

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

XAVIER GABAIX, ARVIND KRISHNAMURTHY, and OLIVIER VIGNERON*

ABSTRACT

“Limits of Arbitrage” theories hypothesize 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 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, consistent with the specialized arbitrageur hypothesis.

A NUMBER OF RECENT “limits of arbitrage” theories predict that 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), Geanakoplos (2003), or He and Krishnamurthy (2006)). These theories are often referenced in explaining episodes of market illiquidity and falling asset prices, such as the events during the fall of 1998.

In traditional asset pricing theory, the marginal investor in every asset market is the same broadly diversified representative investor. Thus, a hedge fund liquidating \$200 billion of mortgage-backed securities finds a large pool of ready buyers (comparable to the entire capital market), and 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 marginal investor in a particular asset market is an investor who specializes in that market. In this case, liquidations can have large effects on prices given that the relevant set of buyers is a smaller, specialized pool of investors.

*Authors, respectively, are at MIT and NBER, Northwestern University, and BNP Paribas. We thank Kobi Boudoukh, Donald Bronstein, John Campbell, Kent Daniel, Thomas Davidoff, 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, and Bill Wheaton for valuable comments. We are particularly grateful to the referee for many constructive suggestions. We thank participants at the NBER Asset Pricing meeting, Econometric Society Summer meeting, Gerzensee-CEPR meeting, WFA, Utah Finance Conference, and at 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-Breeden, and UBS for providing data. Kripa Freitas provided excellent research assistance. All errors are our own.

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 the marginal investor in the MBS market is a risk-averse fund manager who has all of his wealth invested in this market. The novelty of the theory lies in its implications for the pricing of risk specific to the MBS market. In contrast to traditional asset pricing theory, we predict that MBS market-specific risk will carry a positive risk premium. Moreover, we predict that the market price of this risk will be high when the total riskiness of the MBS market is also high. We argue that prepayment risk is a good example of an MBS market-specific risk, and we 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 MBS investor. 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 that drive both aggregates and homeowner prepayments. In 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 implies a sign that is opposite that required to match the observed prices of prepayment risk under traditional asset pricing theory. This suggests that the marginal investor in the MBS market is not the representative investor 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 exclusively in the MBS market. We show that when the fund manager has a coefficient of relative risk aversion of five, 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 and 1994 (collapse of the Askins Capital Management MBS fund) and 1998 and 1999 (LTCM crisis) in which MBS spreads widened 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. Specifically, the marginal investor must be an MBS specialist rather than a broadly diversified representative investor.

The main difficulty we encounter in establishing our results arises from a measurement problem. In order to draw a relation between prepayment risk and the risk premium on an MBS, we need to measure the risk premium of an MBS. We proxy for this risk premium using the security's option-adjusted spread (OAS). The OAS measures the yield on an MBS in excess of Treasuries, after accounting for the value of the homeowner's prepayment option. However, accounting for the prepayment option depends on the specific pricing model used to compute the OAS. Thus, the OAS is a model-dependent measure of the risk premium of an MBS. To the extent that the pricing model is misspecified, the OAS is a noisy measure of the true risk premium, potentially rendering the interpretation of our results suspect (e.g., see Kupiec and Kah (1999)).

One way to bypass the misspecification issue is to use actual MBS returns, a proxy for expected returns, as the dependent variable in our regressions. We decide against this approach because of data limitations. Because actual bond returns are a very noisy estimate of the expected return on the securities, we need more data than we have to implement these regressions. Thus, while using the OAS greatly reduces this measurement error problem, it does so at the cost of raising the possibility of misspecification.

We address the possibility of misspecification in the OAS in three ways. First, our theory predicts that securities with more prepayment risk will have differentially higher risk premia at times when the market-wide price of prepayment risk is higher. We provide support for the theory from panel data regressions in which the independent variable is the interaction between constructed measures of security-specific prepayment risk and market-wide price of prepayment risk. Since the theory is based on the interaction between a security measure and a time measure, we are able to apply both time dummies and security dummies as controls in our regressions. The profusion of dummy variables controls for all possible linear time-specific and security-specific sources of OAS misspecification. Indeed, our results are strengthened when these controls are added.

Second, as we explain in the next sections, our theoretical measure of the market-wide price of prepayment risk is partly based on the *average coupon of all traded MBS*. While a misspecification explanation of the OAS may involve the coupon of a specific MBS, it is unclear why it would involve the average market-wide coupon on mortgages. We present regressions that isolate the average coupon measure and show that it has independent explanatory power for the OAS. Such explanatory power is consistent with our theory of the determination of risk premia, but inconsistent with a pure misspecification explanation of the OAS.

Lastly, we introduce controls in our regressions for specific alternative hypotheses for the determination of the OAS, and show that our results continue to hold up. Two of the hypotheses we consider are, (1) prepayment models underpredict prepayments relative to the true model, and (2) prepayment models

undervalue the homeowner's prepayment option. We show that these hypotheses can explain positive OAS. However, introducing controls that proxy for the underprediction of the prepayment or the undervaluation of the option does not change our conclusions.

The theory we develop in this paper belongs 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. Papers in the literature include Dow and Gorton (1994), Shleifer and Vishny (1997), Kyle and Xiong (2001), Gromb and Vayanos (2002), Geanakoplos (1997, 2003), Gabaix et al. (2003, 2006), and He and Krishnamurthy (2006). Gabaix et al. (2003, 2006) relate spikes in returns and trading volume to the trades of large institutional investors in illiquid markets. Allen and Gale (1994) study an environment in which traders must specialize *ex ante* in a certain asset market, which implies that *ex post* there is limited market participation and the specialized traders become marginal in setting prices.¹ Caballero and Krishnamurthy (2001, 2002) present a model of emerging market crises in which crises are events where the marginal investor switches from a broadly diversified world investor to an investor within the emerging market. Our paper is most similar to He and Krishnamurthy (2006) in that we explicitly model how financial intermediation affects market participation and the preferences of the marginal investor.

