The Cross-Section of Risk-Taking and Asset Prices

Nuno Coimbra
Paris School of Economics

Rustam Jamilov
London Business School

Hélène Rey
London Business School, CEPR and NBER

September 18, 2019

Abstract

The distribution of institutional investor risk-taking carries significant explanatory power for the cross-section of asset returns. We compute an investor-level Value-at-Risk (VaR) measure - our proxy for ex-ante riskiness - from a structural model with stochastic volatility that we estimate with a particle filter. Our pricing factor - CrossRisk - is then constructed from shocks to the procyclical dispersion of the time-varying VaR distribution. CrossRisk is able to price equity, bond, CDS, options, currency, and commodity market portfolios comparably to numerous single and multi-factor benchmarks. We show that the mechanism behind our results is the extensive margin - dynamic entry and exit of investors into the risky market. A synthetic high-minus-low CrossRisk beta pre-sorted equity portfolio built on the full universe of CRSP firms has an annualized returns spread of 5.8%.

Keywords: Cross-Sectional Asset Pricing, Financial Intermediation, Returns Predictability, Institutional Investors, Value at Risk

1 Introduction

The financial sector has played a central role in the run-up to the 2008 Great Financial Crisis and its aftermath. Aggressive risk taking behavior by financial intermediaries and their pricing of risk are widely viewed as having sowed the seeds of the crisis.

We contribute to this debate in two ways. First we propose a novel way to measure the cross-sectional distribution of ex-ante risk-taking in the financial sector. Second, we show that this distribution is an important factor for the pricing of assets. For this we will proceed in the following sequence. First, we derive our measure of risk-taking from a structural model and identify the necessary ingredients to calculate it for each institution. Secondly, since our measure of risk-taking is an ex-ante measure, we need to extract expected returns and expected volatility of market returns. For this, we apply a particle filter to a structural model of returns with stochastic volatility. Thirdly, armed with the necessary ingredients, we calculate the time-varying cross-sectional distribution of our risk-taking measure. Finally, we construct a novel pricing factor - CrossRisk - from shocks to the second moment of this distribution. We then use it to explain the cross-section of asset returns.

Our main asset pricing test involves the standard 2-step Fama-MacBeth estimation of the price of risk. We find that our single-factor model has strong explanatory power for a variety asset classes with an overall adjusted R^2 of 60-69% (depending on whether we include the market factor or not) and the mean average pricing error (MAPE) of 2.06-2.4% percent annually. Our factor is statistically significant for every asset class we explore. Moreover, it is always significant no matter what external factor (such as liquidity or downside-risk CAPM) we include for model comparison. The factor also implies economically significant return spreads across portfolios with different CrossRisk betas. One-standard deviation difference in CrossRisk betas, based on cross-sectional asset pricing results using aggregated sorted portfolios, implies a 2.22% annual return differential. We also construct 10 CrossRisk beta sorted equity portfolios using the full universe of CRSP firms. The high-minus-low

CrossRisk beta pre-sorted portfolio on average earns an 2.87-5.80% excess return with a Sharpe ratio of roughly 0.25.

In order to inspect the mechanism that drives the performance of the CrossRisk factor, we perform three additional exercises. First, we sequentially drop institutional investor types from the sample and see how the performance of the model in terms of fit changes. The largest role in pricing the assets is played by "independent advisors"; excluding them from the sample reduces the R^2 (raises the MAPE) by 16-20%. Second, in an important exercise, we explore the contribution of the dynamic extensive margin, i.e. entry and exit. We re-construct our factor and re-estimate all regressions while varying the degree of panel balance-ness. Model fit falls dramatically as the panel becomes more balanced, i.e. as we shut down the entry-exit margin. Moreover, the model intercept explodes and becomes highly significant. Finally, we perform attribution analysis on the 3 key components that are used for constructing our SDF: stochastic market volatility, institutional investor returns, and investor market betas. Both volatility and betas play an approximately equal role, while individual returns are less important. What this suggests is that our factor performance is not driven by simply shocks to market volatility (VIX shocks, e.g.) or the cross-section of raw market betas. In other words, our model and the model-implied functional form of the pricing factor are necessary ingredients.

In spirit, this paper is similar to Adrian et al. (2014) and He et al. (2017), who argue that we need to shift attention from measuring the SDF of the average household to measuring a "financial intermediary SDF". We expand this idea by considering the possibility that the marginal financial intermediary pricing the assets is potentially time-varying, so information on the cross-sectional distribution of risk-taking can be informative. Consistent with this idea, formalized in Coimbra and Rey (2017), we find that the time variation of the dispersion of our measure of risk-taking is quite informative as a candidate SDF. In most specifications, dispersion has much more explanatory power than the first moment. Our factor can price not only the traditional asset classes like equity and bonds but also more sophisticated ones like currencies, commodities, options, and credit default swaps (CDS). We therefore argue that there is strong empirical evidence for going beyond the representative financial intermediary but also consider information contained in the time variation of the *full* cross-section of intermediary risk-taking .

This paper is related to multiple strands of the literature. First, it builds on the tradition of applying Value-at-Risk methodology to the analysis of financial intermediary risk-taking behaviour. Papers as Adrian and Shin (2014), Adrian and Boyarchenko (2015) and Coimbra (2018) have studied macroeconomic implications of financial intermediaries facing Value-at-Risk constraints. Importantly, Coimbra and Rey (2017) expanded this to consider the effects of heterogeneous constraints across financial intermediaries showing strong links between financial stability and the cross-sectional behaviour of financial intermediaries. As it pertains to our approach to measure individual risk taking in the data, we build on the structural model of Danielsson et al. (2010).

Our approach to the treatment of volatility and uncertainty expands on the big literature on the role of financial uncertainty in macroeconomics and finance. Some of the more relevant papers are Bloom (2009), Bloom et al. (2018), Jurado et al. (2015), Bekaert et al. (2017), Bekaert et al. (2011), Miranda-Agrippino and Rey (2015), Bruno and Shin (2015), and Coimbra and Rey (2018).

Our main empirical application consists of the pricing of asset returns using a financial intermediary Stochastic Discount Factor (SDF). Methodologically we follow most closely Fama and French (1993) and Fama and French (2015). Some of the more recent contributions to intermediary based asset pricing include Adrian et al. (2014), Adrian and Rosenberg (2008), He et al. (2017), and Koijen and Yogo (2018).

Finally, conceptually our paper might also be seen as providing empirical evidence to the modern macro-finance literature that links the real economy and the financial sector. Papers are too numerous to list but some of the more salient contributions are Brunnermeier and Pedersen (2009), Adrian and Shin (2010), Shleifer and Vishny (1997), Brunnermeier and Sannikov (2014), He and Krishnamurthy (2013), Gromb and Vayanos (2002), Gertler and Kiyotaki (2010). We also provide strong evidence to the importance of heterogeneity within the financial sector, as in Coimbra and Rey (2017).

