Do Intermediaries Matter for Aggregate Asset Prices?

Valentin Haddad and Tyler Muir*

November 3, 2017

Abstract

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 causal evidence that intermediaries matter for asset prices. We then provide a simple test: a sufficient condition for intermediaries to matter for asset prices is to document a larger elasticity of the risk premia of intermediated assets to changes in intermediary risk appetite. That is, intermediary health matters more for assets that households have difficulty buying directly. We provide direct empirical evidence that this is the case and hence show that intermediaries matter for a number of key asset classes including CDS, commodities, sovereign bonds, and FX.

*UCLA and NBER. We thank Markus Brunnermeier, Itamar Drechsler, Andrea Eisfeldt, Francis Longstaff, Arvind Krishnamurthy, Adrien Matray, Alan Moreira and participants at UCLA and the Wharton Conference on Liquidity and Financial Fragility for comments.

1. Introduction

A growing number of empirical studies document strong correlations between the health of financial intermediaries and aggregate asset prices.[1] These findings are important and suggestive, in part because they are consistent with models where financial frictions and the health of the financial sector matter for asset prices.[2] However, while the evidence is consistent with the idea that intermediaries matter, it can not rule out that intermediaries simply reflect, or are correlated with, the marginal utility of a representative agent. Consider the 2008 financial crisis where risk premiums rose substantially. While there was indeed a likely drop in intermediary risk-bearing capacity in the crisis, household risk aversion likely also rose, hence it is unclear to what extent, if any, the fall in intermediation mattered for aggregate asset prices.[3] There is also intriguing "micro" evidence that intermediaries matter for particular individual asset prices at particular points in time, though it is unclear what these results imply for aggregate asset price movements.[4] Hence, whether intermediaries are important for understanding broad aggregate asset price movements is an open question.

The goal of this paper is to address this question. First, we show sufficient conditions under which intermediaries *do* in fact matter for asset prices. To do so, we write down a simple, flexible model with financial intermediaries that encompasses both the frictionless view that intermediaries do not matter as well as the possibility that they do, and then we construct empirical tests that help us distinguish these possibilities. A sufficient condition for intermediaries to matter is to document a differential elasticity of risk premia to an intermediary state variable in the cross-section of assets, and this elasticity should be

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

[2] E.g., He and Krishnamurthy (2013), Brunnermeier and Sannikov (2014).

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

[4] E.g., Du et al. (2017), Siriwardane (2016), Fleckenstein et al. (2014), Lewis et al. (2017). For many additional examples, see Duffie (2010).

larger for more intermediated assets. We test this hypothesis by studying risk premia elasticities to intermediary state variables. Specifically, we run predictive regressions across asset classes (stocks, bonds, CDS, currencies, options, etc.) of excess returns in each asset class on measures argued to capture intermediary health and document larger elasticities in more intermediated assets (e.g., CDS markets). Quantitatively, our numbers suggest that a large amount of the variation in returns is linked to intermediary risk aversion. Intuitively, the test exploits that the frictions that make intermediaries matter are larger for some assets than others. Thus, an intermediary risk aversion shock will naturally have a larger impact in CDS markets (more intermediated) and smaller impact in stock markets (less intermediated). In contrast, an aggregate risk aversion shock in the frictionless case affects all risk premia in proportion. Thus, we are able to separate the intermediary risk bearing capacity story from a frictionless risk aversion story.

To lay out these issues and design our empirical tests, we first set up a simple intermediary asset pricing framework with many assets in which intermediaries and households both invest.[5] In the model, households own intermediaries and they take this into account. That is, when making their direct (non-intermediated) investment decisions, households take into account their indirect exposure to the assets through their ownership of the intermediaries.[6] We show in this context that for intermediaries to "matter" requires two things. First, we allow for the substitution by households between direct and indirect holdings to not be one to one, that is we allow for the possibility that households do not frictionlessly "undo" the indirect holdings by intermediaries. We capture this feature with an asset specific quadratic cost of households investing in a given asset

[5]Our 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. See also He and Krishnamurthy (forthcoming).

[6]A related model is Koijen and Yogo (2015) but that paper studies how institutional demand affects individual stock prices and is not able to address the time-series of aggregate asset prices because it does not feature or model the substitution of households direct vs indirect holdings.

directly. When this cost is zero, there is full substitution by households and intermediaries do not matter, as any shocks to intermediary health can be undone by households and thus are not reflected in asset prices. In the other extreme, an infinite cost represents pure intermediary asset pricing and this assumption is made in much of the theoretical literature for convenience where it is simply assumed households can not invest directly (e.g., He and Krishnamurthy (2013)).