Our empirical results are consistent with some existing papers in the MBS literature. Brown (1999) presents evidence that the spreads of mortgage securities (relative to Treasuries) comove with the spreads on corporate bonds (relative to Treasuries). Brown argues that this evidence suggests time variation in the market price of prepayment risk on MBSs. However, Brown acknowledges that the latter evidence is also consistent with time variation in the liquidity premium on Treasury securities (see Krishnamurthy and Vissing-Jorgensen (2006) for evidence on this point). Our study differs from Brown in that Brown analyzes only collaterals while we analyze collaterals and IOs, and he studies a sample of MBSs that is different from ours. However, he arrives at a similar conclusion regarding the existence of a positive market price of prepayment risk.²

Boudoukh et al. (1997) study the pricing of GNMA securities under a benchmark multifactor interest rate model that they propose. They focus on a pricing function that depends only on the yield curve, thus setting aside prepayment information. One of their main findings is that a single (noninterest rate) factor accounts for 80 to 90% of the common variation in the pricing errors. Our study suggests that a candidate for the common factor is prepayment risk³ and the average coupon outstanding in the market.

¹ Merton (1987) presents a model in which segmentation arises endogenously, and he explores the implications of market segmentation for asset prices.

² Cohler, Feldman and Lancaster (1997) and Levin and Davidson (2005) present pricing methodologies in which the OAS is in large part a prepayment risk premium.

³ Boudoukh et al. do hint at this by looking at the prepayment of the different coupons. They find that for lower coupons, which have a lot of relocation-based prepayment, prepayment variables explain a significant fraction of the pricing errors.

Similar common factor phenomena have been documented in other asset markets. Collin-Dufresne, Goldstein, and Martin (2001) study the corporate bond market. They find that a simple Merton (1974) model explains very little of the variation in corporate bond prices. Even after including macro factors, such as the stock market, 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.^{4,5} Unlike us, they are unable to identify either the risk factor or the marginal investor who is pricing the risk.

Froot and O'Connell (1999) demonstrate effects similar to ours in the market for catastrophe insurance. They note that there are times during 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. Therefore, when the capital in the insurance market is low, insurers are less willing to write catastrophe insurance. This drives prices up and quantities down as found empirically.

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 is investigated more recently by Stanton (1995) and Longstaff (2004).⁶ 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 of 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' refinancing costs, for example, are 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 to introduce 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 with market conditions.

⁴ Berndt et al. (2004) present similar evidence based on data from credit default swaps. They find large swings in the risk premia incorporated in default swaps.

⁵ Bates (2003) and Bollen and Whaley (2004) find related effects in the options market. Garleanu, Pedersen, and Poteshman (2005) provide further evidence for options and a model also based on the limited diversification of the marginal options investor.

⁶ See Kau and Keenan (1995) for a survey of this line of research.

Section I formally presents the theoretical model that motivates our empirical tests. The theory links the risk premium on an MBS to the prepayment risk on the MBS and the market-wide price of prepayment risk. Section II presents evidence in support of this theory, using the OAS as a proxy for the security's risk premium. Section III discusses in detail our empirical strategies to deal with the OAS measurement problem, and provides robustness checks of our main results. Finally, Section IV concludes.

I. The Model

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 mortgage, the MBS is a debt security with an amortizing principal value. However, the fact that consumers have the option to prepay their mortgages, makes valuing (and hedging) an 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, and real estate prices. 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, and so on. 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 of these tranches are sophisticated investors such as hedge funds or investment banks that have some expertise in assessing prepayment risk. The success of a securitization of mortgages often hinges 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 pre-specified 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 (1989)). Second, one simulates several interest rate paths, discounting and averaging the cash flows based on a term structure model calibrated to current market risk-free rates.

The model-implied prices under this methodology typically differ from quoted market prices. This difference is termed an option-adjusted spread (OAS). 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 noninterest rate risk premium on the security.

In this section we describe a very simple environment for studying the pricing of MBSs and the determination of 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 MBSs.

A. Mortgage-Backed Securities with No Prepayment Risk

Consider a world with a constant interest rate r and a mortgage pool with constant prepayment rate ϕ and coupon c (and no credit risk). At any date t , the amount 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 issued against this pool. The IO is defined as the claim on all of the coupons (interest payments) 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 mortgage pass through itself—the collateral—is defined as the claim on all cash flows from the mortgage pool, that is, the coupons (interest payments) and the principal repayment. Thus, its value is the value of the IO plus the value of the PO:

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

B. 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 discount rate between periods 0 and 1 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 is 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

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

However, 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.⁷ 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 value with respect to the prepayment rate.

Similarly, 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

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

C. OAS

Let P_{IO}^k and P_{PO}^k be the date 0 prices of one dollar of face value of IO and PO, respectively. The OAS is defined as the premium over the riskless rate r required to set the present 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

⁷ Unlike our abstraction, in practice interest rates are uncertain. The logical extension of our model to the uncertain interest rate case is to write $\bar{\phi}^k(\bar{r})$. 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.

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

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 over r required to recover the actual market price.

There are two ways to interpret the OAS. First, it may simply reflect a misspecified model of the prepayment option. Perhaps informed market participants have a true model of prepayments that 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.⁸

A second way to look at the OAS is as a risk premium. Any time that prices differ from expected values, the OAS will be nonzero. 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 misspecified 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).

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 value bond plus the payoff on $\frac{c^k - r}{c^k}$ of the IO. Thus, by no 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

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

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.

⁸ It is also possible that the OAS is due to a Jensen's inequality term. However, we think that this Jensen's inequality effect is unlikely to be very important. Indeed, since $P_{IO} = E[c/(r + \hat{\phi})] > c/(r + E[\hat{\phi}])$, 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.

Intuitively this is because shocks that lower the value of the IO, that is, 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 about how risks are priced.⁹

D. 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 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. Their lack of expertise renders them vulnerable to exploitation by MBS specialists, who may be able to sell them securities with a value significantly less than their price. As a result, they give their funds to the specialized MBS fund manager who invests on their behalf.

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 on a fund manager for a fund of size w_f .¹⁰ 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

⁹ The OAS for the PO is defined by

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

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

$$OAS_{PO}^k = -\frac{r}{\bar{\phi}^k + OAS_{IO}^k} OAS_{IO}^k. \quad (19)$$

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.

¹⁰ See Holmstrom and Tirole (1997) and He and Krishnamurthy (2006) for models of capital constraints in intermediation.

$$U(w_1) = E[w_1] - \frac{\rho}{2} \text{Var}[w_1], \tag{10}$$

that is, he is a mean-variance maximizer.

