

Financial Crises and Risk Premia*

Tyler Muir

Email: tyler.muir@yale.edu

March 2014

Abstract

I analyze the behavior of risk premia in financial crises, wars, and recessions. I document that risk premia increase substantially in financial crises but not in the other episodes. However, drops in consumption and consumption volatility are similar across financial crises and recessions and are largest during wars, so consumption based models can't explain this variation. Further, I find the equity of the financial sector forecasts returns. The results show that the health of the financial sector is uniquely important for understanding why risk premia vary. I calibrate an intermediary asset pricing model and show it can match the data.

*Yale School of Management, Department of Finance. I would like to thank Andrea Eisfeldt, Dimitris Papanikolaou, Ravi Jagannathan, Gary Gorton, Alan Moreira, Hyun Shin, Sydney Ludvigson, John Cochrane and Nick Roussanov, as well as many seminar participants at many universities and conferences. I would especially like to thank Arvind Krishnamurthy for many valuable discussions. I thank the International Center for Finance at Yale for providing credit spread data, as well as Robert Barro, Jose Ursua, Ken Rogoff, and Carmen Reinhardt for making their data publicly available. I thank Colin Lu for research assistance. All errors are mine.

1 Introduction

Why do risk premia vary over time? This paper explores the behavior of risk premia across financial crises, recessions, and wars and uses variation in the data to distinguish between asset pricing models. I use data on consumption, dividend yields, stock returns, and credit spreads for over 140 years and 14 countries, yielding 180 recessions, 45 financial crises, and 20 major wars. First, I document that risk premia spike dramatically in financial crises – defined specifically as a banking panic or banking crisis – but rise only slightly in recessions or wars. The large increase in risk premia during financial crises puts explaining these episodes to the forefront of the asset pricing literature. To my knowledge, this is the first paper to study and characterize the behavior of risk premia across financial crises, and is the first to compare financial crises to these other events. Second, I examine the ability of leading macro asset pricing models to explain these facts. I find the overall drop in consumption and increase in consumption volatility across financial crises, recessions, and wars is either roughly flat or has the wrong sign, meaning the variation in risk premia can not be explained by standard consumption based asset pricing models. In contrast, theories where risk premia are a function of the health of the financial sector seem to fit the patterns in the data well. The evidence points to the health of the financial sector as being important for aggregate asset price fluctuations. I then calibrate an intermediary asset pricing model and show that the model can replicate the main features of risk premia during financial crises.

Figure 1 plots the increase or decrease of variables of interest across each of the events and provides the main results of the paper. The details, data description, and formal statistical analysis are all left to the main text. Panel A documents large increases in dividend yields and credit spreads, both typical measures of risk premia, during financial crises. Throughout the paper I use the term high risk premia to mean an asset with high expected return, high discount rate, or low price relative to fundamentals or expected cash flows. While these measures spike in financial crises, other episodes show only minor movements in risk premia. Later sections give these same results via regressions. Panel B shows that the key state variables implied by leading macro asset pricing models – the drop in consumption and the conditional volatility of consumption – can not account for the behavior of risk premia across episodes. These variables do not vary meaningfully across financial crises and typical recessions, and each changes the most during wars. The failure of these models can not be saved by changing parameters in their calibrations because the relationship between state variables and risk premia is either flat (across financial crises and recessions) or has the wrong sign (across financial crises and wars). This poses a major challenge for consumption

based asset pricing models and in later sections I show that the calibrated versions of these models substantially miss the behavior of risk premia in financial crises. The facts are instead consistent with models where risk premia depend on the health of the financial sector. In these models, risk premia will naturally be significantly higher in financial crises compared to typical recessions because intermediary balance sheets are more strained.

There are several potential objections to the results that financial crises have much higher risk premia than other events. The first and main objection is that the conclusions are tautological. Of course prices fall and risk premia rise during financial crises if a “crisis” is defined ex-post as a large decline in asset values. In fact, the financial crisis dates are defined as a systemic event – a major bank run or bank failure – and are not defined based on what happened to prices. Therefore, they are not defined ex-post based on drops in asset values. However, this objection also partly misses the point. Regardless of the dating convention, the fact remains that models should explain these episodes. For example, if the habits model of Campbell and Cochrane (1999) is a good description of the world, we can define changes in prices ex-post, but it must still be the case that consumption (or “surplus consumption,” which is consumption relative to habit) drops substantially whenever prices drop substantially. I find that there is variation in risk premia across financial crises and recessions, but little or no variation in these consumption state variables. Equally important, I find that there is no variation in risk premia across recessions and wars despite substantial variation in the consumption state variables. Therefore, the exercise is relevant in distinguishing these models even if one is skeptical that these dates are chosen based on systemic crises and not just price drops. Importantly, I am also comparing the behavior of returns and risk premia in financial crises to recessions, not to normal times. Recessions are also times when the economy is hit by economic shocks of similar magnitudes, as I will show. It is crucial that the prices drop *relative* to dividends in financial crises. The key feature of financial crises relative to typical recessions is the change in the discount rate, not the change in cash flows. In fact, typical recessions look roughly like “bad cash flow news” but no significant change in discount rates. Financial crises, on the other hand, are similar to recessions in terms of their “bad cash flow news” but appear to have a large additional discount rate effect. One should think of the exercise of comparing to recession as a “difference in difference” approach where in both cases the economy is hit with similar shocks to cash flows and fundamentals, but in one episode the financial sector is particularly affected. The larger price drops in financial crises are followed by price rebounds, not dividend changes, while the price drops in recessions feature no large reversal. Similarly, wars feature large negative cash flow shocks and large negative stock returns but show little movement in

discount rates.

The second main objection is that dividend yields and credit spreads are poor measures of risk premia. Maybe the spike in dividend yields in crises is really about expected dividend growth in those episodes, even though unconditionally we don't see dividend yields strongly forecast future dividend growth. I show that dividend yields are valid measures of expected returns because they forecast returns and not cash flows, both *conditional* on financial crises as well as unconditionally. The results show that the standard result that dividend yields forecast future returns and not dividend growth hold during crises and recessions as well. These results are shown in the appendix which conducts a number of other robustness checks. Moreover, I show that *realized* returns after financial crises are very large at around 20%, which is a non-parametric way of showing that risk premia in these episodes are indeed large. This confirms that the key feature of financial crises relative to typical recessions is the discount rate effect or spike in risk premia, as prices drop but then later rebound. In contrast, in recessions subsequent realized returns are not abnormally high. In fact, the extra drop in return in financial crises vs recessions is completely reversed several years out, so there is very little difference in long term cash flow news but substantial difference in discount rates across these episodes.

I then calibrate a simple intermediary asset pricing model and show that it can largely match the behavior of risk premia in financial crises. The model is based on He and Krishnamurthy (2012a) and Brunnermeier and Sannikov (2012). The model features intermediation frictions so that the stochastic discount factor (SDF) depends on intermediary equity rather than on aggregate consumption. I model intermediaries as sophisticated investors subject to an equity capital constraint as in He and Krishnamurthy (2012a) and Brunnermeier and Sannikov (2012). When intermediary equity is high, prices are high and risk premia are low. However, when intermediary equity is low and the equity capital constraint is more binding, risk premia are high as the risk-bearing capacity of intermediaries declines. This generates a "financial crisis." I calibrate the model to match unconditional moments and study the resulting asset pricing dynamics. The model quantitatively matches the spikes in risk premia associated with financial crises as well as the average decline of stock prices in financial crises. I also show that the model generates recessions with and without financial crises, and that, as in the data, risk premia are significantly higher in the latter. The calibrated model matches unconditional risk premia and volatility and generates time-varying risk premia that are tied to the health of the financial sector. Specifically, the model ties movements in risk premia to the equity capital or net worth of the financial sector, and I confirm this prediction in US data. As in the model, the equity of the financial sector divided by GDP has strong

forecasting power for both stock and corporate bond returns, predicting around 17% of the variation in annual returns.

The main takeaway of this paper is that risk premia spike dramatically in financial crises, but *not* in other episodes which feature similar or larger movements in consumption. Relatedly, risk premia are relatively constant across wars and recessions, but movements in consumption are much larger during wars. This is the first paper to systematically document the behavior of risk premia during financial crises. The paper provides an explicit link between aggregate risk premia and the financial sector that is related to similar findings by Adrian *et al.* (2011) and Adrian *et al.* (2012) who find intermediary balance sheets help explain the time-series and cross-section of asset returns. However, this paper goes farther in showing why a “financial crisis” state variable is *necessary* to explain risk premia because of the differential response of risk premia in crises and recessions. It is also related to findings by Greenwald *et al.* (2014) who show that changes in risk premia or risk aversion explain a large fraction of asset price fluctuations, but these changes are uncorrelated to consumption and fundamentals. This paper shows that financial crises are times when these shocks to risk premia are highest, meaning they are uniquely important to understanding risk premia. Thinking of recessions as a control group with similarly bad macroeconomic shocks, this paper shows the additional large increase in risk premia during financial crises despite no additional increase in typical measures of macroeconomic risk. The results show that financial crises are important for understanding variations in risk premia and strongly support models where the health of the financial sector influences risk premia.

2 Data and Empirical Results

2.1 Data Description

The main data spans from 1870-2012 across 14 countries and consists of the following: real per capita consumption and GDP data, dividend yields, real dividend growth, real stock returns, and credit spreads. The countries included in the main sample are the United States, Canada, Australia, Denmark, France, Germany, Italy, Japan, the Netherlands, Norway, Spain, Sweden, Switzerland, and the United Kingdom. Consumption and GDP data are from Barro and Ursua. Dividend yield, stock price, and return data are from Global Financial Data (and these are used to construct the dividend growth series) and are converted to real US dollars. I use 1 year real interest rates in the US from Robert Shiller when constructing excess returns. Credit spreads are from Investor’s Monthly Manual which published bond

prices from 1860-1930. I also add the Moody's BaaAaa default spread from 1930 onwards for US data. Therefore, for countries outside the US, credit spreads are only available from 1860-1930 and this is my most limited data series. The other series are typically continuous (dividend yield data and return data are only occasionally missing). The data are more fully described in the online appendix.

Crisis and event dates come from several sources. Business cycle dates are from Jorda, Schularick and Taylor (Jorda *et al.* (2010)) Table 1 who document business cycle peaks for these 14 countries and document whether each was associated with a financial crisis or not. I will refer to the latter as “non-financial recessions” and I will refer to “recessions” as containing *both* financial and non-financial recessions. Overall, there are 220 non-financial recessions and 70 financial crisis dates in their sample, but these numbers reduce to 180 and 45 when I consider events that have non-missing data (excluding credit spreads which, as mentioned, are missing more frequently). These are my primary dates because they allow me to compare the behavior of asset prices across recessions which contain a financial crisis and those which do not, for a balanced panel of countries. This allows me to estimate the *differential* response of each variable in a recession versus a financial crisis. One can loosely think of the financial crisis group as the treatment group and normal recessions as the control. I will show that both are hit with similar shocks to consumption and dividends (fundamentals) but have different responses in asset prices.

I also use dates from Reinhart and Rogoff (2009) (RR) which only contain dates for financial crises but not recessions in general and the appendix shows robustness to using these dates. The main difference with the RR dates are that they date the crisis when a major bank run or bank failure occurs, rather than using the preceding business cycle peak as in Schularik and Taylor. However, they are also more extensive and cover far more countries with data on 70 countries starting as early as 1800. The main disadvantage of the RR dates is that they do not allow me to directly compare asset price responses in financial crises to those in a typical recession. I analyze the results using the RR dates in the appendix. For both dating conventions, the occurrence of a financial crisis is due to a major bank run or bank failure – therefore financial crisis and banking panic are used synonymously. The crises are not dated ex-post by aggregate stock market declines.

I also run a standard VAR of returns and dividend yields to decompose unexpected returns into discount rate news and cash flow news (see Campbell (1991)). When doing so, I demean the dividend yield and return series within each country and run a single pooled VAR. One major caveat of this approach is that I only have dividend yields historically and no other predictor variables, hence I likely assign too little of unexpected return variation

to discount rate news. This allows me to distinguish return shocks based on changes in dividend yields or dividend growth.

2.2 Empirical Results

I plot the raw data in Figure 1. I plot the five year change in dividend yields and credit spreads one year after the start of the event since this is when risk premia typically peak for both financial crises and recessions. For wars, I compute the five year change at the beginning of the event rather than one year after because this is when risk premia are typically highest, but the results are robust to using the highest possible dividend yield in a five year window around each event. I compute consumption volatility over a forward looking 10 year window beginning at the start of the event to try to pick up persistent increases in consumption volatility. Finally, for drops in consumption I compute the peak to trough change in real consumption over each event. Using peak to trough consumption is natural because consumption often responds with a lag to asset price changes in these events and this captures the overall severity of the event in terms of consumption loss but avoids imposing perfect timing between consumption and asset price changes. In every event, I use the corresponding series for the country experiencing the event (i.e. if Spain experiences a financial crisis or a civil war, I look at changes in asset prices, consumption, etc in Spain at that time). One can easily see the main patterns in the data: financial crises are associated with huge spikes in measures of risk premia, whereas wars and recessions are not. However, drops in consumption and consumption volatility are about the same across financial crises and non-financial recessions. Therefore, we can almost immediately see that consumption based models will struggle to fit the behavior of risk premia.

2.2.1 Main Results: Comparing Financial Crises and Recessions

Table 1 runs regressions of each of these variables on dummies that indicate whether a recession, financial crisis, or war related disaster occurred. The recession dummy is equal to one if a recession occurs for any reason (i.e. it includes both financial *and* non-financial as well as wars). The financial crisis dummy therefore picks up the differential effect of financial crises *relative* to recessions. I include several lags of each dummy variable to account for the fact that the dating is about the business cycle peak and asset prices typically respond after the business cycle peak (Gorton (1988)). Including more lags does not change the results. I also include country fixed effects in every regression to account for differences in returns, growth rates, etc, across countries and I include the lag of each variable as well. For the

dividend yield I also include a linear time trend to account for a secular decline in dividend yields over this time period. The results are not sensitive to including this trend. Figure 2 takes the coefficients from these regressions and plots the implied impulse responses for each variable of interest to a financial crisis and a recession (i.e., I separately set each dummy equal to 1 and simulate the system forward).

We first see a large increase in dividend yields immediately after financial crises relative to typical recessions. Dividend yields increase 4% and 20% after years 1 and 2, respectively. The 20% increase is highly statistically significant. Again it is worth emphasizing this is over and above typical recessions. In typical recessions, dividend yields increase by 9% on impact and then remain relatively stable. Dividend yields in crises are therefore significantly higher than in typical recessions, and massively higher than “normal” times. The cumulative increase in dividend yield in a financial crisis is around 30%. For reference, a one standard deviation change in dividend yield is around 27%. Credit spreads show similar patterns. In financial crises, credit spreads increase by around 5%. This is larger than a 1 standard deviation change (the standard deviation in credit spreads is about 4%). There is no meaningful change in credit spreads in recessions.