Next we show that the second condition required for intermediaries to matter is that their preferences are not perfectly aligned with households. That is, even if there are large costs to direct ownership by households, intermediaries may still act as a veil and invest exactly as households would like them to. If this is the case, even with the friction of direct investment costs, intermediaries will not matter for asset prices in equilibrium. We capture this condition in the model by allowing for possibly different risk aversion between intermediaries and households so that intermediaries may have their own preferences and may make choices that are not identical to what households would choose. Again, we show if intermediary risk aversion is exactly equal to that of households, then again intermediaries do not matter regardless of the direct costs to investing. Intuitively, in this case intermediaries act on behalf of households in a frictionless way. These two conditions: (1) substitution of indirect (intermediated) vs. direct investing, and (2) preference alignment, are the core determinants for whether intermediaries matter and both conditions are jointly required. That is, for intermediation to matter we need a lack of full substitution *and* a lack of preference alignment. 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.[7]

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

We are then able to discuss the key empirical studies in the literature. 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 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. However, we show that, under some additional assumptions, if the household Euler equation does not hold then intermediaries do matter. This speaks to the suggestive evidence on the lack of power of the CCAPM and relative success of intermediary based Euler equations. Yet, in this case, it is unclear whether we have a poor model of household optimization and preferences, bad consumption data (either because of measurement error or because we want to measure consumption of only stockholders), or some other fundamental problem with the model.[8] Similarly, for time-series studies that exploit recessions vs financial crises (e.g., the fact that risk premia are higher in financial crises vs recessions as shown in Muir (2017)), it is unclear whether we have an insufficient model of household risk aversion or whether the financial sector is causing the fluctuations in risk premia.

We show in the model that exploiting both the cross-section and time-series can sort out these issues and provides a much better test of whether intermediaries matter, in the sense of making far fewer assumptions on household behavior. Specifically, we show that when there are many assets, and differential costs to direct investment by households across assets (i.e., substitution rates are not the same across assets), then changes to

[8]For example, Greenwald et al. (2014) argue that movements in aggregate risk aversion appear uncorrelated with standard measures of consumption. Malloy et al. (2009) argue that stockholder consumption lines up better with asset returns, while papers like 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.

intermediary risk aversion will *differentially* affect the cross-section of risk premia. These differential substitutions rates are intuitive: it may be easy for households to invest in the S&P500 directly but difficult for them to invest directly in CDS markets. In this case, an intermediary risk aversion shock will affect CDS risk premia disproportionally more than stock market risk premia. In contrast, a risk aversion shock under the null of a frictionless model will affect all risk premia the same because it simply multiplies the entire covariance matrix. This constitutes the main test in this paper, and we indeed show evidence of such disproportionate effects.

In the model we show that the elasticity of an assets' risk premium to a change in intermediary risk aversion maps directly to the substitution of household's direct vs indirect holdings and hence to the quadratic cost of direct investment. While the basic idea of the test is intuitive, we show that it is only when measured properly through risk premia elasticities that we net out effects of asset supply or differences in risk across asset classes. That is, the correct test is to see the percentage change in risk premium for an asset that is associated with a decline in intermediary health and compare these elasticities across asset classes. Given this, our test provides a lower bound for how much intermediaries matter in each asset class, but the lower bound for our benchmark asset class whether intermediaries matter least is equal to zero and hence uninformative. This is because for this asset class we can not disentangle household vs intermediary risk aversion effects.

We then take our framework to the data. We capture the main test of the model using predictive regressions normalized by unconditional average excess returns to that our coefficients represent elasticities. We use common proxies for intermediary risk aversion such as broker-dealer leverage and intermediary equity capital. We compute risk premia elasticities in various asset classes including stocks, bonds, credit, CDS, options, commodities, and foreign exchange. We expect stocks to have the lowest elasticity consistent with households being most easily able to invest in stocks directly, and hence they are

used as our benchmark asset class.

We find evidence that intermediaries matter for CDS, options, commodities, and foreign exchange and thus we are able to reject the null hypothesis of a frictionless view that intermediaries don't matter for these aggregate asset prices. Specifically, relative to stocks we find much larger elasticities to the same intermediary risk aversion shock in these markets, which says that intermediaries matter relatively more in these markets than they do for the overall stock market. We place a lower bound on the extent to which intermediaries in each asset class and find this lower bound is, at times, fairly large. We stress that the results for stocks are ambiguous: while we *do* see a large elasticity of intermediary risk aversion shocks we can in no way claim causality without taking a stand on household risk aversion behavior. Hence, we can not conclude whether or not intermediaries matter for the overall stock market.

While we emphasize that we do not take any stand on the behavior of household risk aversion, we also include proxies for household risk aversion as additional suggestive evidence for our mechanism. That is, the model says that a household or "aggregate" risk aversion shock should not affect these specialized assets by more than the stock market (assuming the stock market is where direct investment is easiest). We use proxies of consumer sentiment and the aggregate consumption to wealth ratio (cay) as household risk aversion proxies and we find consistent evidence with our hypothesis: while there is evidence that these proxies do forecast asset returns and hence are associated with time-varying risk premiums, we find zero evidence of larger effects in more intermediated assets again consistent with our hypothesis.

Finally, we go through other possibilities that could explain 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, 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.

This paper is the first attempt to lay out and try to tackle the identification challenge associated with intermediary asset pricing. 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.

