

Do Intermediaries Matter for Aggregate Asset Prices?

Valentin Haddad and Tyler Muir*

April 25, 2018

Abstract

Poor intermediary health coincides with low asset prices and high risk premia, but it is unclear how much fluctuations in intermediaries' health matter for aggregate asset prices rather than simply being correlated with aggregate risk aversion. We argue that relative predictability of more vs less intermediated asset classes by intermediary health allows to quantify how much variation in risk premia we can ascribe to intermediaries. Intermediary health should matter relatively more for assets that households are less willing to hold directly, whereas frictionless aggregate risk aversion should, if anything, exhibit the opposite pattern. We provide direct empirical evidence that this is the case and hence argue that intermediaries matter for a number of key asset classes including CDS, commodities, sovereign bonds, and FX. Our findings suggest that a large fraction of variation in risk premia in these asset classes is related to intermediary risk appetite.

*UCLA and NBER. We thank Nina Boyarchenko, Markus Brunnermeier, Mike Chernov, Itamar Drechsler, Andrea Eisfeldt, Zhiguo He, Francis Longstaff, Arvind Krishnamurthy, Adrien Matray, Alan Moreira, Matt Plosser, Pietro Veronesi, and participants at UCLA, the Utah Winter Finance Conference, the Wharton Conference on Liquidity and Financial Fragility, the ASU Sonoran Winter Conference, the LBS AQR Asset Management Conference, LSE, NYU, the San Francisco Fed and the New York Federal Reserve for comments.

1. Introduction

Periods of poor 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 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 provide an answer to this question by comparing variations in risk premia across more and less intermediated asset classes. Specifically, we run predictive regressions across asset classes on variables that proxy for the effective risk bearing capacity of intermediaries such as the book leverage measure of Adrian et al. (2014) and the market equity measure of He et al. (2017). Our measure of the degree of predictability is either the percentage change in risk premium relative to its mean, the change in risk premium relative to the volatility of the asset class, or the R-squared of the predictive regression. Specifically, we run regressions of the form

$$\hat{r}_{i,t+k} = a_i + b_i \gamma_{I,t} + \varepsilon_{i,t+1}, \quad i = 1, \dots, N, \quad t = 1, \dots, T \quad (1)$$

where $\gamma_{I,t}$ is an empirical proxy for intermediary risk-appetite and $\hat{r}_{i,t+k}$ is the return on asset class i , normalized by either the unconditional volatility or the mean of each asset

¹E.g., Adrian et al. (2014), Hu et al. (2013), Haddad and Sraer (2016), Muir (2017), He et al. (2017).

²He and Krishnamurthy (2013) and Brunnermeier and Sannikov (2014) are examples of models of such linkage.

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

class. We find relatively more predictability — larger values of b_i — for asset classes that are more specialized and more intermediated (e.g., CDS, currencies, commodities) and relatively less predictability for asset classes that are less specialized (e.g., stocks). We argue that these differences in predictability of more vs less intermediated asset classes provide a lower bound for how much intermediaries matter in each asset class.

Our fundamental argument is as follows: because variables that proxy for intermediary risk-bearing capacity are likely positively correlated with households willingness to bear risk (or more generally, frictionless aggregate risk aversion), predictive regressions using only one asset class do not help to uncover how much variation in risk premia we can ascribe to intermediaries. Comparisons across asset classes are useful. We argue that the link between risk premia and household risk aversion should have either a flat or decreasing pattern as we move towards asset classes where households are less active (for example, we show that under the null of a frictionless model this pattern is flat). Thus, the increase in predictability between more intermediated and less intermediated assets by measures of intermediary health provide a lower bound on how much intermediaries matter in each asset class. We use these differences to quantify how much intermediaries matter and find substantial variation in risk premia attributable to intermediaries in these asset classes. A similar exercise using proxies for the risk aversion of households provides a lower bound for their role as well. We use this approach to provide a decomposition of variation in risk premium attributable to intermediaries and households. For example, we find that we can attribute about 80% of variation in risk premia in CDS to intermediaries. Similarly, we can attribute about 30% of variation in risk premia of stocks to households. There is still a remaining fraction of variation for each asset class we can not assign to either based on our lower bounds.

We first clarify our argument in a CARA-normal model which households invest di-

rectly in asset markets as well as indirectly through intermediaries.⁴ There are two potential frictions in the model. First, households face quadratic costs of investing directly in some asset classes versus investing indirectly through an intermediary (e.g., the cost could be high for CDS and low for stocks). Second, we allow preferences of intermediaries and households to potentially not be aligned, potentially due to the presence of regulatory or incentive constraints. The presence of both frictions is important for intermediaries to matter. If all costs to investing in all asset classes are zero then households do not need intermediaries and will invest directly in asset classes on their own. If preferences between households and intermediaries are perfectly aligned, then the intermediary is a veil and will not matter for asset pricing.

In the model, we study the response of risk premia to a change in intermediary risk-bearing capacity, and we show this response is increasing in the cost of direct access from households. Intuitively, for asset classes with low cost of direct access, a drop in intermediary risk-bearing capacity is absorbed through quantities as households substitute indirect holdings for direct holdings of the asset. Since households can easily substitute or “step in” to these markets, there is little effect on prices. For assets with high cost of direct investment, the opposite is true. Because households cannot absorb the intermediary positions through direct holdings, prices for these assets fall and risk premia rise. Changes in household risk aversion have essentially the opposite pattern: a larger impact on less intermediated assets and smaller impact on more intermediated assets. This exercise shows how to overcome the empirical challenge of having a proxy for intermediary risk-bearing capacity which is positively correlated with household risk aversion by looking at the differential response of risk premia to the proxy, with the main prediction

⁴Our model is related in spirit to He and Krishnamurthy (2013), Brunnermeier and Sannikov (2014) but includes many assets, and does not assume households can't invest in the assets (this is equivalent to an infinite cost in our setting). Relatedly, Kojen and Yogo (2015) study how institutional demand affects individual stock prices but is not able to address the time-series of aggregate asset prices because their model does not feature or model the substitution of households direct vs indirect holdings.

being a larger response for more intermediated asset classes (ones for which the cost of investing directly by households is high).

We next state more carefully how we measure the degree of predictability across asset classes. Because we predict excess returns across asset classes with different average returns and volatilities, we need to normalize returns appropriately into the same units to make the coefficients comparable. To see why, suppose that one asset class was just a levered version of another asset class. Then any variable which predicts returns in the original asset class will naturally have a larger coefficient on the levered asset class. Normalizing asset returns by their volatility deals with this effect and also provides an intuitive economic interpretation of the coefficient as the degree of predictability relative to the assets volatility (closely related to the R-squared). However, this argument also highlights that we would like to normalize by the amount of risk in the asset class. We thus consider a more general normalization which divides by the unconditional expected return, hence the interpretation of the coefficient is the percentage change in risk premium as risk appetite changes. Because equilibrium expected returns reflect any factor risk exposures, this normalization takes care of any unconditional differences in risk (betas). While we prefer this normalization economically (and indeed show this in the context of our model), there is an empirical tradeoff that average returns are much harder to estimate than standard deviations. Thus, we consider both in our empirical results. We show that both produce qualitatively similar results, though all results are statistically sharper when using volatility due to the uncertainty associated with average returns.

Next, we need to make empirical choices on our proxies of intermediary risk bearing capacity $\gamma_{I,t}$ as well as on the asset classes we use and the degree to which intermediaries are active in each asset class. To measure risk-bearing capacity, we rely extensively on the existing literature on intermediary asset pricing which provides foundations for how risk appetite should be measured empirically and in our main specification we use 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). However, we show robustness to several alternative choices for our proxy of intermediary risk appetite. It is worth emphasizing that our paper does 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), but instead we focus on the identification challenge associated with taking these models to the data.⁵ Further, our goal is simply to obtain proxies that correlate with intermediary health, we do not require that they are “exogenous.”⁶

For our asset classes, ranked from least to most intermediated, we use stocks, bonds, options, sovereign bonds, commodities, foreign exchange, and CDS.⁷ We use several methods to arrive at our ranking. We use quantity and position data from Flow of Funds and the BIS to study the relative holdings of households vs institutions. We also use Value-at-Risk measures from intermediary 10-K’s which give us risk exposure in each asset class. Finally we use ETF expense ratios to gauge household cost of exposure to the asset classes.

We emphasize that we do not need any household risk aversion controls for our main argument. That is, in general, we take no stand on the behavior of household risk aversion. However, having a proxy for household risk aversion helps us on several fronts quantify the role of households for prices. In our framework risk premia respond with an opposite pattern across asset classes to the willingness to bear risk of households relative to intermediaries. We confirm this observation in the data which thus strengthens

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

⁶In fact in most of the theories they are a state variable that responds endogenously to more fundamental shocks.

⁷Some 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.

our mechanism by using the habits measure of Campbell and Cochrane (1999) and the cay measure from Lettau and Ludvigson (2001). 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 in some way: not all return predictors exhibit an increasing pattern as one moves to more intermediated asset classes.

Finally, we go through other possibilities that could explain our results and discuss limitations of our results. Most importantly, in our framework thus far we only consider allowing intermediary and household risk aversion to move (and use the data to separate these two) but other parameters in the model may also change. For example, the covariance of asset payoffs might change and be correlated with the other variables. We address this empirically by including proxies for changing covariances in our regressions (we include time-varying volatilities and betas in our regressions), but we also show that these time-varying covariances would also have to have a very unique factor structure to line up perfectly with our results – in particular it has to be that risk increases differentially for intermediated asset classes at the same time our proxy for risk appetite moves. We find no such relationship when we predict volatility or skewness rather than means of returns, thus we find no evidence that our proxies line up with measures of risk only for intermediated asset classes.

Still, it is also important to acknowledge that because of the joint hypothesis problem we can not make progress in answering how much intermediaries matter without first taking a stand on household behavior in a frictionless model. Our goal is to provide as flexible a model as possible for household behavior while still being able to make sharp empirical predictions. Specifically, our model of household behavior allows 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 freely move around), (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. Thus, our benchmark model is fairly flexible.

We then extend our results to hedge fund returns and note again a large degree of predictability by the intermediary risk aversion proxies for the returns to more complex hedge fund strategy returns – in particular convertible bond arbitrage and fixed income arbitrage – relative to the overall stock market predictability. These results support prior studies that intermediary capital matters for these asset prices (Mitchell et al., 2007; Hu et al., 2013). Again, we show that proxies for household risk aversion do not strongly forecast these returns. These results have the advantage that they directly predict the returns to which intermediaries invest.

This paper is the first attempt to lay out and try to tackle the challenge associated with intermediary asset pricing in terms of attempting to assess how much intermediaries matter. Our goal is to provide a simple framework to lay out the criteria for intermediaries to matter for asset pricing. Existing models with intermediation assume a single risky asset and focus on asset price dynamics in a crisis given the assumption that intermediation matters for this asset. While these papers motivate the empirical literature on intermediary asset pricing, they are insufficient to fully address the identification challenge. We then take a first empirical step at addressing these issues using the cross section and find support that intermediaries do in fact matter for many aggregate asset prices. Our results are important to understand the overall variation in risk premiums and to provide tests for intermediary based models of asset prices. Our results are useful for counterfactuals as well, for example if a given regulation is likely to impact intermediary risk appetite our framework can provide quantitative estimates for how risk premiums in each asset class may change.