But are these changes in dividend yields actually risk premia or just changes in expected dividend growth? We know it must be one of the two based on the work of Campbell and Shiller (see e.g., Campbell (1991)). To answer this, I study the behavior of returns and dividend growth. First, I show that future returns increase dramatically while future dividends do not fall dramatically. The coefficient on the dummy for a crisis at time t gives the contemporaneous effect on returns and dividend growth relative to recessions. Contemporaneously, returns fall by 23% in a financial crisis. After three years, returns rebound by gaining 24% above their mean. The drop in returns is presumably a “shock” due to the fact that we are conditioning on the business cycle peak. The rebound, however, is forecastable by an event several years in the past. This shows that financial crises are associated with large price declines that are subsequently reversed, meaning the crisis is largely about a change in discount rates not in cash flows. This intuition is supported by Figure 2 which shows large differences in discount rate news between the episodes, but not large differences in cash flow news. Further, dividend growth is actually higher in a financial crisis initially compared to a typical recession, which also supports that the change in dividend yield is due to discount rate news and not changes in expected dividend growth rates. In untabulated results, I find no effect on changes in longer term 10 year averages of dividend growth rates. Later, I further confirm this intuition by running standard predictive regressions of returns and dividend growth on dividend yields both unconditionally and

conditional on a financial crisis. I find that the standard relationship that dividend yields forecast returns and not dividend growth also holds during financial crises.

Turning next to macro variables, we see the drop in consumption growth in recessions is around 1.1%, 3.7%, 2.0%, and 1.4% in years 0, 1, 2, and 3 after a recession, showing persistent declines in consumption. It is worth remembering that these are relative to each country's long run average as there are country fixed effects. For financial crises, there is no drop in consumption relative to recessions on impact, but an extra 1.5% and 1.1% drop 1 and 2 years out. This is the sense in which financial crises are deeper and longer than normal recessions. However, note that these differences are not statistically significant, so statistically we can not reject financial crises and recessions being relatively equal in terms of consumption. Additional lags also do not show up as significant. At first glance, it seems obvious that this relatively small extra drop in consumption could generate the large spike in risk premia. The cumulative loss in these recessions is over 10% and the extra cumulative loss in financial crises is estimated around 2%. Not only is this difference not statistically significant, it is not economically large.

The next column of Table 1 looks at the log surplus consumption ratio of Campbell and Cochrane (1999) (CC) using their long run calibration and I also plot the impulse response in Figure 2. Surplus consumption is essentially consumption relative to a slow moving average of past consumption. The patterns look fairly similar to consumption with one key difference. The habit function is non-linear so that a bad shock in bad times decreases the log surplus consumption ratio more than in good times – as you get closer to habit you become more sensitive to drops in consumption. Does this mean the small extra drop in consumption in financial crises can generate the spike in risk premia because it comes when consumption has already fallen? No, for two reasons. First, I look at the *log* surplus consumption ratio, which has a roughly linear relationship with risk premia in their calibrated model. The log surplus consumption ratio declines by only around 10% extra in financial crises relative to recessions. In the CC calibration, that corresponds to an increase in risk premia of up to 6%, which can not match the magnitude of the increase that we observe (around 20%). But there is more convincing evidence. The log surplus consumption ratio in the non-financial recession drops over 18%. The model would imply a spike in risk premia in recessions that would be much *larger* than what we observe in the data. Therefore, the model is not able to *jointly* match the lack of variation in risk premia in recessions and the large variation in risk premia in financial crises. In this case, the lack of variation in risk premia during recessions is nearly as important as the strong increase in risk premia during financial crises.

I next look at the 10 year forward looking variance of consumption growth. This is the

key state variable in Bansal and Yaron (2004) for movements in risk premia. Increases in the volatility of consumption growth increase the volatility of the stochastic discount factor. The recession dummies indicate that indeed consumption becomes more volatile following a recession. The variance increases by 0.2% (equivalently, the volatility increases by about 5%). Compared to recessions, financial crises actually see a *lower* consumption variance with an initial decrease of -0.2% relative to recessions. Therefore, the difference in risk premia clearly can not be accounted for by a change in consumption volatility. This is not a matter of calibration because the point estimate has the wrong sign. The long run risks model would therefore predict a spike in risk premia during recessions, but no additional increase in risk premia during financial crises (and in fact, potentially a decrease).

To summarize, purely consumption based models are faced with the following challenge in matching these episodes: first, relying on changes in consumption volatility can not work because volatility does not meaningfully change across episodes. The only potential way to match the data is using the extra decline in consumption in financial crises of 1.8%. But this has two complications. First, if drops in consumption cause changes in risk premia, one has to explain why there is no substantial increase in risk premia during recessions despite the 8% drop in consumption, and why there is no substantial increase in risk premia during wars when consumption drops can reach 50%.

2.2.2 Wars

So far I have shown that financial crises and recessions are fairly similar among consumption variables but vary greatly in terms of risk premia. I next compare financial crises and recessions to war related disasters. These dates are from Barro and are post 1900. The results are given in Table 1.

There are advantages and disadvantages to looking at the war related disasters. The advantages are that they generate huge drops in consumption and huge increases in the volatility of consumption. This provides good testing grounds for models because it generates large changes in consumption moments. In other words, the increases in the consumption state variables are stark in these episodes. There are also downsides. First, there are relatively fewer of these episodes overall. But the biggest concern is market shut downs during some of the major world wars that prevent dividend yields from being measured. Therefore, data availability is a problem. The markets that shut down are also those that were typically hit the hardest. I will therefore only use consumption data on the countries for which markets did not shut down so my data is balanced. Another approach is to measure

dividend yields in a flexible window both before and after markets shut down. According to the rare disasters story, the probability of a disaster should affect the dividend yield, and the probability is higher both before and after the realization of a disaster. Even looking in a window around the event, I find no increase in dividend yields.

Looking at Table 1, there are virtually no increases in risk premia around wars. In fact changes in dividend yields are negative. Returns are also large and negative, and dividend growth is negative. Returns are more negative than dividend growth, meaning prices decline as well as dividends, but dividends fall by more than prices, making the dividend yield decline overall. The evidence from returns and dividend yields suggests a lack of increase in risk premia.

Consumption, however, falls dramatically by 1%, 17%, and 10% in years 1, 2, and 3, respectively. The surplus consumption ratio naturally collapses while the variance of consumption increases dramatically. Notice the drops in magnitude in consumption, habits, and consumption variance are all drastically larger than in recessions by a factor of 5 or more. Again, the data strongly show that despite massive changes in consumption, risk premia as measured by dividend yields do not change. As mentioned, there are of course data issues with wars, but one has to consider the economic magnitude here. Cumulative consumption losses of near 30% and an increase in consumption variance of several standard deviations produces no measurable increase in risk premia. Dividend yields must be related to expected returns or dividend growth. The only way risk premia could have been high in these wars is if expected dividend growth was equally high, so that price dividend ratios remained constant. Yet, if anything we see (and would expect) negative expected dividend growth rates.

There is one possible defense of the rare disasters model: that these are just the realizations of the disaster, but maybe the probability of disaster did not increase and therefore the risk premium not increasing is natural. However, this explanation does not ultimately work. First, war related disasters are definitely forecastable. There is a “build up” period to wars in which events occur that substantially increase the probability of a disaster. In the appendix, I show that events which almost surely increased the probability of a war related disaster (like the Cuban Missile Crisis or Hitler invading Poland), do not feature large increases in risk premia. Further, when a war actually starts consumption has not yet fallen, but clearly the probability that it will fall has gone up. In the disaster calibrations, even very small increases in the probability of the disaster occurring will have very large implications for risk premia. This is the entire foundation of these models. Even if the probability does not increase before the disaster, it should certainly increase during and after the disaster as these events tend to be clustered. Further, an agent who updates the probability of a

disaster in a Bayesian way will typically expect a higher probability of disaster after one has occurred. But risk premia are not elevated during those times either. Finally, one might still argue that maybe the disasters are not forecastable, so that the probability is a constant. But that would imply constant dividend yields, meaning these models won't be able to match time-variation in risk premia. In any case, the standard consumption disaster model is not likely to fit the data.

One may also worry that the data are poorly measured around wars. For some countries, markets completely shut down during the major world wars. First, there are enough war related episodes without this issue that clearly show large declines in consumption and no change in dividend yield. While there are definitely data issues, the magnitudes of the discrepancies appear too large to be attributed to bad data. However, as mentioned, risk premia in the rare disasters model should increase before the war begins and market shutdown as risk premia depend on the probability of the disaster not its occurrence. These probabilities would increase leading up to wars as the wars became more and more likely. I see no change in risk premia in the lead up to wars or just after wars. Therefore, wars do not seem to be associated with large spikes in risk premia of the magnitude implied by standard disasters models.

The war evidence is consistent with what we found in recessions: large drops in consumption and increases in consumption variance do not seem to be associated with risk premia. Even with potential objections to the data and small sample, the message seems to be clear given the enormous magnitudes.

2.3 Robustness and additional tests

The appendix contains many robustness checks, which I will describe briefly here. First, I re-do all of these results on different dimensions of the data. The general pattern in results qualitatively hold when looking only at post war data (i.e., after 1950), when using the RR dates, and when looking at US data. The appendix confirms that movements in dividend yields correspond to changes in risk premia both unconditionally and during financial crises by running standard predictive regressions of returns on dividend yields with dummies for financial crises and recessions. The unconditional relationship that dividend yields predict returns is unaffected during these times. The appendix also discusses data sources in more detail.

3 Models

3.1 Unifying Framework

This section briefly reviews the state variables that drive risk premia in leading asset pricing models. All asset pricing models specify a stochastic discount factor (SDF) M . The main pricing equation is

$$E_t [R_{t+1}] - R_f = -R_f cov_t (M_{t+1}, R_{t+1}) \quad (1)$$

The SDF, M , is a function of a state variable S . Therefore, the co-variance and risk premia will also depend on S . The volatility of the SDF or “price of risk” is typically a function of state variables S and therefore we end up with an equation of the form¹

$$E_t [R_{t+1}] - R_f = f(S_t) \quad (2)$$

where f is some generic function. Each model proposes a different state variable S that determines risk premia. I will review how each model specifies S and discuss the calibrations of each model. We have already seen variation in the left hand side (risk premia) in the data between recessions and financial crises. The key question is whether that variation can be reasonably explained by variation in the state variables.

3.2 Consumption Based Models

3.2.1 Habits

Habit models specify utility as $U(C) = (C - X)^{1-\gamma}$. I will focus on the external habit model of Campbell and Cochrane (1999) where X depends on past consumption. Here the state variable is the surplus consumption ratio $H_t = (C_t - X_t)$.

$$E_t [R_{t+1}] - r_f \approx \sigma_t (M_{t+1}) \sigma_t (R_{t+1}) \quad (3)$$

$$E_t [R_{t+1}] - r_f \approx \sigma_t \left(\left(\frac{H_{t+1} C_{t+1}}{H_t C_t} \right)^{-\gamma} \right) \sigma_t (R_{t+1}) \quad (4)$$

Consumption growth is i.i.d., and habit is based on past consumption, so the key state variable is H . Therefore high risk premia should be associated with large drops in consumption (or consumption falling relative to habit).

¹The link is clearest in continuous time where the main pricing equation is

$$E_t [dR_{t+1} - r dt] = -\lambda(s_t) \sigma_{R,t}$$

where $\lambda(s_t)$ is the volatility of the pricing kernel and $\sigma_{R,t}$ represents the factor loadings.

To get a sense of magnitudes, I calibrate the habit model based on Campbell and Cochrane (1999). The calibration is difficult because different countries have different expected growth rates and different volatilities of consumption. Campbell and Cochrane (1999) have two calibrations – one based on postwar data and one over a longer sample. In the calibration geared to match the recent sample, habit plays a larger role in the results because consumption volatility is much lower (1.2% vs 3%), and because the Sharpe ratio is larger in the recent sample. Therefore, habit must be “cranked up” dramatically to account for higher volatility of the discount factor (to match the higher Sharpe ratio) with lower consumption volatility. Therefore, in this calibration, the slope coefficient of expected return on log habit is about -3. In the long sample parameterization, it is only about -0.6. This would be further reduced if calibrated to international data because international consumption volatility tends to be even higher. However, I will use the long sample calibration as it best applies to the data I have and is taken directly from their original study. The estimated coefficient and estimated increase in habit in financial crises would imply an increase in expected return of about 3.5% in financial crises relative to recessions. Moreover, the estimated increase in risk premia during recessions imply that expected returns would rise by around 8.5%, whereas they would increase by around 45% during wars. All in all, recessions and financial crises should be about the same in terms of increases in risk premia according to the habit model where risk premia in wars should be dramatically higher than both, therefore the model can not match the data on this dimension.

3.2.2 Long Run Risks

Long run risks models (Bansal and Yaron (2004)) feature Epstein-Zin-Weil utility and slow persistent movements in consumption and consumption volatility. Future consumption enters the SDF and hence

$$E_t [R_{t+1}] - R_f = f(\sigma_{C,t}) \tag{5}$$

The model is lognormal, so more specifically we have

$$E_t [r_{t+1}] - \frac{1}{2}\sigma^2(r_{t+1}) - r_f = \alpha + \lambda\sigma_{C,t}^2 \tag{6}$$

for constants α and λ . Therefore high risk premia should correspond to high consumption volatility. Return volatility also depends on consumption volatility, therefore expected excess returns are only a function of consumption volatility.

Empirically, I look at 10 year rolling estimates of consumption variances. The estimated consumption variance at time t uses the forward 10 years of annual data. This is, of course,

only a proxy for consumption volatility. However, to match the large spike in the observed risk premia, it must be the case that consumption volatility would have to increase dramatically. In the data, consumption volatility is nearly identical across recessions and financial crises. In fact, the point estimate of the variance of consumption on recession and financial crisis indicators implies that consumption is if anything slightly less volatile in financial crises, though statistically the coefficient is not significant. Regardless of magnitudes, this model is not able to account for the variation in risk premia across episodes. Here the calibration is not correctly suited to the data because the original BY study focuses on US data after WWII when consumption volatility is very small. Therefore, rather than using that calibration (which would produce massive increases in risk premia in recessions) I will simply point out that the negative coefficient means the model is not able to match the facts under *any* parameters. Similarly, the model can not explain the lack of an increase in risk premia during wars relative to the other two episodes, despite massive increases in consumption volatility.

3.2.3 Rare Disasters

The rare disasters literature (Barro (2006), Rietz (1988), and Gabaix (2012)) argues that asset prices and risk premia can be explained by rare disasters which are defined as any large decline in consumption and/or GDP. Empirically, most of these disasters are major wars or financial crises. In these models the equity premium is only a function of the probability of the rare disaster, and a 1-2% probability of disaster can match the equity premium with low risk aversion. Gabaix (2012) shows that the expected no-disaster equity premium is approximately given by

$$E_t [R_{t+1}] - r_{f,t} = p_t E_t [B_{t+1}^{-\gamma} (1 - R_{t+1}^{dis})] \quad (7)$$

where p_t is the probability of disaster, B_{t+1} is the size or severity of the disaster (i.e. a 30% loss in output means $B_{t+1} = 0.7$), R_{t+1}^{dis} is the gross return conditional on disaster, and γ is risk aversion. Therefore the equity premium moves one-for-one with an increase in the probability of disaster, and increases exponentially with the size and potential severity of the disaster where the sensitivity depends on the risk aversion parameter γ . Typically, the rare disasters literature exogenously specifies a process for p_t to generate both high unconditional risk premia and time-varying risk premia. In calibrated disaster models, a 2% increase in the probability of disaster would double the equity premium, so even small changes in p will have large changes in risk premia (in fact, this is the point of these models).