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 voluntarily 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

$$\max_{D_I} D_I' (\mu - p) - \frac{\gamma_I}{2} D_I' \Sigma D_I. \tag{1}$$

The household takes as given the investment decision of the intermediary when mak-

ing her choice of direct holding D_H:

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

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

$$D_H + D_I = S. \tag{3}$$

2.2 Equilibrium Portfolios and Prices

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

$$D_I^* = \frac{1}{\gamma_I} \Sigma^{-1} (\mu - p). \tag{4}$$

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:

$$D_H^* = (\gamma_H \Sigma + C)^{-1} (\mu - p) - (\gamma_H \Sigma + C)^{-1} (\gamma_H \Sigma) D_I. \tag{5}$$

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

$$-\frac{\partial D_H^*}{\partial D_I} = (\gamma_H \Sigma + C)^{-1}(\gamma_H \Sigma). \tag{6}$$

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:

$$\mu - p = \gamma_H \Sigma \left(\Sigma + \frac{1}{\gamma_I} C\right)^{-1} \left(\Sigma + \frac{1}{\gamma_H} C\right) S \tag{7}$$

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*

$$\gamma_I \neq \gamma_H \ \textit{and} \ \ C \neq 0 \tag{8}$$

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, $-\partial D_H^*/\partial D_I \neq I$.

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, $-\partial D_H/\partial D_I \neq I$. 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 wether 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 (4). It implies that it forms a pricing kernel: writing R_I the excess return on the intermediary risky portfolio, then for any any risky asset excess

return R_i, we have:

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

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. 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:

$$\frac{\partial(\mu - p)}{\partial \gamma_I} \geq 0 \tag{9}$$

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:

$$\frac{1}{\mu_i - p_i}\frac{\partial(\mu_i - p_i)}{\partial \log(\gamma_I)} = \frac{c_i}{\gamma_I \sigma_i^2 + c_i} \tag{10}$$

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

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 asses 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.

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

Our model guides us to run the following regressions:

$$r_{i,t+k}/E[r_{i,t+k}] = a_i + b_i x_t + \varepsilon_{i,t+1}, \quad i = 1, ..., N, \;\; t = 1, ..., T \tag{12}$$

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.

4.3 Results

We run predictive regressions in each asset class

$$r_{i,t+1}/E[r_{i,t+1}] = a_i + b_i x_t + \varepsilon_{i,t+1}, \quad i = 1, ..., N, \;\; t = 1, ..., T \tag{13}$$

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^2s 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^2s 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 variance, rather than their mean. Specifically, we run

$$r_{i,t+1}/Var[r_{i,t+1}] = a_i + b_i x_t + \varepsilon_{i,t+1}, \quad i = 1, ..., N, \;\; t = 1, ..., T \tag{14}$$

This normalization has some disadvantages – namely that it's justification requires

additional assumptions in our model about covariances – but also some 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. 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' variance 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 3 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.

4.4 Robustness and additional results

4.4.1 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 variance of the asset, and also that there is meaningful variation in R^2s 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.

Variance Normalization

We show how our model implies an alternative normalization for our predictive re-

gressions where we normalize returns by variance rather than means. That is, we run $r_{i,t+1}/Var[r_{i,t+1}]$ 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

$$\epsilon_{\mu_i} = \epsilon_{\gamma_I} \frac{c_i}{\gamma_I \sigma_i^2 + c_i} + \epsilon_{\gamma_H} \frac{\gamma_H \sigma_i^2}{\gamma_H \sigma_i^2 + c_i} \tag{15}$$

We multiply both sides by μ_i/σ_i^2, and then use the equilibrium relationship between μ_i and σ_i^2 (assuming Σ is diagonal), to obtain

$$\epsilon_{\sigma_i} = \left(\epsilon_{\gamma_I} \frac{\gamma_I c_i (\gamma_H \sigma_i^2 + c_i)}{(\gamma_I \sigma_i^2 + c_i)^2} + \epsilon_{\gamma_H} \frac{\gamma_H \sigma_i^2}{\sigma_i^2 + c_i/\gamma_I} \right) S_i \tag{16}$$

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

$$\mu - P = f\left(\gamma_I(x), \gamma_H(x)\right) \tag{17}$$

$$R_{t+1} = f\left(\gamma_I(x_t), \gamma_H(x_t)\right) + \varepsilon_{t+1} \tag{18}$$

$$var(R_{t+1}) = f'(x)^2 var(x) + var(\varepsilon_{t+1}) \tag{19}$$

Assuming without loss of generality that we standardize the variance of the shock, x,

so $var(x) = 1$, and then using that $f'(x) = (\mu - P)\epsilon_\mu$, we can rearrange to obtain

$$R^2 = \frac{(\mu - P)^2 \epsilon_\mu}{(\mu - P)^2 \epsilon_\mu + \sigma^2} \tag{20}$$

$$R^2 = \frac{1}{1 + \epsilon_\mu^{-1} \sigma^2 / (\mu - P)} \tag{21}$$

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.

