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

Xavier Gabaix Arvind Krishnamurthy Olivier Vigneron*

September 18, 2004

Abstract

"Limits of Arbitrage" theories require that the marginal investor in a particular asset market is a specialized arbitrageur rather than a diversified representative investor. We examine the mortgage-backed securities (MBS) market in this light. We show that the risk of homeowner prepayment, which is a wash in the aggregate, is priced in the MBS market. The covariance of prepayment risk with aggregate consumption or wealth implies the wrong sign to match the observed prices of prepayment risk. The price of risk is better explained by a pricing kernel based on the aggregate value of MBS. The evidence suggests that there are limits to arbitrage in the MBS 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 Geanakoplos, Massoud Heidari, Jingzhi Huang, Ravi Jagannathan, Adam Kolasinski, Ananth Krishnamurthy, Francis Longstaff, Andrew Metrick, Anna Pavlova, Tim Riddiough, José Scheinkman, René Stulz, Raman Uppal, Dimitri Vayanos, Bill Wheaton and an anonymous referee. We thank participants at the NBER Asset Pricing meeting, Econometric Society Summer meeting, Gerzensee-CEPR meeting and seminars at the University of British Columbia, University of Chicago, MIT, University of Maryland, and Northwestern for their comments. We also thank Salomon-Smith-Barney and Smith-Breeden for providing 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

A number of recent theories ("limits of arbitrage") have been forwarded wherein the distressed liquidation of assets by hedge funds results in a large drop in asset prices (see, for example, Shleifer and Vishny (1997), Kyle and Xiong (2001), Gromb and Vayanos (2002), or Geanakoplos (2003)). These theories are often referenced in explaining episodes of market illiquidity and falling asset prices, such as the events of fall of 1998.

In traditional asset pricing theory, the marginal investor in every asset market is the same broadly diversified representative investor. Thus, according to traditional theory a hedge fund liquidating $200 billion of mortgage backed securities finds a large pool of ready buyers (comparable to the entire capital market). The liquidation does not affect expected returns because the representative investor acts quickly to eliminate excess returns. According to the limits of arbitrage theory, on the other hand, the relevant set of buyers is a smaller specialized pool of investors and the liquidations can have large effects on prices. The limits of arbitrage theory posit that the marginal investor in a particular asset market is an investor who specializes in that market.

We provide support for the limits of arbitrage theory by examining the mortgage backed securities (MBS) market. MBS securities rise and fall in value based on the exercise of homeowners' prepayment options. When a homeowner prepays a mortgage, the MBS backed by the mortgage is called back at par. Depending on the interest rate environment, prepayment can either hurt or benefit the investor who owns the MBS. Thus, for an investor who specializes in the MBS market, prepayment risk represents a risk to the value of his portfolio. At the aggregate level, prepayments do not cause changes to aggregate wealth or the aggregate endowment, since for every MBS investor who is short a prepayment option, there is a homeowner who is long the prepayment option. Any observed covariance between aggregates and prepayments is due to some common economic factors driving both aggregates and homeowner prepayments. In the traditional asset pricing theory, the covariance between prepayments and aggregate wealth or consumption explains the price of prepayment risk.

We establish three principal results in the paper. First, we show that prepayment risk is priced. Second, we show that the observed covariance between prepayment risk and either aggregate wealth or consumption has a sign opposite to what traditional asset pricing theory implies given the observed prices of prepayment risk. This suggests that the marginal investor in the MBS market is not the representative consumer hypothesized by the traditional CAPM or consumption-CAPM model. Finally, we derive a proxy for the marginal utility of an investor who is wholly specialized to the MBS market and show that the market price of prepayment risk has a systematic relation with this marginal utility proxy.

Taken together, these results support the existence of limits to arbitrage in the MBS market. We argue that the marginal investor in the MBS market is a hedge fund or mutual fund that trades just in the MBS market. We present a simple model of delegated fund management in which a fund manager who has his wealth tied up in the MBS market is the marginal investor in the MBS market. We show that when the manager's coefficient of relative risk aversion is 4, the model's empirical predictions are consistent with what we find in the data.

An important theme in the limits of arbitrage literature is that negative shocks to the capital of hedge funds cause them to liquidate assets, which results in higher expected returns on these assets. In the MBS market the events of 1993-1994 (collapse of the Askins Capital Management MBS fund) and 1998-1999 (LTCM crisis) in which spreads on MBS widen substantially, are usually taken to be a sign of a capital-related shock. Although we do not provide direct evidence of capital-related movements in MBS prices, we do establish a necessary condition for a capital-related shock to effect a large move in MBS prices. Namely, the marginal investor must be an MBS specialist rather than a broadly diversified representative investor.

The evidence presented in this paper is consistent with what is presented in other recent papers. Collin-Dufresne, Goldstein, and Martin (2001) study the corporate bond market. They find that a simple Merton (1973) model explains very little of the variation in corporate bond prices. Even after including macro factors (stock market, etc.) they are only able to explain about 25% of price variation. The tantalizing evidence they present is that the bulk of the remaining variation is due to a single risk factor that is common across all corporate bonds.[1] Unlike us, they are unable to identify either the risk factor or the marginal investor who is pricing the risk.[2]

Boudoukh, Richardson, Stanton, and Whitelaw (1997) provide similar evidence for the MBS market. They 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. Our results suggest that a candidate for this common factor is a prepayment risk factor.

Froot and O'Connell (1999) demonstrate effects similar to the ones we identify by studying the market for 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. Using an argument similar to ours, they assert that the marginal investor in the catastrophe insurance market is a specialized institution (an insurer) rather than the broadly-diversified representative investor. They can therefore explain insurance price spikes as follows: when the capital in the insurance market is low, insurers are less willing to write catastrophe insurance driving prices up and quantities down.

Merton (1987) presents a model in which segmentation arises endogenously, and he explores 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 affect asset prices. Allen and Gale (1994) study an environment in which traders must specialize ex-ante in a certain asset market, which 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), Gromb and Vayanos (2002), and Geanakoplos (1997, 2003). Gabaix et al. (2003) propose to explain the structure of extreme movements in returns and trading volume by the actions of very

[1] Berndt, Douglas, Duffie, Ferguson, and Schranz (2004) present similar evidence based on data from credit default swaps. They find large swings in the risk-premia incorporated in default swaps.

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

large traders in illiquid markets. Caballero and Krishnamurthy (2001, 2002) propose a model of emerging market crises in which the crises is an event in which the marginal investor switches from a broadly diversified world investor to an investor within the emerging market.

The next section of the paper provides a brief description of the mortgage market. Section 3 defines prepayment risk and studies its effect on the value of MBS by presenting a model that formally links a security's risk premium to its prepayment risk and the marginal utility of an investor who is wholly invested in the MBS market. We show that the difference between the average market coupon and interest rates ("the market price of prepayment risk") proxies for the marginal investor's marginal utility. Our model is motivated by a simple agency friction. We show that the form of the market price of prepayment risk arises naturally in this setting.

Section 4 presents evidence that supports our model. We verify empirically the relation, predicted by our model, between a security's risk premium, its measured prepayment risk, and the difference between average market coupon and interest rates. In Section 4.5 we discuss the alternative representative household hypothesis for pricing prepayment risk and show that this hypothesis fails to explain what we find in the data. Section 5 provides robustness checks of our results. The conclusion and an appendix follows.

2 Overview of the mortgage market

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

Mirroring the underlying consumer mortgage, the MBS is a debt security with a declining principal value. The fact that consumers have the option to prepay their mortgages, however, makes valuing (and hedging) MBS very difficult. Consumer 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. Our study focuses on prepayment risk and its pricing.

The securities we study in this paper are known as collateralized mortgage obligations (CMO). A typical CMO has several tiers, known as "tranches," each with a different degree of prepayment risk. All tranches receive interest payments, but principal payments go first to bonds in the top tier until they are entirely repaid, and then to the next tier, etc. Thus prepayment risk is carved up differently among the tranches. The upper tranches have shorter and more certain maturities, and therefore lower prepayment risk. The natural buyers of these tranches are pension funds, insurance companies and other large institutional investors requiring relative safety. The lower tranches have longer maturities and therefore assume greater prepayment risk. These lower tranches ("toxic waste") are especially volatile and hard to price. The natural buyers are sophisticated investors such as hedge funds or investment banks who have some expertise in assessing prepayment risk. The success of a securitization of mortgages often turns on finding sophisticated investors willing to hold toxic waste.

A single CMO tranche typically passes both interest and principal payments of the underlying pool of loans, in some prespecified manner, to the investor. Often a security is created which passes only the principal repayments (PO) or only the interest payments (IO) to the investor. Such a security may be a separate tranche or a derivative stripped from a mortgage.