My findings indicate that consumption disasters can not explain variation in risk premia because the most severe consumption disasters – wars – show little increases in risk premia while financial crises, which are comparably not nearly as severe, have much larger increases in risk premia. Consumption drops an average of 45% in a war related disaster, compared to about 10% in a financial crisis. Therefore, B_{t+1} is clearly highest in wars.

Of course measuring the probability of a consumption disaster is difficult and my analysis so far relies on the idea that the probability of a disaster increases right before the beginning of a disaster. I think this is a reasonable assumption for wars because usually at the start of a war, or just before a major war is lost, it seems reasonable that there is an increased likelihood of consumption falling. My results show that in years right before consumption is about to fall drastically, there is no substantial increase in risk premia, so one would have to believe that the probability of a disaster is constant over that period. However, I also look at the data from several other angles and reach the same conclusion. I find that the results are robust to using the peak dividend yield in a ten year window centered around the disaster, which only assumes that the probability of the disaster increased at some point around the event. Lastly, for US data where I have higher frequency, I check that risk premia do not change substantially in events that clearly increased the probability of a war related disaster such as the attack on Pearl Harbor, Hitler invading Poland, or the Cuban Missile Crisis. These results are in the appendix.

3.3 Intermediary Based Models

In intermediary based theories the SDF depends on the health of the financial sector. This is typically related to how constrained the financial sector is in raising debt and/or equity financing. The constraints affect the risk bearing capacity of the financial sector and therefore affect risk premia. Naturally, these theories imply risk premia will be highest in financial crises.

Generically,

$$E_t [R_{t+1}] - R_f = f(e_t) \tag{8}$$

where one can think of e_t as the health of the financial sector. Times when e are low constitute financial crises when the risk bearing capacity of the financial sector is particularly low and when their balance sheets are weakened. Therefore the central prediction of these models are that risk premia are highest in financial crises, consistent with what we see in the data. Examples of these models include He and Krishnamurthy (2012a), Brunnermeier and

Sannikov (2012), Danielsson *et al.* (2011), and Adrian and Boyarchenko (2012). The next section goes in depth into an intermediary asset pricing model then calibrates the model to the data.

4 Intermediary-based Model

The model is based on the growing literature on intermediaries, asset pricing, and macroeconomics and is most closely related to He and Krishnamurthy (2012a) and Brunnermeier and Sannikov (2012).² Relative to these papers, I use a simplified framework. My goal is to calibrate the model to aggregate asset price data and see how well it can explain the facts presented thus far. My goal is not to focus on micro foundations for where these frictions come from or to fully estimate the model. Instead, the goal is simply that a plausible calibration of these models provides a good way to interpret the data. This allows me to compare the model to the calibrated consumption based models described earlier. As a result, however, the model takes short cuts in the interest of space.

4.1 Model of Financial Crises and Risk Premia

There are two agents in the economy: households who consume and intermediaries who make investment decisions. The key assumptions are that intermediaries are better at making investment decisions than households, but that households can only contribute a limited amount of equity to intermediaries. The first assumption makes it more efficient for households to give funds to intermediaries while the second assumption implies asset prices will depend on the amount of capital households can contribute due to frictions. I refer to intermediary equity capital as the amount of equity households provide to intermediaries at any given time.

Time is continuous and there is a tree which bears fruit, or output, Y that evolves according to:

$$\frac{dY_t}{Y_t} = \mu dt + \sigma dB_t \tag{9}$$

where B_t is a Brownian motion. The growth rate of output, μ , and output volatility, σ , are constant and output shocks are the only source of uncertainty in the model.

²See also Brunnermeier and Sannikov (2012), He and Krishnamurthy (2012a), Rampini and Viswanathan (2012), and Danielsson *et al.* (2011). These also build on earlier work on financial frictions by Holmstrom and Tirole (1997) and Bernanke *et al.* (1996).

Let P denote the (endogenous) price of the tree which is a claim to the stream of dividends $\{Y\}$. The market return is defined as:

$$dR_t = \frac{dP_t + Y_t}{P_t} \quad (10)$$

Given the process for output and definition of returns, I next describe in detail the decisions of households and intermediaries.

4.1.1 Households

Households are risk neutral and discount the future at rate ρ . Households make decisions to maximize

$$E \left[\int_0^{\infty} e^{-\rho t} C_t dt \right] \quad (11)$$

Households make decisions over consumption and investment. They can invest in a risk free asset which earns r_t or they can invest in intermediaries and earn $dR_{E,t}$. Households are bad at managing the tree themselves and if they hold the tree directly it depreciates at constant rate δ forever and they are not able to sell back the tree to intermediaries. Because of this, households would be willing to pay at most $\underline{P} = \frac{Y}{\rho + \delta}$ which follows from the Gordon growth formula. Provided $P \geq \underline{P}$, households will not hold any of the risky asset. This sets a lower bound on the price dividend ratio. Lastly, if households invest in the intermediary, they can invest at most E units of capital where E is taken as given by households and will be discussed in the next section. One can loosely think of this as a moral hazard constraint that limits the amount of equity households can contribute to the intermediary (see He and Krishnamurthy (2012c), He and Krishnamurthy (2012b)).

Households' wealth evolves according to:

$$\frac{dW_t}{W_t} = -\frac{C_t}{W_t} dt + \frac{E_t}{W_t} dR_E + \frac{W_t - E_t}{W_t} r_t dt \quad (12)$$

where I am assuming households give the maximal amount to intermediaries (equivalently the expected return on the intermediaries portfolio is at least as high as the risk free rate: $E[dR_E] \geq r$). I will show that this condition always holds.

4.1.2 Intermediaries

As in He and Krishnamurthy (2012b), intermediaries can only raise a certain amount of equity capital from households. Once intermediaries raise capital from households, they

make a portfolio choice decision over the risky asset and the risk free asset. Thus, their liabilities will be made up of equity from households and any risk free borrowing while their asset side will typically be made up of risky assets. After returns are earned, a fraction ψ of intermediaries die in each period.

There is a continuum of intermediaries who can each raise equity ϵ_t from households. Intermediaries have log preferences over their consumption which is a constant fraction of the equity raised from households $C_t^I = \lambda\epsilon_t$. We should think of λ as an infinitesimal fee intermediaries charge households as a fraction of equity they manage. That is, the fee is small enough that it does not affect the return intermediaries offer households and so that $\lambda\epsilon_t$ does not affect aggregate consumption.

Given their preferences and the stochastic death rate ψ , intermediaries seek to maximize

$$E \left[\int_0^{\infty} e^{-\psi t} \ln(\lambda\epsilon_t) dt \right] \quad (13)$$

The amount of equity capital intermediaries can raise, ϵ , evolves as

$$\frac{d\epsilon_t}{\epsilon_t} = \alpha_t (dR - r dt) + r dt \quad (14)$$

where α_t is the portfolio choice of the intermediary. Intuitively, this says that intermediaries can raise more capital when past returns are high. It captures the idea that households will be less willing or able to invest in the intermediary when past returns are poor. This can be due to informational reasons, or to a moral hazard argument (see He and Krishnamurthy (2012c) for a model which formalizes this). In this paper, I take this constraint as given.

Given the log objective function, the intermediaries' problem is reduced to a simple mean-variance portfolio choice problem

$$\max_{\alpha_t} \alpha_t (\mu_{R,t} - r_t) - \frac{1}{2} (\alpha_t \sigma_R)^2 \quad (15)$$

That is, intermediaries behave "as if" they have preferences over the equity given to them by households directly and optimize the return of this equity in a mean-variance fashion. The first order condition is:

$$\mu_{R,t} - r_t = \alpha_t \sigma_R^2 \quad (16)$$

Define E as the *aggregate* equity raised by intermediaries. E evolves as

$$\frac{dE_t}{E_t} = \alpha_t (dR - r dt) + (r - \psi) dt + d\gamma_t \quad (17)$$

Where α_t is given by the above equation and ψ reflects the death rate. The term $d\gamma_t \geq 0$ reflects entry, which I describe when describing the boundary conditions. Entry happens very rarely when the price falls to the households private value.

The return to households holding a unit of equity in the intermediary is thus:

$$dR_E = \alpha_t (dR - rdt) + rdt \quad (18)$$

It is clear that $\mu_{R_E} \geq r$; hence the households will give maximum possible equity to the intermediary due to risk neutrality of households. Without this assumption, there are regions where the capital constraint doesn't bind and households contribute less capital than the constraint allows (He and Krishnamurthy (2012b)). Since the tightness of the constraint essentially determines risk premia, the assumption here provides a direct link between intermediary equity and risk premia. Moreover, risk neutrality results in a constant risk free rate, whereas the interest rate in He and Krishnamurthy (2012b) can be highly negative and volatile in crises.

4.1.3 Equilibrium and Solution

An equilibrium consists of prices and allocations such that agents' decisions are chosen optimally given prices and the market clears. Given risk neutrality of households, we must have $r = \rho$, which implies the risk free rate is constant. At this interest rate households are indifferent between current and future consumption. As long as $E > 0$, so that intermediaries are able to raise capital, the risky asset is held entirely by the intermediary sector, meaning $\alpha_t E_t = P_t$. For this to hold, it must be that $P \geq \frac{Y}{r+\delta}$, where $\frac{Y}{r+\delta}$ is the households valuation of the risky asset if held directly. I discuss this more fully below.

We must also have that households consume all output $C_t = Y_t$ and own all wealth $P_t = W_t$. Recall that this requires that the fraction λ of households' equity that intermediaries consume is arbitrarily small. Intuitively, intermediaries' consumption makes up a trivial amount of overall consumption, therefore I make this assumption for greater ease in solving the model. One could instead define "total" consumption as $Y_t^* = Y_t + \lambda E_t$. In this case, $C_t = Y_t$ holds and household wealth is the present value of the tree (Y_t) hence $P_t = W_t$.

Equilibrium

I conjecture a price function $P_t = p(e_t)Y_t$, where I define $e_t = \frac{E_t}{Y_t}$, the ratio of intermediary equity to total output, as the main state variable. Given this conjecture, we can calculate the market return using Ito's Lemma as

$$dR_t = \frac{dP_t + Y_t}{P_t} = \frac{dY_t}{Y_t} + \frac{p'}{p} de_t + \sigma \sigma_e \frac{p'}{p} dt + \frac{1}{2} \sigma_e^2 \frac{p''}{p} dt + \frac{1}{p} dt \quad (19)$$

Given our assumption on $p(e_t) \geq \frac{1}{r+\delta}$, the intermediary holds the entire risky asset. Hence by market clearing

$$\alpha_t = \frac{P_t}{E_t} = \frac{p(e_t)}{e_t} \quad (20)$$

I define aggregate “risk aversion” as:

$$\Gamma(e_t) = \frac{p(e_t)}{e_t} \quad (21)$$

Then using market clearing and intermediary optimality

$$\mu_{R,t} - r_t = \Gamma(e_t) \sigma_{R,t}^2 \quad (22)$$

which justifies $\Gamma(e_t)$ as the risk aversion of a fictitious representative agent with mean-variance preferences. We can see two main features of risk premia, $\mu_{R,t} - r_t$. First, risk premia depend on intermediary capital so that low capital implies high risk premia and vice versa. Second, these effects are non-linear due to the $\frac{1}{e_t}$ term in $\Gamma(e_t)$ – when intermediary capital is high, changes in capital will have small effects on risk premia, but when it is low risk premia will spike and be particularly sensitive to further changes in intermediary capital. This means volatility will be high as small changes in intermediary capital can lead to large changes in prices.

It is further useful to define the Stochastic Discount Factor (or pricing kernel), which prices all assets:

$$\frac{d\Lambda_t}{\Lambda_t} = -r dt - \lambda_t dB_t = -\rho dt - \Gamma(e_t) \sigma_{R,t} dB_t \quad (23)$$

This points to the failure of the CCAPM in the model, and highlights dependence of the price of risk on intermediary equity. For the CCAPM to hold the SDF would be based off of aggregate consumption growth, thus the diffusion term on the SDF (λ_t) would need to be $\lambda_t = \alpha \sigma_{C,t} = \alpha \sigma$, which would imply a constant price of risk. Instead, assets are priced off the *intermediary's* marginal rate of substitution, hence $\lambda_t = \alpha \sigma_{E,t}$ and the SDF is a function of intermediary capital. This will make financial crises, defined as times when e_t is low, the times when risk premia are highest.

Solution

We can combine the return equation (19) with optimality (22) to derive the ODE:

$$\mu + \frac{p'}{p} \mu_e + \sigma \sigma_e \frac{p'}{p} + \frac{1}{2} \sigma_e^2 \frac{p''}{p} + \frac{1}{p} - r = \frac{p}{e_t} \sigma_R^2 \quad (24)$$

It remains to solve for the expressions $\mu_e, \sigma_e, \sigma_R$ in the above equation.

We know using equation (19) for dR

$$\sigma_R = \sigma + \frac{p'}{p}\sigma_e \quad (25)$$

Finally, we can apply Ito's Lemma to get the dynamics for $e_t = \frac{E_t}{Y_t}$

$$\begin{aligned} de_t &\equiv \mu_{e,t}dt + \sigma_{e,t}dB_t \\ &= e_t(\mu_E - \mu + \sigma^2 - \sigma\sigma_E)dt + e_t(\sigma_E - \sigma)dB_t \end{aligned}$$

This gives us μ_e and σ_e in terms of μ_E and σ_E . Looking at the dynamics for E

$$\sigma_E = \frac{p}{e}\sigma_R \quad (26)$$

$$\mu_E = \frac{p}{e}(\mu_R - r) - \psi + r = \sigma_E^2 - \psi + r \quad (27)$$

Thus, we can combine these (using σ_R from above) to solve for all three volatilities

$$\sigma_R = \sigma \frac{(p - p'e)}{p(1 - p')} \quad (28)$$

$$\sigma_e = \sigma \frac{(p - e)}{(1 - p')} \quad (29)$$

$$\sigma_E = \sigma \frac{(p - p'e)}{e(1 - p')} \quad (30)$$

We can therefore substitute in $\mu_e, \sigma_e, \sigma_R$ to our ODE in equation (24).

To solve the ODE, we need to specify the boundary conditions. As e_t becomes large, we know prices are no longer dependent on intermediary capital hence $p'(\infty) = 0$. The lower boundary condition is as follows. I assume that new intermediaries enter when the price reaches $\frac{1}{r+\delta}$, which is the households willingness to pay for the asset. I assume that at this price there is an intervention in the economy to prevent the households from buying the risky assets and thus economic growth falling permanently. This can be thought of as new capital coming in because the low price and high Sharpe ratio is attractive (He and Krishnamurthy (2012b)), or as the government injecting new capital into the economy to prevent growth from falling permanently. This means that \underline{e} is a reflecting barrier and $p'(\underline{e}) = 0$ since the price will not change on entry. This condition, together with the condition, $p(\underline{e}) = \frac{1}{r+\delta}$, determines the endogenous entry point \underline{e} . It turns out that the economy rarely ever hits this lower bound. I discuss these conditions in more detail in the appendix and give details on the numerical solution.