2 Model

Our structural model follows closely the discrete-time version of Danielsson et al. (2010) and Miranda-Agrippino and Rey (2015). Consider first the special case of an one-risky-asset economy. Financial intermediaries are endowed with net worth N_t which they need to allocate to the risky asset[1]. The risky asset yields a stochastic return R_t with stochastic volatility σ_t. Allocation into the risky asset as a share of N_t is denoted by ω_t, with the remainder being allocated to the riskless asset. The intermediary faces a variant of the Value-at-Risk constraint where the VaR is $\phi_t > 0$ times the forward-looking volatility of portfolio returns. Crucially, note that ϕ_t - the VaR parameter - is varying both across time and the cross-section. Every intermediary solves the following program:

$$\max_{\omega_t} \mathbb{E}_t \left(\omega_t N_t R_{t+1} + (1 - \omega_t) N_t R^f_t \right) \tag{1}$$

$$\text{s.t.} \quad \phi_t \sqrt{\text{Var}_t(\omega_t N_t R_{t+1})} \leq N_t \tag{2}$$

Denoting λ_t as the Lagrange multiplier of the VaR constraint, the FOC is

$$\omega_t \mathbb{E}_t(R_{t+1} - R^f_t) - \lambda_t \phi_t \omega_t \sigma_t = 0$$

which means λ_t is positive as long as the risk premium is positive, in which case the constraint will bind. Note also that N_t drops out from the expression, indicating that asset holdings are linear in net worth. The solution is then given by the binding constraint: $\phi_t \omega_t N_t \sigma_t = N_t$. This gives

[1]For simplicity, we also abstract from debt financing here. Although information about leverage should also be informative of risk taking by financial intermediaries, it is very correlated with portfolio risk and due to data limitations it would require reducing significantly the number of financial intermediaries in our cross-section.

us the following expression for the allocation:

$$\omega_t = \frac{1}{\sigma_t \phi_t} \tag{3}$$

The allocation into the risky asset will then be larger, the smaller is ϕ, the tightness of the constraint and the smaller is the volatility of returns σ_t. Risky asset holdings are linear in net worth, which simplifies the analysis considerably and allows us to think of simple portfolio shares per unit of investment. Expected returns of intermediary i per unit of investment are:

$$\mathbb{E}(R^i_{t+1}) = \omega^i_t E(R_{t+1}) + (1 - \omega^i_t) R^f_t$$

For simplicity of notation, let r^m_{t+1} denote market expected excess returns and r^i_{t+1} the excess returns of intermediary i. We can then rewrite:

$$\mathbb{E}(r^i_{t+1}) = \frac{\mathbb{E}(r^m_{t+1})}{\sigma_t \phi^i_t} \tag{4}$$

It is possible to extend this reasoning to multiple assets as in Danielsson et al. (2010), where ω becomes a vector and there is a large number of risky assets. One can then show that expected excess returns of the portfolio of intermediary i are:

$$\mathbb{E}(r^i_{t+1}) = \beta^i_t \frac{\mathbb{E}(r^m_{t+1})}{\sigma_t \phi^i_t} \tag{5}$$

where r^i_{t+1} is the expected excess return of investor i and β^i_t is the loading on market risk of the optimal portfolio. ϕ^i_t is the VaR constraint parameter to be backed out. As we describe in the next sections, we have exact data on investor-level equity returns, the extracted volatility state and aggregate market returns. We can also estimate investor-level market betas based on portfolio data. Therefore, we can back out ϕ^i_t directly from equation 5.

$$\phi^i_t = \beta^i_t \frac{\mathbb{E}(r^m_{t+1})}{\sigma_t \mathbb{E}(r^i_{t+1})} \tag{6}$$

From the full time-varying distribution of ϕ_{it} we construct CrossRisk - our pricing factor that equals the log-difference of its time-varying standard deviation. Asset pricing results are the same qualitatively if we use the P9010 measure, i.e. the differential between the 90th and the 10th percentile of ϕ_{it}.

3 Factor Construction and Portfolio Data

3.1 Stochastic Volatility Estimation

One of the key ingredients of both the structural model and our empirical exercises is the time-series of stochastic volatility. We apply standard techniques in time-series econometrics to extract this latent factor. We fit the time series of daily S&P 500 returns into a plain-vanilla stochastic volatility (SV) model. Importantly, the return process features both first- and second-moment shocks:

$$R_t = \rho_r R_{t-1} + \exp(\sigma_t)\epsilon_t \tag{7}$$

$$\sigma_t = (1 - \rho_\sigma)\hat{\sigma} + \rho_\sigma \sigma_{t-1} + \eta u_t \tag{8}$$

Where R_t are the market returns (observable), ϵ_t and u_t are i.i.d. shocks, and σ_t is stochastic volatility. This is a latent factor to be estimated.

This model differs from a Generalized Autoregressive Conditional Heteroskedasticity (GARCH) specification, which has been employed in most existent papers on the subject, e.g. Bekaert et al. (2017). The advantages of SV models over GARCH are numerous. Most importantly, SV specifies a process for volatility that relies on volatility-specific shocks. GARCH, on the other hand, is a one-shock model with first-moment residuals only. As a result, volatility movements in period t+1 depend on first moment shocks in period t. Such specification thus does not allow to separately identify first- and second-moment shocks which are economically associated with very different narratives.

Estimation of the parameters that govern our SV statistical model will follow standard procedures outlined in Schorfheide and Herbst (2015) and Fernandez-Villaverde and Rubio-Ramirez (2010). The strategy is a variant of non-linear likelihood maximization and consists of two parts. First, conditional on some guessed parameter vector, we need to evaluate the likelihood function. The likelihood is a highly non-linear and intractable object. Using a linear filter is ill-advised, and we must resort to non-linear filtering. A very accomplished method is the Particle filter. This filter is essentially an (optimal) importance sampling algorithm that, given an assumed measurement error shock distribution, re-weighs candidate observations (particles) based on their probability of fitting the conditional distribution of the data. To further avoid particle degeneracy we employ Gaussian systematic resampling.

Knowing how to evaluate the non-linear likelihood, we can now combine Particle filtering with a likelihood drawing algorithm in order to maximize it. The route we follow in this paper is a Bayesian Markov Chain Monte Carlo estimation algorithm. The idea is to conjecture a loose prior on the parameters and some initial starting vector. We initialize the chain with initial values obtained from a global minimization routine by Chris Sims. Then, construct a very large Markov Chain with each step consisting of 1 Particle filter application to evaluate the likelihood based on a proposal draw and a standard Metropolis-Hastings step that either rejects or accepts that draw. Having estimated the parameters of the process, we can filter the volatility series with a variant of the Particle Smoother.

Table 1 shows the posterior parameter estimates. All estimates are in line with existent research. Figure 1 plots the S&P500 returns series and the extracted stochastic volatility state. It also plots the variance premium series. The variance premium is conventionally defined as the (modulus) of the residuals from a regression of the risk-neutral squared volatility expectation on fundamental volatility. The dependent variable for us is the actual VIX series. The independent variable is our filtered stochastic volatility. Results are in line with the existing findings in Miranda-Agrippino and Rey (2015).

3.2 Institutional Investor Data and Value-at-Risk

The data on institutional investor stock portfolios is from the Thomson Reuters Institutional Holdings Database (s34 file). Those are in turn compiled from quarterly filings of Securities and Exchange Commission Form 13F.7 All institutional investors that provide investment advice and/or exercise investment discretion over accounts in excess of $100 million in total market value must file the form. The dataset includes all major banks, insurance companies, investment advisors, hedge fund managers, etc. The dataset does not cover short positions or non-equity assets such as cash or bonds.

Our equity data is from the Center for Research in Security Prices (CRSP) Monthly Stock Database. Following Fama and French (2015) and Koijen and Yogo (2018), we restrict our sample to ordinary common shares (i.e., share codes 10, 11, 12, and 18) that trade on NYSE, AMEX, and Nasdaq. We further restrict our sample to stocks with non-missing price and shares outstanding.

We merge the institutional holdings data with the CRSP data by CUSIP and drop any holdings that do not match. We compute the dollar holding for each stock that an institution holds as price times shares held. Investor-level total portfolio daily return is defined as a sum of all daily returns weighted by the allocation size. Our level of aggregation is monthly, and investor-level monthly returns series is compounded based on the daily returns. We drop all investors with less than 80 quarters in the merged dataset, and results are the same for the case of 20 or 40 quarters. A more balanced panel will shrink the valuable entry-exit margin and the model fit falls considerably. We discuss this issue in great detail later in the paper.

