ESSAYS IN EMPIRICAL ASSET PRICING AND MACRO-FINANCE

A Dissertation

Presented to

The Faculty of the C.T. Bauer College of Business

University of Houston

In Partial Fulfillment

Of the Requirements for the Degree

Doctor of Philosophy

By

Mohammad Ghaderi

May, 2021

ESSAYS IN EMPIRICAL ASSET PRICING AND MACRO-FINANCE

Approved:

Kris Jacobs, Professor of Finance (Chair)

C.T. Bauer College of Business, University of Houston

Sang Byung Seo, Assistant Professor of Finance (Chair)

Wisconsin School of Business, University of Wisconsin-Madison

Hitesh Doshi, Associate Professor of Finance

C.T. Bauer College of Business, University of Houston

Mete Kilic, Assistant Professor of Finance

Marshall Business School, University of Southern California

Paul A. Pavlov, Dean

C.T. Bauer College of Business, University of Houston

Acknowledgments

Many people provided me with assistance and support in the completion of my studies and this thesis.

I express my sincere gratitude to Professor Kris Jacobs for his continuous support throughout my studies, both as my advisor and coordinator of the finance Ph.D. program. Professor Sang Byung Seo was always a kind mentor, helpful advisor, wonderful co-author, and dear friend. Undoubtedly, I am lucky for having the chance to know him and cannot be grateful enough for all his guidance. I appreciate Professor Mete Kilic's generous help, reading and discussing my work, and providing me with constructive feedback. Also, I should thank Professor Hitesh Doshi for his encouragement and valuable comments. Frequently, I dropped by Professor Rauli Susmel's office to ask my questions. I am grateful that he was always available to give me research and life advice.

I am indebted to my parents who devoted so much love and supported me through all these years. Lastly, I would like to acknowledge with gratitude the help, patience, and sacrifices of my lovely wife, Mahya. Without her, I could not reach the finish line.

Abstract

This dissertation is composed of three essays in Empirical Asset Pricing and Macro-Finance. In the first essay, titled "Can Time-Varying Risk Premia and Household Heterogeneity Explain Credit Cycles?," I use micro-level data from almost 50 million mortgages to measure the dispersion in the credit quality of borrowers in the housing market. I show that credit dispersion forecasts regional real economic activity and provide empirical evidence that associates the predictive power of dispersion with heterogeneity in the exposure of households' labor income to economy-wide shocks. I explain these observations in a model featuring time-varying risk premia, incomplete markets, and household heterogeneity. Due to risk aversion, the consumption and investment responses of households have a convex association with their labor income exposure to aggregate risks. As a result, dispersion forecasts the aggregate output more strongly in more heterogeneous regions, consistent with the data.

The second essay is joint work with Mete Kilic and Sang Byung Seo. We develop a model that generates slowly unfolding disasters not only in the macroeconomy but also in financial markets. In our model, investors cannot exactly distinguish whether the economy is experiencing a mild/temporary downturn or is on the verge of a severe/prolonged disaster. Due to imperfect information, disaster periods are not fully identified by investors *ex ante*. Bayesian learning induces equity prices to gradually react to persistent consumption declines, which plays a critical role in explaining the VIX, variance risk premium, and

put-protected portfolio returns. We show that our model can rationalize the market patterns of recent major crises, such as the dot-com bubble burst, Great Recession, and COVID-19 crisis, through investors' belief channel.

In the last essay, "Is There a Macro-Announcement Premium?," co-authored with Sang Byung Seo, we argue that the average excess return over macro-announcement days substantially exaggerates the true risk premium. The conditional volatility of returns barely drops at macro-announcements. This is at odds with virtually all models that justify high macro-announcement returns through a high announcement premium. We propose an alternative explanation: macro-announcement days are, on average, with good news in existing sample periods. We develop a novel estimation approach, which reveals that high macro-announcement returns are not a manifestation of high conditional equity premiums but positive return innovations that are not averaged out in-sample. We find that macro-announcement days do not seem to operate with a separate mechanism: the patterns of macro-announcement days are well replicated by random samples from non-announcement days. Our analysis suggests that the large average macro-announcement return might not be compensation for perceived uncertainty.

Table of Contents

List of Tables

List of Figures

1 Can Time-Varying Risk Premia and Household Heterogeneity Explain Credit Cycles?

1.1 Introduction

The United States economy in the first decade of the twenty-first century experienced a cycle of boom and bust in both housing and credit markets that ended up in the 2008 financial crisis. The causal relationship between the expansion in credit markets before the crisis on the one hand, and the rise and fall of prices followed by a wave of defaults in the housing market on the other, has been a hot topic of economic research and policy debates. In the quest to answer this question, two main accounts have emerged in the literature as

I am extremely grateful to my committee chairs Kris Jacobs and Sang Byung Seo, and members Hitesh Doshi and Mete Kilic. I also thank Mahya Amani, Ehsan Azarmsa, Gurdip Bakshi, Yuchen Chen, Ted Juhl, Seyed Mohammad Kazempour, Kevin Roshak, Rauli Susmel, and seminar participants at the University of Houston, University of Kansas, University of Georgia, Cornerstone Research, 2021 SWFA meeting, and 2021 EFA meeting for their helpful comments.

the roots of the crisis: the *subprime view* and the *expectations view*.[1]

In the subprime view of the financial crisis, it is the pre-crisis credit expansion that causes the increase in house prices and the subsequent recession. Deregulations and innovations in the financial sector resulted in an inefficient credit boom: less credit-worthy households were able to access mortgages. The hot housing market ensuing this increased demand for homeownership, resulted in widespread defaults and depressed economic conditions. However, in the expectations view, too optimistic expectations about the growth rate of house prices are the driver of the expansion in credit markets. The boom and bust in the housing market cause the credit boom and bust, and eventually the financial crisis. Both these explanations for the crisis, rely on severely misaligned incentives or irrational behavior in markets' participants.

In this paper, I explore an alternative explanation based on exogenous time-variation in risk premia and household heterogeneity. What appears to be a credit cycle, could originate from optimal consumption and investment decisions of households exposed to time-variation in the risk of economy-wide shocks. I start by presenting empirical evidence consistent with my proposed explanation and complement it with a formal analysis in a theoretical model.

In my empirical analysis, I use loan-level data from Fannie Mae and Freddie Mac to construct a new measure of regional credit risk. The data allows me to use detailed household balance sheet information at the origination of mortgages. Based on Merton (1974) credit risk model, I calculate expected default frequencies (EDF) of individual mortgages and aggregate them at different geographical levels. Moreover, I define a new credit risk measure called "credit dispersion." It measures the difference between the average EDF of households that borrow today and the average EDF of households that borrowed one year

[1] See Piazzesi and Schneider (2016), Guerrieri and Uhlig (2016) and Adelino, Schoar, and Severino (2018) for an in-depth review of this literature including a discussion of these two narratives.

earlier. In other words, credit dispersion compares the current credit qualities of two sets of households: those who are increasing their leverage versus those who are decreasing their leverage. I call this new measure "credit dispersion." The definition of dispersion is motivated by Greenwood and Hanson (2013) and Gomes, Grotteria, and Wachter (2019), who examine the relationship between corporate credit risk and bond excess returns by studying the difference between EDFs of issuing and repaying corporations in the bond market.

I show that this new measure of regional credit risk forecasts regional economic activity both at the state and Metropolitan Statistical Area (MSA) levels. The observed association is both economically and statistically significant. A one standard deviation increase in the regional credit dispersion forecasts a one percent reduction in the growth rate of state-wide GDP per capita and a 0.5 percent reduction in employment growth over the following year.

These observations seem to confirm the above-stated views about the origins of the financial crisis. However, further empirical analysis reveals that they are more consistent with households' response to variation in risk premia. To be more specific, regional credit dispersion is based on the data of high-quality borrowers with sound credit conditions. My analysis reveals the data about credit scores of these borrowers does not support a story based on an inefficient credit boom or deterioration of borrowers' quality during the boom. This finding is consistent with recent evidence about the boom in mortgage originations and subsequent defaults among middle-income households during the financial crisis.[2]

More importantly, I show that the observed forecasting association is closely tied to regional heterogeneity in the exposure of households to the economy-wide shocks. I use publicly available data from the Quarterly Census of Employment and Wages, together with a simple statistical model to construct an index of heterogeneity in labor-income exposures to systematic shocks across different states through time. Using this index, I investigate the

[2] See for example Albanesi, De Giorgi, and Nosal (2017) and Foote, Loewenstein, and Willen (2016).

interaction between regional credit risk and heterogeneity of exposures. The results reveal that higher heterogeneity in a region is associated with more forecasting power for credit dispersion. There is no easy way to interpret this empirical finding using the two prominent narratives of the crisis. Neither the subprime view nor the expectations view provide a clear explanation of why the predictive association between credit risk and economic activity is stronger in more heterogeneous areas. In contrast, this pattern is consistent with risk-averse households deciding rationally about their consumption and investment in response to changes in the amount of aggregate economic risks.

I use these empirical observations to motivate a model in which households are differently exposed to the time-variation in the amount and price of risk in the economy. The time-varying aggregate risk in the economy is driven by the probability of rare economic disasters as in Wachter (2013). In the model, long-lived households earn labor income and are able to use that to invest in a government bond or a housing asset. Due to the incomplete markets assumption of the model, households are unable to insure against their labor income risk. Furthermore, their consumption and investment decisions are made rationally, maximizing their continuation utility. In the model, households' access to credit is only possible by using their houses as collateral, limited by debt-to-income and loan-to-value constraints. This generates a link between housing booms and expansions of credit in the economy.

Next, I calibrate the model and show that under this simple setup, the dynamics of households' leverage reflects the time-variation in risk premia. Households invest less in the government bond and more in the housing asset when the amount of aggregate risk is low. An increase in the total investment of households in the housing asset is accompanied by an increase in the total amount of household debt in the economy. Hence, the amount and risk of household credit are also directly related to the aggregate risk. When the disaster probability increases, the price of the housing asset declines, and households increase their investment in the government bond for precautionary reasons. Unable to meet their financial obligations, some of the households default on their loans. Booms and busts in housing and credit markets are both driven by time-variation in the aggregate amount of risk; none of them is caused by the other as a result of misaligned incentives or behavioral biases.

The model implies that consumption and investment of households with more labor income exposure to rare disasters react more strongly in a non-linear fashion to changes in the systematic risk. Hence, even if average exposure to economic disasters is identical across different regions, comovements between credit risk and aggregate economic variables such as income or consumption are stronger for more heterogeneous regions. A simulation exercise reveals that the forecasting association between the credit risk and macroeconomic aggregates is stronger in regions where heterogeneity among households is larger, consistent with the empirical evidence.

The analysis in this paper does not rule out the possibility of an inefficient credit boom at the beginning of the century, as an outcome of deregulation and financial innovations, resulting in institutional and agency problems. Nor does it reject the possibility that behavioral biases and extrapolative expectations intensified the housing crisis. Rather, this paper fits as a complement to the two prominent narratives about the origins of the financial crisis. The empirical evidence and the theoretical model in the paper provide grounds for considering time-variation in risk premia and household heterogeneity as more significant contributors to the formation and amplification of the boom and bust in both housing and credit markets.

Related Literature

There is vast theoretical and empirical literature that studies the credit fluctuations and their contribution to the dynamics of asset prices. In a seminal theoretical study, Kiyotaki and Moore (1997) emphasized the role that durable assets such as housing play as collateral, in transmitting and exacerbating income and technology shocks.[3] After the Great Recession, many studies aimed to shed light on how changes in the household balance sheet before the financial crisis affected aggregate measures of economic activity during and after the crisis. In a series of influential papers, Mian and Sufi (2009, 2011, 2014) among others showed that an increase in home equity-based borrowing and household leverage prior to the Great Recession contributed to the amplification of the crisis in different regions of the country. The critical role of credit expansion in various regions during the boom period in the housing sector is the central theme in this literature. (See for example Di Maggio and Kermani, 2017; Justiniano, Primiceri, and Tambalotti, 2019).

Several papers in the literature challenge the subprime view by putting more weight on the role of expectations about house prices. Adelino, Schoar, and Severino (2016) exhibit that defaults and delinquencies among high-FICO borrowers increased during the financial crisis. Kaplan, Mitman, and Violante (2017) argue that changes in expectations were a more important contributor to the growth of house prices compared to the credit conditions. Rather than being driven by speculation, variation in tastes of the investors, or extrapolative expectations[4], house price movements in my model are a product of changes in the amount of aggregate risk in the economy.

The role of household heterogeneity is another important area of research. Notably, Mian, Rao, and Sufi (2013) establish the importance of heterogeneity in the balance sheet

---

[3] Also see Bernanke, Gertler, and Gilchrist (1999), Iacoviello (2005), Mendoza (2010), and Eggertsson and Krugman (2012).

[4] See Foote, Gerardi, and Willen (2012), Shiller (2014), and Glaeser and Nathanson (2017).

of households as a determinant of consumption response across different regions. Guerrieri and Lorenzoni (2017) study the effect of tighter credit constraints on output and interest rate dynamics in a heterogeneous incomplete market model. More recently, Beraja, Fuster, Hurst, and Vavra (2019) show that the aggregate effects of the monetary policy on regional economies are related to the time-variation in the distribution of housing equity among households. In this study, I focus on the role of heterogeneity in the exposure of households' labor income to systematic shocks. The close tie between this measure and credit risk motivates modeling decisions of households as rational investment decisions under an incomplete market condition.

A different strand of empirical studies (For example see Gilchrist and Zakrajšek, 2012; Jermann and Quadrini, 2012; Greenwood and Hanson, 2013; Baron and Xiong, 2017) investigates how the quantity and riskiness of corporations' balance sheets are related to the aggregate economic activity and bond excess returns. Recently, Gomes, Grotteria, and Wachter (2019) studies the role of time-variation and heterogeneity in investment opportunities of firms for explaining these observed patterns in a complete markets model. In this paper, I focus on a similar rational explanation to understand the strong association between mortgages' risk of default and the regional economic activity.

The rest of the paper is organized as follows: In Section 1.2, I present data and empirical evidence, then discuss the role of time-varying risk premia and household heterogeneity in generating the observed patterns. Motivated by these results, Section 1.3 introduces a model that features the time-varying risk of rare disasters as the source of risk premia, and households with differential exposures to the realization of disasters. In Section 1.4, I calibrate the model and compare its implications with the empirical results. Section 1.5 concludes.

1.2 Empirical Analysis

In this section, I present and discuss the empirical findings of the paper. Section A discusses the data sources and the construction of different variables used in the empirical analysis. Next, in Section 1.2.2, I introduce the credit dispersion measure implied by mortgage data. In Section 1.2.3, I show that credit dispersion forecasts the state-level growth rate of GDP per capita and employment at various time horizons. Furthermore, I present evidence of the existence of this forecasting ability at the MSA level. Finally, Section 1.2.4 explores the role of heterogeneity in labor income exposures as a driving force for the observed patterns in the data.

1.2.1 Data

The primary sources of mortgage data in this paper are Fannie Mae and Freddie Mac Single-Family Loan-Level Datasets. The two organizations are furnishing these datasets for public use at the direction of the Federal Housing and Finance Agency (FHFA). The datasets contain origination data of the mortgages processed by Fannie Mae and Freddie Mac such as the amount and date of loan origination, the interest rate charged, length of the contracts, credit scores of the recipients, loan to value and debt to income ratios, credit insurance products purchased with the mortgage, etc. It also includes monthly loan performance data, actual loss data, legal costs, etc. which are not used in this study. The data is available from January 1999 to October 2018. The duration of these fixed-rate contracts is 15-35 years. Thus, the dataset does not include adjustable rate mortgages, initial interest mortgages, or other contracts with step rates.

The geographical information in this dataset bears significant importance in my analysis. The data set provides information about the Metropolitan Statistical Area (MSA) of the house. The data also includes a column with the first three digits of the house zip code. I cannot use the information content of this column to assign mortgages to specific counties or MSAs; the 3-digit zip codes can only specify states uniquely. As a result, the most accurate geographical information in the data is the MSA location of the house.

I clean the data by dropping all the mortgages for which the geographical information is missing. Also, I drop the mortgages that are insured using credit insurance products allowing me to rely on mortgages that are only backed by the house value. This choice is also necessary as I rely on the Merton model for measuring the credit risk of the mortgages. What remains after combining the two sources is the data for near 50 million mortgages over the period 1990-2018. The number of observations is lower in the first few months and the last two years of the sample, most likely due to reporting issues.

I utilize Zillow county and MSA level house price data to follow the level and growth rate of house prices. In some cases, the MSA contains more than one county in which case I use a simple average of the growth rates across counties. This data is available from 1994 for most of the counties.

The Bureau of Economic Analysis furnishes the GDP data at various geographical levels. The quarterly state-level GDP data is available from 2005. I use Census population data to turn that into GDP per capita. Also, from December 2019, the BEA started publishing annual county-level GDP data going back to 2001.

For employment figures, I use the data from The Bureau of Labor Statistics. In this paper, I mainly rely on the Quarterly Census of Wages and Employment (QCWE) data that provides employment data at the county level. At each county, the monthly employment and total quarterly wages are classified by the industry using the North American Industry Classification System (NAICS) at 3-6 level digits precision.

1.2.2 Dispersion in Mortgage EDFs

Merton (1974) computed the value of a firm's equity by modeling it as a call option on the asset value with a strike price equal to the debt value. Under this simple modeling framework, the probability of default is calculated as a function of the expected average growth rate of the companies value, the volatility of the growth rate, and the leverage ratio. This probability is called the expected default frequency (EDF) of the company.

Borrowing this framework, I consider a similar approach to modeling the home-equity value of a household and its expected default. This means that the total value of a household's assets or at least the portion that matters for the default decision can be approximated by the house price. Moreover, the household equity is the levered claim on this asset.[5] EDF is calculated as:

$$\text{EDF}_{it} = \mathcal{N}\left(\frac{-\log\frac{V_{it}}{B_{it}} - \left(\mu_{V_{it}} - \frac{1}{2}\sigma^2_{V_{it}}\right)}{\sigma_{V_{it}}}\right). \tag{1.1}$$

In the above V_{it} is the market value of the household i's house. B_{it} is the face value of the mortgage loan. Also, $\mu_{V_{it}}$ and $\sigma_{V_{it}}$ are mean and standard deviation of the growth rate of V_{it}. $\mathcal{N}$ indicates the standard normal cumulative density function.

As described in the previous section, the loan-level data from Fannie Mae and Freddie Mac provides the value of the house at the time of the origination of the mortgage. Also, using the LTV ratio, I can find out B_{it}.

The values of $\mu_{V_{it}}$ and $\sigma_{V_{it}}$ for each house, are not available in the data. I assume that the values of all houses in a given geographical area are perfectly correlated. This assumption allows me to utilize the Zillow regional house price data to compute changes in V_{it}. I

[5]For a majority of households, housing is the most important item in their wealth portfolio (For data about the US and other advanced economies see Jordà, Schularick, and Taylor, 2019). This assumption is also justified if the default decision is unrelated to the value of other household assets.

calculate $\mu_{V_{it}}$ as the growth rate over the last 12 months[6] and use a simple GARCH model to estimate $\sigma_{V_{it}}$. An alternative method is to compute the sum of squared returns over the last 12 months as a benchmark of $\sigma_{V_{it}}$. Both methods produce consistent outcomes.

Consider the pool of all the borrowers in region j at time t and call it $\mathbb{B}_{j,t}$. I define credit dispersion as the difference between the current average EDF of today's borrowers compared to the current average EDF of last year's borrowers:

$$Dispersion_{jt} = \frac{1}{n(\mathbb{B}_{j,t-1})} \sum_{i \in \mathbb{B}_{j,t-1}} EDF_{it} - \frac{1}{n(\mathbb{B}_{j,t})} \sum_{i \in \mathbb{B}_{j,t}} EDF_{it}, \tag{1.2}$$

where $n(.)$ denotes the number of households in the set.

The above definition is inspired by the definition of credit dispersion in the corporate bond market literature such as Greenwood and Hanson (2013) and Gomes, Grotteria, and Wachter (2019). Instead of comparing today's borrowers to last year borrowers, the main focus in the corporate bond market is on the average EDF of issuers compared to repayers. The variable captures how financing decisions of firms with different characteristics are related to the business cycle status or pricing of credit risk.

Due to the nature of housing debt, I cannot follow the same definition here. However, in the same spirit, the above definition allows me to compare households with regards to the timing of their decision to increase their leverage and invest in the housing market. Last year's borrowers are paying off their mortgage and hence decreasing their leverage, similar to repaying companies. In contrast, today's borrowers are increasing their leverage, analogous to those companies that are issuing debt in the bond market.

The major contrast between the set of "Last Year's Borrowers" in this paper and the set

[6]In the corporate literature, it is conventional to use average returns over the past 12 months to calculate expected returns. Given that housing returns show stronger persistence in the data (with an annual AC coefficient of 0.7) compared to the equity returns, the use of this measure as a proxy for expected returns is more justified in the context of this paper.

of “Repayers” in the bond credit risk literature is that the former does not include the entire pool of households that are paying off their debt. Instead, the focus is on the borrowers that are repaying their loan since last year, as they can be considered marginal entrants into the set of all households that are de-levering. Also, given that many households decide to refinance after a few years of mortgage origination, this choice makes sure that the set only contains households that are decreasing their leverage.

The first panel in Figure 1.1 shows the average values of EDF for today’s borrowers versus last year’s borrowers, aggregated by averaging across all the MSAs in the country. Panel B shows the $Dispersion$ at the national level. The average EDF of last year borrowers is much more volatile and is the dominant driver of $Dispersion$. This pattern resembles Gomes, Grotteria, and Wachter (2019) findings about dispersion in corporation’ credit risk.

There is a clear countercyclical pattern apparent in the figures. Both variables in the first panel start increasing prior to the financial crisis and dampen afterward. Since the increase in the EDF of last year’s borrowers is larger, $Dispersion$ also displays countercyclical dynamics.

Table 1.1 reports pairwise correlations between these three variables and measures of aggregate risk in the asset pricing literature. These measures include annual price-dividend ratio, CAY (Lettau and Ludvigson, 2001) , GZ credit spread (Gilchrist and Zakrajšek, 2012), and variance risk premium (Bollerslev, Tauchen, and Zhou, 2009). All variables are in monthly frequency except for CAY that is available at a quarterly frequency. I use a monthly version of CAY in this analysis by assuming it remains constant between quarterly updates. These results show that there is indeed a strong countercyclical pattern in all three variables.

How can we interpret these patterns? One first explanation is that in line with the “subprime view”, these figures confirm the deterioration of borrowers’ credit quality over the business cycle. However, this explanation is not consistent with the fact that the households in this paper are of relatively high quality. To be more precise, the data provides a direct measure of households' creditworthiness over time. The first panel in Figure 1.2 shows 5, 50, and 95 percentiles of FICO score distribution in the sample. Panel B provides the standard deviation of FICO scores. There is a small increase of around $\sim$30-40 points around 2009, partially reversed around 2014. Except for that, there is not much change in the average scores of these borrowers. Also, there is a reduction of $\sim$10 points in the standard deviation of scores. These are in contrast with bold patterns in Figure 1.1 that start well before the crisis.

As is clear from definitions in equations (1.1) and (1.2), the geographical EDF and $Dispersion$ measures are driven by changes in mean and variance of regional housing returns. Figure 1.2 is consistent with Adelino, Schoar, and Severino (2018) findings regarding the prevalence of defaults among middle income and high FICO borrowers. To summarize, at the national level $Dispersion$ seems to be closely related to measures of aggregate risk. Also, it does not seem to be driven by the quality of borrowers. Instead, by construction, it reflects changing growth expectations in the housing market.

1.2.3 Forecasting Macroeconomic Variables

In this section, I present evidence about the information content of the credit dispersion measure for economic activity. One possible weakness of this analysis is that the available time-series for our data is not very long. The 1999-2018 period contains only one major economic recession. Also, the quality of data at the beginning and end of the sample is relatively lower. On the plus side, the mortgage data provides regional variations that could be exploited to understand the forecasting power of credit dispersion. In order to do so, I aggregate the cross-section of available EDFs, up to specific geographical units (i.e. state level, MSA level), then use that to investigate the predictive power in panel regressions.

State Level Evidence

I aggregate loan-level data up to the state level and estimate regressions of the following type at different horizons:

$$\frac{1}{h}\Delta gdp_{i,t\to t+h} = \alpha_i + \alpha_t + \beta Dispersion_{i,t} + \xi\Delta gdp_{i,t-1\to t} + \epsilon_{i,t\to t+h}. \tag{1.3}$$

The regression is run with quarterly data and contains state and time fixed effects.

Table 1.2 reports the estimation results. Credit dispersion forecasts the growth rate of GDP per capita at different horizons. This finding is both economically and statistically significant. The coefficients suggest that a one standard deviation increase in the regional credit dispersion forecasts a one percent reduction in the growth rate of state-level GDP per capita. The statistical significance increases with the horizon. However, one should be careful about the possible role of persistence in the $Dispersion_{i,t}$ variable, and overlapping dependent variable. Adding time fixed effects reduces some of the observed predictive power. Nonetheless, the association remains highly significant; hinting to the fact that aggregate benchmarks of macroeconomic and financial risk fall short of explaining this strong forecasting relationship.

Table 1.3 outlines the results of similar predictive regressions at monthly frequency, where the dependent variable is employment growth:

$$\frac{1}{h}\Delta emp_{i,t\to t+h} = \alpha_i + \alpha_t + \beta Dispersion_{i,t} + \xi\Delta emp_{i,t-1\to t} + \epsilon_{i,t\to t+h}. \tag{1.4}$$

The findings are similar to the Table 1.2 results. Credit dispersion predicts employment growth at various horizons for up to a year. A one standard deviation increase in the dispersion predicts almost a 0.5% reduction in the growth rate of state-level employment in the following year. The predictive association remains significant even after adding time fixed effects. This finding suggests the cross-sectional association is not solely driven by macroeconomic trends.

MSA Level Evidence

The available data allow running similar predictive regressions at the MSA level. The dependent variables in Table 1.4 and 1.5 are GDP and employment growth respectively. Estimation results reaffirm the state-level evidence: credit dispersion strongly forecasts economic activity, this predictive power is not explained by aggregate measures of risk and increases with the horizon. [7]

1.2.4 Role of Heterogeneity

The evidence presented in the last section demonstrates the credit dispersion implied by mortgages as a state variable of the economy that comoves with benchmarks of aggregate risk. Also, credit dispersion predicts measures of macroeconomic activity such as GDP growth and changes in employment. Motivated by these findings, it is natural to ask what economic forces are behind these empirical patterns. In this section, I will explore a possible link between household heterogeneity and the observed predictive patterns in the data. First, I introduce a simple measure of heterogeneity in labor income exposures using the Quarterly Census of Employment and Wages data. Then, I revisit the predictive evidence

[7] It is worth mentioning that the results presented in this section and the previous one, contribute to the so-called "housing is the business cycle" view that identifies a strong relationship between investment in the housing sector and the business cycle. (For example see Leamer, 2007). As an example, these results are in contrast with the findings of Ghent and Owyang (2010). They report that among the MSAs of 51 US cities, house price declines are not followed by declines in the growth rate of employment. This section results show how the credit dispersion measure, links developments in the housing and mortgage markets to the employment growth at the MSA level.

of the last chapter, focusing on the role of heterogeneity in labor income exposures to economy-wide risks.

Measuring Heterogeneity in Labor Income Exposures

As explained in more detail in the data section, the Quarterly Census of Employment and Wages provides county-level employment data for different NAICS codes for each county. I assume that within a county, the labor income exposures of all the employees in a 6-digit NAICS industry are equal to each other. In order to gauge the exposure, I estimate the following regression for each industry-county:

$$\Delta emp_{i,t} = \alpha_i + \phi_i \Delta emp_{s,t} + \text{Seasonal Dummies}_i + \epsilon_{i,t}. \tag{1.5}$$

In the above regression, $\Delta emp_{s,t}$ measures the log change of employment in state s on month t. $\Delta emp_{i,t}$ identifies the same quantity for a give industry-county i. The regressions include month dummies that capture possible seasonalities in the employment data. The coefficient ϕ_i measures the exposure of employment in a certain industry-county, to the economy-wide employment shocks within a state.

If the data contains employment for all industry-counties in a state, then by definition, the average exposure of industry-counties in a state is one. However, since this is not the case, the average exposure can deviate from one. I compute the average exposure weighted by the (average) level of employment in each industry-county as follows:

$$\bar{\phi}_{s,t} = \sum_{i \in s} \frac{emp_{i,t}}{emp_{s,t}} \phi_i \quad , \quad \bar{\phi}_s = \sum_{i \in s} \frac{\overline{emp}_{i,t}}{\overline{emp}_{s,t}} \phi_i.$$

Next, I measure heterogeneity of labor income exposures for state s, using the standard deviation of exposure coefficients:

$$\sigma_{s,t}^{Expo} = \sqrt{\sum_{i \in s} \frac{emp_{i,t}}{emp_{s,t}} \left(\phi_i - \bar{\phi}_{s,t}\right)^2} \quad , \quad \sigma_s^{Expo} = \sqrt{\sum_{i \in s} \frac{\overline{emp}_{i,t}}{\overline{emp}_{s,t}} \left(\phi_i - \bar{\phi}_s\right)^2}. \tag{1.6}$$

Figure 1.3 depicts how heterogeneity of labor income varies among different states. A darker color is associated with a lower level of σ_s^{Expo} in the state. Also, in Figure 1.4 the time series variation in state-wide heterogeneity is presented.

There is a possibility of measurement errors in these calculations. First, the assumption that all the employees in a given industry-county have identical exposures might not always be realistic. The characteristics of different jobs in the same sector of the economy can be very diverse. Nevertheless, this is the most reasonable assumption given the limitations of the available data on the labor income or employment situation of individual households. Second, the employment data on numerous small industries is not available in all the years and months. This issue affects estimations in regression (1.5). To tackle this issue, I use only the data of industry-counties for which I have at least 15 years of data. While this helps reduce this measurement issue problem, it does not eliminate it. Finally, measuring variance or standard deviation is more sensitive to these issues compared to the average or median.

I use both the time-varying and time-invariant versions of this benchmark of heterogeneity of labor income exposures in predictive regressions. Although the existence of these measurement errors might introduce bias in the estimation results, it is natural to think that these problems attenuate the regression coefficients toward zero and hence works against finding a significant role for heterogeneity.

Credit Dispersion and Heterogeneity of Exposures

Table 1.6 presents the estimation results of the following regression:

$$\begin{aligned} \Delta gdp_{i,t\to t+k} &= \alpha_i + \beta Dispersion_{i,t} + \gamma Dispersion_{i,t} \times \sigma_{i,t}^{Expo} + \delta \sigma_{i,t}^{Expo} \\ &\quad + \xi \Delta gdp_{i,t-1\to t} + \epsilon_{i,t\to t+k}. \end{aligned} \tag{1.7}$$

The setting allows me to investigate the forecasting association between credit dispersion and GDP growth in the presence of the time-varying measure of the heterogeneity of exposures $\sigma_{i,t}^{Expo}$. On the left panel, I shut down the interaction term. On the right panel, both the regression coefficient for exposure heterogeneity δ, and the interaction coefficient γ are estimated.

The left panel results show that credit dispersion retains its forecasting power in the presence of $\sigma_{i,t}^{Expo}$ at all horizons. However, on the right panel, the interaction term subsumes all the predicting power of credit dispersion. Interestingly the sign of δ is consistent with theoretical models that rely on time-varying heterogeneity to explain asset prices; an increase in heterogeneity forecasts lower economic growth in the future.

To further investigate the relationship, in Table 1.7, I consider the time-invariant measure of heterogeneity of exposures σ_{i}^{Expo}. Note that the δ coefficient in equation (1.7) is absorbed by the state fixed effects. Similar results are achieved, with the notable exception that the credit dispersion coefficient does not lose all of its statistical significance to the interaction term.

Tables 1.8 and 1.9 present the analogous regression results with employment growth as the dependent variable. The coefficient γ of the interaction term between credit dispersion and heterogeneity of exposures absorbs almost all of the negative forecasting association.

This time the coefficient δ remains insignificant at all horizons.

The results brought up in this section indicate that the negative forecasting association between credit dispersion and measures of economic activity is more pronounced in states where there exists higher heterogeneity of labor income exposures. These results motivate recognizing a more significant role for heterogeneity of exposures in explaining the forecasting power of this measure of credit risk.

How do these results add to the existing evidence about the credit cycles and their relationship with the business cycles? There does not seem to be a straightforward explanation for these findings under the "subprime view" or "expectations view" of the financial crisis. Consider the first narrative. There is no intuitive way to connect increased lending to lower credit-worthy borrowers and the heterogeneity of exposures among households. Similarly, it is not easy to relate the notion of extrapolative expectations about house prices to higher heterogeneity. Why should households in more heterogeneous states rely more on extrapolation for their investment decisions?

In contrast, the most salient feature of rational decision-making is its emphasis on notions of covariation with or exposure to sources of systematic risk. When aggregate risk and its price change, investment and consumption decisions of households and firms change. Households and firms that are more exposed to systematic risk are expected to react more strongly to variation in the amount of aggregate risk. In the next section, I formalize how these patterns arise as a natural consequence of rational decisions of heterogeneous agents to invest in housing wealth in an incomplete market framework.

1.3 Model

1.3.1 Aggregate Economy and Stochastic Discount Factor

I assume that aggregate fluctuations are driven by a time-varying probability that the economy enters into a disaster. I assume that the disaster event is a Bernoulli random variable x_t that takes the value of 1 at time t with probability p_t. The probability of rare disasters follows a square-root process in discrete-time:

$$p_{t+1} = (1-\rho_p)\bar{p} + \rho_p p_t + \sigma_p \sqrt{p_t}\epsilon_{p,t+1}, \quad \text{where } \epsilon_{p,t+1} \overset{i.i.d.}{\sim} N(0,1). \tag{1.8}$$

In the above, $\bar{p}$ is the long-run mean of the disaster probability. The parameters ρ_p and σ_p determine the persistence and volatility of the process, respectively.

Furthermore, I assume that in the event of a rare disaster, aggregate consumption (or output) drops by the amount of ξ_{t+1}. This random variable is distributed according to an independent time-invariant distribution with moment generating function Φ_ξ.

Based on the definition of p_t and following the disaster risk literature, I specify the following stochastic discount factor:

$$\log M_{t+1} = -r_0 - r_p p_t + \sigma_{my}\epsilon_{y,t+1} + \sigma_{mp}\sqrt{p}\epsilon_{p,t+1} + \xi_{m,t+1}x_{t+1}, \tag{1.9}$$

where $\epsilon_{y,t+1} \overset{i.i.d.}{\sim} N(0,1)$. All the real and financial assets in the economy are priced using the above stochastic discount factor.[8] The drift of $\log M_{t+1}$, is an affine function of p_t. Three innovations might impact the SDF through time. First, the SDF is affected

[8]Due to the incomplete market assumption of the model, it is not possible in our settings to derive the SDF from the Epstein-Zin preferences of a representative investor over total consumption or output. However, the adopted specification for the SDF is almost identical to the solution of the Epstein-Zin preferences for a representative agent. This is to make sure that the results of the model do not depend on discount rates being affected by heterogeneity and incomplete markets assumptions.

by the normal shock $\epsilon_{y,t+1}$ that represents the fluctuations in the aggregate consumption or output at time t. Second, the $\epsilon_{p,t+1}$ shock that relates the SDF to the probability of a disaster happening. Third, in the event of a disaster, the SDF jumps upward with the magnitude $\xi_{m,t+1}$. I assume that $\xi_{m,t+1} = -\gamma\xi_{t+1}$, where γ is the relative risk aversion. This assumption relates the magnitude of the jump in the discount rate with the magnitude of the macroeconomic disaster, consistent with CRRA or Epstein-Zin preferences for real consumption.

Based on the above assumptions, the risk-free rate is as follows:

$$\log R_t^f \equiv -\log\left(\mathbb{E}_t\left[M_{t+1}\right]\right) = r_0 + r_p p_t - \frac{1}{2}\sigma_{my}^2 - \frac{1}{2}\sigma_{mp}^2 p_t - \log\left(1 - p_t + \Phi_\xi(-\gamma)p_t\right) \tag{1.10}$$

Following Barro (2006) and Wachter (2013), I allow for the possibility of a partial default by the government on its debt. Conditional on a disaster happening, the government defaults partially on its debt with constant probability q and investors lose an amount equal to the size of the disaster. Denoting the event of government default by L we have:

$$\log R_{t+1}^b = \mu_t^b + L_{t+1}\xi_{t+1}x_{t+1}. \tag{1.11}$$

As a result of this assumption, as proven in the appendix, the face value of government debt is given by:

$$\mu_t^b = \log R_t^f - \log\left[1 - q + (\Phi_\xi(1-\gamma) - \Phi_\xi(\gamma))qp_t\right], \tag{1.12}$$

and the expected log-return on government debt is:

$$\mathbb{E}_t\left[\log R_{t+1}^b\right] = \mu_t^b + \Phi_\xi'(0)qp_t \tag{1.13}$$

1.3.2 Housing Sector

I assume that there is a housing asset available for investment that provides households with housing services (i.e. housing dividends or rents). The dynamics of housing dividends in region j is specified as:

$$\log s^j_{t+1} = \log s^j_t + \mu^j_s + \sigma^j_s \epsilon^j_{s,t+1} + \xi_{s,t+1} x_{t+1}. \tag{1.14}$$

In the above specification, the growth rate and volatility of these dividends are denoted by μ^j_s and σ^j_s respectively. In the event of a disaster, housing revenues encounter a decline equal to $\xi_{s,t+1}$. I assume that the size of the decline in revenues of the housing sector is related to the size of the disaster in the aggregate economy by the equation $\xi_{s,t+1} = \phi_s \xi_{t+1}$. Furthermore, the normal shock to the housing sector revenues is correlated with the shock to aggregate income or consumption, corr$(\epsilon_{y,t+1}, \epsilon^j_{s,t+1}) = \rho^j_s$.