5 Calibration and Comparison to Data

5.1 Calibration and Basic Moments

Table 2 contains the calibrated parameters. I assume standard parameters where available. I set the volatility of aggregate output or consumption growth, σ , to 5%, which is consistent with historical data from Barro *et al.* (2011). In the model there is no difference between consumption and output (GDP) because it is an endowment economy. I will therefore try to calibrate to consumption data, but should note that using GDP instead is not meaningfully different. The parameter μ is chosen at 1.8% to match average real economic growth. The intermediary death rate ψ is set to 8%. This parameter ensures well behaved dynamics and a stationary distribution. Finally, the parameter ρ is set at 3% to roughly match the real interest rate and also to match the average price dividend ratio in the data. The model has only five parameters and these are either pinned down by the data or chosen to match unconditional moments.

The depreciation rate δ governs entry on behalf of households and determines how low the price dividend ratio will fall. I choose a depreciation rate of 13% so that the implied lower bound on the price dividend ratio is about 6. I calibrate this parameter to roughly match the lower bound on the price dividend ratio in the data. For the U.S. the lowest value in the past 100 years is around 10, while when using international data this value fell to as low as 4. If I simulate the model, the average minimum price dividend ratio observed in 100 years of annual data is just over 9, which is close to the lowest value observed in U.S. data.

I plot the model solution in Figure 3. As we can see risk premia and volatility are increasing as intermediary equity falls. These effects are non-linear and the sensitivity of these variables to intermediary equity is significantly higher in bad times. For reference, I plot a dashed vertical line which represents the 7th percentile of the state variable, which will represent the cutoff for a crisis in the economy as I describe later. I also plot the stationary distribution of the state variable (I exclude the high end of the distribution on the plot because all variables are essentially flat there). As we can see in normal times when e is very high, negative shocks will not have large effects on risk premia, but in crisis times (to the left of the dashed line), small shocks will have very large effects.

Basic Moments

Table 3 compares moments in the model to the data. I simulate the model monthly then aggregate all simulated data to the annual frequency to match my sample which contains annual data. The model calibration matches average moments quite well. By design the model matches average economic growth and volatility of consumption. The model also

matches the equity premium (6.9% vs. 7.4% in the data), volatility (20% vs. 19%), and hence the market Sharpe ratio.³ The model also matches the “volatility” of volatility (10.3% in the model vs. 9.2% in the data) so volatility in the model is strongly time-varying. The model is too low on the log price-dividend ratio (2.8 in the model vs. 3.3 in the data). The model has a reasonable real risk free rate at 3%. The real risk free rate is less than 1% in post war US data, but is 2.9% over a longer sample in US data (see Campbell and Cochrane (1999)) and is around 1.5% using an international sample (Barro (2006)). The biggest challenge for the model is the persistence of the dividend yield (0.5 vs. 0.8 in the data), which also results in low volatility of dividend yields. However, the value from the model is much closer to Lettau and Van Nieuwerburgh (2008) who use trend breaks in the dividend yield and find a persistence of 0.6 and volatility of 0.2. In my setting, I could likely add low frequency persistence in the mean dividend growth rate to generate higher dividend yield persistence. This would likely not change price dynamics much due to the log utility, but this isn’t pursued in this paper. Taken as a whole, Table 3 suggests that the model does a good job of matching standard asset pricing moments.

5.2 Crises and Recessions

Defining a crisis in both the model and data is a challenge. Empirically, there is not widespread agreement on what exactly constitutes a financial crisis. In the model it is clear that a crisis should be defined by low realizations of the state variable e_t , but e_t takes on a continuum of values so deciding on the exact cutoff is potentially arbitrary.

I choose to base my cutoff to target the average probability of a financial crisis. Reinhart and Rogoff (2009) estimate the share of years spent in a banking crisis since 1945 to be 7.2% for advanced economies and 11% for emerging markets. I choose the cutoff of e_t that defines a crisis so that the economy spends 7% of its time in the crisis region to match the percentage of time spent in a crisis for advanced economies which make up my sample. My definition of a crisis should not be seen as crucial, but rather a good way to illustrate the effects of a crisis and compare to the data. The results aren’t sensitive to a slightly different cutoff.

Given this definition of a crisis, I show that the model produces declines in consumption and increases in dividend yields during a crisis that are in line with severe crises in the data.

³Note that the Sharpe ratio is computed as the unconditional average return over unconditional volatility of returns, which in the model is somewhat different than the average conditional Sharpe ratio because of the co-movement between expected returns and volatility. The average conditional Sharpe ratio is 0.42 vs an unconditional Sharpe ratio of 0.35, which reflects this co-movement in the model.

I plot this in Figure 4. First, the model matches the data extremely well in the response of dividend yields both in terms of levels and dynamics. In terms of timing, we see in the data dividend yields increasing a year after the onset of the business cycle peak associated with the financial crisis. Therefore in the model, I plot the response of all variables assuming the crisis begins the following year so that the timing of the dividend yield peak matches the data. The model also matches the total drop in consumption in the data well, though this is a sharp decline in the model vs a more persistent decline in the data. In terms of returns, the model matches the initial drop in stock returns and the reversion, though the reversion happens faster in the model than in the data. Some of this could be due to timing in the data as it is hard to measure precisely when things occur, and some could be due to the starkness of the model. Overall, Figure 4 shows that the model dynamics for risk premia and economic growth in a crisis match the data very well.

Table 4 runs regressions of dividend yields, returns, and consumption growth on indicators for financial crises and recessions as in the data. In the model, I define a recession as having occurred if annual economic growth (consumption growth) is negative. I use the timing assumption as mentioned above where both events in the model occur in year 1. We can see that the model does well in matching the data. Dividend yields spike by about 20% more in financial crises than recessions in the model and data. In recessions, dividend yields spike by around 8.6%. Consumption drops also match well across episodes. In the model consumption is about 4% lower than in recessions, which is close to the data. Moreover, recessions generate around an 8-10% decline in consumption in the data and around 8% in the model. One caveat is that the timing in the model is very stark: consumption and asset price drops happen simultaneously and immediately whereas consumption in the data is perhaps slower to respond. For returns, the model generates a large loss in crises and recessions. Notice in the model only the crisis return is substantially reversion in the following period (this is due to the high risk premia). The return two years out is 22% above the unconditional average return, so prices fully rebound. In the data this takes longer, but we do see the exact same effect with returns close to 20% four years out. The bottom panel of this Table reports the results when only considering financial crisis dummies and shows similar results. Overall, the model does a good job capturing the difference between recessions and financial crises, as well as the behavior of consumption, returns, and risk premia during financial crises.

What explains the difference between typical recessions and financial crises in the model? Risk premia are very high in financial crises because the equity of the financial sector is very low and their risk bearing capacity is reduced (equivalently, their balance sheets are

weakened). In times when e is low even small shocks have amplifying effects through the fact that volatility is very high (see Figure 3). In contrast, in normal times, when equity is relatively high, even large negative economic shocks (i.e. a recession) will not change risk premia drastically. Suppose the economy is near the largest possible state for e . A negative economic shock will reduce output Y , but will also reduce equity E . Risk premia and volatility are very low in this state, as can be seen by the model solution, so the market return will be close to dividend growth that period which is equal to the drop in Y . Loosely, $dR \approx dY/Y$ in normal times because there is no amplification through the financial sector. Thus $e = E/Y$ will change very little and the shock will cause little movement in risk premia.

One major caveat of the model is that it currently only has one shock. This is for simplicity and tractability, but admittedly limits some of the analysis. For example, it generates the close timing between consumption and asset price data. It also means that the state variable only depends on the history of economic shocks which affect both consumption and asset prices. There is no way to enter a financial crisis without having negative economic shocks, though the model does amplify those shocks differently depending on the health of the financial sector. What I have shown is that the relationship between the state variable and consumption is much more complicated in this model than other work, therefore the model can generate recessions that are deep in terms of consumption loss but show little effects in risk premia. That accords with the extreme non-linearity of the models solution seen in Figure 3. However, this is also a limit of the model. In the data, it is clear that variables uncorrelated (or at least not perfectly correlated) with consumption shocks can potentially affect intermediary balance sheets and risk bearing capacity. Therefore, a natural extension of the model would be to allow for a separate shock affecting only intermediaries which could induce even more variation between low consumption episodes and high risk premia episodes. In other words, an extension of the model could easily generate independent variation in the state variable. This richer model could help explain the behavior of risk premia during wars, which is currently outside the scope of the model because of the complication it would add. Indeed we should not take the simple model here too literally, the goal is simply to show that it does a reasonably good job at matching the data despite its simplicity. It should be clear that because risk premia are tied to intermediary balance sheets, the model still matches the data in that times when intermediary balance sheets are weak correspond to high risk premia and other measures of “bad times” do not seem to matter. Therefore the core prediction of why risk premia vary in the model is consistent with what we see in the data. The next section tries to make the link between the equity of the financial sector and risk premia even more clear in the data.

5.3 The Link Between Risk Premia and Intermediary Equity

The model says that risk premia fluctuate with the health of the financial sector. While the previous sections have established the link between financial crises and risk premia, this section *directly* shows that intermediary equity measures risk premia by showing that it strongly predicts asset returns and is “priced” in the cross-section of stock returns. This supports the main channel through which risk premia operate in the model, and formalizes the observed link between risk premia and financial crises in the previous sections. One caveat to this section is that financial sector equity data is not available historically across countries and over longer time periods, so I can not systematically analyze its behavior in financial crises.

I measure intermediary equity (e_t) as the total market value of the financial intermediary sector divided by GDP in US data. I calculate market value as price times total shares outstanding of the financial sector in CRSP. I define the financial sector as having an SIC code beginning with 6, though more refined definitions work equally well. For example, one can exclude real estate firms or only focus on commercial and investment banks. A major caveat, however, is that this measure does not include private financial intermediaries such as hedge funds or private equity. I take quarterly GDP from NIPA and create a monthly series by assuming the current monthly growth rate is equal to the previous quarters growth rate so that I do not use any future data in constructing the estimated current months GDP. Monthly data is preferred in order to forecast returns at monthly horizons because in the model there are high frequency movements in risk premia. The analysis using only quarterly data is nearly identical when predicting returns at quarterly or longer horizons. I define $e_t = \ln(\text{FinMktCap}_t/\text{GDP}_t)$.

I run predictive regressions for asset returns as:

$$R_{t+k}^e = \beta_0 + \beta_1 e_t + \beta_2 t + \varepsilon_{t+k}$$

where k is the number of months ahead and R_{t+k}^e is the excess return over the risk free rate. I include a linear time trend t to account for an increasing trend in e_t over time as the financial sector has grown. Alternatively, one can linearly detrend the series, but this technically requires using future data not known at time t . Therefore I simply account for the trend by including it in the regression.

Table 5, Panel A provides the results for forecasting the market excess return for various horizons and shows that the R^2 ranges from 2% (monthly) to 17% (annually) to 44% (5 year horizon) over the 1948-2012 time period. This is a much higher degree of predictability

found using the price dividend ratio alone and when I include the price dividend ratio to the regression it does not, generally, show up as significance or increase the forecasting power. The sign is negative, which is consistent with low intermediary capital corresponding to high risk premia. Intermediary equity also forecasts annual excess corporate bond returns with an R^2 of 17%, and the annual excess return of the financial sector with an R^2 of 20%. I repeat these exercises in the model. All signs are consistent with the model, and many of the values are comparable. One key difference, however, is that in the model predictability is relatively stronger at shorter horizons and relatively weaker at longer horizons, and coefficients decline with horizon. This implies that movements in risk premia are less persistent in the model than in the data. In sum, intermediary equity has substantial forecasting power for asset returns over many frequencies, lending support to the view that it co-moves with risk premia. Figure 5 plots e_t (linearly detrended) in the data along with the subsequent 5-year market return and shows the high correlation between the two series. The lowest realizations of e_t occur in 2008-2009, 1990, and 1982, respectively – all times when the U.S. experienced trouble in the financial sector.

Turning to Panel B, I show that e_t is “priced” in the cross-section of stock returns. In the model e_t enters the SDF along with the market return, motivating a two factor model for expected returns:

$$E[R^e] = a + \beta_{R,mkt}\lambda_{mkt} + \beta_{R,e}\lambda_e + \varepsilon_{t+k}$$

where $\beta_{R,mkt} = \frac{\text{cov}(R_{mkt}, R)}{\text{var}(R_{mkt})}$ and likewise for $\beta_{R,e}$. According to the model, we should see positive prices of risk for exposure to intermediary equity since low e_t states are ones with high marginal value of wealth. Assets that co-vary with e_t are thus risky and must offer high returns. To test this, I use 35 excess returns from Ken French’s website: 25 size and book-to-market portfolios and 10 momentum portfolios. I run standard two-pass regressions where I first estimate β s in a time-series regression and then run a cross-sectional regression of average returns on these β s. Intermediary capital indeed has a positive and statistically significant price of risk at 0.43 with a t-stat of 2.76. I use Shanken (1992) standard errors which correct for first stage estimation of β s. The two-factor intermediary model is able to explain about half of the variation in average returns in these portfolios with an adjust R-squared of 49%, though only intermediary equity carries a significant price of risk. As a benchmark, I compare this to a four factor model which includes the Fama-French and momentum factors. This four factor model explains 86% of the variation in average returns. While there is no true cross-section in the model, these findings still support the implication that intermediary equity enters the pricing kernel.

These results extend previous research linking intermediary balance sheets to risk premia (Adrian *et al.* (2012), Adrian *et al.* (2011)). First, my results use market valuations of intermediary net worth, whereas previous results use book value. My measure of risk premia is explicitly implied by a large number of models, linking my results closely to theory. Second, previous results focus on subsets of the intermediary sector (e.g., broker-dealers or shadow banks) whereas this paper uses the entire sector as a whole. I also use higher frequency data and a much longer sample. Finally, I use the same measure to show both time-series predictive power and cross-sectional pricing power, while these papers study each separately. These results are meant to compliment the previous literature on intermediaries and risk premia and to show that the main implications of the model hold in the data.