3.3 The Cross-Sectional Risk Factor

In order to compute investor-level VaR parameters in 6 we require several pieces of information. First, to construct market betas, we run investor-level 30-month rolling regressions of realized excess returns on the market factor (excess return on the S&P500 index). Our risk-free rate proxy is the 3-month U.S. LIBOR. We then winsorize the resulting distribution of betas at the 2.5% and 97.5% levels in each month. Second, since our measure ϕ is forward looking, we require conditional investor and market returns. For both cases, we construct the conditional returns series as predicted values from monthly time-series regressions of returns with one lag. Third, stochastic volatility series is the particle filtered second-moment state of our model in 8. Similarly to conditional returns, we build the conditional volatility series from one-lag time series regressions. All variables are at the monthly frequency so the VaR measure can be interpreted as being monthly too. Having calculated the ϕ_{it} for each investor and each point in time, we rescale the values so they lie on the $[0, 1]$ interval. Our pricing factor will be constructed from the (time-varying) moments of the resulting cross-sectional distribution of ϕ_{it}. In particular, we focus on (shocks to) the second moment, i.e. the standard deviation.

The raw unconditional second moment of the ϕ_{it} distribution is plotted in Figure 2. The figure also plots, in the bottom panel, our pricing factor defined as the log-difference of the dispersion. For convenience, we label our factor CrossRisk. The second moment is highly procyclical, as evidenced from the picture. Factor piece-wise correlation coefficients are reported in 2. Correlations are just high enough such that CrossRisk proxies aggregate financial dynamics. At the same time, those are not too high in which case, say, fundamental volatility would be essentially driving our factor. Indeed, further dissection of the mechanism reveals that cyclicality of the factor is to a large degree driven by variation in the (dispersion) of investor-level betas (unreported). From the list of other noteworthy pricing factors that we are benchmarking, CrossRisk appears to most closely relate with the downside-risk CAPM of Lettau et al. (2014). We provide more formal evidence on attribution using cross-sectional asset pricing tests later in the paper.

3.4 Asset Portfolio and Factor Data

We collect the data on returns of six asset classes from various sources. For equity, we use the Fama and French (2015) 25 size-book-to-market, 10 profitability, and 10 investment portfolios. We also use 10 momentum sorted equity portfolios. For bonds, we use the CRSP's 6 "Fama bond portfolio returns", where portfolios are sorted according to their maturity.

For options, we leverage the data from Constantinides et al. (2013) who construct 18 portfolios of S&P 500 index options sorted on moneyness and maturity, split by contract type (9 put and 9 call portfolios). For currencies, we use the data from Lettau et al. (2014) who build six currency portfolios sorted on the interest rate differential. For commodities, we use 5 portfolios of commodity futures returns from the Commodities Research Bureau. Finally, for credit default swaps (CDS) we follow He et al. (2017) and construct 25 portfolios sorted by spreads using individual name 5-year contracts.

In order to gauge the asset pricing performance of our factor, we source several well-known external factors such as the market excess return (proxied by the excess return on the S&P 500 index), the Fama and French (2015) HML, SMB, investment, and profitability factors, the Carhart (1997) momentum factor, the Pastor and Stambaugh (2003) liquidity factor, the He et al. (2017) primary dealer capital ratio, the Cochrane and Piazzesi (2005) first principal component of the yield curve, and the Lettau et al. (2014) downside-risk CAPM factor. Our analysis is currently only on monthly data; quarterly aggregation will allow us to use more external factors for model comparison such as the Adrian et al. (2014) intermediary leverage.

4 Asset Pricing Tests

4.1 Time Series Regressions

Our main empirical exercise is the two-step Fama-MacBeth procedure. In this section we report results from the first step. We run time-series regressions of institutional investor-level excess returns on the CrossRisk and market factors by each asset class. We report asset-class specific returns as well as CrossRisk and market betas in Table 3. We note instantly the fact that CrossRisk betas are asset-specific: positive for equity and CDS portfolios, and negative for the rest. In other words, the quantity of risk differs significantly across assets.

4.2 Price of Cross-Sectional Risk

Main empirical results of this paper are reported in Tables 4 and 5. First, Table 4 presents estimates from cross-sectional regressions of average portfolio excess returns on the risk exposures. The price of CrossRisk - our factor - is always statistically significant (at least at the 5% level) for each asset class. We highlight that CrossRisk matters even conditionally on the presence of the market factor. Intercept of the pooled model with 113 portfolios is statistically indistinguishable from zero. The factor is also economically significant. For example, a one-standard deviation difference in CrossRisk betas within all asset portfolios (0.98), given the price of risk of 0.189 implies an annualized returns spread of roughly 2.2%. We measure overall model fit primarily through (adjusted) R^2 and mean average pricing errors (MAPE), the latter reported in annualized percent. Model fit is very high within every individual asset class and for the case of all assets pooled together. The MAPE for the 113 pooled assets sample is just 2.06%, while for 55 equity portfolios it is less that 1%. Contrary to He et al. (2017), we don't find that the price of risk is of the same sign for all asset classes. In a future extension, the number of treasury, currency, and commodity portfolios will be increased.

Second, in Table 5 we compare the performance of our CrossRisk factor against other commonly used factors. We focus on the Fama-French 3 and 5 factor models, the momentum factor, the primary dealer capital ratio, downside CAPM, liquidity, and the first principal component of the Treasury yield curve. We report prices of risk together with model fit statistics for specifications with and without CrossRisk. The price of cross-sectional risk is remarkably robust to the inclusion of other commonly used pricing factors, both economically and statistically. Conditional on other factors, inclusion of CrossRisk always increases (decreases) the aR^2 (MAPE). The only model that marginally outperforms our CrossRisk+market baseline is the Fama-French 5 factor model. However, even in that case the two aR^2 are 0.686 and 0.703 and the MAPEs are 2.06% and 2.00%, respectively.

Interestingly, the second moment of the VaR distribution has (by far) the highest explanatory pricing power among the first three. This can be seen from column 3 (CR3). Conditional on the presence of CrossRisk, including (shocks to the) mean and skewness adds only 10% in the marginal aR^2 which approximately equals the lone market factor's contribution. On the other hand, all three moments are strongly statistically significant, which suggests that there is useful information in the full time-varying *distribution* of risk rather than any of the first three moments individually. Quantitatively, the first three moments of the VaR distribution explain essentially as much of the cross-section of prices as the Fama-French 5-factor model, our most robust benchmark.

We visualize our results in Figures 3 and 4. In figure 3 we compare predicted to realized returns in our baseline CrossRisk+Market model (left panel) and the key benchmark: the Fama-French 5-factor model (right panel). We pool all asset classes (113 portfolios) together. In Figure 4 we plot pricing errors of the CrossRisk+Market model for each individual asset class.

5 Inspecting the Mechanism

5.1 Investor Heterogeneity

In this section we inspect the mechanism behind the strong asset pricing power of the CrossRisk factor. We explore three questions. First, which investor types drive the pricing power of the CrossRisk factor? We are able to distinguish 4 types in the Thomson Reuters database: Banks, Insurance Companies, Investment Companies and their Managers (Mutual Funds Management Companies), and Independent Investment Advisors. Investment companies include institutions whose primary source of revenue and/or assets is derived from regulated portfolio management. Specifically, some brokerage firms with mutual fund subsidiaries will be classified as investment companies if the mutual fund management business is deemed by Thomson to make up more than 50 percent of the total 13f assets for that institutional manager. Otherwise, independent money managers will be classified as independent investment advisors.

Table 6 reports the heterogeneity results. In each column, apart from the baseline, we exclude one investor type. We then re-construct the factor based on the new sample and re-run the time-series and cross-sectional regression steps, as before. The change in model fit statistics - aR^2 and MAPE - relative to the baseline, is our key criterion for attributing model performance to that particular investor type. When we drop independent advisors (type 4), aR^2 (MAPE) drop (increase) by around 16-20%. This is the highest change relative to the baseline. The role of independent investors in pricing assets is not surprising. More broadly, this is very consistent with Adrian et al. (2014) and He et al. (2017) who emphasize the financial intermediary SDF. In their sense, financial intermediaries are actively-trading primary dealers or broker-dealer firms. Passively traded mutual funds, insurance companies, or commercial banks play a lesser role in our estimation.