E. Equilibrium

At date 0, fund managers raise a total of $w_f - w_m$ from investors. This gives fund managers total capital of $\frac{w_m}{\alpha}$. With this sum the fund managers purchase a portfolio of mortgage-backed securities. Let x_{IO}^k and x_{PO}^k be the amount of the k^{th} IO and PO held in a portfolio. Then,

$$w_{f,1} = \frac{w_m}{\alpha} + \sum_k x_{IO}^k (V_{IO}^k - P_{IO}^k) + \sum_k x_{PO}^k (V_{PO}^k - P_{PO}^k) \tag{11}$$

is the date 1 value of the portfolio. Since the fund manager’s wealth, w_1 , increases linearly with $w_{f,1}$ (at slope of α), his problem is to maximize (10) given (11) and subject to the budget constraint

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

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 MBSs.

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 the following expression for the price of the IO:

$$\frac{c^k}{r + \bar{\phi}^k} - P_{IO}^k = -\rho\alpha \text{cov} \left(\frac{c^k}{r + \bar{\phi}^k} \eta^k \Delta\phi^k, R_M \right), \tag{12}$$

where the market is defined as

$$R_M = \sum_j \frac{\theta^j}{(r + \bar{\phi}^j)^2} \Delta\phi^j (r - c^j). \tag{13}$$

The term on the right-hand side of (12) is a risk premium for holding prepayment risk. 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.

F. Covariance Structure

We make the following simplifying assumption on the covariance structure:

$$\Delta\phi^k = \beta^k \Phi + \epsilon^k, \tag{14}$$

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,¹¹

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

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 systematic prepayment risk. To compute the sum requires that one 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 a 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 (FNMA, FHLMC, GNMA), where the weights are the amounts outstanding of each mortgage security. The relation we use in our tests is,¹²

$$OAS_{IO}^k = \overbrace{\beta^k}^{\text{Systematic risk}} \times \underbrace{\rho\alpha(\bar{c} - r)}_{\text{Market price of risk}}, \tag{15}$$

where \bar{c} is the weighted average coupon and $\rho\alpha$ is the effective risk aversion of the MBS fund manager (α 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 (15) 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 MBS market as a whole is at a premium—that is, coupons exceed r —faster prepayments are costly to the representative fund manager. Thus, securities whose value decreases because of faster prepayments command a

¹¹ The 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 (12) with (7), and noting that $\eta^k = 1/(r + \bar{\phi}^k)$.

¹² We also develop a continuous time model to express the relation between the OAS and prepayment risk. The resulting expressions are very similar to those 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>.

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. With some algebraic manipulation, the OAS for the PO is equal to

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

Another example of a security whose value increases with faster prepayment is a *discount* collateral. Collateral with a coupon below the market interest rate increases in value if the mortgage prepays faster than expected. Given relations (15) and (9), we can write the OAS on the collateral as

$$OAS_C^k = \beta^k \times \rho\alpha a (\bar{c} - r)(c^k - r) \times \frac{1}{c^k + \bar{\phi}^k + OAS_{IO}^k}. \quad (16)$$

Thus the OAS on the collateral depends on both whether the MBS market as a whole is at a premium as well as whether 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 MBS market as a whole is at a discount. In this case, faster prepayments increase the value of the market. Hence, the IO has a negative risk premium while the PO commands a positive risk premium.

The dependence of the price of prepayment risk on $(\bar{c} - r)$ is principally a general equilibrium implication. It seems plausible that the relation between β^k and the OAS could be spurious, or due to model misspecification, 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.

G. Testable Empirical Predictions

The main predictions of the model are contained in equation (15), which we can decompose as

$$OAS_{IO}^{kt} = \beta^k \lambda_t \quad (17)$$

$$\lambda_t = \rho\alpha a (\bar{c}_t - r_t), \quad (18)$$

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

- In the cross section, the loading of IO- k on the common component of prepayment uncertainty explains the OAS on the IOs.
- 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 .

- In the cross section, the residual prepayment risk of security k (i.e., $\sigma(\epsilon^k)$) is not priced.
- Equation (16) 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 .

H. Discussion of Assumptions

The model we present above 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 that OAS_{IO} is proportional to $\beta^k \times (\bar{c} - r)$. This relation is likely to be robust to more sophisticated models, although not necessarily in the simple linear form we derived.

On the other hand, the simplification in the derivation means that there are probably other factors that affect 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) documents a positive relation between OAS and implied volatilities on Treasury bond options, which suggests that there is a misspecification in Wall Street prepayment models used to derive the OAS. These issues suggest that the OAS is a 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.

Next, we omit capital constraints from the model, which are significant in the limits of arbitrage literature. Mainly, this omission follows from the fact that 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.

Additionally, in our 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 in which 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 nonmonotonicities in the behavior of fund managers in the MBS market.

Finally, we derive our results in a static CAPM framework. In a dynamic model the current wealth of the fund managers will be an important state

Table I
Description of the IO/POs

	IO/PO Pool Number								
	249	240	252	272	264	237	270	267	268
Coupon ^a	7.08	7.49	7.95	8.07	8.49	8.48	9.01	8.91	9.64
Age ^b	58	60	63	27	50	70	80	47	110
Size ^c	1,375	3,450	1,975	1,055	1,050	1,725	898	1,155	567

^aWeighted average coupon (WAC) in % on underlying mortgage pool (± 5 bps over sample).

^bSeasoning of underlying pool in months as of July 1998.

^cSize of underlying pool at original issue date, in millions.

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.

II. Data and Estimation

We rely on two data sets. Our first data set comprises the OAS for nine IOs and POs (see Table I) furnished by Salomon Smith Barney. These data are daily and cover 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 are fairly representative in age and coupon of the active secondary market. The collateral are all FNMA 30-year conventional loans, uniformly drawn from a mix of loans from across the country. The largest representation is from California, New York, Texas, Florida, and Illinois.

The bulk of our analysis is conducted using the IO data. We check 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.

Our second data set comprises quarterly observations on the OAS for six generic (TBA) FNMA 30-year collateral covering the period from October 1987 to July 1994. The coupons on these securities range from 7.5% to 11% and the data are provided by Smith Breeden. We do not have prepayment information for these pools. We test our model using these data because they cover a period with significant variation in the average coupon. We discuss these tests in further detail in Section F.

We construct time series of monthly OAS data for the IOs by forming simple averages of the daily figures. This reduces microstructure effects. The data are an unbalanced panel, with common last observations but varying initial observations.

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

There are two steps in testing (17) and (18). 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 mortgage pools of FNMA, FHLMC, and GNMA. This variable is computed as the average of the underlying coupons on individual mortgage pools, weighted by the amount outstanding of the pools. We obtain monthly data starting from 1988 to form our estimate of \bar{c}_t . We use the 10-year constant maturity Treasury yield from the Federal Reserve's Board of Governors web site to form r_t .