The price of a unit of housing asset in region j is determined by solving the Euler equation for the price-dividend ratio of the housing asset:

$$\mathbb{E}_t \left[M_{t+1} R^j_{h,t+1} \right] = 1. \tag{1.15}$$

I solve for $pd^j_{h,t}$ as a function of the model's state variable numerically.

1.3.3 Households

The regional economies are populated with long-lived households. Household i is endowed with labor income y^i_t which evolves as follows:

$$\log y^i_{t+1} = \log y^i_t + \mu^i_y + \sigma^i_y \epsilon^i_{y,t+1} + \xi^i_{y,t+1} x_{t+1}. \tag{1.16}$$

Household labor income grows with an average rate of μ_y^i and volatility σ_y^i. The normal shocks to household income are correlated with the aggregate shocks (corr$(\epsilon_{y,t+1}, \epsilon_{y,t+1}^i) = \rho_y^i$). The size of the shock to household labor income in the event of a disaster is $\xi_{y,t+1}^i = \phi_y^i \xi_{t+1}$. This parameter is the main source of heterogeneity in the model. Households are differently exposed to rare disasters as their ϕ_y^i can be different.

In order to simplify the calibration of these parameters throughout the paper, I add a few assumptions. First, instead of calibrating μ_y^i and σ_y^i for households separately, I assume that they are related to aggregate mean and volatility of income growth in the economy. More specifically, I assume that the average growth rate of income for individuals is determined as follows:

$$\mu_y^i = \mu_y + \log \mathbb{E}\left[e^{\xi_{t+1}x_{t+1}}\right] - \log \mathbb{E}\left[e^{\phi_y^i \xi_{t+1}x_{t+1}}\right]. \tag{1.17}$$

This assumption makes sure that the income share of highly exposed households does not shrink over the long-run. Hence the mean growth rate of income remains at μ_y.

Second, I assume that through each region and in the whole economy, the average ϕ_y^i adds up to 1 ($\int_i \phi_y^i = 1$). This makes sure that regions are not different when it comes to their average exposure to rare disasters. However, regions could still be different when it comes to cross-sectional variation in ϕ_y^i.

Third, I assume that $\sigma_y^i = \frac{1}{\rho_y^i}\sigma_y$ and that σ_y^i is equal for all the households in all the regions. While decreasing the number of free parameters, these assumptions make sure that the only source of heterogeneity in the model is the difference in labor income exposures of households to disastrous shocks.

Households can only invest in the government bond and the housing sector. When they decide to invest in the housing sector, the amount of debt that they can get compared to their labor income is restricted by the debt-to-income (DTI) ratio. If they decide to invest in the housing sector, there is a maximum level of leverage they can take which is determined by the LTV ratio. I assume that when investing, households use the maximum possible leverage and house price that is determined by these two ratios and their level of labor income.

To reduce the complexity of the model, I model household debt as a perpetuity contract similar to Beraja, Fuster, Hurst, and Vavra (2019). The required rate of return on household debt is determined by:

$$\log R_t^d = \log \mathbb{E}_t \left[\log R_{t+1}^b\right] + r_{debt} p_t \quad \text{where } r_{debt} > 0. \tag{1.18}$$

Based on the above equation, lenders demand a premium for the risk of defaults by households that is increasing in the probability of rare disasters.

Household derive utility from real consumption according to preferences specified as in Epstein and Zin (1989) and Weil (1990):

$$U_{i,t} = \left[(1-\delta) X_{i,t}^{\frac{1-\gamma}{\theta}} + \delta \mathbb{E}_t [U_{i,t+1}^{1-\gamma}]^{\frac{1}{\theta}}\right]^{\frac{\theta}{1-\gamma}}, \tag{1.19}$$

in which δ is the time discounting parameter, γ is the relative risk aversion, and $\theta = \frac{1-\gamma}{1-\frac{1}{\psi}}$ determines the preference of household toward the timing of uncertainty resolution and is calculated using γ and the IES parameter ψ. $X_{i,t}$ is a Cobb-Douglas function that aggregates the real consumption of household in housing ($s_{i,t}$) and non-housing ($c_{i,t}$) categories:

$$X_{i,t} = s_{i,t}^{\nu} c_{i,t}^{1-\nu}. \tag{1.20}$$

The parameter ν controls the optimal share of housing vs. non-housing consumptions. The specification of household preferences and consumption in this paper is identical to Chen, Michaux, and Roussanov (2020).

Households optimally decide about the level of their consumption and investment in a house or government bond in the presence of budget and debt constraints, trying to maximize their utility. Households are either in the set of homeowners, renters or defaulted households. Homeowners can decide to sell their house or continue owning it. If they cannot meet their financial obligation they default on their mortgage and forfeit all their wealth in government bonds.

The recourse laws about housing debt differ across different states. Many states consider mortgages a non-recourse debt. However, in this paper, I make a simplifying assumption and consider all mortgages recourse meaning that households savings in the government bond will be ceased upon default. Also, when households default, they are excluded from taking a mortgage and purchasing a house for a random period of time. After the default, a household will be eligible again to return to the housing market and become a homeowner with probability ω, independent of other shocks in the model. After being eligible again, households can decide to remain a renter or invest in the housing asset.

Household decisions are determined using Bellman's equation. Appendix B provides details of the model solution equations. Note that the model is solved under partial equilibrium. There is no equilibrium condition relating the sum of households' income or consumption to the stochastic discount factor.

1.4 Model Implications

I calibrate the model and use simulation exercises to study the behavior of different macroeconomic variables, their relationship with time-varying aggregate risk, and the role of heterogeneity. In the next subsection, I present a calibration of the model parameters. Next, I describe the simulation procedure and present some of the results that shed light on the effect of heterogeneity. Also, the impulse response functions of different model variables to changes in disaster risk are analyzed.

1.4.1 Calibration

Table 1.10 reports how model parameters are calibrated at the quarterly frequency. There are six categories of parameters that I need to calibrate in this model.

First, the calibration of preference parameters is relatively standard. Here, I use an annual rate of time preference of 1.2% as in Wachter (2013). This means that δ at quarterly frequency would be 0.997. The relative risk aversion parameter is set to be 4, in line with the equity premium puzzle literature. I set EIS to be 1.5, consistent with preferences for early resolution of uncertainty. The only remaining parameter in the specification of preferences is ν which is set to be 0.15, close to the estimated value of 0.134 in Chen, Michaux, and Roussanov (2020).

For the disaster risk parameters, I choose numbers in line with parameter values used in the literature. Disaster size distribution is according to Barro and Ursúa (2008) data. Their results suggest average rare disaster probabilities of 2.87% and 3.69% respectively for OECD and all countries in their sample. I set $\bar{p}$ to be 3.00% annually (0.75% quarterly). I set ρ_p and σ_p to be 0.98 and 1.75% in quarterly frequency. These numbers are consistent with an annual mean reversion and volatility parameter of 8% and 7% respectively, consistent with Wachter (2013). Finally, the parameter q, the probability of partial government default in the event of a rare disaster is 40%.

Next, I consider the SDF parameters. As explained in the model section, while the SDF in my model is not directly obtained from Epstein-Zin preferences of a representative investor, I try to be consistent with such an outcome while trying to match risk-free rate patterns in the data. I set r_0 and r_p to be 4% and 3 respectively. σ_{my} is calibrated to be -0.02 consistent with aggregate consumption volatility of 0.5% and a relative risk aversion of 4 and σ_{mp} is calibrated to be 1.2.

Now, I turn to the calibration of housing dividends parameters. I use similar parameters across different regions to make sure that the simulation results are only driven by heterogeneity among investors. I assume a growth rate equal to aggregate consumption or income; 2% per year or 0.5% quarterly. I also set the volatility of the housing asset returns to 1% quarterly. I calibrate ϕ_s, the parameter governing the exposure to rare events, to be 3. This choice generates reasonable housing excess return and volatility with a price-dividend ratio consistent with the data from Zillow and long-run housing return data from Jordà, Schularick, and Taylor (2019).

I already explained how I set the average growth rate of labor income. I set σ_y^i to be 20% annually or 5% quarterly consistent with the numbers reported in Gorbachev (2011) about US household income volatility. To be consistent with the aggregate income volatility of 2% I set ρ_y^i of 10% for all households. I allow 3 different values for exposure of labor income to disaster risk: 0.5, 1, and 1.5.

Finally, the parameters associated with household debt are set according to the literature. The values of DTI, LTV and ω are determined to be 14, 0.8 and 0.15 in accordance with Chen, Michaux, and Roussanov (2020). I also use and r_{debt} of 1 to make sure of an average premium of 3% annually for mortgages compared to the government bond.

1.4.2 Simulation Results

Simulation Procedure

I simulate 20-year sample paths of the model after a sufficient burn-in period at a quarterly frequency for 10,000 times. The simulation outcome is then compared with the estimation results presented in Section 1.2. The choice of 20 years represents the length of the available US data in the empirical analysis. Each simulation path contains an economy with three regions. There are 500 households in each region. The probability of rare disaster and house price dynamics are assumed to be similar across the three regions.

However, these three regions are different in terms of heterogeneity in exposures of their households to rare disasters. All the households in region 1, have the same exposure (ϕ_y^i) of 1. Hence, there is no heterogeneity in this region. In region 2, 25% of households have exposure parameter equal to 0.5, the parameter for 50% is equal to 1, and 25% have ϕ_y^i of 1.5. Lastly, in region 3, half of the households are exposed with a ϕ_y^i of 0.5, and the rest are exposed with the parameter equal to 1.5. Due to these simple assumptions, while regions are similar in terms of their average exposure to disaster risk, they are different when it comes to heterogeneity in the exposure of households to rare disasters. Standard deviations of ϕ_y^i, $\sigma(\phi_y^i)$ in these three regions are 0%, 35%, and 50% respectively.

Individuals make consumption and investment decision in the simulation path. I compute EDF for each mortgage following a similar procedure as in the data. Then I aggregate the data at the regional level to have a time-series of dispersion in each region. Also, I define GDP to be the sum of consumption and net investment of households in housing or government bond. Consequently, I can use the consumption and investment decisions of households to create regional macroeconomic time-series.

Impulse Response of Macroeconomic Variables to Disaster Risk

To investigate the role of time-varying risk premia in explaining the dynamics of different model variables, it is insightful to examine the response of variables to a change in the probability of rare disasters.

First, the model is simulated for a long enough (e.g. 100 years) burn-in period for 10,000 times. Then, at time 0 in each of the paths, I conduct two separate exercises. In the first exercise, I shut down all the shocks in the model after this point in time. In the second exercise, I set $\epsilon_{p,0} = 1$ and shut down all the shocks up until the end of the simulation. Hence, the second exercise represents a one standard deviation increase in p_t. The results from these two exercises are then compared by measuring the absolute and relative deviations in macroeconomic and credit variables. These changes are averaged across all the 10,000 sample paths to calculate the impulse response functions to a one standard deviation shift in the probability of disasters.

Figure 1.5 presents the results of this procedure. Panel A in the figure depicts the average absolute change in the quarterly probability of rare disaster in percentages. As a result of the one standard deviation positive shock, the probability of a disaster happening in the next quarter increases by almost 0.13% points ($\sim$ 0.5% annually). Given that all the future shocks are shut down, after time 0 the probability is decaying toward its long-run average. Note that as a result of high persistence in the dynamics of this state variable, even after eight quarters the probability remains elevated.

Panel B plots the absolute deviation in the $Dispersion$ variable. The regional measure of credit risk increases for two quarters, then gradually decreases before reaching zero in the fourth quarter. These changes in regional $Dispersion$ are driven by the increase in expected default frequencies and a decline in house prices as a consequence of increased aggregate risk in the economy. Given that all the shocks, including shocks to house prices, are zero after time 0, there remains no difference between the expected default frequencies of today's borrowers vs last year's borrowers at quarter 4. Hence, the absolute deviation in $Dispersion$ dies out after four quarters. Being the result of subtracting two probabilities, $Dispersion$ can take values in the range $[-1, 1]$. As a result of the one standard deviation shock to p_t, the value of $Dispersion$ goes up almost 0.09 points. This is a relatively large increase given that $Dispersion$ is positive but close to zero in most of the sample.

Panels C to F show a pattern of decline in aggregate macroeconomic and credit variables. In Panel C, the consumption per capita declines more than 2%. Panel D presents a similar response in GDP per capita. These declines are persistent as it takes time for the aggregate risk in the economy to return to its average long-run level. The investment in the risky asset (housing), faces a striking decline of close to 50% in the first two quarters after the shock, before returning to around -20% decline levels. This is driven by a decline in the value function of households and an increase in risk premia, as a result of increased aggregate risk in the economy. in Panel F we see a persistent decline in the amount of per capita household debt in the economy. Given the increased amount of risk, fewer households are willing to take new loans and invest in housing as observed in Panel E. Also, as time passes, some households de-lever or face default as their savings and wages turn out to be insufficient to cover their financial obligations.

How do these results change with various levels of household heterogeneity across different regions? Figure 1.6 answers this question. As discussed earlier, the standard deviation in labor income exposures of individual households is the lowest in region 1 and the highest in region 3. The figure shows that the GDP per capita declines more in response to a one standard deviation increase in p_t in the more heterogeneous region 3 compare to region 1. The consumption and investment responses of households to increases in p_t have a convex association with their exposure to rare disasters. As a result, if there is more heterogeneity among households, the average response is larger compared to a case where all households have similar exposure.

Predictive Regressions

This setting allows me to run panel regressions similar to what I have studied using the US state-level data in Sections 1.2.3 and 1.2.4. I estimate fixed-effect regressions of growth rate in GDP, consumption, and investment in the housing asset and investigate whether dispersion in mortgages has any predictive power for these macro-aggregates in future periods. Furthermore, I explore whether the association is stronger for regions with a higher income exposure heterogeneity.

Table 1.11 reports the results of this exercise in samples that did not experience a disaster over the 20-year period of simulations. The table reports the median value for the coefficient of a fixed-effect predictive regression at different horizons up to 4 quarters. Panel A reports the results for regression analysis similar to equation 1.3. Panel B reports the results for a regression that includes a term that interacts dispersion with our measure of household heterogeneity ($\sigma(\phi_y^i)$) for different regions.

Panel A results show a strong predictive power at different horizons for GDP growth. The coefficient β is very similar in the first quarter to what is observed in the data as reported in Table 1.2. One difference is that the predictive power declines with the horizon in our model while in the data it seems to be increasing. The median value for R^2 is also relatively close to my results from the data. Both the expected growth rates of consumption and housing investment decline as dispersion increases. The highest predictive power is reported for housing investment as expected. Households reduce their investment in risky assets following an increase

Panel B is helpful to understand the role of heterogeneity. As the estimated coefficient for interaction term (γ) shows, the predictive association between GDP growth and dispersion is stronger in a more heterogeneous region. This is in general consistent with the results reported in Tables 1.6 and 1.7. The magnitude of economic activity declines with higher dispersion and the decline is more pronounced wherever heterogeneity is higher. A difference between the model and data is that after adding measures of heterogeneity in exposure, the predictive power of dispersion disappears. In the model, dispersion's predictive power remains while it is stronger for more heterogeneous regions.

Another interesting finding is that the main channel for the strong association of dispersion with the growth rate of GDP is the investment channel. While consumption growth predictability dampens rapidly with the horizon and the role of dispersion disappears, the investment in housing remains closely associated with dispersion and heterogeneity even in longer horizons.

Table 1.12 presents a parallel analysis for all of our sample paths. There is not a significant difference between samples with or without disasters when it comes to the role of dispersion and heterogeneity.

1.5 Conclusion

The relationship between booms and busts in the housing market and the credit market has been a central question in the financial economics literature after the Great Recession. Leading hypotheses have emphasized the role of institutional issues, misaligned incentives, and behavioral biases of investors to explain the causal link between the two cycles. In this paper, I revisit this question by offering an explanation based on rational decision making by households that are exposed to time-variation in aggregate economic risks.

I construct a new measure of regional credit risk by employing loan-level mortgage data capturing the dispersion in the credit quality of borrowers in the housing market. The analysis in this paper shows that dispersion comoves with benchmarks of aggregate risk. While FICO scores of borrowers do not vary considerably in the data over time, this measure strongly forecasts both GDP per capita and employment growth at the regional level. Moreover, the predictive power of dispersion is closely related to regional heterogeneity in households' exposure to systematic risks.

These empirical patterns are consistent with optimal decision-making of households exposed to time-varying risk premia. I formally show this by introducing and solving a model featuring heterogeneous households under the incomplete-markets condition. The primary source of heterogeneity in the model is differential exposure of households' labor income to rare economic disasters. The model generates a predictive association between credit dispersion and regional economic activity that is stronger for more heterogeneous regions as in the data.

Appendix

A Euler Equation and Asset Prices

Return of the Government Bond

Given the dynamics of the stochastic discount factor and the definition of return on government bond in equations (1.9) and (1.11), the Euler equation implies that:

$$\mu_t^b = -\log \mathbb{E}_t \left[\exp\left(\log M_{t+1} + L_{t+1}\xi_{t+1}x_{t+1}\right)\right]. \tag{A.1}$$

Given that the realizations of government default L_{t+1} is independent of tghe rest of the shocks in the model we have:

$$\mu_t^b = -\log\left(\mathbb{E}_t\left[M_{t+1}\right]\left(1 - q + q\frac{\mathbb{E}_t\left[\exp\left(\log M_{t+1} + \xi_{t+1}x_{t+1}\right)\right]}{\mathbb{E}_t\left[M_{t+1}\right]}\right)\right). \tag{A.2}$$

We use the definition of risk-free rate and its value under the model as in equation (1.10) to substitute the value of expectations in the above equation:

$$\mu_t^b = \log R_t^f - \log\left(1 - q + q\frac{\exp\left(-r_0 - r_p p_t + \frac{1}{2}\sigma_{my}^2 + \frac{1}{2}\sigma_{mp}^2 p_t\right) \times \left(1 - p_t + \Phi_\xi(-\gamma)p_t\right)}{\exp\left(-r_0 - r_p p_t + \frac{1}{2}\sigma_{my}^2 + \frac{1}{2}\sigma_{mp}^2 p_t\right) \times \left(1 - p_t + \Phi_\xi(1-\gamma)p_t\right)}\right),$$

and achieve equation (1.12). It is straightforward to prove equation (1.13) given the definition of return to government bond and noting that $\mathbb{E}_t\left[\xi_{t+1}\right] = \Phi'_\xi(0)$.

Return of the Housing Asset

In this subsection I cover the derivation of price-dividend ratio and asset return dynamics for the housing asset. In the following, the region index j is dropped as the derivation is similar for all the regions. I rely on numerical methods as in Lettau, Ludvigson, and

Wachter (2008) to solve for the price-dividend ratio. While achieving a closed-form solution is also possible by using Campbell-Shiller log-linearization, I use numerical methods to be more accurate.

The return to the housing asset can be expressed in terms of the log price-dividend ratio as $R_t^h = \frac{1+e^{pd_{t+1}^h}}{e^{pd_t^h}}.\frac{s_{t+1}}{s_t}$. As a result the Euler equation (1.15) implies:

$$\mathbb{E}_t\left[\exp\left(\log M_{t+1} + \log\left(1 + e^{pd_{t+1}^h}\right) - pd_t^h + \log s_{t+1} - \log s_t\right)\right] = 1. \tag{A.3}$$

Rearranging the above we can find pd_t^h in a recursive equation:

$$pd_t^h = \log \mathbb{E}_t\left[\exp\left(\log M_{t+1} + \log\left(1 + e^{pd_{t+1}^h}\right) + \log s_{t+1} - \log s_t\right)\right]. \tag{A.4}$$

According to equations (1.9) and (1.14) the above can be expressed in terms of our model parameters:

$$\begin{aligned} pd_t^h \quad = \quad & \mu_s - r_0 - r_p p_t + \frac{1}{2}\left(\sigma_{my} + \sigma_s\rho_s\right)^2 + \frac{1}{2}\sigma_s^2(1-\rho_s^2) + \frac{1}{2}\sigma_{mp}^2 p_t \\ & + \log(1 - p_t + \Phi_\xi(\phi_s - \gamma)p_t) + \log\mathbb{E}_t\left[1 + e^{pd_{t+1}^h}\right]. \end{aligned} \tag{A.5}$$

We solve for the above recursively on a fine grid of p_t values.

B Household Problem

Homeowner Problem

In the model, homeowners need to choose their level of housing and non-housing consumptions (s_t, c_t). Also, they decide whether or not to sell the house and become a renter ($I_{t,Sell}$). These decisions are made conditional on the probability of entering a disaster p_t, level of labor income y_t, liquid wealth invested in the government bond w_t, outstanding amount of mortgage loan l_t, and the amount of debt services m_t. The Bellman equation for the homeowner is:

$$U^h_{i,t}(p_t, y_t, w_t, h_t, l_t, m_t) = \max_{c_t, s_t, I_{t,Sell}} \Big[(1-\delta)\left(s_t^{\nu} c_t^{1-\nu}\right)^{\frac{1-\gamma}{\theta}} \tag{B.1}$$
$$+\delta \mathbb{E}_t \left[\left((1-I_{t,Def})\left((1-I_{t,Sell})U^h_{i,t+1} + I_{t,Sell}U^r_{i,t+1}\right) + I_{t,Def}U^d_{i,t+1}\right)^{1-\gamma}\right]^{\frac{1}{\theta}}\Big]^{\frac{\theta}{1-\gamma}}$$

subject to

$$\begin{aligned}
I_{t,Def} &= \mathbb{1}_{y_t + w_t + h_t\left[1+\exp(-pd^h_t)\right] - (s_t + c_t + l_t + m_t) < 0}, \\
w_{t+1} &= (1-I_{t,Def})\left[y_t + w_t + h_t \exp(-pd^h_t) - m_t + I_{t,Sell}(h_t - l_t)\right] R^b_{t+1}, \\
h_{t+1} &= (1-I_{t,Def})(1-I_{t,Sell}) h_t R^{h,ex}_{t+1}, \\
l_{t+1} &= (1-I_{t,Def})(1-I_{t,Sell}) l_t, \qquad m_{t+1} = (1-I_{t,Def})(1-I_{t,Sell}) m_t, \\
& \text{and} \quad c_t, s_t \geq 0.
\end{aligned} \tag{B.2}$$

In the above, the value function of household i at time t owning a house, is denoted by $U^h_{i,t}$. Similarly, value to renter and defaulted households are denoted respectively by $U^r_{i,t}$ and $U^d_{i,t}$. As evident in the above equations, there is a possibility that a homeowner is unable to meet their financial obligations even at the minimum level of consumption ($s_t = c_t = 0$), even after selling the house. In that case the homeowner is in a state of default ($I_{t,Def}$). I assume that rents or dividends from owning a house are liquid such that the homeowner can decide to consume more or less of the housing consumption and use the rest for non-housing consumption and/or investment in government bond.

If the homeowner decides to sell the house, the proceeds are used to pay the outstanding balance of the mortgage and in the next period the balance $l_{i,t+1}$ and debt interest cost $m_{i,t+1}$ become zero.

Note that the parameter ν determines the optimal division of consumption between housing and non-housing goods at a given level of total consumption ($s_{i,t} + c_{i,t} = Cons.$), independent of other variables. To be more clear, in order to maximize $X_{i,t} = s_{i,t}^{\nu} c_{i,t}^{1-\nu}$, we have $\frac{c_{i,t}}{s_{i,t}} = \frac{1-\nu}{\nu}$.

Renter Problem

The problem for a renter household is to make consumption decisions in addition to the decision about increasing her leverage and investing in the housing asset $I_{t,Buy}$. The value for renter household is

$$U_{i,t}^{r}(p_t, y_t, w_t) = \max_{c_t, s_t, I_{t,Buy}} \left[(1-\delta) \left(s_t^{\nu} c_t^{1-\nu} \right)^{\frac{1-\gamma}{\theta}} + \delta \mathbb{E}_t \left[\left((1 - I_{t,Buy}) U_{i,t+1}^{r} + I_{t,Buy} U_{i,t+1}^{h} \right)^{1-\gamma} \right]^{\frac{1}{\theta}} \right]^{\frac{\theta}{1-\gamma}} \tag{B.3}$$

subject to

$$\begin{aligned} w_{t+1} &= \left[y_t + w_t - (s_t + c_t) + I_{t,Buy}(l_t - h_t) \right] R_{t+1}^{b}, \\ h_{t+1} &= I_{t,Buy} y_t \frac{DTI}{LTV} R_{t+1}^{h,ex}, \\ l_{t+1} &= I_{t,Buy} y_t DTI, \qquad m_{t+1} = I_{t,Buy} y_t DTI \left(R_{t+1}^{d} - 1 \right), \\ & \text{and} \quad c_t, s_t \geq 0. \end{aligned} \tag{B.4}$$

If the household decides to remain a renter, all her investment will be in the government bond and she does not have any other chance to increase her leverage. If she decides to invest, the value of the house and her available mortgage are constrained by the DTI and

LTV ratios. To be more specific, the mortgage outstanding balance at origination is set to be equal to $y_t DTI$ and the house value is $y_t \frac{DTI}{LTV}$. As a result, the mortgage interest payment will be $y_t DTI \left(R^d_{t+1} - 1\right)$.

Defaulted Household Problem

A defaulted household, is in essence a renter who is restricted from the mortgage market and as a result unable to invest in the housing asset. This restriction is a result of past default. The problem for a defaulted household is characterized by the following Bellman equation:

$$U^d_{i,t}(p_t, y_t, w_t) = \max_{c_t, s_t} \left[(1-\delta) \left(s_t^{\nu} c_t^{1-\nu}\right)^{\frac{1-\gamma}{\theta}} + \delta \mathbb{E}_t \left[\omega \left(U^r_{i,t+1}\right)^{1-\gamma} + (1-\omega) \left(U^d_{i,t+1}\right)^{1-\gamma} \right]^{\frac{1}{\theta}} \right]^{\frac{\theta}{1-\gamma}} \tag{B.5}$$

subject to

$$\begin{aligned} w_{t+1} &= \left[y_t + w_t - (s_t + c_t)\right] R^b_{t+1}, \\ &\text{and} \quad c_t, s_t \geq 0. \end{aligned} \tag{B.6}$$

With probability ω this household can become a(n unrestricted) renter and regain eligibility to receive a mortgage in order to invest in the housing market.

Numerical Procedures in Model Solution

As characterization of the household problem in previous subsections makes it clear, we can reduce the number of state variables by considering ratio of consumption and investment decisions with respect to labor income $y_{i,t}$ instead of considering their absolute values. This is helpful to reach stationarity in the state variables of the model too. Hence, I define

$$\tilde{s}_t = \frac{s_t}{y_t}, \quad \tilde{c}_t = \frac{c_t}{y_t}, \quad \tilde{w}_t = \frac{w_t}{y_t}, \quad \tilde{h}_t = \frac{h_t}{y_t}, \quad \tilde{l}_t = \frac{l_t}{y_t}, \quad \tilde{m}_t = \frac{m_t}{y_t}. \tag{B.7}$$

Gomes, Grotteria, and Wachter (2019) follow a similar scaling method to construct stationary state variables in their optimization problem.

I conjecture that we can use the scaled household value, $\tilde{U}_{i,t} = \frac{U_{i,t}}{y_{i,t}}$ in our numerical solution. Hence, the scaled household utility is represented as

$$\tilde{U}_{i,t} = \left[(1-\delta) \left(\tilde{s}_t^{\nu} \tilde{c}_t^{1-\nu} \right)^{\frac{1-\gamma}{\theta}} + \delta \mathbb{E}_t \left[\left(\frac{y_{t+1}}{y_t} \tilde{U}_{i,t+1} \right)^{1-\gamma} \right]^{\frac{1}{\theta}} \right]^{\frac{\theta}{1-\gamma}}, \tag{B.8}$$

in terms of scaled housing and non-housing consumption values and scaled utility in the next period.

I consider log-values for each of the scaled state and choice variables and discretize with 7-15 points in each dimension. The values inside this multi-dimensional grid are calculated using linear interpolation. For points outside the grid, I use the value at the nearest grid point.

Figure 1.1: Nationwide Dispersion in Mortgage EDFs

Panel A: Average EDFs across all MSAs

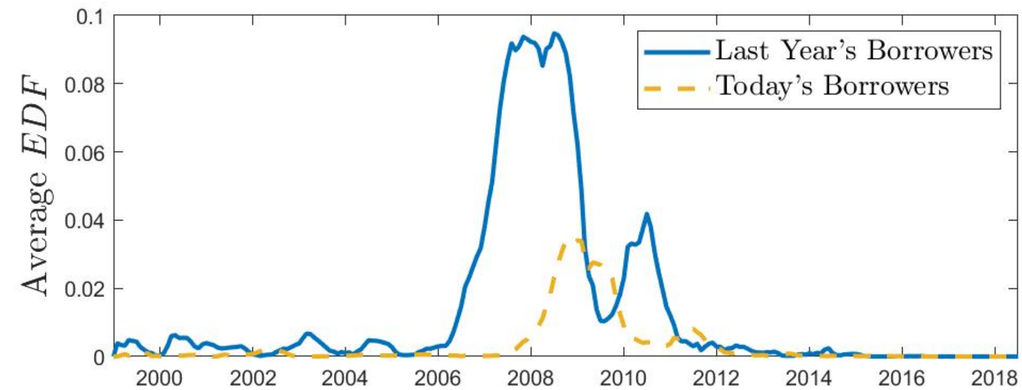

Panel B: Average Dispersion across all MSAs

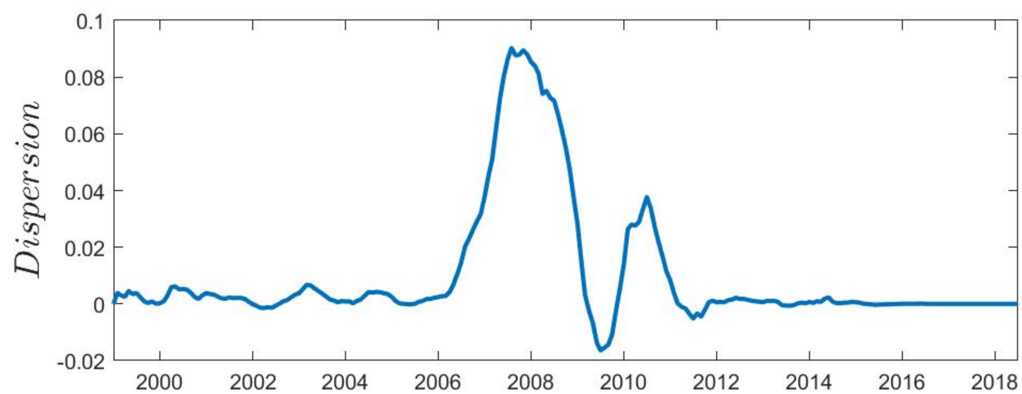

Notes: Panel A in this figure presents average EDF values for today and last year borrowers of mortgages, aggregated to the nation-wide level by averaging the associated EDF values for all the MSAs. Panel B shows dispersion in EDFs by subtracting EDF of borrowers from repayers. Most of the variation in dispersion is driven by changes in EDF of repayers.

Figure 1.2: Distribution of FICO Scores

Panel A: FICO scores

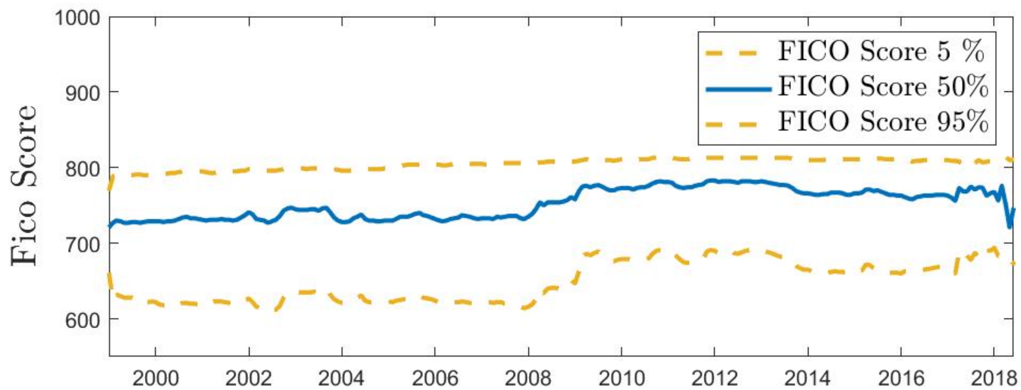

Panel B: Standard Deviation of FICO scores

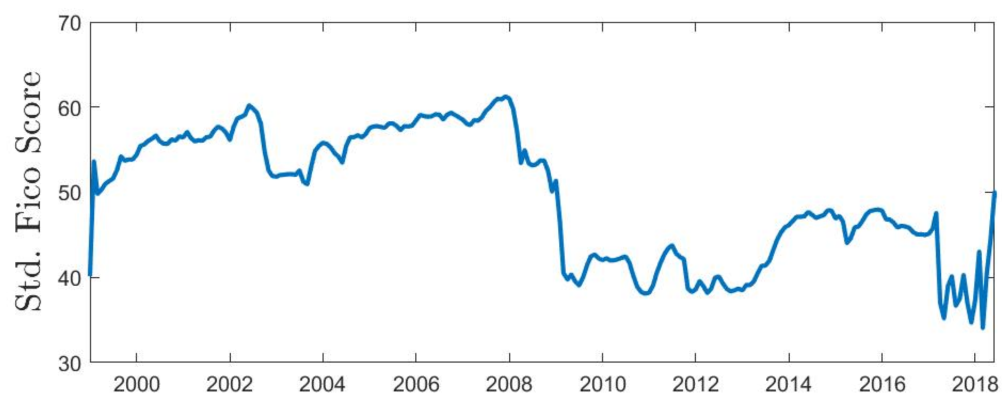

Notes: Panel A in this figure presents 5, 50 and 95 percentiles of FICO credit score distribution for all the borrowers in the data. Panel B depicts the standard deviation of FICO credit scores.

Figure 1.3: Heterogeneity in Labor Income Exposures Across States

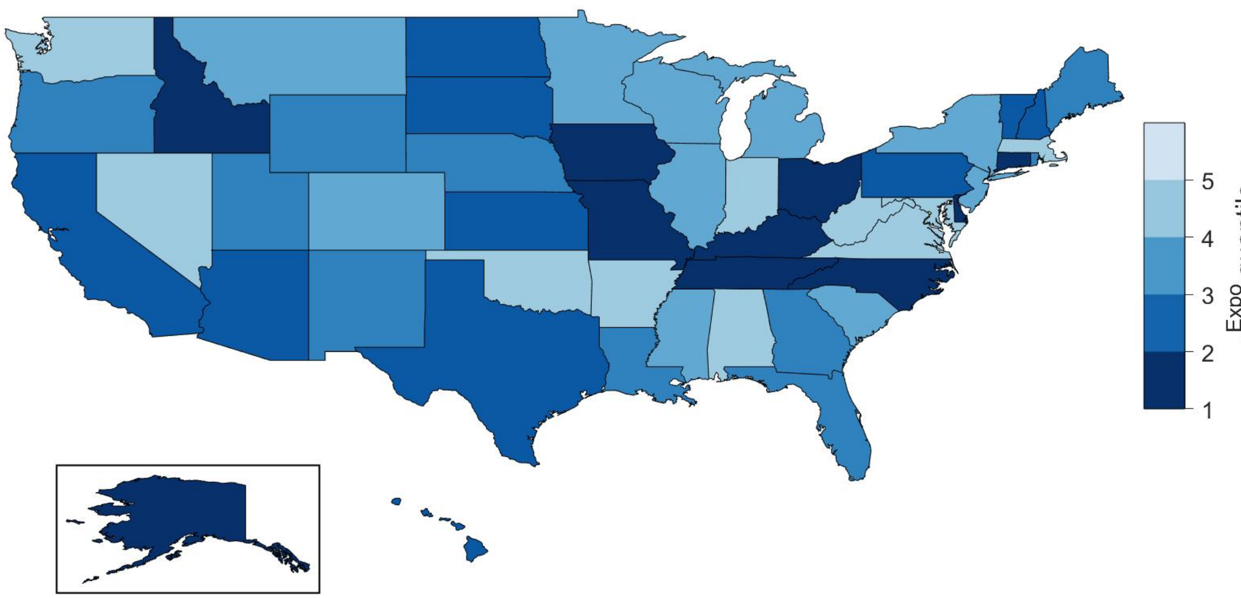

Notes: The figure depicts average heterogeneity in exposures, σ^{Expo}, across different states. Darker colors represent less heterogeneity in exposures.