5.2 The Extensive Margin

We now explore the extensive margin's role in driving our results. CrossRisk is the (log-difference) in the standard deviation of the institutional investor Value-at-Risk (VaR) distribution. Variation in dispersion could be due to either the intensive or extensive margin (or both). The extensive margin in this context stands for dynamic entry and exit of intermediaries into the risky market. Procyclicality of the dispersion suggests, everything else equal, procyclical entry/exit. In the present exercise, we try to capture this dimension more formally by varying the degree of panel balance-ness.

In table 7 we report results with samples of institutional investors which have at least 20 (slightly), 40 (moderately), 80 (baseline), 120 (strongly), and 156 (fully) quarters of non-consecutive data. Numbers clearly show how model performance drops considerably with the degree of panel balance-ness. The model's aR^2 falls and both the model intercept and MAPE increase. For example, relative to the baseline, the strongly balanced panel features a 50% drop in aR^2, a 30% rise in MAPE, and an intercept that is large and statistically different from 0. Overall, the extensive margin plays a very important role for our mechanism, consistent with the theory in Coimbra and Rey (2017). It is also worth to report that we need to drop at least some investors in order to reduce noise. In the case of the slightly balanced panel (at least 20 quarters of data), model fit statistics are quantitatively very similar to the baseline.

5.3 Attribution

An important robustness check for our factor is the following question: to what extent does market volatility explain asset prices by itself? In other words, does our factor simply mimic the VIX or S&P500 realized volatility[2]. Table 8 reports our baseline asset pricing results together with special cases where we use counterfactual values for volatility, investor betas, or investor returns in the construction of the underlying Value-at-Risk measure in Equation 6. In particular, in each column we fix them to the respective time-series average values. We see that when we fix either market volatility or investor betas, model fit worsens. But quantitatively, much more so for betas than for volatility. Furthermore, the model intercept becomes large and statistically different from 0 in both cases.

This suggests that the baseline CrossRisk factor does not simply mimic volatility/uncertainty shocks; there is a large role played by the risk adjustment margin (change in market betas) of individual investors. Overall, these results lend support to our model specification and the functional form of the Value-at-Risk measure. In other words, our particular structure that is imposed on volatility, returns, and betas does much better than any individual component.

6 Cross-Risk Beta Sorting

Our final empirical exercise involves the seminal Fama and French (1992) empirical test. We form equity portfolios based on ex-ante Cross-

[2]The role of second-moment (or indeed skewness) risk has been explored and documented by Ang et al. (2006); Adrian and Rosenberg (2008); Adrian et al. (2017); Gabaix et al. (2006); Bansal et al. (2013), among many others

Risk betas. And then we show that the resulting portfolios deliver an average returns spread. Concretely, we run 120-month rolling CUSIP-level regressions of excess returns on the CrossRisk factor. As usual, CrossRisk is defined as the log-difference of the time-varying standard deviation of the institutional investor model-implied Value-at-Risk density. We require at least 60 observations in these rolling regressions. We don't truncate or winsorize the resulting betas, although doing so doesn't do much to the results. For each month, we then construct 10 CrossRisk beta quantiles and compute quantile-specific average excess returns. We also build a synthetic high-minus-low portfolio that is long the 10th and short the 1st quantile of CrossRisk betas.

We report results from this exercise in Table 9. We present quantile-specific average betas (multiplied by 100 for convenience), average excess returns, standard deviations, the Sharpe ratios, and CAPM αs together with the respective t-stats. CrossRisk beta sorting produces a returns and Sharpe ratio spread as seen from column 3. It is also visible on Figure 5, which plots the quantile-specific average returns. Interestingly, there seems to be a concave relationship between CrossRisk betas and returns, although we currently don't explore this dimension. The high-minus-low portfolio (10-1) in Table 9 earns an average excess return of 5.80% with a Sharpe ratio of 0.24. These numbers are comparable to the returns spread due to leverage beta sorts in Adrian et al. (2014), as well as the more standard SMB or HML portfolios. However, as expected, the CAPM alphas are not significant (for each quantile and the 10-1 portfolio), implying that portfolio construction is quite noisy. It is well understood that sorting on covariances makes it very difficult to capture the true pricing kernel variation.

7 Conclusion

This paper contributes to recent debates on the importance of the financial sector in the economy by first proposing a novel way to measure the cross-sectional distribution of ex-ante risk-taking and then by showing that dispersion of this time-varying distribution is empirically an important factor for the pricing of assets. Overall, this paper provides evidence in support of heterogeneity in the cross-section of financial intermediaries.

First, we derived our measure of risk-taking from a structural model and identified the necessary ingredients to calculate it for each institution. Secondly, we extracted expected returns and expected volatility of market returns by applying a particle filter to a structural model of returns with stochastic volatility. Combining this with information on equity portfolio weights of institutional investors, we calculate the time-varying cross-sectional distribution of our risk-taking measure.

Finally, we construct the CrossRisk factor: a candidate SDF that is the log-difference of the second moment of this distribution. We show that our single-factor model is very informative in pricing a large cross-section of traditional (equity and bonds) and more complex assets (CDS, options, currency, commodities). In particular, the baseline CrossRisk+market model explains 68.6% of the cross-sectional variation in risk quantity (betas) of 113 traditional and exotic asset portfolios. We dissect the mechanism that drives our results and find that independent financial advisors and the dynamic extensive margin play a very important role in our asset pricing performance.

Appendix: Figures and Tables

Figure 1: Market Returns and Stochastic Volatility

This figure plots the time-series of the returns on the S&P500 Index, the (log of) estimated stochastic volatility, and the estimate of a variance premium. Model 8 likelihood is filtered with a Particle Filter that is run on 2500 particles. The likelihood is maximized with a variant of a Bayesian Markov Chain Monte Carlo algorithm with 2000 iterations. The resulting volatility series is filtered (Dashed line) and smoothed with a particle smoother (Bold). Variance premium is the (modulus) of the residuals from a regression of the risk-neutral squared volatility expectation on fundamental volatility. The dependent variable is the CBOE Volatility Index. Data frequency is daily, from 1980:02:01 to 2018:31:12. Returns and volatility are annualized.