6 Conclusion

This paper argues that financial crises are essential for understanding asset price fluctuations and risk premia. I first document this fact empirically by using data on consumption, dividend yields, stock returns, and credit spreads for over 140 years and 14 countries, yielding 180 recessions, 45 financial crises, and 20 major wars. First, I document that risk premia spike dramatically in financial crises – defined specifically as a banking panic or banking crisis – but rise only slightly in recessions or wars. The large increase in risk premia during financial crises puts explaining these episodes to the forefront of the asset pricing literature, but equally interesting is the lack of variation in risk premia across the other episodes. These facts add substantially to the question of why risk premia vary over time. To my knowledge, this is the first paper to study and characterize the behavior of risk premia across financial crises, and is the first to compare financial crises to other events. Next, I examine the ability of leading asset pricing models to explain these facts. The behavior of consumption moments across financial crises, recessions, and wars is either roughly flat or has the wrong sign, meaning the variation in risk premia can not be explained by standard consumption based macro asset pricing models. I then study a simple model similar to He and Krishnamurthy (2012b) and Brunnermeier and Sannikov (2012) that generates financial crises that quantitatively match those in the data. The key feature of the model is that risk premia are tied to the health of the financial sector (specifically equity capital). Consistent with the model, I show that a measure of financial intermediary equity forecasts annual stock and corporate bond returns with a high degree of explanatory power of around 17-20% in US data. Overall, my findings strongly support the idea that risk premia are linked to the health of the financial sector and highlight the role of financial crises for understanding risk

premia.

References

- Tobias Adrian and Nina Boyarchenko. Intermediary leverage cycles and financial stability. *NY Fed working paper*, 2012.
- Tobias Adrian, Emanuel Moench, and Hyun Song Shin. Financial intermediation, asset prices, and macroeconomic dynamics. *working paper*, 2011.
- Tobias Adrian, Erkko Etula, and Tyler Muir. Financial intermediaries and the cross-section of asset returns. *The Journal of Finance*, forthcoming, 2012.
- Ravi Bansal and Amir Yaron. Risks for the long run: A potential resolution of asset pricing puzzles. *The Journal of Finance*, 59(4):1481–1509, 2004.
- Robert Barro, Emi Nakamura, Jon Steinsson, and Jose Ursua. Crises and recoveries in an empirical model of consumption disasters. *working paper*, 2011.
- Robert J. Barro. Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics*, 121(3):823–866, 2006.
- Ben Bernanke, Mark Gertler, and Simon Gilchrist. The financial accelerator and the flight to quality. *Review of Economics and Statistics*, 78:1–15, 1996.
- Markus Brunnermeier and Yuliy Sannikov. A macroeconomic model with a financial sector. *working paper*, 2012.
- John Y. Campbell and John Cochrane. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107(2):205–251, April 1999.
- John Y. Campbell. A variance decomposition for stock returns. *Economic Journal*, 101:157–179, 1991.
- Jon Danielsson, Hyun Song Shin, and Jean-Pierre Zigrand. Balance sheet capacity and endogenous risk. *working paper*, 2011.
- Xavier Gabaix. Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *The Quarterly Journal of Economics*, 127(2):645–700, 2012.

- Gary Gorton. Banking panics and business cycles. *Oxford Economic Papers*, 40(4):751 – 781, 1988.
- Daniel L. Greenwald, Martin Lettau, and Sydney C. Ludvigson. The origins of stock market fluctuations. *working paper*, 2014.
- Zhiguo He and Arvind Krishnamurthy. Intermediary asset pricing. *The American Economic Review*, forthcoming, 2012.
- Zhiguo He and Arvind Krishnamurthy. A macroeconomic framework for quantifying systemic risk. *working paper*, 2012.
- Zhiguo He and Arvind Krishnamurthy. A model of capital and crises. *The Review of Economic Studies*, forthcoming, 2012.
- Bengt Holmstrom and Jean Tirole. Financial intermediation, loanable funds, and the real sector. *The Quarterly Journal of Economics*, 112(3):663–691, 1997.
- Oscar Jordà, Moritz Schularick, and Alan M. Taylor. Financial crises, credit booms, and external imbalances: 140 years of lessons. *NBER working paper*, 2010.
- Martin Lettau and Stijn Van Nieuwerburgh. Reconciling the return predictability evidence. *Review of Financial Studies*, 21(4):1607–1652, 2008.
- Adriano A. Rampini and S. Viswanathan. Financial intermediary capital. *working paper*, 2012.
- Carmen M. Reinhart and Kenneth S. Rogoff. *This time is different: Eight centuries of financial folly*. Princeton University Press, Princeton, NJ, 2009.
- Thomas A. Rietz. The equity risk premium a solution. *Journal of Monetary Economics*, 22(1):117 – 131, 1988.

Table 1: I run regressions of asset prices and macro variables on event indicators. Each indicator takes on the value 1 if the event occurred at the specified lag. The variables: dp is dividend yield, r is real stock return, $r - r_f$ is excess return, Δd is dividend growth, cs is credit spread, cfn and drn are cash flow news and discount rate news extracted from a VAR as in Campbell 1991, Δc is consumption growth, sc is log surplus consumption, σ_c^2 is the variance of consumption growth. See main text for a description of these variables. All regressions include country fixed effects and some include a linear time trend as indicated to deal with decreasing dividend yields and macro volatility over time. All numbers are reported in percent per annum. I also include a lag of the dependent variable in the controls. Stars (*) indicate significance at 10%, 5%, 1% levels. Standard errors clustered by year.

Response of asset prices and macro state variables to events										
$y_t = \gamma_i + \sum_{j=0}^T a_j 1_{fin,t-1} + \sum_{j=0}^T b_j 1_{recession,t-1} + \sum_{j=0}^T c_j 1_{war,t-1} + dy_{t-1} + et + \varepsilon_{t+1}$										
y variables										
	Risk premia and cash flows							C state variables		
	dp	r	$r - r_f$	Δd	cs	cfn	drn	Δc	sc	σ_c^2
Fin Crisis										
0	3.8	3.7	1.6	11.2	5.0***	5.1	1.3	0.7	1.8	-0.2*
1	19.6***	-22.6***	-26.8***	-2.9	-0.4	-10.0**	13.3***	-1.5	-7.0	-0.1
2	-6.9	-0.8	-5.1	-9.0	-0.5	-7.2	-4.8*	-1.1	-6.6	-0.0
3	-6.4	8.1	2.5	8.4	-1.6	2.1	-6.0**	-0.1	0.1	0.0
4	-3.6	16.2***	15.5**	10.4	2.8**	13.5**	-2.0	-0.3	-2.4	0.1
5	1.2	0.8	-1.1	3.2*	-1.3	5.8	1.8	1.1	5.1	0.1
Recession										
0	8.5***	-15.2***	-14.4***	-9.7**	-0.2	-10.3***	5.0***	-1.1**	-4.2	0.2***
1	-0.7	-3.0	-2.9	-4.8	0.3	-6.6**	-1.7	-3.7***	-13.8	0.1**
2	-1.2	1.5	1.6	-3.3	1.3**	-2.7	-0.5	-2.0***	-6.2	0.2***
3	-0.1	-2.1	-1.1	-2.1	0.9	-3.8	-0.6	-1.4***	-4.5	0.1**
4	2.2	4.5	5.2	4.6	-0.7	3.9	-0.2	-0.8	-1.7	0.2**
5	-1.8	2.3	2.3	-4.5	-0.1	-3.3	-1.6	-2.0***	-6.7	0.1**
War										
0	3.3	-2.5	-0.8	3.7	-0.5	0.4	2.7	-1.0	-4.0	1.7***
1	-18.9**	13.7	20.5**	-3.4	-1.7	5.8	-10.4**	-17.6***	-54.8	1.7***
2	-4.9	-33.5***	-26.5**	-29.4*	-1.2	-34.3***	-2.0	-10.0***	-22.0	1.0***
3	-13.1	-32.7***	-22.2**	-30.7*	0.5	-33.2***	-5.4	1.6	3.0	0.4**
4	-12.5	-31.5***	-26.0**	-23.1	0.1	-23.6**	-6.0	2.9	5.3	0.2
5	-6.5	-21.4**	-19.0*	16.5	-2.0	5.9	-0.5	4.9***	16.1	0.2
y_{t-1}	0.79***	0.13***	0.14***	0.46***	0.63***			0.04	0.15	0.05*
t	Y								Y	Y
R^2	76%	13%	13%	16%	74%	15%	9%	21%	75%	27%

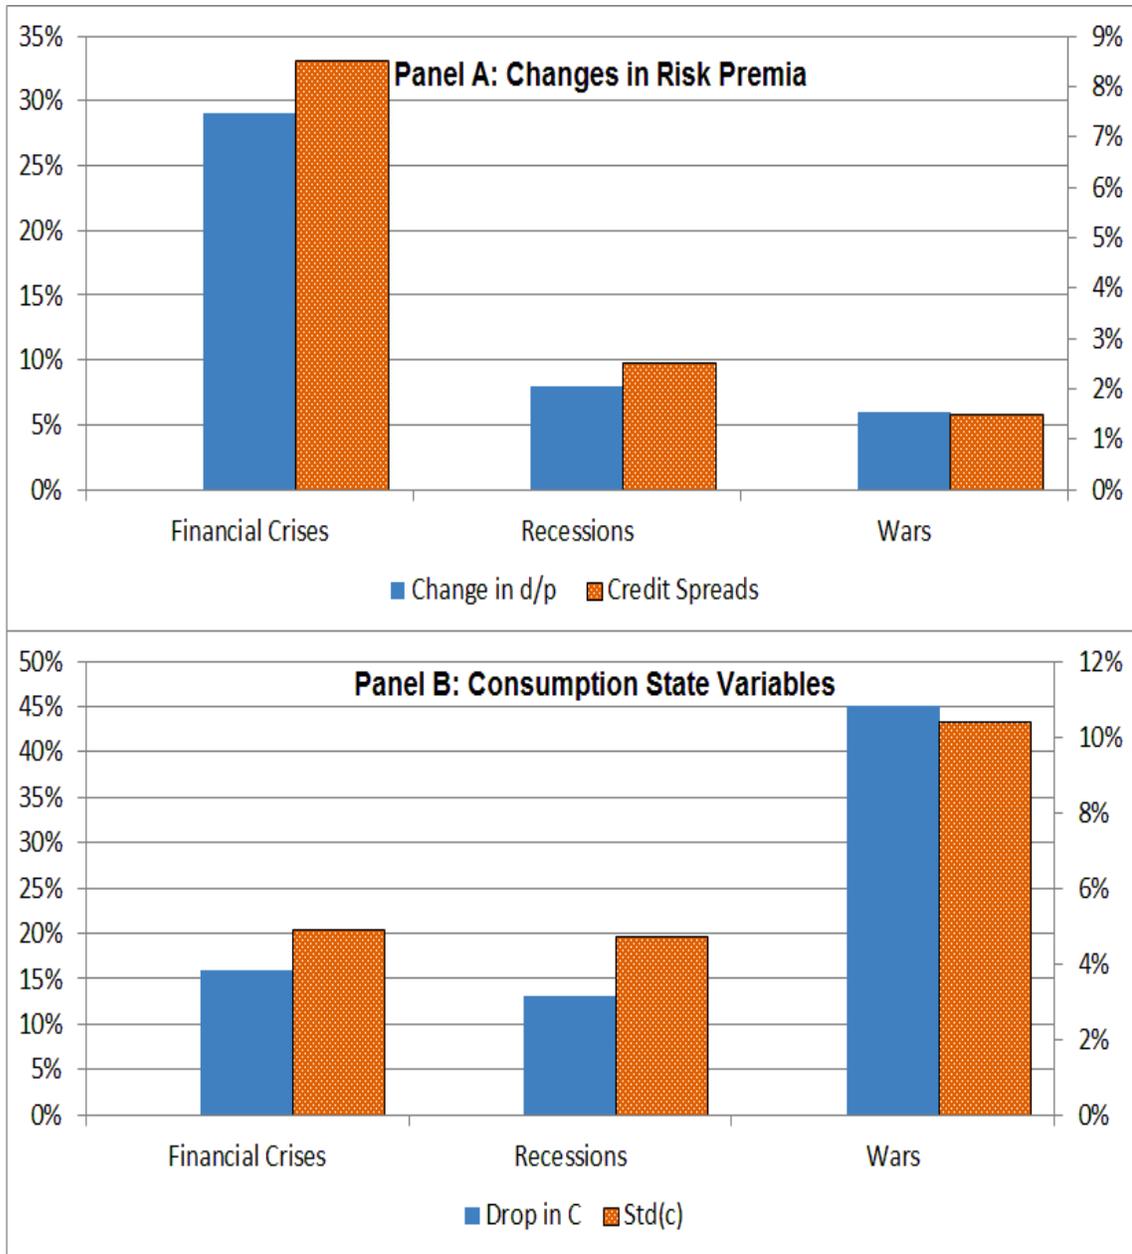


Figure 1: This figure computes changes in risk premia, as measured by dividend yields (left axis) and credit spreads (right axis), in Panel A across financial crises, recessions, and wars. Panel B plots consumption state variables argued to capture variation in risk premia: the peak to trough decline in consumption (left axis) and consumption volatility (right axis).

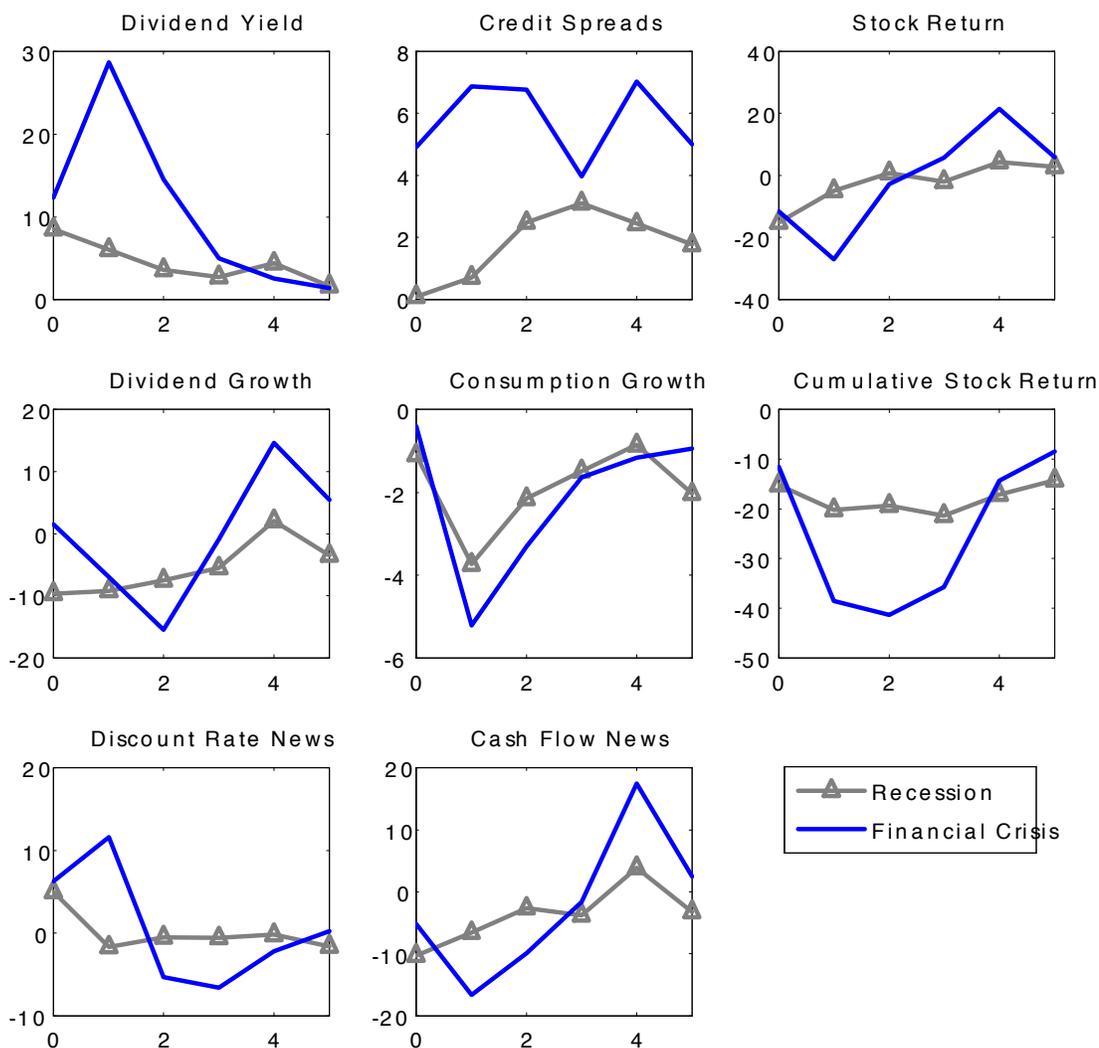


Figure 2: This figure plots the impulse response of each variable conditional on a recession and a financial crisis generated by the regressions in Table 1 where I regress each variable on event dummies and one lag. The plots for discount rate news and cash flow news are implied by a return decomposition using a VAR of returns and dividend yields.