Valuing an MBS typically involves two-steps. 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). Second, 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, the OAS is a spread added to the riskless rate such that the present value of a security's expected cash flows, forecast using the prepayment model and discounted using the riskless rate plus the spread, equals the price of the security (equation (7) below). To the extent that the term-structure model is correct, the OAS constitutes the "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. In our empirical methodology we use the OAS as a proxy for the risk premium on the MBS. If, however, the benchmark prepayment model used to compute the OAS is mis-specified, then the OAS will at least partly reflect the value of the prepayment option as well as the risk premium (e.g., see Kupiec and Kah, 1999). Model misspecification, therefore, would imply that the OAS is a noisy proxy of a security's risk premium, potentially rendering the interpretation of our results suspect.

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 (and no credit risk). 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 $a(0) = 1$.

Suppose that there is a single class of MBS (pass-through, IO-derivative, and PO-derivative) issued against this pool. 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 pass-through 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.[3] 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,

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

3.3 OAS

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

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

[3] 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.

Where, 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.

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.[4]

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,

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

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.[5]

[4] It is also possible that the OAS is due to a Jensen's inequality term. However, we think that this Jensen's inequality effect is unlikely to be very important. Indeed, since, $P_{IO} = E\left[c/\left(r + \widetilde{\phi}\right)\right] > c/\left(r + E\left[\widetilde{\phi}\right]\right)$ this interpretation predicts a negative OAS for the IO. In our sample, the OAS of IOs are almost always positive, which means that this effect is probably small. Also, the Jensen's inequality effect predicts that the OAS depends only on the security-specific factors, not on market-wide factors as we find in our empirical work.

[5] The OAS for the PO is defined by

$$P_{PO}^k = 1 - \frac{r + OAS_{PO}^k}{r + \bar{\phi}^k + OAS_{PO}^k}$$

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,

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

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.

3.4 The marginal investor

The critical assumption that we make – and for which we provide tests – is that a representative and specialized MBS fund manager is the marginal investor in this market.

For the sake of internal consistency, we motivate this assumption using a model of agency and delegated fund management. As will become clear, we do not provide any explicit tests of agency. Thus the agency model should only be viewed as an organizing principle.

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 fund managers (*"fund managers"*) with endowments of w_T. 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 fund manager who invests for them.

Investors require that the fund manager contribute a fraction of his own endowment for every dollar that the investor provides. We think of this as a capital requirement that ensures that the fund manager invests prudently. Let us define αw_F as the capital requirement for a fund manager for a fund of size w_f.[6] Thus for each dollar of his wealth, the fund manager runs a fund of size $\frac{1}{\alpha}$.

The problem is that the fund manager is risk averse. He has utility over date 1 wealth of,

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

i.e. just a mean-variance maximizer.

3.5 Equilibrium

At date 0, fund managers raise a total of $w_F - w_T$ from investors. This gives them total capital of $\frac{w_T}{\alpha}$. With this sum the fund managers 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,

$$W_F = \frac{w_T}{\alpha} + \sum_k x^k_{IO}(V^k_{IO} - P^k_{IO}) + \sum_k x^k_{PO}(V^k_{PO} - P^k_{PO}) \tag{12}$$

is the date 1 value of the portfolio. Since the fund manager's wealth increases linearly with W_F (at slope of α), his problem is to maximize (11) given (12), and subject to the budget constraint that,

$$\frac{w_T}{\alpha} \geq \sum_k x^k_{IO} P^k_{IO} + \sum_k x^k_{PO} P^k_{PO}.$$

In our derivation, we assume that the fund manager has sufficient wealth, or that α is sufficiently low, so that the fund manager is not capital constrained in purchasing his desired portfolio of MBS.

This formulation is a variant of the traditional static CAPM. Deriving the first order condition for the fund manager'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,

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

[6] See Holmstrom and Tirole, 1997, for a model of capital constraints in intermediation.

where the market 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. We note the dependence of the risk premium on α. When $\alpha = 0$ the MBS fund manager is a "veil," and the marginal investor is the risk-neutral investor. When $\alpha = 1$, the MBS fund manager is the only investor in the MBS market.

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,[7]

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

The sum term is difficult to observe empirically. It is a weighted sum of the coupons of all mortgages in the market, where the weights depend on the amounts outstanding 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,[8]

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

where $\bar{c}$ is the weighted average coupon and $\rho\alpha$ is the effective risk aversion of the MBS fund manager (a is a constant of proportionality). 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.

[7] 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 $\eta^k = 1/(r + \bar{\phi}^k)$.

[8] 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. For details, see the Technical Appendix to this paper at $http://econ-www.mit.edu/faculty/xgabaix/papers.htm$.

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 fund manager).

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 fund manager. 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,

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

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,

$$OAS^k_C = \beta^k \times \; \rho\alpha a \left(\bar{c} - r\right)\left(c^k - r\right) \; \times \; \frac{1}{c^k + \bar{\phi}^k + OAS^k_{IO}} \tag{17}$$

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.

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:

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

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

where $\rho\alpha$ is a constant proportional to the risk aversion of the fund managers. Further implications are as follows:

- 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, $\bar{c}_t$.

3.8 Discussion of assumptions

The model we have presented is simplified along many dimensions. We comment on some of these simplifications in this subsection.

At a broad level, the main result of our simplifications is the equation for the IO that OAS_{IO} is proportional to $\beta^k \times (\bar{c} - r)$. This is the relation we test in the empirical work.

From the standpoint of an empirical test, we think that this relation should be robust to more sophisticated models. We predict that when $\bar{c}$ is high relative to r that the aggregate market is very sensitive to prepayment risk and securities with more prepayment risk should carry a high risk premium. On the other hand when $\bar{c}$ is close to r there should be less sensitivity to prepayment risk. While probably not in the linear form that we have assumed, this relation seems like it should survive most models.

On the other hand, the simplification in the derivation means that there are probably other factors affecting the OAS. It is likely the case that in practice the OAS is affected by the optionality of the securities and the history dependence of mortgage prepayment. For example, Brown (1999) notes a positive relation between OAS and implied volatilities on Treasury bond options which suggests that there is a mis-specification in Wall Street prepayment models used to derived the OAS. These issues suggest that the OAS is noisy measure of a security's risk premium and our empirical tests may need to control for these other factors. The controls are discussed in far greater depth in the robustness section.

We have omitted capital constraints from the model, which is a substantive effect in the limits of arbitrage literature. Mainly, this is because we do not provide any direct empirical tests of capital effects. Informally, we can think of capital constraints as raising the effective risk aversion (ρ) of the fund managers. For example, if one-half of the fund managers lose all of their capital so that they are no longer active in the MBS market, the rest of the fund managers will, in equilibrium, bear twice the amount of risk and will therefore demand a higher risk premium. In the next section, we provide some evidence that early on in our sample period the risk premia are higher. The early period also corresponds to the Askins Capital Management hedge fund crisis.

In the model, the fund manager is risk-averse and receives a linear share of profits. In part, we make this assumption because we are interested in exploring the limits of arbitrage. As in Shleifer and Vishny (1997), the effective risk aversion of the manager limits his ability to exploit high returns. In practice, however, monetary compensation contracts are convex which can lead to risk-taking behavior in regions where the fund manager is near the kink in his payoff, and risk averse behavior in other regions. Whether risk-taking behavior is the rule rather than the exception is ultimately an empirical question. Our results suggest that it is the exception, but it would be interesting to empirically explore non-monotonicities in the behavior of fund managers in the MBS market.

We have derived our results in a static CAPM framework. In a dynamic model the current wealth of the fund managers will be an important state variable. To the extent that the aggregate value of the mortgage market is a sufficient statistic for the marginal utility of the representative fund manager, 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.

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 liquid securities and fairly representative in age and coupon of the active secondary market. The collateral are all FNMA 30-year conventional loans. The collateral are uniformly drawn from a mix of loans from across the country. The largest representation is from California, New York, Texas, Florida, and Illinois.

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 the underlying pools.

The bulk of our analysis is conducted using the IO data. We have also checked our results using the PO data and the results are consistent with the IO evidence, albeit a little less strong. The results are not reported but are available upon request.

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 outstanding, 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}). The dependence on past interest rates is a feature of most prepayment models. Longstaff (2004) has shown that this feature arises naturally in a setting with transactions costs of refinancing. We assume that the error follows an AR(1) process,

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

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. The Appendix presents the actual and fitted estimates for two of the securities, along with forecasts from a Wall Street prepayment model.