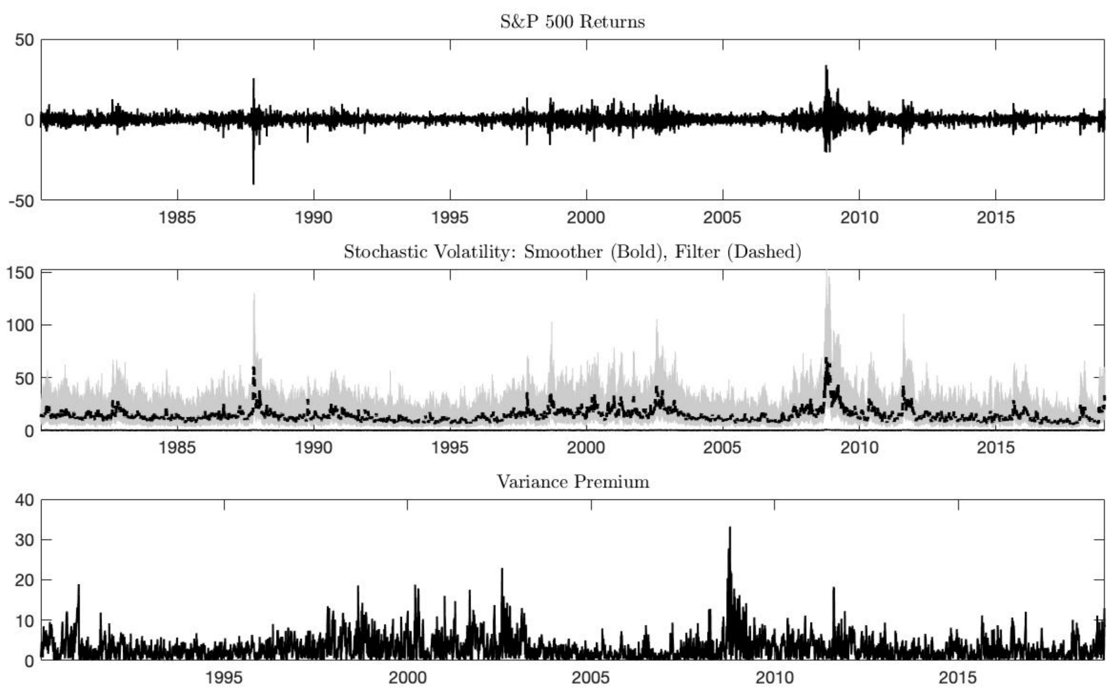

Figure 2: Value-at-Risk Dispersion and the CrossRisk Pricing Factor

This graph shows the dispersion (standard deviation) of the time-varying distribution of model-implied Value-at-Risk (VaR) of individual institutional investors (upper panel) and our pricing factor - CrossRisk - defined as the month-on-month log-difference in the VaR dispersion. VaR dispersion is in percent, annualized. Data frequency is monthly, ranging from 1982:1 to 2018:12

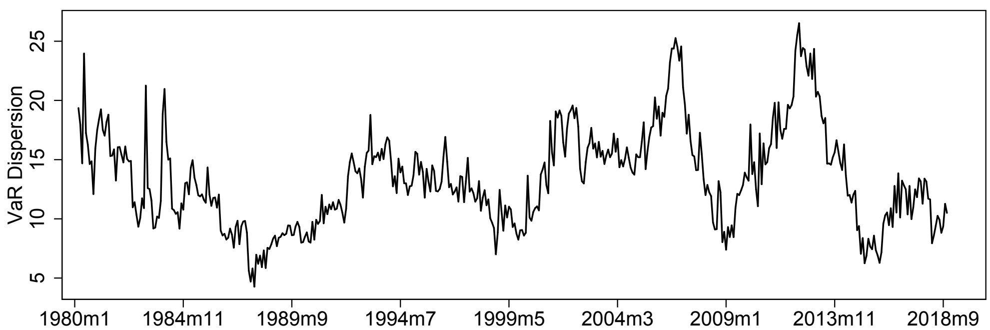

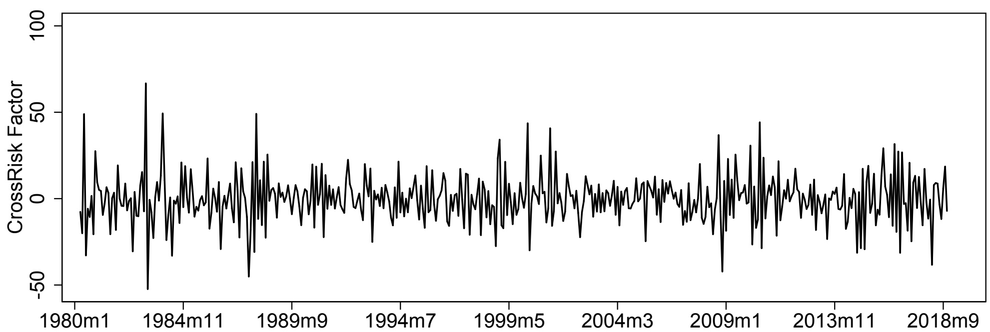

Figure 3: Predicted vs Realized Returns: CrossRisk and FF-5

This graph plots the actual vs predicted expected returns (pricing errors) for all 6 asset classes: 25 size and book-to-market, 10 momentum, 10 profitability, and 10 investment sorted equity portfolios as well as 6 Treasury bond, 24 CDS, 18 options, 5 currency, and 5 commodities portfolios. Left panel represents the CrossRisk + market factor model. CrossRisk is the log-difference of the time-varying standard deviation of institutional investor model-implied Value-at-Risk. The right panel represents the Fama-French 5-factor model (Fama and French, 2015). Adjusted R-squared for the two panels are 68.6% and 70.3% respectively. MAPE are 2.06% and 2.00% annualized, respectively. Sample period of 1982:1-2018:12, data is monthly. Returns are in percent, per year. Shaded lines are 95% confidence intervals. Dashed lines are 45-degree fits. Labels are portfolio identifiers: 1-25 are size and book-to-market, 25-35 momentum, 35-45 investment, and 45-55 profitability sorted equity portfolios; 56-61 are Treasury bond, 62-85 CDS, 86-90 currency, 91-95 commodities, 96-113 options portfolios.

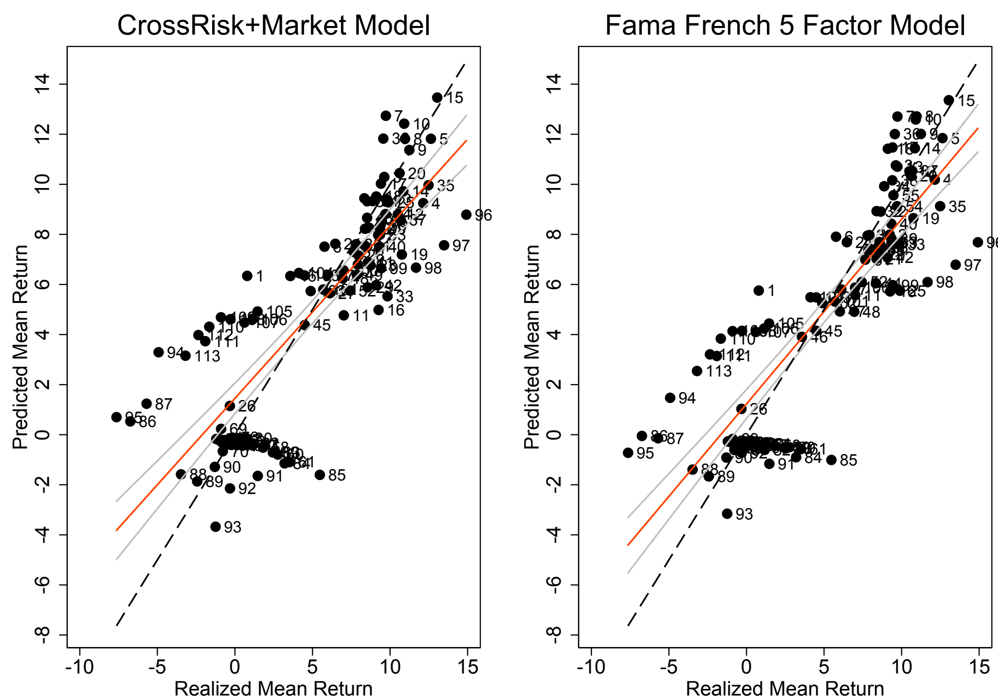

Figure 4: Predicted vs Realized Returns: All Assets

This graph plots the actual vs predicted expected returns (pricing errors) for all 6 asset classes individually: 25 size and book-to-market portfolios separately, all equity that also include 10 momentum, 10 profitability, and 10 investment sorted equity portfolios; 6 Treasury bond, 24 CDS, 18 options, 5 currency, and 5 commodities portfolios. The model in each asset class includes the CrossRisk and market factors. CrossRisk is the log-difference in the cross-sectional standard deviation of institutional investor model-implied value at risk. Returns are in percent, per year. Shaded lines are 95% confidence intervals. Dashed lines are 45-degree fits. Monthly data is 1982:1-2018:12.