Figure 1.4: Heterogeneity in Labor Income Exposures Across States in Time

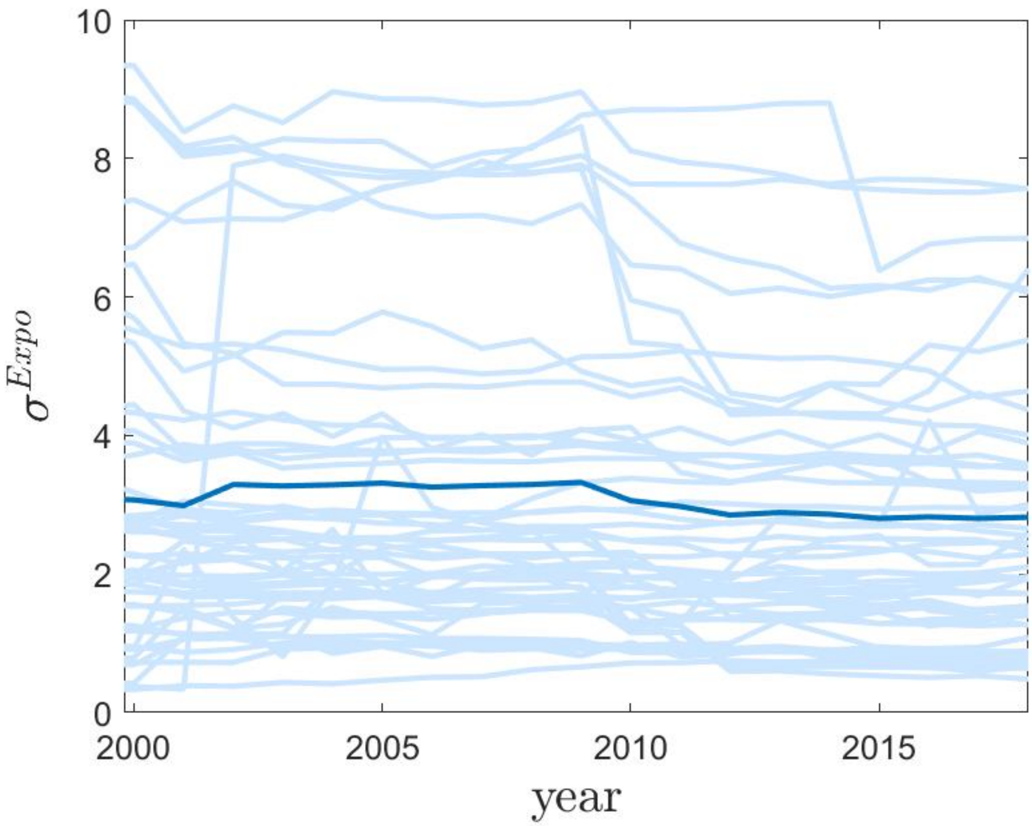

Notes: The figure depicts average heterogeneity in exposures, σ^{Expo}, across different states in time. The darker line represent average σ^{Expo} at each point in time.

Figure 1.5: Impulse Response Functions

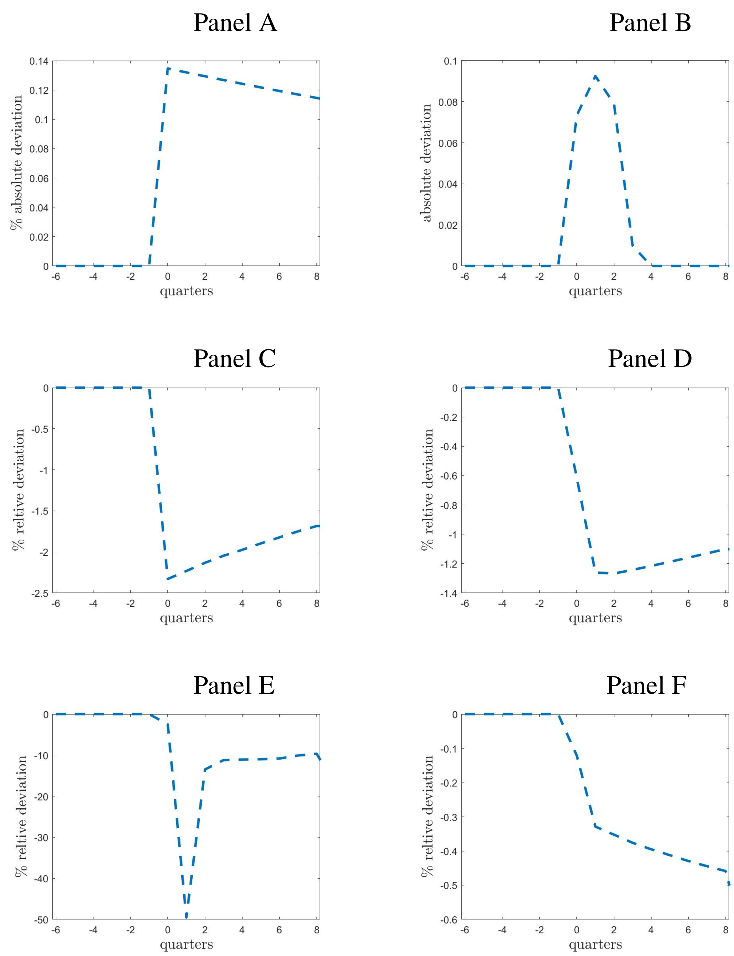

Notes: This figure plots the impulse response function of model variables in the following 8 quarters, to a one standard deviation positive shock to probability of rare events. Panel A shows absolute change in probability of rare disaster in percentages. Panel B depicts the absolute change in $Dispersion$. Panels C, D, E, F show plot the relative change in consumption, GDP, housing investment and mortgage credit.

Figure 1.6: Heterogeneity and GDP Impulse Response Function

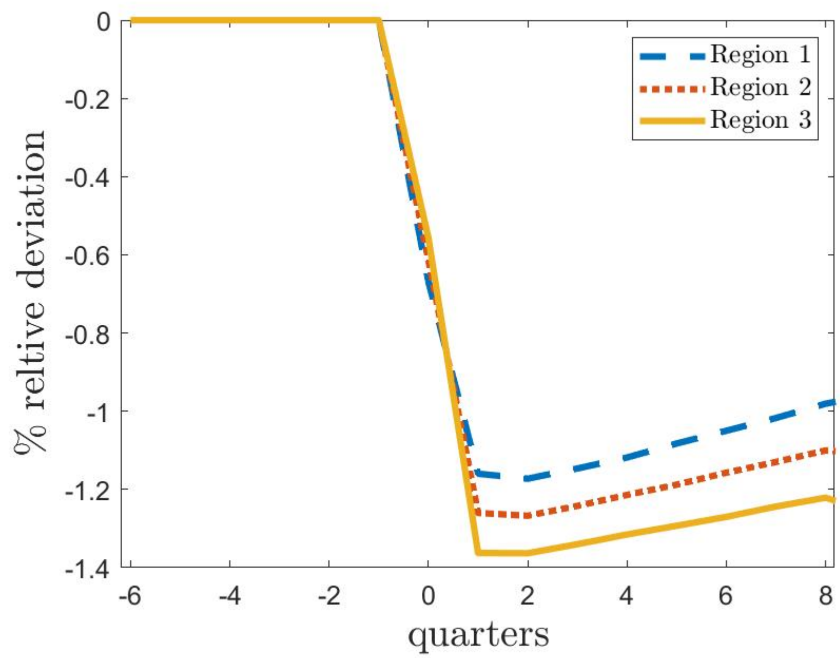

Notes: This figure plots the impulse response function of GDP per capita for the three regions in the simulations. The heterogeneity of labor income exposures increases from region 1 to region 3.

Table 1.1: Dispersion and Measures of Aggregate Risk

Variable							
$Dispersion$	1.00	0.26	0.96	-0.11	0.14	0.17	-0.16
Avergae EDF - today's borrowers	0.26	1.00	0.53	-0.54	0.09	0.74	-0.02
Avergae EDF - last year's borrowers	0.96	0.53	1.00	-0.25	0.15	0.37	-0.14
Price-dividend ratio	-0.11	-0.54	-0.25	1.00	0.48	-0.35	0.14
CAY (Lettau and Ludvigson, 2001)	0.14	0.09	0.15	0.48	1.00	0.32	0.15
GZ spread (Gilchrist and Zakrajšek, 2012)	0.17	0.74	0.37	-0.35	0.32	1.00	-0.02
VRP (Bollerslev, Tauchen, and Zhou, 2009)	-0.16	-0.02	-0.14	0.14	0.15	-0.02	1.00

Notes: this table reports pairwise correlations between $Dispersion$, average EDF of last year's borrowers, average EDF of today's borrowers, price-dividend ratio, CAY (Lettau and Ludvigson, 2001) , GZ credit spread (Gilchrist and Zakrajšek, 2012), and variance risk premium (Bollerslev, Tauchen, and Zhou, 2009). All variables are available in monthly frequency except for CAY which is reported in quarterly frequency. A monthly version of CAY is constructed by assuming it remains constant between quarterly updates.

Table 1.2: Forecasting GDP Growth

Horizon (Quarter)	1	2	3	4	1	2	3	4
β	-0.033	-0.036	-0.040	-0.041	-0.022	-0.023	-0.024	-0.024
	[-5.93]	[-7.35]	[-8.64]	[-9.78]	[-5.11]	[-6.08]	[-6.85]	[-7.58]
State FE	Y	Y	Y	Y	Y	Y	Y	Y
Time FE	N	N	N	N	Y	Y	Y	Y
R^2	0.026	0.060	0.079	0.103	0.224	0.275	0.309	0.333

Notes: the table reports forecasting regression results for GDP growth at the state level using quarterly data. The fixed-effect regression is specified as follows:

$$\Delta gdp_{i,t\rightarrow t+k} = \alpha_i + \alpha_t + \beta Dispersion_{i,t} + \xi \Delta gdp_{i,t-1\rightarrow t} + \epsilon_{i,t\rightarrow t+k}.$$

Under the estimated coefficients, robust t-stats are reported in parenthesis.

Table 1.3: Forecasting Employment Growth

Horizon (Month)	1	3	6	9	12	1	3	6	9	12
β	-0.002	-0.003	-0.005	-0.006	-0.007	-0.001	-0.002	-0.002	-0.003	-0.003
	[-4.22]	[-5.00]	[-5.85]	[-6.72]	[-7.71]	[-2.58]	[-2.91]	[-3.38]	[-3.92]	[-4.56]
State FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Time FE	N	N	N	N	N	Y	Y	Y	Y	Y
R^2	0.691	0.614	0.481	0.406	0.345	0.734	0.725	0.676	0.648	0.628

Notes: the table reports forecasting regression results for employment growth at the state level using monthly data. The fixed-effect regression is specified as follows:

$$\Delta emp_{i,t\rightarrow t+k} = \alpha_i + \alpha_t + \beta Dispersion_{i,t} + \xi \Delta emp_{i,t-1\rightarrow t} + \epsilon_{i,t\rightarrow t+k}$$

Under the estimated coefficients, robust t-stats are reported in parenthesis.

Table 1.4: Forecasting GDP Growth: MSA-Level

Horizon (Year)	1	2	1	2
β	-0.128	-0.110	-0.073	-0.064
	[-10.91]	[-12.55]	[-7.01]	[-7.96]
State FE	Y	Y	Y	Y
Time FE	N	N	Y	Y
R^2	0.0765	0.0763	0.204	0.225

Notes: the table reports forecasting regression results for GDP growth at the MSA level using annual data. The fixed-effect regression is specified as follows:

$$\Delta gdp_{i,t\to t+k} = \alpha_i + \alpha_t + \beta Dispersion_{i,t} + \xi \Delta gdp_{i,t-1\to t} + \epsilon_{i,t\to t+k}$$

Under the estimated coefficients, robust t-stats are reported in parenthesis.

Table 1.5: Forecasting Employment Growth: MSA-Level

Horizon (Month)	1	3	6	9	12	1	3	6	9	12
β	-0.008	-0.011	-0.014	-0.017	-0.022	-0.005	-0.009	-0.012	-0.013	-0.017
	[-4.22]	[-5.00]	[-5.85]	[-6.72]	[-7.71]	[-2.58]	[-2.91]	[-3.38]	[-3.92]	[-4.56]
MSA FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Time FE	N	N	N	N	N	Y	Y	Y	Y	Y
R^2	0.587	0.516	0.383	0.326	0.285	0.625	0.617	0.567	0.552	0.526

Notes: the table reports forecasting regression results for employment growth at the MSA level using monthly data. The fixed-effect regression is specified as follows:

$$\Delta emp_{i,t\to t+k} = \alpha_i + \alpha_t + \beta Dispersion_{i,t} + \xi \Delta emp_{i,t-1\to t} + \epsilon_{i,t\to t+k}$$

Under the estimated coefficients, robust t-stats are reported in parenthesis.

Table 1.6: Forecasting GDP: Role of Heterogeneity in Exposures

Horizon (Quarter)	1	2	3	4	1	2	3	4
β	-0.033	-0.036	-0.039	-0.040	0.005	0.000	-0.006	-0.010
	[-5.76]	[-7.22]	[-8.58]	[-9.81]	[0.62]	[0.02]	[-0.72]	[-1.46]
γ					-0.016	-0.016	-0.014	-0.013
					[-4.50]	[-4.08]	[-4.13]	[-4.23]
δ	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001	-0.001
	[-2.67]	[-2.33]	[-2.21]	[-2.16]	[-2.53]	[-2.20]	[-2.09]	[-2.07]
State FE	Y	Y	Y	Y	Y	Y	Y	Y
R^2	0.028	0.067	0.090	0.117	0.032	0.074	0.099	0.127

Notes: the table reports forecasting regression results for GDP growth at the state level using quarterly data. The fixed-effect regression is specified as follows:

$$\begin{aligned}\Delta gdp_{i,t\rightarrow t+k} &= \alpha_i + \beta Dispersion_{i,t} + \gamma Dispersion_{i,t} \times \sigma_{i,t}^{Expo} + \delta \sigma_{i,t}^{Expo} \\ &\quad + \xi \Delta gdp_{i,t-1\rightarrow t} + \epsilon_{i,t\rightarrow t+k}\end{aligned}$$

Under the estimated coefficients, robust t-stats are reported in parenthesis.

Table 1.7: Forecasting GDP: Role of Heterogeneity in Exposures

Horizon (Quarter)	1	2	3	4
β	0.002	-0.002	-0.009	-0.013
	[0.23]	[-0.26]	[-0.89]	[-1.51]
γ	-0.018	-0.017	-0.015	-0.013
	[-3.54]	[-3.30]	[-3.35]	[-3.42]
State FE	Y	Y	Y	Y
R^2	0.029	0.069	0.092	0.119

Notes: the table reports forecasting regression results for GDP growth at the state level using quarterly data. The fixed-effect regression is specified as follows:

$$\Delta gdp_{i,t\rightarrow t+k} = \alpha_i + \beta Dispersion_{i,t} + \gamma Dispersion_{i,t} \times \sigma_i^{Expo} + \xi \Delta gdp_{i,t-1\rightarrow t} + \epsilon_{i,t\rightarrow t+k}$$

Under the estimated coefficients, robust t-stats are reported in parenthesis.

Table 1.8: Forecasting Employment: Role of Heterogeneity in Exposures

Horizon (Month)	1	3	6	9	12	1	3	6	9	12
β	-0.002	-0.003	-0.005	-0.006	-0.007	-0.000	-0.000	-0.001	-0.001	-0.002
	[-4.20]	[-4.97]	[-5.81]	[-6.66]	[-7.61]	[-0.43]	[-0.58]	[-0.43]	[-0.57]	[-0.87]
γ						-0.001	-0.001	-0.002	-0.002	-0.002
						[-3.47]	[-3.36]	[-3.19]	[-2.97]	[-2.76]
δ	0.000	0.000	0.000	-0.000	-0.000	0.000	0.000	0.000	0.000	-0.000
	[0.40]	[0.24]	[0.10]	[-0.09]	[-0.31]	[0.49]	[0.34]	[0.23]	[0.06]	[-0.16]
State FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
R^2	0.689	0.612	0.479	0.405	0.345	0.689	0.613	0.481	0.409	0.351

Notes: the table reports forecasting regression results for employment growth at the state level using monthly data. The fixed-effect regression is specified as follows:

$$
\begin{aligned}
\Delta emp_{i,t\to t+k} &= \alpha_i + \beta Dispersion_{i,t} + \gamma Dispersion_{i,t} \times \sigma_{i,t}^{Expo} + \delta \sigma_{i,t}^{Expo} \\
&\quad + \xi \Delta emp_{i,t-1\to t} + \epsilon_{i,t\to t+k}
\end{aligned}
$$

Under the estimated coefficients, robust t-stats are reported in parenthesis.

Table 1.9: Forecasting Employment: Role of Heterogeneity in Exposures

Horizon (Month)	1	3	6	9	12
β	-0.000	-0.001	-0.001	-0.001	-0.001
	[-0.37]	[-0.61]	[-0.43]	[-0.52]	[-0.77]
γ	-0.001	-0.001	-0.002	-0.002	-0.003
	[-3.53]	[-3.31]	[-3.13]	[-2.91]	[-2.72]
State FE	Y	Y	Y	Y	Y
R^2	0.689	0.613	0.481	0.409	0.350

Notes: the table reports forecasting regression results for employment growth at the state level using monthly data. The fixed-effect regression is specified as follows:

$$\Delta emp_{i,t\to t+k} \quad = \quad \alpha_i + \beta Dispersion_{i,t} + \gamma Dispersion_{i,t} \times \sigma_i^{Expo} + \xi \Delta emp_{i,t-1\to t} + \epsilon_{i,t\to t+k}$$

Under the estimated coefficients, robust t-stats are reported in parenthesis.

Table 1.10: Calibration of Model Parameters

Preferences				
	δ	γ	ψ	ν
	0.997	4	1.5	0.15
Disaster risk				
	$\bar{p}$	ρ_p	σ_p	q
	0.0075	0.98	0.0175	0.4
SDF				
	r_0	r_p	σ_{my}	σ_{mp}
	0.01	3	-0.02	1.2
Housing sector				
	μ_s	σ_s	ρ_s	ϕ_s
	0.005	0.01	1	3
Household labor income	σ_y^i	ρ_y^i	ϕ_y^i	
	0.1	0.1	[0.5, 1, 1.5]	
Household debt				
	DTI	LTV	ω	r_{debt}
	14	0.8	0.0375	1

Notes: This table reports the parameters of the model, calibrated at quarterly frequency.

Table 1.11: Coefficients of Predictive Regressions: No Disaster Samples

Panel A					
	Horizon (Quarter)	1	2	3	4
GDP	β	-0.035	-0.024	-0.013	-0.008
	R^2	0.058	0.050	0.021	0.009
Consumption	β	-0.080	-0.032	-0.015	-0.008
	R^2	0.109	0.036	0.015	0.010
Investment in housing	β	-0.743	-0.506	-0.216	-0.161
	R^2	0.350	0.239	0.251	0.230

Panel B					
	Horizon (Quarter)	1	2	3	4
GDP	β	-0.032	-0.020	-0.010	-0.006
	γ	-0.012	-0.016	-0.010	-0.007
	R^2	0.056	0.049	0.020	0.009
Consumption	β	-0.077	-0.032	-0.016	-0.008
	γ	-0.019	0.000	0.003	0.004
	R^2	0.107	0.033	0.012	0.007
Investment in housing	β	-0.575	-0.499	-0.211	-0.158
	γ	-0.582	-0.068	-0.067	-0.022
	R^2	0.350	0.239	0.251	0.231

Notes: This table presents median coefficients and R^2s estimated in predictive fixed-effect regressions of the following form:

$$\Delta y_{i,t\to t+k} = \alpha_i + \beta Dispersion_{i,t} + \gamma Dispersion_{i,t} \times \sigma_i^{Expo} + \xi \Delta y_{i,t-1\to t} + \epsilon_{i,t\to t+k},$$

where y stands for growth rate of GDP, consumption, and investment in housing asset. The simulations include a total of 10,000 simulated economies simulated for 20 years, out of which 5,863 economies do not experience a disaster.

Table 1.12: Coefficients of Predictive Regressions: All Samples

Panel A					
	Horizon (Quarter)	1	2	3	4
GDP	β	-0.040	-0.027	-0.016	-0.010
	R^2	0.067	0.056	0.030	0.018
Consumption	β	-0.080	-0.034	-0.018	-0.010
	R^2	0.100	0.037	0.018	0.012
Investment in housing	β	-0.902	-0.586	-0.276	-0.203
	R^2	0.333	0.232	0.235	0.215

Panel B					
	Horizon (Quarter)	1	2	3	4
GDP	β	-0.038	-0.023	-0.013	-0.008
	γ	-0.008	-0.012	-0.007	-0.005
	R^2	0.066	0.055	0.029	0.017
Consumption	β	-0.082	-0.036	-0.019	-0.011
	γ	-0.001	0.006	0.006	0.006
	R^2	0.099	0.035	0.015	0.010
Investment in housing	β	-0.731	-0.587	-0.271	-0.202
	γ	-0.498	-0.006	-0.029	0.002
	R^2	0.333	0.232	0.235	0.215

Notes: This table presents median coefficients and R^2s estimated in predictive fixed-effect regressions of the following form:

$$\Delta y_{i,t\to t+k} = \alpha_i + \beta Dispersion_{i,t} + \gamma Dispersion_{i,t} \times \sigma_i^{Expo} + \xi \Delta y_{i,t-1\to t} + \epsilon_{i,t\to t+k},$$

where y stands for growth rate of GDP, consumption, and investment in housing asset. The simulations include a total of 10,000 simulated economies simulated for 20 years.

2 Learning, Slowly Unfolding Disasters, and Asset Prices

2.1 Introduction

Do investors fear the possibility of another Great Depression? This question encapsulates how models with rare economic disasters account for the high equity premium and the low risk-free rate in the postwar period (Rietz, 1988; Barro, 2006). Typically, disaster risk models assume a small probability of an instantaneous drop in aggregate consumption, whose frequency and size distribution are calibrated to match large peak-to-trough consumption declines compiled across various countries (Barro and Ursúa, 2008).

We thank Bill Schwert (the editor at the Journal of Financial Economics) and the anonymous referee for their valuable comments and suggestions. We are also grateful to John Adams, Hengjie Ai, Daniel Andrei, Jaroslav Borovička, Tom Chang, Hui Chen, Tony Cookson, Dean Corbae, Hitesh Doshi, Redouane Elkamhi, Bjørn Eraker, Xavier Gabaix, Joao Gomes, François Gourio, Jerry Hoberg, Kris Jacobs, Alexandre Jeanneret, Travis Johnson, Chris Jones, Dana Kiku, Hagen Kim, Arthur Korteweg, Howard Kung, Oliver Levine, Ye Li, Lars Lochstoer, Alan Moreira, Kevin Murphy, Sebastien Plante, Vincenzo Quadrini, Erwan Quintin, Nikolai Roussanov, Guillaume Roussellet, Ivan Shaliastovich, Rauli Susmel, Selale Tuzel, Jessica Wachter, Randall Wright, Ben Zhang, and seminar participants at the 2021 AFA meeting, 14th Macro Finance Society Workshop, Finance in the Cloud V, 2020 MFA meeting, 8th HEC-McGill Winter Finance Workshop, 2020 SWFA meeting, 2019 Lone Star Finance Symposium, University of Houston, University of Southern California, and University of Wisconsin-Madison for helpful comments. An earlier version of this paper was circulated under the title "Slowly Unfolding Disasters."

A significant challenge to Rietz-Barro type models with instantaneous disasters is that disasters in the historical data unfold slowly over multiple periods of time not only in the macroeconomy, but also in financial markets. As documented by Barro and Ursúa (2017), the average duration of stock market disasters is about 3 years, as lengthy as macroeconomic disasters. Granted, it is possible to assume slowly unfolding macroeconomic disasters by introducing a persistent low-probability regime with negative growth rates (e.g., Gourio, 2012; Nakamura, Steinsson, Barro, and Ursúa, 2013). However, even in such a model, the stock market fully responds instantaneously as soon as the economy enters the disaster regime, due to the forward looking nature of prices. Once investors become aware of being in the disaster state, they immediately incorporate the low growth expectations into prices, which leads to a large instantaneous price decline in the market.

We argue that this inability to capture gradual stock market declines during disaster periods is actually what is behind the criticism raised by Welch (2016) and Cochrane (2017). Welch (2016) observes that a one-month put-protected strategy with 85% moneyness cannot lose more than 15% in any month, and therefore, should not bear a high risk premium if large instantaneous price declines during disasters are the true source of the equity premium. However, he finds that the put-protected portfolio still earns a premium close to the aggregate equity premium. We claim that this failure is not due to the rare disaster mechanism itself. To illustrate, Figure 2.1 plots the monthly price declines in the CRSP value-weighted portfolio during the Great Depression (from September 1929 to June 1932). While this period represents a cumulative stock market decline of 86%, we observe only a few months with a monthly decline larger than 15%. As a result, rolling over a one-month put option does not provide much insurance against the Great Depression, resulting in a -80% cumulative return, even ignoring the insurance cost. This number is in stark contrast with a loss of 15%, the maximum loss of the put-protected portfolio if the Great Depression were an isolated instantaneous drop as in Rietz-Barro type models. That is, the real issue at hand is the behavior of the stock market during disaster periods, which in turn determines how the possibility of disasters affects asset prices in normal times.

In this paper, we develop a model that generates slowly unfolding disasters both in the macroeconomy and in financial markets. Due to imperfect information, investors cannot exactly distinguish whether the economy is experiencing a mild and temporary downturn or is on the verge of a severe and prolonged disaster. As a result, disaster periods are not fully identifiable by investors *ex ante* at the onset, but only identified *ex post* using the peak-to-trough approach as in the data.

In our model, the consumption jump intensity follows an autoregressive process where its long-run mean is subject to a rare but persistent shift from its low value in the "normal" state to a very high value in the "depression" state. Such an increase can cause a potentially extended period of negative consumption growth due to more frequent negative jump realizations that may cumulatively constitute a macroeconomic disaster. We assume that the representative agent with recursive utility is fully informed about the probability of a jump within the next month. However, the long-run mean of the jump intensity is not observable. Therefore, the agent does not know with certainty whether a high current jump intensity is due to a transitory increase in risk or due to a persistent shift to the depression state that may result in a prolonged disaster period. We use standard Bayesian learning to model the agent's assessment of the probability of being in the depression state. Even after the economy enters the depression state, it takes time for the agent to recognize this and change her belief accordingly. As a result, equity prices react to persistent declines in consumption slowly.

A prime example that illustrates the economic intuition behind our model's mechanism is the bankruptcy of Lehman Brothers in September 2008, which was arguably the most terrifying moment in the U.S. economy since the Great Depression. It seems reasonable to assume that investors were aware of an increase in immediate economic risk upon Lehman's default. However, investors did not know with certainty if this event would trigger a prolonged economic crisis with a similar severity as the Great Depression, or if the economy would recover from the crisis more quickly. In our model, the first scenario is represented by a regime shift to the depression state, and the latter by a transitory increase in the jump intensity. The 2008 financial crisis did not turn out to become a macroeconomic disaster like the Great Depression, but potentially it could have. More importantly, at that moment, investors were not able to know with certainty whether or not it would.

While our model differs from existing disaster risk models in its mechanism and its empirical implications, it is possible to nest them within our framework. To illustrate these differences quantitatively, we calibrate our model and compare it with the time-varying disaster risk model of Wachter (2013), in which stock prices decline instantaneously at the onset of macroeconomic disasters. Both models are consistent with standard asset pricing moments, such as the equity premium, risk-free rate, stock market volatility, and predictability of excess returns as well as consumption and dividend moments.

However, a clear distinction between the two models arises when their implications for put-protected portfolios of Welch (2016) and the variance risk premium are examined. Unlike instantaneous disaster risk models, our model is capable of accounting for the risk premia on put-protected strategies with various moneyness values ranging from 75% to 90%. This is attributable to the slowly unfolding nature of disasters in the stock market, which our model is able to capture through the realistic identification of disaster episodes. The source of the equity premium in our model is not the prospect of an instantaneous disaster but the variation in the anticipation of experiencing a prolonged disaster period in the future.

We also show that disaster risk models with instantaneous equity price declines during disasters imply unrealistically high values of the variance risk premium, and hence, the VIX. This is because abrupt drops in prices give rise to an unrealistically high quadratic variation, which increases the expected payoffs from deep out-of-the-money put options and variance swaps to extreme levels. In our model, the realistic behavior of the equity price path during disasters produces a level of the VIX and the variance risk premium that is consistent with the data.

In our model, information frictions play a significant role in generating realistic stock price dynamics during disasters. We discover that the special case of our model with perfect information struggles to explain the put-protected portfolio premia, VIX, and variance risk premium. This is because in this nested model, the transition to the depression state is directly observable by the agent, causing the stock market to fully react immediately to the regime shift. In sum, a model that accurately depicts how disasters unfold in financial markets as well as in the macroeconomy is crucial to establishing the consistency of short-term contingent claim prices with the rare disaster mechanism. While traditional disaster risk models are inconsistent with the aforementioned moments, we show that this does not invalidate the mechanism itself. Instead, we emphasize the importance of modeling the realistic joint behavior of consumption and financial markets during disasters.

A distinctive feature of our model is that the size distribution of consumption disasters is a model outcome, not a model input. In our model, negative jumps in consumption are calibrated to be much smaller compared to what a typical disaster risk model assumes. However, large peak-to-trough declines in consumption can still be endogenously generated, collectively from small jumps in the depression regime. Model simulations show that if disasters are identified *ex post* using the peak-to-trough approach of Barro and Ursúa (2008), we obtain an average consumption decline of 18% and an average duration of 4.5

years for disaster periods. These results are fairly close to what we observe in the data (21% and 4.1 years). This suggests that our model is immune to the common criticism that disaster risk may overstate consumption risk by treating a peak-to-trough decline in consumption as if it happened within a unit period (Constantinides, 2008). Our resolution is rooted in recognizing that disasters in the data unfold slowly both in consumption and in equity prices.

Lastly, the main message of our paper not only applies to disasters, but also to large stock market declines in non-disaster periods that do not feature a large consumption decline. The main intuition behind our model is that at the beginning of a recession, investors face uncertainty about how long and severe economic downturns will be. This imperfect information impedes investors' ability to distinguish between, say, the Great Depression and the Great Recession *ex ante* at the onset, even though they turn out to be quite different *ex post*. Hence, our mechanism not only allows disaster periods to be in line with the data but also makes non-disaster periods more realistic. In fact, the post-1990 sample (or, dating back even further, the post-war sample) does not include any consumption disasters. However, we do observe severe stock market declines in the same period, such as the dot-com bubble, the Great Recession, and the COVID-19 crisis. We show that our model can rationalize the patterns of recent major crises through investors' belief channel, despite moderate consumption declines during such episodes.

Our paper contributes to the literature on equilibrium asset pricing, with a particular emphasis on extreme market events. Since its introduction by Rietz (1988) and Barro (2006), the rare disaster mechanism has been extended and refined to explain various aspects of macroeconomic and financial data.[1] In particular, to explain dynamic patterns in the data,

[1] Nakamura, Steinsson, Barro, and Ursúa (2013) consider multi-period disasters followed by periods of fast recovery to produce a more realistic representation of consumption dynamics. Hasler and Marfe (2016) demonstrate that considering disaster recovery in disaster models helps generate a downward sloping term structure of equity risk premia and an upward sloping term structure of interest rates. Farhi and Gabaix (2016)

Gabaix (2012), Gourio (2012), and Wachter (2013) introduce variable disaster risk, whose empirical significance for asset pricing is supported by Berkman, Jacobsen, and Lee (2011) and Manela and Moreira (2017). Since the main focus of these models is on asset pricing implications of macroeconomic disasters like the Great Depression, little attention has been paid to market crises with a relatively mild macroeconomic contraction, such as the Great Recession. The central idea behind our model is that the Great Depression and the Great Recession did not appear to be much different *ex ante* at the onset, although the Great Depression showed a far more severe macroeconomic contraction *ex post*. Not only does this allow us to address the shortcomings of traditional disaster risk models, it also offers a unified framework for various extreme market events, ranging from a short-lived crash to a prolonged depression.

Due to the rare nature of economic disasters, incorporating information frictions and investors' learning is particularly relevant for disaster risk models.[2] Gillman, Kejak, and Pakoš (2014) show that uncertainty about the length of a consumption disaster can address important features of equity and bond market data. Wachter and Zhu (2019) consider a model in which investors learn about the current probability of disasters from past realizations of disasters. In our paper, we focus on the implications of investors' learning not just for asset prices but for the joint dynamics of financial markets and the macroeconomy.

In recent work, Collin-Dufresne, Johannes, and Lochstoer (2016) investigate the role of parameter uncertainty in explaining standard asset pricing moments. As one exercise, the authors evaluate how uncertainty about each parameter in a two-state disaster risk model

show that their open economy model with variable disaster risk can address a series of puzzles in exchange rates.

[2]Our paper is also related to the vast literature that studies implications of investors' learning for asset pricing in various contexts. See, for example, Timmermann (1996), Veronesi (2000), Brandt, Zeng, and Zhang (2004), Veronesi (2004), Lettau, Ludvigson, and Wachter (2008), Pastor and Veronesi (2009a), Pastor and Veronesi (2009b), Ai (2010), Benzoni, Collin-Dufresne, and Goldstein (2011), Ju and Miao (2012), David and Veronesi (2014), and Johannes, Lochstoer, and Mou (2016).

contributes to the equity premium. Consequently, they conclude that particularly important is imperfect information about the transition probability from the disaster state to the normal state, which determines the duration of a disaster. Similar to their work, we also highlight investors' uncertainty about how long and severe future economic disasters will be. However, we create this uncertainty through a different mechanism: in our model, investors are unaware of the true economic state, not the model parameters. Therefore, unlike the model of Collin-Dufresne, Johannes, and Lochstoer (2016), investors cannot tell whether the economy is on the verge of a disaster. This imperfect information plays a key role in generating slow responses of equity prices to persistent consumption declines, which helps us account for several dimensions of the data including forward looking volatility, put-protected equity index returns, and the variance risk premium.

The rest of the paper proceeds as follows. Section 2.2 describes our disaster risk model with learning. Section 2.3 discusses the model calibration procedure. Section 2.4 provides results from the model and compares them with the data. Section 2.4.4 examines recent major crises through the lens of our model. Section 2.5 concludes.

2.2 Model

2.2.1 Model Setup

We consider an infinitely-lived representative agent in an endowment economy with complete markets. The agent has recursive preferences of Epstein and Zin (1989) and Weil (1989), which leads to the following stochastic discount factor M_{t+1}:

$$M_{t+1} = \exp\left(\theta \log \delta - \frac{\theta}{\psi}\Delta c_{t+1} + (\theta - 1) r_{c,t+1}\right), \tag{2.2.1}$$

where $\Delta c_{t+1} = \log\left(\frac{C_{t+1}}{C_t}\right)$ represents the logarithm of aggregate consumption growth and $r_{c,t+1}$ denotes the log return on the consumption claim. The coefficient $\theta = \frac{1-\gamma}{1-1/\psi}$ captures the agent's attitudes toward the timing of uncertainty resolution, and the parameters δ, ψ, and γ are the rate of time preference, the elasticity of intertemporal substitution (EIS), and relative risk aversion, respectively.

We assume that aggregate consumption growth evolves according to the following process:

$$\Delta c_{t+1} = \mu_c + \sigma_c \epsilon^c_{t+1} + J_{t+1}, \quad \text{where } \epsilon^c_{t+1} \overset{i.i.d.}{\sim} N(0,1). \tag{2.2.2}$$

That is, log-consumption growth is subject to a constant drift μ_c, an i.i.d. normal shock with a standard deviation of σ_c, and a compound Poisson jump $J_{t+1} = \sum_{j=1}^{N_{t+1}} Z_j$. Each jump Z_j has a time-invariant distribution whose moment generating function is denoted as $\Phi_Z(u) = \mathbb{E}\left[e^{uZ}\right]$. The Poisson process N_{t+1} counts the number of jumps between times t and $t+1$ and depends on time-varying intensity λ_t. Similar to Wachter (2013), this intensity is observable and follows a discrete-time version of a mean-reverting square root process:

$$\lambda_{t+1} = (1-\rho_\lambda)\bar{\lambda}_{t+1} + \rho_\lambda \lambda_t + \sigma_\lambda \sqrt{\lambda_t}\epsilon^\lambda_{t+1}, \quad \text{where } \epsilon^\lambda_{t+1} \overset{i.i.d.}{\sim} N(0,1). \tag{2.2.3}$$

Our setup exhibits two important deviations from the model of Wachter (2013). First, in line with Seo and Wachter (2018a), we allow the long-run mean of λ_t to vary over time. For parsimony, we assume that $\bar{\lambda}_{t+1}$, the long-run mean of the Poisson intensity at time $t+1$, takes a low value $\bar{\lambda}_L$ during "normal" times and a high value $\bar{\lambda}_H$ during "depression" times. By introducing a Markov-switching process s_{t+1}, which switches back and forth between the values of zero (the normal regime) and one (the depression regime), we can express $\bar{\lambda}_{t+1} = (1-s_{t+1})\bar{\lambda}_L + s_{t+1}\bar{\lambda}_H$. The transition probabilities between the two regimes are denoted as $p_{ss'} = P(s_{t+1} = s'|s_t = s)$. Moreover, we assume that the value of $\bar{\lambda}_{t+1}$ is unobservable, which clearly distinguishes our model from conventional disaster risk models. In other words, the agent has imperfect information about the stochastic long-run mean of the jump intensity process and, therefore, tries to learn about it.

A subtle yet critical distinction also exists in the calibration and interpretation of λ_t. In models with instantaneous disasters, this jump intensity represents the risk of disasters. Therefore, λ_t is calibrated to be very small, reflecting the rarity of economic disasters, and the jump size Z_j is calibrated to be large, capturing severe declines in consumption during disasters. In contrast, we do not interpret the jump intensity itself as disaster risk, but as crash risk: λ_t and Z_j are calibrated so that they represent more frequent but much less severe negative shocks. A large drop in aggregate consumption can still be generated under this calibration when the economy experiences a rare transition from the normal regime to the depression regime and remains there for an extended period of time. In this case, the jump intensity λ_t mean-reverts to a higher value, which leads to higher jump probabilities for multiple periods. While each realization of jumps is relatively small in size, it can accumulate over time and can collectively constitute a consumption disaster.

In our model, it is not just consumption that falls in a slow manner; the stock market does so as well. Even if the economy switches to the depression regime, this information is hidden from the agent; she does not know for certain whether a high value of λ_{t+1} is attributable to a transitory shock ϵ^{λ}_{t+1} or attributable to a persistent shift to the depression regime. Therefore, the agent gradually updates her belief about the current state of the economy based on Bayes' rule. Consequently, the stock market reacts to a consumption disaster in a slow manner, rather than sharply declining all at once. As we demonstrate in Section 2.4, this feature is critical in producing a reasonable level of option-based moments such as the VIX, variance risk premium, and risk premia on put-protected portfolios, which traditional disaster risk models struggle to account for.