Related literature. Our main finding is 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, 2016; He et al., 2017). The main difference is that this literature typically tests 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. First, the intermediaries' Euler equation always holds in our framework, and this is true regardless of whether either friction matters. Thus, only knowing whether the intermediary Euler equation holds on its own does not tell us whether, or how much, intermediaries matter for asset prices. For example, the intermediary Euler equation can hold but intermediaries just reflect household preferences, or the Euler equation can hold even with different preferences between intermediaries and households if the household can undo the intermediary choices at no cost. Intuitively, the Euler equation approach is about optimality – it only tells us whether intermediary first order conditions hold. The question in this paper is about whether (and how much) intermediaries matter in equilibrium. Only knowing their first order conditions hold will not tell us whether or how much prices will change if intermediaries become more risk averse. For example, if their preferences and constraints match those of the average investor, then prices need not change at all.

There is also intriguing “micro” evidence that intermediaries matter for particular individual asset prices at particular points in time. For example, Du et al. (2017), Siriwardane (2016), Fleckenstein et al. (2014), Lewis et al. (2017), Krishnamurthy (2010), Mitchell et al. (2007). For many additional examples, see Duffie (2010), Mitchell and Pulvino (2012), and He and Krishnamurthy (forthcoming). While these studies are important in documenting detailed price deviations related to intermediary risk-bearing capacity, it is often unclear what these results imply for the broad behavior of aggregate asset price movements.

2. Framework

We introduce a model of asset pricing with an intermediary. Households can invest directly or through the intermediary, potentially facing two frictions. First, investing directly is costly. Second, households do not control the investment decisions of the intermediary. We show how the interplay of these two frictions is what gives rise to a role of intermediation for asset prices. This simple theory helps understand the limitations to the interpretation of the existing evidence on intermediary asset pricing, but also guides the design of our empirical tests.

2.1 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 the equilibrium price of the assets and assume that all decisions take prices as given.

We assume that the household can invest in the assets in two way. First, the household can buy the assets directly, but at some cost. To do so, we assume the household faces a quadratic cost parametrized by the positive semidefinite matrix C to invest in the various risky asset. 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 γ_I . These two assumptions are volun-

tarily stylized, and we come back to them in more details later in this section. Figure 1 summarizes this setup.

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

$$\max_{D_I} D_I' (\mu - p) - \frac{\gamma_I}{2} D_I' \Sigma D_I. \quad (2)$$

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

$$\max_{D_H} (D_H + D_I)' (\mu - p) - \frac{\gamma_H}{2} (D_H + D_I)' \Sigma (D_H + D_I) - \frac{1}{2} D_H' C D_H. \quad (3)$$

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 (2) and (3) respectively. The market-clearing condition is

$$D_H + D_I = S. \quad (4)$$

2.2 Equilibrium Portfolios and Prices

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

$$D_I^* = \frac{1}{\gamma_I} \Sigma^{-1} (\mu - p). \quad (5)$$

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 of the risk aversion γ_I .

In contrast the household demand is:

$$D_H^* = (\gamma_H \Sigma + C)^{-1}(\mu - p) - (\gamma_H \Sigma + C)^{-1}(\gamma_H \Sigma)D_I. \quad (6)$$

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 is in general not one-to-one. Rather, it is given by

$$-\frac{\partial D_H^*}{\partial D_I} = (\gamma_H \Sigma + C)^{-1}(\gamma_H \Sigma). \quad (7)$$

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, this substitution is the identity. As the investment cost gets larger, the substitution rate converges to 0. If investing directly in the asset becomes too expensive, the household does not offset the decisions of the intermediary anymore.

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

$$\mu - p = \gamma_H \Sigma \left(\Sigma + \frac{1}{\gamma_I} C \right)^{-1} \left(\Sigma + \frac{1}{\gamma_H} C \right) S \quad (8)$$

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 $\left(\Sigma + \frac{1}{\gamma_I} C \right)^{-1} \left(\Sigma + \frac{1}{\gamma_H} C \right)$. This distortion encodes the potential effect of the intermediary on asset prices, through the impact of the parameter γ_I .

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

$$\gamma_I \neq \gamma_H \text{ and } C \neq 0 \quad (9)$$

This proposition states that 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, $-\partial D_H^*/\partial D_I \neq 1$.

Now that we have clarified the importance of our two frictions for the notion of intermediary asset pricing, we come back to more precise motivations for their presence, and relate to how they have been introduced in previous literature. Then, the next section discusses various empirical implications of this model. We explain why some already tested implications do not get exactly at this combination of conditions, and propose a novel empirical test which targets it.

2.3 Interpretation of the Frictions

Intermediary decisions. The first ingredient is that the intermediaries do not invest in a way that reflects the preferences of households. If this is not the case, intermediaries are just a veil. We represent this distinction by allowing the parameter γ_I to differ from γ_H . In practice, multiple reasons can explain that 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)). The presence of costs of financial distress, combined with a limited ability to raise capital also gives rise to a risk management policy specific to the institution. 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.

While these justifications explain a mismatch in investment policies at the micro level, the overall supply of intermediation could adjust so that there are just enough intermediaries to satisfy household's investment needs. One reason this would not be the case is that there are barriers to entry into the intermediation industry, or that raising capital to create an intermediary is difficult. Another reason might be that the private incentives of the managers of intermediaries to enter the market are not lined up with aggregate households' incentives. Haddad (2013) presents a model with free entry into intermediation and shows that even under such conditions, variation in intermediation technology or in aggregate uncertainty gives rise to fluctuation in the aggregate risk appetite of the financial sector.

In this paper, we do not take a stand on the precise micro foundations for this distinction in risk appetite. Instead we highlight this feature as being important for intermediaries to matter for asset prices and devise tests to uncover its presence.

Imperfect substitution. The second ingredient is that households do not offset changes in the decisions of intermediaries through direct investing, $-\partial D_H / \partial D_I \neq 1$. A simple motivation for this feature is that it is difficult for households to access some risky asset markets, for instance for some complex financial products. We materialize this force by

the quadratic cost of direct investing C . 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). This assumption would also generate no direct investing at all in most of the equilibrium. In contrast, a completely frictionless view of direct investing, $C = 0$, completely rules out a role for financial intermediaries. A benefit of our smooth parametrization is that it allows to control the difficulty for households to invest in risky assets, and explore its role empirically.

Other reasons can lead to an imperfect substitution of direct investing against intermediated investment. Households might be less able to manage portfolios of risky assets, making them effectively more risky. Eisfeldt et al. (2017) studies a model along these lines. 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.

3. Empirical Implications

We now consider in more details the implications of our framework. We are particularly interested in the relation between intermediaries and asset pricing. We first revisit two sets of approaches from the existing literature, and highlight their limitations in isolating the impact of intermediaries of asset pricing. We then propose a test to better discriminate whether intermediaries affect asset prices or not. For a related review of empirical work on intermediary asset pricing, see He and Krishnamurthy (forthcoming).

3.1 Euler Equation Approach

A classic approach to study household's optimization in financial markets is by studying whether their Euler equation holds. This corresponds to asking whether their marginal utility of consumption 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.

In our setting, intermediaries have frictionless access to the risky asset market. Therefore their Euler equation holds. Actually, the portfolio of intermediaries is always mean-variance efficient — see Equation (5). It implies that it forms a pricing kernel: writing R_I the excess return on the intermediary risky portfolio, then for any risky asset excess return R_i , we have:

$$\mathbb{E}[R_i] = \beta_{iI} \mathbb{E}[R_I],$$
$$\text{where } \beta_{iI} = \frac{\text{cov}(R_i, R_I)}{\text{var}(R_I)}.$$

Several papers have studied empirically 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.

However, it is worth noting here that the intermediary Euler equation *always* holds in our setting. This result only relies on our specification of the intermediary's demand for risky assets, determined by its objective function. The empirical success of the intermediary Euler equation therefore only validates the specification of a frictionless demand function for intermediaries. In particular, it holds independently of whether intermediaries matter for asset prices.

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. More generally, even if the household risk aversion is left as a free parameter, the CAPM does not hold unless $(\gamma_H \Sigma + C)^{-1} \left(\Sigma + \frac{1}{\gamma_I} C \right)$ is proportional to the identity matrix. This corresponds either to cases where intermediaries do not matter or where the cost of investing C is exactly proportional to the variance Σ .

Going back to Hansen and Singleton (1983), there is a long tradition of evidence inconsistent with particular specifications of the Euler equation for households. It remains 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, Greenwood 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 model that we expect to provide sharper empirical tests, more directly focused on intermediaries.

3.2 Time-Series Predictability Approach

A second approach consists in studying the relation between characteristics of intermediaries and future returns in the time series. There are two broad ways to do so. We discuss them in the context of our model with only one asset.

The first approach consists in assuming that intermediaries have a stable demand

function, that is that γ_I does not change over time. In this case their equilibrium demand directly reveals the risk premium: $(\mu - p)_t = \gamma_I \sigma^2 D_{I,t}$. When intermediaries decide to bear more risk, this reveals a higher market risk premium. Haddad and Sraer (2016) apply this idea by relating the exposure of banks to interest rate risk to expected returns for Treasuries. Similarly, Diep et al. (2016) relate the sign of risk premia on mortgage-backed securities to the direction of the exposure of intermediaries. This approach is based on fluctuations in prices unrelated to changes in the fundamental characteristic of intermediaries, and therefore does not get at the causal effect of changes in intermediary conditions on prices.

The second approach considers implication of changes in intermediary risk appetite. By contrast to the first approach it considers the implications of shifts in intermediaries' demand for risky assets rather than movements along their demand curve. In our model, an decrease in intermediary risk appetite, a higher γ_I , corresponds to a higher risk premium:

$$\frac{\partial(\mu - p)}{\partial\gamma_I} \geq 0 \tag{10}$$

with strict inequality if and only if intermediaries matter for asset prices. Indeed, if intermediary have less risk appetite, they want to decrease their positions in risky assets. If there are direct investment costs, households do not offset this lower demand completely. The risk premium must increase to go back to an equilibrium.

Various papers implement this idea by a regression of future returns on measures of intermediary risk appetite, for instance He et al. (2017), Chen et al. (2016), or Muir (2017). The major limitation of this approach is the following: in order to interpret a significant coefficient as saying that intermediaries matter for risk premia, we need to assume there is no contemporaneous change in household demand. That is, if γ_I and γ_H are positively correlated, then it is unclear whether the predictability coming from our

empirical measure of γ_I is driven causally by intermediaries or whether it simply reflects an change in broad risk aversion. 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. In the language of the model, if both γ_I and γ_H are positively correlated, we can not say whether intermediaries matter from an individual predictive regression.

3.3 Our Approach: Time-Series Predictability Across Assets

Our test builds on this last approach, but aims at disentangling the two conflicting explanation for high risk premia: high overall risk aversion γ_H or high intermediary risk aversion γ_I . To do so, we compare expected returns across asset classes with different direct costs of ownership. Intermediary health matters more for assets that households cannot buy directly, whereas household risk aversion matters less for those assets.

To illustrate this, consider a situation where the asset returns are uncorrelated across asset class, and the cost matrix is diagonal. We note each asset i and c_i its cost of direct holding. In this case we obtain the following result, reflecting the intuition above.

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.*

Figure 2 illustrates this comparison. To understand this proposition, consider the elasticity of the risk premium to changes in household and intermediary risk aversion:

$$\frac{1}{\mu_i - p_i} \frac{\partial(\mu_i - p_i)}{\partial \log(\gamma_I)} = \frac{c_i}{\gamma_I \sigma_i^2 + c_i} \quad (11)$$

$$\frac{1}{\mu_i - p_i} \frac{\partial(\mu_i - p_i)}{\partial \log(\gamma_H)} = \frac{\gamma_H \sigma_i^2}{\gamma_H \sigma_i^2 + c_i} \quad (12)$$

Both of these elasticities are positive, with a role for intermediary risk aversion if and only 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 or flat for household risk aversion. It is increasing for intermediaries because households offset their trades less in asset classes that are harder to invest in directly. In contrast, household reduce their positions less aggressively in asset classes for which it is harder to invest directly when they become more risk averse.