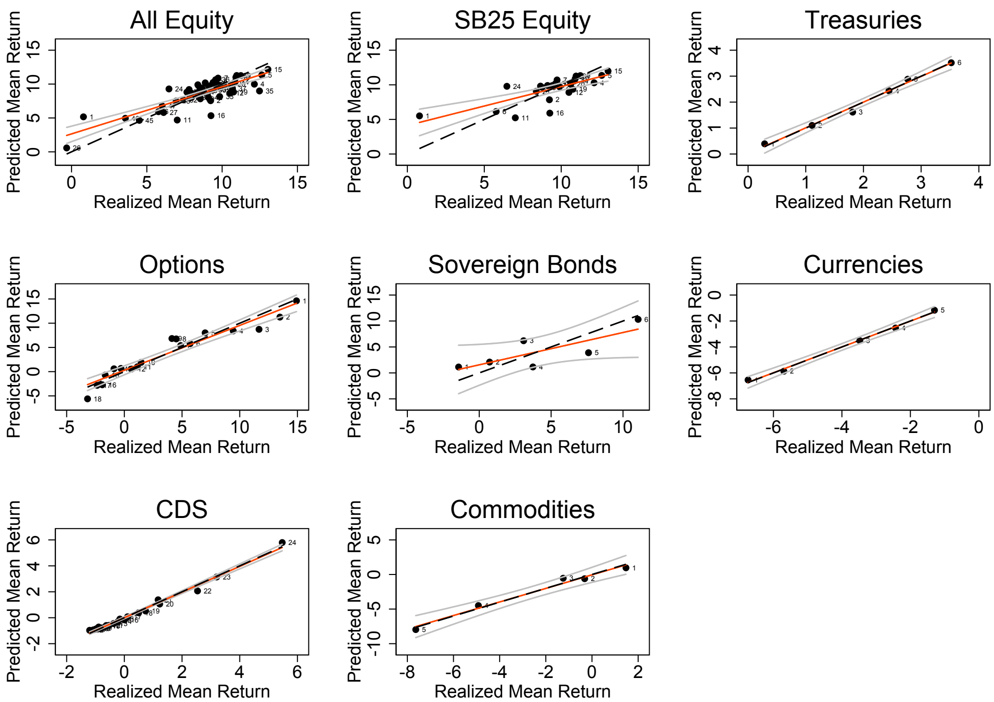

Figure 5: CrossRisk Beta Sorted Portfolio Returns

This graph plots average returns on 10 CrossRisk beta sorted equity portfolios. To construct these portfolios, we run 120-month rolling CUSIP-level regressions of excess returns on the CrossRisk factor for each firm in the CRSP universe. We require at least 60 observations in the rolling regressions. CrossRisk is defined as the log-difference of the time-varying standard deviation of the institutional investor model-implied Value-at-Risk. Returns are in percent, annualized. The monthly data is 1986m1-2018m12.

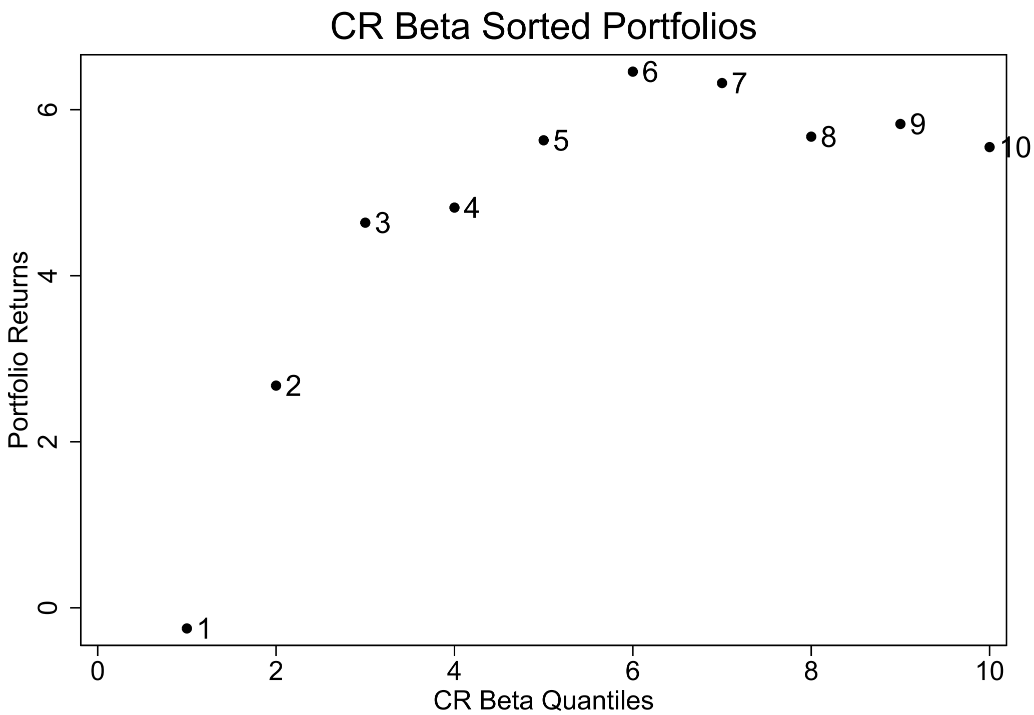

Table 1: Stochastic Volatility Model Posterior Parameter Estimates

This table reports posterior parameter estimates of the stochastic volatility model in (8). Likelihood filtering is done with the particle filter and 2000 particles. Likelihood maximization is done with a loose Bayesian prior and 2500 Markov Chain Monte Carlo iterations. We discard the first 500 draws. Data is daily and ranges from 1980:02:01 to 2018:31:12.

ρ	$\hat{\sigma}$	ρ_{σ}	η
0.99	-4.79	0.96	0.11
(0.99,0.99)	(-4.71,4.82)	(0.95,0.97)	(0.10,0.12)

Table 2: CrossRisk Factor Correlations

This table reports piece-wise correlation coefficients of our factor - CrossRisk - with other noteworthy time series. VaR mean, dispersion, and skewness are the first three moments, respectively, of the model-implied time-varying distribution of Value-at-Risk. Volatility is the extracted stochastic volatility state, market is realized return on the S$P500 index, HML and SMB are the Fama and French (1993) size and book/market factors. MoM is the momentum factor from Kenneth French's website. CP is the first principal component of the yield curve (Cochrane and Piazzesi, 2005). HKM is the primary dealer capital ratio factor (He et al., 2017). PS is the liquidity factor (Pastor and Stambaugh, 2003). LMW is the downside-risk CAPM (Lettau et al., 2014). Monthly data is 1982m1-2018m12.

	VaR Dispersion	CrossRisk Factor	VaR Mean	VaR Skewness
VaR Dispersion	1.00			
CrossRisk Factor	0.21	1.00		
VaR Mean	0.01	0.11	1.00	
VaR Skewness	0.12	0.19	0.00	1.00
Volatility	-0.20	-0.06	-0.22	-0.02
Market	0.02	0.12	0.13	0.08
HML	0.13	0.06	0.07	0.01
SMB	0.05	0.06	0.14	0.06
MoM	0.02	0.04	-0.16	0.00
CP	-0.29	-0.01	-0.02	0.03
HKM	0.00	0.09	0.17	0.10
PS	0.06	0.04	-0.07	-0.03
LMW	0.19	0.18	0.20	0.17

Table 3: Time Series Regressions and Summary Statistics

This table presents the time-series regression of institutional investor level returns on the CrossRisk and the market factors. CrossRisk is defined as the log-difference in the time-varying standard deviation of the institutional investor model-implied Value-at-Risk distribution. Returns are in percent, annualized. CrossRisk betas are multiplied by 100 for convenience. The monthly data is 1982m1-2018m12.

