

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

Xavier Gabaix Arvind Krishnamurthy Olivier Vigneron*

January 1, 2005

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 wealth or consumption implies the wrong sign to match the observed prices of prepayment risk. The price of risk is better explained by a kernel based on MBS-market-wide specific risk. The results support the existence of 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 Kobi Boudoukh, Donald Bronstein, John Campbell, Kent Daniel, John Geanakoplos, Massoud Heidari, Jingzhi Huang, Ravi Jagannathan, Adam Kolasinski, Ananth Krishnamurthy, Yong Liu, Francis Longstaff, Bob McDonald, Andrew Metrick, Anna Pavlova, Matt Richardson, Tim Riddiough, Anthony Sanders, José Scheinkman, Richard Stanton, René Stulz, Raman Uppal, Dimitri Vayanos, Nancy Wallace, 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 UC-Berkeley, University of British Columbia, University of Chicago, HEC-Montreal, LBS, LSE-FMG, MIT, University of Maryland, NYU, and Northwestern for their comments. We also thank Salomon-Smith-Barney, Smith-Breedon and UBS 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 existence of limits of arbitrage in the context of the mortgage-backed securities (MBS) market. We present a simple model of delegated fund management in which a risk-averse fund manager who has all of his wealth tied up in the MBS market is the marginal investor in the MBS market. The novelty of the theory lies in its implications for the pricing of MBS-market specific risk. We predict that MBS-market specific risk will carry a positive risk premium, even if such a risk is not predicted to carry a risk premium in a traditional asset pricing theory. Moreover, we predict that market price of this risk will be high when the total riskiness of the MBS-market is also high. We argue that prepayment risk in the MBS market is a good example of an MBS-market specific risk, and present evidence consistent with our theoretical predictions.

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 empirical results in the paper. First, we show that prepayment risk carries a positive risk premium. 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 riskiness of the MBS market and show that the market price of prepayment risk comoves with this 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 show that when the fund manager in our model of delegated fund management has a coefficient of relative risk aversion of 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

¹Berndt, 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.

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

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

The next section of the paper provides a brief overview of the mortgage market. We also review the main results of the paper in this section. Section 3 formally presents the theoretical model that motivates our empirical tests. Section 4 presents evidence in support of that theory. Section 5 provides robustness checks of our results and discusses alternative hypotheses. The conclusion and an appendix follows.

2 Overview of the mortgage market and the main results

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, real estate prices, 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 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 term-structure such that the present value of a security's expected cash flows, forecast using the prepayment model and discounted under the term structure model, 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 present a theory where the behavior of the risk premium for bearing MBS prepayment risk depends on whether the marginal investor is a specialized MBS fund manager or a representative household. 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 mis-specification, therefore, would imply that the OAS is a noisy proxy of a security's risk premium, potentially rendering the interpretation of our results suspect.

Empirically, we first show that the securities with higher levels of prepayment risk (" β^k ") also have higher OAS. More precisely, we study a cross-section of securities, each with a different underlying mortgage pool. We examine the prepayment experience of each underlying pool and develop a measure of prepayment risk: A pool where a substantial amount of prepayment variability is unrelated to changes in interest rates is classified as having high prepayment risk (β^k). In the cross-section, a regression of the OAS on β^k has a positive and significant slope coefficient and a high R^2 .

We interpret this result as showing that prepayment risk is priced and earns a positive risk premium. But, as mentioned above, it may be that our measure of prepayment risk is correlated with a mis-specification in the prepayment model from which the OAS is derived and the regression correlation we report is spurious.

We should note at the outset that one way to get around this mis-specification issue is to use actual returns on the MBS as the dependent variable in our regressions. We have chosen against this approach because of data limitations. Actual bond returns are a very noisy estimate of the expected return on the securities. Thus, we need more data than we have to implement these regressions. Using the OAS greatly reduces this measurement error problem, at the cost of raising the mis-specification possibility. 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.

In our second set of empirical results, we show that the slope coefficient of the cross-sectional regression of the OAS on β^k varies over time systematically with a measure of the risk-bearing capacity of an investor who is wholly invested in the MBS market. We interpret this result as showing that the market price of prepayment risk moves with investor risk-aversion. The result is strong evidence in favor of our hypothesis that prepayment risk earns a positive risk premium, as opposed to the mis-specification hypothesis.

Our measure of risk-bearing capacity of investors is based on the observation that the dollar-prepayment-risk on a MBS is increasing in the difference between the coupon on that mortgage and prevailing market interest rates. For example, in a low interest rate environment, the unexpected

prepayment of a high coupon MBS results in a large dollar loss to the holder of this MBS. Prepayment risk, measured in terms of prepayment variability on the underlying mortgage pool (β^k as described above), translates into dollar-prepayment-risk based on the difference between the coupon on the MBS and market interest rates. Analogously, the dollar-prepayment-risk across the entire mortgage market is proportional to the difference between the *average coupon* outstanding on all MBS in the market and prevailing market interest rates (“ $\bar{c} - r$ ”).

We show that when $\bar{c} - r$ is high, the slope coefficient of the cross-sectional regression of the OAS on β^k is also high. It is not clear why any mis-specification in the model used to derive the OAS should be related to the average market-wide coupon on mortgages (of course, it is reasonable that mis-specification may be related to the coupon of a specific MBS). On the other hand, we argue that when $\bar{c} - r$ is high, the marginal investor bears a lot of market-wide prepayment risk, so that the price this investor will charge for bearing any given security’s prepayment risk must be proportionately higher.

The result has bearing for limits of arbitrage theories because it suggests that the marginal investor is wholly invested in the MBS market. In other words, this investor resembles the specialized risk-arbitrageur who is the central player in limits of arbitrage theories.