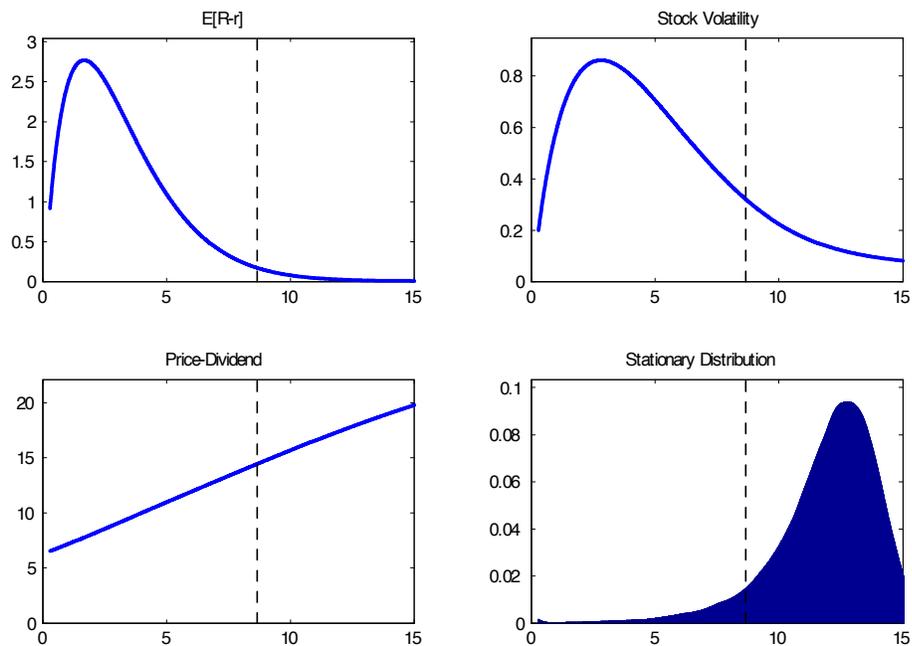


Figure 3: I plot the model solution as a function of the state variable e defined as intermediary equity divided by output. The key feature of the model solution is the spike in volatility and risk premia when e is low. The upper left panel gives the equity premium, the upper right gives equity volatility, the lower left gives the price dividend ratio, and the lower right gives the stationary distribution of the state variable e . Finally, in each panel I draw a dashed line which represents the lowest 7th percentile of realizations of the state variable which constitutes the cutoff for a financial crisis in the model.

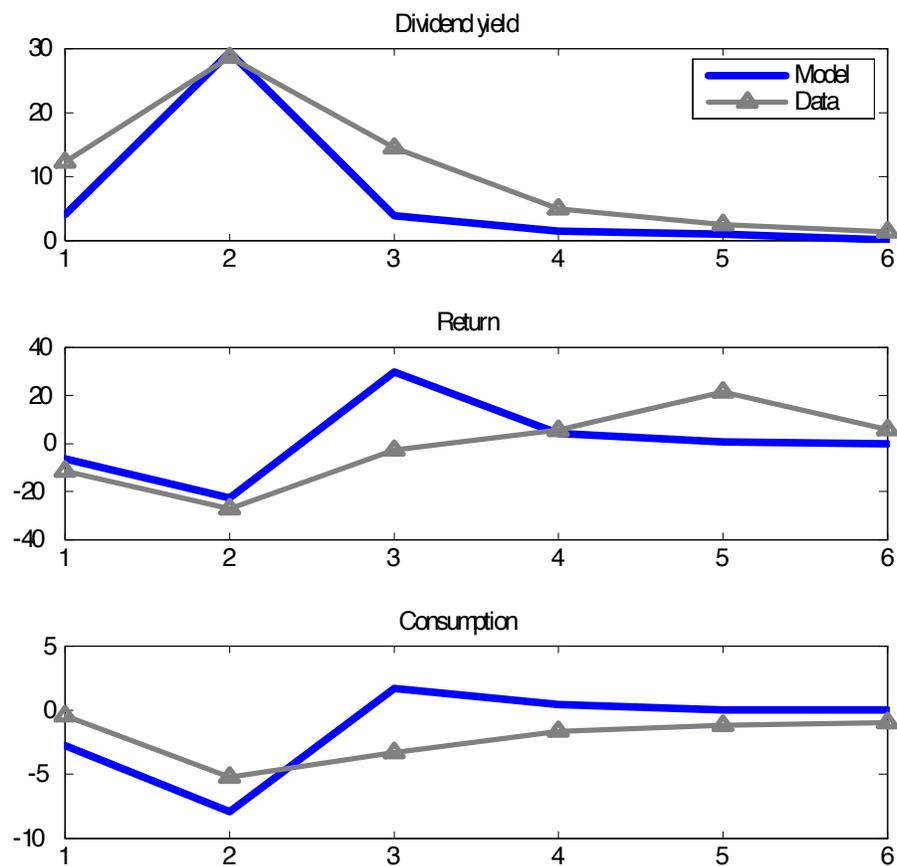


Figure 4: This figure plots the response of dividend yield and consumption growth to a financial crisis in the data and in the model. The top panel matches the timing of the peak of the crisis as occurring in year 2 in the model and data. In the middle panel I plot consumption growth in the data vs the model using the same timing convention as in the top panel. In the bottom panel I use lagged consumption growth in the model instead of current consumption growth. This allows for the fact that in the data consumption may respond with a lag rather than contemporaneously. See text for details.

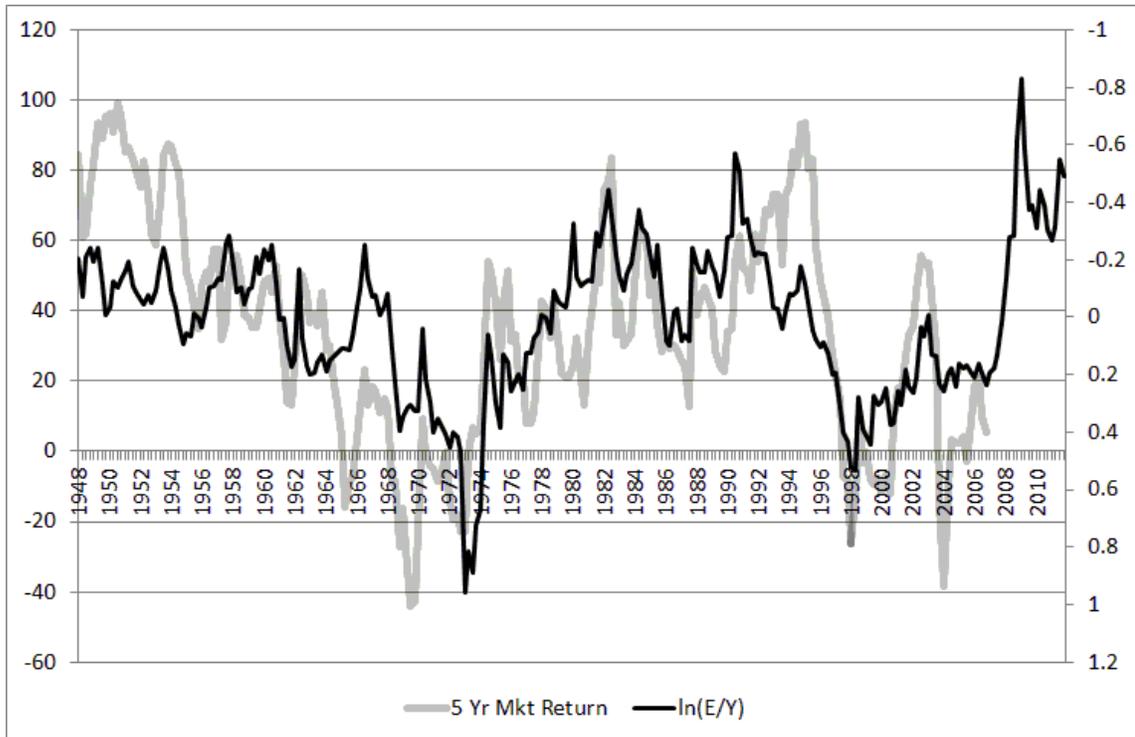


Figure 5: I plot the log of the ratio of intermediary equity to GDP (black line, right axis, decreasing scale) which is the state variable in the model, along with the subsequent 5 year excess return on the market (gray line, left axis). Intermediary equity is defined as the total market capitalization of the financial sector (SIC code of 6). The intermediary equity to GDP series is linearly detrended. The model implies that intermediary equity should forecast returns.

Table 2: This table provides calibrated parameters in the model. All values are annualized.

Parameter	Description	Targeted Moment
σ 5%	Volatility of output	Vol of consumption (or GDP)
ρ 3%	Time discount	Risk free rate, average price dividend
μ 1.8%	Growth rate	Average growth
ψ 8%	Intermediary death rate	Well behaved dynamics
δ 13%	Depreciation for HH	Lowest p , entry

Table 3: This table provides moments on quantities and asset prices implied by the model vs the data. In the model, I form 10,000 100 year long samples and compute corresponding statistics. Simulated data are monthly but reported in annualized numbers. dY/Y represents output or consumption growth in the model. For more details on the series, see data appendix.

Basic Moments (% per year):				
	Data	Model		
		Mean	95%	5%
$E[dY/Y]$	1.80	1.80	2.56	0.90
$\sigma[dY/Y]$	5.17	5.00	5.17	4.84
$P[crisis]$	7.2	7.0	8.3	5.8
$E[r_f]$	0.60	3.00	—	—
$E[R - r_f]$	7.36	6.94	10.24	3.67
$\sigma[R - r_f]$	18.95	19.98	22.03	18.16
$\frac{E[R - r_f]}{\sigma[R - r_f]}$	0.39	0.35	0.50	0.19
$E[\sigma_{R,t}]$	17.98	16.39	16.82	15.83
$\sigma[\sigma_{R,t}]$	9.20	10.30	11.40	9.40
$E[\ln(p/d)]$	3.28	2.84	2.85	2.83
$\sigma[\ln(p/d)]$	0.29	0.12	0.13	0.11
$AC1[\ln(p/d)]$	0.81	0.50	0.55	0.44

Table 4: I run regressions of asset prices and macro variables on event indicators in the Model and in the Data. The variables: dp is dividend yield, r is return, Δd is dividend growth and Δc is consumption growth. In the model, ΔY represents the cash flow shock, but Y represents the aggregate dividend, output (GDP), or consumption because it is an endowment economy. See main text for a description of these variables. All regressions in the data include country fixed effects and some include a linear time trend as indicated to deal with decreasing dividend yields and macro volatility over time. All numbers are reported in percent per annum. I also include a lag of the dependent variable in the controls. Stars (*) indicate significance at 10%, 5%, 1% levels. Standard errors clustered by year.

Response of variables in financial crisis vs typical recession							
$y_t = \gamma_i + \sum_{j=0}^T a_j 1_{fin,t-1} + \sum_{j=0}^T b_j 1_{recession,t-1} + \varepsilon_{t+1}$							
Model				Data			
Risk premia		CF	Risk premia		CF		
dp	r	ΔY	dp	r	Δc	Δd	
Fin Crisis							
0	2.7	-3.0	-1.2	3.8	3.7	0.7	11.2
1	21.6	-13.3	-3.1	19.6***	-22.6***	-1.5	-2.9
2	-4.8	21.8	0.4	-6.9	-0.8	-1.1	-9.0
3	-0.9	2.9	0.0	-6.4	8.1	-0.1	8.4
4	0.0	0.0	0.0	-3.6	16.2***	-0.3	10.4
5	0.0	0.0	0.0	1.2	0.8	1.1	3.2*
Recession							
0	0.0	1.1	0.3	8.5***	-15.2***	-1.1**	-9.7**
1	8.6	-16.5	-7.5	-0.7	-3.0	-3.7***	-4.8
2	2.0	7.1	0.0	-1.2	1.5	-2.0***	-3.3
3	1.1	2.9	0.0	-0.1	-2.1	-1.4***	-2.1
4	0.7	1.4	0.0	2.2	4.5	-0.8	4.6
5	0.5	0.6	0.0	-1.8	23	-2.0***	-4.5
Response in financial crises: $y_t = \gamma_i + \sum_{j=0}^T a_j 1_{fin,t-1} + \varepsilon_{t+1}$							
Model				Data			
Risk premia		CF	Risk premia		CF		
dp	r	ΔY	dp	r	Δc	Δd	
Fin Crisis							
0	4.9	-5.9	-3.0	12.0***	-8.9*	0.6	3.3
1	27.5	-21.6	-7.8	18.7***	-22.1***	-4.0***	-5.3
2	-4.7	29.3	0.0	-7.4*	2.55	-1.6*	-11.1*
3	-1.9	4.9	0.0	-5.5	6.51	-0.5	6.3
4	0.0	0.0	0.0	-0.5	20.7***	-0.6	13.8**
5	0.0	0.0	0.0	3.9	1.5	-0.6	1.2

Table 5: This table compares the predictive power of intermediary capital in the model vs the data by running predictive regressions of returns on lagged intermediary equity. In the data, I use the total market valuation of the financial sector divided by GDP, analogous to the model. When forecasting asset returns, I include a linear time trend for intermediary equity as this variable is increasing over time. I compare performance to the log price-dividend ratio in the data. T-stats computed using Newey-West with lags depending on horizon. Panel B runs a cross-sectional asset pricing test using 35 portfolios (25 size and book to market portfolios and 10 momentum portfolios) to test whether intermediary equity is “priced” in the cross-section of asset returns. Shanken t-stats reported below. Data sources: $R^e mkt$ and $R^e fin$ are the market and financial sector excess returns, respectively. All stock returns are from Kenneth French’s website. $R^e corp$ is an excess corporate bond return from Barclays constructed as the return on maturities between 3 and 5 years. All data are from 1948-2012 except corporate bond returns which are from 1988-2012. See appendix for additional details.