This distinction suggests a test that isolates the role of intermediary risk aversion. Our measures of intermediary risk appetite are positively correlated with household risk appetite. However, the only way they can comove more with risk premia for higher cost of direct holdings is if they capture at least partially intermediary risk aversion and it has a causal impact on asset prices. In other words, the health of financial intermediary is more related to premia for assets which are more difficult for households to invest in only if intermediaries matter for asset prices.

Focusing on elasticities rather than directly 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 that will naturally tend to move more in absolute magnitude with the various risk appetite. Scaling by a baseline level of risk premium cleans out this effect to focus on the role of the financial frictions. In the next section we implement this test empirically.

4. Empirical Results

Having presented the model and discussed the main empirical challenges to assess whether intermediaries matter for broad asset prices, we now provide the main tests of the paper.

4.1 Data Description

We use asset returns and intermediary state variables that are common in the literature. We use excess returns on the market, commodities, CDS, options, sovereign bonds, Treasury bonds, and the currency carry trade, where we take excess returns over the 3 month T-bill where appropriate. These choices are motivated by looking at many markets where we think intermediation may matter. We start by using these asset returns provided by He et al. (2017). For CDS, options, sovereigns, and commodities we take the equal weighted average in each asset class. Treasury bonds (labeled henceforth as just bonds) are longer term Treasury bond returns over the 3 month T-bill rate. CDS is an average across maturities and credit risk. Commodities are simply the equal weighted average across all commodities available in the HKM dataset. Some of the 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. 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 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 large (He et al., 2017).⁸

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.

⁸One 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.

(2017) (HKM), each of which has been argued to capture intermediary distress and each of which is linked to risk premiums. We also consider the noise measure by Hu et al. (2013). 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. Each of these variables has been argued theoretically, and empirically, to capture intermediary distress and risk bearing capacity. Again, we emphasize in our framework that we don't 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 *cay* Lettau and Ludvigson (2001) and the Michigan consumer sentiment measure. We do not take a strong stand on these variables in terms of corresponding perfectly to household risk aversion, though in robustness tests we do consider whether including them in our regressions affects our results. This is important 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.

4.2 Discussion of Costs and 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 crucial for our empirical design because we exploit that assets that are more specialized (held by intermediaries) are will respond more to intermediary changes. In our model this is captured by costs C for household to hold the asset directly, though we emphasize that this may not literally be a physical cost or a literal cost of accessing / trading in a given market, but it could also be costs of complexity of the asset. That is, even if a household could costlessly trade a credit default swap, the

complexity of the instrument may preclude them from doing so. Because of this, we take a multi-pronged approach to identifying which assets are more intermediated: we look at holdings data, we look at volume of trade accounted for by institutions (particularly focusing on dealers), and we also look directly at costs faced by households (we look at 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 the same ranking of which asset classes are more or less intermediated. We report our ranking in Table 10. 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, and CDS. We emphasize that we take a data driven approach to conduct these rankings, though we don't take an overly strong stance on this exact ordering (e.g., one could likely swap some of the adjacent pairs) and we return to this issue when we discuss the results.

Holdings and volume data

We first study holdings data from Flow of Funds (FoF), and the Survey of Consumer Finance (SCF). In FoF we take holdings and stocks, bonds, and foreign and corporate bonds 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 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 bonds and far fewer corporate and foreign bonds. The SCF data gives an alternative way to measure households, and provides a few 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. We confirm the same result with the SCF data. However, neither of these sources list household holdings of more specialized asset classes such as CDS.

Our next source of data is the BIS data on derivatives semiannual report.⁹ 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, though we 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.

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 very high risk or high beta stocks, perhaps the fractions above would miss this. In the model, we really 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 (JP Morgan, Morgan Stanley, Goldman Sachs, etc) and we use data from 2016 for our analysis. 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. The value-at-risk typically reports tail risk – for example, they provide a dollar amount which losses would not be expected to exceed 95% of the

⁹See <https://www.bis.org/statistics/derstats.htm>.

time. We convert this 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 betas. 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 bonds outstanding. This ranking thus gives similar results to just using the position data above.

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.¹⁰ We emphasize that over much of our sample, these products would not have been available – and indeed we are not aware of any data on costs for households to invest in some of the more sophisticated assets in our study. Nonetheless, the approach of 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.

We take the average expense ratio by asset class as our measure of the cost. Our asset classes are Stocks, Government Bonds, Emerging Markets Bonds (our best proxy for the sovereign bonds used in this study), Currency, Commodities, and Volatility. We use volatility to proxy for our option straddle strategy which is a bet on volatility, though we note much of the Volatility ETFs are trading VIX futures directly, and these strategies are

¹⁰<http://etfdb.com/etfdb-categories/>

different though they are exposed to the same underlying risks (Dew-Becker et al., 2017). 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: 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 in 2014 as the first CDS ETFs. Prior to these ETFs, household exposure to CDS would have been nearly impossible as one needs an ISDA master swap agreement to trade CDS.

The expense ratios give us a sense of the costs faced by households to obtaining exposure to each of these assets – however, we emphasize two caveats. First, again, the deep quantity that matters is the willingness of households to step in and invest if there is a shock to institutions demand. that is, we care about the substitution of direct and indirect demand. It may be that, even though physical costs to investing in CDS become low in the future, this household substitution remains low because of the complexity of these contracts.

We need to normalize the expense ratios in each asset class. 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. We show both results in the appendix.

Our ETF ordering implies: stocks, bonds, sovereign bonds, currencies, commodities, options, and CDS. This is largely consistent with our main ranking, we return to the issue of equity options appearing more intermediated further below.

4.3 Empirical Test

Our model guides us to run the following regressions:

$$r_{i,t+k}/E[r_{i,t+k}] = a_i + b_i x_t + \varepsilon_{i,t+1}, \quad i = 1, \dots, N, \quad t = 1, \dots, T \quad (13)$$

where we consider k to be 1 quarter and 1 year (that is, we forecast future 1 year or 1 quarter returns).

Under the null hypothesis that intermediaries don't matter for any of the aggregate asset returns considered in our tests, we have that b_i should be the same for any asset i when we use the predictive variables x_t that are proxies for changes in intermediary risk aversion.

The key feature of the alternative hypothesis, is that the magnitude of b_i should be larger for more intermediated assets i when we use predictive variables x_t that proxy for intermediary risk aversion shocks. If we document such a differential response, then we can reject the null that intermediaries do not matter for some asset classes (those with the highest c).

Further, the exercise allows us to deal with the concern that our x_t variables – which proxy for the health or distress of the financial sector – may also be correlated with γ_H or household risk aversion shocks. Because we use the differential response of risk premia to the shocks, we are able to assign some of the variation to intermediary risk aversion shocks. This is because if all variation was only household risk aversion shocks, we would not see the differential response in the cross-section. Moreover, the model makes the prediction that any differential response should be highest for the most intermediated assets for which we believe the costs c are greatest.

Finally, it is worth noting here that in the lowest c asset class, we can not separate household vs intermediary risk aversion. We can not say in this asset class that the world is one in which c is above zero and intermediary risk aversion matters, or that the world

is one in which $c = 0$ but household risk aversion moves in the right way to make the risk premia of this asset class move. Instead, our tests only apply to the unique predictions of the differential response of risk premia.

4.4 Results

We run predictive regressions in each asset class

$$r_{i,t+1}/E[r_{i,t+1}] = a_i + b_i x_t + \varepsilon_{i,t+1}, \quad i = 1, \dots, N, \quad t = 1, \dots, T \quad (14)$$

We report b_i , where we adjust our standard errors to take into account the uncertainty in the mean of each return as well, i.e., that $E[r_{i,t+1}]$ is estimated and not known. We do this using bootstrap with block length of 8 quarters to deal with autocorrelation of predictor variables.

The results are given in Table 1. We focus on Panel A, the quarterly return results, though we point out that these results typically carry through when using overlapping annual returns in Panel B. We can see larger (in absolute value) and more statistically significant coefficients of the alternative asset classes relative to stocks. For the stock market, the intermediary state variable is not quite significant, whereas the coefficient is negative and strongly significant for all other asset classes.

An alternative way to gauge the degree of predictability by our intermediary state variable across asset classes is to look at the R^2 from the predictive regressions for each asset class. This is another intuitive metric to see if there is more predictability for more intermediated assets. While intuitive, it turns out this measure is not quite as direct as the elasticity measure from the perspective of our model. In the next subsection, we show that our model does in fact say that – all else equal – we should see higher R^2 values for more intermediated assets in response to changes in intermediary risk aversion, justifying this alternative metric.

We find that the regression R^2 is lowest for stocks at 1.6%. All other asset classes have larger predictive R^2 s with the exception of bonds which, at 1.5%, is the same as that of stocks. We show that a prediction of our model is that the least intermediated assets should indeed have lower R^2 values as well. Some of the R^2 s are notable: CDS features a 35% R^2 whereas sovereigns, commodities, and FX are 15%, 5% and 3% respectively, all well above that of stocks. Again, this is consistent with a higher degree of predictability for more intermediated assets.

Next we consider a different normalization that normalizes each of the returns by their volatility, rather than their mean. Specifically, we run

$$r_{i,t+1}/\sigma[r_{i,t+1}] = a_i + b_i x_t + \varepsilon_{i,t+1}, \quad i = 1, \dots, N, \quad t = 1, \dots, T \quad (15)$$

This normalization has some disadvantages – namely that its justification requires additional assumptions in our model about covariances – but also some significant advantages as well – namely that it does not require estimating average returns for the normalization which are notoriously noisy. We show in the next section that this normalization also makes sense in the context of our model when the asset payoffs are orthogonal, and in particular it is identical if Sharpe ratios are roughly equal across asset classes. We note also that it takes care of basic leverage scaling effects – that is if one asset class was simply a levered version of another, this normalization will adjust for this. The interpretation of the coefficient maps directly to the regression R-squared, in particular because we normalize our predictor variable x_t to have unit standard deviation. Thus, the coefficient captures the ratio of the volatility of the conditional expected return relative to the return volatility.

The results are given in Table 2. This normalization makes our main result significantly stronger. That is, the coefficients on alternative asset classes are much larger than they are for stocks (again, in absolute value). The results would suggest that, relative

to stocks, household substitution is lowest for CDS, followed by sovereign bonds, with bonds, commodities, options and FX all around the same level (with coefficients about 3 times as large as stocks in absolute value).

The empirical results are best summarized by Figure 3 which plots the predictive regression coefficient (top panel) and R^2 (bottom panel) when we normalize returns by dividing by the average excess return. The middle panel plots the coefficient when we normalize instead by the assets' volatility rather than its average return. Again, generally speaking, we can see lower coefficients and lower R-squared values for stocks relative to the alternative asset classes. Our results consistently suggest that intermediaries matter the most for CDS markets. The other assets depend on the precise statistic we analyze, but generally we see intermediaries mattering strongly for sovereign bonds, commodities, options, and FX. Treasury bonds produce more mixed results.