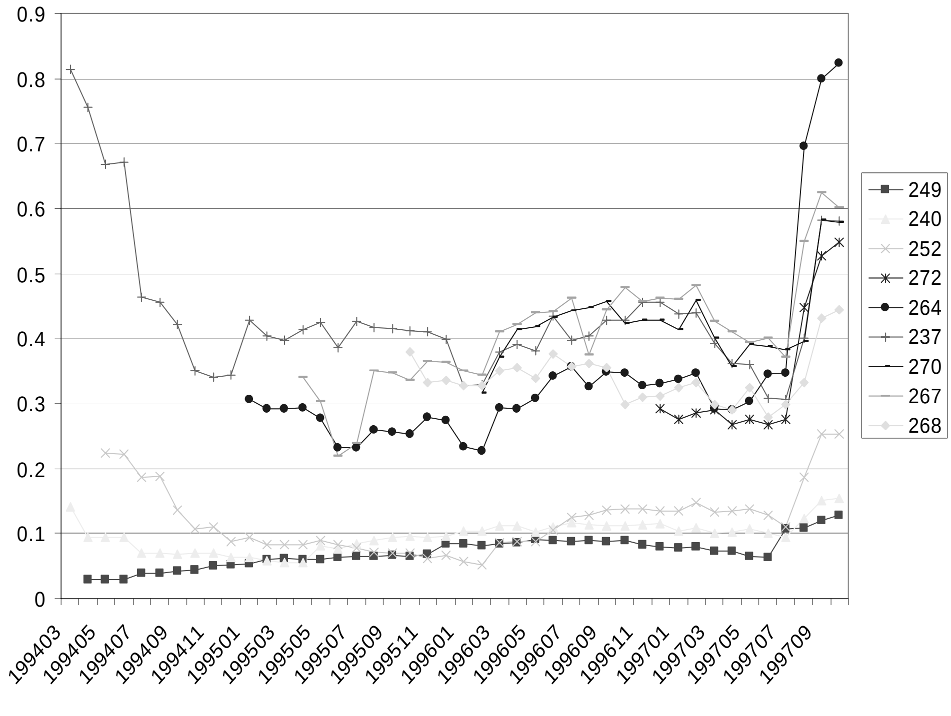

Figure 1: Standard deviation of prepayment errors

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 1 shows the time series of

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

The figure presents 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)	0.095	0.174	0.357	0.603	1.071	0.821	0.890	1.369	0.935

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.[9] 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 (the correlation coefficient is 0.88), 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 Wall Street firm and have used these residuals to form β's. The β's look similar, suggesting that our model is reasonable. See the last line of Table 2.

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 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 2 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

[9] 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.

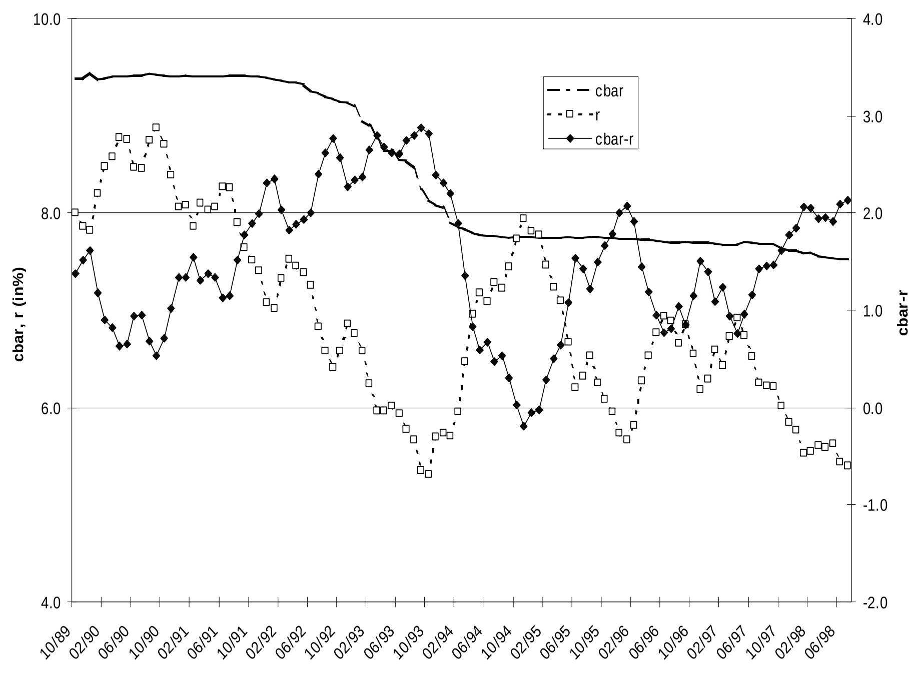

Figure 2: CMT 10-year and Average Coupon

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

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

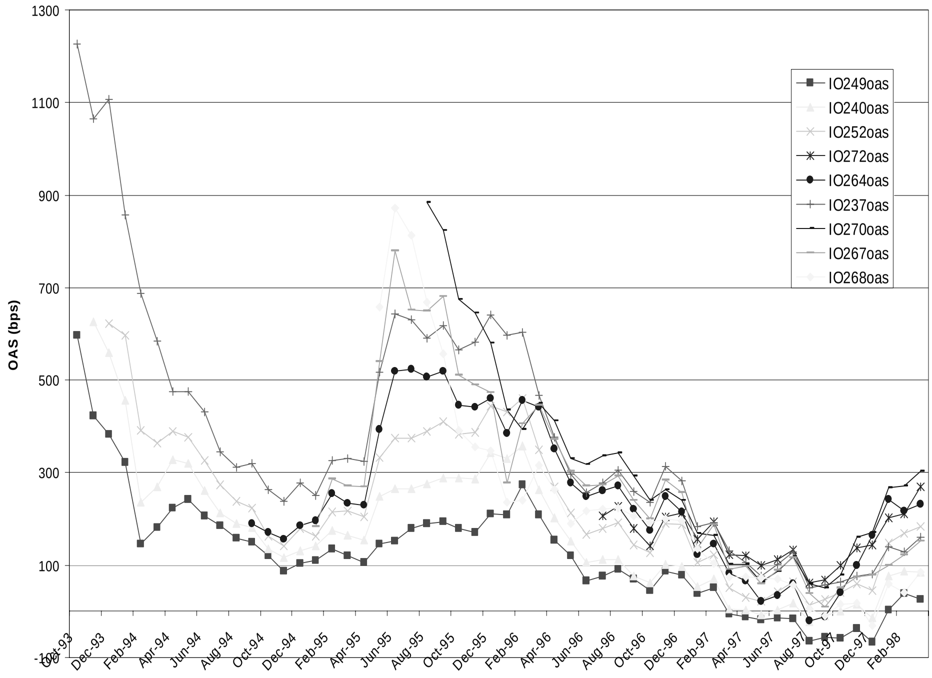

Figure 3: OAS of the IOs

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.

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. The α_t term may pick up any common variation due to mis-specification problems in the OAS. Alternatively, it may pick a time varying interest rate risk premium or a time varying premium due to shortages of arbitrageur capital.

The λ_t estimates along with other sample statistics using β-stdev are in Table 3. The estimation

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

errors are uniformly tight and the R^2 are high.[10] 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. We present the results for two of the months in figures in the appendix.

Figure 4 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

[10] Jan and Feb 96 seem to have unusually low R^2. Our best guess is that there are some data errors in these months.