In our final set of results, we show that if the marginal investor is the representative consumer hypothesized by the traditional CAPM or consumption-CAPM model, the market price of prepayment risk (1) will be unrelated to $\bar{c} - r$; and, (2) should be negative. We make a theoretical argument to prove the first result. MBS are zero net supply assets. A homeowner exercising her prepayment option does not change aggregate wealth as it merely transfers wealth between the holder of the MBS and the homeowner. Thus, in measuring aggregate wealth one should not include the aggregate value of MBS. However, the variable $\bar{c} - r$ only arises if one considers a marginal investor for whom the holdings of MBS constitutes net wealth. Thus, it must be that the marginal investor is not the representative consumer hypothesized by traditional asset pricing theory. To prove the second point, we empirically demonstrate that homeowner prepayments on mortgages, controlling for interest rates, are positively correlated with real estate prices and aggregate consumption. When real estate prices rise, refinancing activity increases and housing turnover rises. Both lead to increased rates of prepayment. Similarly we show that higher rates of prepayment are correlated with positive innovations to aggregate consumption. Since both of these correlations imply that prepayments are highest in “good” states of the world of a representative consumer, prepayment risk should carry a negative risk premium. But, as we show, prepayment risk carries a positive risk premium.

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, where the marginal investor is one who is wholly invested in the MBS market, and present our main hypotheses regarding the pricing of MBS. We formally derive the relation we alluded to in the previous section:

$$OAS_{IO}^k \propto \beta^k (\bar{c} - r).$$

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,

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

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

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

Finally the value of the mortgage pass-through itself – the collateral – is,

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

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,

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

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

³Unlike our abstraction, in practice interest rates are uncertain. The logical extension of our model to the uncertain interest rate case is to write $\bar{\phi}^k(\bar{r})$. 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.

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

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

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

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

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,

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

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,

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

(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,

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

⁴It 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, $P_{IO} = E [c / (r + \tilde{\phi})] > c / (r + E [\tilde{\phi}])$ 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.

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

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_M . 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,

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

i.e. just a mean-variance maximizer.

3.5 Equilibrium

At date 0, fund managers raise a total of $w_F - w_M$ from investors. This gives them total capital of $\frac{w_M}{\alpha}$. With this sum the fund managers purchase a portfolio of mortgage backed securities. Let x_{IO}^k

⁵The 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,

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

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.

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

and x_{PO}^k be the amount of the k -th IO and PO held in a portfolio. Then,

$$W_F = \frac{w_F}{\alpha} + \sum_k x_{IO}^k (V_{IO}^k - P_{IO}^k) + \sum_k x_{PO}^k (V_{PO}^k - P_{PO}^k) \quad (12)$$

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,

$$\frac{w_F}{\alpha} \geq \sum_k x_{IO}^k P_{IO}^k + \sum_k x_{PO}^k P_{PO}^k.$$

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 $x_{IO}^k = x_{PO}^k = \theta^k$, yields an expression for the price of the IO,

$$\frac{c^k}{r + \bar{\phi}^k} - P_{IO}^k = -\rho\alpha \text{cov} \left(\frac{c^k}{r + \bar{\phi}^k} \eta^k \Delta\phi^k, R_M \right) \quad (13)$$

where the market is defined as:

$$R_M = \sum_j \frac{\theta^j}{(r + \bar{\phi}^j)^2} \Delta\phi^j (r - c^j) \quad (14)$$

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,

$$\Delta\phi^k = \beta^k \Phi + \epsilon^k \quad (15)$$

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

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

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

⁷The exact expression is,

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

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

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. As a proxy for the sum term, we use a weighted average coupon across all the agency issued mortgage backed securities in the market (Fannie Mae, Freddie Mac, and Ginnie Mae), where the weights are the amounts outstanding of each mortgage security. The relation we use in our tests is,⁸

$$OAS_{IO}^k = \underbrace{\beta^k}_{\text{Systematic risk}} \times \underbrace{\rho\alpha(\bar{c} - r)}_{\text{Market price of risk}} \quad (16)$$

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 .

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,

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

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

$$OAS_C^k = \beta^k \times \rho\alpha(\bar{c} - r)(c^k - r) \times \frac{1}{c^k + \bar{\phi}^k + OAS_{IO}^k} \quad (17)$$

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.

⁸We 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>.

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

3.7 Testable empirical predictions

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

$$OAS_{IO}^{kt} = \beta^k \lambda_t \quad (18)$$

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

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 derive 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
Size ^d	1,375	3,450	1,975	1,055	1,050	1,725	898	1,155	567
a: Securities are identified by their pool number. b: Weighted average coupon (WAC) on underlying mortgage pool ($\pm 5bp$ over sample.) c: Seasoning of underlying pool in months as of July 98. d: Size of underlying pool at original issue date, in millions.									

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 also have quarterly data on the OAS for six generic (TBA) 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-Breedon. We do not have prepayment information for these pools. We also test our model using this data. We discuss these tests in further detail in Section 4.5.

We construct time series of monthly OAS data for the IO's 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.

4.1 Estimation of β^k and $\bar{c} - r$

There are two steps in testing (18)–(19). We need an estimate of β^k , and we need an estimate of $\bar{c} - r$.

Our estimate of \bar{c} comes from UBS. For the last several years, UBS has tracked the weighted average coupon across all outstanding mortgages pools of Fannie Mae, Freddie Mac, and Ginnie Mae. This variable is computed as the average of the underlying coupons on individual mortgage pools, weighted by the amount outstanding of the pools. UBS kindly provided us with this data going back to 1988 to form our estimate of \bar{c}_t . The series is computed monthly. We use the 10-year constant maturity Treasury yield from the Federal Reserve's BOG website to form r_t .

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 we estimate is,

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

where, age_t is the age of the mortgage. The term $\frac{c^k}{r_{t-1}}$ captures some of the non-linearity of the homeowner's prepayment option. Richard and Roll (1989) use a similar term in their prepayment model. For most of the mortgage pools we study, c^k is above r_t throughout our sample, so that the

option is near-the-money or in-the-money. 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. The term involving age_t captures seasoning effects in mortgage pools. 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 10-year 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. A standard-deviation series is plotted for each of the nine mortgage pools that we study.

