NBER WORKING PAPER SERIES

DO INTERMEDIARIES MATTER FOR AGGREGATE ASSET PRICES?

Valentin Haddad
Tyler Muir

Working Paper 28692
http://www.nber.org/papers/w28692

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2021

We gratefully acknowledge the useful comments and suggestions of Stefan Nagel, an Associate Editor, two anonymous referees, Markus Brunnermeier, Mike Chernov, Itamar Drechsler, Andrea Eisfeldt, Zhiguo He, Arvind Krishnamurthy, Francis Longstaff, Erik Loualiche, Adrien Matray, Alan Moreira, Matthew Plosser, Shrihari Santosh, Matthew Spiegel, Pietro Veronesi, two anonymous referees, and participants at UCLA, the NBER, the Utah Winter Finance Conference, the Wharton Conference on Liquidity and Financial Fragility, the ASU Sonoran Winter Conference, the LBS AQR Asset Management Conference, FIRS, NBER meeting, the Texas Finance Festival, AFA, LSE, NYU, the San Francisco Fed, the New York Federal Reserve, Columbia, Harvard, and Maryland for comments. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Do Intermediaries Matter for Aggregate Asset Prices?
Valentin Haddad and Tyler Muir
NBER Working Paper No. 28692
April 2021
JEL No. G0,G01,G12

ABSTRACT

Poor financial health of intermediaries coincides with low asset prices and high risk premiums. Is this because intermediaries matter for asset prices, or simply because their health correlates with economy-wide risk aversion? In the first case, return predictability should be more pronounced for asset classes in which households are less active. We provide evidence supporting this prediction, suggesting that a quantitatively sizable fraction of risk premium variation in several large asset classes such as credit or MBS is due to intermediaries. Movements in economy-wide risk aversion create the opposite pattern, and we find this channel also matters.

Valentin Haddad
University of California, Los Angeles
Anderson School of Management
Office C4.19
110 Westwood Plaza
Los Angeles, CA 90024
and NBER
vhaddad@ad.ucla.edu

Tyler Muir
University of California at Los Angeles
Anderson School of Management
110 Westwood Plaza
Los Angeles, CA 90024
and NBER
tyler.muir@anderson.ucla.edu

A data appendix is available at http://www.nber.org/data-appendix/w28692

Periods of poor financial health of financial intermediaries such as investment banks, commercial banks or hedge funds coincide with low aggregate asset prices and high risk premia.[1] This correlation suggests the health of the financial sector matters for aggregate asset prices.[2] But, this evidence alone does not rule out the view that intermediaries reflect or are correlated with other frictionless factors driving asset prices. For example, consider the 2008 financial crisis where risk premia rose substantially. While there was indeed a drop in intermediary risk-bearing capacity, household risk aversion likely also rose; hence it is unclear to what extent the fall in intermediation mattered for aggregate asset prices.[3] The goal of this paper is to quantify how much variation in aggregate risk premia can be ascribed to intermediaries rather than to households.

We answer this question by comparing variations in risk premia across more and less intermediated asset classes. We start by regressing the return of each asset class on a proxy for the effective risk-bearing capacity of intermediaries. Specifically, we estimate the predictive regressions

$$r_{i,t+1} = a_i + b_i \gamma_{I,t} + \varepsilon_{i,t+1}, \quad i = 1, ..., N, \quad t = 1, ..., T \tag{1}$$

where $\gamma_{I,t}$ is an empirical proxy for intermediary risk-appetite and $r_{i,t+1}$ is the

[1] E.g., Adrian et al. (2014), Hu et al. (2013), Haddad and Sraer (2020), Muir (2017), He et al. (2017).

[2] He and Krishnamurthy (2013) and Brunnermeier and Sannikov (2014) provide models for how the financial sector can impact asset prices.

[3] Santos and Veronesi (2016) discuss a frictionless model that generates some of the empirical patterns associated with intermediation, leverage, and asset prices.

excess return on asset class i. We then compare the degree of predictability across these asset classes, measured by the R^2 of the predictive regression, the percentage change in risk premium relative to its mean ($b_i/\mathbb{E}[r_i]$), or the change in risk premium relative to the volatility of the asset class ($b_i/\sigma(r_i)$). We find relatively more predictability for asset classes that are more intermediated (e.g., MBS, CDS, currencies, commodities) and relatively less predictability for asset classes that are less intermediated (e.g., stocks). We argue that these differences in the predictability of more versus less intermediated asset classes provide a lower bound for how much intermediaries matter in each asset class.

We clarify our argument in a simple model. When either intermediaries or households are less willing to bear risk, risk premia increase. Because variables that proxy for intermediary risk aversion are likely positively correlated with household or economy-wide risk aversion, the evidence of a predictive relation alone (positive b_i) does not uncover how much variation in risk premia we can ascribe to intermediaries. Comparing across asset classes with different ease of access to households overcomes this challenge. To capture this distinction, we assume households face costs to invest directly in some asset classes relative to investing indirectly through an intermediary.[4] For

[4]Our model is related in spirit to He and Krishnamurthy (2013), Brunnermeier and Sannikov (2014) but includes many assets, and does not assume households are unable to invest directly in the assets (this is equivalent to an infinite cost in our setting). Relatedly, Koijen and Yogo (2015) study how institutional demand affects individual stock prices but their framework does not model the substitution of households' direct versus indirect holdings.

example, this cost is high for CDS and low for stocks. In asset classes with easy direct access, a drop in intermediary risk-bearing capacity does not have a large impact on premia since households can easily substitute or "step in" to these markets. With a high cost of direct investment, households cannot absorb the intermediary positions through direct holdings, so risk premia experience a large increase. Therefore, the predictive relation will be stronger in more intermediated asset classes. Changes in household risk aversion generate the opposite prediction: a more substantial impact on less intermediated assets and a smaller impact on more intermediated assets.

We then make empirical choices to implement our test: we measure a proxy of intermediary risk-bearing capacity ($\gamma_{I,t}$), asset class returns, and the degree to which intermediaries are active in each asset class. To measure risk-bearing capacity, we rely on the existing literature on intermediary asset pricing, which provides foundations for how risk appetite should be measured empirically. Our main specification uses a standardized average of the broker-dealer book leverage variable from Adrian et al. (2014), and the market equity of primary dealers measure from He et al. (2017). We show robustness to alternative proxies for intermediary risk appetite as well. It is worth emphasizing that our exercise does not rely on exactly measuring intermediary risk appetite. We do not model the drivers of intermediary risk-bearing capacity in a micro-founded way as in He and Krishnamurthy (2013) or Adrian and Shin (2014) (i.e., we do not offer a theory of intermediary risk-bearing capacity).[5] We also embrace that variation in our proxy

[5] See also Brunnermeier and Pedersen (2009), Danielsson et al. (2011), Duffie (2010) among many others.

is unlikely to be "exogenous."[6] We overcome these identification challenges by studying a novel dimension of the aggregate data.

We use the following asset classes, ranked from least to most intermediated: stocks, bonds, options, sovereign bonds, commodities, foreign exchange, mortgage-backed securities, and credit (measured using CDS contracts).[7] We use three methods to arrive at this ranking. First, quantity and position data from Flow of Funds and the BIS measure the relative holdings of households and institutions. Second, we gather asset class risk exposure from Value-at-Risk measures taken from intermediary 10-K filings. Third, ETF expense ratios help us gauge households' cost of direct exposure to the asset classes. We also extend our analysis to an alternative set of return series. We compare the predictability of hedge fund returns for strategies of various complexity — convertible-bond arbitrage and fixed-income arbitrage at one extreme, the overall stock market at the other one.[8]

We then compare the degree of predictability across asset classes. The raw predictive coefficient b_i is not an adequate metric, and we need to normalize returns appropriately into the same "units" to draw comparisons. To

[6]In fact, in most theories the health of the financial sector is a state variable that responds endogenously to more fundamental shocks.

[7]Some of our assets are in zero net supply. This is fine as long as the asset return is positively correlated with risk intermediaries are exposed to – e.g., intermediaries on net will be positively exposed to credit risk hence the CDS premium will reflect that credit risk is positive on net. This positive exposure is strongly supported empirically (He et al., 2017). We discuss the issue of asset supply further in the empirical section.

[8]Mitchell et al. (2007) and Hu et al. (2013) provide arguments and evidence as to why intermediary capital matters for the returns of these strategies.

see why, suppose that one asset class is just a levered version of another asset class. Then any variable which predicts returns in the original asset class will mechanically have a larger coefficient on the levered asset class. Dividing the regression coefficient by return volatility ($b_i/\sigma(r_i)$) deals with this effect. It also provides an intuitive interpretation of the coefficient as the degree of predictability relative to the asset's volatility (closely related to the R^2 in the predictive regression). Scaling by unconditional returns ($b_i/\mathbb{E}(r_i)$) also deals with this issue. Our model suggests that this second approach of focusing on the *elasticity* of risk premium to intermediary risk aversion constitutes a better way to capture other dimensions of the differences across asset classes (e.g., differences in unconditional betas). While we prefer this normalization economically, there is an empirical tradeoff because average returns are much harder to estimate than standard deviations. We consider both scalings in our empirical analysis and document a robust pattern of stronger predictability for more intermediated asset classes. An important aspect of statistical inference in the case of elasticities is to account for uncertainty in mean returns estimates, which we divide by. We propose a Bayesian approach to do so and show sharp statistical conclusions in this case under reasonable assumptions. One can either impose that risk premiums of these broad asset classes are not negative or shrink estimates of mean returns toward the assumption of constant unconditional Sharpe ratios across asset classes. We provide economic and empirical arguments in favor of these assumptions.

Our main argument does not rely on any household risk aversion controls, because we need do not need to take a firm stand on the behavior of this quantity. However, having a proxy for household risk aversion allows us to dig deeper and quantify the role of households for prices. In our framework, risk premia respond with an opposite pattern across asset classes to households' willingness to bear risk relative to intermediaries. We confirm this observation in the data, which strengthens our mechanism, by using the *cay* measure from Lettau and Ludvigson (2001) and the habit measure of Campbell and Cochrane (1999). In particular, we find substantially *less* predictability from these measures in more intermediated asset classes. This observation also confirms that our main result is not mechanical: not all return predictors exhibit an increasing pattern as one moves to more intermediated asset classes.

We combine these measured differences in predictability across asset classes to quantify bounds on the role of the two types of investors for risk premia. Our first set of results provides a lower bound on how much intermediaries matter. Predictions with proxies for the risk aversion of households give a lower bound for their role as well. We decompose the variation in risk premium attributable to intermediaries and households for each asset class. For example, we find that we can attribute about 60% of the variation in risk premium in CDS to intermediaries. Similarly, we can attribute about 40% of the variation in risk premium of stocks to households. There is still a remaining fraction of variation for each asset class we cannot assign to either based on our lower bounds.

Finally, we discuss the limitations of the assumptions behind our study and explore other possibilities that could explain our results. Most importantly, our framework thus far considers variation in intermediary and household risk aversions. We allow for (1) arbitrary unobserved time-variation in effective risk aversion (i.e., we do not tie household risk aversion to a specific model but leave it to move around freely), (2) arbitrary unconditional covariances of asset classes with household marginal utility (that is, we do not take a stand on unconditional betas, nor do we tie them to covariance with observables like consumption growth), (3) arbitrary time-varying volatility of the household pricing kernel. However, other factors that drive risk premia may also change. For example, the covariance of asset payoffs might change and be correlated with the other variables. We show that our results are robust to including proxies for changing covariances in the predictive regressions (e.g., time-varying volatilities and betas) or to the possibility that intermediary risk aversion proxies for time-varying loadings on standard risk factors using the framework of Shanken (1990). Further, while we predict differential changes in risk premia across the asset classes, we do not predict differential changes in risk. More broadly, these time-varying covariances would also have to have a unique factor structure to line up perfectly with our results. In particular, it has the be that risk increases more for intermediated asset classes when intermediary risk aversion rises. Across of variety of risk measures, we find no evidence of such a relationship.

Our findings are related to a broader literature studying the link between intermediary balance sheets and asset prices (Adrian et al., 2014; Hu et al., 2013; Haddad and Sraer, 2020; He et al., 2017). Closely related to our work, He et al. (2017) show that an intermediary factor helps explain the cross-section of returns for many asset classes. The main difference is that this literature typically studies intermediary Euler equations, which test for optimality of decisions that link intermediary marginal utility to asset returns, but do not quantify whether intermediaries matter for risk premia. We illustrate this point in our model. Our paper also relates to more "micro" evidence, which provides sharp evidence that intermediaries matter for particular individual asset prices at particular points in time. Du et al. (2017) document stronger violations of covered interest parity at end-of-quarter financial reporting dates, when regulatory constraints are more binding. Some other examples are Siriwardane (2016), Fleckenstein et al. (2014), Lewis et al. (2017), Krishnamurthy (2010), Mitchell et al. (2007).[9] While these studies are important in documenting detailed price deviations related to intermediary risk-bearing capacity in specific periods, it is often unclear what these results imply for the broad behavior of aggregate asset prices.

Section I presents our framework and the model, Section II describes the data, Section III presents the main empirical results. Section IV provides additional analysis, including the results using hedge fund returns. Section V takes stock of our results in the context of the literature on intermediary asset pricing.

I. Separating the Role of Intermediaries and Households for Asset Prices

We introduce our test for the role of intermediaries for risk premia. We first present the basic ideas behind our empirical strategy before moving on to a more formal model. This simple theory guides our empirical implementation, but also helps understand potential limitations to the interpretation

[9] Duffie (2010), Mitchell and Pulvino (2012), and He and Krishnamurthy (forthcoming) offer thorough discussions of this literature.

of the existing evidence on intermediary asset pricing.

A. *The Test*

Economic Question. We are interested in whether the health of the financial sector affects conditional risk premia. The challenge in establishing this link is that intermediaries are not the only ones that can affect risk premia. Households, through changes in how they perceive risk or their risk aversion can also affect risk premia. In terms of a regression equation, this corresponds to the following:

$$\tilde{r}_{i,t+1} = a_i + \beta_{i,H}\gamma_{H,t} + \beta_{i,I}\gamma_{I,t} + \varepsilon_{i,t+1}, \tag{2}$$

where i indexes assets and t time[10] and $\tilde{r}$ denotes the realized excess return on asset i divided on by an asset-specific normalization constant such as the sample average or standard deviation of excess returns. Scaling returns before the regression is identical to regressing unscaled returns and scaling the coefficient and serves to make results comparable across assets. $\gamma_{H,t}$ captures the effective risk aversion of households and $\gamma_{I,t}$ that of intermediaries. We refer to effective risk aversion to generally include anything that effects the willingness to bear risk — for intermediaries this could include losses in net worth, constraints on leverage, and so on (we return to these interpretations later) — and we refer to the negative of effective risk aversion as risk appetite. Naturally, one expects both the coefficients $\beta_{i,H}$ and $\beta_{i,I}$ to be non-negative.