4.4.2 Robustness of empirical results

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

Table 4 compares elasticities of each asset to stocks and uses a balanced sample for each asset. Specifically, we directly test elasticity differences between stocks and other assets classes in balances subsamples. Panel A reports elasticities and runs $r_{i,t+1} / \overline{(r_{i,t+1})} - r_{stock,t+1} / \overline{(r_{stock,t+1})} = a_i + b_i \gamma_{I,t} + \varepsilon_{i,t+1}$, while Panel B normalizes by variances and runs $r_{i,t+1} / \hat{\sigma}^2(r_{i,t+1}) - r_{stock,t+1} / \hat{\sigma}^2(r_{stock,t+1}) = a_i + b_i \gamma_{I,t} + \varepsilon_{i,t+1}$. In Panel A, we find that standard errors are large in our balanced panel data for each asset to where we can not definitively say that the elasticity of each alternative asset is larger than for that of stocks. In Panel B, we show that we can overcome this challenge when we use our variance normalization instead of the elasticity. However, this comes at the cost of requiring additional assumptions from our model.

In Table 5 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 6 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 7 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 8 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.[9] That is, our paper offers no theory of what deter-

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

mines 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 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)).[10] 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.[11]

[10] See also Chen et al. (2016), Hu et al. (2013), Muir (2017), Pasquariello (2014), Baron and Xiong (2017).
[11] See also Lou et al. (2013).

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 causal evidence that intermediaries matter for 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 causally claim that intermediaries matter for a number of key asset classes including CDS, FX, options, and commodities.

References

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

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

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.

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* .

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.

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.

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.[12] 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

[12] The 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:

$$D_H\left(p, D_I(p, x_A, x_I), x_A, x_H\right) + D_I(p, x_A, x_I) = S \tag{22}$$

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

$$\Delta p = \underbrace{\frac{-1}{\frac{\partial D_H}{\partial p} + \left(1 + \frac{\partial D_H}{\partial D_I}\right)\frac{\partial D_I}{\partial p}}}_{\text{demand slope}} \left[\underbrace{\left(\frac{\partial D_H}{\partial x_A} + \left(1 + \frac{\partial D_H}{\partial D_I}\right)\frac{\partial D_I}{\partial x_A}\right)\Delta x_A}_{\text{asset attributes}} + \underbrace{\frac{\partial D_H}{\partial x_H}\Delta x_H + \left(1 + \frac{\partial D_H}{\partial D_I}\right)\frac{\partial D_I}{\partial x_I}\Delta x_I}_{\text{investor attributes}} \right] \tag{23}$$

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 (23).

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.

In the remainder of this section, we present a model of an economy with intermediaries where our two main forces are linked to explicit parameters. Then, we discuss alternative mechanisms shaping those forces.

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) Stocks	(2) Bonds	(3) Sovereign	(4) Commodities	(5) CDS	(6) Options	(7) FX
γ_{Int}	-0.75	-0.49	-1.06	-3.95	-2.68	-1.43	-0.45
	(0.50)	(0.26)	(0.43)	(1.51)	(0.77)	(0.61)	(0.22)
N	167	148	65	105	47	103	116
R^2	1.6%	1.5%	15.4%	5.4%	35.6%	4.4%	3.2%
Panel B: Annual Returns							
	(1) Stocks	(2) Bonds	(3) Sovereign	(4) Commodities	(5) CDS	(6) Options	(7) FX
γ_{Int}	-0.34	-0.35	-0.64	-2.52	-1.08	-0.68	-0.22
	(0.27)	(0.15)	(0.16)	(0.78)	(0.44)	(0.30)	(0.09)
N	164	145	62	102	44	100	113
R^2	1.5%	2.7%	26.2%	7.1%	23.1%	3.8%	3.4%

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

Panel A: Quarterly Returns							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Mkt	Bonds	Sovereign	Commodities	CDS	Options	FX
γ_{Int}	-1.37	-4.74	-7.22	-3.64	-44.16	-2.66	-3.63
	(1.00)	(2.52)	(2.73)	(1.49)	(12.03)	(1.18)	(1.67)
N	167	148	65	105	47	103	116
R^2	1.6%	1.5%	15.4%	5.4%	35.6%	4.4%	3.2%
Panel B: Annual Returns							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Mkt	Bonds	Sovereign	Commodities	CDS	Options	FX
γ_{Int}	-0.02	-3.03	-4.86	-1.95	-17.29	-1.21	-1.91
	(0.02)	(1.31)	(1.24)	(0.61)	(6.88)	(0.54)	(0.81)
N	164	145	62	102	44	100	113
R^2	1.5%	2.7%	26.2%	7.1%	23.1%	3.8%	3.4%

Table 3: 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 until 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 variances and runs $r_{i,t+1}/E[r_{i,t+1}] = a_i + b_i x_t + \varepsilon_{i,t+1}$.