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	0.120	0.181	0.449	0.492	0.549	0.471	0.460	0.383
β^k (PCA)	0.062	0.121	0.198	0.417	0.594	0.503	0.465	0.518	0.334
<i>idiosync</i> ^k (PCA)	0.088	0.074	0.080	0.167	0.263	0.178	0.240	0.178	0.203
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 84% of the variance which suggests that equation (15) is a good representation of the data. The second and third components account for 8.5% and 3% 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 have used a simple prepayment model that captures some of the main features of prepayment behavior – dependence on interest rates relative to mortgage coupon, dependence on lagged interest rates, and seasoning of a mortgage pool. Our aim is to come up with a ranking, across pools, of prepayment variability rather than present a fine-tuned prepayment model. 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. We have experimented with other prepayment models (for example, adding more lags of interest rates, or adding interest rate of other maturities) and while the prepayment estimates change, our rankings of prepayment variability remain relatively unaffected. As a check, we also have prepayment forecasts from a 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 4.3 and 4.4. We discuss alternative hypotheses and robustness in section 5.




Figure 2: CMT 10-year and Average Coupon

The yield on the 10-year constant maturity Treasury note (r) and the weighted average coupon on all outstanding mortgages in the MBS market ($cbar$) are plotted over time.

4.2 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 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 from 1991 through 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

⁹We 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.

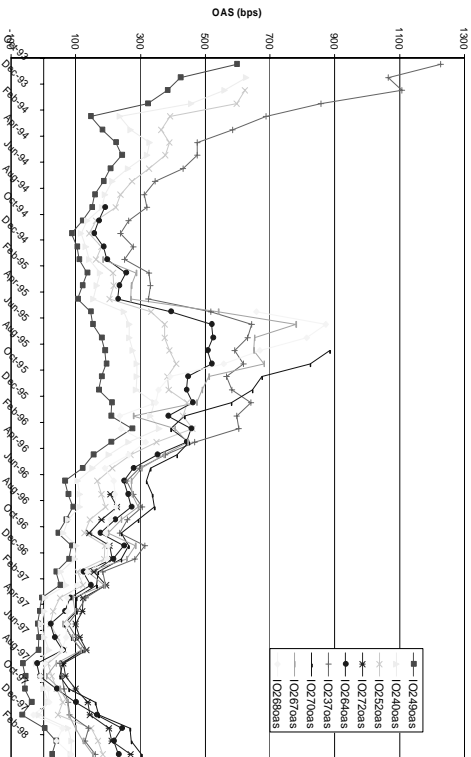


Figure 3: OAS of the IOs

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.

(as predicted by equation (10)). One should also note that the interest rate alone is not enough to understand the relative magnitude of the OAS of the IO between 1993 and 1998 when rates were at the same level: OAS are much higher in 1993 than in 1998. This in fact is not a puzzle in light of our derivations, since equations (18)–(19) tells us that the OAS of the IO is proportional to $\bar{c} - r$ and not r alone. Indeed, when looking at $\bar{c} - r$, we find that it is twice as high in 1993 than in 1998.

4.3 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,

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

The OAS is measured in basis points. β^k is given in Table 2 and is measured in units of percentage-prepayment rates per month. 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 average value of λ_t in the table is 469. As β^k varies from 0.08 to 0.55 across the securities, this coefficient implies a difference between these securities of 220 basis points in the OAS.

The estimation 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

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

Table 3: Regressions of, OAS = constant + λ β -stdev, for each month. λ is the market price of risk. β -stdev is the prepayment risk measure

Month	λ estimate	Std error	R-sq	N (# obs)	Month	λ estimate	Std error	R-sq	N (# obs)
Nov-93	1230	216	0.962	3	Jan-96	342	226	0.246	8
Dec-93	1423	146	0.969	4	Feb-96	317	201	0.263	8
Jan-94	1010	197	0.897	4	Mar-96	486	79	0.843	8
Feb-94	1072	160	0.937	4	Apr-96	456	75	0.841	8
Mar-94	783	118	0.937	4	May-96	371	72	0.792	8
Apr-94	430	143	0.751	4	Jun-96	436	65	0.864	8
May-94	420	105	0.842	4	Jul-96	425	71	0.816	9
Jun-94	425	86	0.890	4	Aug-96	445	67	0.847	9
Jul-94	313	66	0.882	4	Sep-96	407	60	0.850	9
Aug-94	294	56	0.902	4	Oct-96	359	60	0.816	9
Sep-94	208	106	0.492	5	Nov-96	404	59	0.853	9
Oct-94	215	88	0.716	5	Dec-96	357	53	0.850	9
Nov-94	228	61	0.776	5	Jan-97	256	39	0.847	9
Dec-94	264	76	0.751	5	Feb-97	267	32	0.773	9
Jan-95	213	43	0.829	6	Mar-97	254	32	0.887	9
Feb-95	320	51	0.889	6	Apr-97	250	34	0.868	9
Mar-95	324	70	0.812	6	May-97	183	48	0.646	9
Apr-95	344	69	0.831	6	Jun-97	208	49	0.697	9
May-95	730	237	0.614	7	Jul-97	253	55	0.727	9
Jun-95	1110	343	0.635	7	Aug-97	177	59	0.529	9
Jul-95	974	277	0.674	7	Sep-97	140	63	0.385	9
Aug-95	1008	266	0.672	8	Oct-97	204	59	0.600	9
Sep-95	971	209	0.756	8	Nov-97	258	83	0.546	9
Oct-95	725	165	0.739	8	Dec-97	372	114	0.569	9
Nov-95	716	165	0.729	8	Jan-98	316	136	0.405	9
Dec-95	596	167	0.645	8	Feb-98	244	137	0.285	9
					Mar-98	337	139	0.423	9

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 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.