The estimate of β^k is more complex. We first develop a bare-bones statistical prepayment model. For each IO, we have the historical payoff 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) shows that this feature arises naturally in a setting with transaction costs of refinancing. The term involving age_t captures seasoning effects in mortgage pools. We assume that the error follows an AR(1) process, that is,

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

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.

We use two proxies for β^k . In the first line of Table II we present the sample standard deviations of the residuals, \hat{u}_{kt} . In many of our tests we use these standard deviations as β^k .

We also perform a principal components analysis of the errors and use security- k 's loading on the common factor as β^k . We focus only on the overlapping observations (22 months) for this analysis.¹³ The first eigenvector accounts for 84% of the variance, which suggests that equation (14) is a good representation of the data. The second and third components account for 8.5% and 3%, respectively. Table II (second line) presents the loading on the first eigenvector for each security as well as the standard deviation of the residual (third line).

¹³ We have also conducted 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.

Table II
Measures of Prepayment Risk

For each of the nine securities, we present β^k (st. dev.), the measure of prepayment risk based on the standard deviation of prepayment error for security k ; β^k (PCA), the loading of the prepayment error for security k on the first eigenvector of prepayment errors in a principal component analysis; $idiosync^k$ (PCA), the standard deviation of the residual from the principal component analysis; and β^k (mtge. model), the standard deviation of the residual from a Wall Street mortgage model.

	IO/PO Pool Number								
	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
$idiosync^k$ (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

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 the standard deviation of the residual 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 the explanatory power of idiosyncratic risk is not informative.

Our measure of β^k captures how prepayment uncertainty varies across pools. In practice, pools differ in the terms of their underlying mortgages—for example, ages, coupons, and geography. These factors lead to different prepayment uncertainty across pools.

To measure β^k , we use 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 uncertainty 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 nonlinear (prepayment functions are typically complex nonlinear 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 rates 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 use these residuals to form β s. The β s look similar, suggesting that our model is reasonable. See the last line of Table II.

B. Interest Rates, Average Market Coupon, and OAS

Figure 1 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




Figure 1. 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.

average coupon are slow compared to the large movements of market interest rates. Prior to 1993, prevailing mortgage rates were around 10% to 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 from values of 9% to 10% down to 7% to 8%. We follow the evolution of the OAS of the IOs and POs 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 basis points (“bps”) from January 1995 to January 1996. Rates fell continuously from March 1997 to July 1998 by slightly more than 100 bps to reach levels as low as those of November 1993. By the end of our sample period, the outstanding coupon had adjusted down to 7.5%.

Figure 2 shows the variation of the OAS of the IOs in our data over the period Autumn 1993 to Spring 1998. One readily observes the large swings of the OAS of the IOs, from values above 500 bps in the beginning of the period to values close to zero in 1994 and 1996 when interest rates were very high. The OAS of POs give a somewhat symmetric image, although at smaller magnitudes (as predicted by equation (19)). One should also note that the interest rate alone is




Figure 2. 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 are from Salomon Smith Barney.

not enough to understand the relative magnitude of the OAS of the IO between 1993 and 1998 when rates were at the same level: OASs are much higher in 1993 than in 1998. This is not a puzzle in light of our derivations, since equations (17) and (18) tell 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 35% higher in 1993 than in 1998.

C. Cross-sectional Estimates of the Market Price of Risk

We run one cross-sectional regression 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. Note that β^k , given in Table II, 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 misspecification 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.

Figure 3 graphs the estimate of λ_t using β -st. dev. as well as the one-standard deviation envelopes around the estimate.¹⁴ The estimation errors of λ_t are

¹⁴ The graph of the estimate of λ_t using the β -PCA is very similar to Figure 3.




Figure 3. λ_t estimates. The measure of the monthly market price of prepayment risk (λ), obtained from estimating equation (17), is plotted. We use β -st. dev. as the measure of security-specific prepayment risk. To compare this empirical measure to our theoretical prediction of equation (18), we plot $\bar{c}_t - r_t$, where \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.

uniformly tight, and λ_t is significantly different from zero for each month. We interpret these results as supportive of the theory because they suggest that prepayment risk is related to the OAS, and that our measure of β^k is indeed picking up the cross-sectional prepayment risk of the IOs.

The average value of the estimated λ_t 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.

Also pictured in the figure is the difference between \bar{c}_t and r_t . At a broad level, λ_t and $\bar{c}_t - r_t$ 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 result from a falling ρ over the sample period. It is well documented that in the 1993 to 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, causing a higher effective

Table III
Explaining the OAS, Baseline Specification

Regressions based on the OAS of the IOs:

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

where, d_t is a time dummy, d_k is a security dummy, \bar{c}_t is the average coupon outstanding, and r_t is the 10-year interest rate. 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 -st. dev.						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)
$idiosync^k(\bar{c}_t - r_t)$								478 (0.48)	
Security effects	Yes	Yes	Yes	No	Yes	No	Yes	Yes	No
Time effects	Yes	Yes	Yes	Yes	No	No	Yes	Yes	No
R^2	0.94	0.95	0.92	0.93	0.68	0.20	0.93	0.93	0.18
N	383	194	189	383	383	383	383	383	383

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

ρ . As time passed, capital flowed back into these funds and the effective ρ fell. Froot (2001) finds this effect in the catastrophe insurance market.

D. Tests Using the Entire Panel

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

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

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 dummies (d_t and d_k), thereby controlling for any alternative hypothesis 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 regression residuals (ϵ_{kt}). We correct for this in two ways. First, most of the regressions report t -statistics that are corrected for serial correlation in the residuals at the security level (we cluster the residuals at the security level). Second, we run regressions using first-differenced data and report the results in Table III. Another potential problem is correlation in the regression residuals across securities at a single point in time. However, this problem is less severe in our specification because the regressions include a

time dummy that absorbs all common innovations in the OAS. We have also conducted a robustness check using a standard panel data adjustment whereby we assume that ϵ_{kt} is AR(1) at the security level and correlated across securities. We find that our results remain highly significant.

The results in columns (1) to (8) of Table III verify that our model fits the data. The specification in column (1) uses β -st. dev., while the specification in column (7) uses β -PCA. Columns (2) and (3) give the results from two subsamples, formed using June 1996 as the dividing point between the two (there are fewer observations in the early subsample). The coefficient estimate using β -st. dev. is 462, while it is 441 using β -PCA. In the cross section, our measure of β^k -st. dev. 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 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 -st. dev. 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.

