test of prediction (C) is not informative.

We know that the best predictor of s_{kt} given the history of past interest rates is non-linear (prepayment functions are typically complex non-linear functions of the entire path of interest rates), however our simple approach avoids the difficult task of calibrating such a complex model. As a check, we also have prepayment forecasts from the prepayment model of a mortgage trader and have used these residuals to form β's. The β's (subject to scaling factors) look similar, suggesting that our model is reasonable. See the last line of Table 2.

As a preliminary to our main empirical work, we check whether prepayment risk would be priced if markets were integrated. We form a time series of prepayment risk innovations from our estimates of $\hat{u}_{kt}$. For each t we compute,

$$U_t = \frac{1}{K} \sum_{k=1..K} \frac{\hat{u}_{kt}}{\beta^k}$$

where the β^k's are the loading on the first eigenvector from the principal component analysis.

We also form a time series of monthly excess returns on the S&P500 (SP_t). The two series have a correlation coefficient of 0.017. We run a regression of,

$$U_t = A + B \times SP_t.$$

The coefficient estimate for B is 0.482 and the t-statistic is 0.125 ($N = 53, R^2 = 0.0003$). We conclude that the two series are unrelated.

The organization of this empirical section is as follows. In the next subsection we give a brief account of the events that have marked the mortgage market over our sample by

[11] We have also done the principal component analysis dropping the security with the shortest time series. This results in 32 months of overlapping observations. The results are close to what we find for the 22 months.

looking at the evolution of r_t and $\bar{c}_t$. Our main results are in sections 4.2 and 4.3. We discuss alternative hypotheses and robustness in section 5.

4.1 Interest rates, average market coupon, and OAS

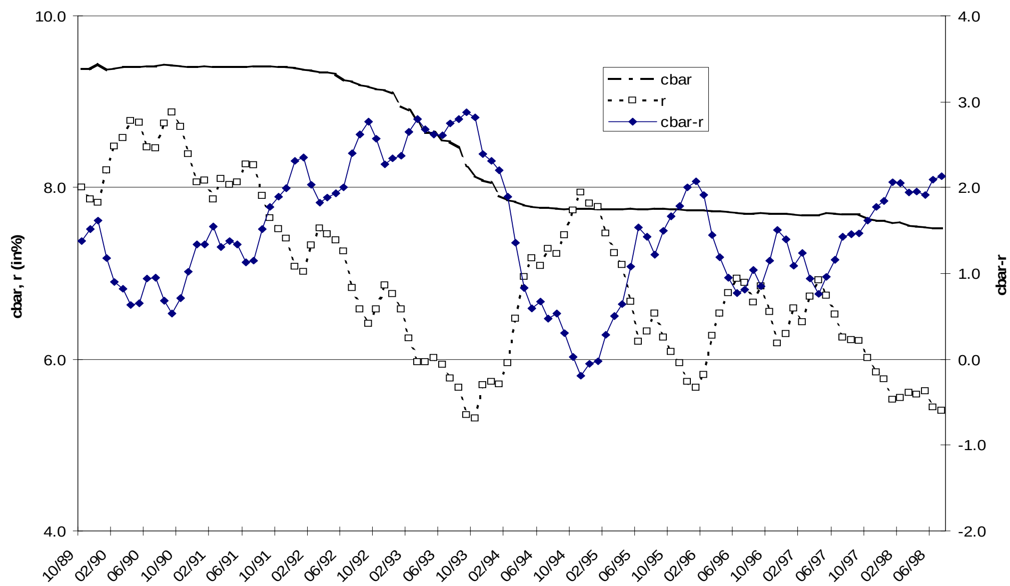

Figure 3: CMT 10-year and Average Coupon

The yield on the the 10-year constant maturity Treasury note and the average coupon on a representative sample of 312 liquid benchmarks in the MBS market are plotted over time.

Figure 3 shows the time-series of the CMT 10 year r_t, and the outstanding average coupon $\bar{c}_t$. It is worth noting that the adjustments of the outstanding average coupon are slow compared to the large movements of market interest rates. Prior to 1993, prevailing mortgage rates were around 10-11%. There was a large prepayment wave as rates fell in 1992 and 1993. As a consequence, the outstanding average coupon $\bar{c}_t$ adjusted down from values of 9-10% to 7-8%. We follow the evolution of the OAS of the IO's and PO's from 1993 to 1998. At the start of this period, interest rates were rising as the U.S. economy was exiting a recession. The Federal Reserve raised their target rate in February of 1994 and followed this move with several others. Interest rates rose dramatically during this period. In 1995, there was another important market rally, as rates fell 200 b.p. from January 1995 to January 1996. Rates fell continuously from March 1997 to July 1998 by slightly more than 100 b.p. to reach levels as low as those of November 1993. It is also worth noting though that by the end of our sample period, the outstanding coupon had adjusted down to 7.5%.