The R^2 are also low during Aug and Sep 97. This latter period was the Asian Crisis, when market observers noted that the liquidity in the MBS market decreased, possibly leading to stale price problems.




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.

4.4 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:

$$OAS_{it}^{st} = \alpha_t + \gamma_k + A \times \beta^k (\bar{c}_t - r_t) + \epsilon_{kit},$$

where the \bar{c}_t and r_t are measured in percentage units, and the OAS is measured in basis points. 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 greater depth in the next section.

Both the OAS series and the $(\bar{c}_t - r_t)$ series are persistent, so there is serial correlation in the standard errors (ϵ_{kit}). We correct for this in two ways. First, most of the regressions report T-statistics which are corrected for serial correlation in the standard errors at the security level (we cluster 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 in the standard errors 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




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.

have also done a robustness check using a standard panel data adjustment where we assume that ϵ_{kt} is 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)
$\beta^k (\bar{c}_t - r_t)$	462.1 (4.89)	547.4 (6.72)	324.1 (5.59)	373.9 (5.99)	441.9 (7.24)	275.7 (4.52)	441.2 (3.19)	282 (1.01)	254.2 (3.64)
$idiosync^k (\bar{c}_t - r_t)$	Yes	Yes	Yes	No	Yes	No	Yes	Yes	No
Security Effects	Yes	Yes	Yes	Yes	No	No	Yes	Yes	No
Time Effects	0.94	0.95	0.92	0.93	0.68	0.20	0.93	0.93	0.18
R^2	383	194	189	383	383	383	383	383	383
N									

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). The coefficient estimate using β -stdev is 462, while it is 441 using β -PCA. In the cross-section, our measure of β^k -stdev varies from 0.08 to 0.55. In the time series, $\bar{c} - r$ varies from a low of -0.1 to a high of 3.06 with an average value of 1.30 . If we consider a security with a β^k of 0.50 , then the coefficient estimate of 462 implies a time-series variation in the OAS from -22 bps to a high of 706 bps as $\bar{c} - r$ varies from -0.1 to 3.06 . If we consider the average level of $\bar{c} - r$ of 1.30 , then the coefficient estimate of 462 implies a cross-sectional variation in the OAS of 285 bps, as β^k -stdev varies from 0.08 to 0.55 .

The time and security effects inflate the R^2 's in the regressions, leading to somewhat misleading R^2 's in the baseline regressions. Dropping both the time and security effects, the variables from our theory explain 20% of the variation in the OAS of the securities (column 6). If we drop the variable from our theory but keep both time and security effects, the R^2 remains high at 89% (not reported). Columns (4), (5), and (9) present other combinations from the results of regressions without the time and/or security effects.

We note the lower (but still highly 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:

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

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:

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

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)
$\Delta \beta^k (\bar{c}_t - r_t)$	173.9 (3.05)	216.9 (2.67)	82.5 (2.40)
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.			

4.5 Average market coupon

There is one further result that is 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. From Figure 2, we note the adjustments in the average coupon outstanding over the period from 1991 to 1998. We exploit this variation in the average coupon to verify that the \bar{c} has independent explanatory power for the OAS. We begin this section by reporting results from the IO sample, and then move on to our main results which come from the passthrough data from Smith-Breeden.

Over the period from September 1993 to April 1998 the average coupon outstanding falls from 8.45% to 7.52% (see Figure 2). This fall occurs in two phases: early in the sample, and again late in the sample. We estimate the following regression for the IO's:

$$OAS_{IO}^{kt} = \alpha_t + \gamma_k + A_1 \times \beta^k \bar{c}_t + A_2 \times \beta^k r_t + \epsilon_{kt}$$

The coefficient estimates are: $A_1 = 378.6(9.78)$, and $A_2 = -193.7(3.66)$ ($R^2 = 96\%$ for $N = 383$). We note that the coefficients are of opposite sign, which is in keeping with the $\bar{c}_t - r_t$ form suggested by our theory. On the other hand, in contrast to the theory, $A_1 + A_2$ is not equal to zero.

The largest variation in the average coupon occurs over the period from mid-1991 to early 1994. We have OAS data for passthroughs over this period from Smith-Breeden, which we also use to test for the explanatory power of the average coupon. 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. From our theory (see equation (17)), the passthrough takes a quadratic form in interest rates. We estimate the following regression:

$$OAS_C^{kt} = \gamma_k + (A_1 \bar{c}_t + A_2 r_t + A_3) \times (c^k - r_t) + \epsilon_{kt},$$

where the \bar{c}_t and r_t are measured in percentage units. Our theory predicts that A_1 is positive and that A_2 is negative, and that $A_1 + A_2 = 0$. A_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.¹¹ It is also encouraging that A_1 and A_2 have opposite signs. While the coefficients, A_1 and A_2 , are of similar magnitudes, it is disappointing that the test of $A_1 + A_2 = 0$ is rejected for most of the specifications. Finally, in contrast to our theory, A_3 is negative and often significant.

A possible explanation for this 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 negative A_3 captures the market's expectation that in the future, the average coupon is going to decrease. Alternatively, this discrepancy could be due to a mis-specification of the interest rate in our simple empirical implementation.

¹¹We can do a rough calculation to see whether the coefficient estimates from the collateral regressions are consistent with the coefficient estimates from the IO regressions that we reported earlier. For the IO's we predict that,

$$OAS_{IO}^k = A\beta^k(\bar{c} - r)$$

and estimated that A was around 460. For the collateral, we predict that,

$$OAS_C^k = A\beta^k(\bar{c} - r)(c^k - r) \frac{1}{c^k + \bar{\phi}^k + OAS_{IO}^k}$$

In our collateral regressions, we estimate \hat{A} where,

$$OAS_C^k = \hat{A}(\bar{c} - r)(c^k - r)$$

We can relate \hat{A} to A using these expressions. Also, keeping track of the unit conversions from bps to %, we find that,

$$\hat{A} = A\beta^k \frac{1}{c^k + \bar{\phi}^k + OAS_{IO}^k} \frac{1}{100}$$

Substituting in typical numbers from our sample of $c^k = 0.09$, $\bar{\phi}^k = 0.11$, $OAS_{IO}^k = 0.04$, and $\beta^k = 0.4$, we find that, $\hat{A} = A \times 0.016$. Thus for an estimate of A of 460, we should expect that \hat{A} is around 7.5. The values are in the neighbourhood of the coefficient estimates for the collateral.