Comparison to aggregate risk aversion variables

Next we consider the implication of our model that this pattern of differential coefficients should only apply to intermediary risk bearing capacity variables and not "generic" risk aversion variables. In fact, in our model, aggregate risk aversion shocks should not differentially affect intermediated assets. In Table 4 we re-run our predictive regressions but we include the *cay* variable of Lettau and Ludvigson (2001) which has been argued to capture aggregate effective risk aversion of a representative agent. Notably, we see none of the same patterns documented for our intermediary state variable. In particular, for quarterly predictive regressions, the coefficient in predicting stock returns is now higher than the coefficient on any other asset class (the only exception being options where the coefficient is just slightly higher than that for stocks). Thus, this variable which is known to predict returns and has been argued to proxy for aggregate risk aversion does indeed look like an aggregate risk aversion state variable. This also suggests there is nothing inherently mechanical in our intermediary state variable predicting returns with the specific

pattern we document. In unreported results, we find similar effects when we replace *cay* with the Michigan consumer sentiment forecast, again a variable that arguably captures aggregate risk aversion rather than intermediary health.

We take these results as supportive of our main conclusion: that intermediary specific state variables should have a differential effect on more intermediated assets. While our main test does *not* rely on identifying or controlling for aggregate risk aversion (in fact, the whole point of our test is that it avoids such measurement), it is nevertheless comforting that aggregate risk aversion proxies do indeed appear to line up with risk premiums as predicted by the model. Figure 4 summarizes these results graphically.

Panel Regressions

We next run our main specification as a panel regression. This allows us to test coefficient differences across asset classes by including an interaction term capturing the degree of intermediation. We pursue this approach in Table 5. Specifically, we test elasticity differences between stocks and other assets classes by including an interaction term for degree of intermediation $\hat{r}_{i,t+1} = a_i + b_1\gamma_{I,t} + b_2INT \times \gamma_{I,t} + \varepsilon_{i,t+1}$ and reports the coefficients. Here *INT* captures how intermediated the asset class is. First, we define *INT* to equal 1 for all asset classes besides stocks and bonds (the two least intermediated according to our rankings). Next, we also include options. Finally, we define *INT* to equal the rank on a scale from 0 to 1 of all seven of our asset classes. In addition, we report a similar interaction term where we use the habit measure for a proxy of γ_H . We find that the coefficient of the interaction of our intermediary risk aversion variable with degree of intermediation is strongly positive regardless of the ranking used, which tells us that the coefficient is significantly larger for more intermediated assets.

To quantify how much intermediaries matter, we use the coefficient $b_2 \times INT$ for the case where we rank assets on a scale from 0 to 1. Because this coefficient already differences out the level of predictability from the least intermediated assets, it provides

our lower bound for how much intermediaries matter. For each asset class, we take $(b_2 \times INT)^2$ as the variance in risk premia attributable to intermediaries (recall we normalized the variance of our predictor to 1). We also run the same regressions for households $\hat{r}_{i,t+1} = a_i + b_1\gamma_{H,t} + b_2INT \times \gamma_{H,t} + \varepsilon_{i,t+1}$ using the habits measure to proxy for γ_H . Now, the lower bound of variation in risk premia from households is measured to be zero for the most intermediated assets, hence we take $b_2(INT - 1)$. Finally, we run a panel regression with both measures γ_I, γ_H jointly (plus interaction terms for each) to assess the total variation in risk premia in each asset class. Our decomposition then looks at the fraction of the lower bound in variation in risk premia we can attribute to intermediaries and households as a fraction of this total variation when all measures are included. Figure 6 gives these results. There is a large degree of gray area in the graph which we are not able to attribute to either. This decomposition suggests that we can attribute at least 80% 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 30% 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 often still quite large and is not easily attributable to either category.

Hedge Fund Returns

We augment our main results with indices on hedge fund returns from Dow Jones Credit Suisse. We argue that hedge fund returns contain returns of specialized strategies and asset classes that will respond more to intermediary health than other assets. We run our predictive regressions again with stocks on the left 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. We argue that equity strategies are likely more accessible to households (e.g., some quant strategies in equities like value and momen-

tum 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 (see Mitchell and Pulvino (2012), Hu et al. (2013)). Event driven is likely in the middle (e.g., merger arbitrage), as is the index of all hedge funds which is weighted by AUM under each asset class.

Figure 5 summarizes our results. We find that predictability is higher for all hedge fund strategies compared to stocks, consistent with our main hypothesis that these constitute more specialized strategies that households would have difficulty investing in. Within hedge fund strategies, we also find convertible bond arb, fixed income arb, and event driven respond more to intermediary health, again consistent with the idea that these appear relatively more specialized. Thus, the results are consistent with our hypothesis using separate data on returns and thus strengthen our main findings. These data also avoid some drawbacks of our main results on using an unbalanced panel with a shorter sample (all these returns begin in 1994).

4.5 Robustness and additional results

4.5.1 Alternative statistics in the model

We showed, empirically, that in addition to meaningful variation in risk premia elasticities, there is also meaningful variation in risk premia across assets when normalized by volatility of the asset, and also that there is meaningful variation in R^2 s from return predictability regressions. While we argued these are intuitively appealing from the perspective of our model, we now formalize the exact prediction of these objects in the model.

Volatility Normalization

We show how our model implies an alternative normalization for our predictive regressions where we normalize returns by volatility rather than means. Essentially, this

will be true when Sharpe ratios are equated across asset classes, which will generally happen if the assets are uncorrelated and the supply of the assets are proportional to their inverse volatility $1/\sigma_i$ – this assumption means the quantity of risk $\sigma_i * S_i$ is the same across the asset classes. Empirically, we run $r_{i,t+1}/\sigma[r_{i,t+1}]$ on the left hand side rather than $r_{i,t+1}/E[r_{i,t+1}]$. One major advantage of this approach is that it avoids estimating the mean return for each asset class which introduces additional noise. However, this comes at a cost of having to assume assets are uncorrelated so that the covariance matrix Σ is diagonal.

Using the elasticity equation from our model, we have

$$\epsilon_{\mu_i} = \epsilon_{\gamma_I} \frac{c_i}{\gamma_I \sigma_i^2 + c_i} + \epsilon_{\gamma_H} \frac{\gamma_H \sigma_i^2}{\gamma_H \sigma_i^2 + c_i} \quad (16)$$

We multiply both sides by μ_i/σ_i , and then use the equilibrium relationship between μ_i and σ_i^2 (assuming Σ is diagonal) as well as the assumption that quantity of risk is equal ($\sigma_i S_i$ is constant), to obtain

$$\epsilon_{\sigma_i} = \left(\epsilon_{\gamma_I} \frac{\gamma_I c_i (\gamma_H \sigma_i^2 + c_i)}{(\gamma_I \sigma_i^2 + c_i)^2} + \epsilon_{\gamma_H} \frac{\gamma_H \sigma_i^2}{\sigma_i^2 + c_i / \gamma_I} \right) \quad (17)$$

We again have in this case that the coefficient multiplying ϵ_{γ_I} is increasing in c_i and hence should be larger for more intermediated assets, while the coefficient multiplying ϵ_{γ_H} is declining in c_i and hence should be smaller for less intermediated assets. These predictions are stronger than those shown before, but force us to make assumptions that Σ is diagonal which is unappealing.

R-squared Predictions

Next we justify looking at differential R-squared values across asset classes as an alternative way to asses the relative degree of predictability

We use a Taylor series approximation of our main equation for the risk premium

$$\mu - P = f(\gamma_I(x), \gamma_H(x)) \quad (18)$$

$$R_{t+1} = f(\gamma_I(x_t), \gamma_H(x_t)) + \varepsilon_{t+1} \quad (19)$$

$$\text{var}(R_{t+1}) = f'(x)^2 \text{var}(x) + \text{var}(\varepsilon_{t+1}) \quad (20)$$

Assuming without loss of generality that we standardize the variance of the shock, x , so $\text{var}(x) = 1$, and then using that $f'(x) = (\mu - P)\epsilon_\mu$, we can rearrange to obtain

$$R^2 = \frac{(\mu - P)^2 \epsilon_\mu}{(\mu - P)^2 \epsilon_\mu + \sigma^2} \quad (21)$$

$$R^2 = \frac{1}{1 + \epsilon_\mu^{-1} \sigma^2 / (\mu - P)} \quad (22)$$

Thus the R^2 will be increasing in $\epsilon_\mu \frac{(\mu - P)^2}{\sigma^2}$ which is again increasing in c . This again means that the R^2 should be higher for assets that the household will be less willing to buy directly.

4.5.2 Robustness of empirical results

We consider alternative stories and various alternative specifications for our main results.

In Table 6 we consider alternative ways to proxy for intermediary risk aversion or risk bearing capacity. The first, in Panel A, uses the log levels of the AEM and HKM factors (again, in levels we average the two after standardizing them). In Panel B, we instead use the Gilchrist and Zakrajek (2012) (GZ) spread 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.

Table 7 shows results when we split our intermediary health measure into the HKM and AEM components separately. Panel A gives our main result using the annual log changes of each measure (as we do in our main result) while Panel B shows results using

the log levels of each variable instead of changes. We find that, generally, both measures contribute to our main result though generally results are slightly stronger for the AEM measure.

Next, we consider time-varying variances and covariances as an explanation for our results. That is, in our main specification, we implicitly assume that the covariance matrix, Σ remains constant. One potential cause for concern is if Σ is changing in ways that are correlated with our intermediary risk-appetite proxies. More specifically, this is only a concern if our variable is correlated with *relative* changes in covariances only for the intermediated assets (for example, a common volatility factor that scales all asset volatilities up and down proportionally does not change our conclusions because we identify off relative changes; similarly, random variation in asset class volatility would not affect the results, only coordinated changes which are not proportional to each other but raise volatility in proportion to our factor would explain our results).

Specifically, we now include lagged individual factor volatilities in all of our regressions as well as lagged conditional market betas, following the work of Lewellen and Nagel (2006). This deals with the issue of time-varying Σ provided that only volatilities and conditional market betas change but not betas with some omitted factor. Table 8 includes these changes in risk as controls for our main results. We capture changes in volatilities and betas using lagged volatility over the previous 12 quarters (3 years) and lagged market beta over the previous 20 quarters (5 years). We use a slightly longer window to estimate conditional betas because we find short window betas are particularly noisy, though the choice of these windows does not affect our results. We include $\hat{\sigma}_{i,t-1}$ and $\hat{\beta}_{i,t-1}$ as controls on the right hand side and find that including these does not dramatically affect our main results.

We acknowledge these proxies will not be perfect in controlling for time-varying risk, but we also point out that a time-varying risk story needs to be very specific to explain our

results. That is, we would not find our main result is time-varying risk moved independently across asset classes, nor would we find it if time-varying risk moved in lock step across asset classes. Instead, what would affect our main conclusions is if time-variation in risk affected the intermediated asset classes by more than the non-intermediated asset classes and that this time-varying proportion in risk exposure exactly lined up with our measure of intermediary health.

Table 9 studies our main result across subsamples. We show results only using data from 1990 onwards, and results that exclude the 2007-2009 financial crisis period. Our main results are generally not changed across these subsamples.

5. Literature review

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

We extend, but also simplify, many models of intermediary asset pricing (He and Krishnamurthy (2013), He and Krishnamurthy (2012), Brunnermeier and Sannikov (2014), Danielsson et al. (2011), Adrian and Shin (2014), Eisfeldt et al. (2017)) but with a goal of interpreting empirical work rather than providing a theory of frictions and theory of intermediation that is micro founded.¹¹ That is, our paper offers no theory of what determines intermediary risk bearing capacity as is done in much of the literature. The main difference with our model is to allow for the possibility of direct investment by households at a cost, and to allow this cost to vary across assets. Whether this is actually a cost as we have modeled it is somewhat irrelevant, what is crucial is the substitution rate of

¹¹This literature fits into earlier models with a financial sector as in Bernanke et al. (1996), Kiyotaki and Moore (1997)

households demand functions to intermediary demand.