Note the lower (but still highly significant) coefficient in specification (3) compared to (2). This result agrees with our conjecture that there is more risk-bearing capacity (i.e., lower ρ) in the latter half of the sample.

Column (8) contains the result of the following regression:

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

Our theory predicts that the idiosyncratic risk should not be priced. Unfortunately, as we mention 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 error on the coefficients to blow up, so neither is significant.

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

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

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) vs (3)).

Table IV
Explaining the OAS, First Differences

Regressions based on the OAS of the IOs:

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

where, β^k is the β -st. dev., d_t is a time dummy, \bar{c}_t is the average coupon outstanding, and r_t is the 10-year interest rate. Results for the full sample are reported in (1). 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 are not reported.

A comforting aspect of the results in Table IV is that monthly changes in OAS spreads correspond more closely to changes in the underlying market prices of the IOs. If interest rates do not change from one month to the next, but the IO price does change, then it must be the case that the OAS changes. In fact, interest rates do change somewhat, but since the OAS is 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.

E. Corporate Bond Spreads

Brown (1999) presents evidence that the spreads of mortgage securities comove with the spreads on corporate bonds (both relative to Treasuries). Although Brown argues that the evidence suggests time variation in the market price of prepayment risk on MBS, he acknowledges that the evidence is also consistent with time variation in the liquidity premium on Treasury securities. Since both mortgage and corporate bond spreads are measured relative to Treasuries, a common time-varying liquidity demand for Treasuries will drive both spreads (see Krishnamurthy and Vissing-Jorgensen (2006) for evidence on this point).

Table V presents our baseline regressions, altered by replacing the time dummies with the AAA-Treasury bond spread. Although the time dummies subsume any variation in the corporate bond spread, the more parsimonious specification offers some insight into the sources of variation in the OAS. The AAA-Treasury spread data are from Lehman Brothers and are measured in percentage units.

Consistent with Brown (1999), we find that the corporate bond spread comoves with the OAS. Over our sample the corporate bond spread varies from 0.37% to 1.05%. Using the coefficient estimate of 674.1 in column (1), the

Table V
Explaining the OAS, Role of the Corporate Bond Spread

Regressions based on the OAS of the IOs:

$$OAS_{IO}^{kt} = \sum_k \gamma_k d_k + A \times \beta^k (\bar{c}_t - r_t) + B \times (AAA - Treasury) + \epsilon_{kt},$$

where the results are reported for β^k -st. dev. in (1) and β^k -PCA in (2), d_k is a security dummy, \bar{c}_t is the average coupon outstanding, and r_t is the 10-year interest rate. *AAA - Treasury* is the high-grade corporate bond spread, from Lehman Brothers.

	(1)	(2)
$\beta^k (\bar{c}_t - r_t)$	318.4 (8.66)	307.5 (6.01)
<i>AAA - Treasury</i>	674.1 (6.90)	682.2 (6.75)
R^2	0.76	0.76
N	383	383

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

corporate bond spread variation implies variation in the OAS of 458 bps. The coefficient estimate on our model is 318.4 in column (1). For a β^k of 0.5, variation in $\bar{c} - r$ from -0.1 to 3.06 implies variation in the OAS of 503 bps. Thus, both the corporate bond spread and our theoretically motivated price of prepayment risk explain similar magnitudes of OAS variation.

F. 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. Although the bulk of the variation in $\bar{c}_t - r_t$ is driven by variation in r_t , from Figure 1 we note that there is some variation 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}_t has independent explanatory power for the OAS. We begin this section by reporting results from the IO sample. Unfortunately, there is only slight variation in \bar{c}_t over the IO sample. We present our main results from the pass-through data from Smith Breeden, which cover a period with more variation in \bar{c}_t .

Over the period from September 1993 to April 1998 the average coupon outstanding falls from 8.45% to 7.52% (see Figure 1). This fall occurs in two phases: early in the sample, and again late in the sample. We estimate the following regression for the IOs:

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

The coefficient estimates are $A_1 = 1474$ (8.61) and $A_2 = -162$ (4.25) ($R^2 = 94\%$ for $N = 383$). The coefficients are significant and of the right sign, but of

(substantially) different magnitudes, contrary to the theory. It seems likely that the large coefficient on $\beta^k \bar{c}_t$ is due to overfitting, driven by the fact that \bar{c}_t moves only a little over this sample period.

The largest variation in the average coupon occurs over the period from mid-1991 to early 1994. We have OAS data for pass-throughs over this period from Smith Breeden, which we also use to test for the explanatory power of the average coupon. The data are for the OAS on FNMA 30-year generic collateral for eight bonds with coupons ranging from 7.5% to 11%. Our data span a period from October 1987 to July 1994. From our theory (see equation (16)), the pass-through takes a quadratic form in interest rates. We estimate the following regression:

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

where the \bar{c}_t and r_t are measured in percentage units, and the OAS is measured in basis points. Our theory predicts that A_1 is positive, A_2 is negative, and $A_1 + A_2 = 0$; A_3 should not have any explanatory power.

The results are reported in Table VI. We run the first set of regressions 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. In terms of magnitudes, if we consider the 9% coupon bond, along with typical values in our sample of an r of 7.5% and $\bar{c} - r$ of 2%, our theoretically motivated variable predicts an OAS of 35 bps. The average OAS over our sample for the

Table VI
Explaining the OAS, Average Market Coupon

Regressions based on the OAS of the collateral:

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

where, d_k is a security dummy, \bar{c}_t is the average coupon outstanding, r_t is the 10-year interest rate, and c^k is the coupon on collateral- k . The last column reports the p -value from testing $A_1 + A_2 = 0$.

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	11.6 (10.45)	-11.7 (-3.18)	-33.3 (-1.36)	0.86	222	0.99

The ALL BONDS regression uses the entire panel, with security fixed effects.

Estimates reported with t -statistics based on robust standard errors in parentheses.

9% coupon bond is 75 bps.¹⁵ It is also encouraging that A_1 and A_2 have opposite signs and similar magnitudes. However, 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. That is, while in our one-period model the OAS is proportional to our current measure of the market price of risk (i.e., $\bar{c} - r$ today), in a richer model, the OAS would reflect all of the future expected market prices of risk, that is, the expected future values of $\bar{c} - r$. As in the 1980s and 1990s, nominal rates were largely declining, a negative A_3 captures the market's expectation that the average coupon will decrease in the future. Alternatively, this discrepancy could be due to a misspecification of the option value in the underlying prepayment model from which the OAS is generated, or a misspecification of the interest rate in our simple empirical implementation.