Figure 4 shows the variation of the OAS of the IO's in our data over the period Autumn

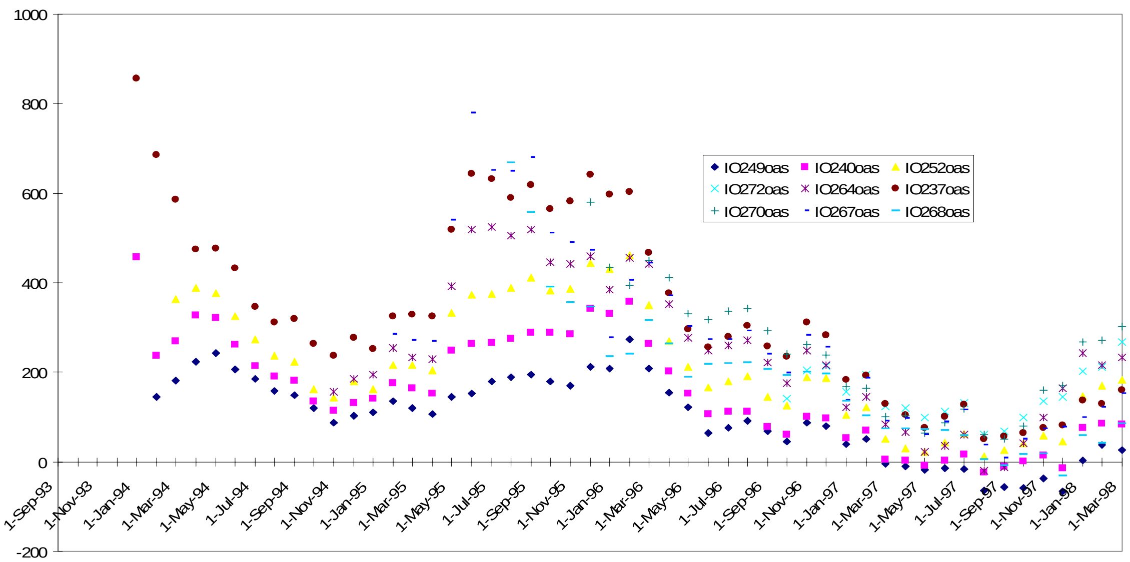

Figure 4: OAS of the IO's

The option-adjusted spreads on a panel of interest only strips is plotted over time. Each monthly data point is formed by taking a daily average of the spreads for that month. The data is from Salomon-Smith-Barney.

1993 to Spring 1998. One readily observes the large swings of the OAS of the IO's, from values above 500 b.p. in the beginning of the period to values close to zero in 1994 and 1996 when interest rates were very high. The OAS of PO's give a somewhat symmetric image, although at smaller magnitudes (as predicted by equation (10)). One should also note that the interest rate alone is not enough to understand the relative magnitude of the OAS of the IO between 1993 and 1998 when rates were at the same level: OAS are much higher in 1993 than in 1998. This in fact is not a puzzle in light of our derivations, since equations (18)–(19) tells us that the OAS of the IO is proportional to $\bar{c} - r$ and not r alone. Indeed, when looking at $\bar{c} - r$, we find that it is twice as high in 1993 than in 1998.

4.2 Cross-sectional estimates of the market price of risk

We run a cross-sectional regression, one for each month, where we estimate λ_t based on,

$$OAS_{IO}^{kt} = \alpha_t + \beta^k \lambda_t + \epsilon_t$$

These estimates exploit only the slope of the OAS. The variation in the level is picked up in α_t (for example, α_t picks up any time varying interest rate risk premium).

The λ_t estimates along with other sample statistics using β-stdev are in Table 3. The estimation errors are uniformly tight and the R^2 are high. In fact, for most of the months the OAS can be clearly ranked by the β^k's. We find this very encouraging for the theory because it suggests that prepayment risk has a lot to do with the OAS, and that our measure of β^k is in fact picking up the cross-sectional prepayment risk of the IO's.

Table 3: Regressions $OAS_{kt}^{IO}=\lambda_t$ β-stdev + $constant_t$, for each month *t*

λ_t is the market price of risk, β-stdev is the prepayment risk measure

Regression Month	λ estimated	Std. Error	R-squared	N (# obs.)	Regression Month	λ estimated	Std. Error	R-squared	N (# obs.)
Nov-93	1200	211	0.941	3	Jan-96	326	223	0.239	8
Dec-93	1384	129	0.968	4	Feb-96	309	175	0.266	8
Jan-94	980	152	0.892	4	Mar-96	472	64	0.848	8
Feb-94	1041	114	0.932	4	Apr-96	443	62	0.845	8
Mar-94	760	91	0.932	4	May-96	361	51	0.799	8
Apr-94	416	116	0.744	4	Jun-96	421	68	0.862	8
May-94	407	84	0.836	4	Jul-96	409	62	0.809	9
Jun-94	412	65	0.884	4	Aug-96	429	56	0.843	9
Jul-94	303	45	0.876	4	Sep-96	391	48	0.842	9
Aug-94	285	40	0.897	4	Oct-96	334	49	0.803	9
Sep-94	201	132	0.488	5	Nov-96	390	56	0.854	9
Oct-94	208	87	0.713	5	Dec-96	344	58	0.849	9
Nov-94	220	76	0.773	5	Jan-97	246	39	0.837	9
Dec-94	256	89	0.746	5	Feb-97	260	38	0.786	9
Jan-95	203	49	0.815	6	Mar-97	245	30	0.885	9
Feb-95	308	53	0.887	6	Apr-97	242	28	0.871	9
Mar-95	312	75	0.808	6	May-97	176	37	0.642	9
Apr-95	331	74	0.828	6	Jun-97	201	43	0.699	9
May-95	669	179	0.603	7	Jul-97	245	50	0.734	9
Jun-95	1068	266	0.629	7	Aug-97	171	57	0.531	9
Jul-95	932	191	0.661	7	Sep-97	135	61	0.382	9
Aug-95	970	240	0.663	8	Oct-97	197	64	0.601	9
Sep-95	939	196	0.754	8	Nov-97	250	72	0.546	9
Oct-95	701	148	0.731	8	Dec-97	360	85	0.572	9
Nov-95	692	146	0.726	8	Jan-98	303	118	0.398	9
Dec-95	576	161	0.642	8	Feb-98	235	108	0.283	9
					Mar-98	324	121	0.421	9

Figure 5 graphs the estimate of λ using β-stdev as well as the one standard deviation envelopes around the estimate. Also pictured is the difference between $\bar{c}_t$ and r_t. At a broad level the two series seem to follow each other. Early in the sample the fit is quite close. Later in the sample, while the ups and downs in the two series seem to track each other, the λ estimates seem like a muted version of $\bar{c}_t - r_t$.