Compared to Wachter (2013), our model has three more parameters: two transition probabilities (p_{01} and p_{10}) and an additional mean jump intensity ($\bar{\lambda}_L$). This is because we model consumption disasters as a persistent low-probability regime with negative growth, rather than negative instantaneous jumps. As discussed in Section 2.3, these extra parameters are pinned down by the three new moments that naturally arise from our model extension.[3] In this sense, they are not free parameters that we can arbitrarily choose. In fact, our model's ability to explain option-based moments comes from investors' learning, which adds no parameters. When we shut off this learning channel, our model with the same consumption process does a poor job in explaining the data, just like traditional disaster risk models (Section 2.4). In the following section, we elaborate on this learning friction, which plays a critical role in capturing the gradual progression of disasters in financial markets.

2.2.2 Learning

As mentioned earlier, we assume that the representative agent is a Bayesian learner who updates her belief about the true state of the economy s_t by observing immediate economic risk λ_t and its past history. In our model, the agent does not directly learn from observing consumption growth. This is because at time t, λ_t is observable by the agent, and it completely determines the conditional distribution of log consumption growth Δc_{t+1}. This particular modeling choice is motivated by the observation that stock market disasters predate consumption disasters (Barro and Ursúa, 2017; Muir, 2017).[4]

[3] Specifically, we calibrate p_{01}, p_{10}, and $\bar{\lambda}_L$ to match the three aspects of the data that Wachter (2013)'s model has to abstract away: (1) average disaster duration (about four years in the data), (2) total duration of disasters relative to the total period (about 7% in the data), and (3) small frequent market crashes (which occur about every other year according to the option pricing literature).

[4] In a model where investors learn about the true economic state solely based on consumption, consumption becomes a leading indicator of the stock market, not the other way around. In such a model, investors put more weight on being in the bad state when consumption falls, whereas they put more weight on being

With a slight abuse of notation, let $\lambda_{\text{-}\infty:t}$ denote the past time series of jump intensities up to time t. We define the agent's time-t belief that the economy is currently in the depression regime as

$$\pi_t \quad \equiv \quad \pi_{t|t} = P(s_t = 1|\lambda_{\text{-}\infty:t}).$$

The belief is updated to π_{t+1} when the agent observes the new jump intensity λ_{t+1} at time $t+1$. It follows from Bayes' rule and the law of total probability that the dynamics of the belief updating process can be described by

$$\pi_{t+1} \quad = \quad P(s_{t+1} = 1|\lambda_{\text{-}\infty:t+1}) = \frac{P(\lambda_{t+1}|s_{t+1} = 1, \lambda_{\text{-}\infty:t})P(s_{t+1} = 1|\lambda_{\text{-}\infty:t})}{\sum_{s\in\{0,1\}} P(\lambda_{t+1}|s_{t+1} = s, \lambda_{\text{-}\infty:t})P(s_{t+1} = s|\lambda_{\text{-}\infty:t})}.$$

Dividing both the numerator and the denominator by the numerator results in the following expression:

$$\pi_{t+1} \quad = \quad \left[1 + \underbrace{\frac{P(\lambda_{t+1}|s_{t+1} = 0, \lambda_{\text{-}\infty:t})}{P(\lambda_{t+1}|s_{t+1} = 1, \lambda_{\text{-}\infty:t})}}_{\substack{\text{(i) belief update due to a} \\ \text{new shock to } \lambda_{t+1}}} \times \underbrace{\frac{P(s_{t+1} = 0|\lambda_{\text{-}\infty:t})}{P(s_{t+1} = 1|\lambda_{\text{-}\infty:t})}}_{\substack{\text{(ii) belief update due to} \\ \text{transition dynamics}}}\right]^{-1}. \tag{2.2.1}$$

Equation (2.2.1) makes it explicit that belief updating comes from two sources. First, the agent updates her belief based on a new shock to the jump intensity. Due to imperfect information, the agent is incapable of exactly determining whether this shock originates from ϵ^{λ}_{t+1} or from $\bar{\lambda}_{t+1}$. A high value of λ_{t+1} can be due to a large transitory shock ϵ^{λ}_{t+1}

in the good state when consumption rises. As a result, consumption declines precede stock market declines. This implies that learning only from consumption does not capture the correct lead-lag relationship. Perhaps, investors learn from other signals that are informative about the true economic state (e.g., critical market events). In our setup, learning from λ_t (immediate economic risk, not disaster risk) is a parsimonious way of modeling this dimension. In Section 2.4.1, we further discuss the economic intuition behind our model using Lehman's default as an example. We also show that an increase in λ_t and the resulting belief updating cause stock prices to fall even before consumption starts to fall.

or due to a high value of $\bar{\lambda}_{t+1}$ under the depression regime. Therefore, the agent considers which scenario is more likely: if the new value of the jump intensity is so much larger than its previous value that it is more likely to be observed under the depression regime, the value of expression (i) in equation (2.2.1) becomes smaller, which, in turn, raises π_{t+1}. In contrast, if a new shock to the jump intensity is small enough to make it more likely to occur under the normal regime, the agent decreases her belief about the depression regime. Specifically, expression (i) reduces to

$$\frac{P(\lambda_{t+1}|s_{t+1}=0,\lambda_{\text{-}\infty:t})}{P(\lambda_{t+1}|s_{t+1}=1,\lambda_{\text{-}\infty:t})} = \exp\left(-\frac{(1-\rho_\lambda)\left(\bar{\lambda}_H-\bar{\lambda}_L\right)\left(\lambda_{t+1}-\rho_\lambda\lambda_t-\frac{(1-\rho_\lambda)(\bar{\lambda}_H+\bar{\lambda}_L)}{2}\right)}{\sigma_\lambda^2\lambda_t}\right) \tag{2.2.2}$$

because conditional on knowing the true regime s_{t+1}, λ_{t+1} follows a normal distribution with a mean of $(1-\rho_\lambda)\left((1-s_{t+1})\bar{\lambda}_L+s_{t+1}\bar{\lambda}_H\right)+\rho_\lambda\lambda_t$ and a variance of $\sigma_\lambda^2\lambda_t$.

We identify three parameters that control the speed of learning from new information. Equation (2.2.2) suggests that π_{t+1} becomes more sensitive to a new shock to λ_{t+1} when $\bar{\lambda}_H-\bar{\lambda}_L$ is larger, σ_λ is smaller, and ρ_λ is smaller. This is intuitive. A large magnitude of $\bar{\lambda}_H-\bar{\lambda}_L$ makes it easier to distinguish between the two regimes. Similarly, it is easier to detect the depression state when the volatility of a transitory shock σ_λ is smaller. Lastly, as ρ_λ gets closer to zero, the jump intensity process mean-reverts to its long-run mean at a faster pace, which helps the agent determine whether the jump intensity process is headed toward the long-run mean under the normal regime or the one under the depression regime.

The second source of belief updating is the transition dynamics of s_{t+1}. Even without any news from the jump intensity, the agent still updates her belief because the true state of the economy alternates between the normal and depression regimes exogenously according to the transition probabilities $p_{ss'}$. If the economy is in the normal regime (depression regime), the probability of staying in the same regime in the next period is p_{00} (p_{11}) and the probability of switching to the other regime is p_{01} (p_{10}). Reflecting these transition dynamics, the belief at time t evolves into the belief at time $t+1$ even when there is no extra information from λ_{t+1}.[5]

This is essentially what expression (ii) in equation (2.2.1) represents: given today's belief π_t, if the transition dynamics suggest that the depression regime is more likely at time $t+1$, expression (ii) becomes smaller, and thus, π_t becomes larger. It follows from the law of total probability that the time-t conditional probability of being in the depression regime at time $t+1$ equals

$$\begin{aligned}\pi_{t+1|t} \equiv P(s_{t+1}=1|\lambda_{\text{-}\infty:t}) &= p_{01}P(s_t=0|\lambda_{\text{-}\infty:t}) + p_{11}P(s_{t+1}=1|\lambda_{\text{-}\infty:t}) \\ &= 1-p_{00}+(p_{00}+p_{11}-1)\pi_t. \end{aligned} \tag{2.2.3}$$

We define this conditional probability as $\pi_{t+1|t}$ in the above equation. Note that $P(s_{t+1} = 0|\lambda_{\text{-}\infty:t})$ is $1-\pi_{t+1|t}$.

Finally, plugging equations (2.2.2) and (2.2.3) into equation (2.2.1) explicitly shows how the agent's belief evolves over time:

$$\pi_{t+1} = \left[1+\exp\left(-\frac{(1-\rho_\lambda)\left(\bar{\lambda}_H-\bar{\lambda}_L\right)\left(\lambda_{t+1}-\rho_\lambda\lambda_t-\frac{(1-\rho_\lambda)(\bar{\lambda}_H+\bar{\lambda}_L)}{2}\right)}{\sigma_\lambda^2\lambda_t}\right)\times\left(\frac{1-\pi_{t+1|t}}{\pi_{t+1|t}}\right)\right]^{-1} \tag{2.2.4}$$

Namely, the agent's future belief (π_{t+1}) is a function of the current belief (π_t) as well as the current and future values of the jump intensity (λ_t and λ_{t+1}).

It is worth highlighting how our learning friction differs from other mechanisms adopted by existing disaster risk models with imperfect information. For example, Wachter and Zhu (2019) incorporate imperfect information about the current probability of dis-

[5] In the absence of new information from λ_{t+1}, π_t mean-reverts to the belief in the steady state: $\mathbb{E}[s_t] = (1-p_{00})\,/\,(2-p_{00}-p_{11})$.

asters; Collin-Dufresne, Johannes, and Lochstoer (2016) consider parameter uncertainty about the average duration of disasters. What separates the two models from our model is that investors do not face uncertainty about the occurrence of a disaster. That is, investors can perfectly tell when a disaster begins. Thus, when a consumption disaster starts (either instantly or slowly), investors can fully incorporate the negative growth expectations into prices immediately, resulting in a counterfactually severe price decline in the market.

In fact, the learning frictions in these models make the problem worse. When a disaster occurs, investors in Wachter and Zhu (2019)'s model raise their belief about being with a higher disaster probability. This amplifies the associated risk premium, leading to even higher discount rates and steeper falls in prices. Similarly, when a disaster starts, an extra layer of uncertainty about the average disaster duration in Collin-Dufresne, Johannes, and Lochstoer (2016)'s model leads to higher discount rates and lower asset prices, compared to the case with perfect information. What our model intends to achieve through investors' learning is the polar opposite: our goal is to generate slow price reactions to consumption declines through state uncertainty surrounding disasters.[6]

2.2.3 Solving the Model

We first solve for the wealth-consumption ratio. Define $P_{c,t}$ as the price of the consumption claim, and let $pc_t = \log(P_{c,t}/C_t)$ denote the log wealth-consumption ratio. Under this notation, the log return on the consumption claim is expressed as $r_{c,t+1} = \log(1 + e^{pc_{t+1}}) -$

[6] A gradual reaction of stock prices to consumption declines can also be obtained in a model á la Veronesi (1999), in which mean consumption growth takes one of the two values (high versus low) but is unobservable. Like in our model, a transition to the "bad" state is not observable, so investors learn about this slowly from observing consumption. This learning friction plays a key role in generating slow stock market responses to persistent consumption declines. As discussed earlier, we take a different approach and model state uncertainty through a multi-layered jump intensity process (not through mean consumption growth) because investors' learning from consumption causes consumption disasters to predate financial disasters.

$pc_t + \Delta c_{t+1}$. In equilibrium, $r_{c,t+1}$ should satisfy the following Euler equation:

$$\mathbb{E}_t \left[\exp \left(\log M_{t+1} + r_{c,t+1} \right) \right] = 1,$$

where $\mathbb{E}_t$ denotes the expectation conditional on the agent's time-t information set. In Appendix A, we show that the Euler equation leads to the following recursive relation between pc_t and pc_{t+1}:

$$pc_t = \frac{1}{\theta} \left[\theta \log \delta + (1-\gamma)\mu_c + \frac{1}{2}(1-\gamma)^2 \sigma_c^2 + \lambda_t \left[\Phi_Z(1-\gamma) - 1 \right] + \log \mathbb{E}_t \left[\left(1 + e^{pc_{t+1}} \right)^\theta \right] \right] \tag{2.2.1}$$

Due to the Markov property, pc_t is a function of λ_t and π_t, namely $pc_t = pc(\lambda_t, \pi_t)$. However, a closed-form expression for this function does not exist in our model due to the nonlinearity of the learning dynamics. Following Lettau, Ludvigson, and Wachter (2008), we numerically find pc over a two-dimensional grid of λ and π. The key idea is that $pc(\lambda, \pi)$ is a fixed point of equation (2.2.1). Under the current values of pc, we can find the "new" value of pc for each set of (λ, π) by calculating the right-hand side of equation (2.2.1). Note that the conditional expectation in equation (2.2.1) is obtained by

$$\mathbb{E}_t \left[\left(1 + e^{pc_{t+1}} \right)^\theta \right] = \sum_{s \in \{0,1\}} P(s_{t+1} = s | \lambda_{\text{-}\infty:t}) \mathbb{E} \left[\left(1 + e^{pc(\lambda_{t+1}, \pi_{t+1})} \right)^\theta | s_{t+1} = s, \lambda_t, \pi_t \right],$$

where $P(s_{t+1} = 0 | \lambda_{\text{-}\infty:t}) = 1 - \pi_{t+1|t}$ and $P(s_{t+1} = 1 | \lambda_{\text{-}\infty:t}) = \pi_{t+1|t}$.[7] We continue updating the values of pc by repeating this procedure until the function converges to its fixed point.

Now we turn to a claim that pays the aggregate dividend. Following Abel (1990), we assume that the aggregate dividend is levered consumption $D_t = C_t^\phi$. We define the price

[7] Given s_{t+1}, conditioning on the agent's full information set simply reduces to conditioning on λ_t and π_t because these two variables, together with s_{t+1}, fully determine the conditional distribution of $(\lambda_{t+1}, \pi_{t+1})$.

of the dividend claim as $P_{d,t}$ and the log price-dividend ratio as $pd_t = \log(P_{d,t}/D_t)$. Then, the log return on the dividend claim (or equity) is expressed as $r_{d,t+1} = \log\left(1 + e^{pd_{t+1}}\right) - pd_t + \phi\Delta c_{t+1}$. It follows from the Euler equation that the log price-dividend ratio satisfies the following equation:

$$pd_t = \theta \log \delta + (\phi - \gamma)\mu_c + \frac{1}{2}(\phi - \gamma)^2\sigma_c^2 + \lambda_t \left[\Phi_Z(\phi - \gamma) - 1\right] \\ - (\theta - 1)pc_t + \log \mathbb{E}_t\left[\left(1 + e^{pc_{t+1}}\right)^{\theta - 1}\left(1 + e^{pd_{t+1}}\right)\right]. \quad (2.2.2)$$

See Appendix A for more details. Since the function pc is already known, equation (2.2.2) recursively characterizes the function for the log price-dividend ratio. We apply the same numerical procedure we use for the wealth-consumption ratio to solve for $pd_t = pd(\lambda_t, \pi_t)$ as a fixed point of equation (2.2.2).

2.2.4 Conditional Moments and Option Prices

Once we obtain the wealth-consumption ratio and the price-dividend ratio as functions of λ_t and π_t, various conditional moments can be computed in a semi-analytical way. Typically, calculating time-t conditional moments requires computing a time-t conditional expectation of a function with the following form:

$$F_{t+1} = F\left(\Delta c_{t+1}, \lambda_{t+1}, \pi_{t+1}, pc(\lambda_{t+1}, \pi_{t+1}), pd(\lambda_{t+1}, \pi_{t+1})\right).$$

As illustrated in Section 2.2.3, the conditional expectation of this general function can be written as

$$\mathbb{E}_t\left[F_{t+1}\right] = \sum_{s \in \{0,1\}} P(s_{t+1} = s | \lambda_{-\infty:t}) \mathbb{E}\left[F_{t+1} | s_{t+1} = s, \lambda_t, \pi_t\right]. \quad (2.2.1)$$

Conditional on $(s_{t+1}, \lambda_t, \pi_t)$, the distribution of F_{t+1} is completely characterized by the distributions of ϵ^{λ}_{t+1} and Δc_{t+1}. Specifically, under this conditioning, a shock to the jump intensity ϵ^{λ}_{t+1} pins down λ_{t+1} and π_{t+1} and, thus, pc_{t+1} and pd_{t+1}, as can be seen in equations (2.2.3) and (2.2.4). Therefore, the value of the expectation $\mathbb{E}\left[F_{t+1}|s_{t+1}=s, \lambda_t, \pi_t\right]$ in equation (2.2.1) can be found using a double integral with respect to ϵ^{λ}_{t+1} and Δc_{t+1}. Note that the conditional density of λ_{t+1} is simply the standard normal density function. In the Internet Appendix, we show that the conditional density of Δc_{t+1} can be derived in terms of the Gauss error function, which is, in turn, expressed in terms of a normal distribution function. Alternatively, this expectation can also be estimated quickly using Monte Carlo simulations.

As our first example, we calculate the risk-free rate $r_{f,t}$. The Euler equation implies that

$$
\begin{aligned}
r_{f,t} &= -\log \mathbb{E}_t\left[M_{t+1}\right] \\
&= -\theta \log \delta + \gamma \mu_c - \frac{1}{2}\gamma^2 \sigma_c^2 - \lambda_t \left[\Phi_Z(-\gamma) - 1\right] + (\theta - 1) pc_t - \log \mathbb{E}_t\left[\left(1 + e^{pc_{t+1}}\right)^{\theta - 1}\right].
\end{aligned}
$$

The expression $\left(1 + e^{pc_{t+1}}\right)^{\theta-1}$ is a special case of F_{t+1} in which the function does not depend on Δc_{t+1}. Hence, its conditional expectation is computed as a one-dimensional integral with respect to ϵ^{λ}_{t+1}. Similarly, the time-t expected equity return is also calculated as a one-dimensional integral because

$$
\mathbb{E}_t\left[r_{d,t+1}\right] = \mathbb{E}_t\left[\log\left(1 + e^{pd_{t+1}}\right)\right] - pd_t + \phi \mu_c + \phi \mu_Z \lambda_t,
$$

where $\mu_Z = \Phi'_Z(0)$ is the mean jump size.

Equation (2.2.1) allows us to calculate not only the first moments of returns, but also their higher moments. For instance, the conditional equity return variance is calculated as:

$$\mathrm{Var}_t\left(r_{d,t+1}\right) = \mathbb{E}_t\left[\left(\log\left(1+e^{pd_{t+1}}\right) - pd_t + \phi\Delta c_{t+1} - \mathbb{E}_t\left[r_{d,t+1}\right]\right)^2\right].$$

That is, it is possible to semi-analytically calculate any moment by first expressing it in terms of a time-t conditional expectation and applying the general formula derived in equation (2.2.1).

2.3 Calibration

Table 2.1 reports the model parameters, which are calibrated at a monthly frequency. Our calibration of the preference parameters is standard. We set risk aversion γ to 5, consistent with Mehra and Prescott (1985) and subsequent papers in the equity premium puzzle literature. The EIS ψ is chosen as 1.5, implying that the representative agent prefers early resolution of uncertainty (see, e.g., Bansal, Kiku, and Yaron, 2012). The time discount factor δ is 0.999, which is equivalent to an annual rate of time preference of 1.2% in Wachter (2013).

We calibrate the transition dynamics, characterized by p_{10} and p_{01}, by targeting the duration of historical consumption disasters in the U.S. The average duration of consumption disasters in the data is roughly four years. This suggests that it is reasonable to choose p_{10}, the transition probability from the depression regime to the normal regime, to be $0.25\Delta t$, where $\Delta t = 1/12$. Then, we set p_{01}, the transition probability from the normal regime to the depression regime, to 2% per annum ($0.02\Delta t$ monthly) so that the unconditional probability of a year being in the depression state $p_{01}/(p_{01} + p_{10})$ is approximately 7%. This is consistent with what we observe in the U.S. consumption time series: over the last 186 years (i.e. 1834-2019), the U.S. experienced three disasters with a total duration of 13 years, indicating that the fraction of disaster periods relative to the total years considered is $13/186 = 6.99\%$. Note that Barro and Ursúa (2008) adopt the same approach to calibrate these switching probabilities, yet with international data, which implies a higher disaster probability (3.63%), mainly due to a larger portion of disaster years (12%). By focusing on U.S. consumption disasters, we pick p_{01} conservatively in an attempt to make it clear that our model's success is not driven by an overstatement of disaster likelihood.[8]

Now we turn to the calibration of jump risk in the model. We assume that the size of each log consumption jump Z_j follows the negative of an exponential distribution with mean μ_Z. Note that in our model, dividends are defined as levered consumption, which causes dividends and equity prices to drop by $\phi\mu_Z$, on average, in response to jumps in consumption. This relation enables us to calibrate μ_Z using the results from the option pricing literature, rather than relying on the historical time series of consumption. Consistent with Eraker (2004), we target -6% stock market jumps that occur every other year on average during non-disaster periods.[9] We choose $\phi = 3$, which is standard in the literature (see, e.g, Bansal and Yaron, 2004). As a result, the mean consumption jump size μ_Z is calibrated as -2%.

The persistence of the jump intensity process is determined by the autoregressive coefficient ρ_λ in equation (2.2.3). We follow Wachter (2013) and set the parameter equal to $1 - 0.08\Delta t$, which corresponds to a mean reversion rate of 8%. As discussed above, we set $\bar{\lambda}_L = 0.5\Delta t$ so that, on average, the equity market experiences negative jumps once in two years under the normal regime. This leaves us two free parameters: $\bar{\lambda}_H$ and σ_λ. We

[8] In our model, shifting to the depression state does not necessarily mean an occurrence of a disaster; the economy can quickly switch back to the normal state, resulting in a consumption drop that is less than 10% in the peak-to-trough sense. Therefore, the unconditional probability of disasters is even lower than 2% under our calibration.

[9] In the option pricing literature, the estimated jump size ranges from -2% to -10%. Typically, when the estimated jump size is small, the frequency of jumps is estimated to be high (2-3 times a year). When the estimated jump size is large, the frequency of jumps is low (once in 2-3 years). Examples include, but are not limited to, Eraker, Johannes, and Polson (2003), Eraker (2004), and Broadie, Chernov, and Johannes (2007).

calibrate these two parameters to match the equity premium and the stock market volatility in the data. Specifically, under our calibration, the long-run mean of the jump intensity increases tenfold when the economy falls into the depression state ($\bar{\lambda}_H = 5\Delta t$). Lastly, based on the calibrated jump process, we pick the values of μ_c and σ_c so that the model matches the postwar mean and volatility of log consumption growth.

2.4 Model Results

In this section, we present and discuss the quantitative implications of our model. Section 2.4.1 describes the model mechanism based on a simulated sample path and examines the properties of the equilibrium solution. Section 2.4.2 evaluates our model's implications for standard asset pricing moments as well as the VIX, variance risk premium, and put-protected portfolio premia. The model-implied interest rate term structure is also examined. Finally, Section 2.4.3 delineates the characteristics of disaster realizations in our model and compares them to the data.

2.4.1 Inspecting the Model Mechanism

How does our model generate slowly unfolding disasters in consumption and in equity prices? In Figure 2.2, we illustrate the mechanism using a sample path of the model that includes a consumption disaster. Panel A plots the path of the true state s_t and the agent's belief π_t. Once the economy switches to the depression state with $s_t = 1$, instantaneous jump risk represented by λ_t enters an upward trend (Panel B). This is because λ_t now mean-reverts to a higher value under the depression regime, compared to the normal regime.

The agent learns about s_t from the historical path of λ_t. Yet, the belief dynamics in Panel A indicate that learning is slow and imperfect. It takes more than a year for the agent to raise her subjective probability of being in the depression state to a level close to one. Furthermore, the belief is not stable as λ_t is a noisy signal and sometimes points in the "wrong" direction. For instance, the perceived depression state probability drops below 50% around month 35 when λ_t is hit by a few consecutive negative shocks ϵ_t^{λ}.

Panel C of Figure 2.2 plots the path of annual consumption. Because it takes time for λ_t to increase due to persistence, the consumption disaster starts later than the onset of the depression state. Lastly, Panel D plots the trajectory of the aggregate equity price along the path. The financial market disaster starts prior to the consumption disaster. The trough of the equity price is observed toward the end of the consumption disaster so that the peak-to-trough duration of the disaster in the equity market is close to the duration of the consumption disaster. This particular disaster represents a cumulative equity index decline of about 80%, similar to the Great Depression, while the consumption decline is about 25%. Once the economy exits the depression state, equity prices start recovering in a slow manner. This is because not only does the jump intensity follow a persistent process, but the speed of learning is also slow for high values of λ_t.

The driver of these realistic equity price dynamics during disasters is the presence of imperfect information. At the beginning of and throughout the disaster period, the representative investor does not know the true state of the economy with certainty. This uncertainty has profound implications for equity prices: while it is a source of risk itself, it drives the extent to which the agent incorporates disaster risk into pricing. If the agent were fully aware of being at the start of a consumption disaster, equity prices would fully react right away. Through the learning mechanism, our model breaks this unrealistic link between a shift to the disaster state and an instant price reaction.

The U.S. economic history features two prime examples in which investors did not immediately recognize the state of the economy at the onset of exceptionally bad periods: the

Great Depression and the Great Recession. During the Great Depression, the cumulative stock market decline of over 80% spans roughly four years. At the beginning of the Great Depression in 1929, however, the stock market decline remained much smaller. The reason is arguably because investors did not know *ex ante* that the recession they were experiencing would become as severe as it turned out to be *ex post*. The Great Recession is an opposite example. At Lehman's default in September 2008, the market feared the possibility of a severe and lengthy economic downturn similar to the Great Depression, which was not quite the case. In sum, at the beginning of a recession, investors face uncertainty about how long and how bad economic downturns will be.

Our model is able to capture this phenomenon. A significant increase in consumption risk in our model can arise due to either a large transitory shock or a shift to the depression regime. Moreover, even if the economy transitions to the depression regime, it is possible to quickly switch back to the normal regime. These features create uncertainty about the duration and severity of economic recessions, which the representative investor incorporates into equity prices.

Figures 2.3 and 2.4 present model-implied quantities as a function of one of the two state variables, λ_t or π_t.[10] Panel A of Figure 2.3 illustrates that the log wealth-consumption ratio and the log price-dividend ratio have an almost linear relationship with λ_t, as in conventional variable disaster risk models such as Wachter (2013). Panel B of Figure 2.3 shows that these valuation ratios decline as π_t increases, but so does the rate of decrease, making the two graphs convex. This convexity appears because uncertainty about the true state of the economy, which the agent dislikes under the model's preference configuration, is highest when π_t is 50%.[11] As a result, for small values of π_t, the valuation ratios decrease with

[10] When examining the relationship between each state variable and a model quantity, we fix the other state variable at its median value.

[11] Note that $\mathrm{Var}_t(s_t) = \pi_t(1 - \pi_t)$ is maximized when π_t equals 0.5.

π_t at a relatively faster speed as (i) the risk of being in the bad regime and (ii) uncertainty about the true economic state both rise. In contrast, for high values of π_t, the valuation ratios decrease with π_t at a relatively lower speed because (ii) reduces, partially offsetting the effect from a rise in (i).

Figure 2.4 plots the risk-free rate, equity premium, and conditional equity return volatility as a function of λ_t or π_t. Panels A, C, and E show that for medium to high values of λ_t, the behaviors of asset prices in our model with respect to λ_t are similar to those in conventional disaster risk models. That is, as λ_t increases, the risk-free rate decreases due to the precautionary savings motive. The conditional equity premium and return volatility increase due to the heightened risk of a joint decline in consumption and dividends as well as higher conditional volatility of λ_t.

However, there is another source of variation that plays a significant role, especially when λ_t is low. As can be seen in equation (2.2.4), λ_t affects the conditional volatility of π_{t+1}. Specifically, as λ_t decreases, the speed of learning becomes faster, which raises the conditional volatility of π_{t+1} by making it more responsive to shocks to the jump intensity process. This channel pushes the risk-free rate downward and the conditional equity premium and return volatility upward, as the agent dislikes higher uncertainty about the future belief π_{t+1}. For small values of λ_t, this effect dominates the effect described in the previous paragraph, producing U-shaped/hump-shaped patterns of the risk-free rate, conditional equity premium, and return volatility in Panels A, C, and E.

Panels B, D, and F of Figure 2.4 illustrate the behaviors of the risk-free rate, equity premium, and conditional equity return volatility with respect to π_t. Under our calibration, π_t remains fairly close to zero most of the time because the depression state is unlikely. In such cases, an increase in π_t, which implies a higher perceived probability of being in the depression state, leads to a decrease in the risk-free rate and an increase in the equity premium and return volatility, like in conventional models. However, if π_t moves significantly away from zero, we observe U-shaped/hump-shaped patterns. As discussed above, this is because uncertainty about the hidden state of the economy is highest when π_t is 50%.

We also observe that in Panels B, D, and F, the risk-free rate dips and the equity premium/return volatility peaks at a value of π_t that is smaller than 50%. For instance, the conditional equity premium peaks around $\pi_t = 30\%$. Why is this the case, given that uncertainty about the hidden state, namely $\mathrm{Var}_t(s_t) = \pi_t(1 - \pi_t)$, peaks at 50%? There are two additional forces at play here. On the one hand, an increase in π_t raises the perceived probability of being in the depression state and leads to a higher equity premium. On the other hand, an increase in π_t also implies a lower sensitivity of the price-dividend ratio with respect to π_t (as can be seen in Figure 2.3) and leads to a lower equity premium. While the first effect pushes the peak of the curve to the right above 50%, the second effect pushes the peak of the curve to the left below 50%. Under our calibration, the latter dominates the former, and, as a result, the peak of the equity premium is located at $\pi_t < 50\%$. A similar argument applies to the risk-free rate and the return volatility.

2.4.2 Asset Pricing Moments

Simulation Procedure and Identification of Disasters

We simulate 70-year-long samples from our model at a monthly frequency and compare the resulting model moments to their data counterparts in the postwar U.S. data. For some moments that involve option prices, we use 30-year-long simulated paths instead, as the options data are only available from 1990 to 2019. We juxtapose the asset pricing implications of our model with those of Wachter (2013)'s variable disaster risk model, referred to as the instantaneous disaster risk model hereafter.[12] To investigate the role of learning in isolation, we also compute the results from the special case of our model with perfect information and report them in the Internet Appendix.

We first simulate the consumption process given in equations (2.2.2) and (2.2.3), and the endogenous belief process given in equation (2.2.4).[13] Once λ_t and π_t are simulated, we find the corresponding paths of the wealth-consumption and price-dividend ratios, the risk-free rate, the equity premium, the conditional return volatility, and option prices by interpolating their values over the grid. For each simulation, we compute our moments of interest using the simulated time series of the model quantities.

This procedure is repeated 10,000 times so that we obtain the model-implied distribution of each moment. Since the postwar U.S. data do not include any macroeconomic disasters, we report the results from no-disaster samples. No-disaster samples constitute 57% of the entire simulated samples in the case of 70-year simulations and 77% in the case of 30-year simulations. We also investigate the population properties of the model by simulating a long path of 1,000,000 years.

A distinguishing feature of our approach is in the identification of disaster periods. In conventional disaster risk models, an occurrence of a disaster is identified as soon as a Poisson process jumps or as soon as the economy enters into a disaster state, depending on the model. In contrast, disaster periods in our model are not fully identified *ex ante* using a state variable, but only *ex post*.

The reason is twofold. First, entering the depression state with $s_t = 1$ does not ensure that the economy will remain in the depression state long enough to lead to a large consumption decline; it is possible for the economy to quickly switch back to the normal

[12] To facilitate comparison, we implement the model of Wachter (2013) in discrete time with monthly intervals.

[13] In each simulation, we set the starting values of λ_t and π_t to their long-run means and apply a 100-year burn-in period.

state. Furthermore, the true economic state is not directly observable, and the agent cannot perfectly distinguish the onset of a disaster from an increase in transitory risk. We establish full consistency between the model and the data by following Barro and Ursúa (2008) and identifying macroeconomic disasters in our model as peak-to-trough consumption declines that are larger than 10%.

As a side note, the way in which disasters are identified in the data and in our model is reminiscent of the identification of stock price bubbles in Pastor and Veronesi (2009b). The main intuition behind their model is that technological revolutions can only be identified *ex post*: investors living through a technological revolution do not know *ex ante* whether the adoption of a new technology will eventually happen. In this sense, their model and our model share a similarity, even though the underlying economic mechanisms behind the two models are quite different. In both models, investors are not fully aware of the true underlying driver of the economy, and the resulting learning dynamics generate distinctive patterns in asset prices that can only be identified *ex post*.

Standard Asset Pricing Moments

Table 2.2 reports the standard asset pricing moments from our model (Panel A) and from the instantaneous disaster risk model (Panel B). A general observation from comparing Panels A and B is that both models perform fairly well in terms of explaining the high equity premium, the high stock market volatility, and the low and smooth risk-free rate, while matching the low consumption growth volatility in the postwar U.S. data. The model with time-varying instantaneous disaster risk is designed to account for these aspects of the data, providing a joint solution to the equity premium, risk-free rate, and stock market volatility puzzles. While our model operates through a different mechanism, it is reassuring that it can address these puzzles.

A striking difference between our model and the model with instantaneous disasters is in the discrepancy between asset pricing moments in no-disaster samples versus in population. For instance, the median no-disaster consumption growth volatility in our model is 1.80% and its population value is 2.50%. The difference is much larger in the instantaneous disaster risk model: a median of 1.62% in no-disaster samples versus 5.47% in population. This gap is also manifested in the volatilities of the equity return and the risk-free rate.

Why does the instantaneous disaster risk model exhibit a larger discrepancy between no-disaster sample moments and population moments? In such a model, consumption, dividends, and equity prices all fall within a unit period (i.e. one month) during a disaster. As a result, volatilities and other higher moments are greatly amplified due to the instantaneous nature of disasters. In contrast, our model generates disasters that unfold slowly over multiple periods of time both in consumption and prices, consistent with the data. This brings the asset pricing moments in population closer to those in no-disaster samples.

Now we turn to the predictability of returns by the price-dividend ratio. Panel B of Table 2.3 illustrates that the instantaneous disaster risk model produces the empirical predictive relation between equity returns and the price-dividend ratio. This is possible because the price-dividend ratio is decreasing in disaster risk whereas the conditional equity premium is increasing in disaster risk, resulting in a negative predictive relation between the price-dividend ratio and future excess returns.

We find that our model also generates predictable returns, consistent with the data. However, the relation between the price-dividend ratio and the equity premium is not as straightforward: as illustrated in Figures 2.3 and 2.4 and explained in Section 2.4.1, the price-dividend ratio is monotonically decreasing in λ_t and π_t whereas the conditional equity premium is not. Therefore, it is a question of calibration whether the model would spend enough time in a region where the correlation between the price-dividend ratio and the conditional equity premium is negative. Panel A of Table 2.3 shows that this is indeed the case: the price-dividend ratio predicts future excess equity returns with a negative coefficient. The magnitude of model-implied coefficients and R^2 values are close to the data.

Lastly, Table IA.1 in the Internet Appendix demonstrates that the two models of interest account for the absence of consumption growth predictability in the data. Note that in both models, expected consumption growth exhibits persistent variation due to time variation in jump risk. In the case of the instantaneous disaster risk model, however, this does not lead to consumption growth predictability because no-disaster samples, by definition, do not contain jump realizations (i.e. disasters). In contrast, in our model, disaster samples can potentially include a series of jumps that are not large enough to constitute a disaster. Despite this possibility, we discover that the predictability of consumption growth by the price-dividend ratio is weak, containing the postwar U.S. data within the confidence bands of our model.

All in all, our model and Wachter (2013)'s instantaneous disaster risk model are capable of explaining standard asset pricing facts. In what follows, we discuss additional moments that can sharply contrast these two models: the VIX and the variance risk premium (Section 2.4.2), the risk premia on put-protected portfolios (Section 2.4.2), and the term structure of interest rates (Section 2.4.2).

The VIX and the Variance Risk Premium

An important distinction between our model and the instantaneous disaster risk model is the source of the equity premium, which brings about crucial empirical implications. In our model, the agent fears disasters because disasters are prolonged depression periods that are slowly revealed over time. On the contrary, the instantaneous disaster risk model attributes the equity premium almost entirely to the risk of a joint tail event in consumption, dividends, and equity prices. We argue that an unrealistically quick financial market response to macroeconomic disasters significantly distorts the model implications for short-term contingent claim prices.

We start with the model implications for implied variance and the variance risk premium. Following Drechsler and Yaron (2011), we calculate one-month implied variance (IV_t) and the variance risk premium (VRP_t) as follows:

$$\begin{aligned} IV_t &= \mathbb{E}_t^Q[\mathrm{Var}_{t+1}^Q(r_{d,t+2})] \\ VRP_t &= \mathbb{E}_t^Q[\mathrm{Var}_{t+1}^Q(r_{d,t+2})] - \mathbb{E}_t[\mathrm{Var}_{t+1}(r_{d,t+2})], \end{aligned}$$

where Q denotes the risk-neutral measure. Consistent with the standard convention in the literature, both quantities are expressed in monthly percentage squared terms. Note that under this unit, the VIX is equal to $\sqrt{12 \times IV_t}$.

Table 2.4 compares the resulting implied variance and variance risk premium from the two models with their data counterparts.[14] Panel B shows that the instantaneous disaster risk model significantly overstates the level and volatility of implied variance as well as those of the variance risk premium.[15] For instance, the average implied variance in the data between 1990 and 2019 is 35.15. The corresponding value in the median no-disaster sample is 170.92 in the instantaneous disaster risk model.