Our model allows us to speak to macro asset pricing studies that link intermediary balance sheets to risk premia (Adrian et al. (2014), Haddad and Sraer (2016), He et al. (2017)).¹² However, we discuss the limitations of these papers in saying whether or not intermediaries “matter” for asset prices, and use our model to come up with better tests to distinguish this from the alternative frictionless view. See also Santos and Veronesi (2016) as an example of a model where intermediary balance sheets and leverage relate to risk premia in equilibrium but the economy remains frictionless.

We also relate to “micro” studies which study 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 affect 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. This 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) provides a host of similar examples, and has a model related to our to explain these facts.¹³

6. Conclusion

We propose a simple framework for intermediary asset pricing. Two elements shape if and how intermediaries matter for asset prices: how they make investment decisions (preference alignment), and the extent to which final investors offset their decisions by direct trading (substitution). We show that existing empirical evidence has *not* provided

¹²See also Chen et al. (2016), Hu et al. (2013), Muir (2017), Pasquariello (2014), Baron and Xiong (2017).

¹³See also Lou et al. (2013).

direct evidence that intermediaries matter for aggregate asset prices and we discuss the specific reasons why. We then provide a simple test: a sufficient condition for intermediaries to matter for asset prices is that the elasticity the risk premium of relatively more intermediated assets responds more to changes in intermediary risk appetite. We provide direct empirical evidence that this is the case and hence argue that intermediaries matter for a number of key asset classes including CDS, FX, options, and commodities.

References

- Adrian, Tobias, Erkki 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.
- Baron, Matthew, and Wei Xiong, 2017, Credit expansion and neglected crash risk, *The Quarterly Journal of Economics* 132, 713–764.
- Bernanke, Ben, Mark Gertler, and Simon Gilchrist, 1996, The financial accelerator and the flight to quality, *Review of Economics and Statistics* 78, 1–15.
- 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.
- Chen, Hui, Scott Joslin, and Sophie X Ni, 2016, Demand for crash insurance, intermediary constraints, and risk premia in financial markets .
- 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, 2016, Prepayment risk and expected mbs returns .
- Du, Wenxin, Alexander Tepper, and Adrien Verdelhan, 2017, Deviations from covered interest rate parity, Working paper, National Bureau of Economic Research.
- 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, Working paper, National Bureau of Economic Research.

- 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.
- Gilchrist, Simon, and Egon Zakrajek, 2012, Credit spreads and business cycle fluctuations, *American Economic Review* 102, 1692–1720.
- Greenwald, Daniel L., Martin Lettau, and Sydney C. Ludvigson, 2014, The origins of stock market fluctuations, *working paper* .
- Haddad, Valentin, 2013, Concentrated ownership and equilibrium asset prices, *working paper* .
- Haddad, Valentin, and David Sraer, 2016, The banking view of bond risk premia, Working paper, Citeseer.
- 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, 2012, A model of capital and crises, *The Review of Economic Studies* forthcoming.
- He, Zhiguo, and Arvind Krishnamurthy, 2013, Intermediary asset pricing, *The American Economic Review* 103, 732–770.
- He, Zhiguo, and Arvind Krishnamurthy, forthcoming, Intermediary asset pricing and the financial crisis, *Annual Review of Financial Economics* .
- Hu, Grace Xing, Jun Pan, and Jiang Wang, 2013, Noise as information for illiquidity, *The Journal of Finance* 68, 2341–2382.
- Kiyotaki, Nobuhiro, and John Moore, 1997, Credit cycles, *Journal of Political Economy* 105, pp. 211–248.
- Koijen, Ralph SJ, and Motohiro Yogo, 2015, An equilibrium model of institutional demand and asset prices, Working paper, National Bureau of Economic Research.
- 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.
- Lewellen, Jonathan, and Stefan Nagel, 2006, The conditional capm does not explain asset-pricing anomalies, *Journal of Financial Economics* 82, 289 – 314.
- 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.
- Muir, Tyler, 2017, Financial crises and risk premia, *The Quarterly Journal of Economics* .
- Pasquariello, Paolo, 2014, Financial market dislocations, *The Review of Financial Studies* 27, 1868–1914.
- Santos, Tano, and Pietro Veronesi, 2016, Habits and leverage, Working paper, 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* .
- Siriwardane, Emil, 2016, Concentrated capital losses and the pricing of corporate credit risk, *working paper, Harvard University* .

7. Appendix

We generalize the results in the main text by introducing the simplest price-theoretic framework of an asset market that includes an intermediary. This setting highlights the two basic forces determining the role of intermediaries for asset prices. The first element is their demand for the asset, how they make investment decisions. The second element is how final investors substitute between holding the assets through the intermediary and directly. We then flesh out a particular model that fits this framework before discussing alternative foundations for those two key elements.

7.1 General Setting

Consider the market for one asset, in supply S , that will trade in equilibrium at price p .¹⁴ The asset is characterized by a vector of attributes x_A , e.g. the mean and variance of its final payoff. There are two market participants, households, and intermediaries.

Intermediaries are characterized by a vector of attributes x_I , e.g. their size, leverage, or manager. Their demand for the asset depends on the characteristics of the asset x_A , their own attributes x_I , and the price of the asset p , summarized by the function $D_I(p, x_A, x_I)$.

Households are characterized by a vector of attributes x_H , e.g. their wealth, risk aversion, or beliefs. Importantly they own the intermediaries. Therefore, their demand for the asset depends not only on the price, their attributes and the attributes of the asset, but also of how much of the asset they is owned by the intermediary D_I^* . This is summarized by the function $D_H(p, D_I^*, x_A, x_H)$.

These demand functions map to the notation of our setting in the main text. The attributes of the assets are μ , σ and c . The attributes of the household and the intermediary

¹⁴The case of a non-fixed supply function, for instance $S(p)$ does not affect our conclusions.

are γ_H and γ_I respectively. The intermediary chooses the standard mean variance optimal, the ratio of the expected return $\mu - p$ to the product of payoff variance σ^2 and its risk aversion σ^2 . The households targets a similar optimum total portfolio, but offsets her own trading to take into account the assets she already holds through the intermediary. The inside-outside substitution rate is $-\partial D_H / \partial D_I = \frac{\gamma_H \sigma^2}{\gamma_H \sigma^2 + c}$. Finally, there will be a difference in preferences (and hence a separate notion of intermediary demand), when $\gamma_I \neq \gamma_H$. The case where they are equal essentially means the intermediary simply acts on the households behalf with no friction.

The equilibrium price is determined by market clearing, plugging into households demand the intermediary demand for the asset:

$$D_H(p, D_I(p, x_A, x_I), x_A, x_H) + D_I(p, x_A, x_I) = S \quad (23)$$

To understand price determination, consider the local change in price in response to a change in the various attributes:

$$\Delta p = \underbrace{\frac{-1}{\frac{\partial D_H}{\partial p} + \left(1 + \frac{\partial D_H}{\partial D_I}\right) \frac{\partial D_I}{\partial p}}}_{\text{demand slope}} \left[\underbrace{\left(\frac{\partial D_H}{\partial x_A} + \left(1 + \frac{\partial D_H}{\partial D_I}\right) \frac{\partial D_I}{\partial x_A} \right) \Delta x_A + \frac{\partial D_H}{\partial x_H} \Delta x_H}_{\text{asset attributes}} + \underbrace{\left(1 + \frac{\partial D_H}{\partial D_I}\right) \frac{\partial D_I}{\partial x_I} \Delta x_I}_{\text{investor attributes}} \right] \quad (24)$$

The first term in the product is the slope of the aggregate demand curve, the second term is the shift in demand curves coming from a change in the attributes. From this relation, we can immediately see that two ingredients shape the impact of intermediaries on asset prices. The first element is not surprisingly the *intermediary demand* for the asset. In particular, how their investment decisions respond to changes in their environment affects the aggregate demand for the asset, and in equilibrium the price. This effect manifests itself through the partial derivatives of D_I in equation (24).

The second element is how households substitute between holdings through the in-

termediary and direct holdings. This corresponds to what we call the *inside-outside substitution rate*, the sensitivity $-\partial D_H/\partial D_I$. This sensitivity controls the extent to which households offset intermediaries trade by directly trading the asset.

To highlight the separate importance of those two elements, let us consider the particular cases where intermediaries do not affect prices. For our first element, it could be that the investments of intermediary do not have depend at all on their attributes, but rather only on households attributes. In this case there wouldn't be a meaningful notion of intermediary demand curve. Intuitively, this occurs if intermediaries simply reflect the preferences of households and act on their behalf with no friction. For our second element, it might be that households substitute exactly one-to-one between the assets they hold directly and those held through intermediaries, perfectly offsetting these decisions ($-\partial D_H/\partial D_I = 1$). Intuitively, this occurs if households can invest directly in asset markets with no cost and there is no advantage to investing through intermediaries. The next section makes this more explicit in a simple example.

Table 1: Main predictive regressions. Predictive regressions of future excess returns in each asset class on our proxy for intermediary risk aversion, γ_{Int} . Our proxy is the average of the standardized versions of the AEM and HKM intermediary factors. We run: $r_{i,t+1}/E[r_{i,t+1}] = a_i + b_i x_t + \varepsilon_{i,t+1}$ and report b_i which gives the elasticity of the risk premium of asset i to x . See text for more details. Bootstrapped standard errors are in parenthesis and adjust for the fact that unconditional expected returns ($E[r_{i,t+1}]$) are estimated. See text for more details.

Panel A: Quarterly Returns							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Stocks	Bonds	Options	Sovereigns	Commodities	FX	CDS
γ_{Int}	0.71 (0.57)	0.48** (0.21)	1.30** (0.64)	1.03** (0.40)	3.49** (1.69)	0.43* (0.25)	2.67*** (0.74)
N	167	148	103	65	105	116	47
R^2	1.4%	1.4%	3.6%	14.0%	4.1%	3.0%	33.1%
Panel B: Annual Returns							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Stocks	Bonds	Options	Sovereigns	Commodities	FX	CDS
γ_{Int}	0.31 (0.25)	0.36** (0.16)	0.57* (0.33)	0.65*** (0.14)	2.07** (0.82)	0.19** (0.10)	1.12*** (0.38)
N	164	145	100	62	102	113	44
R^2	1.2%	2.9%	2.6%	25.7%	4.7%	2.7%	23.3%

Table 2: Predictive regressions using volatility normalization. We repeat the previous regressions but we normalize by volatility instead of means. We run: $r_{i,t+1}/\sigma[r_{i,t+1}] = a_i + b_i x_t + \varepsilon_{i,t+1}$ and report b_i .