We conjecture that the more muted relationship later in the period may have to do with a falling ρ over the sample period. It is well documented that in the 1993/1994 period a number of mortgage hedge funds suffered losses, and many went out of business. We conjecture that this led to a loss of capital in the mortgage market and lower capacity for risk taking. As time passed, capital flowed back into these funds and ρ fell. Froot (2001) finds this effect in the catastrophe insurance market. We intend to investigate this more in future work.

Figure 6 graphs the estimate of λ using the β-PCA. The results are similar to those of Figure 5.

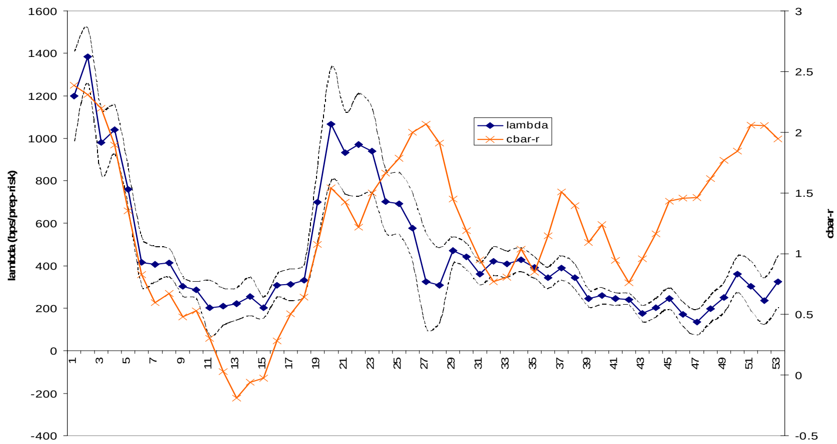

Figure 5: λ Estimates using β-Stdev

The estimates of the market price of prepayment risk (λ) using the β-Stdev as the measure of security-specific prepayment risk, is plotted over time.

4.3 Tests using the entire panel

We now report the results of testing our model using the entire panel. Table 3 reports regressions based on the following model.

$$OAS^{kt}_{IO} = \alpha_t + \gamma_k + A \times \beta^k \left(\bar{c}_t - r_t\right) + \epsilon_{kt}$$

The regression includes both fixed effects and time effects.

Both the OAS series and the $(\bar{c}_t - r_t)$ series are persistent, so there is correlation in the standard errors. We correct for this in two ways. First, most of the regressions report standard errors using the clustering option on STATA (cluster by security). In Table 5, we also report regressions based on first-differencing the data.

Table 4

Regressions based on the OAS of the IO's:

$$OAS^{kt}_{IO} = \alpha_t + \gamma_k + A \times \beta^k(\bar{c}_t - r_t) + \epsilon_{kt}.$$

We also consider $idiosync^k\,(\bar{c}_t - r_t)$ as an explanatory variable.

	β^k-stdev			β^k-PCA	
	(1)	(2)	(3)	(4)	(5)
$\beta^k(\bar{c}_t - r_t)$	438.1 (4.13)	546.5 (5.81)	315.1 (5.57)	405.9 (2.73)	218 (0.81)
$idiosync^k\,(\bar{c}_t - r_t)$					576 (0.59)
R^2	0.93	0.95	0.92	0.93	0.93
N	383	194	189	383	383

Estimates reported with T-statistics based on clustered (by security) standard errors in parentheses. All regressions have fixed-effects and time-effects (not reported).

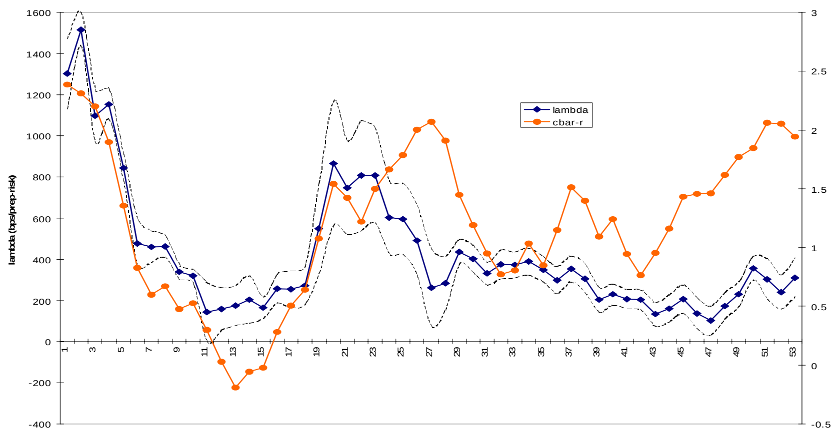

Figure 6: λ Estimates using β-PCA

The estimates of the market price of prepayment risk (λ) using the β-PCA as the measure of security-specific prepayment risk, is plotted over time.

Columns (1) - (4) of Table 4 contain the verification that our model fits the data. Column (1) is based on β-stdev, while column (4) is based on β-PCA. Column (2) and (3) gives the results from two sub-samples. We divided the sample into roughly equal number of observations. June 1996 is the dividing point between the two sub-samples (there are less observations early in our sample).