[10] To see that this equation defines risk premia dynamics, take the conditional expectation: $\mathbb{E}_t\left[\tilde{r}_{i,t+1}\right] = a_i + \beta_{i,H}\gamma_{H,t} + \beta_{i,I}\gamma_{I,t}$.

We want to test if $\beta_{i,I}$ is strictly positive and quantify it.

Measurement Challenge. If one can perfectly measure the two risk appetites, then running this simple predictive regression immediately provides estimates of these coefficients. However, this is in general not possible, because risk appetite is imprecisely measured. What is available to econometricians are proxies for these variables, which we denote $\hat{\gamma}_{I,t}$ and $\hat{\gamma}_{H,t}$. It is reasonable to assume we have valid proxies, i.e. that these variables are positively correlated with their actual counterparts. But, because the risk appetite of households and intermediaries are likely positively correlated, it is natural to expect these proxies to also positively correlate with the risk appetite of the other group. In short, we have: $\mathrm{cov}(\hat{\gamma}_{I,t}, \gamma_{I,t}) > 0$ and $\mathrm{cov}(\hat{\gamma}_{I,t}, \gamma_{H,t}) \geq 0$. An implication of these properties is that the reduced-form estimate $b_{i,I}$ in the regression:

$$\tilde{r}_{i,t+1} = a_i + b_{i,I}\hat{\gamma}_{I,t} + \varepsilon_{i,t+1} \tag{3}$$

is positively influenced by both $\beta_{i,H}$ and $\beta_{i,I}$. In other words, measures of intermediary health can forecast returns because the health of intermediaries matter for expected returns ($\beta_{i,I} > 0$) or because it proxies for households' risk appetite and this appetite matters for expected returns ($\beta_{i,H} > 0$). The example of the 2008 financial crisis is useful: while risk premia did spike substantially, and the financial sector was in poor shape, it is also reasonable that aggregate risk aversion increased in the same period. Hence, it is unclear whether the changes in risk premia were due to the collapse in intermediation or not.

Using the Cross-Section. A simple remark allows us to overcome this challenge: the effects of intermediaries and households *across* asset classes should vary in opposite directions. This assumption is intuitively appealing. When an asset is more specialized or more difficult to access directly for households, households play a weaker role for its risk premium (low $\beta_{i,H}$), and intermediaries play a larger role (high $\beta_{i,I}$). The behavior of the estimate $b_{i,I}$ across asset classes combines these opposite patterns. Therefore, if $b_{i,I}$ increases as one moves to more intermediated asset classes, we can conclude that intermediaries affect prices. In addition, the strength of this relation offers a lower bound on their importance. Indeed, if the proxy for intermediary risk appetite captures household risk aversion, this will lead to a smaller slope across asset classes than the actual effect of intermediaries. Internet Appendix Section IA.I.A derives these conclusions formally.

We now turn to a simple model which serves three different purposes. First, it establishes a clear economic motivation behind the structural relation of equation (2). Second, it provides a justification for our assumption of the pattern of predictive coefficients across asset classes. An important ingredient of this comparison is how to appropriately make returns comparable across asset classes – the tilde in our return regressions. Third, we also use the model to understand the limitations to the interpretation of the existing evidence of intermediary asset pricing.

B. An Asset Pricing Model with Intermediaries and Households

1. Setup

There are two periods, 0 and 1, and a representative household. There is a risk-free saving technology with return 1, and n risky assets with supply given by the vector S. Investment decisions are made at date 0 and payoffs are realized at date 1. The payoffs of the risky assets are jointly normally distributed, with mean μ and definite positive variance-covariance matrix Σ. The household has exponential utility with constant absolute risk aversion coefficient γ_H. We write p to denote the vector of equilibrium prices of the assets and assume that all decisions take prices as given.

The household can invest in the assets in two ways. First, the household can buy the assets directly, but at some cost. We assume the household faces a quadratic cost per unit of risk parametrized by the diagonal non-negative matrix C to invest in the various risky assets. This corresponds to a cost $\frac{1}{2}D'\Sigma_{diag}CD$ of investing in a vector D of the risky assets, with Σ_{diag} a matrix containing the diagonal elements of Σ. A simple motivation for this feature is that it is difficult for households to access some risky asset markets, for instance for complex financial products. Existing models of intermediation such as He and Krishnamurthy (2013) typically assume that households cannot invest at all in risky assets, $C = \infty$. A slightly different version is that there is a discretely lower value to risky assets when in the hands of households, for instance in Brunnermeier and Sannikov (2014). Households might also be less able to manage portfolios of risky assets, making them effectively more risky as in Eisfeldt et al. (2017). It might also be that households are only imperfectly informed about the trades that intermediaries do, and therefore do not completely undo changes in their balance sheets through direct trading. More generally, households might have preferences for some asset classes over others for reasons beyond risk and reward. These interpretations highlight that the cost is a stand in for willingness or ability to directly invest in an asset class.

Second, the household can invest through an intermediary which it owns. The intermediary can access markets at no cost, and pass through the payoffs to the household. However, the household cannot completely control the intermediary's investment decisions. We model this distinction by assuming the intermediary invests as if it has exponential utility with risk aversion parameter γ_I. We assume that $\gamma_I \geq \gamma_H$, that is intermediaries are not willing to bear all the risks households want in the first place.[11] In practice, there can be many reasons for why the risk-taking decisions of intermediaries differ from those of households. Managers of financial institutions might have different preferences from their investors and limits to contracting prevent going around this difference. This approach is pursued, for example, in He and Krishnamurthy (2013) and Brunnermeier and Sannikov (2014) (see also He and Krishnamurthy (forthcoming)). Financial institutions also face regulations explicitly limiting their risk-taking. For example the Basel agreements specify limits on risk-weighted capital, measured by pre-specified risk weights or Value-at-Risk. Adrian and Shin (2014) explore this channel.

These two assumptions are voluntarily stylized, and we discuss them in

[11]This assumption is distinct from the assumption in many models that the *relative* risk aversion of intermediaries is lower than that of households. The two assumptions can coexist as long as the intermediary sector is not too large.

more detail in Internet Appendix Section IA.I.D. Figure 1 summarizes this setup.

[Insert Figure 1 about here.]

Because of exponential utility, initial endowments do not affect the demand for risky assets, so we ignore them hereafter. The intermediary problem determining its demand D_I for the risky assets is therefore

$$\max_{D_I} \; D_I' \left(\mu - p\right) - \frac{\gamma_I}{2} D_I' \Sigma D_I. \tag{4}$$

The household takes as given the investment decision of the intermediary when making her choice of direct holding D_H:

$$\max_{D_H} \left(D_H + D_I\right)' \left(\mu - p\right) - \frac{\gamma_H}{2} \left(D_H + D_I\right)' \Sigma \left(D_H + D_I\right) - \frac{1}{2} D_H' \Sigma_{diag} C D_H. \tag{5}$$

An equilibrium of the economy is a set of prices p and demands D_I^* and D_H^* so that the intermediary and household decisions are optimal, and risky asset market clears. The first two conditions are that D_I^* and D_H^* solve problems (4) and (5) respectively. The market-clearing condition is

$$D_H + D_I = S. \tag{6}$$

2. Equilibrium Portfolios and Prices

We now characterize the equilibrium. The intermediary demand follows the classic Markowitz result:

$$D_I^* = \frac{1}{\gamma_I}\Sigma^{-1}\left(\mu - p\right). \tag{7}$$

It invests in the the mean-variance efficient portfolio: the product of the inverse of the variance Σ^{-1}, and the expected returns $(\mu - p)$. The position is more or less aggressive depending on the risk aversion γ_I.

In contrast the household demand is:

$$D_H^* = (\gamma_H \Sigma + \Sigma_{diag} C)^{-1}(\mu - p) - (\gamma_H \Sigma + \Sigma_{diag} C)^{-1}(\gamma_H \Sigma) D_I. \tag{8}$$

The first term of this expression reflects the optimal demand absent any intermediary demand. It balances the expected returns with the quadratic risk and investment costs of buying the assets. The second term represents an adjustment for the fact that the household already owns some assets through the intermediary. Importantly, an asset held through the intermediary does not have the same value as an asset held directly as it avoids the trading costs, and therefore the substitution between direct and intermediated investment is in general not one-to-one.[12] Rather, it is given by

$$-\frac{\partial D_H^*}{\partial D_I} = (\gamma_H \Sigma + \Sigma_{diag} C)^{-1} (\gamma_H \Sigma). \tag{9}$$

The role of the investment cost for this substitution is clear in this expression. Without investment costs, $C = 0$, assets in and out have the same value, so this substitution is the identity. As the investment cost gets larger, the substitution rate converges to 0. If investing directly in the asset is too expensive, the household does not offset the decisions of the intermediary.

We obtain an expression for prices clearing the market by combining the demand from the household and the intermediary:

$$\mu - p = \gamma_H \Sigma \left(\Sigma + \frac{1}{\gamma_I} \Sigma_{diag} C \right)^{-1} \left(\Sigma + \frac{1}{\gamma_H} \Sigma_{diag} C \right) S \tag{10}$$

This relation is the nonlinear counterpart to equation (2), which posited a role for intermediary and household risk appetites, γ_I and γ_H on risk premia. It is interesting to compare these risk premia to those obtained in an economy without any friction. In this case, one would obtain $\mu - p = \gamma_H \Sigma S$. The prices in our economy are distorted relative to this benchmark by a factor $\left(\Sigma + \frac{1}{\gamma_I} \Sigma_{diag} C \right)^{-1} \left(\Sigma + \frac{1}{\gamma_H} \Sigma_{diag} C \right)$. This distortion encodes the potential effect of the intermediary on asset prices, through the impact of the parameter γ_I. The following proposition highlights conditions under which a meaningful

[12] This distinction can also have implications for the price of the financial institutions holding the assets, see for example Garleanu and Pedersen (2011) or Chodorow-Reich et al. (2021).

notion of "intermediary asset pricing" arises.

PROPOSITION 1: *The intermediary matters for asset prices, that is* $\partial(\mu - p)/\partial\gamma_I \neq 0$*, if and only if*

$$\gamma_I \neq \gamma_H \quad \text{and} \quad C \neq 0 \tag{11}$$

The combination of the two frictions of the model is necessary to obtain a role for intermediaries. The first condition captures the idea that, at least in part, intermediary decisions must not exactly reflect the desires of the household. In our simple model, this discrepancy is captured by a distinct investment goal, $\gamma_I \neq \gamma_H$. But this condition is not sufficient for intermediaries to matter. It must also be that households are limited in their ability to reach their investment objectives on their own. Our model materializes this limitation by a non-zero investment cost C. More generally, the key feature of investment policies to obtain this limitation is that households do not exactly offset decisions of intermediaries, $-\partial D_H^*/\partial D_I \neq I$.[13]

Now that we have clarified the importance of our two frictions for the notion of intermediary asset pricing, we derive empirical implications of this framework, including the assumption behind the test of Section I.A.

C. Empirical Implications

We have already discussed how linearizing the model leads to a structure like equation (2): higher values of either intermediary or household risk aver-

[13]See Internet Appendix Section IA.I.C for a more general discussion.

sion yield higher risk premia. The following proposition goes one step further and shows that the pattern of predictability across asset classes arises in our model, and therefore validates our test. To simplify, we assume that Σ is diagonal. We index assets by i and denote c_i the elements of C.

PROPOSITION 2: *The elasticity of risk premium to intermediary risk aversion γ_I is increasing in the cost of direct holding c_i, strictly if the intermediary matters for asset prices. The elasticity to household risk aversion γ_H is decreasing in the cost of direct holding.*

To understand this proposition, consider the elasticity of the risk premium to changes in risk aversion:

$$\beta_{i,I} = \frac{1}{\mu_i - p_i}\frac{\partial(\mu_i - p_i)}{\partial \log(\gamma_I)} = \frac{c_i}{\gamma_I + c_i}, \tag{12}$$

$$\beta_{i,H} = \frac{1}{\mu_i - p_i}\frac{\partial(\mu_i - p_i)}{\partial \log(\gamma_H)} = \frac{\gamma_H}{\gamma_H + c_i}. \tag{13}$$

Both of these elasticities are positive, with a role for intermediary risk aversion if and only if there is a non-zero cost of direct investment $c_i > 0$. However, the elasticity is increasing in the cost c_i for intermediary risk aversion while it is decreasing for household risk aversion, and flat if there are no frictions (e.g., $C = 0$). It is increasing for intermediaries because households offset their trades less in asset classes that are harder to invest in directly. In contrast, the opposite is true for changes in household risk aversion. Figure 2 illustrates this comparison.

[Insert Figure 2 about here.]

Focusing on elasticities rather than simply the derivative of the risk premium with respect to the risk appetite quantities is a useful scaling. Indeed, assets in higher supply or with higher risk have higher risk premium, and therefore will tend to move more in absolute magnitude with risk appetite. Scaling by a baseline level of risk premium cleans out this effect to focus on the role of the financial frictions. Internet Appendix Section IA.I.B discusses a few mechanical properties of this elasticity.

This empirical implication contrasts with other work on intermediaries and predictability which typically focuses on a single asset class or does not explicitly consider *relative* predictability (e.g., Haddad and Sraer (2020), Diep et al. (Forthcoming), Chen et al. (2019), He et al. (2017), and Muir (2017), among others). It also differs from the Euler equation approach of linking intermediaries marginal value of wealth to the cross-section of risk premiums (e.g., Adrian et al. (2014) and He et al. (2017)). We discuss this issue in more depth in Section V, but note in our model that the intermediary Euler equation always holds, because they hold the mean-variance efficient portfolio — equation (7). This holds true regardless of whether intermediaries matter.

II. Data and Empirical Approach

We now turn to empirical measurement of asset returns, cost rankings across asset classes in terms of more vs less intermediated assets, and measurement of proxies for intermediary health.