Panel A: Predicting Excess Returns: $R_{t+k}^e = \beta_1 x_t + \beta_2 t + \varepsilon_t$							
Return	Model: $\ln(e_t)$		Data: $\ln(e_t)$		Data: [$\ln(e_t), \ln(P/D)$]		
	β_e	R^2	β_e	R^2	β_e	β_{pd}	R^2
$R^e mkt_{t+1}$	-0.42	3%	-0.26***	2%	-0.25***	-0.02	1%
$R^e mkt_{t+3}$	-0.39	7%	-0.27***	5%	-0.26***	-0.02	5%
$R^e mkt_{t+12}$	-0.28	21%	-0.28***	17%	-0.25***	-0.04	18%
$R^e mkt_{t+60}$	-0.10	32%	-0.27***	44%	-0.22***	-0.08*	47%
$R^e fin_{t+12}$	-0.92	51%	-0.30***	20%	-0.28***	-0.04	19%
$R^e corp_{t+12}$	-0.16	32%	-0.09***	17%	-0.08***	-0.01	17%

Panel B: Cross-Sectional Asset Pricing: $E[R^e] = a + \lambda_f \beta_f + \varepsilon_t$							
	a	mkt	smb	hml	mom	$\ln(e_t)$	Adj R^2
Intermediary Model	5.96	0.63				0.43	49%
t-stat	(0.87)	(0.34)				(2.76)	
Benchmark Model	7.55	0.15	1.45	3.53	13.91		86%
t-stat	(1.99)	(0.04)	(1.01)	(2.08)	(4.55)		

Online Appendix: Financial Crises and Risk Premia

Tyler Muir

Email: tyler.muir@yale.edu

March 2014

Abstract

This appendix supplements the main text. It describes the data sources in detail and performs various robustness checks.

A0.1 Robustness

I show that the results in the main table (Table 1 in the main text) are robust to subsamples where I look specifically at data from 1950 onwards. The similarity of these results to the main results suggests that the main findings are still relevant in more modern economies. This is important because the nature of the banking sector, regulation, and government intervention in crises has changed over time. Table A1 gives these results. In fact, many of the results are stronger in this subsample. However, one caveat is that there are only 12 crises in the post-1950 sample. Next, I look at using just the Reinhart and Rogoff (2009) dates in Table A2. Again, the results are qualitatively similar. However, I can no longer control for typical recessions. Another big difference with the RR dates is the timing. Whereas my main dates come from a business cycle peak, the RR dates are at the occurrence of the crisis. Therefore, some variables may have already responded since by the time the crisis hits there has typically been bad news (crises tend to occur after business cycle peaks). This is likely also why dividend yields respond more on impact, rather than with a lag as in the main table, because we are conditioning on a later date. Lastly, Figure A1 plots variables in US data using higher frequency (monthly) data. The dates used are from Gorton (1988) and Bordo and Haubrich (2012). We see the same patterns in US data and also see stock volatility being much higher in financial crises. This is also consistent with the model but without high frequency historical data I can not compute conditional stock volatility.

Next, I assess whether dividend yields forecast returns or dividend growth during financial crises. We know that changes in dividend yields must represent either changes in expected dividend growth or changes in expected returns. Standard predictive regressions indicate that it is the latter. However, maybe in financial crises dividend yield changes are about dividend growth and not returns. Table A3 runs standard predictive regressions including an interaction term of lagged dividend yields with an indicator for whether a financial crisis or recession has occurred in the past 5 years. The coefficients on the interaction terms are small in magnitude and not statistically significant. This indicates that there is no differential effects in these times so the relationship is unchanged. It is therefore appropriate to interpret changes in dividend yields based on the standard unconditional forecasting regressions which indicate these are mostly related to expected returns.

As mentioned in the main text, the rare disasters story is perhaps the hardest to evaluate because measuring the probability of a disaster is difficult. Here I show that three events that were unexpected increases to the probability of disaster did not have significant effects on risk premia. These results are shown in Figure A2 which plots log dividend yields and the

BaaAaa default spread during select episodes. In the top panel I plot a shaded indicator for Hitler invading Poland and for the attack on Pearl Harbor, both events that clearly increased the probability of a war related disaster for the US. Note in the rare disasters calibrations, even very small increases in disaster probabilities should have large effects on risk premia. An increase in disaster probability of 1 or 2% should roughly double the conditional equity premium. Also, the model would predict a one time sharp increase in risk premia, so we want to look for a spike during only those events, rather than focus on the trends. The evidence does not point to a clear increase and risk premia: credit spreads and dividend yields appear to go down when Hitler invades Poland (the catalyst for WWII). Dividend yields barely increase during Pearl Harbor, while credit spreads are roughly flat or possibly decline. Moving to the lower panel which looks at the Cuban Missile Crisis, we also see a slight increase in dividend yields in that month (about 5%), but it is basically flat. Credit spreads are completely flat. The long slow increase in the early part of 1962 is known as the Kennedy Slide where the stock market experienced large declines after a long period of growth, but this is prior to the Cuban Missile Crisis. For the rare disasters story to work, one has to believe that during the month the Cuban Missile Crisis was occurring there was no meaningful increase in the probability of a war related disaster, which seems implausible. This evidence is therefore consistent with the main text that the war related disasters are unlikely to explain variation in risk premia. Of course one can never know exactly what investors believed at the time or exactly how much the probability of disaster would have increased in those events. Figure A3 plots the time-series of the dividend yield and BaaAaa spread monthly in US data and highlights events as financial crises, wars, or recessions. It shows that the largest increases in these series occur during financial crises (the Great Depression and 2008 crisis).

A0.2 Details on VAR

I run a VAR of returns and dividend yields following Campbell (1991). Linearizing the return equation and iterating forward yields

$$\begin{aligned} r_{t+1} &= k + dp_t - \rho dp_{t+1} + \Delta d_{t+1} \\ dp_t &= \frac{k}{1-\rho} + \sum \rho^j r_{t+j} - \sum \rho^j \Delta d_{t+j} \end{aligned}$$

Rearranging, and applying $\Delta E_{t+1} \equiv E_{t+1} - E_t$ to both sides,

$$r_{t+1} - E_t[r_{t+1}] = \Delta E_{t+1} \sum \rho^j \Delta d_{t+j} + \Delta E_{t+1} \sum \rho^j r_{t+j+1}$$

We can compute these forecasts using a VAR of returns on dividend growth. Let $x = [r_{t+1}, dp_{t+1}]$ and assume both variables are de-meant

$$x_{t+1} = \Gamma x_t + \Psi_{t+1}$$

One can show

$$\Delta E_{t+1} \Sigma \rho^j r_{t+j} = e1' \rho \Gamma (I - \rho \Gamma)^{-1} \Psi_{t+1}$$

where $e1$ is a vector whose first element is 1 and the rest are zeros. This equation summarizes discount rate news. Cash flow news is defined as the residual, or the movement in unexpected returns not related to discount rate news:

$$\Delta E_{t+1} \Sigma \rho^j \Delta d_{t+j} = r_{t+1} - E_t[r_{t+1}] - \Delta E_{t+1} \Sigma \rho^j r_{t+j}$$

Following Campbell (1991) I use $\rho = 0.96$ for annual data. I also de-mean price dividend ratios and returns separately for each country.

A0.3 Details of Model Solution

The ODE to solve is

$$p'' = \frac{2 \left[\frac{p^2}{e} \sigma_R^2 - p\mu + pr - p'\mu_e - p'\sigma\sigma_e - 1 \right]}{\sigma_e^2}$$

Where we can substitute in the means and volatilities (in terms of first order terms) using the expressions in the text.

I use matlabs `bvp4c` function to solve the ODE on a grid $[e, \bar{e}]$ by specifying the boundary conditions. At \underline{e} we have the price falling down to $\frac{1}{r+\delta}$ which is the condition for entry, hence, $p(\underline{e}) = \frac{1}{r+\delta}$. We also know that the price will not change on entry, thus $p'(\underline{e}) = 0$. I search for the endogenous value of entry \underline{e} that satisfies these two boundary conditions by imposing the lower boundary $p(e^*) = \frac{1}{r+\delta}$ and running through values of e^* until $p'(e^*) = 0$. It turns out that the economy very rarely hits this boundary. The other condition is $p(\infty)$. Intuitively, we know when E goes to ∞ prices no longer depend on intermediary equity and hence $p'(\infty) = 0$. In solving the ODE numerically, I choose a finite upper bound \bar{e} and ensure that the process rarely reaches this bound. For higher values of e , the drift is increasingly negative since the equity premium goes to zero, thus $\mu_e = e(\sigma^2 - \psi) < 0$. I verify that the solution is not dependent on the choice of this upper bound. Lastly, we need to verify ex-post that $p(e) \geq \frac{1}{r+\delta}$ for $e > \underline{e}$ so that the household never steps in to buy the asset. This is easily verified ex-post by showing that $p'(e) \geq 0$ for $e > \underline{e}$ which intuitively just says that prices are increasing in intermediary capital.

A0.4 Details on Data Sources

BaaAaa spread: Federal Reserve’s FRED database series AAA and BAA represent Moody’s corporate bond yields. I begin using this data monthly from 1930-present for the US only. This supplements my credit spread data which end in 1930.

Credit spreads: Credit spread data provided by the International Center for Finance at Yale. The data set is constructed from the Investor’s Monthly Manual from 1860-1930 and contains bond price data. I form credit spreads by taking all corporate bond yields over UK government yields.

War and War Related Disaster Dates: I use war-related disasters from Barro (2006) (see Table I Part A. I use the 20 OECD countries only due to lack of availability of historical dividend yields for any of the Latin or Asian countries). Note: in Barro (2006) every disaster is war related (WWI, WWII, or aftermath) or related to financial crises (Great Depression). Results are robust if augmented with dates the U.S. entered – or nearly entered – into a major war: 1898 (Spanish-American), 1916 (WWI), 1941 (WWII), 1950 (Korea), 1955 (Vietnam), 2001-02 (Afghanistan, Iraq), 1962 (Cuban Missile Crisis).

Financial Crisis Dates: My main dates come from Jorda *et al.* (2010) (ST) and Reinhart and Rogoff (2009) (RR). The difference is the ST dates date business cycle peaks whereas the RR dates use the occurrence of the actual crisis. When focusing only on US data, Gorton (1988) and Bordo and Haubrich (2012) contain a history of US business cycles categorized as banking crises or not (much of their categorization is based on Friedman and Schwartz (1971), and the resulting dates are similar to Jorda *et al.* (2010)). When using Bordo and Haubrich (2012) and Gorton (1988) I date the financial crisis based on their banking crisis dates (see Gorton (1988) Tables 1 and 6, and Bordo and Haubrich (2012) Table 2). For Bordo and Haubrich (2012) I drop 1975 as most authors do not consider this a crisis (i.e. Reinhart and Rogoff (2009), Jorda *et al.* (2010)). The results are robust to including additional dates used in other studies: for example, 1973-1975, 1988-1991 (Lopez-Salido and Nelson (2010)), 1984 (Reinhart and Rogoff (2009)). For US recession dates I use NBER.

GDP and Consumption Data: Robert Barro’s website (see, e.g., Barro *et al.* (2011)).

Country Level Price, Return, and Dividend Yield Data: All indices for all countries are from Global Financial Data. All price series are real values of stock indices in U.S. dollars. Real dividends are constructed using prices and returns.

U.S. Equity Returns: From Kenneth French’s website when possible. When calculating monthly volatility of the stock market and “vol of vol,” however, I use daily S&P500

observations from CRSP. I also pull the momentum, size and book-to-market portfolios from Kenneth French's website.

References

Robert Barro, Emi Nakamura, Jon Steinsson, and Jose Ursua. Crises and recoveries in an empirical model of consumption disasters. *working paper*, 2011.

Robert J. Barro. Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics*, 121(3):823–866, 2006.

Michael D. Bordo and Joseph G. Haubrich. Deep recessions, fast recoveries, and financial crises: Evidence from the american record. *NBER working paper*, 2012.

John Y. Campbell. A variance decomposition for stock returns. *Economic Journal*, 101:157–179, 1991.

Milton Friedman and Anna Jacobson Schwartz. *A monetary history of the United States, 1867-1960*. Princeton University Press, Princeton, NJ, 1971.

Gary Gorton. Banking panics and business cycles. *Oxford Economic Papers*, 40(4):751 – 781, 1988.

Oscar Jorda, Moritz Schularick, and Alan M. Taylor. Financial crises, credit booms, and external imbalances: 140 years of lessons. *NBER working paper*, 2010.

David Lopez-Salido and Edward Nelson. Postwar financial crises and economic recoveries in the united states. *working paper*, 2010.

Carmen M. Reinhart and Kenneth S. Rogoff. *This time is different: Eight centuries of financial folly*. Princeton University Press, Princeton, NJ, 2009.

Table A1: I repeat the results from the main text using the sample post 1950. I run regressions of asset prices and macro variables on event indicators. Each indicator takes on the value 1 if the event occurred in any of the past three years. The variables: dp is dividend yield, r is return, $r - r_f$ is excess return, Δd is dividend growth, cs is credit spread, cfn is “cash flow news”, drn is discount rate news, Δc is consumption growth, sc is log surplus consumption, σ_c^2 is consumption variance. See main text for a description of these variables. All regressions include country fixed effects and some include a linear time trend as indicated to deal with decreasing dividend yields and macro volatility over time. All numbers are reported in percent per annum. For persistent variables (dp and habits), I also include a lag in the controls. Stars (*) indicate significance at 10%, 5%, 1% levels. Standard errors clustered by year.