Panel A: Quarterly Returns							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Mkt	Bonds	Options	Sovereigns	Commodities	FX	CDS
γ_I	0.12 (0.09)	0.13 (0.08)	0.19* (0.10)	0.38*** (0.13)	0.20** (0.10)	0.18** (0.08)	0.57*** (0.12)
N	167	148	103	65	105	116	47
R^2	0.0144	0.0142	0.0360	0.140	0.0409	0.0299	0.331
Panel B: Annual Returns							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Mkt	Bonds	Options	Sovereigns	Commodities	FX	CDS
γ_I	0.11 (0.10)	0.18** (0.07)	0.16 (0.13)	0.50*** (0.13)	0.21** (0.08)	0.16* (0.08)	0.47*** (0.13)
N	164	145	100	62	102	113	44
R^2	0.0123	0.0285	0.0259	0.257	0.0469	0.0265	0.233

Table 3: Predictive regressions including *habit*. We repeat our predictive regressions at the quarterly horizon but now include $-\ln(\textit{habit})$ as a control (defined as the surplus consumption ratio from Campbell and Cochrane), which is used as a potential proxy for movement in household risk aversion. If this is true, it should not display the increasing absolute magnitudes of predictive coefficients across assets to the degree that the intermediary variables do. Both predictive variables are standardized to have mean zero and unit standard deviation. Panel A reports elasticities and runs $r_{i,t+1}/E[r_{i,t+1}] = a_i + b_i x_t + \varepsilon_{i,t+1}$ and reports b_i while Panel B normalizes by volatilities and runs $r_{i,t+1}/\sigma[r_{i,t+1}] = a_i + b_i x_t + \varepsilon_{i,t+1}$.

Panel A: Annual Returns, Elasticities							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Mkt	Bonds	Options	Sovereigns	Commodities	FX	CDS
γ_I	0.17 (0.26)	0.23 (0.17)	0.60* (0.33)	0.62*** (0.17)	2.17** (0.88)	0.26** (0.11)	1.01*** (0.38)
γ_H	0.44* (0.24)	0.37 (0.27)	-0.14 (0.55)	0.07 (0.19)	-0.43 (1.45)	-0.28** (0.12)	0.22 (0.42)
N	164	145	100	62	102	113	44
R^2	0.0350	0.0589	0.0271	0.259	0.0484	0.0751	0.237

Panel B: Annual Returns, Volatility Normalization							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Mkt	Bonds	Options	Sovereigns	Commodities	FX	CDS
γ_I	0.06 (0.10)	0.12 (0.07)	0.17* (0.10)	0.48*** (0.14)	0.23** (0.10)	0.22*** (0.06)	0.43*** (0.16)
γ_H	0.19* (0.10)	0.18 (0.11)	-0.04 (0.11)	0.05 (0.11)	-0.04 (0.15)	-0.24** (0.10)	0.09 (0.17)
N	164	145	100	62	102	113	44
R^2	0.0350	0.0589	0.0271	0.259	0.0484	0.0751	0.237

Table 4: Predictive regressions including *cay*. We repeat our predictive regressions at the quarterly horizon but now include *cay* as a control, which is sometimes used as a potential proxy for movement in household risk aversion. If this is true, it should not display the increasing absolute magnitudes of predictive coefficients across assets to the degree that the intermediary variables do. Both predictive variables are standardized to have mean zero and unit standard deviation. Note: *cay* has a positive coefficient as it positively predicts returns consistent with prior studies. Panel A reports elasticities and runs $r_{i,t+1}/E[r_{i,t+1}] = a_i + b_i x_t + \varepsilon_{i,t+1}$ and reports b_i while Panel B normalizes by volatilities and runs $r_{i,t+1}/\sigma[r_{i,t+1}] = a_i + b_i x_t + \varepsilon_{i,t+1}$.

Panel A: Quarterly Returns, Elasticities							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Mkt	Bonds	Options	Sovereigns	Commodities	FX	CDS
γ_I	0.90*	0.49**	1.31*	1.00***	3.50*	0.43**	2.74***
	(0.48)	(0.24)	(0.67)	(0.32)	(1.84)	(0.22)	(1.00)
γ_H^{cay}	1.25***	0.35	1.98***	0.61	-0.65	0.30	-0.26
	(0.37)	(0.37)	(0.68)	(0.45)	(2.34)	(0.30)	(1.39)
N	167	148	103	65	105	116	47
R^2	0.0557	0.0197	0.0842	0.171	0.0417	0.0394	0.332
Panel B: Quarterly Returns, Volatility Normalization							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Mkt	Bonds	Options	Sovereigns	Commodities	FX	CDS
γ_I	0.15*	0.13*	0.19*	0.36***	0.20**	0.18*	0.59***
	(0.08)	(0.07)	(0.11)	(0.13)	(0.09)	(0.09)	(0.19)
γ_H^{cay}	0.21***	0.09	0.29***	0.22	-0.04	0.12	-0.06
	(0.07)	(0.09)	(0.09)	(0.15)	(0.14)	(0.09)	(0.27)
N	167	148	103	65	105	116	47
R^2	0.0557	0.0197	0.0842	0.171	0.0417	0.0394	0.332

Table 5: Regressions with interaction term. We test elasticity differences between stocks and other assets classes by including an interaction term for degree of intermediation $r_{i,t+1}/\overline{(r_{i,t+1})} = a_i + b_{i,1}\gamma_{I,t} + b_{i,2}INT \times \gamma_{I,t} + \varepsilon_{i,t+1}$ and reports the coefficients. Here INT captures how intermediated the asset class is. First, we define INT to equal 1 for all asset classes besides stocks and bonds (the two least intermediated according to our rankings). Next, we also include options. Finally, we define INT to equal the rank on a scale from 0 to 1 of all seven of our asset classes. In addition, we report a similar interaction term where we use the habit measure for a proxy of γ_H .

$INT =$	$1_{\neq Stock/Bond}$	$1_{\neq Stock/Bond/Opt}$	Rank $\in [0,1]$			
γ_I	0.33* (0.18)	0.20 (0.18)	0.39** (0.20)	0.31 (0.19)	0.36* (0.21)	0.23 (0.21)
$INT \times \gamma_I$	0.56** (0.27)	0.76*** (0.28)	0.61** (0.30)	0.77** (0.34)	0.75** (0.37)	1.04** (0.41)
γ_H		0.41** (0.19)		0.29 (0.21)		0.40* (0.22)
$INT \times \gamma_H$		-0.61* (0.36)		-0.53 (0.37)		-0.85* (0.45)
N	730	730	730	730	730	730
R^2	0.0288	0.0335	0.0296	0.0330	0.0280	0.0320

Table 6: Alternative Measures of Intermediary Health. We consider alternative ways to proxy for intermediary risk aversion or risk bearing capacity. The first, in Panel A, uses the log levels of the AEM and HKM factors (again, in levels we average the two after standardizing them). In Panel B, we instead use the Gilchrist and Zakrajek (2012) (GZ) spread to proxy for intermediary risk aversion instead of the AEM or HKM measures.

Panel A: Levels of AEM & HKM instead of changes							
	Stocks	Bonds	Options	Sovereigns	Commodities	FX	CDS
γ_I^{level}	0.51** (0.21)	0.01 (0.19)	1.35*** (0.38)	1.16*** (0.29)	1.81** (0.81)	-0.18 (0.11)	1.44** (0.68)
N	168	145	100	62	102	113	44
R^2	0.0331	3.28e-05	0.110	0.212	0.0275	0.0204	0.137
Panel B: GZ Spread as proxy for γ_I							
	(1) Stocks	(2) Bonds	(3) Options	(4) Sovereigns	(5) Commodities	(6) FX	(7) CDS
GZ	0.10 (0.31)	0.60*** (0.12)	-0.50 (0.76)	0.49*** (0.16)	0.32 (1.03)	0.10 (0.11)	1.22*** (0.28)
N	156	145	100	62	102	113	44
R^2	0.00118	0.0899	0.0217	0.220	0.00126	0.00827	0.404

Table 7: Splitting the Intermediary Health Measure. We split our measure into the HKM and AEM components separately. Panel A gives our main result using the annual log changes of each measure (as we do in our main result) while Panel B shows results using the log levels of each variable instead of changes.

Panel A: Annual Changes							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Stocks	Bonds	Sovereign	Commodities	CDS	Options	FX
γ_I^{AEM}	0.42 (0.26)	0.22* (0.12)	0.50*** (0.15)	3.44*** (0.58)	0.79** (0.38)	0.90*** (0.26)	0.26*** (0.08)
γ_I^{HKM}	0.04 (0.27)	0.27 (0.18)	0.39** (0.16)	-1.12 (0.93)	0.71* (0.39)	-0.25 (0.37)	-0.01 (0.10)
N	164	145	62	102	44	100	113
R^2	0.020	0.029	0.262	0.201	0.234	0.094	0.056
Panel B: Levels							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Stocks	Bonds	Sovereign	Commodities	CDS	Options	FX
γ_I^{AEM}	0.01 (0.39)	-0.31 (0.20)	0.75* (0.39)	1.75 (1.49)	0.80 (0.76)	1.00** (0.49)	0.22* (0.12)
γ_I^{HKM}	0.59 (0.37)	0.32 (0.22)	0.63*** (0.20)	0.23 (1.52)	0.78 (0.49)	0.45 (0.54)	-0.42*** (0.16)
N	168	145	62	102	44	100	113
R^2	0.041	0.020	0.214	0.035	0.137	0.117	0.095

Table 8: Changes in risk. We include changes in risk as controls for our main results. Specifically, we use trailing 3 year (12 quarter) rolling estimates of the volatility of each asset return and trailing 5 year (20 quarter) rolling market betas in each regression (we find we need a slightly longer time period to estimate the betas accurately). We report our main regression $r_{i,t+1}/\overline{(r_{i,t+1})} = a_i + b_i\gamma_{I,t} + \lambda_i\sigma_{i,t} + \delta_i\beta_{i,t} + \varepsilon_{i,t+1}$ using annual forecast horizon. Panel B repeats this using the volatility normalization $r_{i,t+1}/\sigma(r_{i,t+1})$.

Panel A: Elasticity							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Stocks	Bonds	Options	Sovereign	Commodities	FX	CDS
γ_I	0.35 (0.31)	0.33** (0.16)	0.76** (0.37)	0.64*** (0.20)	1.94* (1.05)	0.18* (0.10)	1.01** (0.42)
σ_t	-0.17 (0.20)	22.94 (29.00)	-26.04 (18.70)	18.33*** (5.85)	-69.29 (52.94)	9.56 (6.66)	241.18** (117.82)
β_t		-1.71 (2.03)	-2.02 (3.68)	-0.00 (0.12)	10.07 (6.23)	-0.09 (0.67)	-33.16 (22.23)
N	157	133	88	50	90	101	32
R^2	0.0154	0.0368	0.0725	0.359	0.0687	0.0367	0.360
Panel B: Volatility Norm							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Stocks	Bonds	Options	Sovereign	Commodities	FX	CDS
γ_I	0.13 (0.08)	0.16* (0.09)	0.21** (0.10)	0.50*** (0.12)	0.20** (0.10)	0.15* (0.09)	0.42*** (0.15)
σ_t	-0.06 (0.06)	11.31 (12.75)	-7.20* (4.34)	14.14*** (4.17)	-7.19 (7.06)	8.07 (7.80)	101.24*** (37.20)
β_t		-0.85 (0.99)	-0.56 (1.23)	-0.00 (0.09)	1.05 (0.72)	-0.07 (0.58)	-13.92* (7.50)
N	157	133	88	50	90	101	32
R^2	0.0154	0.0368	0.0725	0.359	0.0687	0.0367	0.360

Table 9: Subsamples. We run our main regression $r_{i,t+1}/\overline{(r_{i,t+1})} = a_i + b_i\gamma_{I,t} + \varepsilon_{i,t+1}$ across subsamples. Panel A excludes the 2007-2009 financial crisis period while Panel B uses only data from 1990 onwards.