Table 6						
Regressions based on the OAS of the collateral: $OAS_C^{kt} = \gamma_k + (A_1 \bar{c}_t + A_2 r_t + A_3) \times (c^k - r_t) + \epsilon_{kt}$ The last column reports the p-value from testing $\alpha_1 + \alpha_2 = 0$.						
Bond Coupon	A_1	A_2	A_3	R^2	N	p-value
7.5	19.4 (2.11)	-10.1 (-2.18)	-104.1 (-1.85)	0.27	26	0.13
8	24.4 (5.26)	-14.0 (-4.41)	-120.0 (-3.22)	0.45	28	0.03
8.5	18.6 (4.32)	-13.1 (-3.83)	-78.6 (-2.24)	0.45	28	0.21
9	9.6 (4.46)	-9.2 (-4.67)	-29.9 (-1.54)	0.47	28	0.88
9.5	11.0 (6.65)	-8.2 (-3.81)	-47.2 (-2.84)	0.51	28	0.30
10	11.1 (5.85)	-7.3 (-2.65)	-53.0 (-2.85)	0.48	28	0.23
10.5	14.0 (6.8)	-6.3 (-2.37)	-85.9 (-5.29)	0.53	28	0.01
11	9.4 (3.57)	-6.0 (-1.61)	-42.9 (-1.82)	0.35	28	0.44
ALL BONDS(1)	11.6 (10.45)	-11.7 (-3.18)	-33.3 (-1.36)	0.86	222	0.99
ALL BONDS(2)	6.87 (3.41)	-4.73 (-7.55)		0.99	222	0.35
ALL BONDS(1) regression uses the entire panel, with security fixed effects. ALL BONDS(2) regression has both security effects and time fixed effects. Estimates reported with T -statistics based on robust standard errors in parentheses.						

4.6 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 wealth or aggregate consumption 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 first measure the correlation between prepayment innovations and measures of a representative household's wealth. We note that mortgage backed securities are zero net supply assets. A homeowner exercising her prepayment option does not change aggregate wealth as it merely transfers wealth between the holder of the MBS and the homeowner. Thus, in measuring aggregate wealth one should not include the aggregate value of MBS. We check the correlations between prepayment innovations and the stock market and real estate prices.

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 0.032. We run a regression of,

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

The coefficient estimate for B is 0.873 and the t -statistic is 0.23 ($N = 53, R^2 = 0.001$). 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. Empirically, there is extensive evidence that prepayment rates fall when real estate falls in value, holding interest rates constant (see for example, Caplin, Freeman, and Tracy, 1997, Bennett, Peach, and Peristiani, 2000, Matthey and Wallace, 2001, or Downing, Stanton and Wallace, 2003). Caplin, Freeman, and Tracy argue that this phenomenon is due to collateral constraints; the homeowner is unable to refinance a mortgage when the equity value in the home is small. Hurst and Stafford (2003) present evidence that households refinance their mortgages and borrow against the equity value in their homes to boost consumption when real estate prices rise. Stein (1995) argues that down-payment effects lead to more housing turnover when real estate prices rise. This activity naturally creates higher rates of prepayment.

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. Our evidence is not as strong as other evidence presented in the literature (cited above). However, in line with 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)	0.01	0.15	0.03	0.06	0.19	0.18	-0.03	0.09	0.04	0.08
corr(-1)	-0.12	-0.11	-0.07	0.11	0.12	-0.03	-0.17	0.15	-0.21	-0.11
corr(-2)	0.04	0.21	0.21	0.26	0.18	0.12	0.03	0.40	0.08	0.18

We note that the IO falls in value with faster prepayment shocks, and rises with slower prepayment shocks. 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 (i.e., it is a hedge against falling real estate prices). On the other hand, the traditional theory 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.

Table 7			
Prepayment shocks and Aggregate Consumption: $U_t = A + B \times \text{Consumption-Growth.}$ 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	32.5 (0.19)	278.5 (1.49)	0.13
Services		320.4 (2.06)	0.21
Housing		150.7 (0.94)	0.05
OLS estimates reported with T -statistics in parentheses. $N = 18$			

We also check the implications of the consumption-based CAPM for the pricing of MBS. We aggregate the monthly series of prepayment innovations 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 0.03. The correlation between one-quarter lagged consumption growth and prepayment shocks is 0.36. 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).

The positive correlation between consumption and prepayment innovations means that (as with the case of real estate prices), under the traditional theory, the IO should command a negative risk premium while the PO should command a positive risk premium. Neither is true in practice.

We should note that the positive correlation between consumption and prepayment innovations is from a small sample of a boom period during the 1990's. Unlike the link we demonstrate between real estate prices and prepayment innovations, we are not aware of prior work that documents this correlation. However the result seems believable on a priori grounds. As noted above, prepayments are positively correlated with real estate prices. As consumption is also positively correlated with real estate prices (see Case, Quigley, and Shiller, 2003), prepayments are thereby correlated with consumption. A second possible explanation is that homeowners are only able to take advantage of lower interest rates and refinance if they qualify for a new mortgage (see Longstaff, 2004). Thus, when income and consumption grows, the credit quality of the average homeowner improves, thereby allowing more homeowners to take advantage of refinancing options and increasing prepayment rates.

4.7 Preferences of the marginal investor

The coefficient estimates on our model range from 82 to 547, depending on specification and sub-sample. 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,

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

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,

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

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,

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

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.¹²

¹² 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.