Response of asset prices and macro state variables to events									
$y_t = \gamma_i + \sum_{j=0}^T a_j 1_{fin,t-1} + \sum_{j=0}^T b_j 1_{recession,t-1} + cy_{t-1} + dt + \varepsilon_{t+1}$									
y variables									
	Risk premia and cash flows					C state variables			
	dp	r	$r - r_f$	Δd	cfn	drn	Δc	sc	σ_c^2
Fin Crisis									
0	2.4	8.2	8.8	12.0**	10.8	2.2	0.3	-0.1	-0.00
1	48.4***	-36.8***	-36.8***	-2.0	-10.6	26.5***	-1.6	-6.3**	-0.00
2	-21.2**	14.8*	17.2*	-16.9***	-4.5	-12.0***	-1.9***	-12.9***	-0.01
3	-4.7	-6.3	-6.6	7.0	-14.8	-6.6	-0.6	-2.0	-0.02
4	3.9	0.6	-0.5	0.2	4.1	4.4	-1.3	-7.0	-0.02*
5	7.6	-11.1	-10.8	-6.8	-7.3	5.2	0.1	5.0	-0.02*
Recession									
0	20.3***	-26.3***	-25.4***	-6.6**	-18.1***	9.1***	-1.8***	-7.4***	0.01**
1	-4.2	-7.3*	-7.5*	-12.9***	-14.5***	-3.7*	-3.0***	-12.9***	0.01**
2	1.2	1.5	1.5	-7.7***	-0.6	-1.9	-0.8**	-2.7*	0.01
3	3.5	-3.9	-3.8	-2.6	-3.8	1.1	-0.5	-2.6*	0.01
4	3.8	2.4	2.7	2.5	1.9	0.9	0.1	0.2	0.01
5	-2.2	9.0**	8.7**	0.7	6.0*	-2.0	-0.4	-2.1	0.00
y_{t-1}	0.74	0.04	0.04	0.14***			0.25***	0.25***	0.01***
t	Y						Y		Y
R^2	76%	14%	14%	16%	11%	15%	48%	93%	40%

Table A2: I run regressions of asset prices and macro variables on event indicators. Each indicator takes on the value 1 if the event occurred in any of the past three years. The variables: dp is dividend yield, r is return, $r - r_f$ is excess return, Δd is dividend growth, cs is credit spread, cfn is “cash flow news”, drn is discount rate news, Δc is consumption growth, sc is log surplus consumption, σ_c^2 is consumption variance. See main text for a description of these variables. All regressions include country fixed effects and some include a linear time trend as indicated to deal with decreasing dividend yields and macro volatility over time. All numbers are reported in percent per annum. For persistent variables (dp and habits), I also include a lag in the controls. Stars (*) indicate significance at 10%, 5%, 1% levels. Standard errors clustered by year. Dates used are from [ReinRog2012](#) (RR) and differ from those in the main text. Please read the data description for details.

Response of asset prices and macro state variables to events									
$y_t = \gamma_i + \sum_{j=0}^T a_j 1_{fin,t-j} + b y_{t-1} + c t + \varepsilon_{t+1}$									
y variables									
Risk premia and cash flows						C state variables			
	dp	r	$r - r_f$	Δd	cfn	drn	Δc	sc	σ_c^2
Fin Crisis									
0	15.9***	-24.6***	-30.1***	-7.1	-14.8***	11.0***	-0.2	-2.6	-0.06
1	2.6	-0.9	-2.2	0.5	-5.8	-0.8	-2.4***	-10.0***	-0.07
2	-2.4	7.6	5.5	8.6	3.5	-2.7	-0.4	1.8	-0.06
3	-2.6	3.1	0.5	4.4	6.0	0.4	-0.7	-3.5	0.06
4	-1.8	0.7	-2.2	-2.1	0.1	-1.9	-0.6	0.2	0.07
5	-1.6	1.2	0.5	7.9	3.1	0.0	0.2	0.8	0.01
y_{t-1}	0.84***	0.19***	0.18***	0.49***			0.11***	0.22***	0.12***
t	Y							Y	Y
R^2	80%	6%	7%	14%	3%	8%	3%	72%	12%

Table A3: I run predictive regressions of returns and dividend growth on lagged dividend yields. I include indicators for financial crises and recessions. The indicator takes the value 1 if the event has occurred in the past 5 years. This tests whether the relationship between dividend yields, returns, and dividend growth changes during these periods. Small and insignificant coefficients indicate that it does not. See main text for data details.

Panel A: Return predictability in crises and recessions					
	$R_{t+1} - r_{f,t} = a_i + b \ln(d/p)_t + c 1_{fin} \ln(d/p)_t + d 1_{non,fin} \ln(d/p)_t + et + \varepsilon_{t+1}$				
	b	c	d	e	Adj R^2
Raw dp	0.06				2%
(t-stat)	(3.25)				
Raw dp	0.080	-0.003	0.018		2%
(t-stat)	(3.80)	(-0.20)	(-1.68)		
With time trend	0.096			0.12	3%
(t-stat)	(4.57)			(3.70)	
With time trend	0.11	0.00	-0.01	0.11	3%
(t-stat)	(4.70)	(0.00)	(-1.06)	(3.32)	

Panel A: Dividend growth predictability in crises and recessions					
	$\ln\left(\frac{d_{t+1}}{d_t}\right) = a_i + b \ln(d/p)_t + c 1_{fin} \ln(d/p)_t + d 1_{non,fin} \ln(d/p)_t + \varepsilon_{t+1}$				
	b	c	d		Adj R^2
Raw dp	-0.001				3%
(t-stat)	(-4.78)				
Raw dp	-0.001	0.000	-0.001		3%
(t-stat)	(-4.07)	(0.37)	(-0.67)		
With time trend	-0.001			-0.00	3%
(t-stat)	(-4.29)			(-0.14)	
With time trend	-0.001	0.000	-0.000	-0.00	3%
(t-stat)	(-3.86)	(0.34)	(-0.69)	(-0.21)	

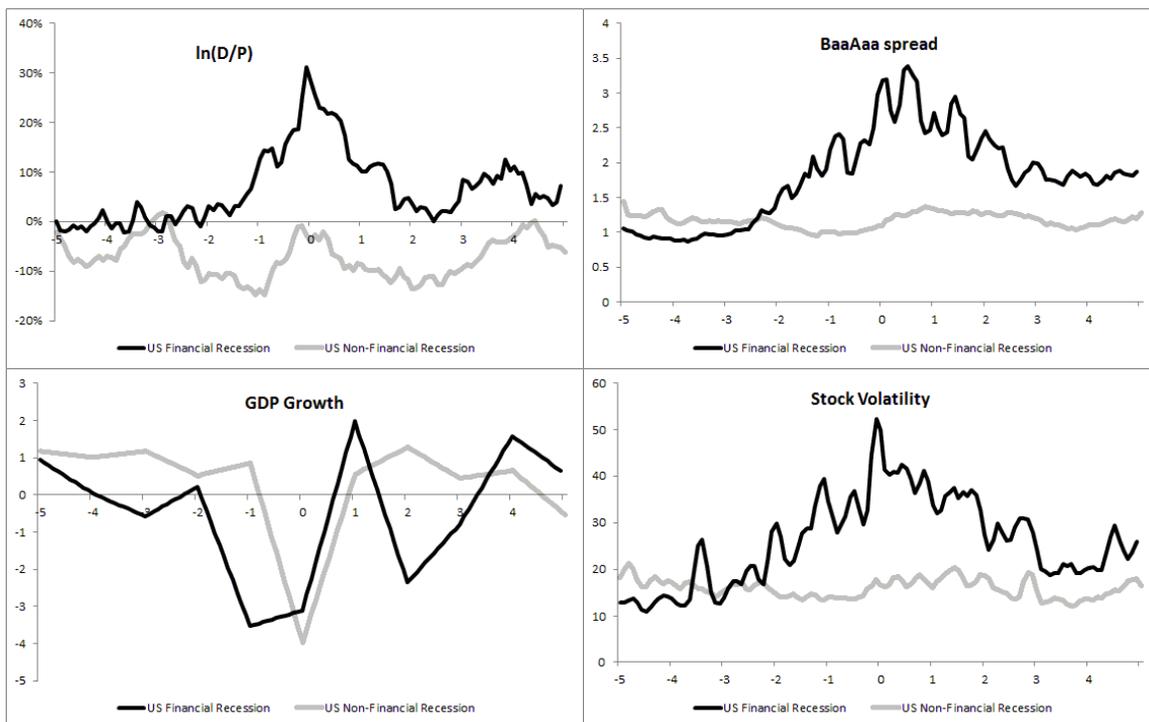


Figure A1: I compare financial recessions (a recession for which a financial crisis occurs) in black lines to non-financial recessions (those without a financial crisis) in gray lines in U.S. data. I use recession dates from NBER and financial crisis dates from Gorton (1988).

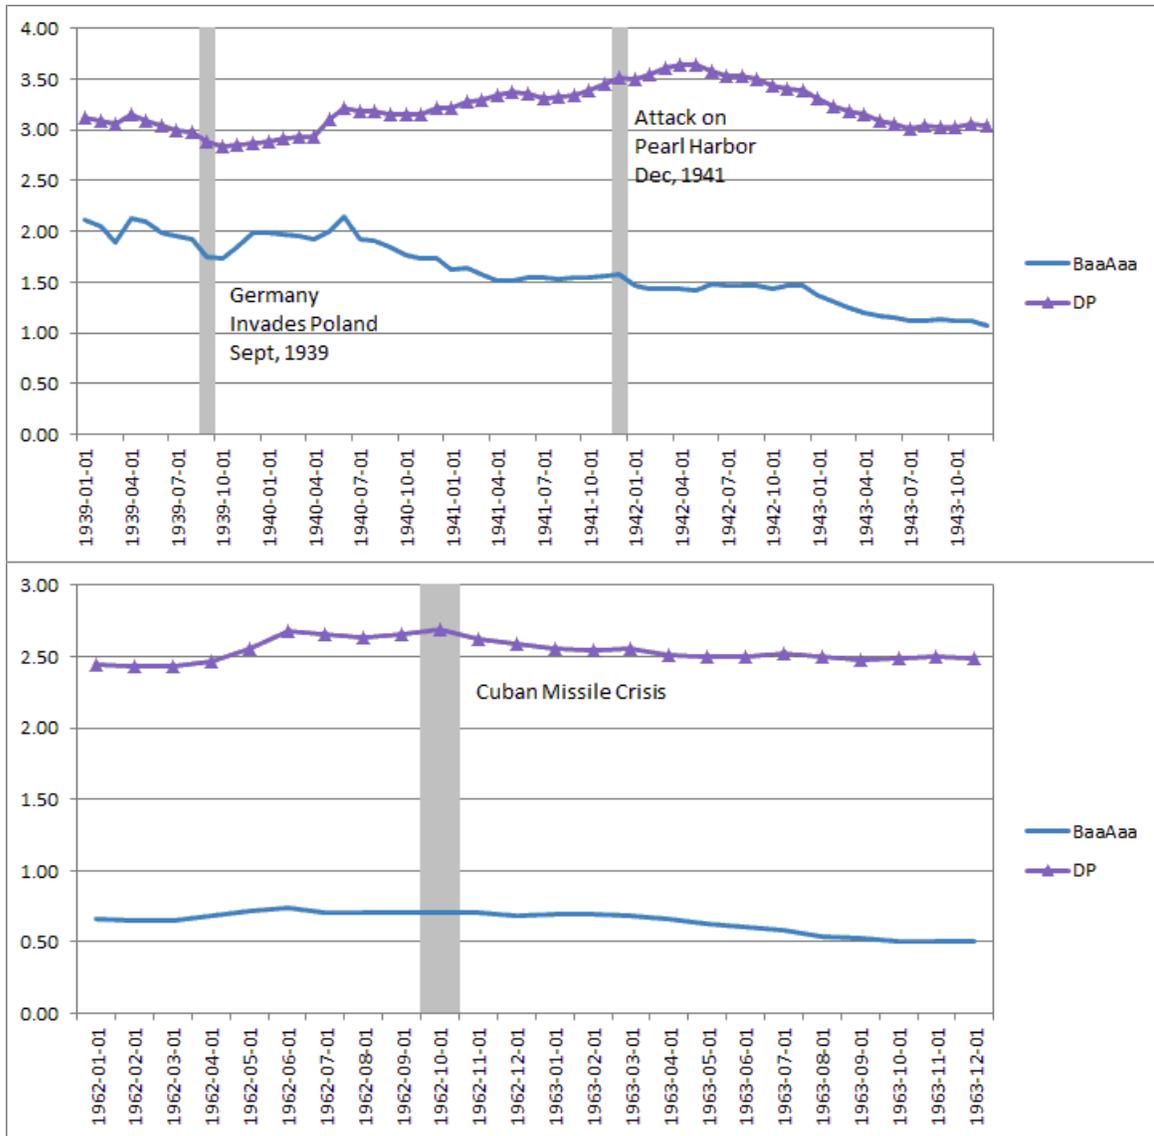


Figure A2: This figure plots the behavior of log dividend yields and the BaaAaa default spread using monthly US data. I analyze the behavior around events which were unexpected and increased the probability of a war related disaster. The top panel shows Germany invading Poland and the Attack on Pearl Harbor, while the bottom shows the Cuban Missile Crisis.

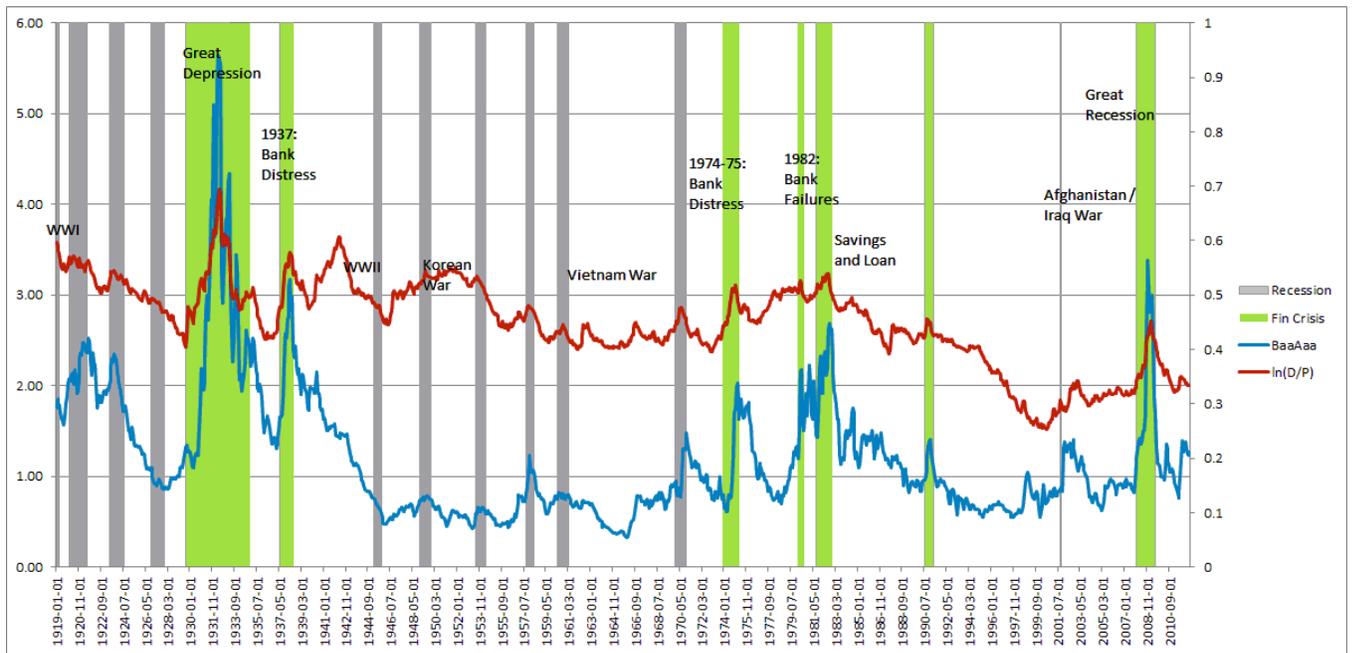


Figure A3: This figure plots the log dividend yield (red) and BaaAaa spread (blue) monthly from 1919 to 2012. Green shaded areas indicate financial crises or events considered a “credit crunch” involving bank failures. Grey areas indicate recessions and also label times the US entered into a major war.