We note the lower (but as significant) coefficient in specification (3) compared to (2). This accords with our conjecture that there was more risk-bearing capacity (i.e lower ρ) in the latter half of the sample.

Column (5) contains the result of the following regression:

$$OAS_{IO}^{kt} = \alpha_t + \gamma_k + A \times \beta^k \left(\bar{c}_t - r_t\right) + B \times idiosync^k \left(\bar{c}_t - r_t\right) + \epsilon_{kt}$$

Our theory predicts that the idiosyncratic risk should not be priced. Unfortunately, as mentioned earlier, there is not enough independent variation in $idiosync^k$ and β^k to fashion a meaningful test of this prediction. Because the two series are so collinear, the standard errors on the estimates blow-up and neither coefficient is significant.

The persistence in the two data series may raise concerns that the correlation we find is spurious. Table 4 reports the result of a regression based on first-differencing the series:[12]

$$\Delta OAS_{IO}^{kt} = \alpha_t + A \times \Delta\beta^k \left(\bar{c}_t - r_t\right) + \epsilon_{kt}$$

The coefficients estimates do fall, but the results remain highly significant. Again there is a fall in the coefficient in the second half of the sample (specification (2) versus (3)).

[12] Adding a fixed effect to this regression has little effect since we are first-differencing the data.

Table 5			
Regressions based on the OAS of the IO's: $$\Delta OAS^{kt}_{IO} = \alpha_t + A \times \Delta\beta^k(\bar{c}_t - r_t) + \epsilon_{kt}.$$ β^k is the β-stdev.			
	(1)	(2)	(3)
$\Delta\beta^k(\bar{c}_t - r_t)$	175.1 (3.07)	222.1 (2.73)	76.3 (2.21)
R^2	0.71	0.68	0.81
N	374	186	180
Estimates reported with T-statistics based on robust standard errors in parentheses. Time dummies not reported.			

5 Robustness checks

5.1 Is the OAS due to a mis-specified interest rate model?

As we have noted before, the OAS may be due to a mis-specified term structure model. In this case, the OAS will reflect an interest rate risk premium.

There are a few reasons to think that this is not the case. If the OAS is picking up an interest rate risk premium, then we would expect that the OAS on the IO and the PO would have the same sign, and changes in these OAS would be positively correlated. This is not the case in the data. To see this, we return to our original valuation expressions for the IO and PO ((1) and (2)). Previously, in order to arrive at expressions for the pricing of prepayment risk, we linearized around $\bar{\phi}^k$. We also linearize around the interest rate of r.

$$V^k_{IO} = \frac{c^k}{r + \bar{\phi}^k}\left(1 - \eta^k \Delta\phi^k - \eta^k \Delta r\right).$$

Likewise the date 1 value of the k-th PO is,[13]

$$V^k_{PO} = 1 - \frac{r}{r + \bar{\phi}^k}\left(1 - \eta^k \Delta\phi^k + \frac{\bar{\phi}^k}{r}\eta^k \Delta r\right).$$

Note that the effect of a change in the prepayment rate has the *opposite* sign on the value of the IO and the PO. On the other hand, changing r has the *same* sign on the value of the IO and the PO. The latter statement is simply because both the IO and the PO are bonds, and increasing (decreasing) the market interest rate, lowers (raises) the value of a bond.

[13]Note that we have constructed $\Delta\phi^k$ to be the innovation to prepayments that is orthogonal to interest rates. Thus $\Delta\phi^k$ and Δr are uncorrelated.

This observation helps us to reject the alternative that the OAS is an interest rate risk premium. Suppose that the market price of interest rate risk was λ_r. Then deriving the equilibrium prices of IO and PO lead us to,

$$\frac{c^k}{r+\bar{\phi}^k} - P^k_{IO} = -\rho \frac{c^k}{r+\bar{\phi}^k} cov\left(\eta^k \Delta\phi^k, r_m\right) + \frac{c^k}{r+\bar{\phi}^k}\eta^k \lambda_r \sigma_r$$

while,

$$1 - \frac{r}{r+\bar{\phi}^k} - P^k_{PO} = \rho \frac{r}{r+\bar{\phi}^k} cov\left(\eta^k \Delta\phi^k, r_m\right) + \frac{\bar{\phi}^k}{r+\bar{\phi}^k}\eta^k \lambda_r \sigma_r.$$

It is clear that if the prepayment risk term was zero, then the OAS on the IO and the PO would have the same sign and changes in the OAS for both IO and PO would be positively correlated. In fact, the OAS on the IO and the PO typically have opposite signs. Moreover, changes in these OAS almost always have opposite signs. Figure 7 plots the time series of the OAS for IO's and PO's for four of our securities. The evidence suggests that the interest rate risk term is subordinate in the determination of the OAS.