A. Returns

We use asset returns and intermediary state variables that are common in the literature. We use excess returns on the market, commodities, credit (CDS), options, sovereign bonds, Treasury bonds, the currency carry trade, and MBS, where we take excess returns over the 3 month T-bill where appropriate. These choices are motivated by looking at many large markets where we think intermediation may matter. We start by using these asset returns provided by He et al. (2017), and refer to that paper for a thorough description of these series. For CDS, options, sovereigns, and commodities we take the equal weighted average in each asset class. Treasury bonds are longer term Treasury bond returns over the 3 month T-bill rate. The credit return is an average across maturities and credit risk. MBS is the Barclay's hedged MBS return index.[14] We use the hedged return to remove exposure to interest rate risk, just as CDS isolates credit risk exposure. Commodities are the equal weighted average across all commodities available in the HKM dataset. The carry trade data are from Adrien Verdelhan. Some of the assets are in zero net supply. This is not an issue as long as the asset return is positively correlated with risk intermediaries are exposed to (e.g., intermediaries on net will be positively exposed to credit risk hence the CDS premium will reflect that credit risk is positive on net). In general, our assumption is that the intermediary sector has positive exposure to the asset returns in question such that if their effective risk aversion increases they will be less willing to bear

[14] We thank Peter Diep for help with this data. This series is also available on Bloomberg LUMSER.

this risk unless the premium also rises. This positive exposure is strongly supported empirically because betas for these assets classes with respect to the intermediary sector are positive and align with their risk premiums (He et al., 2017).[15] For credit, we favor CDSs over corporate bonds because they provide a pure exposure to credit risk, and contribute to more variation in intermediation cost across asset classes.[16] Table I gives summary statistics for the asset class returns including means, standard deviations, and Sharpe ratios. All units are quarterly.

[Insert Table I about here.]

B. Intermediary Health

Next, we use variables in the literature that are argued to proxy for intermediary distress or risk-bearing capacity. That is, we want variables that we believe are correlated with γ_I in our framework. We use two primary measures; the broker-dealer leverage factor from Adrian et al. (2014) (AEM) and the intermediary equity measure by He et al. (2017) (HKM), each of which has been argued, theoretically and empirically, to capture intermediary

[15] One can also accommodate a fixed demand from outside investors in our model to generate the supply that the intermediaries are exposed to. That is, suppose some investors want to hedge oil prices or some other risk. Then this demand creates effectively positive supply. Thus, even though there is zero net supply the intermediaries' risk exposure is positive.

[16] Internet Appendix Section IA.V repeats the analysis for corporate bonds and we return to this comparison later in the text.

distress and each of which is linked to risk premiums. We take annual log changes of each state variable. In our main results, we standardize each of the AEM and HKM measures and take the average, so as to take the average of the risk bearing capacity measures used in the literature. We refer to this measure as risk bearing capacity, and the negative of this measure to be effective risk aversion. Again, we emphasize that we do not provide a deep theory for what determines intermediary distress or risk bearing capacity, though these variables are motivated in such a way elsewhere. Our goal instead is to take off-the-shelf measures from the literature to test our main hypothesis.

Finally, we also include variables we think may capture aggregate or household risk aversion, such as the consumption-wealth ratio proxy of Lettau and Ludvigson (2001) and the habit measure (surplus consumption) of Campbell and Cochrane (1999). Specifically, Lettau and Ludvigson (2001) construct the variable *cay* using a combination of aggregate consumption, labor income, and asset wealth. This last component in particular relies on valuations, and therefore would naturally also capture aggregate fluctuations in the willingness to invest in risky assets. We do not take a strong stand on these variables in terms of corresponding perfectly to household risk aversion, though we consider whether including them in our regressions affects our results. This is useful because our theory does have a differential prediction about how household risk aversion shocks should interact with risk premia so this provides a nice additional test of the model.

Figure 3 plots the time-series behavior of the intermediary effective risk aversion proxy (green line), as well as the effective household risk aversion series from the surplus consumption measure (red line). According to these measures, periods of high intermediary risk aversion often coincide with periods of high household risk aversion, the financial crisis of 2008 being a striking example. The two series also exhibit some amount of independent variation which we will exploit later on in the paper.

[Insert Figure 3 about here.]

C. Ranking of Assets by Degree of Intermediation

Our empirical tests require us to rank assets by the willingness of households to hold them. Dispersion in this dimension is important for our empirical design because we exploit that assets that are more specialized (held by intermediaries) will respond more to intermediary health. We take a multi-pronged approach to identifying which assets are more intermediated: we look at holdings data, volume of trade accounted for by institutions (particularly focusing on dealers), and we also consider directly the costs faced by households (we use the fees charged by ETFs by asset class and we also discuss other physical costs households would face in each market).

Importantly, all of these approaches yield roughly similar rankings of which asset classes are more or less intermediated. We report our ranking in Table II. Stocks always appear least intermediated. On the other extreme, credit default swaps appear most intermediated. This makes sense: one needs an ISDA master swap agreement to trade CDS which a household would find close to impossible. The remainder of the ordering, from less to more intermediated, is roughly government bonds, options, sovereign bonds (emerging market), commodities, foreign exchange, MBS, and CDS. We emphasize that we take a data-driven approach to conduct these rankings, though we do not take an overly strong stance on the exact ordering (e.g., one could swap some of the adjacent pairs) and we return to this issue in our empirical tests.

[**Insert Table II about here.**]

1. Holdings and volume data

By revealed preferences, relative holdings of assets directly by households and through intermediaries offer a proxy for the cost of intermediation. For example, in our simple model, we obtain $\frac{D_{I,i}}{D_{H,i}} = \frac{c_i+\gamma_H}{\gamma_I-\gamma_H}$, which is increasing in the cost c_i. We first study holdings data in 2016 from the Flow of Funds (FoF), though we find similar results using the Survey of Consumer Finance (SCF).[17] In FoF we take holdings of stocks, Treasuries, foreign and corporate bonds, and mortgage-backed securities as a percentage of total assets for households and non-profits (HH) as well as for broker-dealers and commercial banks. We compare relative fractions of each of these asset classes, that

[17]The SCF data gives an alternative way to measure households, and provides some advantages. First, we can focus on higher income households which participate more actively in asset markets. Second, FoF lumps households with non-profits, some of which have significant assets, which SCF does not do. However, SCF is a survey, so is subject to other issues.

is the ratio of HH holdings of stocks relative to either broker dealers or banks and likewise for the other assets. Households hold far more equities relative to intermediaries, while households hold fewer Treasuries and far fewer corporate and foreign bonds and MBS. However, this does not give us holdings of more specialized asset classes such as CDS. A limitation of this exercise is that FoF computes household holdings as a residual from other sectors. In particular, this measure includes the holdings of hedge funds, which fit more as intermediaries in our approach.

Our next source of data is the BIS data on derivatives semiannual report.[18] We use data from the end of 2016. The data provide total gross notional positions in each market, the total gross positions by reporting dealers, other financial institutions, and non-financial institutions. We use the sum of reporting dealers and other financial institutions relative to totals; we obtain similar results when using reporting dealers as a fraction of total. These positions are available for commodities, CDS, foreign exchange, and equity derivatives which we use to proxy for equity index options in our sample. Our ranking suggests equity options, commodities, foreign exchange, and CDS as least to most intermediated.

2. Value-at-risk data

One issue with the previous rankings is that they may not capture true "exposure" to the various asset classes, which is what the theory dictates. For example, if households held very low risk stocks, and intermediaries held

[18] See https://www.bis.org/statistics/derstats.htm.

very high risk or high beta stocks, perhaps the fractions above would miss this. In the model, we want relative wealth betas to each asset class which we call exposures.

For intermediaries, we can get a window into exposures by looking at large primary dealers who report value-at-risk across four asset classes on their annual 10Ks. We have this data for the largest dealer banks and we use data from 2016 for our analysis.[19] The 10Ks report value-at-risk for commodities, equities, interest rates, and foreign exchange, giving us the effective relative dollar exposures to each asset class. Value-at-risk reports tail risk – for example, it provides a dollar amount which losses would not be expected to exceed 99% of the time. We convert this number to exposures by assuming a normal distribution for each asset class and normalizing by the standard deviation of the asset class returns from our sample. This gives us relative exposure. We then normalize each asset class by a measure of total supply. For equities and bonds we use the relative sizes of the equity and fixed income markets in the US, roughly 15 trillion and 50 trillion respectively. For bonds our numbers are unchanged if we use only US Treasuries outstanding. For commodities and FX we use the gross market value numbers from the BIS to normalize exposures. We again find consistent results: relative to the sizes of the markets, dealer exposures are smallest for equities, then bonds, then commodities, and then FX. The absolute exposures are largest for fixed income, but importantly this market is quite large and much of this risk is born by other investors so that it is not as large relative to total quantity of

[19]Internet Appendix Section IA.II lists these banks.

bonds outstanding. This ranking thus gives similar results to just using the position data above.

3. Direct Measures of Costs

We next study household ease of access to asset classes by analyzing fees for ETFs in terms of expense ratios from ETF database.[20] While these products would not have been available to households over much of our sample, looking at ETF expense ratios helps gauge the current cost of households investing in these asset classes, and it is likely that they reflect the historical difficulty of investing as well. Another caveat is that this takes the physical cost of accessing asset classes more literally. In reality households may not invest in some assets due to complexity or other features not captured by physical costs.

We take the average expense ratio by asset class: Stocks, Government Bonds, Emerging Markets Bonds (our best proxy for the sovereign bonds), Currency, Commodities, Volatility, and MBS. We use volatility to proxy for our option straddle strategy which is a bet on volatility.[21] There is no category for CDS, since there are very few ETFs trading CDS, though we supplement this by studying two ETFs that specialize in CDS.[22]

[20] http://etfdb.com/etfdb-categories/

[21] We note much of the Volatility ETFs are trading VIX futures directly, and these strategies are different though they are exposed to the same underlying risks (Dew-Becker et al., 2017).

[22] ProShares North American high yield CDS. ProShares offers both a long and short ETF for this product (e.g., you can effectively buy or sell protection). These were launched

We need to normalize the expense ratios in each asset class, just like we do in our simple model. For example, government bond ETFs are safer, lower return funds and hence a high expense ratio here means post fee returns are likely to be particularly low. This is less critical for equity ETFs. We choose to normalize expense ratios by standard deviation of returns in each asset class. Another option is to normalize by the mean return in each asset class, this gives similar results though is subject to the issue that means are much less precisely estimated than standard deviations.

Our approach based on ETF data implies the following ordering, from easier to hardest to access: stocks, Treasuries, sovereign bonds, currencies, commodities, options, MBS, and CDS. This is largely consistent with our main ranking.

4. Other Differences Across Asset Classes

One should not forget that, while we focus on heterogeneity in ease of access by households, these asset classes differ on other dimensions. For example some of the more intermediated asset classes such as MBS and CDS are less liquid than the stock market or currency markets, with for example larger transaction costs. To the extent that liquidity risk is reflected in expected return, this heterogeneity would manifest itself in the return data. Intermediaries would still play a central role in this mechanism, but for different reasons than in our theory. More broadly, the various asset classes also trade different sources of economic risk.

in 2014 as the first CDS ETFs.

The presence of heterogeneity across asset classes beyond the cost of direct access is not a challenge for our empirical strategy in itself. These other forces could confound our results only if they change disproportionately for the harder-to-access asset classes in times of high intermediary risk aversion. For example it would have to be that not only liquidity risk is more pronounced for the more intermediated asset classes, but also that it increases more in periods of high intermediary risk aversion for these asset classes. We empirically address the possibility of such patterns in Section IV.A.

III. Empirical Results

A. *Intermediary Health Forecasts Returns*

1. Methodology

We estimate the following linear equation using quarterly data for each asset class:

$$r^{\sigma}_{i,t+1} = a_i + b_i \times \tilde{\gamma}_{I,t} + \epsilon_{i,t+1}. \tag{14}$$

$\tilde{\gamma}_{I,t}$ is our standardized measure of intermediary risk aversion. $r^{\sigma}_{i,t+1}$ is the excess return of asset class i between quarter t and $t+1$, divided by its full-sample volatility. As we have discussed, this scaling makes predictive coefficients comparable across asset classes. We would ideally prefer to normalize returns by their unconditional risk premium. However, this approach poses additional statistical challenges, which we come back to in the next section. Notice that while scaling returns by a constant affects the magnitude of coefficient estimates, it does not influence t-stats and R^2s.

We use the reverse regression approach of Hodrick (1992) to compute standard errors for our coefficient estimates. Additionally, we account for potential small sample bias, such as the Stambaugh (1999) bias, by computing p-values for the predictive coefficient from a parametric bootstrap procedure. Precisely, we first estimate a restricted VAR for quarterly excess returns and intermediary health under the null of no return predictability by intermediary health .[23] We assume the joint distribution of innovations in the VAR corresponds to their empirical distribution. Then, we draw 5,000 samples from this estimated process to obtain a distribution of reverse regression t-statistics. We report the p-value of our estimated t-statistic relative to this bootstrapped distribution. Both the asymptotic standard error and the p-value are informative: the asymptotic standard error is robust to the specifics of the data-generating process, while the p-value handles finite-sample issues conditional on a parameterized data-generating process.[24]

2. Main Predictive Regressions

Table III gives our predictability results. First, we note that intermediary effective risk aversion generally positively predicts risk premiums across these asset classes. When intermediary health is poor, and their effective risk

[23]When we add additional controls to the regression, such as in Internet Appendix Table IAXV, we allow these other variables to predict returns in the VAR estimation.

[24]We report in Internet Appendix Table IAIII estimates of equation (14) using Newey-West standard errors allowing for eight quarter lags and show statistical significance is generally stronger with this approach. However, this procedure has been found to over-reject the null hypothesis in small samples (see, e.g., Ang and Bekaert (2006)).

aversion is high, risk premia going forward are generally higher.[25] Since we normalize by asset class volatility, and because our predictor variable is standardized to have unit variance, the coefficients indicate the increase – in Sharpe ratio units – to a one standard deviation increase in the intermediary risk aversion measure. Sharpe ratios for the asset classes are fairly similar and typically around 0.25 quarterly. Therefore, the typical coefficient of 0.2 implies a bit less than a doubling of Sharpe ratios, which is economically large.

[Insert Table III about here.]

Second, the degree of predictability, measured either by the coefficient, significance level, or (adjusted) R^2 generally increases as we go from left to right, that is, as we go from less to more intermediated asset classes. There is relatively less predictive power for stocks (coefficient of 0.12, R^2 of 0.8%), and relatively more predictive power for more intermediated asset classes such as MBS (coefficient 0.30, R^2 of 7.8%) or credit (coefficient of 0.57, R-square of 31.6%). Of the eight asset classes, six of the slope coefficients are significant at conventional levels, and all of these six are the more intermediated asset classes. Further, the magnitude of the coefficients of all these six asset classes are larger than that of stocks or Treasuries. Figure 4 shows the increasing pattern of slope coefficients b_i in a scatter plot where the x-axis is our ranking of less vs more intermediated asset classes. Panel B of Figure 5 illustrates a similar increasing pattern for the R^2 of the regression. This is not surprising:

[25] He et al. (2017) also provide evidence of common predictability across asset classes.

because the volatility of returns is one with our normalization, and the intermediation measure is standardized, the (unadjusted) R-square is the square of the predictive coefficient.[26]

[Insert Figure 4 about here.]