This discrepancy is deeply rooted in how disasters unfold in financial markets. When the economy enters the disaster state with a, say, 20% drop in consumption, equity prices

[14] In the data, implied variance, which is the squared VIX multiplied by 12, is computed as the value of a portfolio of S&P 500 index option prices. We obtain the time series of implied variance and the variance risk premium from Hao Zhou's website. For details, see Zhou (2018).

[15] To mitigate the issue of extremely high values of the VIX in disaster risk models, Dew-Becker, Giglio, Le, and Rodriguez (2017) and Seo and Wachter (2018c) adopt an ad-hoc approach and apply an upper bound on the maximum instantaneous decline in equity prices during disasters. There is no need to make such an assumption in our model as it generates an endogenously slow decline in prices, consistent with the data.

immediately fall by 60% (20% times a leverage parameter of 3) in the model. As a result, the expected variance within a month is extremely high since the entire 60% decline in equity is expected to happen during that month. This is inconsistent with the way in which disasters actually develop in the real world, and results in unrealistically high values of the expected variance.[16]

The variance risk premium is also entangled in the same issue. States with abrupt disasters are associated with extremely high marginal utility, which leads to a very high average variance risk premium in the instantaneous disaster risk model. The average variance risk premium in the data is 15.51, but its counterpart in the model is much higher: the median of the average variance risk premium in the model is 139.00, and the lower end of the confidence band is 54.42, which is still substantially larger than the data value.

Our model resolves this inconsistency by generating realistic disaster dynamics in financial markets. While disasters still occur, they unfold over multiple periods, resulting in a lower expected variance. Panel A of Table 2.4 reveals that our model implies realistic levels of implied variance and the variance risk premium. Specifically, the medians of the average implied variance and variance risk premium in no-disaster samples are 40.23 and 16.35, respectively, which are very close to the data. While the volatilities of implied variance and the variance risk premium are unrealistically high in the instantaneous disaster risk model, our model performs well in these aspects of the data: 39.61 for the volatility of implied variance (versus 32.65 in the data) and 17.95 for the volatility of the variance risk premium (versus 19.94 in the data).

In the Internet Appendix, we also consider the special case of our model with perfect information in an attempt to highlight the importance of learning. This model features slow

[16]For instance, the largest one-month decline in the aggregate equity price during the Great Depression is 30% while the cumulative decline is 86%. Assuming that the cumulative decline occurs within a month causes a major distortion in the underlying price process and implied one-month option prices.

consumption disasters but not slow financial disasters; since the transition to the depression state is directly observable, the agent can fully react to this bad news. While this model produces a lower level of implied variance and the variance risk premium compared to the instantaneous disaster risk model, the model values are still too high compared to the data (Panel A of Table IA.2). That is, the instantaneous reaction of prices to the depression state still results in unrealistically high values of short-term variance and its market price. We conclude that information frictions and learning dynamics are crucial to generating a realistically slow reaction of prices to macroeconomic disasters, which translates into realistic values of expected short-term variance and the risk premium attached to it.

Given the slow nature of disasters in our model, it is worth revisiting the criticism against the rare disaster pricing mechanism regarding the timing of shocks to the stochastic discount factor and returns. Constantinides (2008) points out that although the Euler equation is based on concurrent moves in consumption and asset prices, Rietz-Barro type models implicitly violate this premise by treating slow consumption disasters in the data as one-shot shocks. This simplifying assumption may artificially amplify risk premia in these models due to a counterfactually strong connection between consumption declines and financial markets.

Our analysis shows that the assumption of instantaneous disasters is indeed not innocuous when one considers the consistency between the equity premium and the variance risk premium. As discussed above, the instantaneous disaster risk model implies a reasonable equity premium yet a too large variance risk premium. Assuming slowly unfolding consumption disasters alone does not resolve this issue: in the special case of our model with perfect information, the variance risk premium is still high.

The solution lies in modeling a reasonable joint behavior of consumption and equity prices during disasters. Our model achieves an equity premium close to the data due to the pricing of the time variation in the prospect of future disasters. The variance risk premium is realistically low because disasters are not identifiable *ex ante*, preventing equity prices from plummeting too abruptly. Essentially, addressing the criticism of Constantinides (2008) is at the heart of reconciling the equity premium and the variance risk premium under the rare disaster mechanism: both can be resolved by introducing slowly unfolding consumption and financial disasters.

Risk Premia on Put-Protected Portfolios

Another testable implication of disaster risk models concerns the source of the equity premium. Welch (2016) provides an intuitive empirical exercise to test whether the equity premium can be explained by the possibility of rare disasters. If the equity premium is due to the likelihood of very large jumps in the equity market, a portfolio that is protected against the risk of large jumps should not carry a significant premium. Welch (2016) builds a portfolio that consists of the S&P 500 index and a one-month out-of-the-money (OTM) put option with 85% moneyness and finds that this portfolio still earns a significant premium close to the full equity premium. This result leads him to conclude that a major portion of the equity premium cannot be attributed to rare disaster events.

Our framework provides a way to reconcile Welch (2016)'s evidence with the presence of rare disasters in the data. Before we discuss the pricing mechanism, we point out that one-month put options with moneyness 85% would not have provided significant protection against the most prominent disaster, the Great Depression. Panel B of Figure 2.1 shows that a put-protected strategy with 85% moneyness would have lost 80% throughout the Great Depression, even ignoring the cost of purchasing put options. Given that the stock market experienced a cumulative decline of 86%, the insurance against a monthly drop larger than 15% would not have served as an effective hedge against the disaster. From this episode, we can understand that modeling slowly unfolding disasters in equity prices, rather than assuming instantaneous price reactions, is critical in explaining the risk premia on put-protected portfolios.

Let $O_t(K)$ denote the one-month OTM put option price with strike price K. Then, the return on a put-protected portfolio is calculated as $\frac{\max\left(P_{d,t+1},K\right)+D_{t+1}}{P_{d,t}+O_t(K)}$. This expression is intuitive. The denominator represents the initial investment in the portfolio at the beginning of each month at time t, which is the sum of the equity value $P_{d,t}$ and the option price $O_t(K)$. The numerator represents the payoff from this portfolio after a month at time $t+1$, which consists of (i) the equity value $P_{d,t+1}$, (ii) the dividend payment D_{t+1}, and (iii) the put option payoff $\max(K - P_{d,t+1}, 0)$. By dividing both the denominator and the numerator by $P_{d,t}$, we express the log return on the put-protected portfolio in terms of moneyness $k = K/P_{d,t}$:

$$r^k_{d,t+1} = \log\left(\frac{\max\left[P_{d,t+1}/P_{d,t}, k\right] + D_{t+1}/P_{d,t}}{1 + O^n_t(k)}\right).$$

The expression for the normalized option price $O^n_t(k)$ is provided in the Internet Appendix.

Table 2.5 presents the resulting risk premia in the model and in the data for various moneyness values k = 0.75, 0.80, 0.85, and 0.90. To evaluate the relative performance among different models, we report the results in the following two ways: (i) the difference between the full equity premium and put-protected portfolio premium, $\mathbb{E}[r_d - r^k_d]$, and (ii) the put-protected portfolio premium as a fraction of the full equity premium, $\frac{\mathbb{E}[r^k_d - r_f]}{\mathbb{E}[r_d - r_f]}$.[17]

[17] The moments of deep OTM option returns are extremely noisy even in a long sample (Broadie, Chernov, and Johannes, 2009). In most cases, these options expire out of the money and result in -100% returns. If they expire in the money, however, their returns are typically massive because deep OTM options are cheap. Our analysis does not inherit this problem. Unlike the return on a naked option position, the return on a put-protected portfolio well behaves, just like an equity position with no options. This is because OTM put options occupy only a tiny portion of the entire portfolio construction cost. In addition, the payoff from put options limits the loss when there is a sharp decline in the equity price. In fact, the volatility of each put-protected portfolio is about 18% in the data across all of the four moneyness values, which is very close to the stock market volatility.

In the data, the premium difference between the equity index and the put-protected portfolio with $k = 0.85$ is approximately 2% per annum, which is consistent with the evidence from Welch (2016). The premium difference further shrinks to 0.72% when the moneyness value is 0.75. As can be seen from Table 2.5, this put-protected portfolio bears 90% of the entire equity premium, suggesting that monthly price declines that are larger than 25% are not the main source of the equity premium, at least over a one-month horizon.

As anticipated by Welch (2016), the instantaneous disaster risk model has difficulty in capturing the high premium on put-protected portfolios: the put-protected portfolio with $k = 0.75$ earns a premium that is 58% of the full equity premium in the median no-disaster sample (Panel B). That is, the instantaneous disaster risk model relies too much on extremely large shocks to explain the equity premium, and this is inconsistent with the data.

Now we turn to the results from our model with slow consumption and financial disasters. Panel A of Table 2.5 demonstrates that our model performs well in explaining put-protected portfolio premia across various moneyness values. For instance, in the median no-disaster sample, the risk premium on the put-protected portfolio with moneyness $k = 0.75$ is 0.72% lower than the full equity premium (versus 0.72% in the data), occupying 90% of the full equity premium (versus 90% in the data). This is due to realistic price dynamics during disasters in our model, which, in turn, results in put-protected portfolio returns that are consistent with the data. In our model, it is the prospect of slow and prolonged disasters in equity prices that gives rise to the equity premium, not the prospect of an extremely severe crash in the equity market.

Note that our results on the VIX and put-protected portfolios implicitly suggest that the model's option pricing is sensible.[18] Nevertheless, it is still worth checking implied

[18]This is because the VIX is expressed as a weighted average of option prices, mostly driven by out-of-the-money puts; each put-protected portfolio consists of the S&P 500 and out-of-the-money puts.

volatilities, as they are more direct measures of our model's ability to price options. In the Internet Appendix, we show that the model indeed well explains the average implied volatility skew in the data. As can be seen in Figure **??**, the data values fall within the model's 90% confidence band.

Term Structure of Interest Rates

Another well-known shortcoming of dynamic models with instantaneous disasters is that long-term real bond yields are remarkably low. As shown in Figure 2.5 (dashed yellow line), a 30-year default-free zero-coupon bond has an average yield below -20% under the instantaneous disaster risk model. This is because long-term real bonds are extremely valuable under the presence of instantaneous disasters. These securities provide a hedge against rare-yet-possible scenarios with multiple disasters, where consumption suffers a large decline and investors' marginal utility surges. Since consumption instantly crashes at disaster arrivals, the state prices of such scenarios rapidly increase with horizon. This results in extremely high long-term real bond prices, which correspond to severely negative yields.

Backus, Chernov, and Zin (2014) call this phenomenon "horizon dependence." They define a diagnostic measure for horizon dependence as the difference in entropies over horizons of 10 years and one month. Since this measure is equivalent to the mean monthly yield spread between one-month and 10-year maturity real bonds, Backus, Chernov, and Zin (2014) argue that it should be bounded between -0.1% and 0.1%. That is, in order for a model to have reasonable predictions for the interest rate term structure, the long-term dispersion of the pricing kernel (i.e. long-term entropy) needs to be quite limited. They find, however, that a stochastic intensity model similar to Wachter (2013) implies too much horizon dependence. In fact, the instantaneous disaster risk model examined in this paper has horizon dependence of 0.1569%, which falls outside of the diagnostic bounds proposed by Backus, Chernov, and Zin (2014). Note that the problem would have been much more severe if the diagnostic bounds were defined in terms of longer-horizon entropies.

As clearly illustrated in Figure 2.5 (solid blue line), our model offers a significant improvement on the term structure dimension. The term structure of interest rates is nearly flat, albeit still downward sloping. Unlike the instantaneous disaster risk model, the term structure slope remains small over all horizons. As a result, average yields are positive up until 15 years. Beyond that, yields become negative, but their magnitudes are still modest: the 30-year yield is only about -2.41%. Behind this improvement is the slow progression of disasters. In our model, it takes a long time for disasters to unfold, unlike the instantaneous disaster risk model. Therefore, the likelihood of experiencing multiple disasters stays quite small even in very long horizons. Thus, the state prices of extreme scenarios with multiple disasters increase at a much moderate pace as horizon increases. In other words, our model exhibits lower horizon dependence. In fact, our model's horizon dependence is 0.0963%, which falls within Backus, Chernov, and Zin (2014)'s diagnostic bounds.[19]

In sum, the extremely downward sloping term structure of interest rates under traditional dynamic disaster risk models is another artifact of modeling disasters as instantaneous jumps. Our analysis shows that modeling slowly unfolding disasters in consumption is critical in generating more realistic term structure implications.

[19] As another diagnostic measure, Backus, Chernov, and Zin (2014) argue that the one-month entropy of the pricing kernel should exceed 1%. Under our calibration, the model's one-month entropy is 0.8580%, barely missing the lower bound. This can be improved by slightly raising risk aversion, as the entropy is highly sensitive to it. The one-month entropy of the instantaneous disaster risk model is 0.5534%.

2.4.3 Model-Implied Disaster Statistics

The size distribution of consumption disasters, constructed by Barro and Ursúa (2008) using the peak-to-trough approach, plays an important role in the calibration of disaster risk models. Panel A of Figure 2.6 shows the empirical distribution of consumption declines during disasters. In typical models with instantaneous disasters, this distribution is used as a time-invariant jump size distribution (see Barro and Jin, 2011).

In our model, the size distribution of consumption disasters is not an input, but an important outcome. We do not assume that disasters occur exogenously following a certain size distribution; rather, disasters occur slowly over time as a series of negative shocks to consumption accumulates. Therefore, the size distribution of consumption disasters endogenously arises as an implication of the model.

To examine the severity of consumption disasters under the model, we simulate the consumption process for 1,000,000 years and identify disasters using the peak-to-trough approach. Panel B of Figure 2.6 displays the model-implied disaster size distribution in population. We can observe that this implied distribution is fairly close to its empirical counterpart in Panel A. While the average consumption decline during disasters is 21% in the data, it is 18% in our model. Furthermore, our analysis shows that the model accounts for the average duration of consumption disasters: 4.5 years in the model versus 4.1 years in the data. In sum, we conclude that our model produces consumption disasters that are consistent with historical consumption data.

A key feature of our model is that not only consumption but also equity prices fall gradually in periods of disasters. We show that generating realistic equity price dynamics during disasters is essential in explaining a number of empirical patterns, as discussed in Sections 2.4.2 and 2.4.2. We investigate whether the average size and duration of stock market disasters in the model are indeed comparable to the data. Following Barro and

Ursúa (2017), we identify stock market disasters as peak-to-trough price declines that are larger than 25%. In our model, stock market disasters exhibit a 46% cumulative decline and a 2.5-year duration, on average. These numbers are close to their data counterparts: 46% and 3.2 years.[20] Hence, we confirm that our model's characterization of financial disasters is in line with the data.

2.4.4 Recent Major Crises in the Stock Market

The main takeaway from our results thus far is that our model can explain the unconditional moments of option-based quantities, such as the VIX (equivalently, implied variance) and variance risk premium, over the last three decades. This recent sample, during which options data exist, does not include any consumption disasters, defined as peak-to-trough declines that are larger than 10% (Barro and Ursúa, 2008). However, we do observe severe stock market declines in the same period. Such events have clear counterparts in our model: stock market disasters without a consumption disaster. To see if our model can rationalize the crisis patterns in the data, we study the model's conditional dynamics when the economy experiences a large stock market decline that is not accompanied by a large consumption decline. Here, we briefly summarize our main results; the quantitative details of our analysis are relegated to the Internet Appendix.

In the period of 1990-2019, we identify two prominent stock market disasters: the dot-com bubble burst and the Great Recession (Table IA.3). To obtain the model counterparts, we collect 19,214 stock market disasters out of 30-year-long simulation paths without consumption disasters (Section 2.4.2). We find that the model-implied stock market dis-

[20] Barro and Ursúa (2017) report 71 cases of stock market disasters that are associated with consumption disasters (a 53% average decline and a 3.8-year average duration), and 161 cases of stock market disasters that are not associated with consumption disasters (a 43% average decline and a 2.9-year average duration). By aggregating these two types of events, we obtain an average decline of 46% and an average duration of 3.2 years for all historical stock market disasters.

asters well capture the first and second moments of implied variance and the variance risk premium during the two crises, containing the data values within the model's confidence bands. In fact, the model creates a path that resembles the two crisis episodes (Figure IA.2).

In February and March 2020, global financial markets witnessed an unprecedented economic shock associated with the outbreak of the coronavirus disease, the so-called COVID-19 crisis. What makes the COVID-19 crisis so unique compared to the aforementioned crises is its extremely short-lived stock market disaster. The S&P 500 decline over the two months (measured month-end over month-end) is 19.87%, which is actually smaller than the 25% threshold of Barro and Ursúa (2017). This is because the market rapidly recovered toward the end of March. If the decline is measured over a one-month sub-period from February 21st to March 23rd (with 21 trading days), the decline becomes much larger with 32.97%. Hence, accounting for an event like the COVID-19 crisis in our model framework boils down to producing very short-lived stock market disasters that last only a month (i.e. a unit period in our monthly calibration). As discussed in the Internet Appendix, not only does our model generate one-period stock market disasters, but the average stock market decline and the average VIX implied by the model are quite close to what we observe in the COVID-19 crisis.

How does our model produce these stock market crises and address their conditional dynamics without a large consumption decline? Investors' learning plays a key role. Even though the economy remains in the normal state, investors may perceive a significant probability of being in the depression state due to imperfect information. This, by itself, is enough to result in a disaster in the stock market as well as spikes in implied variance and the variance risk premium, consistent with the data. Overall, our model provides a unified framework that can account for a wide range of market extreme events, from a short-lived crash to a prolonged depression.

2.5 Conclusion

What is the source of the high equity premium in the postwar sample? As alluded by Welch (2016), the risk premia on put-protected portfolios can shed light on answering this question. By considering one-month put-protected portfolios with different moneyness levels, it is possible to decompose the equity premium into multiple pieces. Specifically, we find that 48%, 27%, 17%, and 10% of the entire equity premium in the data originate from monthly price declines that are larger than 10%, 15%, 20%, and 25%, respectively. On the surface, this seems to suggest a rejection of the rare disaster mechanism because only a small portion of the equity premium over a month is attributable to shocks that are larger than 25%.

In this paper, we show that modeling realistic equity dynamics during disasters resolves this issue. As exemplified by the Great Depression, macroeconomic disasters entail gradual but prolonged declines in the stock market, which makes it ineffective to hold a short-maturity put option as a hedge against disasters. To demonstrate this idea, we propose a model in which the true state of the economy is hidden. Investors form their beliefs about the current economic state in a Bayesian fashion, as they do not know with certainty whether today's change in immediate risk is due to a transitory shock or a persistent regime shift. The information structure and learning dynamics of our model create a slow response of equity prices to consumption declines during disasters.

Our results emphasize that how disasters unfold in the stock market has important asset pricing implications not only for disaster periods, but also for normal periods with high valuation. While typical models with instantaneous stock market disasters struggle to account for the VIX, variance risk premium, and risk premia on put-protected portfolios, we find that our model can naturally explain these quantities when it is calibrated to generate reasonable stock market dynamics during disasters. After all, investors fear disasters not because they expect sharp declines in equity prices within a short period of time, but because they are concerned about an extended period of stock market depression with no clear end in sight.

The slowly unfolding disaster mechanism developed in this paper can be applied to other asset markets. For instance, a multi-country version of our model can potentially reconcile rare disasters with high put-protected carry trade returns in currency markets, documented by Jurek (2014). We leave this for future work.

Appendix

A Model Derivations

Valuation Ratios

We start from the Euler equation for the return on the consumption claim:

$$\mathbb{E}_t\left[\exp\left(\log M_{t+1} + r_{c,t+1}\right)\right] = 1.$$

The definition of the stochastic discount factor in equation (2.2.1) and the expression $r_{c,t+1} = \log\left(1 + e^{pc_{t+1}}\right) - pc_t + \Delta c_{t+1}$ lead to the following equation:

$$\mathbb{E}_t\left[\exp\left(\theta \log \delta - \frac{\theta}{\psi}\Delta c_{t+1} + \theta \log\left(1 + e^{pc_{t+1}}\right) - \theta pc_t + \theta \Delta c_{t+1}\right)\right] = 1.$$

We can take the term $\exp\left(\theta \log \delta - \theta pc_t\right)$ out of this conditional expectation, as it is time-t measurable. In addition, since log consumption growth consists of the i.i.d. normal shock ϵ^c_{t+1} and the jump process J_{t+1} that are both independent of the price-consumption ratio pc_{t+1} conditional on λ_t and π_t, we can split this conditional expectation into the following three:

$$\exp\left(\theta \log \delta - \theta pc_t\right) \mathbb{E}_t\left[e^{(1-\gamma)\left(\mu_c + \sigma_c \epsilon^c_{t+1}\right)}\right] \mathbb{E}_t\left[e^{(1-\gamma)J_{t+1}}\right] \mathbb{E}_t\left[\left(1 + e^{pc_{t+1}}\right)^{\theta}\right] = 1.$$

The first conditional expectation equals $e^{(1-\gamma)\mu_c + \frac{1}{2}(1-\gamma)^2 \sigma_c^2}$, because it is the moment generating function of a normal distribution, evaluated at $1-\gamma$. Similarly, the second conditional expectation is the moment generating function of a Poisson jump process J_{t+1} evaluated at

$1-\gamma$, and is expressed as $e^{\lambda_t(\Phi_Z(1-\gamma)-1)}$. Therefore, the Euler equation becomes:

$$\exp\left(\theta\log\delta-\theta pc_t+(1-\gamma)\mu_c+\frac{1}{2}(1-\gamma^2)\sigma_c^2+\lambda_t\left[\Phi_Z(1-\gamma)-1\right]\right)\mathbb{E}_t\left[\left(1+e^{pc_{t+1}}\right)^{\theta}\right]=1.$$

Taking the logarithm of both sides leads to the recursive expression for the log wealth-consumption ratio in equation (2.2.1).

The recursive expression for the log price-dividend ratio in equation (2.2.2) can also be obtained using a similar approach. The log return on the dividend claim (or equity) satisfies the Euler equation:

$$\mathbb{E}_t\left[\exp\left(\log M_{t+1}+r_{d,t+1}\right)\right]=1.$$

Plugging the expression $r_{d,t+1}=\log\left(1+e^{pd_{t+1}}\right)-pd_t+\phi\Delta c_{t+1}$ into the Euler equation results in:

$$\begin{aligned}\mathbb{E}_t\big[\exp\big(\theta\log\delta+(\phi-\gamma)\Delta c_{t+1}+(\theta-1)\log\left(1+e^{pc_{t+1}}\right)\\ -(\theta-1)pc_t+\log\left(1+e^{pd_{t+1}}\right)-pd_t\big)\big]=1.\end{aligned}$$

We re-express this equation using the fact that pd_t is time-t measurable and Δc_{t+1} is independent of pc_{t+1} and pd_{t+1}, given λ_t and π_t:

$$\exp\left(\theta\log\delta-(\theta-1)pc_t-pd_t\right)\mathbb{E}_t\left[\exp\left([\phi-\gamma]\Delta c_{t+1}\right)\right]\mathbb{E}_t\left[\left(1+e^{pc_{t+1}}\right)^{\theta-1}\left(1+e^{pd_{t+1}}\right)\right]=1,$$

where, as explained earlier, $\mathbb{E}_t\left[\exp\left([\phi-\gamma]\Delta c_{t+1}\right)\right]$ equals the multiplication of the moment generating function of ϵ_{t+1}^c and that of J_{t+1}, evaluated at $\phi-\gamma$. By taking the logarithm of both sides of the above equation, we obtain the recursive expression for the log price-dividend ratio in equation (2.2.2).

Zero-Coupon Bonds

Let $B_t^{(n)} \equiv B(n; \lambda_t, \pi_t)$ denote the time-t price of a real zero-coupon bond maturing in n periods. At time $t+1$, this bond becomes a zero-coupon bond maturing in $n-1$ periods. Since zero-coupon bonds have no intermediate cash flow, the Euler equation implies:

$$B_t^{(n)} = \mathbb{E}_t \left[M_{t+1} B_{t+1}^{(n-1)} \right]. \tag{A.1}$$

By plugging the expression for the stochastic discount factor in equation (2.2.1) into equation (A.1), we obtain the following relation:

$$\begin{aligned} B(n; \lambda_t, \pi_t) = \exp\left(\theta \log \delta - (\theta - 1) pc_t - \gamma \mu_c + \frac{1}{2} \gamma^2 \sigma_c^2 + \lambda_t \left[\Phi_Z(-\gamma) - 1 \right] \right) \\ \times \mathbb{E}_t \left[\left(1 + e^{pc_{t+1}} \right)^{\theta - 1} B(n-1; \lambda_{t+1}, \pi_{t+1}) \right]. \end{aligned}$$

This relation recursively defines the zero-coupon bond price $B(n; \lambda, \pi)$ from $B(n-1; \lambda, \pi)$. Since the bond pays one unit of consumption at maturity, the boundary condition for the recursion is given as $B(0; \lambda, \pi) = 1$. We calculate the n-period continuously compounded zero-coupon yield as:

$$y_t^{(n)} = -\frac{1}{n} \log \left(B(n; \lambda_t, \pi_t) \right).$$

Since we calibrate our model at a monthly frequency, each period represents one month. Thus, $y_t^{(n)}$ also represents a monthly log yield. The annualized yield is obtained by multiplying $y_t^{(n)}$ by 12.

Figure 2.1: Monthly Returns During the Great Depression

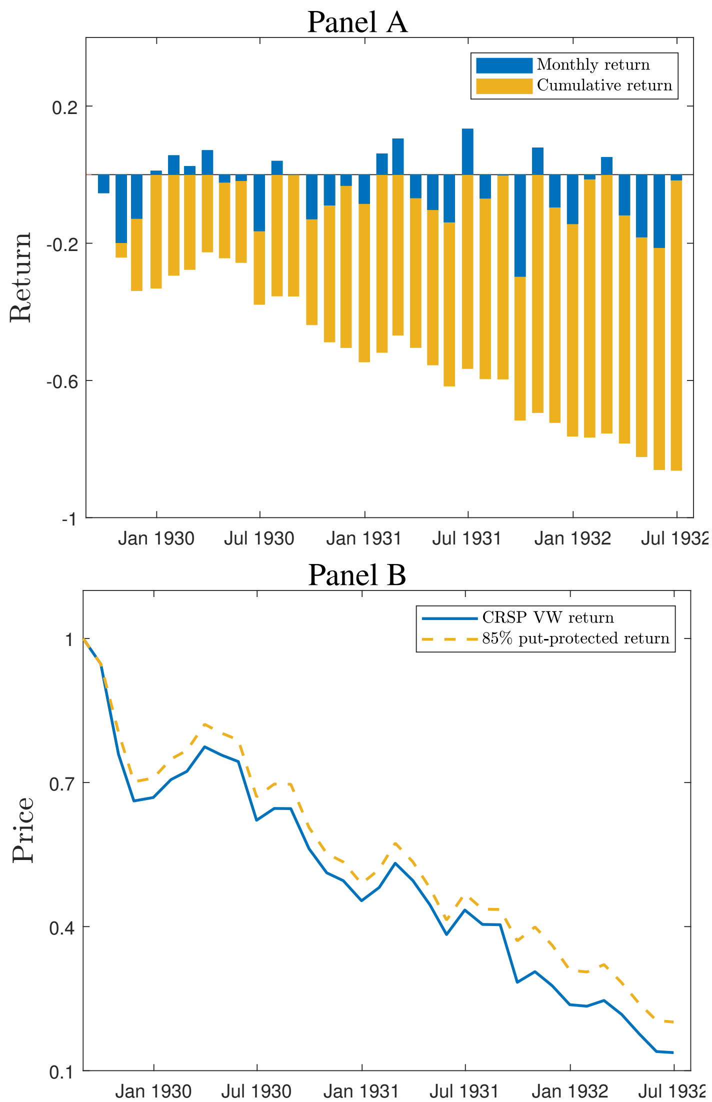

Notes: Panel A depicts the monthly returns on the CRSP value-weighted index excluding dividends, from September 1929 to June 1932. The blue bars represent month-over-month returns and the yellow bars represent cumulative returns. In Panel B, the solid blue line represents the path of the equity index while the dashed yellow line represents the value of the put-protected portfolio with 85% moneyness, ignoring the cost for acquiring put options. Both positions are normalized to one as of September 1929.

Figure 2.2: Learning and Slowly Unfolding Disasters

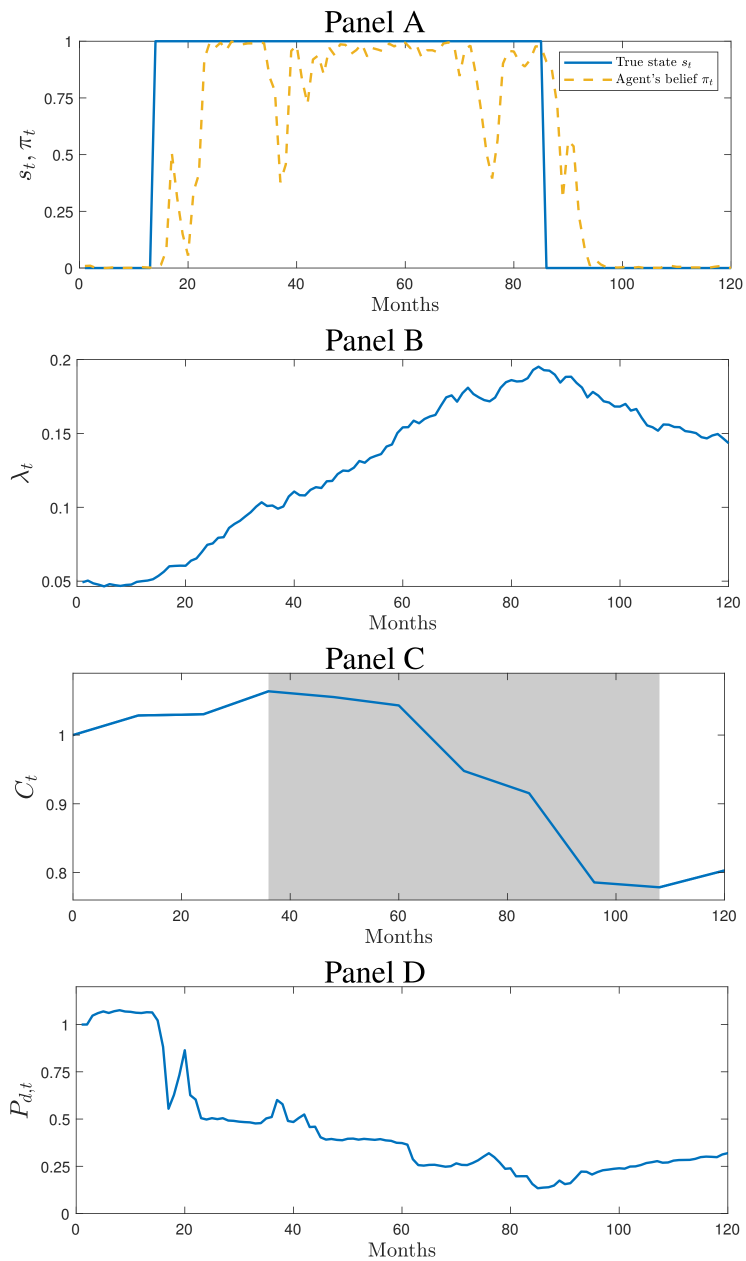

Notes: This figure plots the dynamics of our model in a sample path that includes a consumption disaster. Panel A plots the path of the state s_t and belief π_t. Panel B plots the path of jump intensity λ_t. Panels C and D plot the paths of the equity price $P_{d,t}$ and annual consumption growth C_t, respectively, both of which are normalized to one at month 0.

Figure 2.3: Valuation Ratios

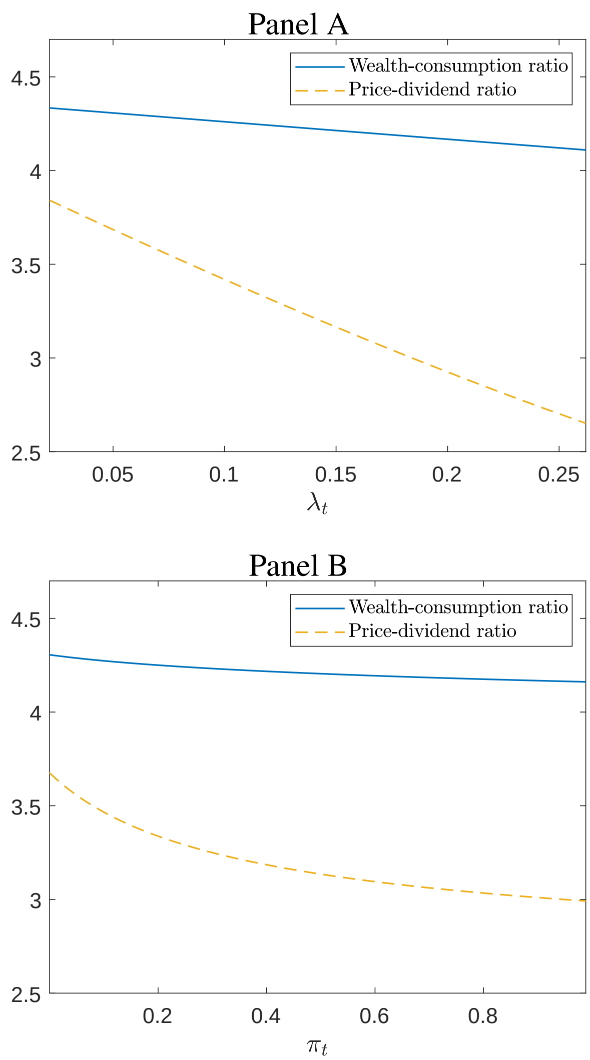

Notes: This figure plots the annualized log wealth-consumption ratio and the annualized log price-dividend ratio as a function of jump intensity λ_t (Panel A) or belief π_t (Panel B). We set π_t to its median value in Panel A. Likewise, we set λ_t to its median value in Panel B.

Figure 2.4: Risk-Free Rate, Equity Premium and Conditional Volatility

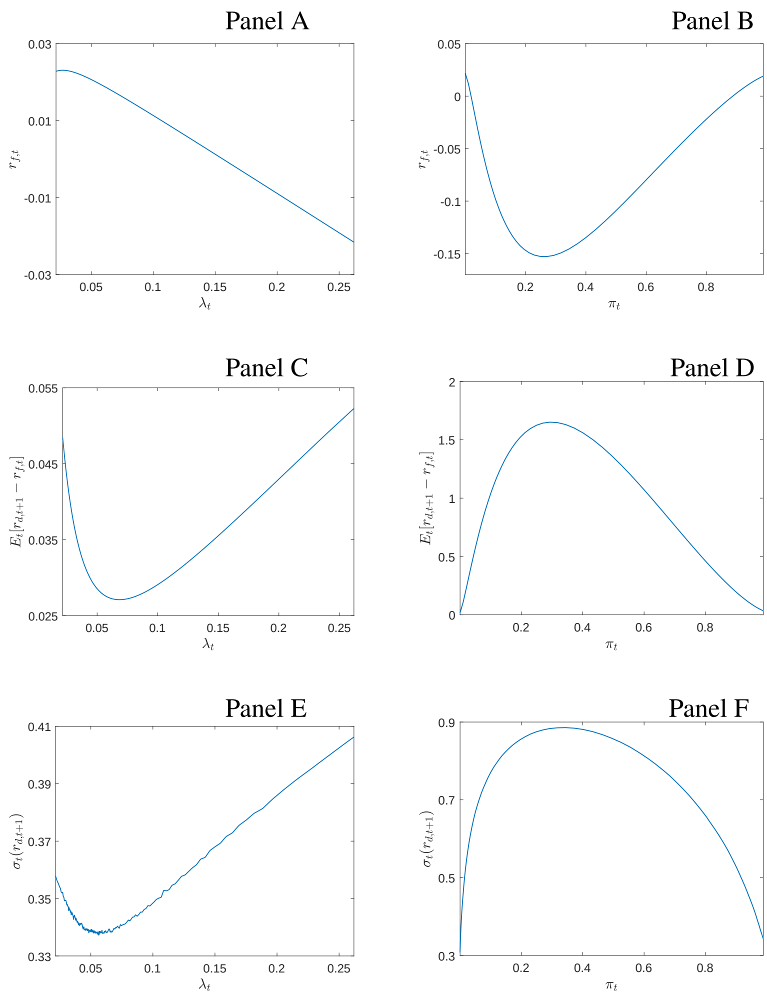

Notes: This figure plots the log risk-free rate $r_{f,t}$, the conditional equity premium $\mathbb{E}_t[r_{d,t+1} - r_{f,t}]$, and the conditional volatility of log equity return $\sigma_t(r_{d,t+1})$ as a function of jump intensity λ_t or belief π_t. We set π_t to its median value in Panels A, C, and E. Likewise, we set λ_t to its median value in Panels B, D, and F.

Figure 2.5: Term Structure of Interest Rates

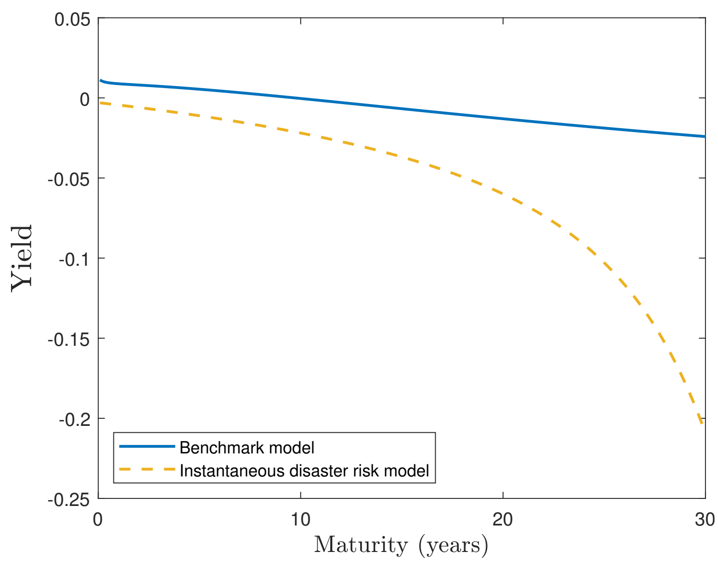

Notes: This figure plots the average log yields to maturity on risk-free zero-coupon bonds as a function of maturity. The solid blue line represents our model, and the dashed yellow line represents the instantaneous disaster risk model. Yields are in annual terms.