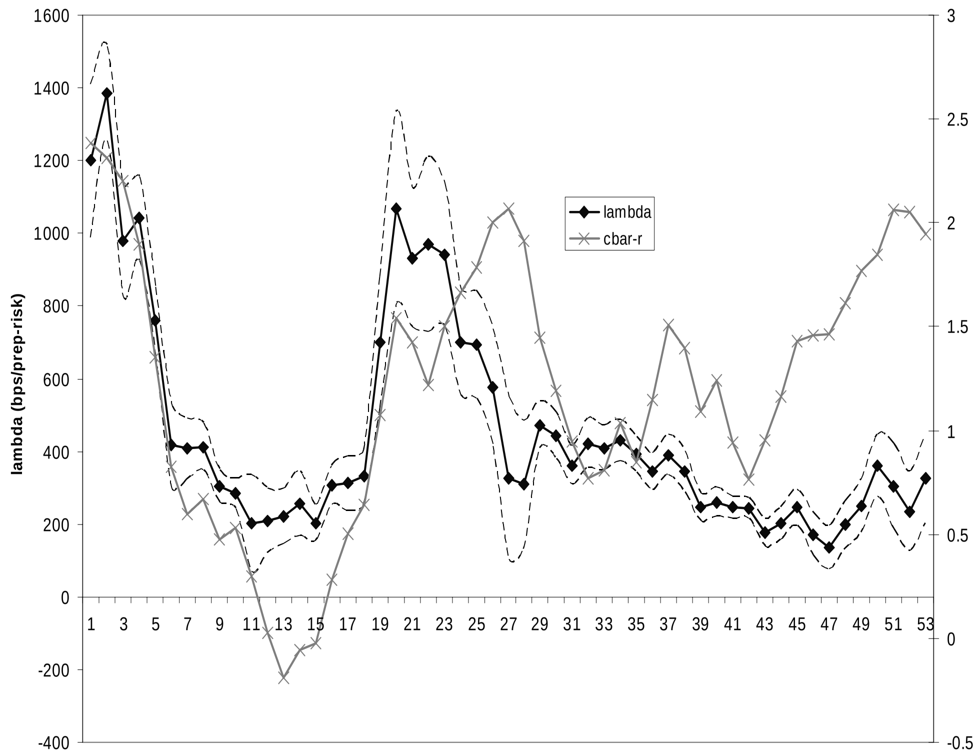

Figure 4: λ estimates using β-stdev

The measure of the monthly market price of prepayment risk (λ), obtained from estimating Eq. (18), is plotted. We use β-stdev as the measure of security-specific prepayment risk. To compare this empirical measure to our theoretical prediction of Eq. (19), we plot $\bar{c}_t - r_t$. $\bar{c}_t$ is the average coupon outstanding in the MBS market, and r_t is the 10 year constant maturity Treasury (CMT) interest rate. The empirical and theoretical value of the market price of risk move together. This is confirmed by empirical tests reported in the rest of the paper.

back into these funds and ρ fell. Froot (2001) finds this effect in the catastrophe insurance market. We intend to further investigate this effect in future work.

Figure 5 graphs the estimate of λ using the β-PCA. The results are similar to those of Figure 4.

4.3 Tests using the entire panel

We now report the results of testing our model using the entire panel. Table 4 reports regressions based on the following model.

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

The regression includes both time and security effects, thereby controlling for any alternative that involves either security specific effects or time specific effects. We discuss alternative hypotheses in

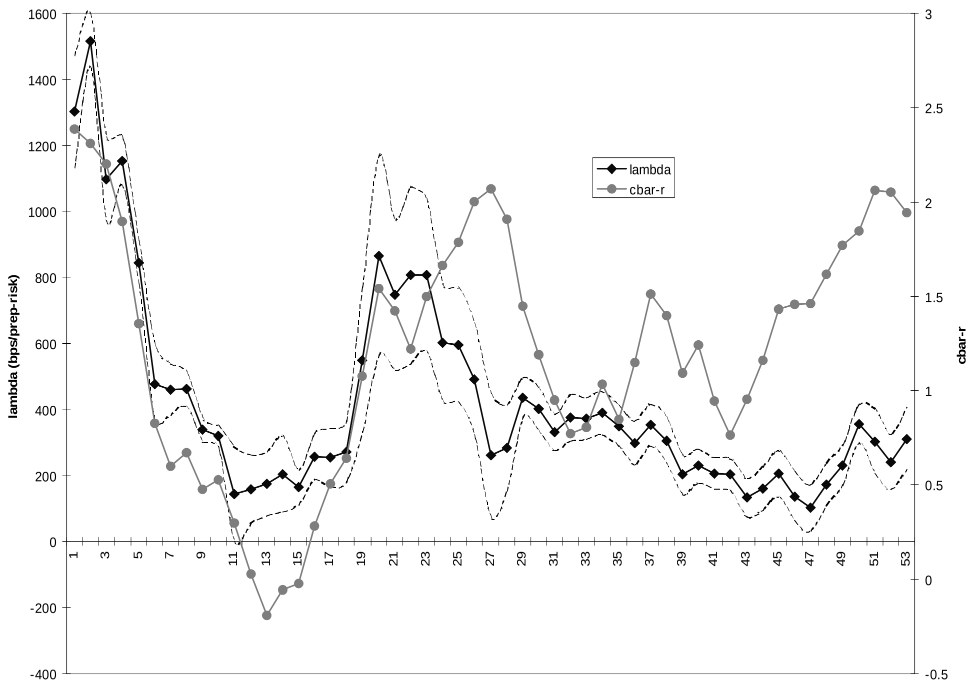

Figure 5: λ estimates using β-PCA

The measure of the monthly market price of prepayment risk (λ), obtained from estimating Eq. (18), is plotted. We use β-PCA as the measure of security-specific prepayment risk. To compare this empirical measure to our theoretical prediction of Eq. (19), we plot $\bar{c}_t - r_t$. $\bar{c}_t$ is the average coupon outstanding in the MBS market, and r_t is the 10 year constant maturity Treasury (CMT) interest rate. The empirical and theoretical value of the market price of risk move together. This is confirmed by empirical tests reported in the rest of the paper.

greater depth in the next section.

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 T-statistics which are corrected for clustering in the standard errors at the security level. Second, we run regressions using first-differenced data and report the results in Table 5. Another potential problem is correlation across securities at a single point in time. This problem, however, is less severe in our specification because the regressions include a time dummy that absorbs all common innovations in the OAS. We have also done a robustness check using a standard panel data adjustment where we assume that errors are AR(1) at the security level and correlated across securities. We find that our results remain highly significant.

Table 4

Regressions based on the OAS of the IO's:

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

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

Results by subsample are reported in (2) (first-half) and (3) (second-half). The break point is June 1996.

	β^k-stdev						β^k-PCA		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$\beta^k(\bar{c}_t - r_t)$	438.1 (4.13)	546.5 (5.81)	315.1 (5.57)	366.8 (5.53)	420.7 (5.55)	242.9 (3.98)	405.9 (2.73)	218 (0.81)	221.1 (3.36)
$idiosync^k\,(\bar{c}_t - r_t)$								576 (0.59)	
Security Effects	Yes	Yes	Yes	No	Yes	No	Yes	Yes	No
Time Effects	Yes	Yes	Yes	Yes	No	No	Yes	Yes	No
R^2	0.93	0.95	0.92	0.92	0.65	0.15	0.93	0.93	0.14
N	383	194	189	383	383	383	383	383	383

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

The results in columns (1) - (8) of Table 4 verify that our model fits the data. The specification in column (1) uses β-stdev, while the specification in column (7) uses β-PCA. Column (2) and (3) give the results from two sub-samples, where June 1996 is the dividing point between the two (there are fewer observations early in the early subsample). Columns (4), (5), (6) and (9) present the results of regressions without the time and/or security effects. The variables from our theory explain 15% of the variation in the OAS of the securities.

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

Column (8) 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. The two series have a correlation coefficient of 0.88 and their near collinearity causes the standard errors on the coefficients to blow-up, so neither is significant.

The persistence in the two data series may raise concerns that the correlation we find is spurious. Table 5 reports the result of a regression run using first-differenced data:

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

The coefficients estimates are lower than those obtained in the other regressions, but the results remain highly significant. As before, the coefficient estimate for the second half of the sample is lower than that of the first half (specification (2) versus (3)).

A comforting aspect of the results in Table 5 are that monthly changes in OAS spreads correspond more closely to changes in the underlying market prices of the IO's. That is to say, if one imagines that interest rates don't change from one month to the next, but that the price of the IO changes, then it must be the case that the OAS changes. In fact, interest rates do change somewhat, but since the OAS is basically a spread over Treasuries, part of the interest rate change is accounted for. Therefore, our results will be less sensitive to the particular OAS prepayment model when we run regressions using first-differenced data.

Table 5

Regressions based on the OAS of the IO's:

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

β^k is the β-stdev.

Results by subsample are reported in (2) (first-half) and (3) (second-half).

The break point is June 1996.

	(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.

4.4 The quadratic dependence on interest rates

There is one further result that seems unique to our equilibrium theory. We predict that the market price of risk should vary with the *average market coupon.* Plausibly, alternative hypotheses will only link security specific attributes (e.g., the coupon of the specific security being studied) and the market interest rate to the OAS, but not the average market coupon.