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 462, and a leverage of 10, this implies a risk aversion parameter, $\hat{\rho}$, of about 5. 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 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.¹³

¹³We can reduce the noise in bond returns if we take a stand on the mortgage prepayment model and calculate

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.

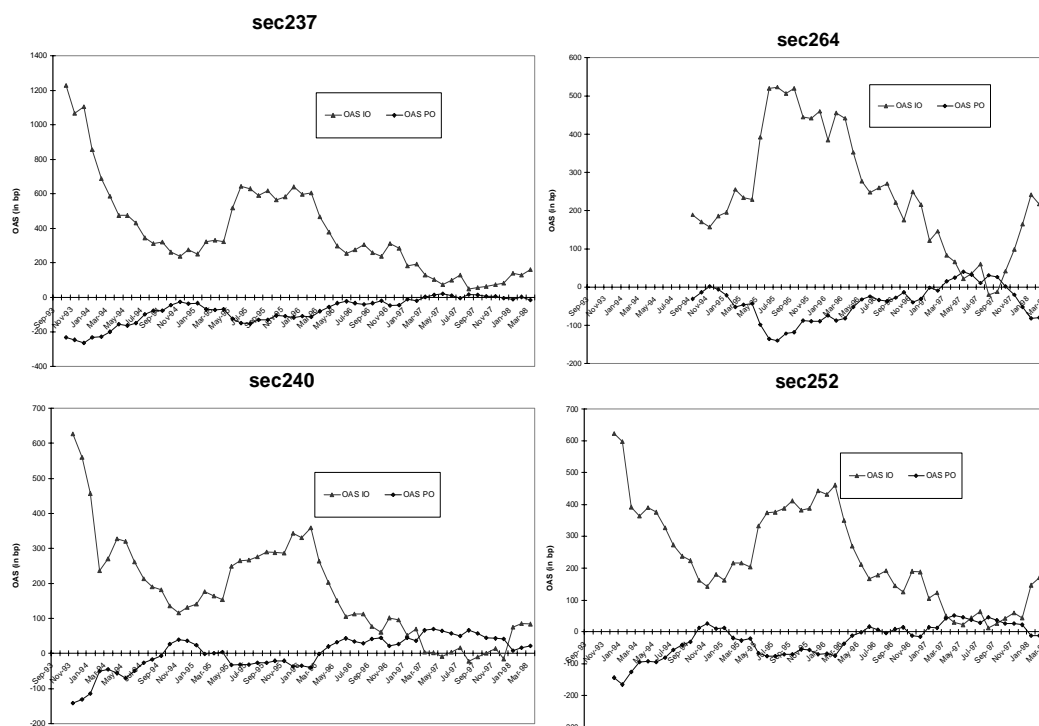


Figure 6: Time series of OAS for IOs and POs

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

There are two main reasons why we discount this possibility. First, the IO almost always has a negative duration, while the PO has a positive duration. Since term premia are typically positive, 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.

one would expect the OAS on the PO's (IO's) to be positive (negative) if the OAS mainly reflects interest rate risk. On the other hand, from Figure 6, we can see that the PO typically has a negative OAS while the IO has a positive OAS. Moreover, while the data we use in our analysis does not extend into the Fall of 1998, market participants note that during this period the OAS on the IO's rose while the OAS on the PO's fell. This observation is surprising if one considers that term premia probably rose during this period.

Second, it is possible, that a small miscalibration in the interest rate model may lead to an OAS of 50bp, but seems highly unlikely that the miscalibration will produce an OAS of 500bp – a trader is likely to fix the term structure model if it is this far off. The OAS on our IO's are quite high and seem unlikely to be due to a miscalibrated interest rate model.

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.

There is a plausible hypothesis why underprediction may have been the norm over our sample. Banks typically calibrate their prepayment functions to historical experience. Consumers have in the past been slow to exercise their refinancing option. However, over the 1990's, increased competition and the explosion of internet-based lending greatly reduced the costs of refinancing. As a result, consumers were quicker to take advantage of refinancing possibilities and prepayments increased.

Suppose that over our sample, smart investors forecast this decrease in refinancing costs. Then, a smart investor would have also forecast a higher level of prepayments. However, if the OAS was based on a model calibrated to historical experience, then the OAS will be based on a model that “underpredict” prepayments.

If underprediction is a phenomena that only affected the general level of prepayments, then the time-effects in our regression specification will handle the mis-specification. For example, an increase in the general level of prepayment may occur if the reduction in refinancing costs led to more trading-up or relocation based prepayments across all mortgage pools. Thus, our regression

specification controls for an underprediction that can be written as,

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

where the γ_t is a common mis-specification in the rate of prepayment across all mortgages.

Alternatively, suppose that some types of mortgage pools, say for geographic reasons, had a higher level of underprediction, then the security effects in our specification will handle the mis-specification. Thus, our regression specification controls for all underprediction that can be written as,

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

Thus the only case that poses a problem for our results is if the underprediction were a function of both security and time.

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)
$\beta^k (\bar{c}_t - r_t)$	547.4 (5.18)	273.5 (4.35)	254.6 (4.73)
s_{kt}	-30.1 (-0.73)		
\bar{s}_{kt}		-11.4 (-0.32)	
$c^k \times r_t$			16.0 (6.82)
R^2	0.94	0.96	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.			

Suppose that the underprediction is proportional to $\bar{\phi}^k$, the current level of prepayments on mortgage- k . That is, suppose that $\hat{\phi}^k$ is equal to $\bar{\phi}^k$ times a constant. For example, this situation may arise if $\bar{\phi}^k$ proxied for refinancing desires of the average household in a particular pool, so that the lower costs of refinancing particularly sped up prepayments in that pool. 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,

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

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 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.