Panel A: Dropping 2007-2009							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Stocks	Bonds	Options	Sovereigns	Commodities	FX	CDS
γ_I	0.23 (0.32)	0.23 (0.17)	0.67** (0.33)	0.75*** (0.17)	2.92*** (0.81)	0.29*** (0.10)	1.20*** (0.27)
N	152	133	88	50	90	101	32
R^2	0.00665	0.00953	0.0361	0.326	0.104	0.0457	0.456
Panel B: Post 1990							
γ_I	0.51* (0.30)	0.52*** (0.16)	0.41 (0.40)	0.65*** (0.14)	2.03** (0.96)	0.25** (0.10)	1.12*** (0.35)
N	85	81	81	62	85	85	44
R^2	0.0323	0.169	0.0123	0.257	0.0412	0.0377	0.233

Table 10: 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 to 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	CDS
<i>FoF</i>	Stocks	Bonds		Sov Bonds			
<i>VaR</i>	Stocks	Bonds			Comm	FX	
<i>BIS</i>		Bonds	Options		Comm	FX	CDS
<i>Expense</i>	Stocks	Bonds	Sov Bonds	FX	Comm	Options	CDS

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).

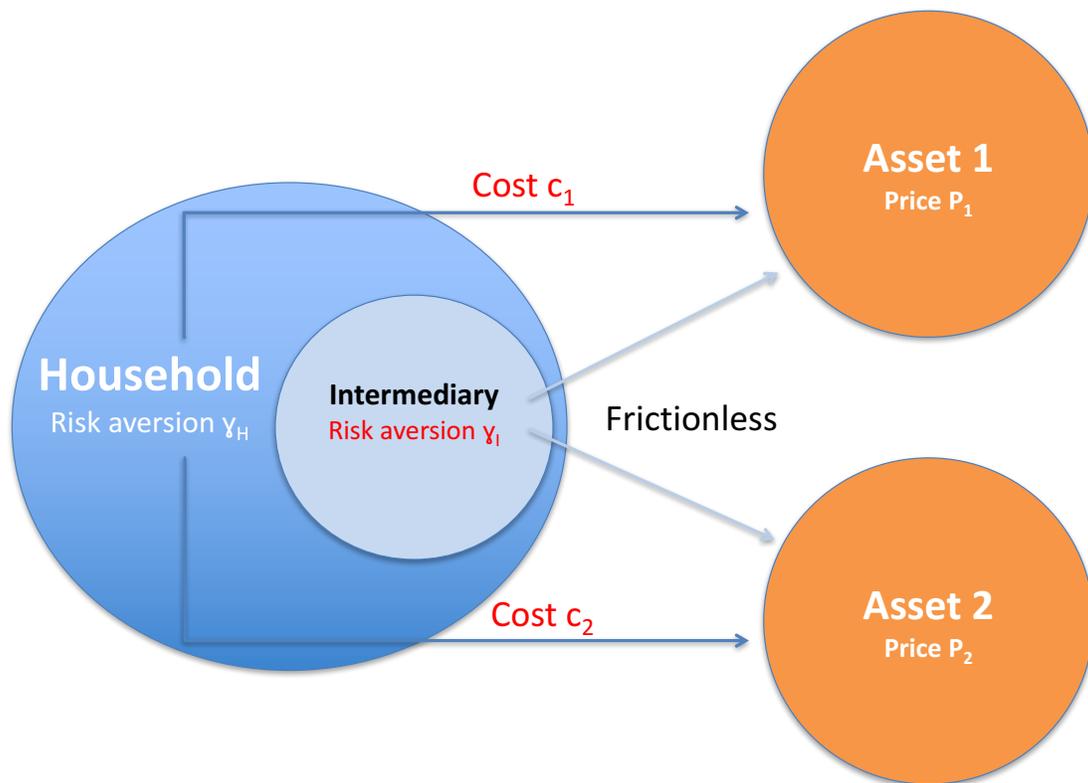
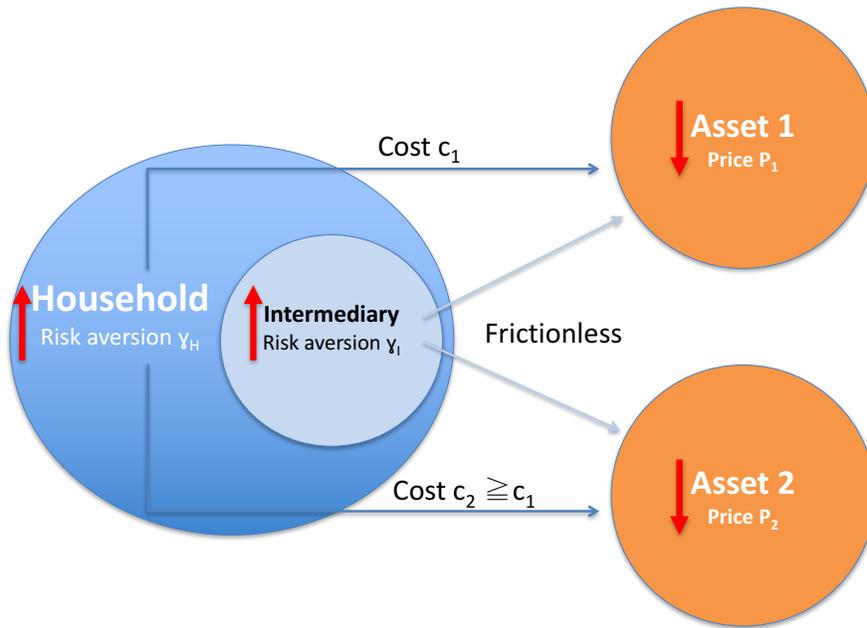


Figure 2: Model Shocks. This figure describes the response of asset prices to risk aversion changes. In Panel A, we show the response of a risk aversion shock under the null that intermediaries *don't* matter (either because $c = 0$ for all assets *or* because $\gamma_I = \gamma_H$) and in this case 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.

Panel A: Response to Aggregate Risk Aversion Shock Under Null



Panel B: Response to Intermediary Risk Aversion Shock

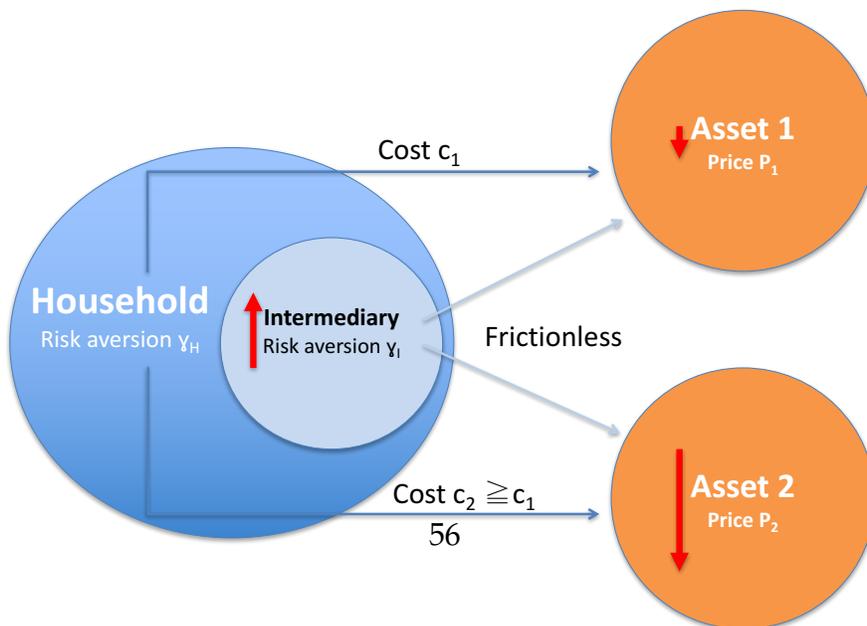


Figure 3: Main Empirical Results. This figure reports the behavior of risk premiums across asset classes (stocks, bonds, options, currencies, commodities, sovereigns, and cds) associated with a change in intermediary distress. We first plot the risk premia elasticity found by running $r_{i,t+k}/E[r_{i,t+k}] = a_i + b_i x_t + \varepsilon_{i,t+k}$ and we report b_i as we change the asset class i (units are $b_i \times 100$). The middle panel reports the R^2 , in percent, from this predictive regression as another measure of the degree of predictability by asset class. The bottom panel repeats this regression but normalizes by volatility $r_{i,t+k}/\sigma[r_{i,t+k}] = a_i + b_i x_t + \varepsilon_{i,t+k}$. The forecast horizon is 1 year ($k = 4$ quarters). All units are reported in percent. The right hand side variable x_t that measures intermediary health is an equal weighted average of the AEM and HKM factors (each are first standardized). We choose these measures because they have been argued to pick up health of the financial sector and have been shown to predict returns. See text for more details.