Panel A: Quarterly Returns, Elasticities							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Mkt	Bonds	Sovereign	Commodities	CDS	Options	FX
Int Fac	-0.90*	-0.49**	-1.00**	-3.50**	-2.74***	-1.31**	-0.43**
	(0.50)	(0.24)	(0.40)	(1.69)	(0.92)	(0.58)	(0.19)
cay	1.25***	0.35	0.61*	-0.65	-0.26	1.98***	0.30
	(0.38)	(0.37)	(0.37)	(2.30)	(1.42)	(0.58)	(0.26)
N	167	148	65	105	47	103	116
R^2	0.056	0.020	0.171	0.042	0.332	0.084	0.039
Panel B: Quarterly Returns, Variance Normalization							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Mkt	Bonds	Sovereign	Commodities	CDS	Options	FX
Int Fac	-1.64*	-4.71**	-6.75**	-3.23**	-45.07***	-2.43*	-3.53*
	(0.85)	(2.28)	(2.69)	(1.41)	(14.33)	(1.42)	(1.96)
cay	2.26***	3.39	4.11*	-0.60	-4.30	3.69***	2.44
	(0.63)	(2.91)	(2.38)	(2.20)	(23.43)	(1.31)	(2.41)
N	167	148	65	105	47	103	116
R^2	0.056	0.020	0.171	0.042	0.332	0.084	0.039

Table 4: Differences in elasticities. We directly test elasticity differences between stocks and other assets classes in balances subsamples. Panel A reports elasticities and runs $r_{i,t+1}/\overline{(r_{i,t+1})} - r_{stock,t+1}/\overline{(r_{stock,t+1})} = a_i + b_i\gamma_{I,t} + \varepsilon_{i,t+1}$ and reports b_i while Panel B normalizes by variances and runs $r_{i,t+1}/\hat{\sigma}^2(r_{i,t+1}) - r_{stock,t+1}/\hat{\sigma}^2(r_{stock,t+1}) = a_i + b_i\gamma_{I,t} + \varepsilon_{i,t+1}$.

Panel A: Elasticity Difference						
	(1)	(2)	(3)	(4)	(5)	(6)
	Bonds	Sovereign	Commodity	CDS	Options	FX
γ_I	0.26	-0.01	-1.40	-0.08	-0.06	0.09
	(0.20)	(0.36)	(0.85)	(0.36)	(0.15)	(0.28)
N	145	62	102	44	100	113
R^2	0.016	0.002	0.028	0.002	0.003	0.050
Panel B: Variance Norm Difference						
	(1)	(2)	(3)	(4)	(5)	(6)
	Bonds	Sovereign	Commodity	CDS	Options	FX
γ_I	-2.22*	-3.11***	-0.87	-14.88**	-0.14	-1.79**
	(1.21)	(1.13)	(0.68)	(6.66)	(0.20)	(0.76)
N	145	62	102	44	100	113
R^2	0.013	0.191	0.011	0.238	0.004	0.139

Table 5: 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	Sovereign	Commodities	CDS	Options	FX
γ_I	-0.53**	-0.01	-1.16***	-1.72*	-1.40**	-1.29***	0.18
	(0.22)	(0.18)	(0.28)	(0.89)	(0.58)	(0.34)	(0.13)
N	168	145	62	102	44	100	113
R^2	0.033	0.000	0.212	0.027	0.137	0.110	0.020
Panel B: GZ Spread as proxy for γ_I							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Stocks	Bonds	Sovereign	Commodities	CDS	Options	FX
GZ	-0.01	-6.14***	-3.10***	0.83	-12.35***	0.86	-0.38
	(0.28)	(1.09)	(1.01)	(1.05)	(4.09)	(0.84)	(0.98)
N	156	145	62	102	44	100	113
R^2	0.000	0.129	0.204	0.016	0.253	0.024	0.002

Table 6: 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) Stocks	(2) Bonds	(3) Sovereign	(4) Commodities	(5) CDS	(6) Options	(7) FX
γ_I^{AEM}	-0.42	-0.22*	-0.50***	-3.44***	-0.79**	-0.90***	-0.26***
	(0.26)	(0.12)	(0.15)	(0.58)	(0.38)	(0.26)	(0.08)
γ_I^{HKM}	-0.04	-0.27	-0.39**	1.12	-0.71*	0.25	0.01
	(0.27)	(0.18)	(0.16)	(0.93)	(0.39)	(0.37)	(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) Stocks	(2) Bonds	(3) Sovereign	(4) Commodities	(5) CDS	(6) Options	(7) FX
γ_I^{AEM}	-0.01	0.31	-0.75*	-1.75	-0.80	-1.00**	-0.22*
	(0.39)	(0.20)	(0.39)	(1.49)	(0.76)	(0.49)	(0.12)
γ_I^{HKM}	-0.59	-0.32	-0.63***	-0.23	-0.78	-0.45	0.42***
	(0.37)	(0.22)	(0.20)	(1.52)	(0.49)	(0.54)	(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 7: 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 quarterly and annual forecast horizons.