Figure 2.6: Disaster Size Distribution

Panel A: Data

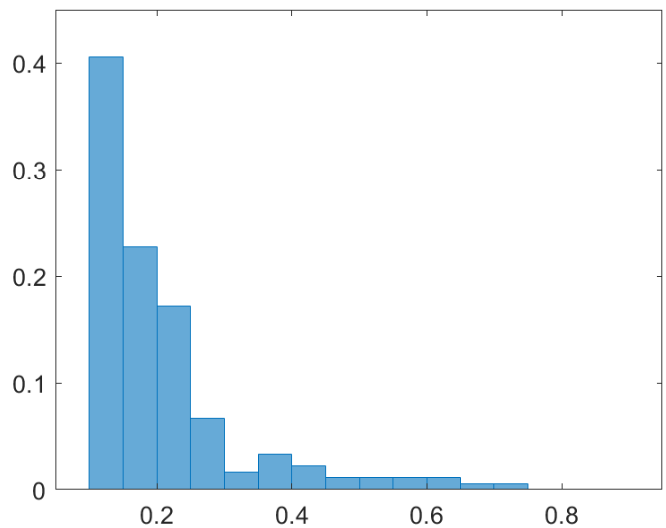

Panel B: Model

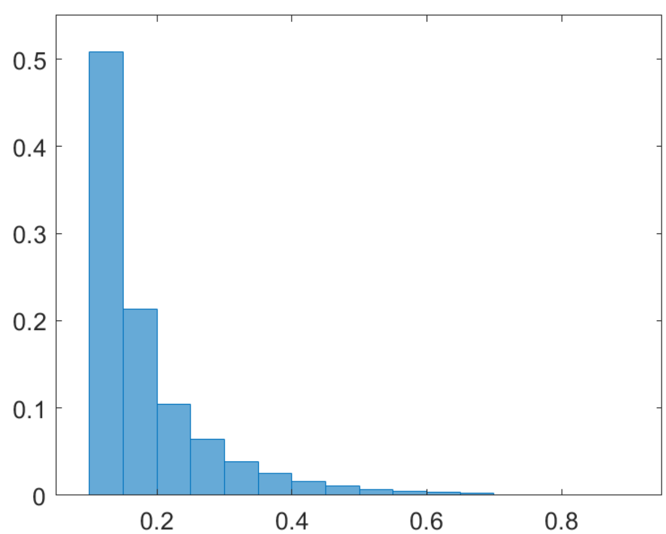

Notes: This figure plots the distribution of cumulative consumption declines during disasters in the data (Panel A) and in the model (Panel B). Disasters are identified as peak-to-trough consumption declines that are larger than 10%, consistent with Barro and Ursúa (2008).

Table 2.1: Model Parameters

Risk aversion, γ	5
Elasticity of intertemporal substitution, ψ	1.5
Time discount, δ	0.9990
Mean consumption growth in the absence of jumps, μ_c	0.0026
Volatility of consumption growth in the absence of jumps, σ_c	0.0020
Leverage parameter, ϕ	3
Persistence of jump intensity, ρ_λ	0.9933
Conditional volatility of shocks to jump intensity, σ_λ	0.0083
Average jump intensity in normal times, $\bar{\lambda}_L$	0.0417
Average jump intensity in depression times, $\bar{\lambda}_H$	0.4167
Mean jump size, μ_Z	-0.0200
Transition probability from the normal state to depression state, p_{01}	0.0017
Transition probability from the depression state to normal state, p_{10}	0.0208

Notes: This table reports the parameters for the benchmark model, calibrated at a monthly frequency.

Table 2.2: Consumption and Asset Pricing Moments

	$\sigma(\Delta c)$	$\sigma(\Delta d)$	$\mathbb{E}[r_d - r_f]$	$\sigma(r_d)$	$\mathbb{E}[r_f]$	$\sigma(r_f)$
Data	1.40	6.76	6.67	17.13	0.63	2.23
Panel A: Benchmark model						
Median	1.80	5.40	7.13	17.06	1.51	0.61
5%	1.23	3.71	5.48	12.45	0.67	0.42
95%	2.38	7.14	8.47	21.36	1.97	0.77
Population	2.50	7.50	6.51	18.17	1.06	0.68
Panel B: Model with instantaneous disaster risk						
Median	1.62	4.21	6.15	14.58	2.12	0.32
5%	1.38	3.58	4.31	11.09	0.84	0.17
95%	1.87	4.87	8.61	19.37	2.82	0.60
Population	5.47	14.36	5.38	24.30	0.85	4.10

Notes: This table reports the annual consumption and equity market statistics in the data and in the model. $\sigma(\Delta c)$ is the standard deviation of log consumption growth. $\sigma(\Delta d)$ is the standard deviation of log dividend growth. $\mathbb{E}[r_d - r_f]$ is the average excess log return on the market. $\sigma(r_d)$ is the standard deviation of the log market return. $\mathbb{E}[r_f]$ is the average and $\sigma(r_f)$ is the standard deviation of the log risk-free rate. Data values are for the period from 1950 to 2019. In the model, we simulate 10,000 70-year-long samples and report the 50th, 5th, and 95th percentiles of the model statistics from no-disaster samples. The population statistics are obtained from a long simulation path of 1,000,000 years. All values are reported in percentage terms.

Table 2.3: Return Predictability

	β_{1y}	β_{3y}	β_{5y}	R^2_{1y}	R^2_{3y}	R^2_{5y}
Data	-0.11	-0.07	-0.07	6.84	12.03	17.46
Panel A: Benchmark model						
Median	-0.24	-0.15	-0.11	9.27	15.32	17.29
5%	-0.69	-0.30	-0.21	0.30	0.33	0.27
95%	-0.02	0.01	0.02	30.63	40.82	46.71
Population	-0.05	-0.04	-0.03	1.58	2.99	3.62
Panel B: Model with instantaneous disaster risk						
Median	-0.23	-0.20	-0.17	13.66	34.16	46.41
5%	-0.39	-0.31	-0.24	6.05	14.11	17.76
95%	-0.12	-0.11	-0.10	22.38	50.92	66.70
Population	-0.10	-0.09	-0.08	3.68	9.80	14.51

Notes: This table reports the statistics from the predictability regressions of the following form:

$$\frac{1}{h}\sum_{j=1}^{h} r_{d,t+j} - r_{f,t+j-1} = \beta_0 + \beta_{hy} pd_t + \epsilon_{t+h},$$

where $r_{d,t+j} - r_{f,t+j-1}$ is the excess log return on the market from year $t+j-1$ to year $t+j$, and pd_t is the log price-dividend ratio of the aggregate equity. The results are reported for h = 1, 3, and 5 years and the regressions are run at an annual frequency from 1950 to 2019. In the model, we simulate 10,000 70-year-long samples and report the 50th, 5th, and 95th percentiles of the model statistics from no-disaster samples. The population statistics are obtained from a long simulation path of 1,000,000 years. In the last three columns, R^2 values are reported in percentage terms.

Table 2.4: Implied Variance and the Variance Risk Premium

	$\mathbb{E}[IV]$	$\sigma(IV)$	$\mathbb{E}[VRP]$	$\sigma(VRP)$
Data	35.15	32.65	15.51	19.94
Panel A: Benchmark model				
Median	40.23	39.61	16.35	17.95
5%	34.87	27.17	8.94	12.87
95%	47.41	51.80	21.19	23.67
Population	40.71	39.68	14.49	18.47
Panel B: Model with instantaneous disaster risk				
Median	170.92	106.94	139.00	88.82
5%	68.69	48.75	54.42	40.60
95%	400.37	229.57	328.44	188.92
Population	270.43	239.28	221.13	197.96

Notes: This table reports the statistics for implied variance and the variance risk premium. $\mathbb{E}[IV]$ is the average and $\sigma(IV)$ is the standard deviation of implied variance. $\mathbb{E}[VRP]$ is the average and $\sigma(VRP)$ is the standard deviation of the variance risk premium. All quantities are expressed in monthly percentage squared terms. Data values are for the period from 1990 to 2019. In the model, we simulate 10,000 30-year-long samples and report the 50th, 5th, and 95th percentiles of the model statistics from no-disaster samples. The population statistics are obtained from a long simulation path of 1,000,000 years.

Table 2.5: Put-Protected Portfolio Moments

	Premium difference $\mathbb{E}[r_d - r_d^k]$				Premium ratio $\frac{\mathbb{E}[r_d^k - r_f]}{\mathbb{E}[r_d - r_f]}$			
Moneyness	75%	80%	85%	90%	75%	80%	85%	90%
Data	0.72	1.13	1.89	3.27	0.90	0.83	0.73	0.52
Panel A: Benchmark model								
Median	0.72	1.15	1.77	2.53	0.90	0.84	0.75	0.64
5%	-0.31	-0.24	-0.01	0.36	0.79	0.68	0.55	0.41
95%	1.36	2.07	2.95	3.99	1.05	1.04	1.00	0.94
Population	0.53	0.88	1.40	2.07	0.92	0.87	0.79	0.68
Panel B: Model with instantaneous disaster risk								
Median	2.56	2.88	3.23	3.61	0.58	0.53	0.47	0.41
5%	1.01	1.14	1.27	1.42	0.12	0.01	-0.11	-0.24
95%	6.02	6.79	7.63	8.54	0.81	0.78	0.76	0.73
Population	2.66	2.96	3.29	3.64	0.51	0.45	0.39	0.32

Notes: This table reports the statistics for the returns on put-protected portfolios with various moneyness values $k = 0.75$, 0.80, 0.85, and 0.90. The premium difference, $\mathbb{E}[r_d - r_d^k]$, is the difference between the average equity premium and the average put-protected portfolio premium. The premium ratio, $\frac{\mathbb{E}[r_d^k - r_f]}{\mathbb{E}[r_d - r_f]}$, is the ratio of the average put-protected portfolio premium to the average equity premium. Data values are for the period from 1990 to 2019. In the model, we simulate 10,000 30-year-long samples and report the 50th, 5th, and 95th percentiles of the model statistics from no-disaster samples. The population statistics are obtained from a long simulation path of 1,000,000 years.

3 Is There a Macro-Announcement Premium?

3.1 Introduction

Lately, the literature has discovered striking return patterns in the equity market. The average excess return over macro-announcement days is about 10 basis points (bp) per day whereas that over non-announcement days is merely 1-2 bp (e.g., Savor and Wilson, 2013). Recent state-of-the-art models, such as Ai and Bansal (2018) and Wachter and Zhu (2020), justify this return gap by hypothesizing a macro-announcement premium. Important economic uncertainty resolves on macro-announcement days, and investors demand extra compensation for bearing such uncertainty. Simply put, a high macro-announcement return in these models is a manifestation of a high conditional equity premium. Uncertainty per-

We thank Hengjie Ai, Ehsan Azarmsa, Mikhail Chernov, Anna Cieslak, Francois Cocquemas, Christian Dorion, Hitesh Doshi, Casey Dougal, Bjørn Eraker, Stephen Figlewski, Thomas Gilbert, Kris Jacobs, Mete Kilic, Jung-Wook Kim, Jongsub Lee, Michael Junho Lee, Xinyang Li, Dmitriy Muravyev, Neil Pearson, Ivan Shaliastovich, Xiaoxiao Tang, Jessica Wachter, Randall Wright, and seminar participants at the Virtual Derivatives Workshop, 2021 MFA meeting, 2021 EFA meeting, 2021 SWFA meeting, SunTrust Speaker Series at the Florida State University, Seoul National University, University of Wisconsin-Madison, and University of Houston for valuable suggestions and helpful comments.

ceived by investors increases so much prior to a macroeconomic announcement that the conditional equity premium becomes as high as 10 bp per day (or 25% per annum).

Despite the empirical success of these models, a question remains. Why does such a high degree of uncertainty only resolve on macro-announcement days? As evident in survey data, investors can form their expectations about macroeconomic outcomes far ahead of official releases. Although it is true that some degree of uncertainty has to resolve exactly at announcements, effective communication by government agencies like the Fed may allow investors to become confident and, hence, resolve most uncertainty well before actual announcements. Unless the Fed (whose role centers around managing market expectations) and other government agencies are doing something terribly wrong in communicating with the market, it is difficult to imagine that the entire 10 bp average excess return is compensation for perceived uncertainty.

The premise that the high average macro-announcement return originates from the high average equity premium faces another challenge when the patterns of the ex-ante second moment of returns are examined. In typical dynamic equilibrium models (including the aforementioned models), the conditional equity premium and the conditional return volatility arise from the same mechanism. Thus, if macro-announcement returns spike up due to a significant increase in the conditional equity premium, we should also expect a significant increase in the conditional volatility over macro-announcement days. However, this is not the case in the data. The average VIX barely changes from 19.50% to 19.27% before and after announcements. Volatility indices with shorter maturities (say, 1 or 2 weeks), constructed using option prices like the VIX, feature a small decline.[1] The daily conditional volatilities implied by several GARCH models also show little difference. If high

[1] We construct the volatility indices for 1-week and 2-week horizons using options data from OptionMetrics. In the period of 1996-2018, the 1-week VIX and the 2-week VIX change on average from 19.44% to 18.50% and from 19.37% to 18.77%, respectively, before and after macro-announcements.

macro-announcement returns are due to high uncertainty perceived by investors, why isn't the ex-ante conditional volatility, which directly reflects such uncertainty, also high?

In this paper, we propose an alternative explanation: macro-announcement days happen to be, on average, with good news in existing sample periods. That is, the high average excess return over macro-announcement days is simply the outcome of positive return innovations; it is not because of high conditional equity premiums associated with heightened uncertainty. This explanation is consistent with the view that since the 1980s, the Fed has constantly surprised the market by taking more aggressive actions than what the market anticipated ex-ante (Cieslak, 2018; Cieslak, Morse, and Vissing-Jorgensen, 2018).

Our claim that return innovations do not average out in-sample relies on a small sample argument. In fact, samples with macro-announcement days are small. In the period of 1990-2018, there are a total of 915 observations (equivalent to less than 4 years of data). Even if we extend our sample period to 1961, there are 1,903 macro-announcement days (equivalent to less than 8 years of data). This small sample issue is exacerbated by the empirical regularities of daily return data: returns are extremely volatile and fat-tailed at the daily frequency both in the macro-announcement and non-announcement samples. Hence, it is completely possible that the average excess return significantly deviates from the true risk premium in-sample.

Instead of relying on the average excess return, we develop a novel approach to accurately estimate the macro-announcement premium by exploiting so-called "asymmetric volatility." Also known as the "leverage effect," asymmetric volatility is one of the most salient features of the equity market.[2] Essentially, this market phenomenon indicates that

[2]For a further discussion about asymmetric volatility, see Black (1976), Christie (1982), French, Schwert, and Stambaugh (1987), Schwert (1989), Campbell and Hentschel (1992), Cheung and Ng (1992), Duffee (1995), Bekaert and Wu (2000), Christoffersen and Jacobs (2004), Bollerslev, Litvinova, and Tauchen (2006), Bollerslev and Zhou (2006), Christoffersen, Heston, and Jacobs (2006), Bandi and Reno (2012), Bollerslev, Sizova, and Tauchen (2012), Wang and Mykland (2014), and Kalnina and Xiu (2017).

return innovations are strongly negatively correlated with variance innovations (i.e. unexpected changes in variance). Options data imply that this conditional correlation is fairly close to -1, ranging from -0.6 to -0.8. This suggests that a variance innovation can serve as a highly informative signal about the contemporaneous return innovation: if the market variance is hit by a negative (positive) innovation, it is probable that the market return experiences a positive (negative) innovation. Therefore, even if return innovations do not average out in-sample due to a small sample size, we can correct for them using variance innovations and obtain a more robust estimate for the macro-announcement premium.

Specifically, we introduce a simple statistical model featuring asymmetric volatility. We estimate the parameters of the model via maximum likelihood estimation (MLE), based on the sample period of 1990-2018. To simplify the estimation procedure, we avoid a direct filtering of daily conditional variances. Instead, we assume that they share a linear relationship with the squared VIX. Our estimation results reveal that macro-announcement days are with on-average negative variance innovations during the sample period, which translate into on-average positive return innovations through asymmetric volatility. Although the model is designed to accommodate the potential existence of a macro-announcement premium, it is estimated to be statistically insignificant. It turns out that the high average macro-announcement return does not come from a high macro-announcement premium, but is mostly due to on-average positive return innovations.

Our estimation results are robust in various dimensions. The estimation results are similar when we choose only one specific type of macro-announcement news, such as FOMC decisions, employment figures, or price indices. Applying different risk premium specifications to the model does not change our conclusion either. Moreover, extending our sample back to 1961 and filtering out daily conditional variances via Markov Chain Monte Carlo (MCMC) provide similar estimation results. The main takeaway of these exercises is that the average excess return over macro-announcement days substantially exaggerates the true risk premium.

We compare and contrast our asymmetric-volatility-based explanation with an alternative explanation based on uncertainty resolution. These two explanations differ in their accounts of high excess returns and negative changes in the VIX over macro-announcement days. In our explanation, they are the result of asymmetric volatility: positive return innovations coincide with negative variance innovations (i.e. unexpected falls) like any other trading day. In contrast, the alternative explanation interprets them as high conditional equity premiums and negative volatility drifts (i.e. expected falls) that are associated with uncertainty resolution. To empirically assess this alternative story, we consider a different statistical model that can accommodate an expected increase and a subsequent decrease in variance around macro-announcements. The estimation results show that the volatility drift does absorb a portion of the average variance innovation in the macro-announcement sample, resulting in a higher macro-announcement premium. However, the estimated macro-announcement premium is still statistically insignificant and much smaller than what the average excess return suggests.

The advantage of our explanation is that it offers a joint account of the return and VIX patterns over macro-announcement days. The small average decrease in the VIX is not anomalous if large macro-announcement returns originate from large return innovations (rather than large conditional premiums). At the daily frequency, return innovations are extremely volatile, and, therefore, a small unexpected fall in variance is associated with a large positive return realization through asymmetric volatility. In fact, we show that the small drop in the average VIX in the data translates into positive return innovations that are large enough to justify the average macro-announcement return in the data.

The main intuition behind our explanation is that the patterns of the macro-announcement sample are not so different from those of the non-announcement sample, if we properly take asymmetric volatility into account. We find that the empirical relation between excess returns and changes in the squared VIX is almost identical in the two sub-samples. Thus, non-announcement days with a similar change in the squared VIX exhibit a similar excess return on average compared to the average macro-announcement day. Moreover, the large average excess return and the small average decrease in the VIX over the macro-announcement sample are well replicated by taking random samples of non-announcement days. Lastly, we discover that a long representative sample only consisting of non-announcement days features the average excess return that is quite close to the unconditional equity premium in the postwar sample. This finding goes directly against the claim that the equity premium is mostly realized over macro-announcement days. Overall, the data is inconsistent with the view that macro-announcement days are special and operate with a separate mechanism.

Our conclusion bears important implications for macro-finance models. Under our explanation, the large average macro-announcement return in the data does not represent a large risk premium. We do not need a new model to justify it; traditional models are just fine. As an example, we show that the model of Seo and Wachter (2018b), which does not feature an announcement premium, fully rationalizes the return and VIX patterns over macro-announcement days. Considering all the evidence from empirical and modeling contexts, we believe that researchers should take a step back and ponder the need for a more complex model. At the least, new models should be calibrated to produce a much smaller macro-announcement premium, as the average excess return greatly overstates the true macro-announcement premium.

As a precaution, we emphasize that our findings do not imply that macro-announcements themselves are not important. Without a doubt, the Fed's decisions and economic indicators have an immense impact on asset markets including the equity market. However, the timing at which uncertainty actually resolves is an empirical question. It might be the case that most uncertainty resolves exactly at announcements. Or it might be the case that most uncertainty effectively resolves far in advance due to investors' expectation formation. Perhaps the reality is a mixture of both, but the data suggest that the latter dominates the former, at least when it comes to the realization of the equity premium.[3]

Related Literature

Large macro-announcement returns are first documented by Savor and Wilson (2013). More recently, Brusa, Savor, and Wilson (2020) find that global equity markets also exhibit high equity returns over U.S. macro-announcement days.[4] Typically, large macro-announcement returns are interpreted as large conditional equity premiums, and recent theoretical work is built on this premise. Ai and Bansal (2018) establish the necessary and sufficient conditions for the set of intertemporal preferences to produce a high macro-announcement premium. Wachter and Zhu (2020) propose a model in which macroeconomic announcements convey information about the latent probability of disaster. Our work is in sharp contrast with these papers: we argue that the large average macro-announcement return does not necessarily translate into a large macro-announcement premium.

Estimating the equity premium using the sample average of excess returns makes an implicit assumption that the given sample is long and representative. There is an extensive

[3]Consistent with our results, Bernanke and Kuttner (2005) document that the stock market responds strongly to unanticipated Fed policy changes whereas it responds little, if at all, to anticipated ones.

[4]The scope of the literature on macro-announcements and their asset pricing implications extends beyond the equity market. Andersen, Bollerslev, Diebold, and Vega (2007) compare the responses of global equity, bond, and foreign exchange markets to the release of real-time U.S. macroeconomic news. Nakamura and Steinsson (2018) study how real interest rates, expected inflation, and expected output growth react to scheduled monetary policy announcements within 30-minute windows around them. Balduzzi and Moneta (2017) and Mueller, Tahbaz-Salehi, and Vedolin (2017) investigate the impact of macro-announcements on the markets for Treasury bonds and exchange rates, respectively.

literature questioning this very assumption, even in samples longer than 60-70 years. For example, Avdis and Wachter (2017) adopt a statistical model and estimate it not only with the time series of the excess return, but also with that of the dividend-price ratio.[5] They find that incorporating additional information contained in dividends and prices results in a significantly lower unconditional equity premium, compared to the sample average excess return. Given that the macro-announcement sample collects a small number of non-random days, its representativeness is even more in question. Our approach can be compared with Avdis and Wachter (2017): we adopt a statistical model to obtain a more precise estimate for the macro-announcement premium. Instead of the price-dividend ratio, we exploit the shock to the variance as an informative signal, motivated by asymmetric volatility.

Our explanation of high macro-announcement returns relies on the discrepancy between the average return and the expected return in short samples of macro-announcements. One might wonder why the average return happens to be larger than the expected return, not the other way around. While this can purely be due to statistical reasons, recent empirical evidence surrounding Fed policies provides some economic insights. Cieslak (2018) documents large and persistent errors in investors' expectations about the future path of interest rates. Specifically, she points out that the Fed has often surprised the public by easing more aggressively than expected. Building on this, Cieslak, Morse, and Vissing-Jorgensen (2018) find that the high average return in the post-1994 FOMC announcement sample was, in part, unexpected. These results are exactly consistent with our findings: the Fed's more-than-anticipated easing policies result in positive return innovations and negative variance innovations that do not average out in-sample. Our small sample argument still plays a critical role; the sample is too short to accommodate investors' subsequent belief revisions in response to positive surprises from the Fed.

[5]Other examples include, but are not limited to, Gebhardt, Lee, and Swaminathan (2001), Constantinides (2002), Fama and French (2002), Ibbotson and Chen (2003), Donaldson, Kamstra, and Kramer (2010), Martin (2017), and Chabi-Yo and Loudis (2020).

Unusual and intriguing return patterns surrounding macro-announcements are not limited to the presence of high macro-announcement returns. Savor and Wilson (2014) show that the CAPM well explains the cross section of equities, bonds, and currencies when it is restricted to macro-announcement days. Lucca and Moench (2015) find a positive price drift that begins a few hours prior to FOMC announcements, the so-called pre-FOMC drift. Although examining these additional issues is beyond the scope of our paper, we believe that our small-sample argument bears some empirical relevance. Due to the small sample size, the empirical findings on macro-announcements are rather sensitive to the choice of event definitions and sample periods.[6] Specifically, Ernst, Gilbert, and Hrdlicka (2021) show that depending on how a macro-announcement sample is constructed, it can account for more than 100% of the entire equity premium. They also find that the higher CAPM slope on macro-announcement days is a mechanical result of high ex-post returns. Consistently, our paper raises some concerns and doubts about the representativeness of existing macro-announcement samples. We point out that daily return data are particularly susceptible to small-sample issues as return innovations are extremely volatile and fat-tailed.

The rest of the paper proceeds as follows. In Section 3.2, we cast doubt on the representativeness of a macro-announcement sample and raise the possibility that the average excess return may substantially overstate the true risk premium over macro-announcement days. In Section 3.3, we introduce our novel approach of estimating the macro-announcement premium, which exploits asymmetric volatility. A statistical model and its estimation are introduced, and the estimation results are analyzed. In Section 3.4, we investigate an alternative explanation based on uncertainty resolution and compare it with our explanation.

[6]For instance, using the data from 1980 to 2011, Lucca and Moench (2015) conclude that a pre-announcement drift is only observed before FOMC decisions. However, Hu, Pan, Wang, and Zhu (2020) discover that other major macroeconomic news also feature comparable drift patterns in the period of 1994-2018. Furthermore, Boguth, Grégoire, and Martineau (2019) find that after 2011, both the pre-FOMC drift and the drop in VIX only appear for the FOMC meetings that are followed by a press conference.

The implications of our results for macro-finance models are also provided. Section 3.5 checks the robustness of the benchmark estimation by extending the sample period via MCMC. Section 3.6 concludes.

3.2 What Is the Issue With the Average Excess Return?

In the existing literature, the macro-announcement premium is typically estimated using the average excess return over macro-announcement days. This estimation approach makes an implicit assumption that a given sample is long enough to be representative of the entire population. Specifically, averaging over macro-announcement days should effectively eliminate the impact of return innovations, leaving behind the true risk premium.

How realistic is this assumption in a sample only with macro-announcement days? Examining the validity of this assumption is especially relevant because the sample is small. Although the average excess return over macro-announcement days is large in the data, the true macro-announcement premium can be small if macro-announcement days coincide, on average, with positive return innovations (i.e. good news) in-sample.

To illustrate, let rx_{t+1} denote the log excess return on the equity market index between times t and $t+1$:

$$rx_{t+1} \equiv \log(R_{e,t+1}) - \log(R_{f,t}),$$

where R_{t+1} is the gross return on the index and $R_{f,t}$ is the one-period gross risk-free rate. We define $\mu_t \equiv \mathbb{E}_t\left[rx_{t+1}\right]$ as the time-t conditional mean of the excess return (or simply, conditional equity premium) and $v_t \equiv \mathrm{Var}_t\left(rx_{t+1}\right)$ as the time-t conditional variance of the excess return. Under this notation, the excess return rx_{t+1} can be expressed as

$$rx_{t+1} = \mu_t + \sqrt{v_t}\epsilon_{t+1}, \tag{3.2.1}$$

where $\mathbb{E}_t\left[\epsilon_{t+1}\right] = 0$ and $\mathrm{Var}_t\left(\epsilon_{t+1}\right) = 1$. Equation (3.2.1) makes it clear that the realized excess return can be decomposed into two components: the conditional equity premium (μ_t) and the contemporaneous return innovation ($\sqrt{v_t}\epsilon_{t+1}$). Therefore, a high realized excess return on a certain day does not necessarily imply that the conditional equity premium was high that day; instead, it might be the case that the return innovation was realized as a high value.

Whereas the existence of the return shock ϵ_{t+1} makes it difficult to identify the conditional equity premium, it does not affect the estimation of the unconditional equity premium, at least in theory. Let $\mathbb{U}$ denote the set of all trading days in a given sample. Then, $n(\mathbb{U})$, the cardinality of $\mathbb{U}$, represents the sample size. Under the reasonable assumption that the excess return rx_{t+1} follows an ergodic process, the average excess return over the full sample converges in probability to the unconditional equity premium $E\left[\mu_t\right]$, as the sample size increases:

$$\frac{1}{n(\mathbb{U})}\sum rx_{t+1} \quad \xrightarrow{p} \quad \mathbb{E}\left[\mu_t\right] + \underbrace{\mathbb{E}\left[\sqrt{v_t}\epsilon_{t+1}\right]}_{=0}. \tag{3.2.2}$$

This is because $\mathbb{E}\left[\sqrt{v_t}\epsilon_{t+1}\right] = \mathbb{E}\left[\sqrt{v_t}\cdot\mathbb{E}_t\left[\epsilon_{t+1}\right]\right] = 0$. As long as we have a lengthy representative sample, positive and negative realizations of $\sqrt{v_t}\epsilon_{t+1}$ would balance out, and the effect of $\sqrt{v_t}\epsilon_{t+1}$ would vanish from the sample average.

We can expect the same result if we take the sample mean of excess returns only over macro-announcement days. We define $\mathbb{M}$ as the subset of $\mathbb{U}$ that collects macro-announcement days. Namely, if $t+1 \in \mathbb{M}$, it represents a day with a pre-scheduled macroeconomic announcement. As the sample size increases, the average excess return over macro-announcement days indeed converges in probability to the average equity premium over macro-announcement days $\mathbb{E}[\mu_t \mid t+1 \in \mathbb{M}]$:

$$\frac{1}{n(\mathbb{M})} \sum_{t+1 \in \mathbb{M}} rx_{t+1} \quad \xrightarrow{p} \quad \mathbb{E}[\mu_t \mid t+1 \in \mathbb{M}] + \underbrace{\mathbb{E}[\sqrt{v_t}\epsilon_{t+1} \mid t+1 \in \mathbb{M}]}_{=0}, \quad (3.2.3)$$

where $n(\mathbb{M})$, the cardinality of $\mathbb{M}$, is the number of macro-announcement days in the sample.

Equations (3.2.2) and (3.2.3) indicate that the average excess return converges to the true risk premium, provided the sample size goes to infinity. However, its finite-sample performance critically depends on how quickly return innovations die out as the sample period expands. Note that $n(\mathbb{M}) \ll n(\mathbb{U})$ because there are only a few macro-announcement days each year (8 to 30 days depending on the types of macro-announcement events). Therefore, the convergence of the average return innovation to zero can be much slower in the macro-announcement sample than in the full sample. That is, even if the sample is long enough for return innovations to average out over the full sample $\left(\text{i.e. } \frac{1}{n(\mathbb{U})} \sum_{t+1 \in \mathbb{U}} \sqrt{v_t}\epsilon_{t+1} \approx 0\right)$, this might not be the case for the macro-announcement sample $\left(\text{i.e. } \left|\frac{1}{n(\mathbb{M})} \sum_{t+1 \in \mathbb{M}} \sqrt{v_t}\epsilon_{t+1}\right| \gg 0\right)$.

In fact, the macro-announcement sample is quite small in the data. As discussed in the data section (Appendix A), we construct our main macro-announcement sample $\mathbb{M}_{\text{All}}$ based on three types of pre-scheduled releases of macroeconomic news: (i) Federal Open Market Committee (FOMC) decisions, (ii) employment figures, and (iii) price indices. In the period of 1990-2018, this results in a total of 915 macro-announcement days (equivalent to less than 4 years of data). Even if we consider the longest sample period starting from 1961, the number of macro-announcement days remains small. In the period of 1961-2018, there are 1,903 macro-announcement days (equivalent to less than 8 years of data).[7]

The empirical regularities of daily return data make this small sample issue worse. First, returns are extremely volatile at the daily frequency. Figure 3.1 presents the histograms of daily excess returns over the macro-announcement sample (blue bars) vs. the non-announcement sample (yellow bars). The dotted lines and the dashed lines represent the averages over the macro-announcement sample and the non-announcement sample, respectively. In the data, the average excess returns over the macro-announcement and non-announcement samples are 9.93 bp and 1.29 bp. This mean difference of 8.64 bp seems substantial, as it corresponds to an annualized return gap of 22%. However, Panel A of Figure 3.1 shows that this difference is completely dwarfed by the large variances of the distributions; the standard deviations are 112.04 bp and 108.62 bp. The two vertical lines that represent the sub-sample averages are not even distinguishable in Panel A when the histograms are plotted with their full ranges. When we zoom into the middle 99% (from 0.5 to 99.5 percentiles) in Panel B, we can finally separate out the two lines, but the gap is still extremely small. The main takeaway from this figure is that comparing the two sub-samples based on daily return data warrants extra caution.

Second, returns exhibit not only a large variance, but also a large kurtosis value at the daily frequency (10.43 and 11.97 for the macro-announcement and non-announcement samples, respectively). Hence, the average excess return converges to the true equity premium very slowly. Under such a fat-tailed distribution, three-sigma events where daily return innovations even exceed 300 bp do take place occasionally. For instance, suppose that the Fed's more-than-anticipated easing policies surprised the market (Cieslak, 2018; Cieslak, Morse, and Vissing-Jorgensen, 2018), which translated into large positive return innovations. To wash away the effect of such large shocks, we need many data points.

[7]The maximum sample period is currently about 60 years, as the U.S. Bureau of Labor Statistics began pre-releasing its announcement days from 1961.

However, the size of the macro-announcement sample might not be big enough.

In sum, with the macro-announcement sample being so small, it is difficult to determine whether the average excess return is a good estimate for the true equity premium over macro-announcement days. If macro-announcement days are, on average, with positive (negative) return innovations in-sample, simply taking the average of excess returns over these days will exaggerate (understate) the true risk premium.

3.3 Estimation Based on Asymmetric Volatility

In this paper, we propose an alternative approach to estimate the macro-announcement premium, which does not rely on the assumption that return innovations average out over the macro-announcement sample. As discussed in Section 3.2, the average excess return is not a good proxy for the true risk premium when the sample is not long and representative. A typical solution to this sample representativeness issue is to adopt a statistical model and estimate the model not only with the return time series but also with some other informative series (e.g., Avdis and Wachter, 2017). The key idea is that although μ_t and $\sqrt{v_t}\epsilon_{t+1}$ are not separately observed, we are able to statistically disentangle them if we observe an informative signal about either of the two components.

The challenge of adopting this approach in the context of macro-announcement returns is that it requires a signal that is informative at the daily frequency. We have no choice but to work with the daily return data, as we compare announcement days vs. non-announcement days. Commonly used informative signals such as the the price-dividend ratio and macroeconomic indicators are observed at a much lower frequency (monthly or quarterly) and, hence, are not particularly helpful.

The innovation of our approach is in exploiting asymmetric volatility. Prior studies find that asymmetric volatility, also known as the "leverage effect," is one of the most salient features of the equity market. Asymmetric volatility is typically measured as the conditional correlation between the market index return and its variance. Options data strongly support the existence of asymmetric volatility at the daily frequency, implying that this correlation is significantly negative, ranging from -0.6 to -0.8.

Why is asymmetric volatility useful in separating out μ_t and $\sqrt{v_t}\epsilon_{t+1}$? It essentially captures the conditional correlation between the return innovation and the variance innovation:

$$\mathrm{Corr}_t\left(rx_{t+1}, v_{t+1}\right) = \mathrm{Corr}_t\left(\sqrt{v_t}\epsilon_{t+1}, v_{t+1} - E_t\left[v_{t+1}\right]\right).$$

If the market variance is hit by a negative innovation (i.e. $v_{t+1} - E_t\left[v_{t+1}\right] < 0$), it is highly probable that the market return experiences a positive innovation (i.e. $\sqrt{v_t}\epsilon_{t+1} > 0$), and vice versa. Therefore, the variance innovation actually serves as a highly informative signal about the return innovation.[8]

Taking a step further, this suggests that the variance innovation contains important information about the conditional equity premium μ_t, as it can be obtained by subtracting $\sqrt{v_t}\epsilon_{t+1}$ from rx_{t+1}. In sum, if macro-announcement days tend to coincide with negative (positive) innovations to the market variance in-sample, return innovations should be, on average, positive (negative) over those days. As a result, the average of excess returns over macro-announcement days would overestimate (underestimate) the true macro-announcement premium.

[8] Asymmetric volatility is a common feature of most dynamic equilibrium models with Epstein-Zin preferences, such as Bansal and Yaron (2004), Drechsler and Yaron (2011), Bansal, Kiku, and Yaron (2012), Wachter (2013), and Seo and Wachter (2019). Under recursive utility, shocks to a model's variance-related state variables are negatively priced, creating a negative correlation between variance shocks and return shocks.

3.3.1 Statistical Model and Estimation

We introduce a simple statistical model, which puts more structure into equation (3.2.1). The right-hand side of the equation has three components: (i) the conditional equity premium μ_t, (ii) the conditional variance v_t, and (iii) the return shock ϵ_{t+1}.

First, we specify the conditional equity premium in such a way that can accommodate the existence of an extra premium over macro-announcement days. For parsimony, our benchmark setup assumes that μ_t can take only two values:

$$\mu_t \quad = \quad \mu + \gamma I_t, \quad \text{where} \quad I_t = \begin{cases} 1 & \text{if } t+1 \in \mathbb{M}, \\ 0 & \text{otherwise.} \end{cases} \tag{3.3.1}$$

Namely, the conditional equity premium is $\mu_t = \mu$ on non-announcement days whereas it is $\mu_t = \mu + \gamma$ on macro-announcement days. Here, I_t is a dummy variable that takes a value of 1 if time t is a pre-macro-announcement day (the day immediately preceding a macro-announcement day) and a value of 0 otherwise. Therefore, the term γI_t acts like a fixed effect, allowing for the potential existence of an added premium, the so-called macro-announcement premium. If the conditional equity premium over macro-announcement days is, on average, higher than that over non-announcement days, the coefficient γ should bridge the gap and be estimated as a significantly positive number. In Section 3.3.4, we show that our results are robust to different specifications of μ_t; we reach the same conclusion when μ_t varies with a proxy for market uncertainty (e.g., v_t).