The last row of Table III reports the elasticity of the risk premium to a one-standard deviation increase in intermediary risk aversion, given by $b_i/\mathbb{E}(r_i^\sigma)$ — Panel A of Figure 5 plots the elasticities.[27] Here again, we observe a strong increase of elasticities as we move from less to more intermediated assets. We revisit this pattern of increasing predictability and its statistical properties after discussing robustness of the predictive regressions.

[Insert Figure 5 about here.]

3. Robustness of Predictability

We assess the robustness of the predictive regressions on several dimensions. To save space, we describe these results in the text but leave the tables to the Internet Appendix. First, one concern is whether our particular sample drives the results. We explore this by dropping the financial

[26]Formally, $R^2 = var(b_i\tilde{\gamma}_{I,t})/var(r^\sigma_{i,t+1}) = b_i^2$.

[27]Note that it is irrelevant that we first normalize returns by their volatility when computing this elasticity. That is, the elasticity can be equivalently measured as the coefficient b_i^r in $r_{i,t+1}/E[r_{i,t+1}] = a_i^r + b_i^r \times \tilde{\gamma}_{I,t} + \epsilon_{i,t+1}$ where $r_{i,t+1}$ is the raw excess return.

crisis (2007-2009) and find similar patterns (Internet Appendix Table IAIV). This alleviates the concern that this one period of more severe intermediary distress drives our findings. We also show results using only data after 1990 (Internet Appendix Table IAV). Doing so reduces the heterogeneity in sample length across our asset classes. Finally, we also ask whether the predictability changes during periods of low intermediary health by estimating different coefficients for values of $\tilde{\gamma}_I$ above and below its mean (Internet Appendix Table IAVI). Splitting the sample in half results in limited statistical power and no firm result emerges across periods of high or low intermediary health.

Next, another concern relates to our choice of intermediary health measures. Our baseline uses an average of the HKM and AEM intermediary factors in our main results without taking a strong stand on the drivers, and measurement, of intermediary health. We show that several other ways of measuring intermediary health lead to similar conclusions in the pattern of coefficients. Internet Appendix Figures IA3 and IA4 show results when we split our intermediary health measure into the HKM and AEM components separately — Internet Appendix Tables IAVII and IAVIII report the regressions. We find that, generally, both measures contribute to our main result though results are stronger for the AEM measure. Internet Appendix Figure IA5 uses the log levels of the AEM and HKM factors rather than annual log changes (we still average the two after standardizing them). In Internet Appendix Figure IA2 we show the pattern of predictability when we use the GZ spread of Gilchrist and Zakrajek (2012) to proxy for intermediary risk aversion instead of the AEM or HKM measures. Gilchrist and Zakrajek (2012) argue that this spread captures the health of the financial sector and show it closely follows dealer CDS spreads in their sample.

Finally, in Internet Appendix Section IA.V, we discuss corporate bond returns as a an alternative measure of credit returns to CDS. We find strong return predictability for corporate bonds, albeit somewhat lower in magnitude than for CDS — see Internet Appendix Table IAI. This makes sense in part because corporate bonds contain both credit and duration risk.

B. Do Intermediaries Matter?

1. Interpretation of the Pattern across Asset Classes

Figures 4 and 5 consistently show an increasing pattern of predictability: more intermediated asset classes are more predictable by intermediary health. Fitting a regression through these estimates — the red line — implies that the most intermediated asset classes have a predictive coefficient about 0.25 greater than the least intermediated asset classes.[28] In Internet Appendix Table IAII we show that capturing the degree of predictability across asset classes in a panel regression with an interaction term for more intermediated asset classes gives the same estimate for this slope. Elasticities experience an increase of a comparable order of magnitude (remember that the typical normalized return mean is about 0.25, so elasticities should be about 4 times larger than the predictive coefficients). Finally, it could be that this increasing pattern is driven by a single asset class with an extreme value. Internet Appendix Figure IA6 shows this is not the case: the result

[28]This slope across asset classes is the coefficient B in a regression as $b_i = A + B \times c_i + u_i$, where c_i grows linearly from 0 to 1.

holds even if we remove any single asset class.

This consistent increase in predictability is our main empirical result. This observation supports the view that intermediaries matter for risk premia, in particular for the most intermediated asset classes. If the predictability we measure reflected changes in household risk aversion, we would observe instead less predictability for more intermediated asset class. Importantly, our estimates do not rule out that intermediary health partially proxies for household risk aversion and that this variation also matters for asset prices. Rather, they tell us that the increasing effect of intermediaries on risk premia as we move from least to more intermediated asset classes dominates the decreasing effect from households. In this way, we are not only rejecting the null hypothesis that intermediaries do not matter, but offering a quantitative lower bound on their effect. It is tempting to go one step further and use equation (12) to recover intermediation costs, but we actually only recover bounds on the elasticity to actual intermediary risk appetite up to a multiplicative constant. This limitation arises because we only know that our proxy for intermediary health is positively related to the actual intermediary risk appetite γ_I, but not the strength of this relation.

2. Statistical Properties of our Test: a Bayesian Approach

Naturally, this bound comes with standard errors. We turn to assessing the statistical properties of this increasing pattern of predictability. We ask with which degree of statistical confidence we can conclude that more intermediated asset classes are more predictable by intermediary health. The main issue we deal with is uncertainty associated with estimating the mean returns.

We take a Bayesian approach to this question; Internet Appendix Section IA.III provides technical details. We start by the counterpart to equation (14) where excess returns are not scaled:

$$r_{i,t+1} = a_i + b_i \times \tilde{\gamma}_{I,t} + \epsilon_{i,t+1}. \tag{15}$$

We choose a prior on the set of coefficients $\{(a_i, b_i)\}_{i=1,\ldots,N}$ in this regression, and consider various properties of their posterior distribution given our sample. Specifically, the joint distribution of these coefficients implies a joint distributions of the elasticities b_i/a_i.[29] We ask what is the posterior probability that the slope across elasticities – the red line in our figures – is positive, or that the average elasticity for more intermediated asset classes is large than for less intermediated asset classes. For this second type of comparison, we consider the case of Stocks and Treasuries relative to CDS and MBS, Stocks and Treasuries relative to all other asset classes, and Stocks, Treasuries and Options relative to all other asset classes.

We assume that the errors $\epsilon_{i,t+1}$ are normally distributed, uncorrelated over time and have a known cross-sectional variance-covariance matrix Σ_ϵ given by the variance-covariance matrix of OLS residuals.[30] These assump-

[29]We demean $\tilde{\gamma}_{I,t}$ for the sample of each asset class so that the coefficient a_i estimates the unconditional mean return. Demeaning the right-hand-side of the regression does not affect statistical inference in the Bayesian approach below, as well as in the frequentist view; therefore doing so is without loss of generality.

[30]In Internet Appendix Section IA.III.C, we show that uncertainty about this variance-covariance matrix does not play a quantitatively meaningful role in our exercise.

tions correspond to the Seemingly Unrelated Regression framework of Zellner (1962). We focus on truncated multivariate normal priors, which are conjugate with our assumptions on residuals. For the vectors $a = (a_1, ..., a_N)$ and $b = (b_1, ..., b_N)$, we assume means of $\bar{a}\sqrt{diag(\Sigma_\epsilon)}$ and $\bar{b}\sqrt{diag(\Sigma_\epsilon)}$, and variances $\sigma_a^2\Sigma_\epsilon$ and $\sigma_b^2\Sigma_\epsilon$, where $\bar{a}$, $\bar{b}$, σ_a and σ_b are scalar. Finally, we assume that the vectors a and b are independent from each other. Intuitively this prior is expressed in units of Sharpe ratio.[31] For the unconditional mean a, Pástor (2000), Pástor and Stambaugh (2000) and Kozak et al. (2020) show that this choice of prior — scaled by the variance of realized returns — ensures reasonable properties of returns.[32] The arguments in Haddad et al. (2020) show that these ideas extend to the case of predictability, justifying the use of this prior for b as well. In the absence of truncation, and when σ_a and σ_b are large, the Bayesian posterior has mean and variance which correspond to the point estimate and standard errors of the frequentist approach. Throughout, we assume the prior is centered around a quarterly Sharpe ratio of 0.25 (corresponding to 0.5 annually) and no predictability: $\bar{b} = 0$, $\bar{a} = 0.25$. We choose a loose $\sigma_b = 1$ and impose no truncation for b_i. These assumptions are mild and imply that the answer to our questions on the pattern of elasticities is always 50% for the prior: we are never building in an increasing pattern of predictability through the choice of prior. We now turn to two sets of more meaningful assumptions for statistical inference.

First, we impose a positive lower bound on the a_i. The computation of

[31] A small distinction with our previous analysis is that these scaling are in terms of the average conditional Sharpe ratio rather than the unconditional Sharpe ratio.

[32] For example, the expected Sharpe ratio will be bounded above irrespective of Σ_ϵ.

elasticities involves dividing by the unconditional mean a_i. If the distribution of a_i goes through 0, this will yield arbitrarily large positive and negative values of the elasticity.[33] It therefore appears necessary to bound a_i away from 0 to obtain reasonable statistical properties. We do so by imposing a lower bound $\underline{a}\sqrt{diag(\Sigma_\epsilon)}$, with $\underline{a}$ a scalar, for these coefficients. This assumption corresponds to a lower bound $\underline{a}$ on the average Sharpe ratio for each of the asset classes. Naturally, it is important to ask what such an assumption means economically. There are three ways to interpret this bound. First, it is an economic restriction: broad asset classes are all likely to receive a positive risk premium because they capture aggregate risk.[34] Second, it is an empirically motivated assumption: across many samples, an extensive literature has documented positive premia for these asset classes. We discuss this evidence in Internet Appendix Section IA.IV. Third, the doubtful reader can also stress that our conclusions are part of a joint hypothesis framework: we draw economic conclusions under the assumption that unconditional Sharpe ratios have a lower bound.

Figure 6 reports the answer to our questions as a function of the truncation level for the quarterly Sharpe ratio. The thick black line is the posterior probability that the slope across elasticities is negative; lower values favor our theory, in the spirit of a p-value for the null of no pattern. Without getting

[33]Besides, if the distribution of a_i is non-zero and continuous at the point 0, the posterior mean of the elasticity does not exist.

[34]Campbell and Thompson (2008) argues for this type of restrictions. Our approach is less stringent: they bound conditional expected returns at each point in time, while we only impose a bound on unconditional expected returns.

rid of negative Sharpe ratios, in particular the values close to 0, the inference about this slope is rather imprecise, with probabilities around 11%. However, as soon as one imposes a positive bound on the Sharpe ratio, the probability drops sharply, with values under the 5% threshold. For example, a reasonable lower bound of 0.05 give posterior probabilities of around 3%. The pattern for our other three questions is similar. Comparing the extremes of Stocks and Treasuries with Credit and MBS, the solid black line, yields even tighter conclusions, while comparing Stocks, Treasuries and Options with the rest of the sample, the dashed line, gives slightly higher probabilities. The only noticeably large values, while still around 10%, occur when comparing Stocks and Treasuries with the remainder of asset classes. In summary, we conclude with a high degree of confidence that more intermediated asset classes are more predictable by intermediary health.

[Insert Figure 6 about here.]

One might want to go further in terms of statistical regularization. Even abstracting from the issue of dividing by 0, uncertainty about unconditional means contributes to the uncertainty about patterns of predictability. This phenomenon is amplified by the fact that we estimate means for multiple asset classes. To regularize our estimates, a natural approach is to impose some form of shrinkage towards a common value. We follow this approach and shrink Sharpe ratios of all of our asset classes towards a common value.[35]

[35]Pástor (2000) and Pástor and Stambaugh (2000) follow this approach. Kozak et al. (2020) assume that Sharpe ratios are proportional to variances in order to impose the

In our Bayesian framework, this corresponds to tightening the priors on average Sharpe ratios around a common value, our mean $\bar{a} = 0.25$. Why 0.25? This number corresponds to the well-known estimate of an annual Sharpe ratio of 0.5 for the equity market, is roughly the average Sharpe ratio across asset classes in our sample, and is justified by an extensive body of empirical work studying long historical samples for these asset classes. For example Gorton and Rouwenhorst (2006) find that commodities have Sharpe ratio similar to that of stocks, or Asvanunt and Richardson (2016) stress the comparable Sharpe ratios of stocks, Treasuries, and credit; see Internet Appendix Section IA.IV for more discussion. We also note this specific choice of value for $\bar{a}$ is not critical for our results. Shrinking all the way to any common value corresponds to comparing directly predictive coefficients scaled by volatility $b_i/\sqrt{\Sigma_{\epsilon,ii}}$ in terms of the pattern of predictability across asset classes.[36] This extreme provides an economic interpretation to the "model-free" comparison of coefficients we were reporting in the previous section: we are comparing elasticities under the assumption that all asset classes have the same unconditional Sharpe ratio.

[Insert Figure 7 about here.]

absence of near-arbitrage across the thousands of available stocks. Our setting is different: we focus on a small set of asset class index returns, which could all have sizable Sharpe ratios.

[36]Relative to our earlier results, the shrunken estimates just multiply all predictive coefficients by the constant $1/\bar{a}$ so that all coefficients are four times larger, but does not change the relative pattern across asset classes.

Figure 7 illustrates what happens when we implement the shrinkage by bringing σ_a towards 0. Panel A reports the median as well as 5th, 10th, 90th, and 95th percentiles for the posterior of the cross-sectional slope of elasticities. First, one can observe the phenomenon of shrinkage. For large values of σ_a, the estimate corresponds to what one obtains using the OLS estimates of b_i and a_i. As σ_a goes to 0, the median estimate is brought down to a lower value, which corresponds to the slope when replacing a_i, by $0.25\sqrt{\Sigma_{\epsilon,ii}}$. In this case we estimate a slope of around 1, which is about four times larger than the slope in Figure 4 since we divide coefficients by $\bar{a} = 0.25$. Second, this approach brings regularization because shrinkage reduces the uncertainty caused by estimation of mean returns. The distribution of the slope estimate tightens. As a byproduct, inference about patterns of predictability will be much stronger when using some shrinkage. Panel B confirms this view, by reporting posterior probability for our measures of increasing slope. If we assume that all Sharpe ratios are equal to 0.25, then the probabilities of negative slope are all well below 1%. Of course, this case is an extreme form of prior, but imposing weaker assumptions already achieves meaningful regularization. For example, assuming $\sigma_a = 0.05$, that is that Sharpe ratios mostly lie between 0.15 and 0.35, already bring posterior probabilities of a negative slope to very low values. In summary, we find that if one is willing to impose views on unconditional Sharpe ratios to regularize elasticity estimates, the conclusions on patterns of predictability are reinforced. It is important to repeat that these stronger views are not building in the pattern of predictability in any way, but rather bring in extraneous plausible economic assumptions on unconditional properties of returns.