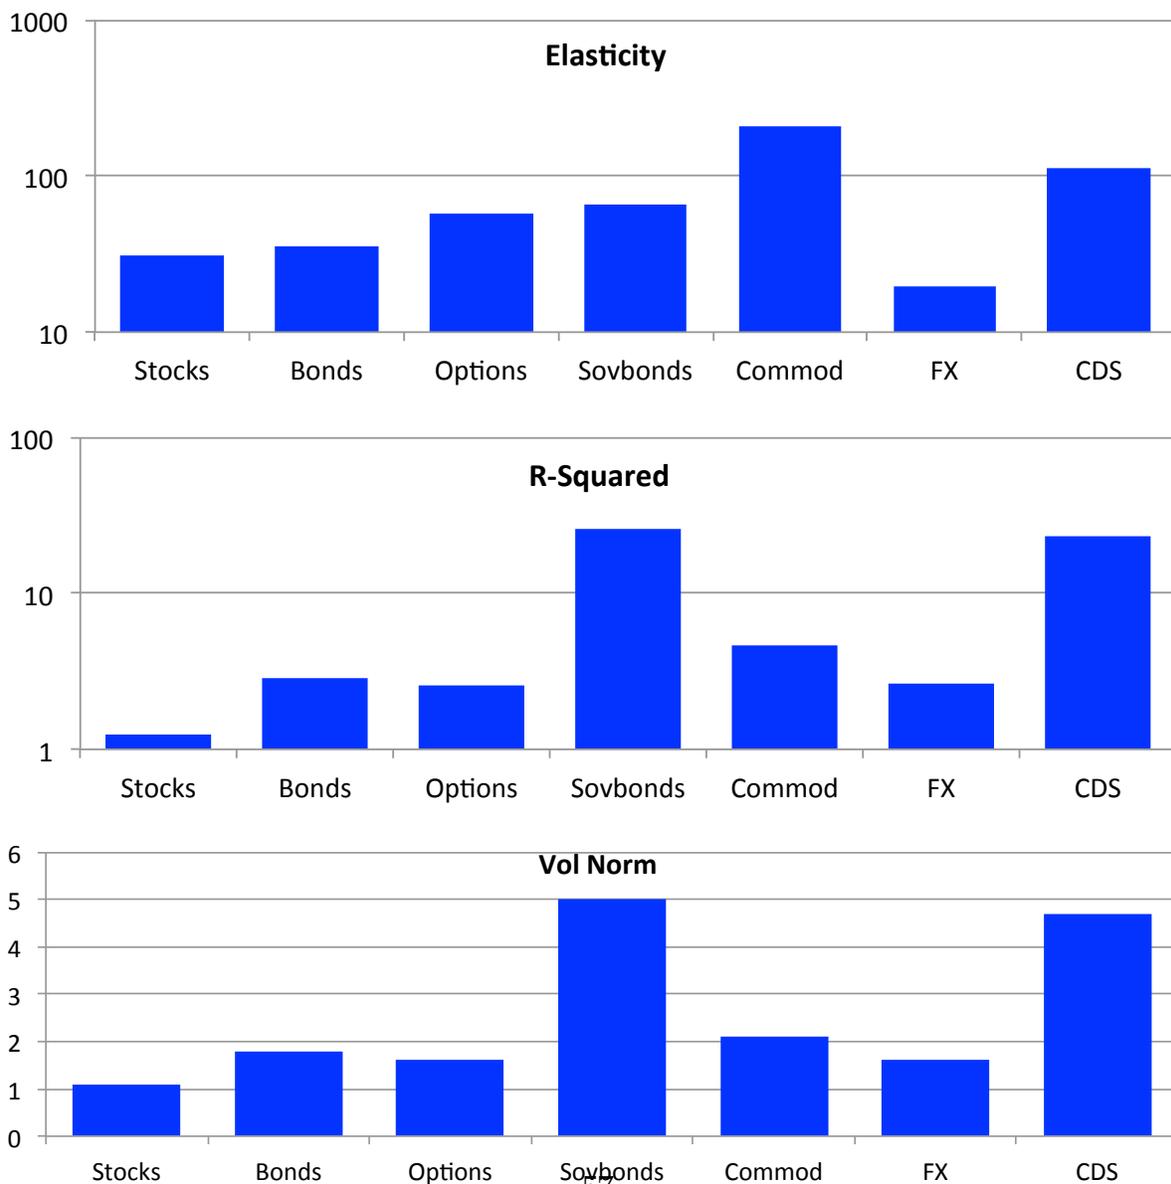


Figure 4: Replacing with HH Risk Aversion. This figure reports the behavior of risk premiums across asset classes (stocks, bonds, options, currencies, commodities, sovereigns, and cds) associated with a change in household risk aversion, proxied for by *cay*. We first plot the risk premia elasticity found by running $r_{i,t+k}/E[r_{i,t+k}] = a_i + b_i x_t + \varepsilon_{i,t+k}$ and we report b_i as we change the asset class i (units are $b_i \times 100$). The middle panel reports the R^2 , in percent, from this predictive regression as another measure of the degree of predictability by asset class. The bottom panel repeats this regression but normalizes by volatility $r_{i,t+k}/\sigma[r_{i,t+k}] = a_i + b_i x_t + \varepsilon_{i,t+k}$. The forecast horizon is 1 year ($k = 4$ quarters). All units are reported in percent. The right hand side variable x_t is the consumption to wealth ratio (*cay*) from Lettau and Ludvigson (2001). See text for more details.

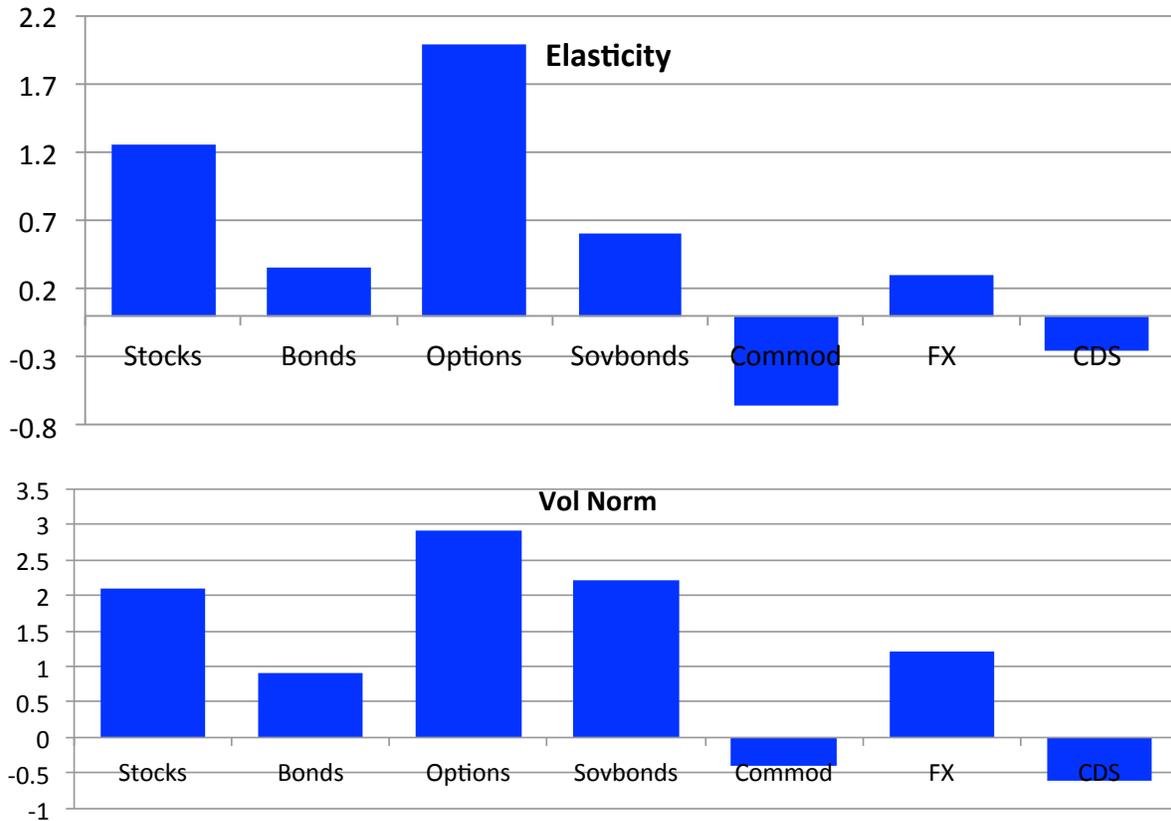


Figure 5: Hedge Fund Strategy Returns. This figure reports the behavior of risk premiums across stocks and hedge fund returns by category: long short equity, market neutral equity, the DJCS hedge fund index weighted across all hedge fund styles, event driven, convertible bond arbitrage, and fixed income arbitrage. We first plot the risk premia elasticity found by running $r_{i,t+k}/E[r_{i,t+k}] = a_i + b_i x_t + \varepsilon_{i,t+k}$ and we report b_i as we change the asset class i (units are $b_i \times 100$). The middle panel reports the R^2 , in percent, from this predictive regression as another measure of the degree of predictability by asset class. The bottom panel repeats this regression but normalizes by volatility $r_{i,t+k}/\sigma[r_{i,t+k}] = a_i + b_i x_t + \varepsilon_{i,t+k}$. The forecast horizon is 1 year ($k = 4$ quarters). All units are reported in percent. The right hand side variable x_t that measures intermediary health is an equal weighted average of the AEM and HKM factors (each are first standardized). We choose these measures because they have been argued to pick up health of the financial sector and have been shown to predict returns. See text for more details.

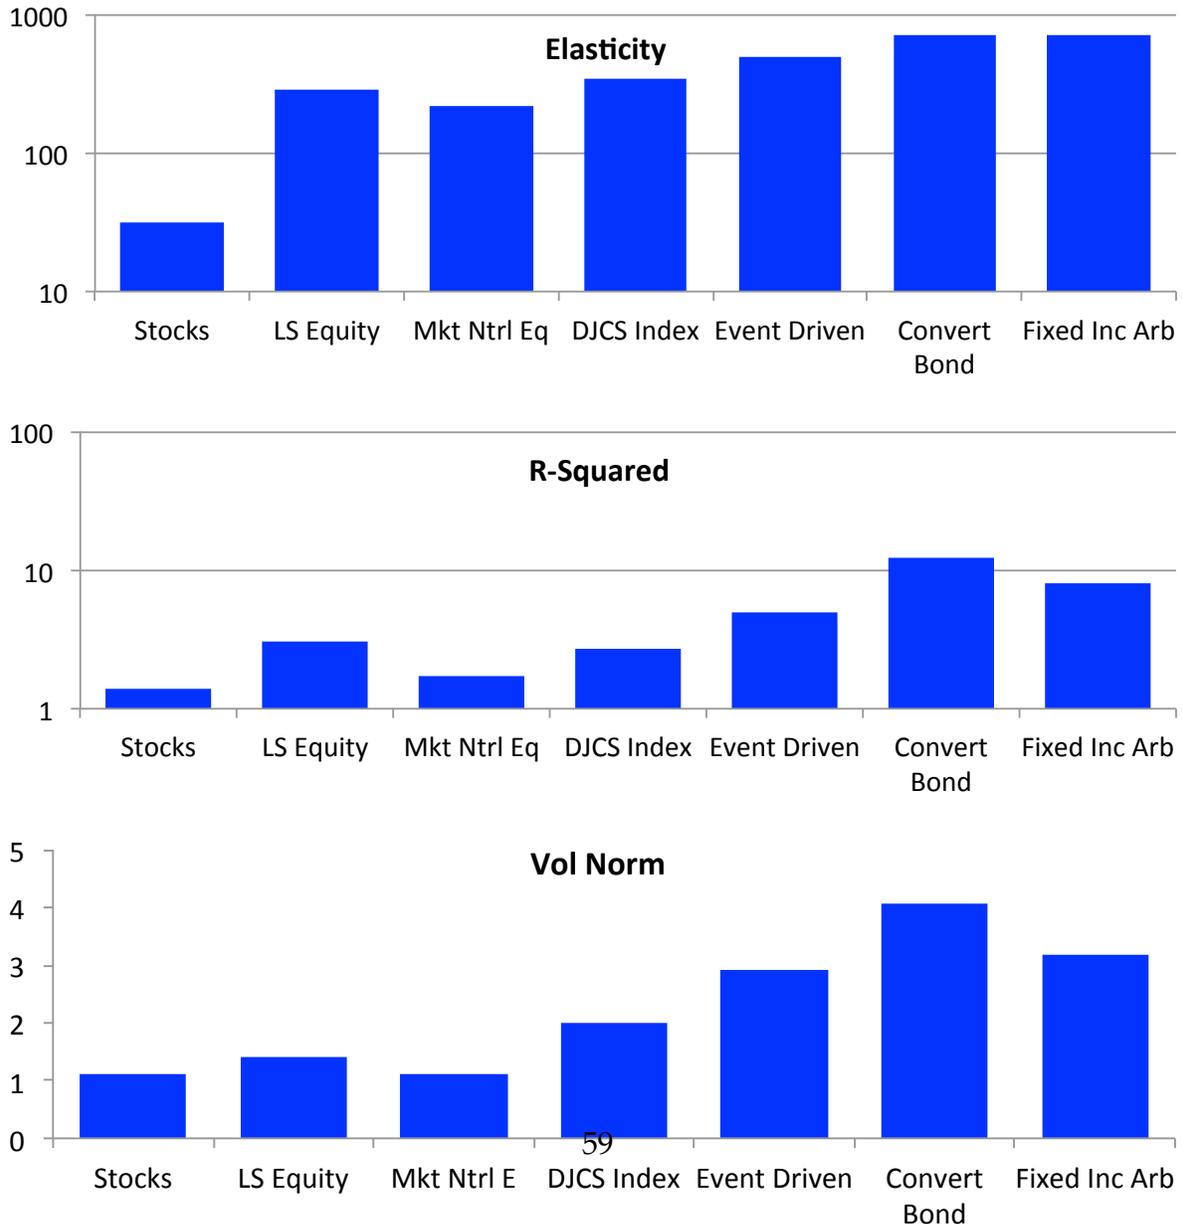


Figure 6: Decomposition of risk premia variation. This figure plots lower bounds of variation in risk premia coming from households vs intermediaries for each asset class. See text for detail.