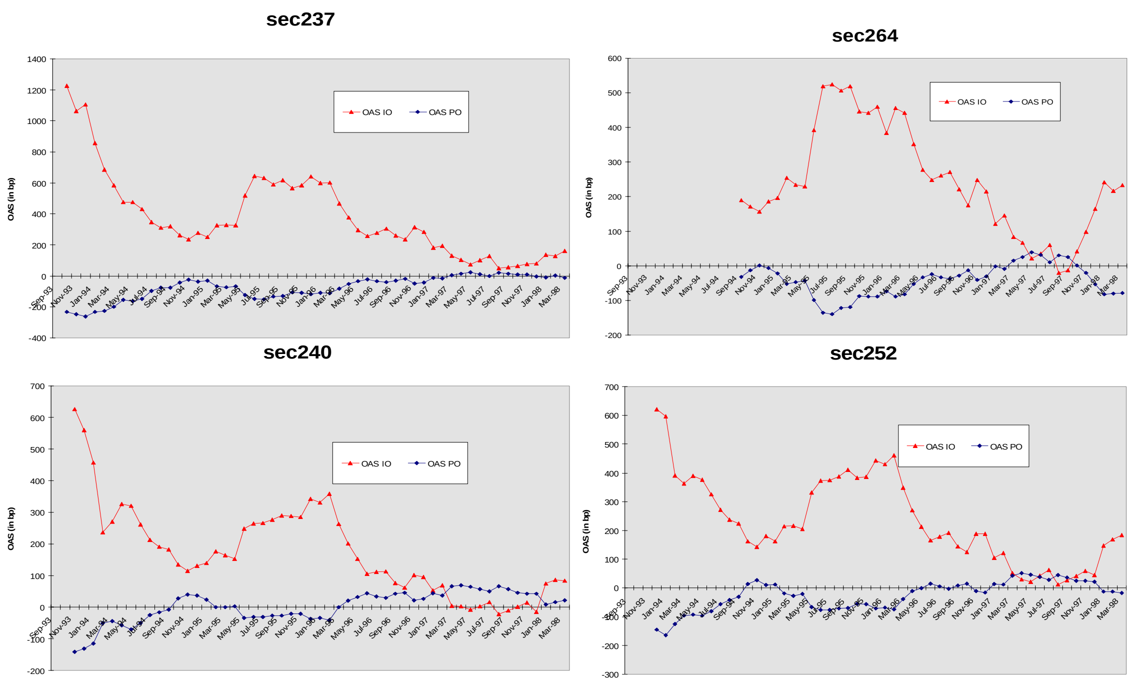

Figure 7: Time series of OAS for IO's and PO's

The option adjusted spreads of both IO and PO for four of the mortgage pools are plotted over time.

It may still be possible that there is a small interest rate risk premium in the OAS. But this should not pollute our estimates. Our basic regressions have time-dummies that will absorb all time variation in the risk premium. Thus as long as the interest risk premium is not higher for the securities with higher measured prepayment risk (i.e. β's), our estimates will be unaffected.

5.2 Is the OAS due to a mis-specified model of the prepayment option?

The other leading candidate explanation for the OAS is that it is due to an incorrect model of the homeowner's prepayment option. This possibility is harder to dismiss than the interest risk premium hypothesis. The reason – following on the logic of the previous subsection – is that faster prepayments lower the value of the IO while raising the value of the PO. This will naturally give rise to the result that the OAS on the IO's and the PO's have opposite signs (and have changes that have opposite signs). Thus we cannot use a similar argument to that of the previous subsection.

Let us revisit equation (7), rewritten below:

$$P_{IO}^{k} = \frac{c^k}{r + \bar{\phi}^k + OAS_{IO}^{k}}$$

Suppose that informed market participants have a true model of prepayments which is actually $\hat{\phi}^k$. If the average market participant quotes the OAS based on an incorrect assessment of prepayment and uses $\bar{\phi}^k$, then an additional discount rate of $\hat{\phi}^k - \bar{\phi}^k$ is required in order to recover the true market prices. In this case, the OAS is equal to $\hat{\phi}^k - \bar{\phi}^k$.

Note that the OAS on the IO's in our sample are for the most positive, while those on the PO are negative. Thus, under the mis-specified model hypothesis, the OAS must be based on a model which consistently *underpredicts* prepayments.

It is not clear why underprediction should be a regularity of prepayment models, but we attempt to control for this possibility in a few ways. If the underprediction can be written as,

$$\hat{\phi}^k - \bar{\phi}^k = \alpha_k + \gamma_t,$$

then the fixed-effects/time-effects regression specification we have used should take care of mis-specification. That is, for example, if the underprediction was constant across time, but varying by security, then the α_k's should pick up this mis-specification.

Thus the only case that poses a problem for our results is if the underprediction was a function of both security and time. This is more generally a virtue of our estimation strategy. Since our estimation is based on the interaction between a security effect (β^k) and a time effect ($\bar{c}_t - r_t$), all alternative hypotheses that have implications either for the time series or the cross-section, separately, have been controlled for.

Suppose that the underprediction is proportional to $\bar{\phi}^k$. That is, suppose that $\hat{\phi}^k$ is equal to $\bar{\phi}^k$ times a constant. This will clearly match the sign patterns of the OAS of the IO and PO. We run the following regression to account for this possibility,

$$OAS_{IO}^{kt} = \alpha_t + \gamma_k + A \times \beta^k \left(\bar{c}_t - r_t\right) + B \times s_{kt} + \epsilon_{kt}$$

where, s_{kt} is the actual single month mortality (SMM) for month t. We also run the same specification using an average of s_{kt} where the average is for 13 months centered around month t.

The results are reported in the first two columns of Table 6. The s_{kt} variables are not significant, while the coefficient on our model does drop but remains highly significant.

Table 6

Regressions based on the OAS of the IO's:

$$OAS_{IO}^{kt} = \alpha_t + \gamma_k + A \times \beta^k(\bar{c}_t - r_t) + \epsilon_{kt}.$$

β^k is the β-stdev. Additional explanatory variables are:

s_{kt}: SMM for security-k, month-t,

$\bar{s}_{kt}$: one-year moving average of s_{kt},

$c^k \times r_t$: coupon of security-k interacted with r_t.

	(1)	(2)	(3)
$\beta^k(\bar{c}_t - r_t)$	228.0 (3.74)	214.5 (4.0)	226.9 (4.27)
s_{kt}	-12.8 (-0.66)		
$\bar{s}_{kt}$		5.7 (0.15)	
$c^k \times r_t$			16.9 (6.81)
R^2	0.93	0.94	0.96
N	374	344	383

All regressions have fixed-effects and time-effects (not reported).