G. 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.

¹⁵ We 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 report earlier. For the IOs we predict that

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

and estimate that A is 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 neighborhood of the coefficient estimates for the collateral.

We form a time series of prepayment risk innovations from our estimates of \hat{u}_{kt} . For each t , we compute

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

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 doing so 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 (over the short-term rate) on the S&P500 (SP_t) (a proxy for aggregate wealth). The excess return series has a correlation coefficient of 0.032 with the prepayment innovation series. 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 (2005)). Caplin et al. 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 (2004) 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 VII 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 FHLMC'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.

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

Table VII
Correlations between Shocks to Prepayment and Shocks to Real Estate Values

Correlations between prepayment shocks (U_t) and measures of real estate price appreciation for both regional indices and a national index. $\text{corr}(0)$ is the contemporaneous correlation. $\text{corr}(-1)$ 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
$\text{corr}(0)$	0.01	0.15	0.03	0.06	0.19	0.18	-0.03	0.09	0.04	0.08
$\text{corr}(-1)$	-0.12	-0.11	-0.07	0.11	0.12	-0.03	-0.17	0.15	-0.21	-0.11
$\text{corr}(-2)$	0.04	0.21	0.21	0.26	0.18	0.12	0.03	0.40	0.08	0.18

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 IOs are positive, while those on POs are negative.

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 are from the NIPA accounts, and are 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 VIII presents these results in the form of standard ordinary least squares 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,

Table VIII
Prepayment Shocks and Aggregate Consumption

We regress:

$$U_t = A + B \times \text{Consumption-Growth},$$

where, U_t are prepayment shocks. C_t is measured household consumption. Results are presented for aggregate (nondurables 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. $N = 18$.

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 derives from a simplified prepayment model, and obtains over a small sample of a boom period during the 1990s. 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 (2005)), prepayments are thereby likely to be 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 grow, the credit quality of the average homeowner improves, thereby allowing more homeowners to take advantage of refinancing options and increasing prepayment rates.

H. Preferences of the Marginal Investor

The coefficient estimates on our model range from 82 to 547, depending on the specification and subsample. These numbers are not readily interpretable as corresponding to preferences. In this subsection, we provide a “back-of-the-envelope” calibration to assess these numbers. We show that our findings are in the range of what one would expect if the marginal investor is a leveraged mortgage fund manager.

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

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

We translate this into preferences for an agent with *CRRA* preferences with parameter $\hat{\rho}$ and wealth of w , that is,

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

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

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

Hence, 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 the fraction α of fund size, so that a fund manager who starts a fund by contributing w_m of his own wealth has fund capital of $\frac{w_m}{\alpha}$. 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

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

where P_C^j is the price of the j^{th} collateral. We can use this expression to solve for w_m , 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. 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 $w = \frac{w_m}{\kappa}$, 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 a factor of 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.¹⁶

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,

¹⁶ Note that 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.

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

For the leverage number, we conduct an informal poll of MBS traders and find that typical leverage ranges from 5 to 20 for funds that trade IOs and POs. 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 make sense under a limits of arbitrage view according to which the marginal investor is a specialized institution. The specialized mortgage fund manager bears disproportionate amounts of mortgage risk. Leverage magnifies this effect, enabling us to match the data with reasonable risk aversion parameters. 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 capital constraints, for example.

Finally, MBS are zero net supply securities. However, if we ignore this fact and suppose that the relevant measure of aggregate household wealth 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—that is, 100 times as high, resulting in numbers in the range of 70 to 500. Reconciling this with consumer preferences would be very problematic. We conclude that to make sense of the prepayment risk premia in MBS, one needs a “limits of arbitrage” view.

III. Model Misspecification and Robustness Checks

As mentioned above, observed OAS behavior might be explained by a misspecification 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 misspecification possibility.

We should note at the outset that one way to sidestep the misspecification issue is to use actual bond returns as the dependent variable in our regressions. The primary reason we do not follow this procedure is that actual bond returns are a very noisy estimate of the expected return on the securities.¹⁷ Thus, we need more data than we have to implement these regressions. Using

¹⁷ We can reduce the noise in bond returns if we take a stand on the mortgage prepayment model and calculate interest rate hedge ratios. Then we can strip out the interest rate component of actual bond returns. However, this seems no better than using the OAS from the prepayment model of a dealer, as we have done.

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 draw attention to Exhibit 74 in particular, which demonstrates that the strongest relation is between the OAS for IOs and subsequent returns.¹⁸

A. Is the OAS Due to a Misspecified Model of the Prepayments?

If the Wall Street pricing model uses an incorrect model of homeowner prepayment, then the OAS will be biased.

Let us revisit equation (7), which we rewrite as

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

Suppose that informed market participants have a true model of prepayments that 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 nonzero even if prepayment risk is not priced.

Note that the OAS on the IOs in our sample are for the most part positive, while those on the PO are negative. Thus, under the misspecified model hypothesis, the OAS must be based on a model that consistently *underpredicts* prepayments.

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

Suppose that in 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 were based on a model calibrated to historical experience, then the OAS would be based on a model that “underpredicts” prepayments.

¹⁸ For each MBS security, Breeden (1994) calculates optimal interest rate hedges based on a prepayment model and then uses these hedge ratios to construct hedged returns on the MBS security. He regresses the realized hedged returns on the OAS of these securities and shows that the regression coefficients are positive. The regression coefficient of the hedged returns on the OAS is: 1.1 (s.e. 0.34), 1.35 (s.e. 0.33), and 0.87 (s.e. 0.58), for collaterals, IOs, and POs, respectively. The result holds true for collaterals, IOs, and POs, and is supportive of our working hypothesis that the OAS is not solely driven by model misspecification.

To the extent that underprediction is a phenomenon that only affects the general level of prepayments, then the time effects in our regression specification should handle the misspecification. For example, an increase in the general level of prepayments may occur if a reduction in refinancing costs leads 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 misspecification in the rate of prepayment across all mortgages.