Unfortunately, there is little variation in the average market coupon over the 1993 to 1998 sample. However, from Figure 2 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:

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

where the $\bar{c}_t$ and r_t are measured in percentage units. 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 6. 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 6						
Regressions based on the OAS of the collateral: $OAS_C^{kt} = \gamma_k + (\alpha_1 \bar{c}_t + \alpha_2 r_t + \alpha_3) \times (c^k - r_t) + \epsilon_{kt}$ The last column reports the p-value from testing $\alpha_1 + \alpha_2 = 0$.						
Bond Coupon	α_1	α_2	α_3	R^2	N	p-value
7.5	13.44 (1.9)	-9.37 (-2.08)	-54.6 (-1.38)	0.26	26	0.65
8	19.45 (5.3)	-14.17 (-4.48)	-73.45 (-2.52)	0.46	28	0.85
8.5	14.91 (4.51)	-13.51 (-3.96)	-41.97 (-1.51)	0.44	28	0.30
9	7.69 (5.09)	-9.54 (-4.77)	-10.35 (-0.72)	0.45	28	0.59
9.5	9.30 (7.5)	-9.02 (-4.31)	-27.29 (-2.02)	0.51	28	0.09
10	9.66 (6.79)	-8.37 (-3.12)	-34.02 (-2.31)	0.49	28	0.36
10.5	11.76 (6.83)	-7.73 (-2.78)	-59.23 (-4.1)	0.51	28	0.83
11	7.68 (3.61)	-6.99 (-1.83)	-24.13 (-1.19)	0.33	28	0.13
ALL BONDS(1)	10.03 (10.41)	-11.88 (-3.17)	-17.25 (-0.83)	0.87	222	0.38
ALL BONDS(2)	6.53 (3.61)	-4.78 (-8.37)		0.98	222	0.59

ALL BONDS(1) regression uses the entire panel, with security fixed effects.
ALL BONDS(2) regression has both security effects and time fixed effects, and therefore drops α_3.
Estimates reported with T-statistics based on robust standard errors in parentheses.

4.5 Representative household model

Our theory and tests lend support to the view that a specialized mortgage investor sets prices in the MBS market. Thus the delegation of fund management has important effects on asset prices. This view contrasts with traditional asset pricing theory which sees institutions as a "veil." In this section we provide further support for our view by showing that the correlation between prepayment risk and aggregate consumption or wealth has, given the observed values of the OAS, a *sign opposite to that which traditional asset pricing theory predicts.* The reason for this phenomenon seems to be that, controlling for interest rates, households are more likely to prepay mortgages in good states than in bad states.

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. This procedure results in a monthly series of prepayment innovations.

We aggregate the monthly series up to a quarterly level for comparison to aggregate consumption data (data from Q4 1993 to Q1 1998). The consumption data is from the NIPA accounts, and is in real terms. The contemporaneous correlation between the quarterly growth in consumption and the prepayment shocks series is 0.06. The correlation between one-quarter lagged consumption growth and prepayment shocks in 0.35. Table 7 presents these results in the form of standard OLS regressions. The correlations are uniformly positive (but only statistically different from zero for services).

Table 7			
Prepayment shocks and Aggregate Consumption: $U_t = A + B\times$ Consumption-Growth. U_t are prepayment shocks. C_t is measured household consumption. Results are presented for aggregate (non-durables plus services), services, and housing.			
Consumption Series	$\frac{C_t}{C_{t-1}} - 1$	$\frac{C_{t-1}}{C_{t-2}} - 1$	R^2
Aggregate	55.9 (0.31)	279.3 (1.46)	0.13
Services		327.2 (2.06)	0.21
Housing		142.6 (0.87)	0.05
OLS estimates reported with T-statistics in parentheses. $N = 18$			

We note that the IO falls in value with faster prepayment shocks, and rises with slower prepayment shocks. Since prepayment shocks are positively correlated with consumption growth, a representative household pricing model predicts a negative IO risk premium (i.e., it is a hedge against falling consumption). On the other hand, it predicts a positive premium for the PO. However, in the data, the spreads on IO's are positive, while those on PO's are negative. See Figures 3 and 6.

We think there are two reasons for this positive correlation between prepayment and consumption. Households exercise their prepayment option in part based on their ability to pay down their mortgage. Thus when they receive more income, they both increase consumption as well as prepay their mortgage. A second reason is that prepayments are positively correlated with real estate prices (see Table 8, below). As consumption is also positively correlated with real estate prices (see Case, Quigley, and Shiller, 2003), prepayments are thereby correlated with consumption.

We also check whether a CAPM (as opposed to CCAPM) can explain our findings. We form a time series of monthly excess returns on the S&P500 (SP_t) (as a proxy for the bulk of aggregate wealth). 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 portion of prepayment risk that is orthogonal to interest rates is unrelated to the stock market.

The other major part of aggregate wealth is real estate. Mortgage backed securities are zero net supply assets. Thus the relevant correlation for pricing MBS is between prepayment shocks and the value of real estate as a part of the representative household's wealth.

Empirically, it is fairly well established that prepayment rates fall when real estate falls in value, holding interest rates constant (see for example, Caplin, Freeman, and Tracy, 1997, or Downing, Stanton and Wallace, 2003). This phenomenon seems to be due to collateral constraints; the homeowner is unable to refinance a mortgage when the equity value in the home is small.

Table 8 presents the correlation of the U_t series with measures of (real) house price appreciation in different regions of the U.S. The data is from Freddie Mac's index of home prices. In line with most other empirical studies, the correlations are for the most part positive.

Table 8

Correlations between prepayment shocks (U_t) and measures of real estate price appreciation for both regional indices and a national index. corr(0) is the contemporaneous correlation. corr(-1) is the correlation between the one-quarter lagged real estate appreciation and U_t.

	New Eng.	Mid Atl.	South Atl.	E.So Cent.	W.So Cent.	W.No Cent.	E.No Cent.	Mount.	Pac.	United States
corr(0)	0.03	0.18	0.06	0.06	0.21	0.13	-0.09	0.09	0.05	0.08
corr(-1)	-0.08	-0.07	-0.02	0.15	0.18	0.02	-0.12	0.20	-0.18	-0.05
corr(-2)	0.07	0.26	0.24	0.30	0.22	0.16	0.07	0.42	0.11	0.23

Since the IO rises in value when prepayment rates fall, the former acts as a hedge against real estate and should command a negative risk premium if the representative household model is correct. Again, in the data we observe the opposite.

4.6 Preferences of the marginal investor

The coefficient estimates on our model range from 76 to 546, depending on specification and subsample. These numbers are not readily interpretable as corresponding to preferences. In this subsection, we provide a "back-of-the-envelope" calibration to assess these numbers. We show that our findings are in the range of what one would expect if the marginal investor is a leveraged mortgage fund manager.

Previously we found that for a mean-variance investor with risk tolerance of ρ, the OAS is,

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

Let us translate this into preferences for an agent with $CRRA$ preferences with parameter $\hat{\rho}$ and wealth of w,

$$U(w)=\frac{w^{1-\hat{\rho}}-1}{1-\hat{\rho}}.$$

Taking the Taylor expansion around a point w_0 and retaining the first two terms gives us,

$$U(w)-U(w_0)\approx u'(w)\left(\Delta w-\frac{1}{2}\frac{\hat{\rho}}{w}(\Delta w)^{2}\right).$$

So, locally, this agent is a mean-variance investor with risk tolerance of $\hat{\rho}/w$, where w is the fund manager's wealth. Substituting this into the OAS expression gives,

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

We have assumed that the capital requirement for fund managers is α fraction of fund size, so that a fund manager who starts fund by contributing w_T of his own wealth actually has fund capital of $\frac{w_T}{\alpha}$. Now, mortgage funds typically also leverage up this capital via the repo market. Suppose that the typical mortgage fund manager has leverage of L. Then, market clearing (i.e. the fund managers, via leverage, hold the entire mortgage market) requires,

$$L\frac{w_T}{\alpha} = \sum P_C^j \theta^j,$$

where P_C^j is the price of the j-the collateral. We can use this expression to solve for w_T which is the amount of wealth that fund managers have at stake in the mortgage market.

For a typical hedge fund, it is plausible that the largest share of the manager's wealth is tied up in the fund. But more generally, let us suppose that the representative fund manager has a portfolio of κw ($\kappa > 0$) in the mortgage market and $(1-\kappa)w$ in a riskless bank account. Then, $w = \frac{w_T}{\kappa}$, and we can substitute for w into the OAS expression to find,

$$OAS_{IO}^k \approx L\kappa\hat{\rho}\beta^k \frac{\left(\sum \frac{\beta^j}{(r+\phi^j)^2}\theta^j(c^j - r)\right)}{\sum P_C^j \theta^j}.$$

We see that leverage increases the effective risk aversion of the fund manager by L. The reason is that leverage implies that a fund manager with little wealth is taking a large position in the market. In order to compensate the fund manager for bearing this risk, the risk premium must be correspondingly large.