	All Equity	CDS	Options	Treasuries	Currencies	Commodities	All Assets
Mean Return	8.67	0.17	3.81	1.99	-3.93	-2.53	4.68
Sd Return	0.72	0.46	1.62	0.34	0.65	1.06	1.51
Mean CR Beta	35.78	1.86	-30.25	-19.96	-30.29	-48.53	8.45
Sd CR Beta	127.62	7.56	30.68	14.17	73.01	122.17	98.04
Mean Market Beta	1.03	-0.03	0.87	0.00	0.04	0.09	0.64
Sd Market Beta	0.15	0.04	0.13	0.01	0.05	0.04	0.49
Assets	55	24	18	6	5	5	113

Table 4: Cross-Sectional Asset Pricing Tests by Asset Class

This table presents the price of risk on the CrossRisk and excess market returns factors. CrossRisk is defined as the log-difference in the time-varying standard deviation of the institutional investor model-implied Value-at-Risk distribution. Risk prices are the slopes of linear cross-sectional regressions of average portfolio excess returns on risk exposures. Betas are estimated from the first-stage time-series regression. Mean absolute pricing error (MAPE) is in annual percentage points. Shanken SEs are in the parentheses. The monthly data is 1982m1-2018m12

	SB25	All Equity	CDS	Options	Treasuries	Currencies	Commodities	All Assets	
CrossRisk	0.0970***	0.117***	0.0717*	0.470***	-0.768***	-0.135**	-0.178**	0.189***	0.189***
	(0.0309)	(0.0145)	(0.0356)	(0.112)	(0.0387)	(0.0180)	(0.0279)	(0.0195)	(0.0194)
Market	-0.337	-0.218	-3.013***	2.685***	-5.504**	1.895**	-6.279**	0.614***	0.597***
	(0.343)	(0.157)	(0.116)	(0.404)	(0.999)	(0.242)	(0.676)	(0.0476)	(0.0261)
Intercept	1.064***	0.904***	-0.0795***	-1.874***	0.0140	-0.447***	0.289*	-0.0170	
	(0.362)	(0.158)	(0.00290)	(0.380)	(0.00990)	(0.0126)	(0.0724)	(0.0416)	
Assets	25	55	24	18	6	5	5	113	113
aR2	0.522	0.683	0.988	0.917	0.985	0.990	0.959	0.686	0.826
RMSE	0.144	0.117	0.0147	0.135	0.0121	0.0186	0.0625	0.245	0.244
MAPE	1.18	0.96	0.12	1.16	0.07	0.13	0.45	2.06	2.04

Table 5: Cross-Sectional Asset Pricing Tests: Model Comparison

This table compares model performance with and without CrossRisk across a wide range of commonly used pricing factors. CrossRisk is defined as the log-difference in the time-varying standard deviation of the institutional investor model-implied Value-at-Risk distribution. CR Mean and Skew are the mean and skewness of the same VaR distribution. FF3 is the Fama and French (1992) model with market, SMB, and HML factors; FF5 is the Fama and French (2015) model with RMW (profitability) and CMA (investment) factors, in addition to SMB, HML, and the market; MoM is the momentum factor from Kenneth French's website. PC is the first principle component of the yield curve, i.e. the Cochrane and Piazzesi (2005) factor; HKM is the primary dealer capital ratio (He et al., 2017); LMW is the downside CAPM model (Lettau et al., 2014); and PS is the liquidity factor (Pastor and Stambaugh, 2003). Betas are estimated from the first-stage time-series regression. Mean absolute pricing error (MAPE) is in annual percentage points. Shanken SEs are in the parentheses. The monthly data is 1982m1-2018m12

	CR	CR+MKT	CR3	FF3	MoM	PC	FF5	HKM	LMW	PS
CR	0.152***	0.189***	0.165***	0.182***	0.149***	0.159***	0.107***	0.164***	0.0907***	0.0983***
	(0.0094)	(0.0195)	(0.0248)	(0.0228)	(0.0103)	(0.0113)	(0.0372)	(0.0202)	(0.0123)	(0.0332)
Market		0.614***		0.605***			0.600***			
		(0.0476)		(0.0506)			(0.0472)			
CR Mean			0.103***							
			(0.0211)							
CR Skew			-8.031***							
			(2.119)							
SMB				0.242***			0.310***			
				(0.0700)			(0.0669)			
HML				0.314***			0.179*			
				(0.0806)			(0.104)			
MOM					0.382***					
					(0.104)					
PC1						-0.274				
						(0.397)				
RMW							0.344***			
							(0.104)			
CMA							0.440***			
							(0.117)			
Dealer Capital								0.00639***		
								(0.00132)		
Downside Risk									0.583***	
									(0.0547)	
Liquidity										0.0142**
										(0.00679)
Intercept	-0.0131	-0.0170	0.0239	-0.0180	0.00402	-0.0458	-0.0279	-0.0264	0.0544	-0.0407
	(0.0368)	(0.0416)	(0.0415)	(0.0414)	(0.0348)	(0.0592)	(0.0382)	(0.0403)	(0.0415)	(0.0419)
Assets	113	113	113	113	113	113	113	113	113	113
aR2	0.608	0.686	0.697	0.694	0.560	0.611	0.724	0.638	0.711	0.632
RMSE	0.273	0.245	0.240	0.241	0.289	0.272	0.229	0.262	0.235	0.264
MAPE	2.40	2.06	2.20	2.03	2.57	2.39	1.92	2.27	2.08	2.32
aR2 w/o CR				0.572	0.102	0.106	0.703	0.515	0.555	0.583
RMSE w/o CR				0.285	0.413	0.412	0.238	0.304	0.291	0.282
MAPE w/o CR				2.51	4.16	4.45	2.00	2.77	2.69	2.54

Table 6: Cross-Sectional Asset Pricing Tests: Investor Heterogeneity

This table reports the price of risk and model fit in the baseline specification as well as for special cases when we exclude individual institutional investor types. For example, "ex-banks" stands for a model where we construct the CrossRisk factor from a sample of investors that excludes banks. We then re-run the time-series and cross-sectional regression steps, as before. Mean absolute pricing error (MAPE) is in annual percentage points. Shanken SEs are in the parentheses. The monthly data is 1982m1-2018m12.

	Baseline	ex-Banks	ex-Insurance	ex-Managers	ex-Advisors
CrossRisk	0.152***	0.207***	0.230***	0.146***	0.174***
	(0.0094)	(0.0135)	(0.0143)	(0.0092)	(0.0142)
Intercept	-0.0131	-0.003	0.00564	-0.0266	0.0258
	(0.0368)	(0.0338)	(0.0333)	(0.0373)	(0.0390)
Assets	113	113	113	113	113
aR2	0.608	0.532	0.541	0.598	0.493
RMSE	0.273	0.298	0.296	0.277	0.31
MAPE	2.40	2.57	2.51	2.44	2.80

Table 7: Cross-Sectional Asset Pricing Tests: Extensive Margin Role

This table reports the price of risk and model fit in the baseline specification as well as for special cases when we we vary the degree of panel balance-ness. In the baseline, we only include institutional investors with at least 80 (non-consecutive) quarters of data. Slightly, moderately, strongly, and fully balanced samples include investors with at least 20, 40, 120, and 156 quarters, respectively. We then re-construct our factor and re-run the time-series and cross-sectional regression steps, as before. Mean absolute pricing error (MAPE) is in annual percentage points. Shanken SEs are in the parentheses. The monthly data is 1982m1-2018m12.

	Slightly Balanced	Moderately Balanced	Baseline	Strongly Balanced	Fully Balanced
CrossRisk	0.0829***	0.0859***	0.152***	0.266***	0.00464
	(0.0061)	(0.0064)	(0.0094)	(0.0514)	(0.0592)
Intercept	-0.0018	0.0112	-0.0131	0.127***	0.390***
	(0.0403)	(0.0397)	(0.0368)	(0.0414)	(0.0447)
Assets	113	113	113	113	113
aR2	0.599	0.58	0.608	0.318	-0.009
RMSE	0.276	0.283	0.273	0.36	0.438
MAPE	2.568	2.652	2.4	3.168	4.68

Table 8: Cross-Sectional Asset Pricing Tests: Attribution

This table reports the price of risk and model fit in the baseline specification as well as for special cases where we fix the conditional market volatility, investor market betas, and investor returns. For example, Ex Volatility means that we construct our investor-level Value-at-Risk measure while fixing market volatility to its time-series average. We then re-run the time-series and cross-sectional regression steps, as before, using the (log-difference) of the standard deviation of the distribution of this new measure. Mean absolute pricing error (MAPE) is in annual percentage points. Shanken SEs are in the parentheses. The monthly data is 1982m1-2018m12.