Alternatively, suppose that some types of mortgage pools, say for geographic reasons, realize a higher level of underprediction. Then the security effects in our specification should handle the misspecification. Thus, our regression specification controls for all underprediction that can be written as

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

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

Suppose that the underprediction at time t is proportional to $\bar{\phi}_t^k$, the current level of prepayments on mortgage k . That is, suppose that $\hat{\phi}_t^k$ is equal to $\bar{\phi}_t^k$ times a constant. This situation may arise, for example, if $\bar{\phi}_t^k$ proxies for the refinancing desires of the average household in a particular pool, so that the lower costs of refinancing speed up prepayments in that pool in particular. Such an underprediction implies an OAS whose sign patterns match those in the data. We run the following regression to account for this possibility:

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

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 IX. 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. This may be due in part to different sample sizes. However, 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 (c^k) and the market interest rate. This situation may occur if the lower costs of refinancing lead 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. Note that if

Table IX
Explaining the OAS, Robustness Checks

Regressions based on the OAS of the IOs:

$$OAS_{IO}^{kt} = \sum_t \alpha_t d_t + \sum_k \gamma_k d_k + A \times \beta^k (\bar{c}_t - r_t) + x_{kt} + \epsilon_{kt},$$

where, β^k is the β -st. dev., and additional explanatory variables x_{kt} are: s_{kt} , the SMM for security k , month t ; \bar{s}_{kt} , half-year moving average of s_{kt} , centered at t ; and $c^k \times r_t$, the 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.

the bias is simply proportional to $c^k - r_t$ then the security/time fixed effects specification will control for this possibility. Thus, the only possibility that we need to address is if the bias depends on both c^k and r_t . The results of Table IX control for this possibility by including a term such that c^k is interacted with r_t . While the interaction term is significant, the coefficient on our model continues to be large and significant.

B. Is the OAS Due to a Misspecified 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 should not reflect interest rate risk. However, any miscalibration or misspecification in the interest rate model will pollute the OAS. In practice, state-of-the-art term structure models produce relatively small pricing errors for bonds. However, when pricing interest rate derivatives, model-implied prices can deviate significantly from market prices. Longstaff, Santa-Clara, and Schwartz (2001) report deviations of 2% to 4% when pricing swaptions using a four-factor string market model. These deviations increased substantially during the fall of 1998 crisis period. Since the MBS is a combination of a bond and interest rate option, the OAS will partly reflect the sort of deviations that Longstaff et al. report.

C. Agnostic Interest Rate Misspecification Controls

On a priori grounds, it is not obvious what specific alternative hypothesis to control for. While there is likely to be a bias in the OAS, neither the direction of

the bias, nor how the bias affects inference in our regressions, is clear. In this subsection we present “agnostic” controls, where the controls are guided by the literature and the tests we perform.

If the bias varies across time and equally affects all of the IOs we study, then the time fixed effects in our regressions will serve as a control. For example, time dummies may control for spikes in the pricing deviations similar to those reported by Longstaff et al. during the fall of 1998.

Andersen and Andreasen (2000) report a skew in pricing deviations for caps and floors. Lower strike options have larger deviations than higher strike options. The strike in an MBS is related to the underlying coupon. Our regressions include security fixed effects that can proxy for the strike and thereby control for this sort of cross-sectional pricing bias.

The type of bias our security/time fixed effects do not control for is a bias that varies over time but affects different strikes differentially. Recall that we test our theory by interacting the prepayment risk of the security and a term involving the level of interest rates ($\bar{c}_t - r_t$). Thus, a candidate for the bias that seems pertinent for our inference is one that depends on the interaction between the moneyness of the option and the level of interest rates.

We can control for this possibility by introducing a regressor that depends on $(c^k - r_t) \times r_t$, where c^k is the coupon on the MBS. This control is subsumed by introducing a regressor that is quadratic in c^k and r_t :

$$A(c^k)^2 + Bc^k r_t + Cc^k + Dr_t + Er_t^2.$$

As our basic regression already has time and security fixed effects, the terms involving only c^k or only r_t are already controlled for. Thus, we only need to include an interaction term between c^k and r_t to control for this hypothesis. The results reported in the last column of Table IX confirm that our results are robust to this type of misspecification possibility.

D. Option Undervaluation Hypothesis

Brown (1999) documents a positive relation between the OAS and implied volatilities on Treasury bond options, which suggests that pricing models may be using an incorrect interest rate volatility. The OAS on IOs are usually positive, while those for POs are usually negative. If the model interest rate volatilities are too low, the prepayment option will be undervalued in the pricing models, giving rise to a positive OAS for the IOs (and a negative one for the POs).

We investigate the option undervaluation hypothesis in Table X. If models 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,

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

Table X
Explaining the OAS, the Option Undervaluation Hypothesis

Regressions based on the OAS of the IOs:

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

where OV_{kt} is the option value, minus intrinsic value, of a 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).

where r_t is the 10-year CMT at time t and c^k is the coupon underlying mortgage k . Thus, 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 (81 bps). Finally, we scale this option value by 100 for ease of comparison.

The above is obviously 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 are being conservative.

The results are in Table X. 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.

IV. Conclusion

We provide theory and evidence that the marginal investor in the mortgage-backed securities market is a fund manager who is principally 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 riskiness of the entire mortgage market. Our empirical findings support the notion of limits to arbitrage in the MBS market, a large market in the United States.

The MBS market is a highly specialized market that requires a great deal of expertise on the part of the active 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 et al. (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 that this evidence also supports limits of arbitrage theories. In line with the theory we present in this paper, we predict that when the standard deviation of returns of the corporate bond market increases, the price of noninterest risk in the bond market will increase. Similarly, when credit risk becomes large, the price of credit risk itself will increase. More work needs to be done to evaluate these conjectures.

There are several avenues along 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. Additionally, while we try to rule out the alternative explanation that our results are driven by a misspecification in the model from which the OAS is derived, this remains a potential alternative. Collecting sufficient data on MBS prices so that we can directly estimate MBS returns instead of relying on the OAS is an important step to take in this line of research.

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 in 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.

REFERENCES

- Allen, Franklin, and Douglas Gale, 1994, Limited market participation and volatility of asset prices, *American Economic Review* 84, 933–955.
- Andersen, Leif, and Jesper Andreasen, 2000, Volatility skews and extensions of the LIBOR market model, *Applied Mathematical Finance* 7, 1–32.
- Bates, David, 2003, Empirical option pricing: A retrospection, *Journal of Econometrics* 116, 387–404.
- Bennett, Paul, Richard Peach, and Stavros Peristiani, 2000, Implied mortgage refinancing thresholds, *Real Estate Economics* 28, 405–434.