Another possible time and security effect interaction arises if the underprediction is related to the difference between the coupon on a particular mortgage pool and the market interest rate. This situation may occur if the lower costs of refinancing leads consumers to exercise their refinancing option more optimally in reaction to falling interest rates. This alternative may predominantly affect high coupon mortgages in low interest rate environments. 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 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 by including a term where c^k is interacted with r_t . While the interaction term is significant, the coefficient on our model continues to be large and significant.

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. We only need to include an interaction term between c^k and r to control for this hypothesis. The results reported in the last column of Table 9 confirm that our results are robust to this type of option mis-specification possibility.

We investigate the option mis-specification hypothesis 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,

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

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. Specifications (2) and (4) interact the option value with β^k , based on the 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)
$\beta^k (\bar{c}_t - r_t)$	495.3 (6.50)	570 (5.32)	477.6 (5.83)	590 (8.68)
OV_{kt}	-3.8 (-2.50)		-2.0 (-1.18)	
$\beta^k \times OV_{kt}$		17.7 (2.82)		15.3 (6.23)
R^2	0.94	0.95	0.94	0.95
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).				

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).

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

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 risk¹⁵ and the average coupon outstanding in the market are this common factor.

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

¹⁵Boudoukh *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.

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.

There are several avenues in which to extend our research. First, the evidence we present is from a data sample that is small, especially in relation to the MBS market. Including more securities in our cross-sectional analysis and extending our sample beyond 1998 are important steps to take in order to verify the generality of our results.

On the theoretical side, the result that a limited amount of capital sets prices in the MBS market 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 capital constraints in the form of a value-at-risk requirement. 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.

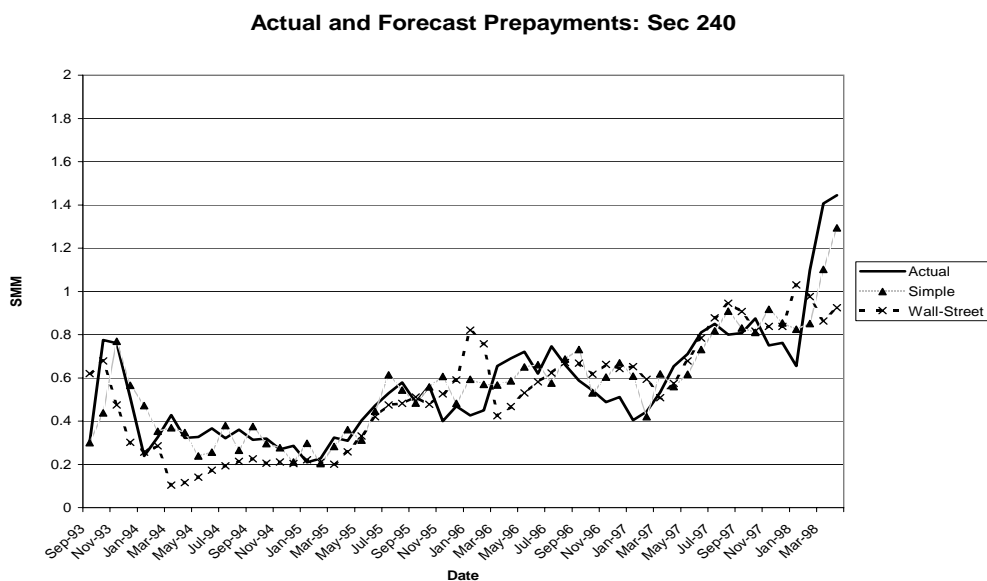


Figure 7: Actual and Forecast Prepayments: Sec 240

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.