C. The Role of Households

Our main argument also makes predictions about proxies for household risk aversion in influencing risk premia. Specifically, when intermediaries matter, an increase in household risk aversion should, if anything, have larger effect in assets that are more easily accessed and more directly held by households. While we emphasize that our main tests do not require us to take a stand on the behavior of household risk aversion, we explore this hypothesis using two proxies for household willingness to bear risk often used in the literature: the aggregate consumption-wealth ratio, *cay* from Lettau and Ludvigson (2001), and the habit measure of Campbell and Cochrane (1999).

[Insert Figure 8 about here.]

Figure 8 plots predictive coefficient on intermediary risk aversion and household risk aversion, measured using *cay*; we report the regressions in Table IV. Panel A illustrates that controlling for household risk aversion does not affect the strongly increasing pattern of predictability by intermediary health as we move towards asset classes which are more difficult to access. In Panel B, we see that the coefficients on household risk aversion if anything exhibit the opposite pattern. They are mildly decreasing as we move from assets that are more to less directly held by households. As an aside, this result suggests there is nothing inherently mechanical in finding stronger predictability as we move along our ranking. Internet Appendix Figure IA7 shows the pattern of coefficients in our predictive regressions when we include the habit measure in our predictive regressions; Internet

Appendix Table IAXI reports the regressions. Again, we observe a distinctly decreasing pattern of predictability. Finally, we consider the dividend-price ratio on the CRSP value-weighted stock portfolio. While less directly related to aggregate household conditions, this measure offers an alternative proxy for valuations in the more frictionless asset class, stocks. A similar pattern emerges: a strongly increasing response to intermediary risk aversion, and a mild decrease in response to this proxy for household risk aversion; Internet Appendix Figure IA8 and Internet Appendix Table IAXII report the results.

[**Insert Table IV about here.**]

These results suggest that households also matter in influencing risk premiums, but in a way that is distinct from the role of intermediaries. Here again, it is worth pointing out that our proxies are imperfect: this household variation might be affected by changes in intermediary health, in the same way that the intermediary variation was potentially affected by changes in household risk aversion. It is nevertheless comforting that aggregate risk aversion proxies do indeed appear to line up with risk premiums as predicted by the model. The independent variation we have between the two types of proxy allows us to tease out the distinction between the two sides of our mechanism. And, this pattern highlights two distinct components of risk premium cycles.

D. Decomposing Variations in Expected Returns

We use these results to decompose variation in risk premiums into parts due to intermediaries and due to households. For each asset class, the predictive regression provides a baseline estimate of risk premium variation. Then, we use the lower bound implied by coefficient comparison across asset classes to quantify separately the role of intermediaries and households. These bounds, even when combined, do not explain all the variation in risk premium: there is some remaining measured variation in risk premium that we cannot trace specifically to one of these two sources.

We implement the decomposition as follows. We start from the predictive regression estimates coefficients $b_{i,I}$ and $b_{i,H}$ for each asset class — the points in Figure 8. To take into account of the patterns across asset classes, we fit linear slopes across these coefficients:

$$b_{i,I} = A_I + B_I \times c_i + u_{I,i} \tag{16}$$

$$b_{i,H} = A_H + B_H \times c_i + u_{H,i}, \tag{17}$$

where c_i increases linearly from 0 to 1. This linear fit corresponds to the red lines in Figure 8; statistical significance for the pattern of coefficients on intermediary risk aversion across assets was established in Section III.B. Using this linear model, we obtain an estimate of the total variance of risk premium for each asset class:

$$\sigma^2_{Total} = \sigma^2\left(\mathbb{E}_t(r^{\sigma}_{i,t+1})\right) = \sigma^2\left(\left(A_I + B_I \times c_i\right)\tilde{\gamma}_{I,t} + \left(A_H + B_H \times c_i\right)\tilde{\gamma}_{H,t}\right), \tag{18}$$

where we compute the variance of the right-hand-side quantity in our sample. Of course, this regression might include an incomplete set of predictors, and there might be more overall variation in risk premia in these asset classes. That said, we are reluctant to include more variables, because their choice would be somewhat arbitrary, and might artificially inflate risk premium variation through overfitting. In addition the high R^2s we obtain suggest it is unlikely that much more predictability is out there. Alternatively, one should just interpret our decomposition as the fraction of *our measured variation* in risk premium that can be traced back to intermediaries or households.

We then quantify how much variation we can attribute to intermediaries. In the framework of Section I.A, the lower bound on the role of intermediaries holds even if we have an imperfect proxy and comes from the stronger predictability of the more intermediated asset classes — see also the derivation in Internet Appendix Section IA.I.A. This pattern is driven by the positive slope B_I of the fit across asset classes: $\sigma^2_{Intermediaries} \geq \sigma^2\left(\left(B_I \times c_i\right)\tilde{\gamma}_{I,t}\right)$. Conversely, the role of households is driven by the decreasing predictability across asset classes, the negative of B_H: $\sigma^2_{Households} \geq \sigma^2\left(\left(-B_H \times (1 - c_i)\right)\tilde{\gamma}_{H,t}\right)$. For example, we cannot identify any variation due to intermediaries in the least intermediated asset class with $c_i = 0$, Stocks. Conversely we cannot identify any variation due to households in the most intermediated asset class, with $c_i = 1$, Credit.

Figure 9 presents the results. The red bars indicate the fraction of variation in risk premia for each asset class attributable to households $\sigma^2_{Households}/\sigma^2_{Total}$, the blue is that which is attributable to intermediaries $\sigma^2_{Intermediaries}/\sigma^2_{Total}$, and the gray region indicates variation we cannot confidently ascribe to either. This decomposition suggests that we can attribute at least 60% of the variation in CDS risk premia to intermediaries, and this declines by asset class but is still substantial for FX, commodities, and sovereign bonds. For households, we can attribute at least 40% of risk premia variation in stocks, and this declines as we move to more intermediated asset classes. Again, we stress that these are *lower* bounds: the gray area in the middle is not necessarily the sign of a third force, just variation not attributable to either category using our empirical strategy.

[**Insert Figure 9 about here.**]

Internet Appendix Figure IA9 repeats this exercise using elasticities rather than the predictive coefficients normalized by volatility. The decomposition is similar, with a somewhat larger role for intermediaries. Internet Appendix Figure IA10, assesses the sensitivity of our inference to the assumption of a linear structure across c_i. We use a local regression, a quadratic specification, and a cubic specification. The pattern and magnitude of variation in the fitted curves does not change much, with a stronger increase for intermediary risk aversion than for household risk aversion; these estimates lead to variance decomposition similar to our linear baseline. Finally, while the decomposition we have reported is based on our best estimates of the lower bound, one ought to remember these estimates come with uncertainty. Internet Appendix Table IAXIII builds on the approach of Section III.B to measure this uncertainty. The strong role of intermediaries comes with a relatively high precision, with for example an interquartile range for Credit between about 45% and 90%. In contrast the effect of households, in addition to be weaker, is also more imprecise.

Finally, note that we implement an *unconditional* decomposition: our empirical model assumes that predictive coefficients and the variance of predictors is constant over time. With more data — recall our sample did not allow us to detect variation in the predictive coefficient — or by imposing a more structural view one can imagine entertaining variation over time in these quantities. For example, periods of poor financial health of intermediaries, $\tilde{\gamma}_I > 0$, experience about 60% more standard deviation in $\tilde{\gamma}_I$ than periods of good health, while the standard deviation of $\tilde{\gamma}_H$ decreases by about 10%. This suggests a larger role of intermediaries in variations in risk premium during episodes of poor financial health of the financial sector, in line for example with the model of He and Krishnamurthy (2019).

IV. Additional Evidence

A. The Role of Variations in Risk

So far we only consider movements in intermediary and household risk aversion (and use the data to separate these two) but other determinants of returns may also change. One salient possibility is that risks vary over time, and this may be a particular concern for some of the assets we study with non-linear payoffs (e.g., options or credit). For example, in our model, the covariance matrix of asset payoffs might fluctuate. Variation in risk is a concern if more intermediated asset classes become riskier when intermediary distress increases. We explore this possibility in several ways empirically, by studying the behavior of various notions of risk including volatility, skewness, or time varying betas on standard asset pricing factors that include market risk and liquidity risk among others.

First, we ask whether intermediary health predicts *future* risk in a way that lines up with the pattern of risk premia. To do so, we run the same predictive regressions as our main table (Table III) but instead of putting future excess returns on the left-hand-side we use squared future returns. A positive coefficient indicates higher expected variance when intermediary risk aversion is high. Table V Panel A reports the coefficients from this regression to assess how intermediary health predicts risk; Internet Appendix Figure IA11 plots their values. Most importantly, these coefficients do not exhibit an increasing pattern. The more intermediated asset classes do not appear relatively riskier in times of high intermediary distress, so future risk does not qualitatively explain the increasing pattern of risk premiums that constitute our main result. In addition, the coefficients are, on average, slightly positive but much smaller in magnitude than the elasticities documented earlier. In particular, in response to a one-standard deviation increase in intermediary risk aversion, variance increases by about 20% while the risk premium increases by close to 100%. This implies that changes in risk cannot account for the increase in the average risk premium. Measures of risk beyond variance could be relevant to risk premia; we explore some of these possibilities. In Internet Appendix Figure IA13, we study downside risk and skewness in addition to variance. We find no support for the idea that crash risk or left skewness increases for the more intermediated asset returns when intermediary risk aversion is high. If anything, the results go slightly in the opposite direction, though are not large quantitatively.

[Insert Table V about here.]

Our assets might exhibit nonlinear exposures to other aggregate risk factors that variance and skewness do not capture. This would imply variation in conditional betas, and therefore expected returns, even absent variation in total risk. To assess whether such a behavior is driving our results, one must take a stand on what the relevant *economic* risk factors are. Indeed, without such a stance, one could just reverse engineer the stochastic discount factor that prices these 8 asset classes and seemingly explain risk premia, even though this would be a mechanical result.[37] Rather, the interesting question is whether changes in exposure to some economic sources of risk is the reason expected returns are changing. One such risk could be captured by the aggregate stock market return. For example, because stocks are the asset class with the lowest cost, their return might reveal frictionless economic risk. In Panel B of Table V, we ask whether intermediary risk aversion predicts the

[37]See for example Kozak, Nagel, and Santosh (2018) for a discussion of the lack of interpretability of reduced-form factor models.

covariance of returns with the market.[38] While the covariance with the market appears to increase somewhat across the asset classes in times of high intermediary risk aversion, the magnitudes are small relative to the change in expected returns we have documented. And more importantly, the increases are *less* pronounced for more intermediated asset classes, ruling out these variations as an explanation for our main results. In Panel C, we consider the exposure to the liquidity factor of Pastor and Stambaugh (2003). As we have discussed earlier, differential exposures to liquidity could be an alternative mechanism by which intermediaries affect risk premia. However, the changes in exposure to liquidity are tiny and do not exhibit an increasing pattern. So time-varying exposure to liquidity risk does not appear to explain our findings.

An alternative way to assess whether time-varying exposures are behind our findings is to absorb these variations in the predictive regression of returns. This approach allows to consider multiple sources of conditional risk simultaneously. We do so following the framework in Shanken (1990) where we explicitly allow intermediary risk aversion $\gamma_{I,t}$ to proxy for time-varying loadings on pre-specified standard risk factors.[39] Specifically, we estimate the following specification:

$$r_{i,t+1} = a_i + b_{I,i} \times \tilde{\gamma}_{I,t} + \sum_k (\beta_{0,i,k} + \beta_{1,i,k}\tilde{\gamma}_{I,t}) f_{k,t+1} + \epsilon_{i,t+1}, \qquad (19)$$

[38] This approach does not rule out richer variation in risk exposure uncorrelated to $\tilde{\gamma}_I$. Rather, it just asks if there is variation in the risk exposure correlated with $\tilde{\gamma}_I$.

[39] See also Kelly et al. (2019).

with a set of returns on K risk factors $f_{k,t+1}, k = 1, ..., K$. If time-varying exposure to these risk factors is not driving our main results, then the inclusion of the interaction $\tilde{\gamma}_{I,t} \times f_{k,t+1}$ (that is, entertaining the possibility that $\beta_{1,i,k} \neq 0$) will not affect the estimates of $b_{I,i}$.[40] In Internet Appendix Table IAXIV, we compare these specifications for various sets of factors to our baseline. In addition to the market and the liquidity factor, we include the value and size factors of Fama and French (1993), the momentum factor, and the short-term reversal factor which has been argued to proxy for liquidity provision (Nagel (2012)). While the connection of all these factors to specific sources of economic risk is not always well-established, they have been used extensively in the reduced-form literature to capture variation in expected returns. We find that none of these specifications meaningfully affect the coefficients $b_{i,I}$ and in particular the cross-sectional slope across regression coefficient moves very little across specifications. Taken together, these results do not support an explanation of the increasing pattern of predictive coefficients based on time-varying quantities of risk.

Finally, controlling for ex-ante (rather than ex-post) measures of risks in the predictive regressions has little effect on the results. Specifically, we include time-varying volatilities and market betas in our regressions constructed using rolling 5-year regressions for each asset class. Internet Appendix Table IAXV reports the estimates. The pattern in coefficients on intermediary health is similar to before, or if anything slightly stronger and more significant due to the addition of controls. Internet Appendix Figure

[40] Including the factors in themselves could affect $b_{i,I}$ within our model if they capture a sizable fraction of shocks to returns.

IA12 confirms this result visually, plotting the predictive coefficients and elasticities.

Taken together, these results suggest that fluctuations in the quantity of risk do not drive out the role of intermediary health. We acknowledge that a time-varying risk story can never be completely ruled out – for example if risk is captured by some other unobserved time-varying covariances we cannot measure. However, across a large variety of typical risk measures we find no evidence that time-varying risk goes in the right direction to qualitatively explain our results.