Estimates reported with T-statistics based on clustered (by security) standard errors in parentheses.

We have mentioned earlier that another explanation behind the OAS may be that traders use too low an interest rate volatility in their prepayment model. In this case, the prepayment option will be undervalued in the traders' models, giving rise to a positive OAS for the IO's (and a negative one for the PO's). This effect will vary across the moneyness of the option, possibly in a non-linear fashion. We control for this possibility by introducing a term that is quadratic in c^k and r:

$$A(c^k)^2 + Bc^k r + Cc^k + Dr + Er^2.$$

Since our basic regression already has a time-effect and a fixed-effect the terms involving only c^k or only r are already controlled for. So, the regression we run is,

$$OAS_{IO}^{kt} = \alpha_t + \gamma_k + A \times \beta^k\left(\bar{c}_t - r_t\right) + B \times c^k r_t + \epsilon_t$$

The result is in the last column of Table 6. The interaction term is significant suggesting that there is some mis-specification that distorts the OAS. The coefficient on our model still remain significant.

Another possible reason for underprediction that has been suggested to us is the "rational prepayment" hypothesis.[14] Banks typically calibrate their prepayment functions to historical experience. Consumers have in the past exercised their refinancing options suboptimally. In particular, in cases where consumers should refinance, they refinance too little – so low interest rate environments, for high coupon mortgages will be particularly affected. Now suppose that a smart trader knows that looking to the future, consumers will exercise their refinancing options more optimally. Then there clearly will be a bias in a model that is calibrated to historical experience. One side of this bias is that a smart trader will require a faster prepayment rate for high coupon mortgages in low interest rate environments. Thus there is underprediction for these mortgages. The underprediction should by related to the difference between c^k and r. We control for this effect in our specifications.

As a comment, the only problem with this hypothesis is that it also implies overprediction for low coupon mortgages. That is, in the past there has been suboptimal prepayment of low coupon mortgages. To the extent that consumers shift to more optimal exercise of their prepayment option, the smart trader will also predict slower prepayment for these low coupon mortgages. We do not encounter this overprediction in our data. But it is possible that this effect is small and that our data sample is too short.

Notice that if the bias is simply proportional to $c^k - r$ then the fixed-effects/time-effects specification will control for this possibility. So again, the only possibility that we need to address is if the bias depends on both c^k and r. The results of Table 6 control for this possibility.

5.3 The quadratic dependence on interest rates

At some level, the mis-specification possibility will always cast a shadow on our results. Our only defense is to demonstrate results that fit our theory easily, and only fit a very contorted theory of mis-specification. So far, we have shown that any alternative which predicts purely a fixed effect or a time effect will not affect our results. An alternative involving the interaction between coupon and market interest rates will also not affect our results.

There is one further result of our theory that a mis-specification theory would have to work very hard to replicate. Our equilibrium theory predicts that the market price of risk should vary with the average market coupon. Plausibly, a mis-specification theory will only

[14]We thank John Geanokoplos for this suggestion.

link security specific attributes (e.g., the coupon of the specific security being studied) and the market interest rate to the OAS.

Unfortunately, there is little variation in the average market coupon over the 1993 to 1998 sample. However, from Figure 3 we note the large change in the average coupon over the 1992 to 1993 period. We have OAS data for passthroughs over this period from Smith-Breeden. We use this data to test the quadratic relation for the collateral (see equation (17)) as well as to check whether the average coupon has explanatory power for this data.

The data is for the OAS on FNMA 30-year generic collateral for 8 bonds with coupons ranging from 7.5% to 11%. Our data spans a period from October 1987 to July 1994. We estimate the following regression:

$$OAS_C^{kt} = \gamma_k + (\alpha_1 \bar{c}_t + \alpha_2 r_t + \alpha_3) \times (c^k - r_t) + \epsilon_{kt}$$

The main prediction of our theory is that α_1 is positive and that α_2 is negative. We will take a positive value of α_1 to imply that the average market coupon is an important explanatory factor. The theory also predicts that $\alpha_1 + \alpha_2 = 0$, and that α_3 should not have any explanatory power.

The results are reported in Table 7. The first set of regressions are run separately by bond. The last regression combines all of the data in a panel, and implicitly sets the β^k loadings for each security equal to each other.

The coefficients on $\bar{c}_t$ are uniformly positive and significant, as predicted. It is also encouraging that the magnitudes of α_1 and α_2 have opposite signs, and the $\alpha_1 + \alpha_2$ is close to zero often. α_3 is negative and sometimes significant.

A possible explanation for this occasional discrepancy is that a true measure of the model's $\bar{c}$ would include expected values of future coupons. As in the 1980's and 1990's nominal rates were largely declining, a weakly negative α_3 captures the market's expectation that in the future, the average coupon is going to decrease. Alternatively, this occasional discrepancy could be due to a mis-specification of the interest rate in our simple empirical implementation.