Panel A: Elasticity							
	(1) Stocks	(2) Bonds	(3) Sovereign	(4) Commodities	(5) CDS	(6) Options	(7) FX
γ_I	-0.31	-0.32*	-0.68***	-1.30	-1.07**	-0.71*	-0.17
	(0.28)	(0.18)	(0.17)	(0.95)	(0.45)	(0.38)	(0.11)
σ_t	-1.51	22.87	15.02**	-59.60	182.22*	-24.59	15.66*
	(8.41)	(20.40)	(7.20)	(47.18)	(107.50)	(17.25)	(8.83)
β_t		-1.77	-0.01	13.33**	-23.35	-2.39	-0.34
		(1.95)	(0.38)	(6.05)	(16.43)	(4.09)	(0.54)
N	157	133	50	90	32	88	101
R^2	0.011	0.035	0.303	0.093	0.311	0.077	0.053
Panel B: Variance Norm							
	(1) Stocks	(2) Bonds	(3) Sovereign	(4) Commodities	(5) CDS	(6) Options	(7) FX
γ_I	-0.57	-3.06**	-4.58***	-1.20	-17.62***	-1.32**	-1.42*
	(0.51)	(1.50)	(1.31)	(1.03)	(6.59)	(0.58)	(0.84)
σ_t	-2.75	219.02	101.88**	-55.03	3002.02*	-45.75	127.15**
	(15.24)	(216.70)	(47.69)	(46.60)	(1677.10)	(33.70)	(60.21)
β_t		-16.95	-0.04	12.31**	-384.73	-4.45	-2.80
		(19.71)	(2.17)	(5.62)	(287.65)	(8.20)	(4.69)
N	157	133	50	90	32	88	101
R^2	0.011	0.035	0.303	0.093	0.311	0.077	0.053

Table 8: 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.

	(1) Stocks	(2) Bonds	(3) Sovereign	(4) Commodities	(5) CDS	(6) Options	(7) FX
Panel A: Dropping 2007-2009							
γ_I	-0.22	-0.26	-0.73***	-2.74***	-0.90***	-0.49*	-0.25**
	(0.30)	(0.17)	(0.18)	(0.75)	(0.15)	(0.27)	(0.11)
N	141	126	46	79	21	81	90
R^2	0.007	0.010	0.354	0.170	0.628	0.037	0.057
Panel B: Post 1990							
γ_I	-0.42	-0.41***	-0.64***	-4.25**	-1.07***	-0.42	-0.23**
	(0.30)	(0.10)	(0.17)	(1.98)	(0.38)	(0.50)	(0.11)
N	84	80	62	84	44	80	84
R^2	0.025	0.163	0.254	0.038	0.231	0.008	0.035

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).

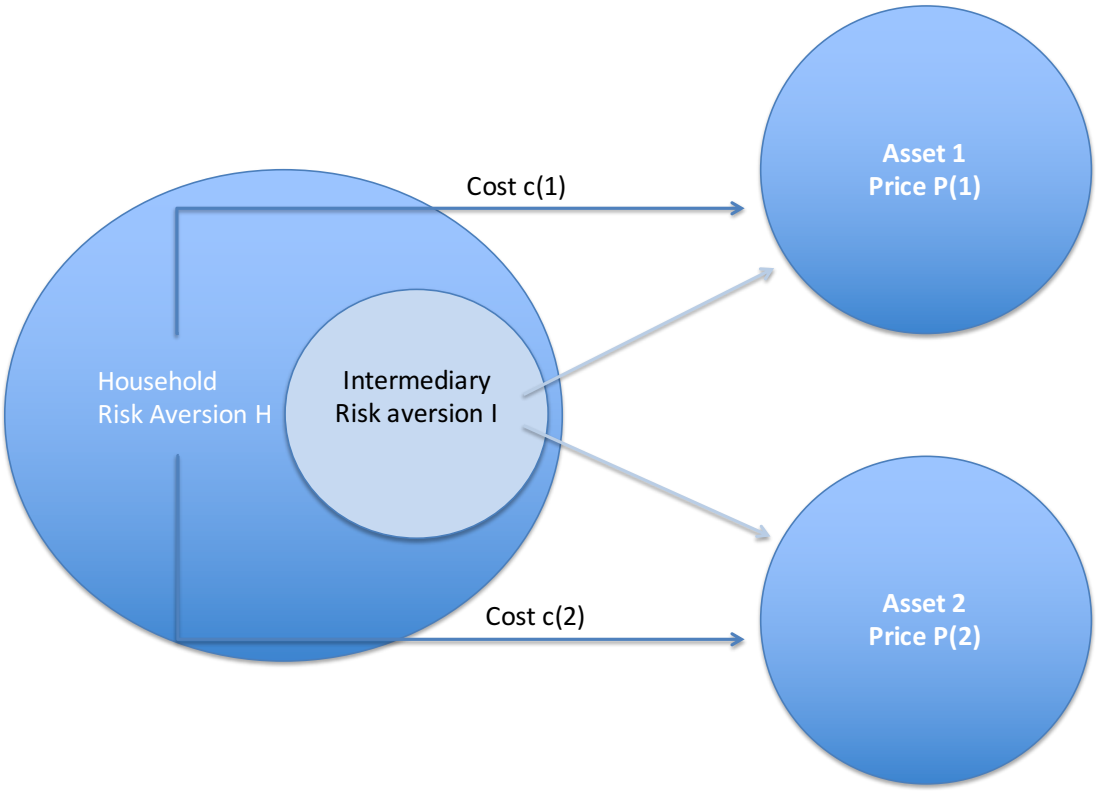

Figure 2: Model Shocks. This figure describes the response of asset prices to risk aversion changes. In Panel A, we show the response of a risk aversion shock under the null that intermediaries *don't* matter (either because $c = 0$ for all assets *or* because $\gamma_I = \gamma_H$) and in this case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