B. Evidence from Hedge Fund Returns

As an alternative to comparing predictability by intermediary health across asset classes, we compare the properties of more or less complex strategies. We use indices of hedge fund returns from Dow Jones Credit Suisse. Hedge fund returns are the returns of specialized strategies and asset classes. We argue that they should respond more to intermediary health than other assets, more so for more complex strategies. We run our predictive regressions again with stocks on the left of our ranking and various hedge fund return strategies on the right. We consider long short equity, equity market neutral, an overall hedge fund index from DJCS of all funds, event driven, convertible bond arbitrage funds, and fixed income arbitrage funds. While we acknowledge a detailed ranking of strategies as we pursued earlier across asset classes is not available, we use guidance from the previous literature on the complexity of these strategies and the degree sophisticated intermediaries are involved in each. We argue that equity strategies are likely more accessible to households. For example, some quant strategies in equities like value and momentum could be implemented by households though at likely higher costs. On the other hand, convertible bond arbitrage and fixed income arbitrage are likely the most difficult for households to engage in. Indeed, intermediary capital effects have been argued to play an important role in both of these strategies. Mitchell and Pulvino (2012) state that hedge funds and arbitrageurs make up 75% of the convertible bond market and document significant dislocations in prices following hedge fund redemptions. Hu et al. (2013) suggest intermediary capital effects cause deviations in fixed income along the yield curve. Duarte et al. (2007) study specialized fixed income arbitrage pursued by hedge funds, arguing that these strategies require significant intellectual capital and leverage. Event-driven strategies (e.g., merger arbitrage) likely fit in the middle and also exhibit price pressure effects (Mitchell et al., 2004); so does the index of all hedge funds which is weighted by AUM under each asset class.

Figure 10 and Internet Appendix Table IAXVI report our results. We find that predictability is higher for all hedge fund strategies compared to stocks, consistent with out main hypothesis that these constitute more specialized strategies that households would have difficulty investing in. Within hedge fund strategies, we also find that convertible bond arbitrage, fixed income arbitrage, and event-driven strategies respond more to intermediary health. Interestingly, the magnitude of predictive coefficients normalized by volatility for more specialized strategies is comparable to that of the more sophisticated asset classes of our main sample. Overall, these results are consistent with the idea that these strategies are more complex and specialized. These results further support the view that intermediaries matter, using separate data on returns than our main analysis. Thus, they strengthen our conclusions.

[**Insert Figure 10 about here.**]

V. Discussion and Relation to Literature

Having documented our main results, it is useful to contrast our approach with the existing work on intermediary asset pricing. We find this discussion more useful ex-post so that we can relate to the literature the particular aspects of our empirical work and our model.

A. Contrast with the Euler Equation Approach

A classic approach to study households' optimization in financial markets is by studying whether their Euler equation holds. This corresponds to asking whether their marginal utility is a stochastic discount factor that can price the cross-section of expected returns. A natural counterpart to this approach for a view that intermediaries are central to asset pricing is to ask whether their Euler equation also holds. Several papers empirically evaluate the intermediary Euler equation. For instance Adrian et al. (2014) and He et al. (2017) construct empirical counterparts of intermediaries' marginal utility and find empirical success in using these variables to explain the cross-section of expected returns.

In our setting, intermediaries have frictionless access to the risky asset market and therefore their Euler equation always holds. In fact, the portfolio of intermediaries is always mean-variance efficient — see Equation (7) – which implies their marginal utility is a valid pricing kernel for all assets. However, in our model this is true independently of whether intermediaries matter for asset prices. The empirical success of the intermediary Euler equation, while very useful, only validates the specification of a frictionless demand function for intermediaries.

Tests of the household Euler equation can complement this evidence. In our setting, intermediaries do matter if and only if the household Euler equation fails. This is a direct consequence of the observation that when intermediaries do not matter, prices coincide with the frictionless benchmark. Specifically, in our model the CAPM does not hold unless $\left(\Sigma + \frac{1}{\gamma_I}\Sigma_{diag}C\right)^{-1}\left(\Sigma + \frac{1}{\gamma_H}\Sigma_{diag}C\right)$ is the identity matrix.

Since Hansen and Singleton (1983), there is a long literature providing evidence inconsistent with particular specifications of the Euler equation for households. It remains somewhat unclear if this empirical failure reflects the fact that the household Euler equation does not hold, or that we have insufficient models of household marginal utility, or that data on quantities like aggregate consumption are poor for these purposes. For example, Greenwald et al. (2014) argue that movements in aggregate risk aversion appear uncorrelated with standard measures of consumption. Malloy et al. (2009) argue that stockholder consumption lines up better with asset returns, while papers such as Constantinides and Duffie (1996) and Schmidt (2015) focus on household heterogeneity and idiosyncratic risk. Savov (2011) and Kroencke (2017) argue that measurement of NIPA consumption plays a role in the failure of the CCAPM. These papers point to failures of the CCAPM for specific reasons related to preferences or measurement. The approach of this paper is to go beyond these shortcomings and instead to discuss alternative predictions of the theory, more directly focused on intermediaries.

B. Micro Evidence

Our results also relate to "micro" studies which show intermediary frictions mattering in a particular asset class or at a particular point in time. For example, Siriwardane (2016) shows price dispersion in CDS contracts that relates to dealer net worth. That is, losses for a particular dealer on other contracts affect the CDS price that dealer is willing to offer, that is it affects their risk-bearing capacity. Similarly, Du et al. (2017) document that end of quarter regulatory constraints for banks affect their risk-bearing capacity and spill over into FX markets. These end-of-quarter constraints result in large violations of covered interest parity for short periods of time. Gabaix et al. (2007) provide evidence that banks are marginal investors in mortgage backed securities (MBS). Duffie (2010) and He and Krishnamurthy (forthcoming) provide a host of similar examples.[41] These studies are extremely useful in documenting clean effects of intermediaries on asset prices, by getting as directly as possible to the mechanisms behind intermediary decisions. However, a limitation is that they typically study specific relative price effects at particular points in time but do not give a sense of more aggregate effects of intermediaries. In particular, it could be that the local disruptions they document "wash out" in the aggregate. By zooming out at a very aggregate

[41] See also Lou et al. (2013).

level, we get directly at broad asset class variation in risk premiums.

The natural next step is to relate magnitudes of intermediary effects in the microeconomic studies with those in the aggregate evidence. To do so, one has to step out of our simple, single intermediary framework, and account for the structure of the intermediary sector. In Internet Appendix Section IA.I.E, we extend the model of Section I.B to a richer organization of the intermediary sector. Using this setting, one can consider both broad shocks to the entire intermediary sector, or shocks specific to a single intermediary, and understand the determinants of the price response they imply. However, we also show that doing so relies crucially on quantifying specific properties of the intermediation sector. One must measure the degree of substitutability across various intermediaries and the assets they specialize in.[42] For example, to aggregate the results of Siriwardane (2016), one needs to know how easy it is for another dealer to access the same CDS contract, and how closely related are the risks of various CDS contracts. In addition, one needs to be able to relate the shocks used in various empirical settings to common measures of risk appetite. Doing so is beyond the scope of this paper, and we leave this promising avenue open for future research.

VI. Conclusion

A sufficient condition for intermediaries to matter for asset prices is that the risk premium of more intermediated assets responds relatively more to

[42]Morelli et al. (2019) is an example of using such an exercise in the context of emerging market debt.

changes in intermediary risk appetite. This prediction is valid even if intermediary risk appetite is positively correlated with other aggregate drivers of risk appetite. We provide direct empirical evidence of this pattern of variation in risk premia. Hence, we argue that intermediaries matter for a number of key asset classes including CDS, FX, MBS, and commodities. Quantitatively, a sizable amount of variation in risk premium in these asset classes is attributable to intermediaries. We view this study as a first step in quantifying the effect of intermediaries on variation in aggregate asset prices.

REFERENCES

Adrian, Tobias, Erkko Etula, and Tyler Muir, 2014, Financial intermediaries and the cross-section of asset returns, *The Journal of Finance* 69, 2557–2596.

Adrian, Tobias, and Hyun Song Shin, 2014, Procyclical leverage and value-at-risk, *Review of Financial Studies* 27, 373–403.

Ang, Andrew, and Geert Bekaert, 2006, Stock return predictability: Is it there? *The Review of Financial Studies* 20, 651–707.

Asvanunt, Attakrit, and Scott Richardson, 2016, The credit risk premium, *The Journal of Fixed Income* 26, 6–24.

Brunnermeier, Markus, and Lasse Pedersen, 2009, Market liquidity and funding liquidity, *Review of Financial Studies* 22, 2201–2238.

Brunnermeier, Markus K, and Yuliy Sannikov, 2014, A macroeconomic model with a financial sector, *The American Economic Review* 104, 379–421.

Campbell, John Y., and John Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251.

Campbell, John Y, and Samuel B Thompson, 2008, Predicting excess stock returns out of sample: Can anything beat the historical average? *The Review of Financial Studies* 21, 1509–1531.

Chen, Hui, Scott Joslin, and Sophie Xiaoyan Ni, 2019, Demand for crash insurance, intermediary constraints, and risk premia in financial markets, *The Review of Financial Studies* 32, 228–265.

Chodorow-Reich, Gabriel, Andra Ghent, and Valentin Haddad, 2021, Asset insulators, *The Review of Financial Studies* 34, 1509–1539.

Constantinides, George M, and Darrell Duffie, 1996, Asset pricing with heterogeneous consumers, *Journal of Political economy* 219–240.

Danielsson, Jon, Hyun Song Shin, and Jean-Pierre Zigrand, 2011, Balance sheet capacity and endogenous risk, *working paper* .

Dew-Becker, Ian, Stefano Giglio, Anh Le, and Marius Rodriguez, 2017, The price of variance risk, *Journal of Financial Economics* 123, 225 – 250.

Diep, Peter, Andrea L Eisfeldt, and Scott A Richardson, Forthcoming, Prepayment risk and expected mbs returns, *Journal of Finance* .

Du, Wenxin, Alexander Tepper, and Adrien Verdelhan, 2017, Deviations from covered interest rate parity, Technical report, National Bureau of Economic Research.

Duarte, Jefferson, Francis A Longstaff, and Fan Yu, 2007, Risk and return in fixed-income arbitrage: Nickels in front of a steamroller? *The Review of Financial Studies* 20, 769–811.

Duffie, Darrell, 2010, Presidential address: Asset price dynamics with slow-moving capital, *The Journal of finance* 65, 1237–1267.

Eisfeldt, Andrea L, Hanno N Lustig, and Lei Zhang, 2017, Risk and return in segmented markets with expertise, Technical report, National Bureau of Economic Research.

Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bond, *Journal of Financial Economics* 33, 3–56.

Fleckenstein, Matthias, Francis A. Longstaff, and Hanno Lustig, 2014, The tips-treasury bond puzzle, *The Journal of Finance* 69, 2151–2197.

Gabaix, Xavier, Arvind Krishnamurthy, and Olivier Vigneron, 2007, Limits of arbitrage: Theory and evidence from the mortgage-backed securities market, *The Journal of Finance* 62, 557–595.

Garleanu, Nicolae, and Lasse Heje Pedersen, 2011, Margin-based asset pricing and deviations from the law of one price, *The Review of Financial Studies* 24, 1980–2022.

Gilchrist, Simon, and Egon Zakrajek, 2012, Credit spreads and business cycle fluctuations, *American Economic Review* 102, 1692–1720.

Gorton, Gary, and K Geert Rouwenhorst, 2006, Facts and fantasies about commodity futures, *Financial Analysts Journal* 62, 47–68.

Greenwald, Daniel L., Martin Lettau, and Sydney C. Ludvigson, 2014, The origins of stock market fluctuations, *working paper* .

Haddad, Valentin, Serhiy Kozak, and Shrihari Santosh, 2020, Factor timing, *The Review of Financial Studies* 33, 1980–2018.

Haddad, Valentin, and David Sraer, 2020, The banking view of bond risk premia, *The Journal of Finance* 75, 2465–2502.

Hansen, Lars Peter, and Kenneth J Singleton, 1983, Stochastic consumption, risk aversion, and the temporal behavior of asset returns, *Journal of political economy* 91, 249–265.

He, Zhiguo, Bryan Kelly, and Asaf Manela, 2017, Intermediary asset pricing: New evidence from many asset classes, *Journal of Financial Economics* forthcoming.

He, Zhiguo, and Arvind Krishnamurthy, 2013, Intermediary asset pricing, *The American Economic Review* 103, 732–770.

He, Zhiguo, and Arvind Krishnamurthy, 2019, A macroeconomic framework for quantifying systemic risk, *American Economic Journal: Macroeconomics* 11, 1–37.

He, Zhiguo, and Arvind Krishnamurthy, forthcoming, Intermediary asset pricing and the financial crisis, *Annual Review of Financial Economics* .

Hodrick, Robert J, 1992, Dividend yields and expected stock returns: Alternative procedures for inference and measurement, *Review of Financial Studies* 5, 357–86.

Hu, Grace Xing, Jun Pan, and Jiang Wang, 2013, Noise as information for illiquidity, *The Journal of Finance* 68, 2341–2382.

Kelly, Bryan T, Seth Pruitt, and Yinan Su, 2019, Characteristics are covariances: A unified model of risk and return, *Journal of Financial Economics* 134, 501–524.

Koijen, Ralph SJ, and Motohiro Yogo, 2015, An equilibrium model of institutional demand and asset prices, Technical report, National Bureau of Economic Research.

Kozak, Serhiy, Stefan Nagel, and Shrihari Santosh, 2018, Interpreting factor models, *The Journal of Finance* 73, 1183–1223.

Kozak, Serhiy, Stefan Nagel, and Shrihari Santosh, 2020, Shrinking the cross-section, *Journal of Financial Economics* 135, 271–292.

Krishnamurthy, Arvind, 2010, How debt markets have malfunctioned in the crisis, *Journal of Economic Perspectives* 24, 3–28.

Kroencke, Tim A, 2017, Asset pricing without garbage, *The Journal of Finance* 72, 47–98.

Lettau, Martin, and Sydney Ludvigson, 2001, Resurrecting the (c)capm: A cross-sectional test when risk premia are time-varying, *Journal of Political Economy* 109, 1238–1287.

Lewis, Kurt F., Francis A. Longstaff, and Lubomir Petrasek, 2017, Asset mispricing, Working Paper 23231, National Bureau of Economic Research.

Lou, Dong, Hongjun Yan, and Jinfan Zhang, 2013, Anticipated and repeated shocks in liquid markets, *The Review of Financial Studies* 26, 1891–1912.

Malloy, Christopher J., Tobias J. Moskowitz, and Annette Vissing-Jorgensen, 2009, Long-run stockholder consumption risk and asset returns, *Journal of Finance* 64, 2427–2479.

Mitchell, Mark, Lasse Heje Pedersen, and Todd Pulvino, 2007, Slow moving capital, *American Economic Review* 97, 215–220.

Mitchell, Mark, and Todd Pulvino, 2012, Arbitrage crashes and the speed of capital, *Journal of Financial Economics* 104, 469 – 490, Market Institutions, Financial Market Risks and Financial Crisis.

Mitchell, Mark, Todd Pulvino, and Erik Stafford, 2004, Price pressure around mergers, *The Journal of Finance* 59, 31–63.