Table 7					
Regressions based on the OAS of the collateral: $OAS_C^{kt} = \gamma_k + (\alpha_1 \bar{c}_t + \alpha_2 r_t + \alpha_3) \times (c^k - r_t) + \epsilon_{kt}$					
Bond Coupon	α_1	α_2	α_3	R^2	N
7.5	13.44 (1.9)	-9.37 (-2.08)	-54.6 (-1.38)	0.26	26
8	19.45 (5.3)	-14.17 (-4.48)	-73.45 (-2.52)	0.46	28
8.5	14.91 (4.51)	-13.51 (-3.96)	-41.97 (-1.51)	0.44	28
9	7.69 (5.09)	-9.54 (-4.77)	-10.35 (-0.72)	0.45	28
9.5	9.30 (7.5)	-9.02 (-4.31)	-27.29 (-2.02)	0.51	28
10	9.66 (6.79)	-8.37 (-3.12)	-34.02 (-2.31)	0.49	28
10.5	11.76 (6.83)	-7.73 (-2.78)	-59.23 (-4.1)	0.51	28
11	7.68 (3.61)	-6.99 (-1.83)	-24.13 (-1.19)	0.33	28
ALL BONDS	10.03 (10.41)	-11.88 (-3.17)	-17.25 (-0.83)	0.87	222
ALL BONDS regression uses the entire panel, with bond fixed effects. Estimates reported with T-statistics based on robust standard errors in parentheses.					

5.4 Actual MBS bond returns

We should note that one way to get around the mis-specification issue is to use actual bond returns as the dependent variable in our regressions. There are a few reason we have not done this. Actual bond returns are a very noisy estimate of the expected return on the securities.[15] We need more data than we have to implement these regressions. Using the OAS greatly reduces this measurement error problem. Breeden (1994) provides support for this approach. He studies a large panel of GNMA securities and finds that the OAS has strong predictive power for the subsequent returns. The results are reported in Exhibits 72, 73, and 74 of Breeden (1994). We note Exhibit 74 in particular which demonstrates that the strongest relations is between the OAS for IO's and subsequent returns.

6 Relation to the MBS literature

The academic work on MBS valuation is primarily concerned with prepayment modeling. In one line of research, prepayment stems from rational choice by homeowners. This "rational" prepayment approach was pioneered by Dunn and McConnell (1981) and investigated more

[15]We can reduce the noise in bond returns if we take a stand on the mortgage prepayment model and calculate interest rate hedge-ratios. Then we can strip out the interest rate component of actual bond returns. But this seems no better than using the OAS from the prepayment model of a dealer, as we have done.

recently by Stanton (1995) and Longstaff (2002)[16]. In the other main line of research (and in the practitioner approach), prepayment behavior is modeled statistically. The justification for this approach is that, given the complexity of the constraints faced by consumers, prepayment behavior on a pool of consumer mortgages is better captured statistically than by modeling these complex constraints. Examples of the latter approach include Schwartz and Torous (1989), Richard and Roll (1989), and Patruno (1994).

Our research suggests that it is also necessary to model the uncertainty surrounding prepayment behavior, which arises naturally once we recognize that homeowners' cost of refinancing, for example, will be subject to innovations. In our approach, we directly model this prepayment uncertainty as an error around a mean prepayment forecast. However there are many other ways of introducing this prepayment uncertainty, in both the rational as well as the statistical approach. The important point we make however is that this uncertainty is priced and that the market price of this uncertainty varies in a systematic way.

Boudoukh *et al.* (1997) use a novel approach to directly estimate the pricing function for a panel of GNMA passthroughs as a function of one or two factors (long rate, long rate and spread). They estimate this function non-parametrically using kernel techniques and find that the two-factor pricing model performs better (the level captures the refinancing incentive and the spread proxies for the future behavior of the discount factor used to discount cash-flows). They only use information in market prices, and focus on a pricing functional which depends on the yield curve[17], thus setting aside prepayment information. One of their main findings is that the pricing errors under the model have a large common factor. Our study suggests that prepayment risk[18] and the average coupon outstanding in the market are this common factor.

While the bulk of the academic literature does not address the OAS, market participants and academics writing in more applied journals recognize that there are important time patterns in the OAS. Breeden (1994) provides extensive documentation of how the OAS methodology performs in rich/cheap analysis and hedging the interest rate risk of MBS portfolios. Breeden concludes that effective durations (duration keeping the OAS constant) help to reduce risk of pass-throughs and PO's by 40% and 25%, but perform poorly for the more risky IO's. For the latter, hedging is improved by using an empirical duration (i.e. a statistical estimate of the price elasticity). The correlation between mean returns and OAS is, however, higher for IO's than for PO's and pass-throughs. A closer look at the data

[16]See Kau and Keenan (1995) for a survey of this line of research.

[17]Their choice of variables is also constrained by the trade-off between adding variable to the function and increasing pricing errors.

[18]Boudoukh *et al.* do hint at this by looking at the prepayment of the different coupons. They find that for lower coupons, which have a lot of relocation-based prepayment, prepayment variables explain a significant fraction of the pricing errors.

gathered by Breeden reveals that the downward bias of the effective duration in measuring real duration for collateral is systematic across securities in the discount period of 1992-1993. Our approach recognizes that OAS should not be considered as a random pricing residual, and proposes a model that links the uncertainty of prepayment to the OAS and, solving the market equilibrium, derives their dynamic evolution. For example, we suggest a change in the common practice of computing duration. The usual practice is to compute durations under a prepayment model by changing the interest rate, while holding the OAS fixed. Our works suggests that the OAS is also a function of interest rates and therefore current practice will yield a biased duration.

7 Conclusion

We provide theory and evidence that the marginal investor in the mortgage-backed securities market is a trader who is wholly invested in the mortgage market. The theory predicts that prepayment risk – a risk that is diversifiable in the aggregate – is priced and that the pricing of this risk depends on the value of the entire mortgage market.

The evidence is consistent with recent "limits of arbitrage" theories. If the marginal investor is a mortgage trader, then other sources of variation in the mortgage trader's marginal utility will also affect MBS prices. For example, if hedge funds are capital constrained then the capital deployed by hedge funds in the MBS market will be an important state variable in determining prices. Thus the increase in MBS spreads in the Fall of 1998 also fits with our evidence.