	Baseline	Ex Volatility	Ex Betas	Ex Returns
CrossRisk	0.152***	0.133***	0.249***	0.116***
	(0.0094)	(0.0158)	(0.0803)	(0.0103)
Intercept	-0.0131	0.0934***	0.354***	0.0714*
	(0.0368)	(0.0337)	(0.0420)	(0.0372)
Assets	113	113	113	113
aR2	0.61	0.38	0.12	0.55
RMSE	0.27	0.34	0.41	0.29
MAPE	2.40	3.19	4.21	0.23

Table 9: CrossRisk Beta Sorted Portfolio Returns

This table presents CrossRisk beta pre-sorting of all stocks in CRSP. We run 120-month rolling CUSIP-level regressions of excess returns on the CrossRisk factor. CrossRisk is defined as the log-difference of the time-varying standard deviation of the institutional investor model-implied Value-at-Risk. We require at least 60 observations in the rolling regressions. The table reports quantile-specific betas, mean excess returns, standard deviation, Sharpe ratio, and CAPM α. The 10-1 portfolio is short the first and long the tenth quantile over time. Returns are in percent, annualized. The monthly data is 1986m1-2018m12.

CR Beta Quantiles	CR Beta	$\mathbb{E}$(R)	σ(R)	$\frac{\mathbb{E}(R)}{\sigma(R)}$	CAPM α	CAPM α - tstat
1	-19.26	-0.25	68.54	0.00	-3.4	-1.4
2	-7.02	2.68	52.48	0.05	-2.4	-0.9
3	-2.49	4.64	46.21	0.10	0	-0.1
4	0.57	4.82	44.06	0.11	0	0.1
5	3.12	5.63	42.22	0.13	1.2	0.6
6	5.51	6.46	42.32	0.15	1.3	0.9
7	7.96	6.32	43.00	0.15	1.2	0.7
8	10.81	5.68	44.76	0.13	0	0.2
9	14.90	5.83	49.26	0.12	0	0.1
10	25.79	5.55	62.61	0.09	1.4	0.9
10-1	45.05	5.80	24.71	0.24	4.8	1.1

References

Adrian, T. and N. Boyarchenko (2015): "Intermediary leverage cycles and financial stability," *Staff Reports, Federal Reserve Bank of New York*, 567.

Adrian, T., N. Boyarchenko, and D. Giannone (2017): "Vulnerable Growth," *Federal Reserve Bank of New York Staff Report*, 794.

Adrian, T., E. Etula, and T. Muir (2014): "Financial Intermediaries and the Cross-Section of Asset Returns," *Journal of Finance*, 69(6), 2557–2596.

Adrian, T. and J. Rosenberg (2008): "Stock Returns and Volatility: Pricing the Short-Run and Long-Run Components of Market Risk," *Journal of Finance*, 63(6).

Adrian, T. and H. S. Shin (2010): "Liquidity and Leverage," *Journal of Financial Intermediation*, 19(3), 418–437.

——— (2014): "Procyclical Leverage and Value-at-Risk," *Review of Financial Studies*, 27(2), 373–403.

Ang, A., R. Hodrick, Y. Xing, and X. Zhang (2006): "The Crossâ€Section of Volatility and Expected Returns," *Journal of Finance*, 61.

Bansal, R., D. Kiku, I. Shaliastovich, and A. Yaron (2013): "Volatility, the Macroeconomy, and Asset Prices," *Journal of Finance*.

Bekaert, G., E. C. Engstrom, and N. R. Xu (2011): "Whats Vol Got to Do with It," *Review of Financial Studies*, 24(1).

——— (2017): "The Time Variation in Risk Appetite and Uncertainty," *Manuscript*.

Bloom, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77(3), 623–685.

Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten, and S. J. Terry (2018): "Really Uncertain Business Cycles," *Econometrica*, Forthcoming.

Brunnermeier, M. and L. Pedersen (2009): "Market Liquidity and Funding Liquidity," *Refiew of Financial Studies*, 22, 2201–2238.

Brunnermeier, M. and Y. Sannikov (2014): "A Macroeconomic Model with a Financial Sector," *American Economic Review*, 104(2), 379–421.

Bruno, V. and H. Shin (2015): "Capital flows and the risk-taking channel of monetary policy," *Journal of Monetary Economics*, 71.

Carhart, M. (1997): "On Persistence in Mutual Fund Performance," *Journal of Finance*, 52(1).

Cochrane, J. and M. Piazzesi (2005): "Bond Risk Premia," *American Economic Review*, 95.

Coimbra, N. (2018): "A dynamic model of sovereign debt and banking leverage," *Manuscript*.

Coimbra, N. and H. Rey (2017): "Financial Cycles with Heterogeneous Intermediaries," *NBER Working Paper*, 23245.

——— (2018): "Financial Cycles and Credit Growth across Countries," *AEA Papers and Proceedings*, 108.

Constantinides, G., J. Jackwerth, and A. Savov (2013): "The puzzle of index option returns," *Review of Asset Pricing Studies*, 3.

Danielsson, J., H. S. Shin, and J-PZigrand (2010): "Risk Appetite and Endogenous Risk," *AXA Working Paper Seris 2*, 647.

Fama, E. and K. French (1992): "The Cross Section of Expected Stock Returns," *Journal of Finance*, 47(2).

——— (1993): "Common risk factors in the returns on stocks and bonds," *Journal of Financial Economics*, 33.

——— (2015): "A five-factor asset pricing model," *Journal of Financial Economics*, 116(1).

Fernandez-Villaverde, J. and J. Rubio-Ramirez (2010): "Macroeconomics and Volatility: Data, Models, and Estimation," *NBER Working Paper*, 16618.

Gabaix, X., P. Gopikrishnan, V. Plerou, and H. E. Stanley (2006): "Institutional Investors and Stock Market Volatility," *Quarterly Journal of Economics*, 121(2).

Gertler, M. and N. Kiyotaki (2010): "Financial Intermediation and Credit Policy in Business Cycle Analysis," *Handbook of Monetary Economics*, 3, 547–599.

Gromb, D. and D. Vayanos (2002): "Equilibrium and welfare in markets with financially constrained arbitrageurs," *Journal of Financial Economics*, 66.

He, Z., B. Kelly, and A. Manela (2017): "Intermediary asset pricing: New evidence from many asset classes," *Journal of Financial Economics*, 126.

He, Z. and A. Krishnamurthy (2013): "Intermediary Asset Pricing," *Journal of Financial Economics*, 103(2), 732–770.

Jurado, K., S. Ludvigson, and S. Ng (2015): "Measuring Uncertainty," *American Economic Review*, 105(3), 1177–1216.

Koijen, R. and M. Yogo (2018): "A Demand System Approach to Asset Pricing," *Journal of Political Economy*, Forthcoming.

Lettau, M., M. Maggiory, and M. Weber (2014): "Conditional risk premia in currency markets and other asset classes," *Journal of Financial Economics*, 114(2).

Miranda-Agrippino, S. and H. Rey (2015): "World Asset Markets and the Global Financial Cycle," *NBER Working Paper*, N 21722.

Pastor, L. and R. Stambaugh (2003): "Liquidity risk and expected stock returns," *Journal of Political Economy*, 111(3).

Schorfheide, F. and E. Herbst (2015): "Bayesian Estiamtion of DSGE Models," *Princeton University Press*.

Shleifer, A. and R. W. Vishny (1997): "The Limits of Arbitrage," *Journal of Finance*, 52(1).