Morelli, Juan M, Pablo Ottonello, and Diego J Perez, 2019, Global banks and systemic debt crises, *New York University and University of Michigan* .

Muir, Tyler, 2017, Financial crises and risk premia, *The Quarterly Journal of Economics* .

Nagel, Stefan, 2012, Evaporating Liquidity, *The Review of Financial Studies* 25, 2005–2039.

Pástor, L'uboš, 2000, Portfolio selection and asset pricing models, *The Journal of Finance* 55, 179–223.

Pástor, L'uboš, and Robert F Stambaugh, 2000, Comparing asset pricing models: an investment perspective, *Journal of Financial Economics* 56, 335–381.

Pastor, Lubos, and Robert F. Stambaugh, 2003, Liquidity risk and expected stock returns, *Journal of Political Economy* 111, 642–685.

Santos, Tano, and Pietro Veronesi, 2016, Habits and leverage, Technical report, National Bureau of Economic Research.

Savov, Alexi, 2011, Asset pricing with garbage, *The Journal of Finance* 66, 177–201.

Schmidt, Lawrence, 2015, Climbing and falling off the ladder: Asset pricing implications of labor market event risk, *working paper* .

Shanken, Jay, 1990, Intertemporal asset pricing: An empirical investigation, *Journal of Econometrics* 45, 99–120.

Siriwardane, Emil, 2016, Concentrated capital losses and the pricing of corporate credit risk, *working paper, Harvard University* .

Stambaugh, Robert F, 1999, Predictive regressions, *Journal of Financial Economics* 54, 375–421.

Zellner, Arnold, 1962, An efficient method of estimating seemingly unrelated regressions and tests for aggregation bias, *Journal of the American statistical Association* 57, 348–368.

Table I. Summary statistics of asset returns.

We report means, standard deviations, and Sharpe ratios of excess returns for each of the asset classes. All numbers are quarterly. The text describes the returns and sources in detail.

	(1) Stocks	(2) Treas.	(3) Options	(4) Sov.	(5) Comm.	(6) FX	(7) MBS	(8) Credit
$E[r_t]$	1.50%	0.46%	2.35%	1.97%	1.29%	2.06%	0.09%	0.28%
$\sigma(r_t)$	9.10%	1.91%	4.78%	5.39%	6.28%	5.04%	0.68%	1.31%
$E[r_t]/\sigma(r_t)$	0.17	0.24	0.49	0.37	0.21	0.41	0.13	0.22
Observations	167	160	103	65	105	116	97	47

Table II. Ranking of asset classes.

Ranking by degree of intermediation by source, with our chosen ranking on the top row. From left to right is less intermediated asset classes, with relatively easier access of investing by households, to more intermediated asset classes, with lower participation by households. The sources for the rankings are: the Flow of Funds (FoF), BIS derivatives positions, Vale-at-Risk (VaR), and ETF expense ratios. The text explains these sources and rankings in detail.

Our Ranking	Stocks	Bonds	Options	Sov Bonds	Comm	FX	MBS	CDS
FoF	Stocks	Bonds		Sov Bonds			MBS	
VaR	Stocks	Bonds			Comm	FX		
BIS		Bonds	Options		Comm	FX		CDS
Expense	Stocks	Bonds	Sov Bonds	FX	Comm	Options	MBS	CDS

Table III. Intermediary health and excess returns.

Predictive regressions of future excess returns in each asset class on our proxy for intermediary risk aversion, $\tilde{\gamma}_{I,t}$. We run: $r^{\sigma}_{i,t+1} = a_i + b_i \times \tilde{\gamma}_{I,t} + \epsilon_{i,t+1}$ and report b_i. Excess returns $r^{\sigma}_{i,t+1}$ are normalized by their full sample volatility. $\tilde{\gamma}_{I,t}$ is the standardized average of the AEM and HKM intermediary factors. Standard errors are computed using the reverse regression approach of Hodrick(1992). *, **, and *** means statistically different from zero at 10, 5 and 1% level of significance, where the p-values are computed using the bootstrap approach described in Section III.A. The last row, elasticity, computes the elasticity of expected returns as $b_i/E[r^{\sigma}_{i,t+1}]$. See text for more details.

	(1) Stocks	(2) Treas.	(3) Options	(4) Sov.	(5) Comm.	(6) FX	(7) MBS	(8) Credit
γ_I	0.12	-0.01	0.29***	0.38**	0.18*	0.18*	0.30**	0.57***
	(0.09)	(0.07)	(0.10)	(0.17)	(0.10)	(0.09)	(0.13)	(0.22)
Boots. p-value	0.198	0.904	0.005	0.019	0.083	0.056	0.016	0.006
Observations	167	160	103	65	105	116	97	47
Adjusted R^2	0.008	-0.006	0.075	0.126	0.022	0.021	0.078	0.316
Elasticity								
	0.71	-0.04	0.58	1.03	0.87	0.43	2.34	2.67

Table IV. Predicting returns with intermediary and household risk aversion.

Predictive regressions of future excess returns in each asset class on our proxy for intermediary risk aversion, $\tilde{\gamma}_{I,t}$ and household risk aversion, $\tilde{\gamma}_{H,t}$. We run: $r^{\sigma}_{i,t+1} = a_i + b_{I,i} \times \tilde{\gamma}_{I,t} + b_{H,i} \times \tilde{\gamma}_{H,t} + \epsilon_{i,t+1}$ and report coefficients b_i. Excess returns $r^{\sigma}_{i,t+1}$ are normalized by their full sample volatility. $\tilde{\gamma}_{I,t}$ is the standardized average of the AEM and HKM intermediary factors. $\tilde{\gamma}_{H,t}$ is proxied by the consumption wealth ratio (cay) from Lettau and Ludvigson (2001). Standard errors are computed using the reverse regression approach of Hodrick(1992). *, **, and *** means statistically different from zero at 10, 5 and 1% level of significance, where the p-values are computed using the bootstrap approach described in Section III.A. The last row, elasticity, computes the elasticity of expected returns as $b_i/E[r^{\sigma}_{i,t+1}]$. See text for more detail.

	(1) Stocks	(2) Treas.	(3) Options	(4) Sov.	(5) Comm.	(6) FX	(7) MBS	(8) Credit
γ_I	0.15*	-0.00	0.29***	0.36**	0.18*	0.18*	0.31**	0.59**
	(0.09)	(0.07)	(0.10)	(0.17)	(0.10)	(0.09)	(0.13)	(0.27)
γ_H^{cay}	0.21***	0.06	0.12	0.22	0.01	0.12	0.20*	-0.06
	(0.07)	(0.07)	(0.12)	(0.14)	(0.12)	(0.11)	(0.11)	(0.34)
Observations	167	160	103	65	105	116	97	47
Adjusted R^2	0.044	-0.009	0.075	0.144	0.013	0.022	0.092	0.302
Elasticity								
γ_I	0.91	-0.02	0.58	1.00	0.87	0.43	2.39	2.74
γ_H^{cay}	1.23	0.25	0.24	0.59	0.05	0.29	1.53	-0.26

Table V. Predicting risk with intermediary risk aversion.

Predictive regressions of future risk measures in each asset class on our proxy for intermediary risk aversion, $\tilde{\gamma}_{I,t}$. We run: $Y_{i,t+1} = a_i + b_{I,i} \times \tilde{\gamma}_{I,t} + \epsilon_{i,t+1}$ and report coefficients $b_{I,i}$. Excess returns $r_{i,t+1}$ are normalized by their full sample volatility. Panel A predicts the square returns, $Y_{i,t+1} = r^2_{i,t+1}$. Panel B predicts the exposure to market returns, $Y_{i,t+1} = r_{i,t+1} \times r_{MKT,t+1}$. Panel C predicts the exposure to the liquidity factor of Pastor and Stambaugh (2003), $Y_{i,t+1} = r_{i,t+1} \times r_{LIQ,t+1}$. Standard errors are computed using the reverse regression approach of Hodrick(1992). *, **, and *** means statistically different from zero at 10, 5 and 1% level of significance, where the p-values are computed using the bootstrap approach described in Section III.A.

Panel A. Variance ($r^2_{i,t+1}$)

	(1) Stocks	(2) Treas.	(3) Options	(4) Sov.	(5) Comm.	(6) FX	(7) MBS	(8) Credit
γ_I	0.34**	0.17	0.10	0.23	0.30	-0.09	0.15	0.35
	(0.16)	(0.13)	(0.17)	(0.34)	(0.26)	(0.11)	(0.20)	(0.54)
Observations	167	160	103	65	105	116	97	47
Adjusted R^2	0.041	-0.002	-0.007	-0.005	0.004	-0.004	-0.004	-0.000

Panel B, Market Risk Exposure ($r_{i,t+1} \times r_{MKT,t+1}$)

	(1) Stocks	(2) Treas.	(3) Options	(4) Sov.	(5) Comm.	(6) FX	(7) MBS	(8) Credit
γ_I	0.38**	0.07	0.27*	0.24	0.09	0.14	0.11	0.22
	(0.16)	(0.08)	(0.14)	(0.28)	(0.11)	(0.12)	(0.17)	(0.39)
Observations	167	160	103	65	105	116	97	47
Adjusted R^2	0.057	-0.003	0.046	0.009	-0.004	0.006	-0.003	-0.005

Panel C. Liquidity Risk Exposure ($r_{i,t+1} \times r_{LIQ,t+1}$)

	(1) Stocks	(2) Treas.	(3) Options	(4) Sov.	(5) Comm.	(6) FX	(7) MBS	(8) Credit
γ_I	0.07	-0.03	0.27**	0.06	0.14	0.04	0.13	0.00
	(0.14)	(0.06)	(0.14)	(0.14)	(0.18)	(0.10)	(0.16)	(0.27)
Observations	167	160	103	65	105	116	97	47
Adjusted R^2	-0.004	-0.006	0.026	-0.013	-0.004	-0.008	-0.002	-0.022

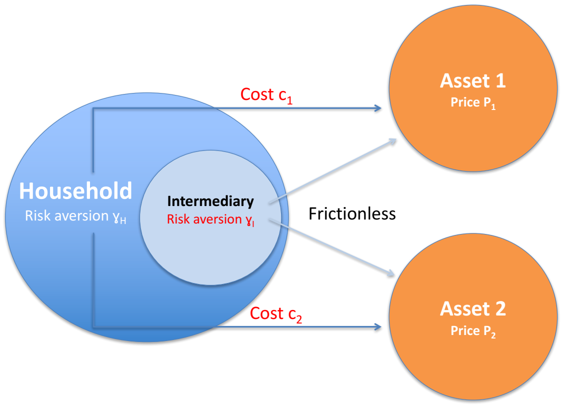

Figure 1. Model setting.

This figure describes the model with two risky assets but this picture easily generalizes to N assets. We highlight that the household owns the intermediary in the model (though they may have differing risk aversions) and that the household can also invest directly into various assets at different costs c_1, c_2. The costs might be higher in some assets (e.g., CDS markets) than others (e.g., the stock market).

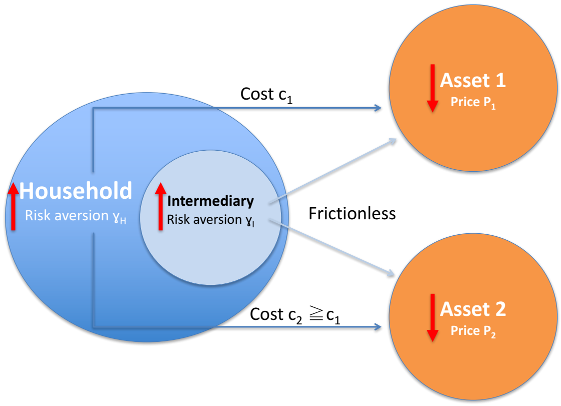

Panel A. Response to aggregate risk aversion shock under null

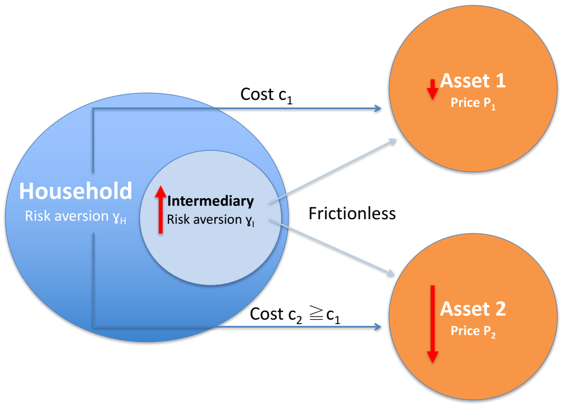

Panel B. Response to intermediary risk aversion shock

Figure 2. Model shocks.

This figure describes the response of asset prices to risk aversion changes. In Panel A, we show the response of a risk aversion shock under the null that intermediaries *do not* matter (either because $c = 0$ for all assets *or* because $\gamma_I = \gamma_H$) and in this case all risk premia move proportionally when risk aversion changes. In Panel B, we show the response of an intermediary risk aversion shock in the case where there are differential costs c across assets and show how the cross-section of risk premia change.

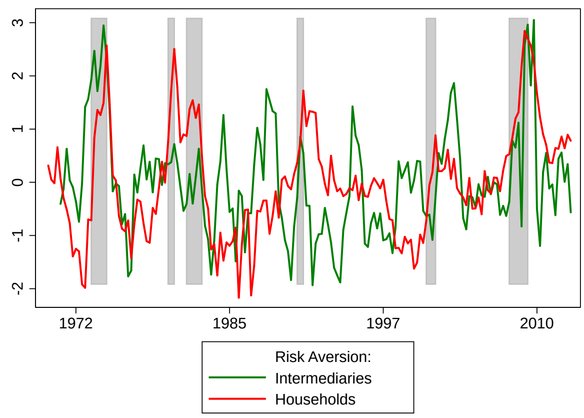

Figure 3. Intermediary and household risk aversion.
This figure plots our proxy for intermediary risk aversion taken from AEM and HKM in green. For reference, it also plots aggregate risk aversion implied by a habit model using aggregate consumption in red.