If a limited amount of capital sets prices in the MBS market then changes in supply will have important price effects. Customer sales/mortgage securitizations have to be absorbed by this limited capital, and one would naturally expect prices to fall in such events. That is, the market will exhibit liquidity effects. The liquidity of the market however has nothing to do with bid-ask spreads; It reflects the limited capital of MBS traders.

Although our evidence suggests market segmentation, it suggests a fairly subtle form of segmentation. MBS traders are able to hedge out all interest rate risk via the Treasury bond market. They only retain the prepayment risk in mortgages. That is the interest rate component of MBS is priced in an integrated market, only the prepayment risk is priced locally to the mortgage market.

Our approach also raises a host of issues that could be of interest to academics, both empirical and theoretical. First, and most importantly, more empirical work should be done to investigate the issues we have highlighted. We have established our results in a fairly small sample. Secondly, better prepayment models should be constructed that pay attention to the variance of the errors in the prepayment forecasts. Models developed thus far just assume, because they assume rational prepayments, a zero uncertainty in the prepayment side. Our analysis highlights that when prepayment uncertainty is non-zero, prices will be affected.

Finally, the success of our approach for MBS also suggests new directions to explore in other asset markets. We think that this approach may hold some promise in the bond market. We conjecture that the term premium in the bond markets should be an increasing function of the average maturity outstanding in the bond market. The intuition is that when the average maturity outstanding is high, the market is riskier (the duration is higher), so that the risk premia, and hence the slope of the yield curve, should be higher. More work is needed to assess these conjectures.

References

[1] Allen, F. and D. Gale, 1994, "Limited Market Participation and Volatility of Asset Prices," *American Economic Review* 84(4), 933-955.

[2] Boudoukh, J., M. Richardson, R. Stanton, R. Whitelaw, 1997, "Pricing Mortgage-backed Securities in a Multifactor Interest Rate Environment: A Multivariate Density Estimation Approach," *Review of Financial Studies*, 10, 405-446.

[3] Breeden, D. T., 1991, "Risk, Return and Hedging of Fixed Rate Mortgages," *Journal of Fixed Income*, 1, 85-107.

[4] Breeden, D.T., 1994, "Complexities of Hedging Mortgages", *Journal of Fixed Income*, 4, 6-41.

[5] Caballero, R. and A. Krishnamurthy, 2001, "International and Domestic Collateral Constraints in a Model of Emerging Market Crise" *Journal of Monetary Economics* 48(3), 513-548.

[6] Caballero, R. and A. Krishnamurthy, 2002, "A Dual Liquidity Model of Emerging Markets," *American Economic Review*, 92 (2), 33-37.

[7] Collin-Dufresne, P., B. Goldstein, and S. Martin, 2001, "The Determinants of Credit Spread changes," *The Journal of Finance*, 66(6).

[8] Dow, J. and G. Gorton, 1994, "Arbitrage Chains," *Journal of Finance* 49(3), 819-849

[9] Dunn, K.B. and J.J. McConnell, 1981, "Valuation of GNMA Mortgage-Backed Securities," *Journal of Finance*, 36, 599-616.

[10] Froot, K. 2001. "The Market for Catastrophe Risk: A Clinical Examination," *Journal of Financial Economics*, 60, 529-571.

[11] Froot, K., and P. O'Connell, 1999, "The Pricing of US Catastrophe Reinsurance." In *The Financing of Catastrophe Risk,* edited by K. Froot, University of Chicago Press, 195-232.

[12] Gabaix, X., 2003, "Bond Pricing when Supply and Demand Matter", MIT mimeo.

[13] Gromb, D. and D. Vayanos, 2002, "Equilibrium and Welfare in Markets with Financially Constrained Arbitrageurs, *Journal of Financial Economics*, 66, 361-407.

[14] Holmstrom, B. and P. Milgrom, 1987. "Aggregation and Linearity in the Provision of Intertemporal Incentives," *Econometrica*, 55(2), 303-28.

[15] Kau, J., and D. Keenan, 1995, "An Overview of the Option-Theoretic Pricing of Mortgages", *Journal of Housing Research*, 217-44.

[16] Krishnamurthy, A., 2002, "The Bond/Old-Bond Spread," *Journal of Financial Economics*, 66(2), 463-506.

[17] Kyle, A. and W. Xiong, 2001, "Contagion as a Wealth Effect", *Journal of Finance* 56, 1401-1440.

[18] Longstaff, F., (2002), "Optimal Recursive Refinancing and the Valuation of Mortgage-Backed Securities", UCLA Mimeo.

[19] Merton, R.C., 1987, "A Simple Model of Capital Market Equilibrium with Incomplete Information", Journal of Finance, 42 (3), 483-510.

[20] Patruno, G., 1994, "Mortgage Prepayments: A New Model for a New Era," *Journal of Fixed Income*, December.

[21] Richard, S. F., and R. Roll, 1989, "Prepayment on Fixed-Rate Mortgage-backed Securities," *Journal of Portfolio Mangement*, 15, 375-392.

[22] Schwartz, E. S., and W.N. Torous, 1989, "Prepayment and the Valuation of Mortgage-Backed Securities," *Journal of Finance*, 44, 375-392.

[23] Shleifer, A. and R. Vishny, 1997, "The Limits of Arbitrage," *Journal of Finance*, 35-55.

[24] Stanton, R. H., 1995, "Rational Prepayment and the Valuation of Mortgage-Backed Securities," *Review of Financial Studies*, 8, 677-708.