In the data, the patterns of the second moment of returns do not appear to be distinct between macro-announcement days vs. non-announcement days. As discussed in the introduction, various ex-ante conditional variance measures (including daily GARCH volatilities as well as model-free risk-neutral volatilities over 1-week to 1-month horizons) barely change over macro-announcement days on average. Hence, we take a simple approach and assume that the daily conditional variance v_t follows a discrete-time version of a mean-reverting square-root process (Cox, Ingersoll, and Ross, 1985):

$$v_{t+1} \quad = \quad (1-\phi)\bar{v} + \phi v_t + \sigma\sqrt{v_t}u_{t+1}, \tag{3.3.2}$$

where u_{t+1} is an iid standard normal shock; $\bar{v}$, ϕ, and σ represent the long-run mean, persistence, and volatility of the conditional variance process, respectively. In Section 3.4.1, we consider an alternative volatility specification where the conditional variance is expected to rise and fall before and after macro-announcements.

Lastly, we assume that the return shock ϵ_{t+1} follows an iid normal shock, similar to u_{t+1}. These two shocks, however, may be contemporaneously correlated. The correlation coefficient between the two shocks is denoted as $\mathrm{Corr}_t\left(\epsilon_{t+1}, u_{t+1}\right) = \rho$. Similar to the model of Heston (1993) and its many extensions, asymmetric volatility is generated when the correlation between the return shock and the volatility shock is negative:

$$\mathrm{Corr}_t\left(rx_{t+1}, v_{t+1}\right) = \mathrm{Corr}_t\left(\epsilon_{t+1}, u_{t+1}\right) = \rho \quad < \quad 0.$$

We estimate this model via maximum likelihood estimation (MLE). The main challenge is in estimating the time series of the conditional variance. In our benchmark setup, we simplify our estimation by assuming that the daily conditional variance v_t shares a linear relationship with the squared VIX:

$$\mathrm{VIX}_t^2 = k_0 + k_v v_t. \tag{3.3.3}$$

This relation is true under many models with an affine structure (Duffie, Pan, and Singleton, 2000). In Section 3.5, we show that this assumption is not critical to our results. We obtain essentially the same results even if we filter out the conditional variance directly via Markov Chain Monte Carlo (MCMC) methods.

We define the set of model parameters to be estimated as $\Theta = [\rho, \mu, \gamma, \bar{v}, \phi, \sigma, k_0, k_v]$. To construct the log-likelihood function, we first calculate the following transition probability:

$$\mathcal{L}_{t+1} = P\left(rx_{t+1}, \mathrm{VIX}^2_{t+1} \mid rx_t, \mathrm{VIX}^2_t; \Theta\right).$$

Since the excess log return and the squared VIX conditionally follow a bivariate normal distribution under our model, the transition probability can be expressed as:

$$\mathcal{L}_{t+1} = \frac{1}{2\pi k_v \sigma v_t \sqrt{1-\rho^2}} \exp\left(-\frac{\hat{\epsilon}^2_{t+1} - 2\rho\hat{\epsilon}_{t+1}\hat{u}_{t+1} + \hat{u}^2_{t+1}}{2\left(1-\rho^2\right)}\right),$$

where

$$\hat{\epsilon}_{t+1} = \frac{rx_{t+1} - \mu - \gamma I_t}{\sqrt{v_t}} \quad \text{and} \quad \hat{u}_{t+1} = \frac{v_{t+1} - (1-\phi)\bar{v} - \phi v_t}{\sigma\sqrt{v_t}}.$$

Given a set of parameter values, the values of v_t and v_{t+1} are obtained from the squared VIX, according to equation (3.3.3). Finally, the log-likelihood function is expressed as

$$\log \mathcal{L} = \sum_{t=0}^{T-1} \log \mathcal{L}_{t+1},$$

where T represents the sample size or the total number of days in our sample. We estimate our model parameters by maximizing the log-likelihood function $\log \mathcal{L}$.

3.3.2 Main Estimation Results

Panel A of Table 3.1 reports the parameter estimates of our model with $\mathbb{M} = \mathbb{M}_{\text{All}}$ based on the sample period of 1990-2018. Their robust standard errors are also provided in square brackets.

First of all, we confirm the phenomenon of asymmetric volatility. The correlation coefficient ρ is estimated to be -0.711. This value is not only statistically significant, but also largely consistent with prior studies on option pricing. This strong correlation suggests that the variance innovation $(v_{t+1} - \mathbb{E}_t[v_{t+1}]) = \sigma\sqrt{v_t}u_{t+1}$ is highly informative about the return innovation $(rx_{t+1} - \mathbb{E}_t[rx_{t+1}]) = \sqrt{v_t}\epsilon_{t+1}$, allowing us to separate out the conditional equity premium μ_t from the realized excess return rx_{t+1}.

The estimated parameters governing the return variance dynamics are reasonable. As shown in the table, $\bar{v}$, the long-run mean of v_t, is estimated to be $1.043\%^2$. This corresponds to an annual volatility of $\sqrt{1.043 \times 252} = 16.21\%$. The variance process is persistent with a daily autocorrelation ϕ of 0.985. The volatility of variance parameter σ is estimated to be 0.192%. The coefficients k_0 and k_v determine the relation between the squared VIX and the daily ex-ante physical return variance v_t. In our estimation, we divide the squared VIX by (252×100^2) to express it in daily variance terms. This facilitates the interpretation of coefficients k_0 and k_v by making the units for v_t and the squared VIX identical. We find that the slope coefficient k_v is 1.2, namely, larger than 1. This makes sense since the squared VIX captures the risk-neutral variance, which is larger and more volatile than its physical counterpart. The intercept k_0 is estimated to be $0.301\%^2$.

The two parameters μ and γ concern the conditional equity premium $\mu_t = \mu + \gamma I_t$. The point estimate for μ is 2.953 bp and is statistically significant. In contrast, our estimation indicates that γ is not statistically different from zero. The point estimate is actually negative: -0.815 bp. That is, according to our estimated model, macro-announcement days do not feature an extraordinarily high equity premium, beyond the level we observe on non-announcement days.

In Panel B of Table 3.1, we compare and contrast the average excess returns and conditional equity premiums over different sub-samples. First, the average excess return over macro-announcement days is 9.93 bp per day (equivalently, 25.04% per year), whereas the one over the entire sample period of 1990-2018 is 2.38 bp per day (5.99% per year). The average excess return over non-announcement days is only 1.29 bp per day (3.26% per year). As a result, we can see that more than 50% of the whole equity premium is attributable to announcement days. Given that there are only 915 macro-announcement days out of the 7,306 trading days (approximately 13%) in our sample, this is a considerably large share.

Prior studies have made an implicit assumption that a high average macro-announcement return translates into a high macro-announcement premium. Panel B, however, shows that this might not be the case. The average conditional equity premium is similar to the average excess return over the entire sample (2.73 vs. 2.38 bp). However, the average conditional equity premium is much smaller than the average excess return over the macro-announcement sample (1.21 vs. 9.93 bp). Consequently, only 5% of the equity premium actually originates from macro-announcement days. Our results signify that a high average excess return and a high equity premium do not necessarily go hand in hand.

Then, how can we explain the gap between the small equity premium and the extremely large average excess return over macro-announcement days? Table 3.2 shows that the variance shock u_{t+1} is, on average, negative on macro-announcement days in-sample: the average variance shock is -0.15. Due to the presence of strong asymmetric volatility, this tells us that the return shock ϵ_{t+1} is, on average, positive on macro-announcement days: the average return shock is 0.09. Since both u_{t+1} and ϵ_{t+1} are assumed to be standard normal shocks, both their averages over the macro-announcement sample follow a mean-zero normal distribution with a standard deviation of $\sqrt{1/n(\mathbb{M})} = \sqrt{1/915} \simeq 0.03$. Hence, we can see that the average shock values of -0.15 and 0.09 are significant and statistically different from zero.

Consequently, the return innovation $\sqrt{v_t}\epsilon_{t+1}$ is estimated to be significantly positive on average over macro-announcement days. Since the return innovation does not average out to zero but remains as a positive quantity, the average excess return exaggerates the true risk premium, as discussed in Section 3.2. In fact, Table 3.2 indicates that the average return innovation $\sqrt{v_t}\epsilon_{t+1}$ is 8.72 bp, exactly accounting for the gap between the average excess return (9.93 bp) and the equity premium (1.21 bp) over macro-announcement days.[9]

So far, our analysis has relied on the main macro-announcement sample $\mathbb{M} = \mathbb{M}_{\text{All}}$, which consists of the three types of macroeconomic events. Do we reach the same conclusion if we choose only one specific type of macroeconomic news among $\mathbb{M}_{\text{FOMC}}$ (FOMC decisions), $\mathbb{M}_{\text{Empl}}$ (employment figures), and $\mathbb{M}_{\text{Price}}$ (price indices)? In Table 3.3, we provide the answer by estimating the model under these three alternative definitions of $\mathbb{M}$.[10]

Panel A of Table 3.3 reports the parameter estimates for μ and γ together with their robust standard errors in square brackets. We omit other parameters, as they are practically identical to the ones under the benchmark estimation in Table 3.1: the asymmetric volatility parameter ρ as well as variance-related parameters $\bar{v}$, ϕ, σ, k_0, and k_v are determined by the return and volatility time series over the entire sample period and, thus, are not critically affected by the definition of macroeconomic announcements. The results demonstrate that different types of macroeconomic news do not lead to different outcomes in the estima-

[9] Note that we do not observe a similar issue in the entire sample. Both variance and return shocks roughly average out to zero: the average variance and return shocks are 0.01 and -0.01, respectively. The resulting average return innovation is only -0.35 bp, making the average excess return (2.38 bp) and the equity premium (2.73 bp) fairly close in the entire sample.

[10] See Appendix A for more details about the sample construction.

tion. Similar to our benchmark case, the point estimates for μ and γ are around 3 bp and -2 bp. Most importantly, the values for γ are still statistically insignificant for all three types of macroeconomic news. This implies that conditional equity premiums on macro-announcement days are not statistically different from those on non-announcement days, regardless of how we define the macro-announcement sample.

Panel B of Table 3.3 juxtaposes average excess returns and conditional equity premiums over the three types of macroeconomic announcements. The average excess return over days with FOMC announcements is extremely high, even exceeding 25 bp per day (or 63% per year). These days account for more than 30% of the total sum of excess returns, accumulated over the entire period of 1990-2018. Relatively speaking, days with employment figures and price indices exhibit low average excess returns (6.38 and 4.29 bp, respectively), but they are still much larger than the full-sample average (2.38 bp).

Once again, we find that high average excess returns over macro-announcement days are not a manifestation of high realizations of the conditional equity premium over those days. Under all three definitions of $\mathbb{M}$, the average equity premiums are estimated to be lower than 1 bp, sharply contrasting themselves with the average excess returns. These results corroborate what we find in Table 3.1 and Table 3.2: high excess returns over macro-announcement days are mostly attributable to on-average positive returns innovations that are associated with on-average negative variance innovations. The true risk premium over macro-announcement days is, in fact, not abnormally large.

3.3.3 The Role of Asymmetric Volatility

In our estimation, asymmetric volatility plays a key role in measuring the true magnitude of the macro-announcement premium. To highlight its importance, we re-estimate our model by fixing the value of ρ at 0, -0.25, -0.50, and -0.75 and check how the contribution of macro-announcement days to the full equity premium varies. Table 3.4 summarizes the results.

In this exercise, we put the restriction that the average excess return and the average equity premium over the entire sample are identical:

$$\frac{1}{n(\mathbb{U})}\sum rx_{t+1} = \frac{1}{n(\mathbb{U})}\sum [\mu + \gamma I_t].$$

As a result, the full equity premium is consistently estimated as 2.38 bp regardless of the ρ value.[11] This helps us better understand how the equity premium is divided into macro-announcement days vs. non-announcement days, depending on the value of ρ.

From Table 3.4, we find that the macro-announcement premium monotonically increases as the magnitude of ρ gets smaller. When ρ is set to -0.75, we can see that the macro-announcement premium is less than 1 bp per day, merely occupying 2.27% of the entire equity premium. However, the macro-announcement sample's share of the total equity premium goes up to 16.06%, 28.89%, and 40.82% when ρ takes -0.5, -0.25, and 0. In particular, when ρ is set to 0, the macro-announcement premium is 7.75 bp per day, which is quite close to the average excess return over macro-announcement days in the data.

These results are intuitive. The magnitude of the correlation coefficient ρ captures the informativeness of a variance innovation as a signal about the contemporaneous return innovation. During our sample period, the return variance, estimated based on the VIX, experiences on-average negative innovations on macro-announcement days. If ρ is larger in size, these negative variance innovations are more likely to translate into positive return innovations, increasing the discrepancy between the average excess return and the average

[11] Without this restriction, the average equity premium varies with ρ. In our estimation, the return and VIX time series are directly observed. Hence, no matter which ρ value we choose, the actual correlation between the return and variance time series in the data still remains unchanged. If we force ρ to be an arbitrary value that is far away from the true correlation implied by the return and VIX time series, the MLE can result in fairly different average equity premiums.

equity premium. In other words, when the degree of asymmetric volatility becomes greater, the true macro-announcement premium ends up shrinking further, relative to the average excess return over macro-announcement days.

It is worth emphasizing that the macro-announcement premium and the average excess return become fairly close to each other in value when asymmetric volatility does not exist (i.e. $\rho = 0$). In this hypothetical case, the variance process has no information about return innovations. Therefore, despite the fact that variance innovations are, on average, negative in the macro-announcement sample, the model is estimated in a way that return innovations are still symmetric around zero on macro-announcement days.

Conversely, our finding suggests that estimating the macro-announcement premium using the average excess return only makes sense if asymmetric volatility is weak in the data. Clearly, this is not the case. Asymmetric volatility is one of the most robust features of the equity market. Given that ρ is much closer to -1 than it is to 0, Table 3.4 leads us to conclude that the true macro-announcement premium should be much lower than what the average excess return implies.

3.3.4 Different Risk Premium Specifications

In this section, we check the robustness of our results to different specifications of the conditional equity premium μ_t. In our benchmark case, the conditional equity premium only takes two values: μ on non-announcement days and $\mu + \gamma$ on macro-announcement days. What would happen if we adopt alternative risk premium specifications?

First, we consider a setup where the conditional equity premium varies with the level of uncertainty in the market. A common proxy for market uncertainty is the market variance v_t. Thus, we extend our benchmark specification for μ to include an additional term that is proportional to v_t:

$$\mu_t \quad = \quad \mu + \eta v_t + \gamma I_t, \quad \text{where} \quad I_t = \begin{cases} 1 & \text{if } t+1 \in \mathbb{M}, \\ 0 & \text{otherwise.} \end{cases}$$

The conditional equity premium in this setup is affine in v_t. This is a standard assumption that captures the intuition that the equity premium is a compensation for taking on uncertainty (typically proxied by the return variance). As in the benchmark case, the term γI_t accommodates the potential existence of an extra premium over macro-announcement days. If macro-announcement events are special and bring about an extra premium beyond the level that the conditional variance v_t can explain, the coefficient γ would be estimated as a significantly positive value.

Panel A-(1) of Table 3.5 presents the parameter estimates via MLE under this alternative specification. The estimate for μ is 2.941 bp, which is very close to the benchmark case. The estimate for γ also takes on a similar value: the point estimate is -1.743 bp and still insignificant. As can be seen in Panel B-(1), the resulting average equity premium over macro-announcement days is only 1.23 bp, accounting for less than 6% of the total equity premium.

It is worth noting that the estimated value for η is positive but statistically indistinguishable from zero. Most theoretical frameworks, exemplified by the ICAPM of Merton (1973), imply that the conditional equity premium and the conditional variance share a strong positive relation. This intertemporal risk-return relation is, however, not as straightforward in the data as theory predicts because the conditional variance process is latent. In fact, empirical evidence is mixed: some studies find the relation to be positive whereas others find it to be insignificant or even negative.[12] Our estimation reveals that the relation is positive

[12]There is a vast literature on the empirical examination of this relation. Examples include, but are not limited to, Glosten, Jagannathan, and Runkle (1993), Harvey (2001), Brandt and Kang (2004), Ghysels,

but statistically insignificant at the daily frequency, consistent with French, Schwert, and Stambaugh (1987) and Campbell and Hentschel (1992).

In Panels A-(2) and B-(2) of Table 3.5, we consider another alternative specification for μ_t by setting μ in equation (3.3.1) to zero:

$$\mu_t \quad = \quad \eta v_t + \gamma I_t, \quad \text{where} \quad I_t = \begin{cases} 1 & \text{if } t+1 \in \mathbb{M}, \\ 0 & \text{otherwise.} \end{cases} \tag{3.3.1}$$

That is, we assume that the conditional equity premium μ_t is directly proportional to the conditional variance v_t with zero intercept, provided that macro-announcement days are not special (i.e. $\gamma = 0$). Motivated by the ICAPM, the restriction of zero intercept is frequently imposed in a reduced-form option pricing model (see, e.g., Heston and Nandi, 2000).

Panel A-(2) indicates that fixing μ at zero not only makes the estimate for η much larger (a point estimate of 2.386) but also statistically significant (a standard error of 0.945). This implies that if the return is expected to be more volatile between time t and time $t+1$ (i.e. high v_t), the excess return rx_{t+1} is driven by a higher drift (i.e. higher μ_t). However, γ still remains statistically insignificant. Panel B-(2) shows that the resulting macro-announcement premium is only 2.08 bp, occupying less than 10% of the total equity premium over the entire sample.

In sum, we conclude that our results are robust to different specifications of μ_t. Regardless of how we specify the conditional equity premium, the macro-announcement premium is small; the large average excess return over macro-announcement days that we observe in the data is simply attributable to the on-average positive realizations of return innovations.

Santa-Clara, and Valkanov (2005), Guo and Whitelaw (2006), Ludvigson and Ng (2007), Lundblad (2007), and Pastor, Sinha, and Swaminathan (2008).

3.4 Discussion

3.4.1 Uncertainty Resolution Around Macro-Announcements

The gist of Section 3.3 is that the macro-announcement premium is actually not abnormally large if we correctly take asymmetric volatility into account in the estimation. Our statistical model reveals that over the macro-announcement sample, the average variance innovation is negative, which suggests that the average return innovation is positive. As a result, we conclude that the average excess return overstates the true risk premium over macro-announcement days.

Under what scenario might our conclusion fail to hold? Hypothetically, the conditional variance v_t is expected to increase before macroeconomic announcements due to heightened uncertainty associated with the announcement outcomes and is expected to fall afterward as uncertainty resolves. That is, the small average decline in the conditional variance over macro-announcement days might not be due to negative variance innovations (i.e. unexpected falls), but due to negative drifts (i.e. expected falls). If variance innovations average out in-sample, it is likely that return innovations do so as well. In such a scenario, the average excess return no longer exaggerates the average equity premium.

In this section, we take this alternative story based on uncertainty resolution seriously and check if it is in line with the data. To this end, we need a different specification of v_t that can accommodate an expected rise and a subsequent fall in the conditional variance around macro-announcements. For parsimony, we assume that when the conditional equity premium μ_t goes up and down by γ before and after macro-announcements (equation (3.3.1)), the conditional variance v_t is also subject to a positive and a negative drift, say, ξ:

$$v_t = (1-\phi)\bar{v} + \phi v_{t-1} + \sigma\sqrt{v_{t-1}}u_t + \xi\tilde{I}_t, \quad \text{where} \quad \tilde{I}_t = \begin{cases} 1 & \text{if } t+1 \in \mathbb{M}, \\ -1 & \text{if } t \in \mathbb{M}, \\ 0 & \text{otherwise.} \end{cases} \tag{3.4.1}$$

Suppose that time t represents today. There are three cases. (i) If $t+1 \in \mathbb{M}$, today is a pre-announcement day. The conditional variance is expected to increase by ξ (aside from mean reversion) and the conditional equity premium rises by γ. (ii) If $t \in \mathbb{M}$, today is an announcement day. As uncertainty resolves, the conditional variance is expected to decrease by ξ and the conditional equity premium falls by γ. (iii) For t in all other cases, the conditional variance follows exactly the same square-root process as in equation (3.3.2).[13]

We estimate this model via MLE. Table 3.6 reports the parameter estimates for μ, γ, $\bar{v}$, and ξ along with their robust standard errors in square brackets. First, we consider the case with the main macro-announcement sample $\mathbb{M} = \mathbb{M}_{\text{All}}$, which consists of all three types of macroeconomic events. We find that the estimate for the variance drift ξ is 0.013$\%^2$ and is statistically significant (a standard error of 0.005$\%^2$). This implies that under the new variance specification, the average drop in v_t over macro-announcement days is partly attributable to the negative drift. Note that the macro-announcement premium γ is estimated to be 2.666 bp, which is larger compared to the benchmark case (Table 3.1). The increase in γ is not surprising. Allowing for a drift in v_t over macro-announcement days reduces the size of the average variance innovation, which, through asymmetric volatility, translates into a smaller average return innovation in-sample. This decreases the discrepancy between the average excess return and the average equity premium. That is, the estimate for γ becomes larger.

[13] We also consider a slightly different specification where the conditional variance is subject to a negative drift (ii) but not subject to a positive drift (i). The estimation results are similar.

However, γ is still statistically insignificant; the standard error is 2.960 bp. Why is the announcement premium γ still much smaller than what the average excess return implies? This is because the volatility drift ξ cannot capture the entire variance drop in the macro-announcement sample. Variance innovations are still, on average, negative, implying that return innovations are still, on average, positive. Therefore, the main message of our paper still remains: the true risk premium over macro-announcement days is meaningfully lower than the average excess return. This conclusion is robust to using various definitions of the macro-announcement sample ($\mathbb{M}_{\text{FOMC}}$, $\mathbb{M}_{\text{Empl}}$, or $\mathbb{M}_{\text{Price}}$), as shown in Table 3.6.

Recall that our analysis relies on a statistical model. The estimation procedure freely chooses the parameter values that fit the data best (or maximize the likelihood); no restrictions are imposed on model quantities such as μ, γ, $\bar{v}$, and ξ, unlike the case in equilibrium models. Nonetheless, we can still compare the relative magnitudes of the parameter estimates to see whether an equilibrium model with uncertainty resolution around macro-announcements can potentially support the estimation outcome in Table 3.6.

Of interest is what the parameter estimates imply about the market prices of variance risk, which measure how much equity premium investors demand per each unit of variance. On non-announcement days, the equity premium of μ is associated with the average variance of $\bar{v}$, implying that the market price of variance risk is about $\mu/\bar{v} \simeq 2.32$. That is, investors demand a 2.32 bp equity premium per $1\%^2$ unit of variance. What is puzzling is that the market price of variance risk surges on macro-announcement days. Since the extra equity premium of γ is associated with the extra variance of ξ, the market price of variance risk is $\gamma/\xi \simeq 205$, which is nearly 90 times larger.

The staggering difference in the market prices of variance risk epitomizes the problem faced by models with uncertainty resolution such as Savor and Wilson (2013), Ai and Bansal (2018), and Wachter and Zhu (2020). The change in the conditional variance on macro-announcement days is simply too small to justify high macro-announcement returns through the equity premium channel. For example, the average VIX on pre-announcement days is only 19.50%, which is not so much higher than 19.27%, the average VIX on announcement days. Equivalently, the average squared VIX only falls about $10\%^2$, from $442.20\%^2$ to $432.83\%^2$ on announcement days. This volatility gap is too small to justify the large gap in the conditional equity premium (8.64 bp daily or 22% annually). In other words, it is too big of an extra premium for such little change in volatility.

One might suspect that this extremely steep relation between the conditional equity premium and volatility is a result of investors' different risk attitudes over macro-announcement days. Perhaps, investors are more risk-averse those days, creating a high equity premium despite a low level of uncertainty proxied by return variance. However, this argument would not work in an equilibrium model where the return dynamics are endogenously determined. A higher market price of risk not only raises the conditional equity premium, but also the conditional volatility (and the risk-neutral volatility, such as the VIX).[14] Unless we rely on a behavioral explanation, this poses a considerable challenge to virtually all dynamic equilibrium models in which the high equity premium and the high return volatility originate from the same mechanism.

In sum, the main challenge with the uncertainty-resolution-based explanation lies in the ex-ante second moment of returns. If the high average macro-announcement return is a manifestation of a high average equity premium, why isn't the forward-looking volatility also high? Although our statistical setup, which features an expected rise and fall in variance around macro-announcements, is well estimated with the data, the resulting parameter values are not well supported by equilibrium models with uncertainty resolution. In Section 3.4.3, we further discuss the model implications of our results.

[14] In fact, the implied VIX under the models of Savor and Wilson (2013), Ai and Bansal (2018), and Wachter and Zhu (2020) all well exceed 30% or even more prior to macro-announcements.

3.4.2 Joint Explanation of Excess Returns and Changes in the VIX

In this section, we provide a joint explanation of the first and second moments of returns over macro-announcement days. As discussed in Section 3.4.1, recent state-of-the-art models, which generate large macro-announcement returns using a large announcement premium, struggle to justify the small average decline in the VIX on announcement days. We show that high macro-announcement returns are no longer at odds with the patterns of the conditional variance, however, if we assume that they originate from on-average positive return innovations. Moreover, the patterns of excess returns and changes in the VIX over the macro-announcement sample are well explained by those over the non-announcement sample when we control for asymmetric volatility.

The main intuition from Section 3.3 is that macro-announcement days are not so much different from non-announcement days. This is in sharp contrast with the uncertainty-resolution-based explanation, in which macro-announcement days are special and operate with a separate mechanism. To illustrate, Figure 3.2 examines the relationship between daily excess returns (in bp) and contemporaneous changes in the squared VIX (in $\%^2$) in the 1990-2018 period.[15] As expected, these two variables are strongly negatively correlated because of the presence of asymmetric volatility. Surprisingly, the scatter plot in Panel A reveals that macro-announcement days (yellow circle markers) and non-announcement days (blue cross markers) show similar patterns. If we plot the trend lines of the two sub-samples, they almost coincide with each other.

Specifically, Panel A of Table 3.7 presents the results of regressing excess returns on contemporaneous changes in the squared VIX over the two sub-samples as well as over the

[15] Like in Panel B of Figure 3.1, we plot the middle 99% to facilitate comparison between the macro-announcement and the non-announcement samples.

full sample:

$$rx_{t+1} = \text{const.} + \beta \Delta \text{VIX}^2_{t+1} + e_{t+1}.$$

Over macro-announcement days, the slope coefficient β is estimated to be -0.75, implying that a $1\%^2$ decrease in ΔVIX^2 is associated with a 0.75 bp increase in the excess return. These two variables are significantly correlated with a correlation coefficient of -0.73, resulting in an R^2 value of $(-0.73)^2 \simeq 0.53$. Interestingly, we obtain very similar results from the other two samples: the slope coefficient β is estimated to be about -0.72 and the R^2 value is estimated to be about 0.52, both in the non-announcement sample and in the full sample. In other words, splitting the data into the two sub-samples does not seem to make any meaningful difference.

More formally, Panel B of Table 3.7 runs the following fixed effects regression:

$$rx_{t+1} = \text{const.} + \beta_1 \Delta \text{VIX}^2_{t+1} + \beta_2 I_t + \beta_3 \Delta \text{VIX}^2_{t+1} \times I_t + e_{t+1},$$

where, like we define in Section 3.3.1, I_t is a dummy variable that takes a value of 1 if time t is a pre-scheduled macro-announcement day. While the loading on ΔVIX^2 is still significant and remains at a similar level (-0.71), the loading on I and the loading on $\Delta \text{VIX}^2 \times I$ are statistically insignificant. This regression result has two clear implications. First, consistent with our estimation results in Section 3.3, excess returns are not abnormally high over macro-announcement days, after controlling for changes in the squared VIX. Second, the distinction between macro-announcement days vs. non-announcement days does not play a role in the empirical relation between excess returns and changes in the squared VIX.

Figure 3.2 makes this point clear. In Panel A, the two points that represent the averages of the macro-announcement and non-announcement samples are located so closely that they are barely distinguishable. To compare these two points, Panel B zooms into the center of Panel A: the averages of the macro-announcement and non-announcement samples are indicated by the yellow circle and the blue cross, respectively. Note that the point representing the macro-announcement sample is right next to the regression line obtained from the non-announcement sample (black solid line). This suggests that non-announcement days with a similar change in the squared VIX compared to the average macro-announcement day exhibit a similar level of excess returns on average. The predicted excess return from the non-announcement trend is 8.99 bp (red triangle), which is as high as the average excess return of 9.93 bp over macro-announcement days.

Going a step further, Figure 3.3 confirms that the patterns of macro-announcement days can be well replicated by randomly drawing from non-announcement days. There are a total of 915 macro-announcement days during the 1990-2018 period (Appendix A). Thus, we randomly choose 915 days out of non-announcement days and calculate their average excess return and average change in the squared VIX. We repeat this one million times and display the results as a scatter plot with blue cross markers in Panel A. The other three panels show the results when we construct one million random samples of non-announcement days with a sample size of 232 (FOMC decisions; Panel B), 348 (employment figures; Panel C), and 348 (price indices; Panel D). In each panel, we can see that the black hollow circle, which indicates the actual data point from the given macro-announcement sample, is located not only within the scatter plot, but also in close proximity with the black solid line representing the trend line.

Lastly, we point out that the y-intercept of the regression line from the non-announcement sample in Figure 3.2 (which is one of the constant coefficients found in Table 3.7) is 2.29 bp. This implies that if we consider a long representative sample of non-announcement days whose variance innovations average out (i.e. $\Delta\mathrm{VIX}^2 \approx 0$ on average), its average excess return would be 2.29 bp daily, or 5.78% annually, which is close to the 6% unconditional equity premium in the postwar sample. This counters the claim that the equity premium is mostly realized over macro-announcement days. Even if we exclude all macro-announcement days and construct a sample only consisting of non-announcement days, the average excess return can capture the full equity premium, as long as the sample is long and representative. Overall, we do not find any definitive empirical evidence that macro-announcement days and non-announcement days operate with separate mechanisms.

3.4.3 Implications for Macro-Finance Models

Our results have important implications for macro-finance models. Our analysis alludes that the large average macro-announcement return might not be a puzzle nor compensation for risk that resolves at macro-announcements; instead, it is an artifact of the small size of the macro-announcement sample. Under our explanation, there is no need for a new model to account for the large average macro-announcement return, as it does not represent a large risk premium. Traditional models are just fine.

To be more concrete, Figure 3.4 explicitly shows that the return and VIX patterns over macro-announcement days are rationalized by the model of Seo and Wachter (2018b), which does not feature an announcement premium. We simulate the model to generate one million 29-year paths at the daily frequency. Then, we randomly select 915 days from each no-disaster path and calculate the average excess return and the average change in the squared VIX (each blue cross), just like we do for the macro-announcement sample in the data (black hollow circle). Figure 3.4 clearly shows that the average excess return of 9.93 bp that we observe over macro-announcement days is not extraordinarily large under Seo and Wachter (2018b)'s model. Samples with 915 random days from the model frequently generate a similar or even higher level of average excess returns. Furthermore, due to the presence of asymmetric volatility in the model, such samples with high average excess returns exhibit average declines in the squared VIX that are comparable to the data.

Considering our findings in Section 3.3, this is not surprising. It is likely that other traditional dynamic equilibrium models, such as Bansal and Yaron (2004), Drechsler and Yaron (2011), Bansal, Kiku, and Yaron (2012), and Wachter (2013), can also rationalize the first and second return moments over macro-announcement days, if they are properly calibrated. Here, we choose Seo and Wachter (2018b) as an example because it already well matches the properties of daily returns with its original calibration. In the median 29-year no-disaster sample, the average excess return is 2.44 bp and the return standard deviation is 112.98 bp, which are very close to their data counterparts (2.38 bp and 109.08 bp). The model also generates a reasonable degree of asymmetric volatility (-0.66).

In sum, there is nothing puzzling about the large average excess return and the small decrease in the VIX over macro-announcement days, as discussed in Section 3.4.2. While a $10\%^2$ gap in the squared VIX (or a 0.23% gap in the VIX) is not large enough to justify a 8.64 bp gap in the daily conditional equity premium, it is still large enough to justify a 8.64 bp return innovation through asymmetric volatility. In fact, the estimated slope coefficients in Table 3.7 suggest that the $10\%^2$ drop in the squared VIX in the data tends to be associated with a 7-8 bp increase in the excess return, which accounts for the majority of the excess return gap. We reiterate that our explanation relies on the following two empirical facts: (i) the return distribution, regardless of whether it is for macro-announcement days or non-announcement days, is with a very large variance and kurtosis, and (ii) there is a limited number of macro-announcement days. Due to these reasons, it is possible that return and variance innovations do not average out over the macro-announcement sample.

Although we can fully account for the data using a model without a macro-announcement premium, this does not necessarily rule out the existence of such a premium. It is still possible that some portion of high macro-announcement returns, if any, arises from a higher conditional equity premium associated with heightened uncertainty around announcement events. Even if this were the case, our estimation results (especially in Section 3.4.1) suggest that the true macro-announcement premium is likely to be much smaller than what the average excess return indicates. At the least, existing models should be calibrated to produce a much smaller macro-announcement premium.

If one still believes that the high average macro-announcement return comes from a high conditional equity premium over macro-announcement days, he or she would need a model that captures the following three stylized facts: (i) the conditional equity premium is extremely high before macro-announcements, and (ii) economic uncertainty is expected to resolve at macro-announcements, leading to a drop in ex-ante return volatility, but (iii) ex-ante return volatility should not be too high before macro-announcements so that the volatility drop is quite small. Unfortunately, creating such a model would be almost impossible unless we walk away from the existing asset pricing paradigm and start believing that the high equity premium and the high volatility actually arise from different sources of risk. Given all the evidence from empirical and modeling dimensions, isn't it more natural to believe that the high average return over the small sample of macro-announcements is not an indication of a high macro-announcement premium?

3.5 Robustness: Results From a Longer Sample

So far we have examined the true magnitude of the macro-announcement premium based on the sample period of 1990-2018. This is because the time series of the VIX starts from 1990; by assuming that the squared VIX and the ex-ante daily return variance v_t share a linear relationship, we have avoided the necessity of filtering out the latent process v_t.

Not only does this allow us to estimate our model quickly, but it also significantly reduces computational burden for various robustness tests.

What happens if we extend our sample period back to 1961? Even in this longer period, the macro-announcement sample is still small. (See Appendix A.) Is it still the case that the average conditional equity premium over macro-announcement days is not particularly larger than that over non-announcement days? In this section, we answer this question by re-estimating our model based on the sample of 1961-2018 via MCMC estimation.[16]

3.5.1 MCMC Estimation

We define $\Theta = [\Theta_1, \cdots, \Theta_K]$ as the set of model parameters we want to estimate. For notational convenience, let $v = \{v_t\}_{t=0}^{T}$ denote the time series of the latent conditional variance v_t over the sample period. At each point in time t, we observe the data Y_t: we only observe the excess return (i.e. $Y_t = rx_t$) prior to 1990 whereas we observe the excess return together with the VIX (i.e. $Y_t = \left[rx_t, \text{VIX}_t^2\right]$) after 1990. The time series of Y_t is simply abbreviated as $Y = \{Y_t\}_{t=1}^{T}$.

In a Bayesian approach, the ultimate goal is to characterize the joint posterior distribution. According to Bayes' rule, it follows that

$$P(\Theta, v \mid Y) \propto P(Y \mid v, \Theta) P(v \mid \Theta) P(\Theta), \tag{3.5.1}$$

where $P(Y \mid v, \Theta)$ represents the likelihood function, $P(Y \mid v, \Theta)$ captures the distribution of the latent process v, and $P(\Theta)$ denotes the prior distribution. It is clear that drawing from the posterior distribution directly using equation (3.5.1) is technically not possible. Under our model, the expressions for the three distribution functions on the right-hand side

[16]For an overview of MCMC estimation methods for asset pricing models, see Johannes and Polson (2010).

of equation (3.5.1) are extremely complicated, making it impossible to map the posterior into a standard/known distribution. Not to mention that the posterior distribution is with an extremely high dimension (9 parameters and 14,599 latent conditional variances); directly sampling from it is infeasible.

This is where MCMC methods can help us. The Clifford-Hammersley theorem dictates that under some mild regularity conditions, a joint distribution is completely determined by its conditional distributions. The direct implication of this theorem is that the joint posterior $P(\Theta, v \mid Y)$ can be perfectly characterized by the two conditional posteriors $P(\Theta \mid v, Y)$ and $P(v \mid \Theta, Y)$. This theorem can be used more than once: by repetitively applying the Clifford-Hammersley theorem, we can ultimately break our 14,608-dimensional joint posterior into 14,608 one-dimensional marginal conditional posteriors (or so-called full conditional posteriors):

$$P(\Theta_k \mid \Theta_{-k}, v, Y) \qquad k = 1, \cdots, K, \tag{3.5.2}$$

$$P(v_t \mid v_{-t}, \Theta, Y) \qquad t = 0, \cdots, T, \tag{3.5.3}$$

where $\Theta_{-k} = \Theta \backslash \{\Theta_k\}$ is the collection of parameters except for Θ_k and $v_{-t} = v \backslash \{v_t\}$ is the time series of latent conditional variances except for the one at time t. Based on equations (3.5.2) and (3.5.3), MCMC algorithms allow us to build a Markov chain $\left\{\Theta^{(n)}, v^{(n)}\right\}_{n=1}^{N} = \left\{\Theta_1^{(n)}, \cdots, \Theta_K^{(n)}, v_0^{(n)}, \cdots, v_T^{(n)}\right\}_{n=1}^{N}$ that converges to the joint posterior distribution $P(\Theta, v \mid Y)$. Below, we summarize the steps for our entire estimation procedure.