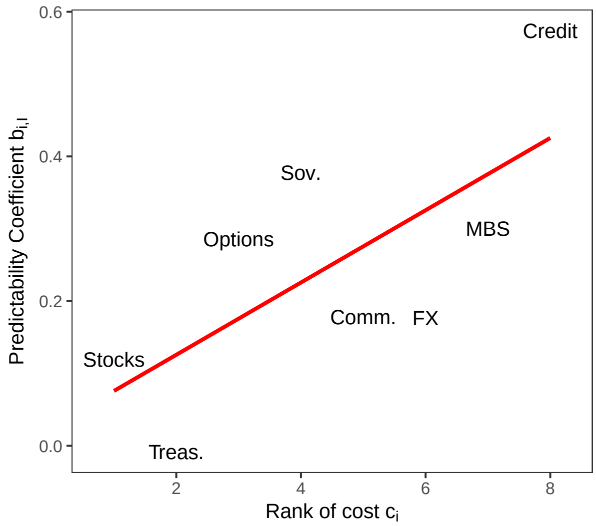

Figure 4. Predictability across asset classes: predictive coefficients.
We plot the predictive coefficients from Table III, which runs predictive regressions of excess returns on intermediary effective risk aversion. The x-axis is our ranking for how intermediated each asset class is. The red line is a linear regression fit through these points. An upward slope indicates more predictability in more intermediated asset classes.

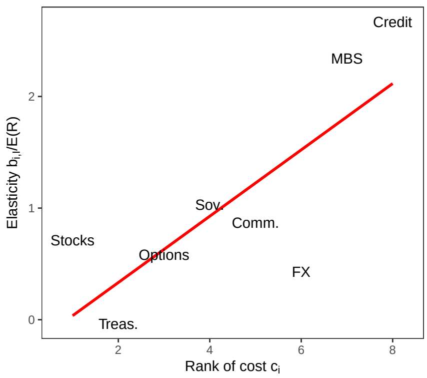

Panel A. Elasticity

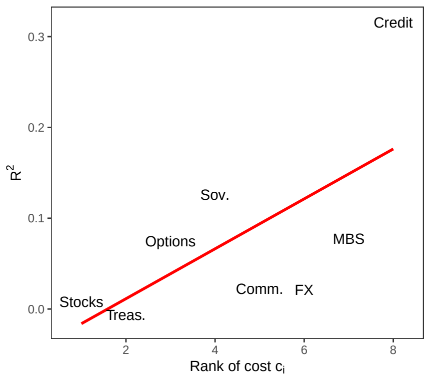

Panel B. R^2

Figure 5. Predictability across asset classes: risk premium elasticity and R^2.
Panel A repeats the previous figure using the elasticity of the risk premium to intermediary risk aversion: the predictive coefficient divided by the sample mean of excess returns in each asset class. Panel B measures the degree of predictability using the R^2 in each predictive regression across asset classes (see Table III). The x-axis is our ranking for how intermediated each asset class is. The red line is a linear regression fit through these points. An upward slope indicates more predictability in more intermediated asset classes.

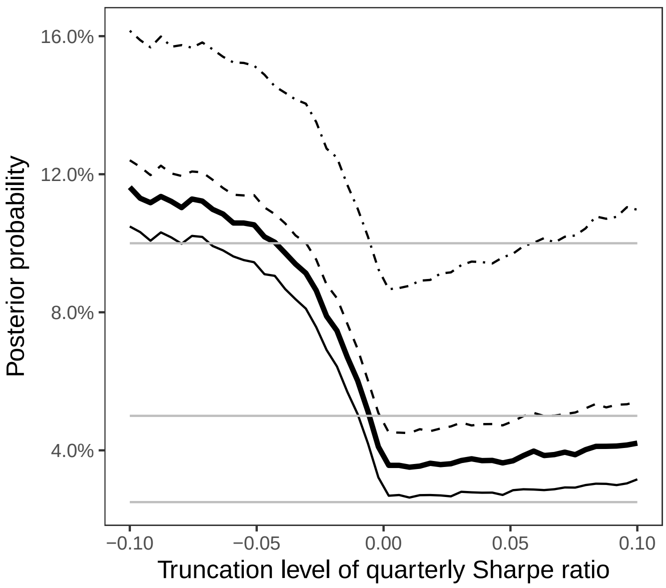

Figure 6. Posterior probability of less predictability for intermediated asset classes: the effect of truncation.
We report the posterior probability that the elasticity of risk premium to intermediary health is lower for more intermediated assets. The x-axis varies the lower bound on the prior of the unconditional quarterly Sharpe ratio. The different lines correspond to different criteria for less predictability. The thick black line is for the slope of a linear regression across elasticities. The thin black line compares the average elasticity of Credit and MBS relative to the average elasticity of Stocks and Treasuries. The dotted-dash line compares Stocks and Treasuries to all other asset classes. The dotted line compares Stocks, Treasuries, and Options to all other asset classes.

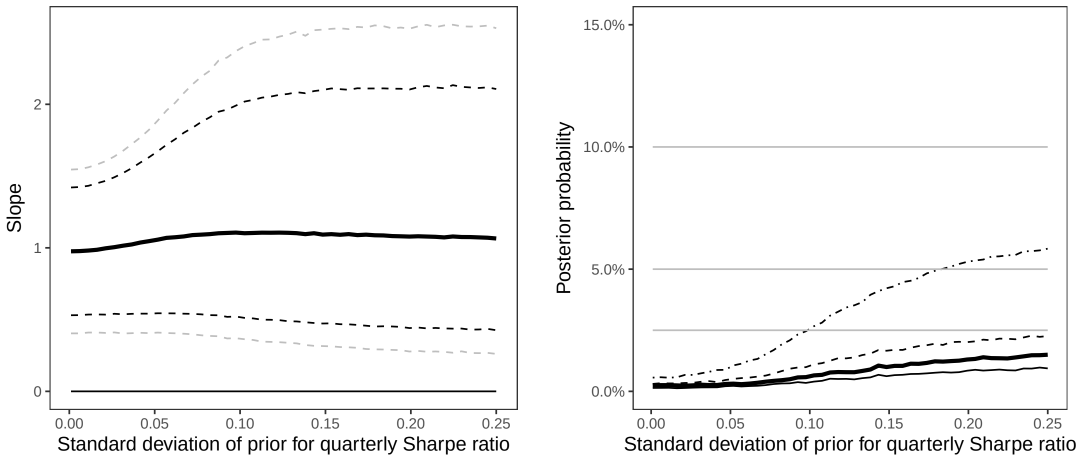

Panel A. Slope across asset classes

Panel B. Probability of less predictability for intermediated assets

Figure 7. Statistical significance of predictability across asset classes.

The x-axis varies the standard deviation on the prior of the unconditional quarterly Sharpe ratio, lower values imply a shrinkage towards the assumption of a constant unconditional Sharpe ratio across asset classes equal to 0.25. Panel A reports the distribution of the slope of a linear regression across elasticities: median, 5th, 10th, 90th, and 95th percentile. Panel B reports the posterior probability that the elasticity of risk premium to intermediary health is lower for more intermediated assets. The different lines correspond to different criteria for less predictability. The thick black line is for the slope of a linear regression across elasticities. The thin black line compares the average elasticity of Credit and MBS relative to the average elasticity of Stocks and Treasuries. The dotted-dash line compares Stocks and Treasuries to all other asset classes. The dotted line compares Stocks, Treasuries, and Options to all other asset classes.

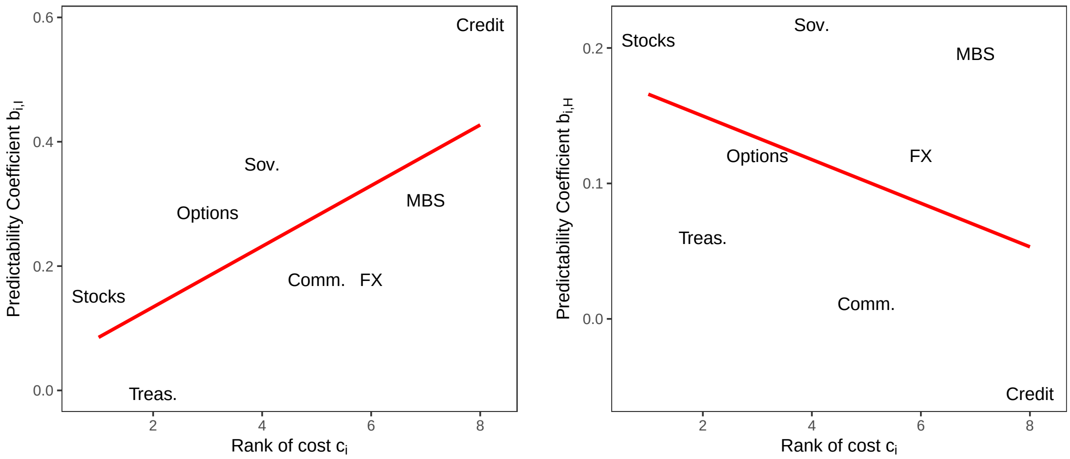

Panel A. Intermediaries

Panel B. Households

Figure 8. Predictability across asset classes: households versus intermediaries.
We plot coefficients from a predictive regressions of excess returns on intermediary effective risk aversion and household risk aversion, proxied by the consumption wealth ratio *cay*. The x-axis is our ranking for how intermediated each asset class is. Panel A shows the pattern of coefficients on intermediary risk aversion, and Panel B shows this for household risk aversion. The red line is a linear regression fit through these points. An upward slope indicates more predictability in more intermediated asset classes, and vice versa. See text for more details.

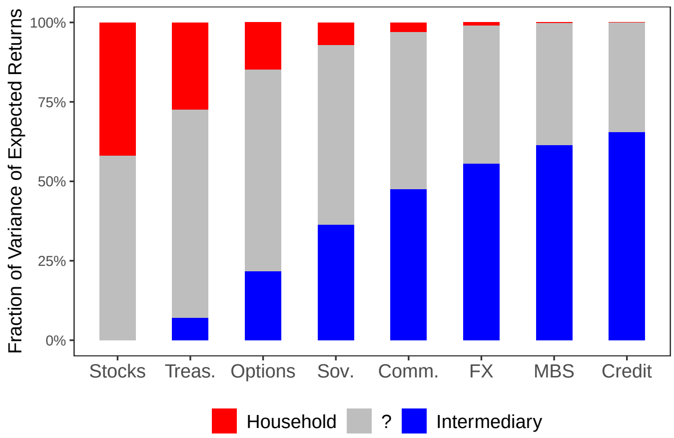

Figure 9. Decomposition of risk premium variation.
This figure plots lower bounds of variation in risk premia coming from households and intermediaries for each asset class using the pattern of predictability across the asset classes. See text for details.

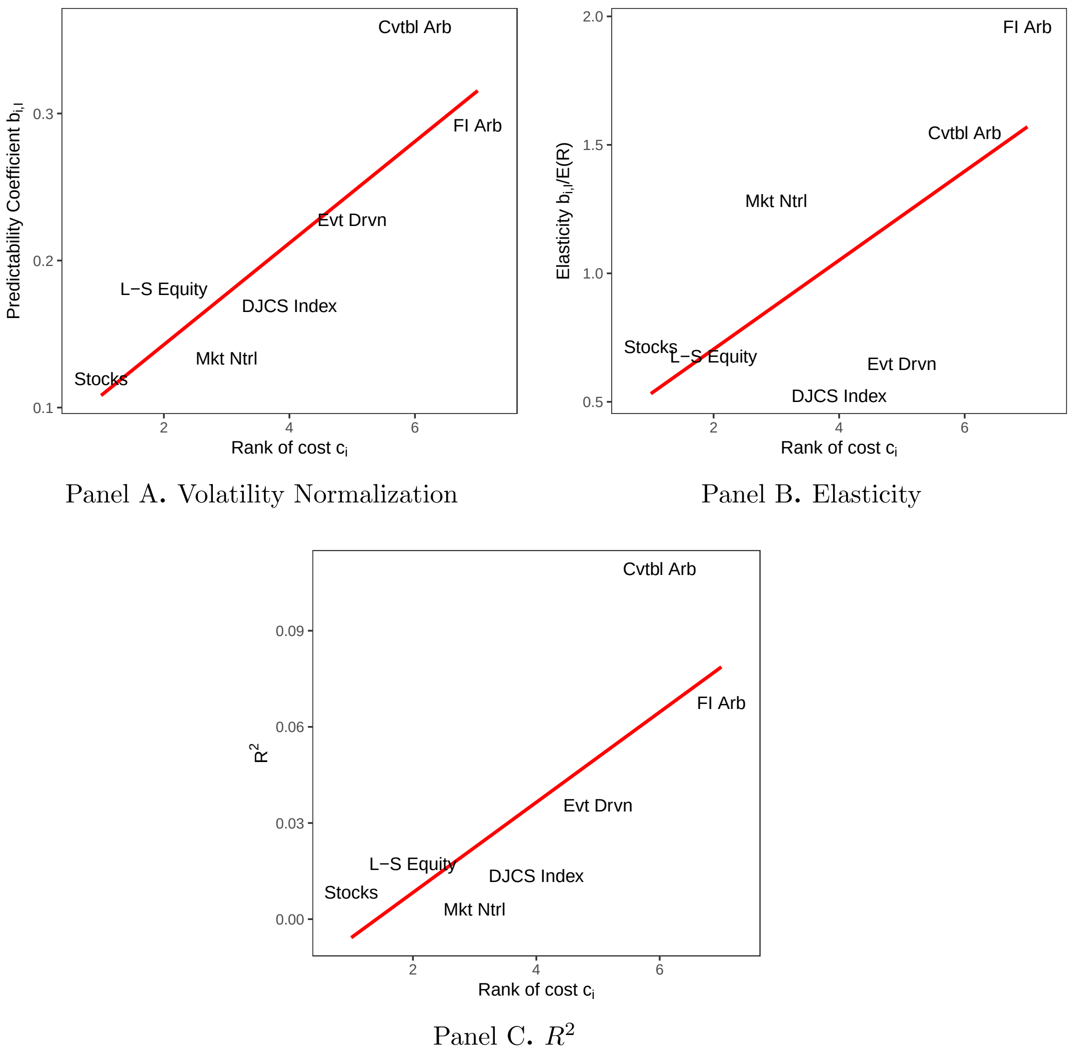

Figure 10. Hedge fund strategy returns.

This figure reports the behavior of risk premiums across stocks and hedge fund returns by category: long short equity, market neutral equity, the DJCS hedge fund index weighted across all hedge fund styles, event driven, convertible bond arbitrage, and fixed income arbitrage. The top panel runs: $r^{\sigma}_{i,t+1} = a_i + b_i \times \tilde{\gamma}_{I,t} + \epsilon_{i,t+1}$ and reports b_i across fund categories. Excess returns $r^{\sigma}_{i,t+1}$ are normalized by their full sample volatility. The lower right panel reports the R^2 in this predictive regression. The lower left panel gives the risk premia elasticity found by running $r_{i,t+k}/E[r_{i,t+k}] = a_i + b_i\tilde{\gamma}_{I,t} + \varepsilon_{i,t+k}$. The right hand side variable $\gamma_{I,t}$ that measures intermediary health is an equal weighted average of the AEM and HKM factors. See text for more details.