We also see that a lower κ decreases the effective risk aversion of the manager. This is because a fund manager whose wealth is more diversified is less risk averse with respect to shocks in the mortgage market.[11]

We now calibrate this expression based on data from the mortgage market. Over our sample, the average 10-year CMT rate is 6.5%. The average annual prepayment rate (across all of the mortgage pools) is 11.8%. The average β^k is 0.38. If we approximate the OAS formula as,

$$OAS_{IO}^k \approx L\kappa\hat{\rho}\beta^k \frac{\left(\frac{\bar{\beta}}{(\bar{r}+\bar{\phi})^2}\sum \theta^j(c^j - r)\right)}{\sum P_C^j \theta^j},$$

and further take $P_C^j = 1$ (i.e. no discount or premium on the underlying collateral in the market), then,

$$OAS_{IO}^k \approx L\kappa\hat{\rho}\beta^k \frac{\bar{\beta}}{(\bar{r}+\bar{\phi})^2}(\bar{c} - r) = L \times \kappa \times \hat{\rho} \times 11 \times \beta^k(\bar{c} - r).$$

For the leverage number, we have conducted an informal poll of MBS traders, and have found that typical leverage ranges from 5 to 20 for funds that trade IO's and PO's. For the κ number if we use one (i.e. 100% of wealth tied up in the MBS market), then for a coefficient estimate on the model of 438, and a leverage of 10, this implies a risk aversion parameter, $\hat{\rho}$, of about 4. Thus our estimates of $\hat{\rho}$ range from 0.7 to 5 for an L of 10 and κ of one. If κ is one-half, the corresponding risk aversion parameter ranges from 1.4 to 10.

These preference parameters can make sense under a limits of arbitrage view that the marginal investor is a specialized institution. The specialized mortgage fund manager bears disproportionate amounts of mortgage risk. Leverage magnifies this effect making the risk aversion parameters look

[11] α drops out because it has two offsetting effects. On the one hand, low α means that fund managers will run bigger funds and be exposed to more risk. On the other hand, a low α means that the fund manager's exposure to this risk is smaller. The model only requires that α be positive.

reasonable. We should also note that other institutional features of delegation, such as capital constraints, open-ending, minimum benchmarks, etc. may also affect risk aversion. For example, Grossman and Zhou (1996) have shown how institutional demand for portfolio insurance can end up having important effects on aggregate risk aversion and prices. It would be interesting to study further the effects of, for example, capital constraints.

Finally, MBS are zero net supply securities. However if we ignore this fact and suppose that the relevant measure of aggregate wealth of the household included the securitized value of MBS, then the corresponding risk aversion parameters will be too high. As the MBS market is about 10% of aggregate wealth, the risk aversion parameters will be $10L$ times bigger – i.e. 100 times as high, so numbers of 70 to 500. Reconciling this with consumer preferences would be very problematic. We conclude that, to make sense of the risk premia in MBS, one needs a "limited arbitrage" view.

5 Model mis-specification and robustness checks

As we have mentioned, observed OAS behavior might be explained by a mis-specification in the Wall Street mortgage model from which the OAS are derived. We have shown that our results hold using OAS from both Salomon-Smith-Barney as well as Smith-Breeden. Thus our results are not driven solely by peculiarities of one firm's prepayment model. We provide a number of other robustness checks in this section to address the mis-specification possibility.

We should note at the outset 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 reasons we have not done this. Actual bond returns are a very noisy estimate of the expected return on the securities.[12] Thus, 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 our 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 relation is between the OAS for IO's and subsequent returns.

5.1 Is the OAS due to a mis-specified interest rate 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. By construction, therefore, the OAS cannot reflect interest rate risk. However, if the model is miscalibrated or mis-specified, the OAS may reflect an interest rate risk premium.

There are two main reasons why we discount this possibility. First of all, it is possible, but unlikely, for a small miscalibration the model to account for an OAS as high as 500bp – a trader is likely to fix the term structure model if it is this far off. See Figure 6 for the behavior of the OAS of a couple of the securities in our sample. Second, the PO usually has a negative OAS. Recall that the PO is a bond that becomes more valuable when the mortgage prepays. As result, its value increases

[12]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.

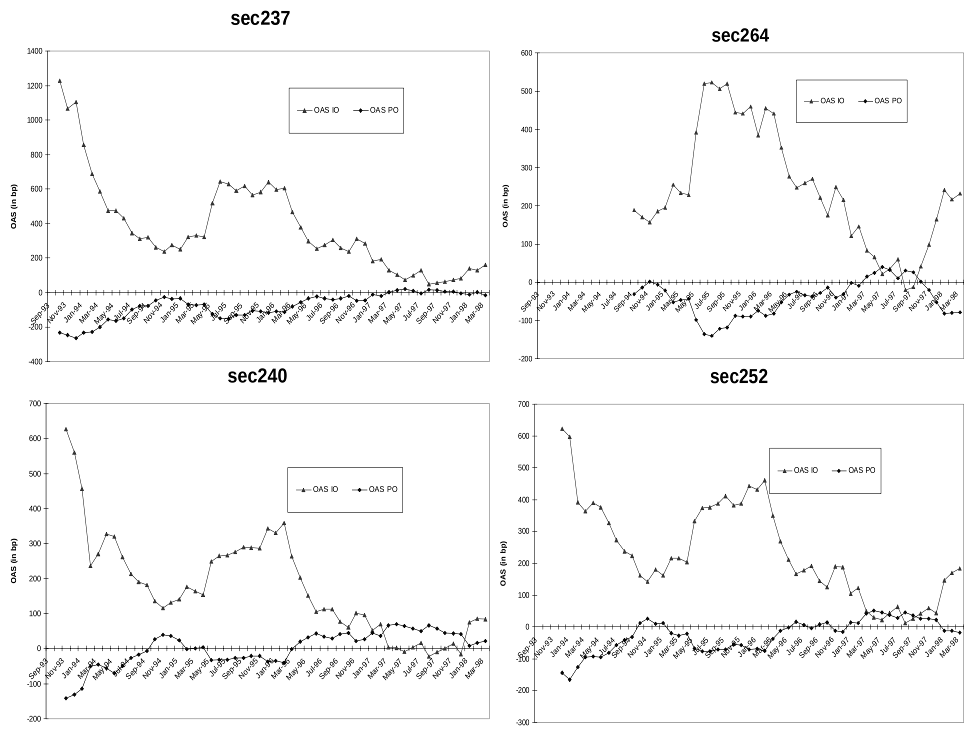

Figure 6: Time series of OAS for IOs and POs
The option adjusted spreads of both IO and PO for four of the mortgage pools are plotted over time

when interest rates fall (both through the effect on the discount rate, as well as the implied faster prepayments). The PO, therefore, always has a positive duration. Since term premia are typically positive, one would expect the OAS to likewise be positive if it mainly reflects interest rate risk. On the other hand, from Figure 6, we can see that the PO has a negative OAS. Moreover, while our graph does not extend into the Fall of 1998, during this period the OAS on the PO's actually became more negative, a highly surprising occurrence if one believes the OAS mainly reflects interest rate risk.

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

The other alternative explanation for observed OAS behavior is that the model of homeowner prepayment is incorrect.

5.2.1 Underprediction hypothesis

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 $\hat{\phi}^k - \bar{\phi}^k$, which is non-zero even if prepayment risk not priced.

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 prepayment models should regularly underpredict prepayment, 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 security/time-fixed effects regression specification we have used should take care of mis-specification. That is, for example, if the underprediction were 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 were a function of both security and time.

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. Such an underprediction implies an OAS whose sign patterns match those found in the data. 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 7 months centered around month t.

The results are reported in the first two columns of Table 9. The s_{kt} variables are not significant (and are negative), while the coefficient on our model remains large and significant.

The coefficient on our model does drop in the second specification, while the R^2 rises. Part of this may be due to different sample sizes. But there are also seems to be an interaction with the fixed effects, as the coefficients (not reported) on some of the bonds change in the second specification.

Table 9			
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, $\bar{s}_{kt}$: $\frac{1}{2}$-year moving average of s_{kt}, centered at t, $c^k \times r_t$: coupon of security-k interacted with r_t.			
	(1)	(2)	(3)
$\beta^k(\bar{c}_t - r_t)$	492.1 (4.56)	199.2 (3.65)	226.9 (4.27)
s_{kt}	-16.9 (-0.37)		
$\bar{s}_{kt}$		-22.48 (-.83)	
$c^k \times r_t$			16.9 (6.81)
R^2	0.93	0.95	0.96
N	374	337	383
All regressions have security and time fixed effects (not reported). Estimates reported with T-statistics based on clustered (by security) standard errors in parentheses.			