Actual and Forecast Prepayments: Sec 237




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,

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

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

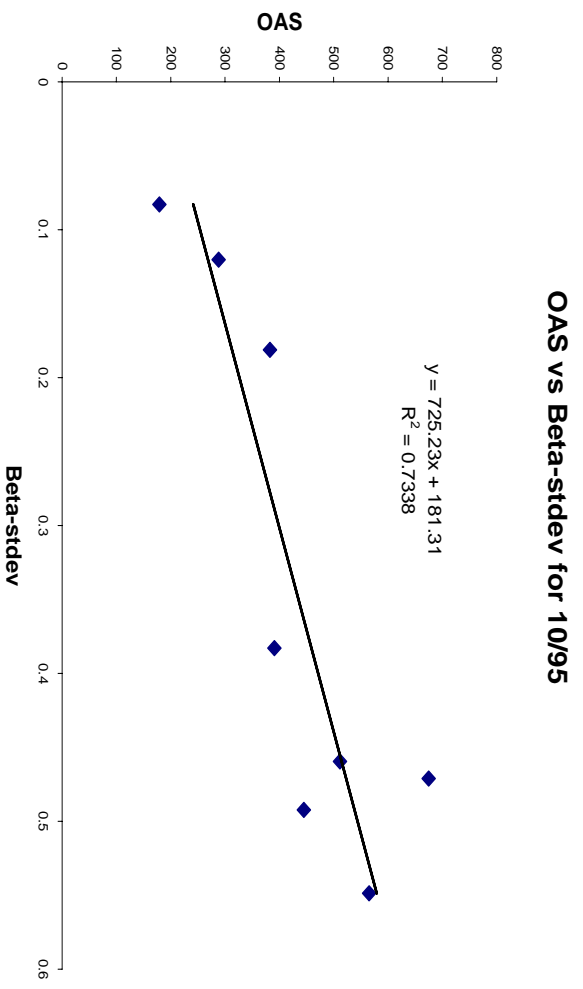


Figure 9: OAS versus β -stdev for 10/95




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

References

- [1] Allen, Franklin and Douglas Gale, 1994, Limited Market Participation and Volatility of Asset Prices, *American Economic Review* 84, 933-955.
Bennett, Paul, Richard Peach, and Stavros Peristiani, 2000, Implied Mortgage Refinancing Thresholds, *Real Estate Economics* 28, 405-434.
- [2] Berndt, Antje, Rohan Douglas, Darrell Duffie, Mark Ferguson, and David Schranz, 2004, Measuring Default Risk Premia from Default Swap Rates and EDFs, mimeo, Stanford University.
- [3] Boudoukh, Jacob, Matthew Richardson, Richard Stanton and Robert 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, Douglas, 1991, Risk, Return and Hedging of Fixed Rate Mortgages, *Journal of Fixed Income* 1, 85-107.
- [5] Breeden, Douglas, 1994, Complexities of Hedging Mortgages, *Journal of Fixed Income* 4, 6-41.
- [6] Caballero, Ricardo and Arvind Krishnamurthy, 2001, International and Domestic Collateral Constraints in a Model of Emerging Market Crises, *Journal of Monetary Economics* 48, 513-548.
- [7] Caballero, Ricardo and Arvind Krishnamurthy, 2002, A Dual Liquidity Model of Emerging Markets, *American Economic Review* 92, 33-37.
- [8] Caplin, Andrew, Charles Freeman and Joseph Tracy, 1997, Collateral Damage: How Refinancing Constraints Exacerbate Regional Recessions, *Journal of Money, Credit and Banking* 29, 496-516.
- [9] Case, Karl E., John M. Quigley, and Robert J. Shiller, 2003, Comparing Wealth Effects: The Stock Market Versus the Housing Market, Berkeley Mimeo.
- [10] Collin-Dufresne, Pierre, Robert Goldstein, and J. Spencer Martin, 2001, The Determinants of Credit Spread changes, *The Journal of Finance*, 66, 2177-2207.
- [11] Dow, James and Gary Gorton, 1994, Arbitrage Chains, *Journal of Finance* 49, 819-849.
- [12] Downing, Chris, Richard Stanton and Nancy Wallace, 2003, An Empirical Test of a Two-Factor Mortgage Prepayment and Valuation Model: How Much Do House Prices Matter?, Berkeley mimeo.
- [13] Dunn, Kenneth B. and John J. McConnell, 1981, Valuation of GNMA Mortgage-Backed Securities, *Journal of Finance* 36, 599-616.
- [14] Froot, Kenneth, 2001. The Market for Catastrophe Risk: A Clinical Examination, *Journal of Financial Economics* 60, 529-571.

- [15] Froot, Kenneth A., and Paul G. O'Connell, 1999, The Pricing of US Catastrophe Reinsurance in *The Financing of Catastrophe Risk*, edited by Kenneth Froot, University of Chicago Press.
- [16] Gabaix, Xavier, Parameswaran Gopikrishnan, Vasiliki Plerou, H. Eugene Stanley, 2003, A Theory of Power Law Distributions in Financial Market Fluctuations, *Nature* 423, 267-70.
- [17] Geanakoplos, John, 2003, Liquidity, Default, and Crashes in Mathais Dewatripont; Lars Peter Hansen; Stephen J. Turnovsky, eds. *Advances in Economics and Econometrics II* (Cambridge University Press)
- [18] Geanakoplos, John, 1997, Promises Promises, in W.Brian Arthur, Steven N. Durlauf, and David A Lane, eds. *The Economy as an Evolving Complex System II*, (Sante Fe Institute Studies in the Sciences of Complexity, Vol 5)
- [19] Gromb, Denis and Dimitri Vayanos, 2002, Equilibrium and Welfare in Markets with Financially Constrained Arbitrageurs, *Journal of Financial Economics* 66, 361-407.
- [20] Grossman, Sanford J. and Zongquan Zhou, 1996, Equilibrium Analysis of Portfolio Insurance, *Journal of Finance* 51, 1379-1403.
- [21] Holmstrom, Bengt and Jean Tirole, 1997, Financial Intermediation, Loanable Funds, and the Real Sector, *Quarterly Journal of Economics* 112, 663-691.
- [22] Hurst, Erik and Frank Stafford, forthcoming, Home is Where the Equity Is: Liquidity Constraints, Refinancing and Consumption, *Journal of Money, Credit and Banking*.
- [23] Kau, James B., and Donald C. Keenan, 1995, An Overview of the Option-Theoretic Pricing of Mortgages, *Journal of Housing Research* 6, 217-44.
- [24] Kupiec, Paul and Adam Kah, 1999, On the Origin and Interpretation of OAS, *Journal of Fixed Income* 9, 82-92.
- [25] Kyle, Albert S. and Weu Xiong, 2001, Contagion as a Wealth Effect, *Journal of Finance* 56, 1401-1440.
- [26] Longstaff, Francis A., 2004, Optimal Recursive Refinancing and the Valuation of Mortgage-Backed Securities, NBER Working Paper # 10422.
- [27] Matthey, Joe and NancyWallace, 2001, Housing-Price Cycles and Prepayment Rates of U.S. Mortgage Pools, *The Journal of Real Estate Finance and Economics* 23, 161-184.
- [28] Merton, Robert C., 1987, A Simple Model of Capital Market Equilibrium with Incomplete Information, *Journal of Finance* 42, 483-510.
- [29] Patruno, Gregg N., 1994, Mortgage Prepayments: A New Model for a New Era, *Journal of Fixed Income* 4, 42-56.
- [30] Richard, Scott F., and Richard Roll, 1989, Prepayment on Fixed-Rate Mortgage-backed Securities, *Journal of Portfolio Management* 15, 375-392.

- [31] Schwartz, Eduardo S., and Walter N. Torous, 1989, Prepayment and the Valuation of Mortgage-Backed Securities, *Journal of Finance* 44, 375-392.
- [32] Shleifer, Andrei and Robert W. Vishny, 1997, The Limits of Arbitrage, *Journal of Finance* 36, 35-55.
- [33] Stanton, Richard H., 1995, Rational Prepayment and the Valuation of Mortgage-Backed Securities, *Review of Financial Studies* 8, 677-708.
- [34] Stein, Jeremy C., 1995, Prices and Trading Volume in the Housing Market: A Model with Down-Payment Effects, *The Quarterly Journal of Economics* 110, 379-406.