- Berndt, Antje, Rohan Douglas, Darrell Duffie, Mark Ferguson, and David Schranz, 2004, Measuring default risk premia from default swap rates and EDFs, Working paper, Stanford University.
- Bollen, Nicholas B., and Robert E. Whaley, 2004, Does net buying pressure affect the shape of implied volatility functions? *Journal of Finance* 59, 711–753.
- 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.
- Breeden, Douglas, 1994, Complexities of hedging mortgages, *Journal of Fixed Income* 4, 6–41.
- Brown, David, 1999, The determinants of expected returns on mortgage-backed securities: An empirical analysis of option-adjusted spreads, *Journal of Fixed Income* 9, 8–18.
- 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.
- Caballero, Ricardo, and Arvind Krishnamurthy, 2002, A dual liquidity model of emerging markets, *American Economic Review* 92, 33–37.
- 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.
- Case, Karl E., John M. Quigley, and Robert J. Shiller, 2005, Comparing wealth effects: The stock market versus the housing market, *Advances in Macroeconomics* 5, 1–32.
- Cohler, Gene, Mark Feldman, and Brian Lancaster, 1997, Price of risk constant (PORC): Going beyond OAS, *The Journal of Fixed Income* 23, 6–15.
- Collin-Dufresne, Pierre, Robert Goldstein, and J. Spencer Martin, 2001, The determinants of credit spread changes, *The Journal of Finance* 66, 2177–2207.
- Dow, James, and Gary Gorton, 1994, Arbitrage chains, *Journal of Finance* 49, 819–849.
- Downing, Chris, Richard Stanton, and Nancy Wallace, 2005, An empirical test of a two-factor mortgage prepayment and valuation model: How much do house prices matter? *Real Estate Economics* 33, 681–710.
- Dunn, Kenneth B., and John J. McConnell, 1981, Valuation of GNMA mortgage-backed securities, *Journal of Finance* 36, 599–616.
- Froot, Kenneth, 2001, The market for catastrophe risk: A clinical examination, *Journal of Financial Economics* 60, 529–571.
- Froot, Kenneth A., and Paul G. O’Connell, 1999, The pricing of US catastrophe reinsurance, in Kenneth Froot ed. *The Financing of Catastrophe Risk* (University of Chicago Press).
- Gabaix, Xavier, Parameswaran Gopikrishnan, Vasiliki Plerou, and H. Eugene Stanley, 2003, A theory of power law distributions in financial market fluctuations, *Nature* 423, 267–270.
- Gabaix, Xavier, Parameswaran Gopikrishnan, Vasiliki Plerou, and H. Eugene Stanley, 2006, Institutional investors and stock market volatility, *Quarterly Journal of Economics* 121, 461–504.
- Garleanu, Nicolae, Lasse Heje Pedersen, and Allen M. Poteshman, 2005, Demand-based option pricing, Working paper, Wharton School.
- 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* (Reading, Mass.: Addison-Wesley).
- Geanakoplos, John, 2003, Liquidity, default, and crashes in Mathias Dewatripont, Lars Peter Hansen, and Stephen J. Turnovsky, eds. *Advances in Economics and Econometrics II* (Cambridge, UK; New York: Cambridge University Press).
- Gromb, Denis, and Dimitri Vayanos, 2002, Equilibrium and welfare in markets with financially constrained arbitrageurs, *Journal of Financial Economics* 66, 361–407.
- Grossman, Sanford J., and Zongquan Zhou, 1996, Equilibrium analysis of portfolio insurance, *Journal of Finance* 51, 1379–1403.
- He, Zhiguo, and Arvind Krishnamurthy, 2006, Intermediation, capital immobility, and asset prices, Working paper, Northwestern University.
- Holmstrom, Bengt, and Jean Tirole, 1997, Financial intermediation, loanable funds, and the real sector, *Quarterly Journal of Economics* 112, 663–691.
- Hurst, Erik, and Frank Stafford, 2004, Home is where the equity is: Liquidity constraints, refinancing and consumption, *Journal of Money, Credit and Banking* 36(6), 985–1014.

- Kau, James B., and Donald C. Keenan, 1995, An overview of the option-theoretic pricing of mortgages, *Journal of Housing Research* 6, 217–244.
- Krishnamurthy, Arvind, and Amette Vissing-Jorgensen, 2006, The demand for Treasury debt, Working paper, Northwestern University.
- Kupiec, Paul, and Adama Kah, 1999, On the origin and interpretation of OAS, *Journal of Fixed Income* 9, 82–92.
- Kyle, Albert S., and Wei Xiong, 2001, Contagion as a wealth effect, *Journal of Finance* 56, 1401–1440.
- Levin, Alexander, and Andrew Davidson, 2005, Prepayment risk- and option-adjusted valuation of MBS, *Journal of Portfolio Management* 31, 73–85.
- Longstaff, Francis A., 2004, Optimal recursive refinancing and the valuation of mortgage-backed Securities, NBER Working paper 10422.
- Longstaff, Francis A., Pedro Santa-Clara, and Eduardo S. Schwartz, 2001, The relative valuation of caps and swaptions: Theory and empirical evidence, *Journal of Finance* 56, 2067–2109.
- Mattey, Joe, and Nancy Wallace, 2001, Housing-price cycles and prepayment rates of U.S. mortgage pools, *The Journal of Real Estate Finance and Economics* 23, 161–184.
- Merton, Robert C., 1974, On the pricing of corporate debt: The risk structure of interest rates, *Journal of Finance* 29, 449–470.
- Merton, Robert C., 1987, A simple model of capital market equilibrium with incomplete information, *Journal of Finance* 42, 483–510.
- Patruno, Gregg N., 1994, Mortgage prepayments: A new model for a new era, *Journal of Fixed Income* 4, 42–56.
- Richard, Scott F., and Richard Roll, 1989, Prepayment on fixed-rate mortgage-backed securities, *Journal of Portfolio Management* 15, 375–392.
- Schwartz, Eduardo S., and Walter N. Torous, 1989, Prepayment and the valuation of mortgage-backed securities, *Journal of Finance* 44, 375–392.
- Shleifer, Andrei, and Robert W. Vishny, 1997, The limits of arbitrage, *Journal of Finance* 36, 35–55.
- Stanton, Richard H., 1995, Rational prepayment and the valuation of mortgage-backed securities, *Review of Financial Studies* 8, 677–708.
- Stein, Jeremy C., 1995, Prices and trading volume in the housing market: A model with down-payment effects, *Quarterly Journal of Economics* 110, 379–406.