5.2.2 Undervaluation of interest rate option hypothesis

We have mentioned earlier that another explanation behind observed OAS behavior 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:

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

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

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

The result is in the last column of Table 9. The interaction term is significant suggesting that there is some mis-specification that distorts the OAS. The coefficient on our model still remain significant.

We investigate this further in Table 10. If traders are undervaluing the interest rate option, then a simple way to control for this effect is to introduce a regressor that is equal to the value of the interest rate option minus intrinsic value.

We take the following approach. For each security we compute,

$$OV_{kt} = E[\max(c^k - \tilde{r}, 0)] - \max(c^k - r_t, 0) \qquad \text{where} \;\; \tilde{r} \sim \mathcal{N}(r_t, \sigma^2\tau).$$

r_t is the 10 year CMT at time t. c^k is the coupon underlying mortgage-k. So the computation is of the value of a European "floor" on the 10 year CMT, minus the intrinsic value of the option. Our distributional assumption is that the 10 year CMT is distributed normally around the current value of the 10 year CMT. Finally the "time to maturity" is τ. We use two different values of τ, 5 years and 10 years. The σ is the sample standard deviation of changes in the 10 year CMT ($81bps$). Finally, we scale this option value by 100 for ease of comparison.

The above is obviously just a crude representation of the value of the option. However, to the extent that we are assuming that *none* of this option value is accounted for in the trader's model, we think we are being fairly conservative.

The results are in Table 10. The coefficient on our model remains significant and of the same order of magnitude as in other specifications. The coefficient on the option also comes out positive and significant, suggesting that there is merit in the idea that prepayment models underestimate the value of the prepayment option. Specifications (2) and (4) interact the option value with β^k, based on the, admittedly atheoretical, idea that perhaps β^k is picking up the number of options embedded in the mortgage.

Table 10

Regressions based on the OAS of the IO's:

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

OV_{kt} is the option value, minus intrinsic, of an European floor on the 10-year CMT, struck at c^k. We also use $\beta^k \times OV_{kt}$ as an independent variable. Results are reported for 5 and 10 year maturities for the option valuation.

	5 year		10 year	
	(1)	(2)	(3)	(4)
$\beta^k(\bar{c}_t - r_t)$	331.9 (4.65)	435.5 (9.96)	279.5 (4.20)	371.2 (8.38)
OV_{kt}	21.7 (3.60)		24.5 (5.27)	
$\beta^k \times OV_{kt}$		41.5 (8.54)		40.2 (8.07)
R^2	0.94	0.96	0.95	0.96
N	383	383	383	383

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

5.2.3 Increasing sophistication hypothesis

Another possible explanation for underprediction of prepayment that has been suggested to us is the "increasing sophistication " hypothesis.[13] Banks typically calibrate their prepayment functions to historical experience. Consumers have in the past been slow to react to interest rate changes in exercising their refinancing option, perhaps because of the fixed costs associated with refinancing. Now suppose that a smart investor knows that looking to the future a decrease in the costs of refinancing will lead consumers to be more reactive in exercising their refinancing options. Then the smart investor will expect a higher rate of prepayments for high coupon mortgages in low interest rate environments. This clearly creates a bias in a model that is calibrated to historical experience.

[13] We thank John Geanokoplos for this suggestion.

There is underprediction for the high coupon mortgages in low interest rate environments. The underprediction should be related to the difference between c^k and r.

Notice that if the bias is simply proportional to $c^k - r$ then the security/time fixed 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 9 control for this possibility.

An alternative way to think about the increasing sophistication effect is that prepayment models calibrated to historical experience assign too few prepayment options to the homeowner. In other words, the OAS is biased because it undervalues the interest rate option of the homeowner. By accounting for the interest rate option explicitly, the results in Table 10 control for this possibility.

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 recently by Stanton (1995) and Longstaff (2004)[14]. 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 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 function which depends on the yield curve, 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[15] 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

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

[15]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.

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 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 fund manager who is largely invested in the mortgage market, as opposed to a well diversified household. The theory predicts that prepayment risk is priced and that the pricing of this risk depends on the value of the entire mortgage market. Our empirical findings support the notion of limits to arbitrage in the MBS market.

The MBS market is a large market in the U.S., and so the evidence we provide is a strong argument in favor of limits of arbitrage theories. The MBS market is also a highly specialized market that requires a great deal of expertise on the part of the investor. We conjecture that the limits of arbitrage effects are most pronounced in markets that require a great deal of expertise. For example, recent evidence from the corporate bond market (Collin-Dufresne, Goldstein, and Martin 2001) and the credit default swap market (Berndt et al. 2004) suggests that market-specific risk factors have important effects on risk premia in these markets. We believe this evidence also supports limits of arbitrage theories. In line with the theory we present in this paper, we predict that when the standard deviation of returns of the corporate bond market increases, the price of non-interest risk in the bond market will increase. Likewise, when credit risk becomes large, the price of credit risk itself will increase. More works needs to be done to assess these conjectures.

The result that a limited amount of capital sets prices in the MBS market also has bearing for models of capital constraints. Liquidations induced by low capital can be expected to have large price effects in such a market as a small set of investors have to absorb the sales.

We have investigated an extension of the model of this paper in which some fund managers face value-at-risk constraints. Our investigations so far confirm that under natural assumptions the MBS risk effect we identify in this paper and the capital constraint effect proposed by others should reinforce each other. Capital scarcity affects MBS prices by first affecting the price of prepayment risk according to the factor structure we identify. This theoretical result suggests an interesting avenue for empirically identifying capital effects: Capital constraints should raise the spreads between high and low prepayment risk securities, in addition to raising the level of all spreads to Treasury securities. We intend to pursue this extension in future work.

A Appendix

A.1 Actual and forecast prepayments

We present the actual prepayment rates for two of the securities (security 237 and 240). Sec 237 is from a high coupon mortgage pool and has a relatively higher rate of prepayments. Sec 240 is from low coupon mortgage pool with a lower rate of prepayment. We also present the fitted estimates from our “simple” statistical prepayment model, as well as the prepayment forecasts from a Wall Street prepayment model.

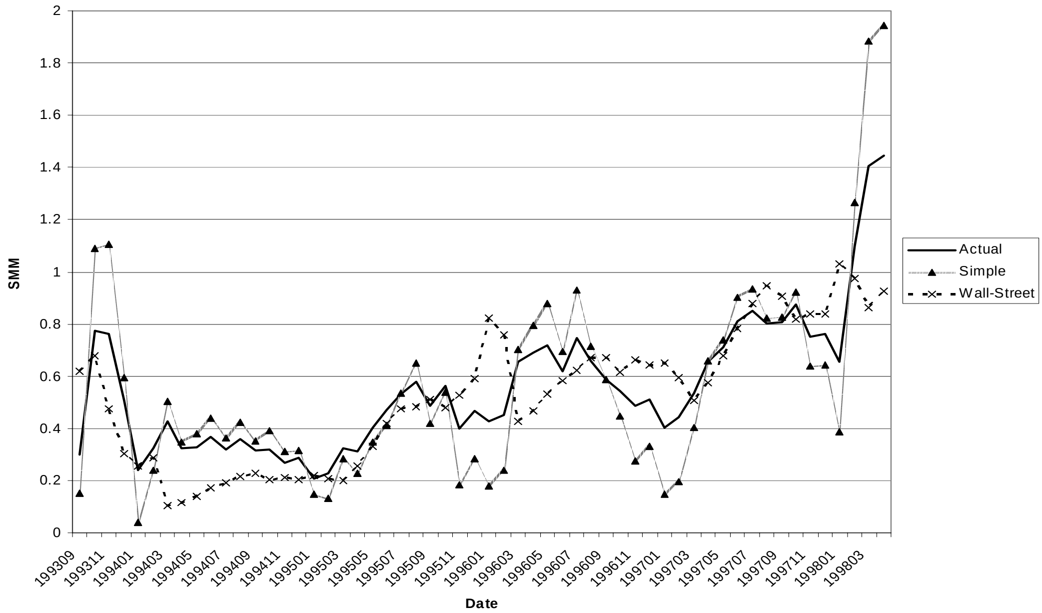

Figure 7: Actual and Forecast Prepayments: Sec 240
The fitted estimates from our “simple” statistical prepayment model, as well as the prepayment forecasts from a Wall Street prepayment model are plotted for security 240.