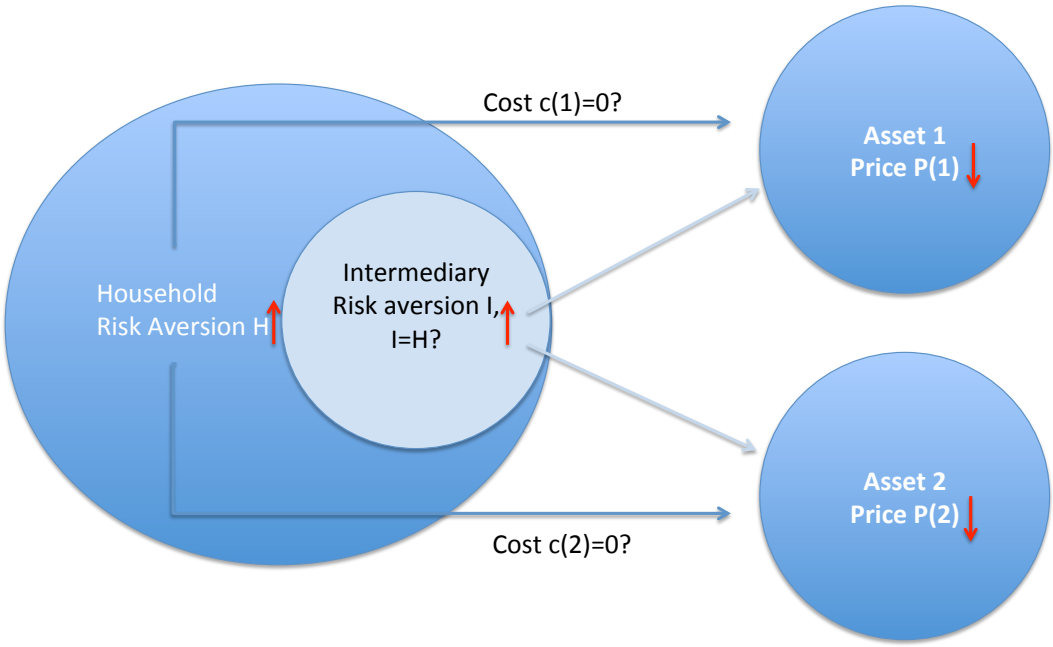

Panel B: Response to Intermediary Risk Aversion Shock

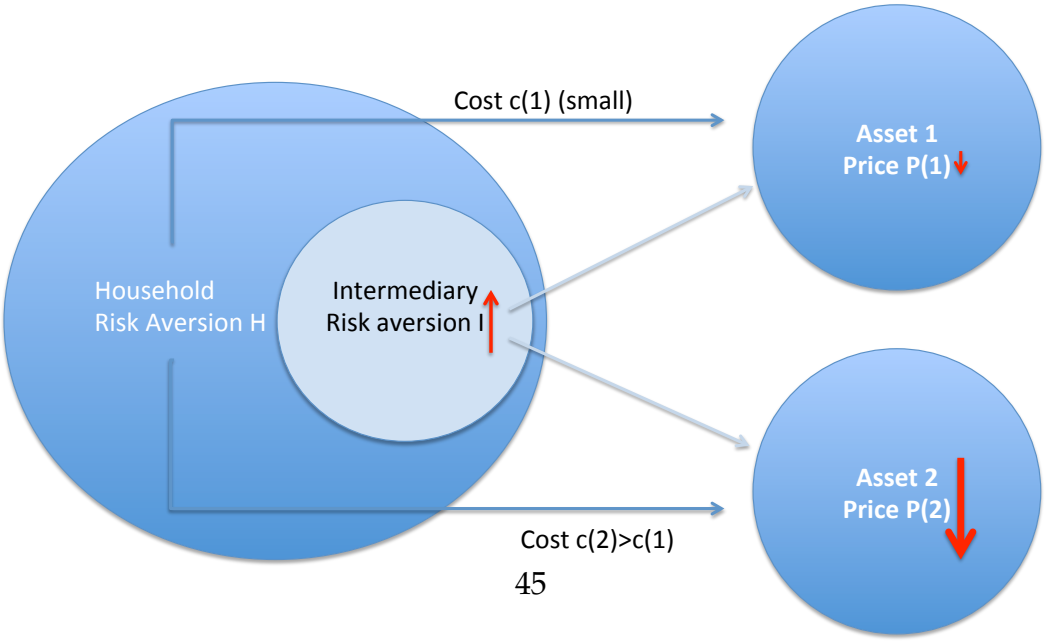
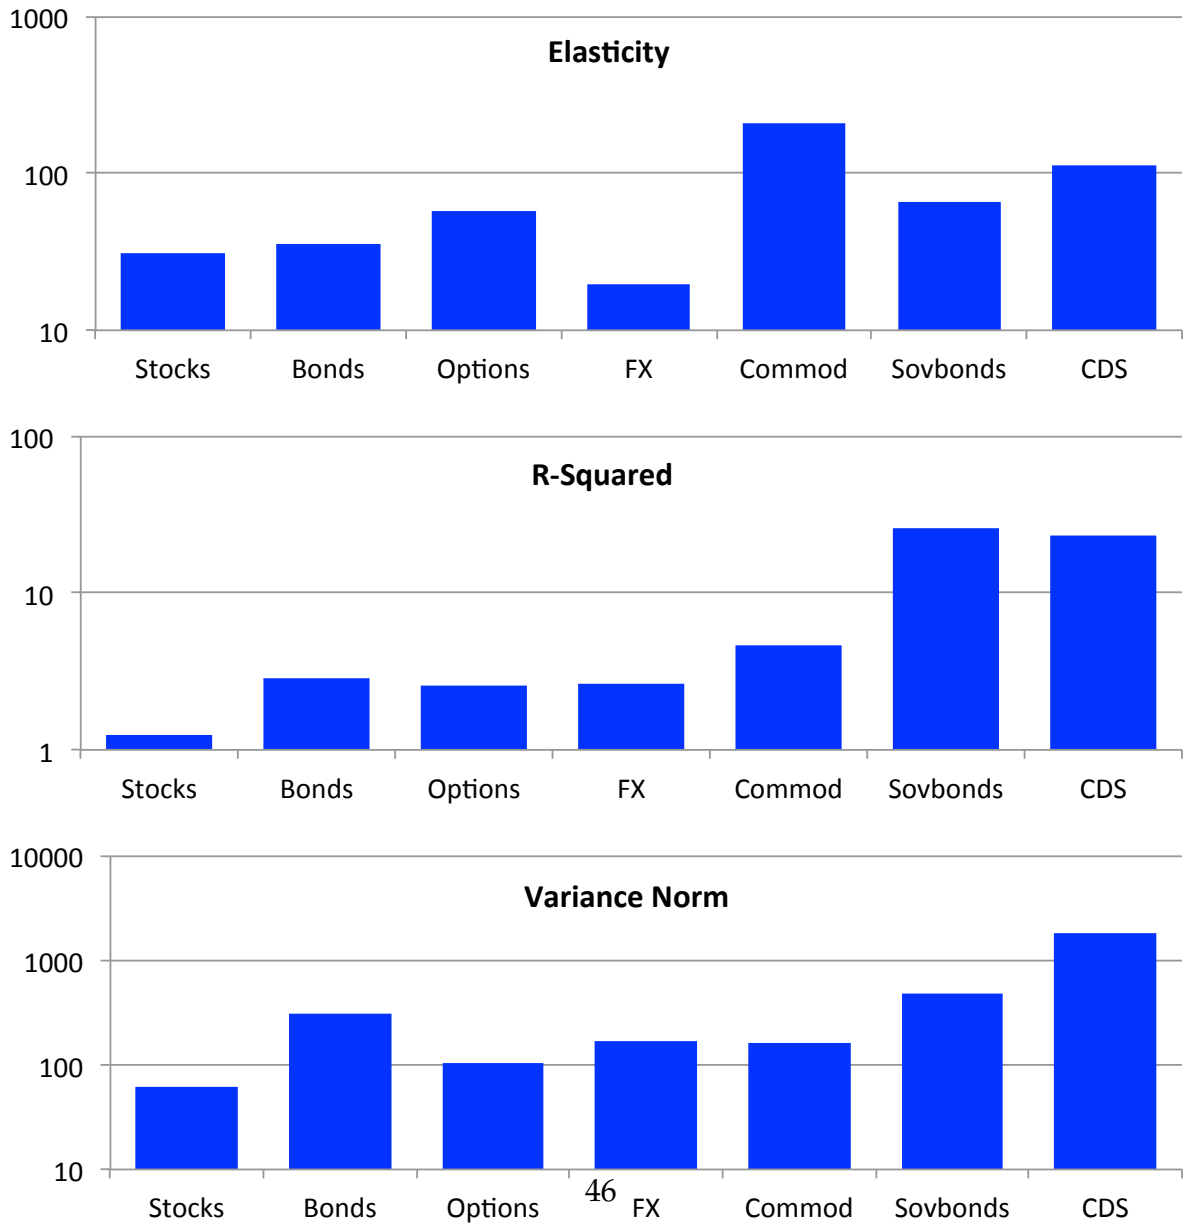
Figure 3: Main Empirical Results. This figure reports the behavior of risk premiums across asset classes (stocks, bonds, options, currencies, commodities, sovereigns, and cds) associated with a change in intermediary distress. We first plot the risk premia elasticity found by running $r_{i,t+k}/E[r_{i,t+k}] = a_i + b_i x_t + \varepsilon_{i,t+k}$ 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 variance $r_{i,t+k}/Var[r_{i,t+k}] = a_i + b_i x_t + \varepsilon_{i,t+k}$. 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.