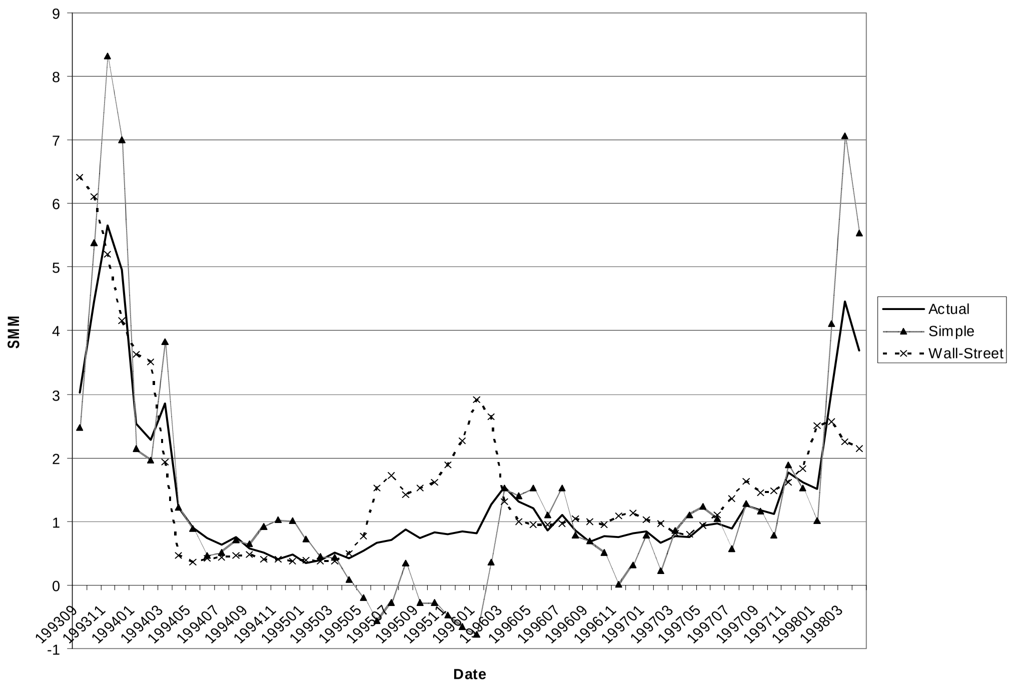

Figure 8: Actual and Forecast Prepayments: Sec 237

The fitted estimates from our “simple” statistical prepayment model, as well as the prepayment forecasts from a Wall Street prepayment model are plotted for security 237.

A.2 Cross-sectional estimates of λ_t

This subsection supplements the results presented in Table 3. We present the results of regressions of the form,

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

for 10/95 and 10/97 in graphical form. The slope coefficient is the estimate of λ_t.

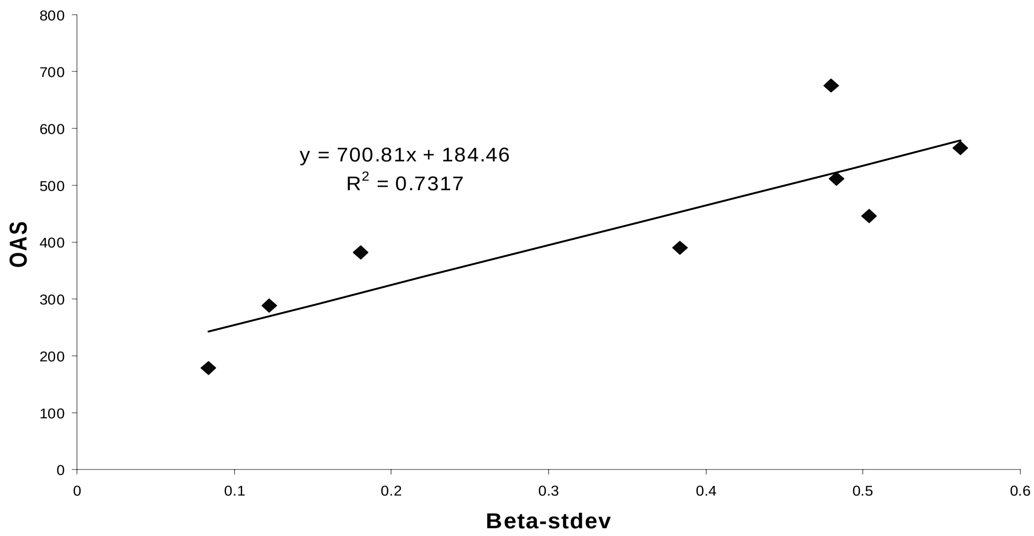

Figure 9: OAS versus β-stdev for 10/95

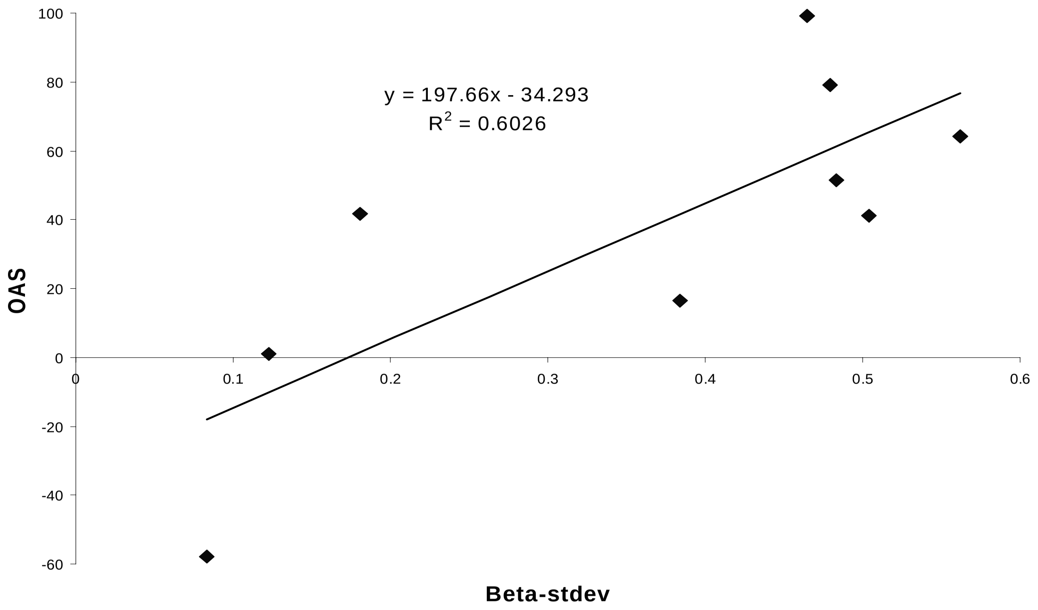

Figure 10: OAS versus β-stdev for 10/97

References

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

[2] Berndt, A., R. Douglas, D. Duffie, M. Ferguson, and D. Schranz, 2004, Measuring Default Risk Premia from Default Swap Rates and EDFs, mimeo, Stanford University.

[3] 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.

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

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

[6] 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.

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

[8] Caplin, A., C. Freeman and J. Tracy, 1997, "Collateral Damage: How Refinancing Constraints Exacerbate Regional Recessions," *Journal of Money, Credit and Banking*, 29, 496-516.

[9] Case, K. , J. Quigley, and R. Shiller, 2003, "Comparing Wealth Effects: The Stock Market Versus the Housing Market," Berkeley Mimeo.

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

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

[12] Downing, C. , R. Stanton and N. Wallace, 2003, "An Empirical Test of a Two-Factor Mortgage Prepayment and Valuation Model: How Much Do House Prices Matter?", Berkeley mimeo.

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

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

[15] 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.

[16] Gabaix, X., P. Gopikrishnan, V. Plerou, H.E. Stanley, 2003, "A Theory of Power Law Distributions in Financial Market Fluctuations", *Nature*, 423, 267-70.

[17] Geanakoplos, J., 2003, "Liquidity, Default, and Crashes." In *Advances in Economics and Econometrics II,* edited by M. Dewabtripont, L.P. Hansen, and S.J. Turnovsky.

[18] Geanakoplos, J., 1997, "Promises Promises", In *The Economy As an Evolving Complex System II*, edited by B. Arthur, S. Durlauf, and D. Lane.

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

[20] Grossman, S. and Z. Zhou, 1996,"Equilibrium Analysis of Portfolio Insurance," *Journal of Finance*, 51(4), 1379-1403.

[21] Holmstrom, B. and J. Tirole, 1997. "Financial Intermediation, Loanable Funds, and the Real Sector," *Quarterly Journal of Economics*, 112, 663-691.

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

[23] Kupiec, P., and A. Kah, 1999, "On the Origin and Interpretation of OAS," *Journal of Fixed Income* 9(3), 82-92.

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

[25] Longstaff, F., (2004), "Optimal Recursive Refinancing and the Valuation of Mortgage-Backed Securities", NBER Working Paper # 10422.

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

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

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

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

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

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