Step 0: Initialization

We initialize the chain $\left\{\Theta^{(0)}, v^{(0)}\right\}$. For each parameter Θ_k, we use its prior mean as the initial value. In order to speed up convergence, we estimate the GARCH model and take the resulting time series of conditional variances as the starting values for $v = \{v_t\}_{t=0}^{T}$.

Step 1: Gibbs sampler based on equation (3.5.2)

Given $\left\{\Theta^{(n-1)}, v^{(n-1)}\right\}$, we draw new values of model parameters $\Theta^{(n)}$ using the Gibbs sampler. To this end, the full conditional posterior of each parameter in equation (3.5.2) should be a known distribution, which we can directly sample from. We adopt the technique of Jacquier, Polson, and Rossi (2004) and re-parametrize the variance process in equation (3.3.2) by introducing $\psi \equiv \rho\sigma$ and $\omega \equiv \sigma^2(1-\rho^2)$:

$$\begin{aligned} v_{t+1} &= (1-\phi)\bar{v} + \phi v_t + \sigma\sqrt{v_t}\left(\sqrt{\rho}\epsilon_{t+1} + \sqrt{1-\rho^2}u_{t+1}^{\perp}\right) \\ &= (1-\phi)\bar{v} + \phi v_t + \psi\sqrt{v_t}\epsilon_{t+1} + \sqrt{\omega v_t}u_{t+1}^{\perp}, \end{aligned}$$

where $u_{t+1}^{\perp}$ is a component of u_{t+1} that is orthogonal to ϵ_{t+1}. The benefit of replacing ρ and σ by ψ and ω is that ρ's conditional posterior has a nonstandard distribution whereas ψ's and ω's do not. Once ψ and ω are sampled, the two original parameters can be recovered by $\rho = \psi/\sqrt{\psi^2+\omega}$ and $\sigma = \sqrt{\psi^2+\omega}$.

Similar to our benchmark setup, we assume that the squared VIX and the ex-ante conditional variance share a linear relationship with some noise e_t:

$$\text{VIX}_t^2 = k_0 + k_v v_t + e_t, \quad \text{where } e \sim N\left(0, s^2\right). \tag{3.5.4}$$

Given that the VIX time series is not available prior to 1990, this specification provides us a simple solution to deal with unequal lengths of our observable data. When the VIX is not available, the above equation is irrelevant and the extraction of v_t entirely relies on the return time series. When the VIX is available, both the VIX and the return data contribute to the filtering of v_t. The magnitude of the noise variance s^2 reflects the informativeness of the VIX.

In sum, we have a total of 9 parameters to be sampled: $\Theta = [\mu, \gamma, \phi, \bar{v}, \psi, k_0, k_v, \omega, s^2]$. In Appendix B, we show that the full conditional posteriors for the first 7 parameters are normal and those for the last 2 parameters are inverse-gamma if their respective priors are chosen to be normal and inverse-gamma (i.e. conjugate priors). We choose the parameters for the priors in a way that the priors are non-committal/non-informative. Appendix B discusses how we sequentially sample these parameters together. Further information about the priors as well as the detailed derivations of the full conditional posteriors are also provided.

Step 2: Metropolis-Hastings based on equation (3.5.3)

Given $\left\{\Theta^{(n)}, v^{(n-1)}\right\}$, we draw a new sample for latent conditional variances $v^{(n)}$ using the Metropolis-Hastings algorithm. We cannot apply the Gibb sampler to equation (3.5.3) because the full conditional posterior of v_t follows a nonstandard distribution with a complicated density function, making it difficult to directly sample from it. Under random walk Metropolis-Hastings, we propose a candidate value for $v_t^{(n)}$ as

$$v_t' = v_t^{(n-1)} + \kappa \xi_t,$$

where ξ_t follows a t-distribution with 1 degree of freedom. The coefficient κ is a scale factor that guarantees that $\kappa\xi_t$ is with a reasonable size relative to conditional variances. Typically, κ is chosen in such a way that the acceptance rate of a proposal is in the range of 20-40% (see Johannes and Polson, 2010). We accept this candidate with probability $\alpha\left(v_t', v_t^{(n-1)}\right)$, or reject it and keep the previous value:

$$v_t^{(n)} = \begin{cases} v_t' & \text{with probability } \alpha\left(v_t', v_t^{(n-1)}\right), \\ v_t^{(n-1)} & \text{otherwise.} \end{cases}$$

The expression for the acceptance probability $\alpha\left(v_t', v_t^{(n-1)}\right)$ is derived in Section B. We obtain $v^{(n)}$ by applying this algorithm to all t values, sequentially from 0 to T.

Last step: Iterations

We repeat Step 1 and Step 2 100,000 times to generate a long chain of $\left\{\Theta^{(n)}, v^{(n)}\right\}$. We discard the first quintile of the chain as a burn-in period. The steady-state distribution represented by the remainder of the chain finally characterizes the joint posterior distribution $P\left(\Theta, v \mid Y\right)$.

3.5.2 MCMC Results

Before expanding our sample period, we run our MCMC procedure with the main sample of 1990-2018 and report the results in Table 3.8. Comparing these results with our main MLE results in Section 3.3.2 can serve as a quick sanity check, as both estimations are based on the same sample period.

Table 3.8 reports the posterior mean as well as 5%, 50%, and 95% percentile values for each parameter. The MCMC results are largely consistent with the MLE results in Table 3.1. The posterior means of $\bar{v}$, ρ, σ, k_0, and k_v are similar to their MLE counterparts, implying similar variance dynamics.[17] Furthermore, we obtain a similar degree of asymmetric volatility ρ (the posterior mean of -0.685 vs. the MLE point estimate of -0.711). Although we assume a non-informative prior with a large variance for each of these parameters, the return and VIX time series make it possible to generate narrow posterior distributions, as can be seen in Table 3.8.

[17] Our estimation results indicate that the value for s^2 is nearly zero at each step of the chain, placing its entire posterior distribution near zero. Due to space constraints, we omit s^2 from Tables 3.8 and 3.9. The tiny magnitudes of s^2 in our estimation suggests that the VIX is highly informative about the daily return variance.

Our utmost interest is in μ and γ, which are the parameters that are associated with the equity premium. Again, our MCMC procedure generates results that are consistent with the main MLE results. According to Table 3.8, the non-announcement premium μ has a posterior mean of 2.480 bp per day (equivalently, 6.25% per year) and is credibly different from zero. In contrast, γ, which captures the gap between the macro-announcement premium and the non-announcement premium, is not credibly different from zero: its 90% credible interval is [-3.805 bp, 1.864 bp], which contains zero. This is also the case for the 95% and 99% credible intervals, as they are even wider. The posterior mean of γ is, in fact, negative (-0.968 bp), in line with the MLE results. From these, we draw the same conclusion as in Section 3.3.2. Although the average excess return is much higher over macro-announcement days than over non-announcement days, the equity premium over macro-announcement days is not statistically different from that over non-announcement days, at least based on the data time series of 1990-2018.

Given that our MCMC procedure generates sensible results, we apply it to a longer sample: we expand our sample back to 1961 and re-estimate the model. The results are reported in Table 3.9. Panel A shows that extending the sample period generates qualitatively similar estimation results. In the 1961-2018 sample, the posterior mean of γ is now positive (1.537 bp). However, we still cannot reject the hypothesis that it is identical to zero: its 90% credible interval is [-0.609, 3.667], which includes zero. This suggests that the macro-announcement premium and the non-announcement premium are not statistically different from each other.

In the 1961-2018 sample, the average excess return over macro-announcement days is 8.64 bp whereas the full sample average is 1.94 bp. As a result, macro-announcement days, which occupy only 13% of all trading days, contribute nearly 60% of the average excess return over the full sample. Similar to our conclusion from the 1990-2018 sample, this does not mean that the equity premium arises mostly over macro-announcement days. Panel B indicates that only 20% of the equity premium is realized over macro-announcement days.

3.6 Conclusion

Macroeconomic news constitutes an important part of investors' information set, making a substantial impact on the equity market. This does not necessarily mean that investors' perceived uncertainty about macroeconomic outcomes resolves only on days with announcements. Market participants are not passive learners who simply wait for announcements. Instead, they acquire information and begin to form their expectations far ahead of time through various channels. The government and the Fed are not passive communicators either. To reduce unnecessary noise, they actively send signals to the market. Hence, it is possible that a significant portion of uncertainty gradually resolves well before announcements, if investors are confident about future announcement outcomes.

Our estimation results are consistent with this story. The equity premium is not realized disproportionately more on macro-announcement days. Although the average macro-announcement return is high, it is not a manifestation of high equity premiums that are associated with uncertainty resolution at macro-announcements. Rather, it is simply the result of return innovations that are not averaged out in-sample due to a small sample size.

In this paper, we offer a joint explanation of the high average excess return and the small average decrease in the VIX over macro-announcement days. At the daily frequency, return innovations are extremely volatile and fat-tailed in the data. Thus, even a small drop in the VIX, like we observe on macro-announcement days, translates into a sizable return innovation through asymmetric volatility. Accounting for both patterns is virtually impossible if high macro-announcement returns originate from high conditional equity premiums:

if the conditional equity premium spikes up due to heightened perceived uncertainty, the conditional volatility, which directly reflects such uncertainty, should also spike up.

The main message of our paper to the macro-finance literature is straightforward. If high ex-post macro-announcement returns are not a manifestation of high ex-ante conditional equity premiums, we do not need a complex model that treats announcement and non-announcement days differently. As we demonstrate, traditional equilibrium models can rationalize the data without hypothesizing an announcement premium. Before creating a new model, we should step back and take the econometrics of macro-announcement returns more seriously.

Appendix

A Data

Our full sample period is from January 1961 to December 2018, including a total of 14,598 trading days. The estimation of our model requires specifying a sample of macro-announcement dates, $\mathbb{M}$. We first construct three samples of macro-announcement dates based on the following three types of pre-scheduled releases of macroeconomic news: (i) FOMC decisions, (ii) employment figures, and (iii) price indices.

The first sample, $\mathbb{M}_{\text{FOMC}}$, collects a total of 582 FOMC announcement dates. We obtain these dates from the Federal Reserve Board website. Since 1994, the FOMC has released a statement outlining its monetary policy decisions after each meeting. In this case, we recognize the days the statements are issued as announcement days. Prior to 1994, however, the FOMC did not directly announce its decisions to the general public. Instead, investors were able to recognize the decisions the next trading day by observing open market operations carried by the Fed's domestic trading desk (the so-called "open market desk"). In this case, we recognize the trading days after the FOMC meetings as announcement days. We do not include conference calls in our sample, as they are not pre-scheduled far ahead of time; they are typically held under special circumstances, such as national emergencies.

The second sample, $\mathbb{M}_{\text{Empl}}$, consists of the first Friday of each month when the Bureau of Labor Statistics (BLS) publishes a monthly report called "Employment Situation." This report draws investors' attention because it contains important employment figures such as the unemployment rate and non-farm payroll. The total number of observations in this sample is 696 (i.e. once a month for 58 years).

Another important type of macroeconomic news is about inflation and price level fluctuations. Each month, the BLS also announces the Consumer Price Index (CPI) and the

Producer Price Index (PPI). Our third sample, $\mathbb{M}_{\text{Price}}$, is constructed based on these two price indices. For the period 1961-1970, we collect CPI announcement dates, as the PPI is not yet available. From 1971, we only collect PPI announcement dates, following Savor and Wilson (2013) and Ai and Bansal (2018). This is because the information revealed by the CPI and the PPI largely overlaps and the announcement of the PPI precedes that of the CPI for most of the months (more than 95% of the time in our sample period).[18] Since we have only one announcement per month, the total number of observations is also 696.

Our main macro-announcement sample is constructed by combining these three types of macroeconomic news: $\mathbb{M}_{\text{All}} = \mathbb{M}_{\text{FOMC}} \cup \mathbb{M}_{\text{Empl}} \cup \mathbb{M}_{\text{Price}}$. Due to overlapping dates, the total number of announcement days in the sample is $1,903 < 582 + 696 + 696$.

Although the full sample period is 1961-2018, Section 3.3 first focuses on a sub-sample period starting from 1990. This is because the time series of the VIX dates back to 1990; as discussed in Section 3.3.1, we exploit the VIX to simplify the estimation of our benchmark model. In this sub-sample period, there are a total of 7,306 trading days, and 915 of them are macro-announcement days (232 for FOMC decisions, 348 for employment figures, and 348 for price indices). The estimation with the full-sample period via MCMC is provided in Section 3.5. We find that our main conclusion remains unchanged regardless of the sample period we choose.

B Detailed Derivations for MCMC Estimation

Gibbs Sampler

In this section, we derive the full conditional posterior in equation (3.5.2) for each parameter $\Theta_k \in \Theta = [\mu, \gamma, \phi, \bar{v}, \psi, k_0, k_v, \omega, s^2]$. We start with the case in which $\Theta_k = \mu$.

[18] Constructing the sample only with CPI announcement dates makes our results even stronger.

According to Bayes' rule, μ's full conditional posterior is expressed as:

$$\begin{aligned} P(\mu \mid \Theta_{-k}, v, Y) &\propto P(Y \mid v, \Theta)\, P(v \mid \Theta)\, P(\mu \mid \Theta_{-k}), \\ &= P(rx \mid v, \Theta)\, P\left(\mathrm{VIX}^2 \mid rx, v, \Theta\right) P(v \mid \Theta)\, P(\mu \mid \Theta_{-k}). \end{aligned}$$

We assume mutually independent priors, which implies that $P(\mu \mid \Theta_{-k}) = P(\mu \mid \Theta \backslash \{\mu\}) = P(\mu)$. Moreover, since $P\left(\mathrm{VIX}^2 \mid rx, v, \Theta\right)$ and $P(v \mid \Theta)$ are constant with respect to μ under our model specification, it follows that

$$P(\mu \mid \Theta_{-k}, v, Y) \propto P(rx \mid v, \Theta)\, P(\mu).$$

Conditional on observing the time series of conditional variances v, the excess return time series rx follows a normal distribution with the following density function:

$$P(rx \mid v, \Theta) = \prod_{t=0}^{T-1} \frac{1}{\sqrt{2\pi(1-\rho^2)v_t}} \exp\left\{ -\frac{\left(rx_{t+1} - \mu - \gamma I_t - \sqrt{v_t}\rho\hat{u}_{t+1}\right)^2}{2(1-\rho^2)v_t} \right\}. \quad \text{(B.1)}$$

Equation (B.1) suggests that the expression for $P(rx \mid v, \Theta)$ can also be viewed as a normal density function with respect to μ with mean $\left[\sum_{t=0}^{T-1} \frac{rx_{t+1} - \gamma I_t - \sqrt{v_t}\rho\hat{u}_{t+1}}{(1-\rho^2)v_t}\right] \left[\sum_{t=0}^{T-1} \frac{1}{(1-\rho^2)v_t}\right]^{-1}$ and variance $\left[\sum_{t=0}^{T-1} \frac{1}{(1-\rho^2)v_t}\right]^{-1}$. We choose a normal prior for μ, which is a conjugate prior for the normal likelihood function, so that the posterior also follows a normal distribution.[19] To make sure that the data speak for themselves, we choose a non-informative prior with a large variance: $N(0, 100)$.

Following the same logic, we are able to write down the full conditional posterior dis-

[19] If the likelihood function is $N(\mu_\ell, \sigma_\ell^2)$ and the prior is $N(\mu_0, \sigma_0^2)$, the posterior follows $N(\mu_p, \sigma_p^2)$ where $\mu_p = \sigma_p^2 \left[\frac{\mu_0}{\sigma_0^2} + \frac{\mu_\ell}{\sigma_\ell^2}\right]$ and $\sigma_p^2 = \left[\frac{1}{\sigma_0^2} + \frac{1}{\sigma_\ell^2}\right]^{-1}$.

tribution for the case in which $\Theta_k = \gamma$:

$$P(\gamma \mid \Theta_{-k}, v, Y) \quad \propto \quad P(rx \mid v, \Theta)\, P(\gamma)\,.$$

The functional form in equation (B.1) also indicates that the likelihood $P(rx \mid v, \Theta)$ is a normal density with mean $\left[\sum_{t=0}^{T-1} \frac{(rx_{t+1}-\mu-\sqrt{v_t}\rho\hat{u}_{t+1})I_t}{(1-\rho^2)v_t}\right]\left[\sum_{t=0}^{T-1} \frac{I_t}{(1-\rho^2)v_t}\right]^{-1}$ and variance $\left[\sum_{t=0}^{T-1} \frac{I_t}{(1-\rho^2)v_t}\right]^{-1}$. We use a diffuse normal conjugate prior with a variance of 100. We conservatively choose the prior mean as 8.64 bp; if the conventional belief that the average excess return is a good proxy for the macro-announcement premium is indeed correct, γ should be close to 8.64 bp, the difference between the average excess return over macro-announcement days and that over non-announcement days.

Now we turn to variance-related parameters Θ_k = ϕ, $\bar{v}$, ψ, or ω. By Bayes' rule, we can show that:

$$P(\Theta_k \mid \Theta_{-k}, v, Y) \propto P(v \mid rx, \Theta)\, P(\Theta_k)\,,$$

where

$$P(v \mid rx, \Theta) = \prod_{t=0}^{T-1} \frac{1}{\sqrt{2\pi\omega v_t}} \exp\left\{-\frac{\left(v_{t+1} - (1-\phi)\bar{v} - \phi v_t - \psi\sqrt{v_t}\hat{\epsilon}_{t+1}\right)^2}{2\omega v_t}\right\}. \quad \text{(B.2)}$$

Equation (B.2) reveals that while the likelihood $P(v \mid rx, \Theta)$ is normal with respect to ϕ,

$\bar{v}$, and ψ, it is inverse-gamma with respect to ω:

$$
\begin{aligned}
\text{w.r.t } \phi : \quad & N\left(\left[\sum_{t=0}^{T-1}\frac{\left(v_{t+1}-\bar{v}-\psi\sqrt{v_t}\hat{\epsilon}_{t+1}\right)(v_t-\bar{v})}{\omega v_t}\right]\left[\sum_{t=0}^{T-1}\frac{(v_t-\bar{v})^2}{\omega v_t}\right]^{-1}, \left[\sum_{t=0}^{T-1}\frac{(v_t-\bar{v})^2}{\omega v_t}\right]^{-1}\right),\\
\text{w.r.t } \bar{v} : \quad & N\left(\left[\sum_{t=0}^{T-1}\frac{\left(v_{t+1}-\phi v_t-\psi\sqrt{v_t}\hat{\epsilon}_{t+1}\right)(1-\phi)}{\omega v_t}\right]\left[\sum_{t=1}^{T-1}\frac{(1-\phi)^2}{\omega v_t}\right]^{-1}, \left[\sum_{t=0}^{T-1}\frac{(1-\phi)^2}{\omega v_t}\right]^{-1}\right),\\
\text{w.r.t } \psi : \quad & N\left(\left[\sum_{t=0}^{T-1}\frac{\left(v_{t+1}-(1-\phi)\bar{v}-\phi v_t\right)\hat{\epsilon}_{t+1}}{\omega\sqrt{(v_t)}}\right]\left[\sum_{t=0}^{T-1}\frac{\hat{\epsilon}_{t+1}^2}{\omega}\right]^{-1}, \left[\sum_{t=0}^{T-1}\frac{\hat{\epsilon}_{t+1}^2}{\omega}\right]^{-1}\right),\\
\text{w.r.t } \omega : \quad & IG\left(\left[\frac{T}{2}-1\right], \left[\sum_{t=0}^{T-1}\frac{\left(v_{t+1}-(1-\phi)\bar{v}-\phi v_t-\psi\sqrt{v_t}\hat{\epsilon}_{t+1}\right)^2}{2v_t}\right]\right),
\end{aligned}
$$

where $IG(a,b)$ represents the inverse-gamma distribution with shape parameter a and scale parameter b.

For ϕ, $\bar{v}$, and ψ, we take a diffuse normal conjugate prior with a variance of 100. Based on the estimation results from the GARCH model, we choose the prior mean of ϕ as 0.97 and that of $\bar{v}$ as $1.2\%^2$. In order to maintain the stationarity of the variance process, we reject the draws of ϕ larger than 0.9999. The prior mean of ψ (which determines ρ) is set to zero, implying no asymmetric volatility *a priori*. For ω, we adopt an inverse-gamma prior, which is conjugate to the inverse-gamma likelihood function.[20] Specifically, we use a non-informative prior $IG(2^{-52}, 2^{-52})$.[21]

Lastly, we consider a VIX-related parameter $\Theta_k = k_0$, k_1, or, s^2. In this case, Bayes' rule implies that

$$
P(\Theta_k \mid \Theta_{-k}, v, Y) \quad \propto \quad P\left(\text{VIX}^2 \mid v, \Theta\right) P(\Theta_k).
$$

[20] If the likelihood function is $IG(a_\ell, b_\ell)$ and the prior is $IG(a_0, b_0)$, the posterior follows $IG(a_p, b_p)$ where $a_p = a_0 + a_\ell - 1$ and $b_p = b_0 + b_\ell$.

[21] We choose 2^{-52} as it is the smallest positive number in a 64-bit IEEE floating-point system.

It follows from equation (3.5.4) that:

$$P\left(\mathrm{VIX}^2 \mid v, \Theta\right) = \prod_{t\in\mathbb{U}_{\mathrm{VIX}}} \frac{1}{\sqrt{2\pi}s} \exp\left\{-\frac{\left(\mathrm{VIX}^2 - k_0 - k_v v_t\right)^2}{2s^2}\right\}, \tag{B.3}$$

where $\mathbb{U}_{\text{VIX}}$ represents the second half of our sample in which the VIX time series is available. From equation (B.3), we observe that the likelihood $P\left(\mathrm{VIX}^2 \mid v, \Theta\right)$ is normal with respect to k_0 and k_v but inverse gamma with respect to s^2:

$$\begin{aligned}
\text{w.r.t } k_0 : \quad & N\left(\left[\frac{s^2}{n(\mathbb{U}_{\mathrm{VIX}})}\sum_{t\in\mathbb{U}_{\mathrm{VIX}}}\frac{\mathrm{VIX}_t^2 - k_v v_t}{s^2}\right], \left[\frac{s^2}{n(\mathbb{U}_{\mathrm{VIX}})}\right]\right), \\
\text{w.r.t } k_v : \quad & N\left(\left[\sum_{t\in\mathbb{U}_{\mathrm{VIX}}}\frac{\left(\mathrm{VIX}_t^2 - k_0\right)v_t}{s^2}\right]\left[\sum_{t\in\mathbb{U}_{\mathrm{VIX}}}\frac{v_t^2}{s^2}\right]^{-1}, \left[\sum_{t\in\mathbb{U}_{\mathrm{VIX}}}\frac{v_t^2}{s^2}\right]^{-1}\right), \\
\text{w.r.t } s^2 : \quad & IG\left(\left[\frac{n(\mathbb{U}_{\mathrm{VIX}})}{2} - 1\right], \left[\sum_{t\in\mathbb{U}_{\mathrm{VIX}}}\frac{\left(\mathrm{VIX}_t^2 - k_0 - k_v v_t\right)^2}{2}\right]\right).
\end{aligned}$$

The priors for k_0, k_v, and s^2 are assumed to be $N(0, 100)$, $N(1, 100)$, and $IG(2^{-52}, 2^{-52})$, respectively, all of which are non-informative.

Metropolis-Hastings

In this section, we derive the expression for the acceptance probability $\alpha(\cdot,\cdot)$ in Section 3.5.1. Random walk Metropolis-Hastings indicates that for each time t, this probability is calculated as

$$\alpha\left(v_t', v_t^{(n-1)}\right) = \min\left[1, \frac{P\left(v_t' \mid v_1^{(n)}, \cdots, v_{t-1}^{(n)}, v_{t+1}^{(n-1)}, \cdots, v_T^{(n-1)}, \Theta^{(n)}, Y\right)}{P\left(v_t^{(n-1)} \mid v_1^{(n)}, \cdots, v_{t-1}^{(n)}, v_{t+1}^{(n-1)}, \cdots, v_T^{(n-1)}, \Theta^{(n)}, Y\right)}\right].$$

Therefore, in order to determine this acceptance probability, it suffices to determine the full conditional posterior $P\left(v_t \mid v_{-t}, Y, \Theta\right)$ up to a constant multiplication.

Note that under our model structure, the full conditional posterior of v_t depends only on the previous and next variances (i.e. v_{t-1}, v_{t+1}), the excess returns in current and next periods (i.e. rx_t, rx_{t+1}), the current level of the squared VIX (i.e. VIX_t^2), and the model parameters (i.e. Θ). Hence, the full conditional posterior can be written as:

$$\begin{aligned} P\left(v_t \mid v_{-t}, \Theta, Y\right) &= P\left(v_t \mid v_{t-1}, v_{t+1}, rx_t, rx_{t+1}, \text{VIX}_t^2, \Theta\right) \\ &\propto P\left(v_{t+1}, rx_{t+1}, \text{VIX}_t^2 \mid v_{t-1}, v_t, rx_t, \Theta\right) P\left(v_t \mid v_{t-1}, rx_t, \Theta\right). \end{aligned}$$

Since VIX_t^2 is independent of v_{t+1} and rx_{t+1} given v_t under our model, the full posterior of v_t reduces to

$$P\left(v_t \mid v_{-t}, \Theta, Y\right) \propto \underbrace{P\left(\text{VIX}_t^2 \mid v_t, \Theta\right)}_{\text{(i)}} \underbrace{P\left(v_{t+1}, rx_{t+1} \mid v_t, \Theta\right)}_{\text{(ii)}} \underbrace{P\left(v_t \mid v_{t-1}, rx_t, \Theta\right)}_{\text{(iii)}} \tag{B.4}$$

This equation explicitly identifies three sources of information through which v_t is updated. Part (i) of equation (B.4) signifies that observing the contemporaneous level of the VIX provides direct information about v_t. The VIX and the daily conditional variance share a tight relationship through equation (3.5.4), which implies:

$$P\left(\text{VIX}_t^2 \mid v_t, \Theta\right) = \frac{1}{\sqrt{2\pi s^2}} \exp\left\{-\frac{\left(\text{VIX}_t^2 - k_0 - k_1 v_t\right)^2}{2s^2}\right\}.$$

Part (ii) captures the influence of the observation from the next period (v_{t+1}, rx_{t+1}). Observing v_{t+1} indirectly affects the full conditional posterior of v_t because transitioning from v_t to v_{t+1} is governed by the square-root dynamics in equation (3.3.2). The excess return in the next period r_{t+1} also contains information about v_t due to two reasons: not only is v_t the conditional variance of r_{t+1}, but variance shocks and return shocks are correlated at ρ. It follows from our model setup that:

$$P\left(v_{t+1}, rx_{t+1} \mid v_t, \Theta\right) = \frac{1}{2\pi\sigma v_t\sqrt{1-\rho^2}} \exp\left(-\frac{\hat{\epsilon}_{t+1}^2 - 2\rho\hat{\epsilon}_{t+1}\hat{u}_{t+1} + \hat{u}_{t+1}^2}{2\left(1-\rho^2\right)}\right).$$

Lastly, part (iii) represents the likelihood of observing v_t given v_{t-1} and rx_t. By Bayes' rule, this probability equals $P\left(v_t, rx_t \mid v_{t-1}, \Theta\right)/P\left(rx_t \mid v_{t-1}, \Theta\right)$. Since $P\left(rx_t \mid v_{t-1}, \Theta\right)$ is constant with respect to v_t, it is sufficient to calculate $P\left(v_t, rx_t \mid v_{t-1}, \Theta\right)$ instead:

$$\begin{aligned} P\left(v_t \mid v_{t-1}, rx_t, \Theta\right) &\propto P\left(v_t, rx_t \mid v_{t-1}, \Theta\right) \\ &= \frac{1}{2\pi\sigma v_{t-1}\sqrt{1-\rho^2}} \exp\left(-\frac{\hat{\epsilon}_t^2 - 2\rho\hat{\epsilon}_t\hat{u}_t + \hat{u}_t^2}{2\left(1-\rho^2\right)}\right). \end{aligned}$$

Note that in the first half of our sample, the VIX is unavailable. In such days, part (i) is omitted from equation (B.4). Similarly, on the last day of the sample period (i.e. $t = T$), part (ii) disappears. Unlike these two cases, part (iii) does not drop when the first day of the sample period (i.e. $t = 0$) is considered. Instead, it simply reduces to $P(v_0)$, the stationary distribution of the conditional variance. We approximate this distribution by a gamma distribution: when time intervals get smaller, our variance process converges to a Cox, Ingersoll, and Ross (1985) process whose steady-state distribution is gamma.

Figure 3.1: Histograms of Daily Excess Returns

Panel A

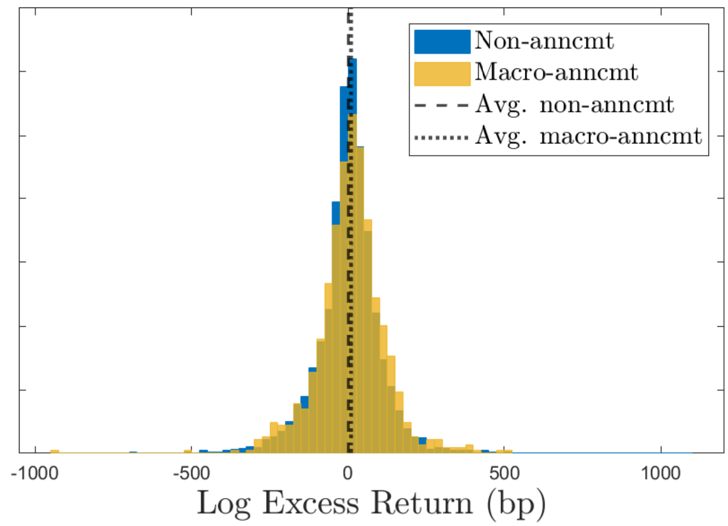

Panel B

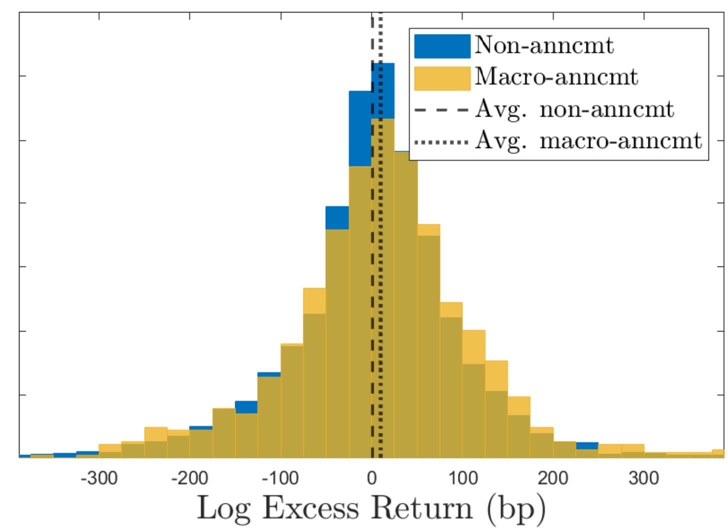

Note: This figure presents the histograms of daily excess returns over the macro-announcement sample (blue bars) vs. the non-announcement sample (yellow bars). The dotted lines and the dashed lines represent the averages over the macro-announcement sample and the non-announcement sample, respectively. Panel A displays the full distribution while Panel B focuses on the middle 99% (from 0.5 to 99.5 percentiles) in order to facilitate the visual comparison between the two sub-samples. The sample period is from 1990 to 2018.

Figure 3.2: Daily Excess Returns and Changes in VIX^2

Panel A

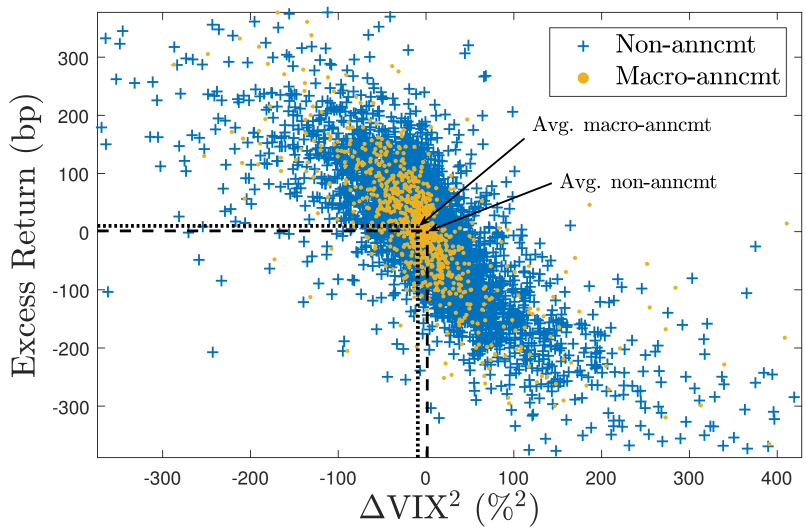

Panel B

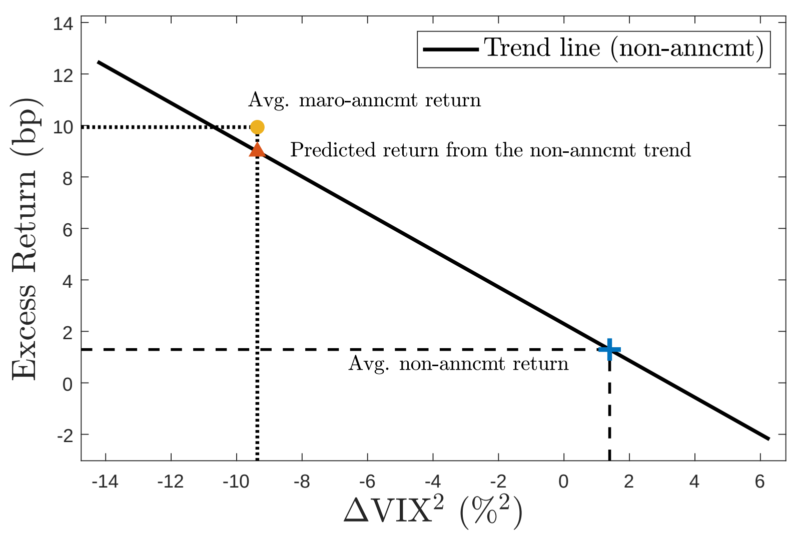

Note: This figure displays the relation between excess returns and changes in the squared VIX in the sample period of 1990-2018. In the scatter plot shown in Panel A, the blue cross markers represent the macro-announcement sample and the yellow circle markers represent the non-announcement sample. Panel B zooms into the center of Panel A to facilitate comparison between the averages of the two sub-samples. The averages of the macro-announcement and non-announcement samples are indicated by the yellow circle and the blue cross, respectively. The trend line for the non-announcement sample is added to Panel B as a black solid line.

Figure 3.3: Macro-Announcement Samples vs. Simulated Non-Announcement Samples

Panel A

Panel B

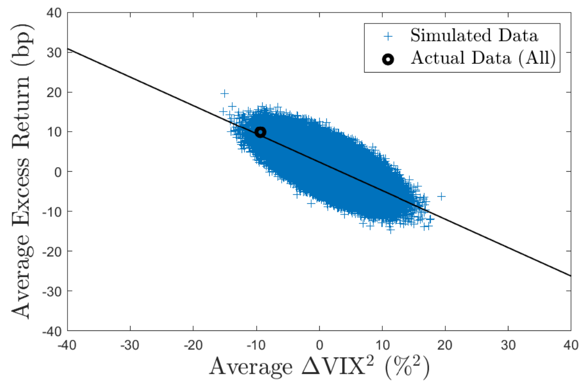

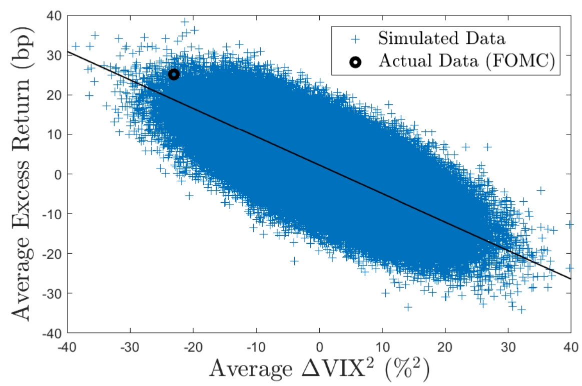

Panel C

Panel D

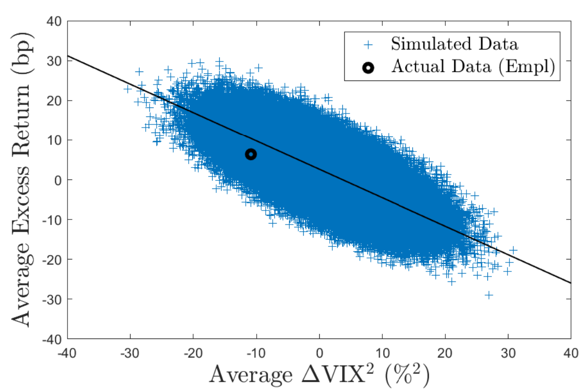

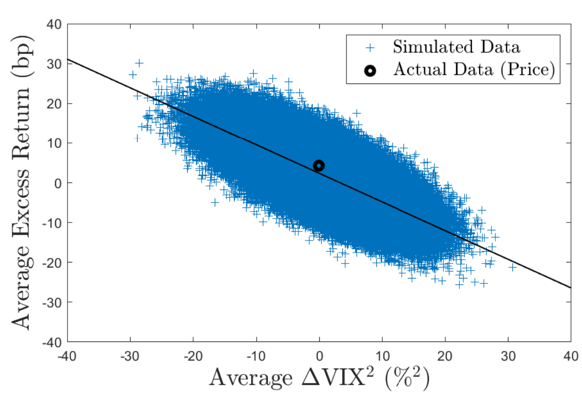

Note: This figure displays the relation between average excess returns and average changes

Figure 3.4: A Model That Does Not Feature an Announcement Premium

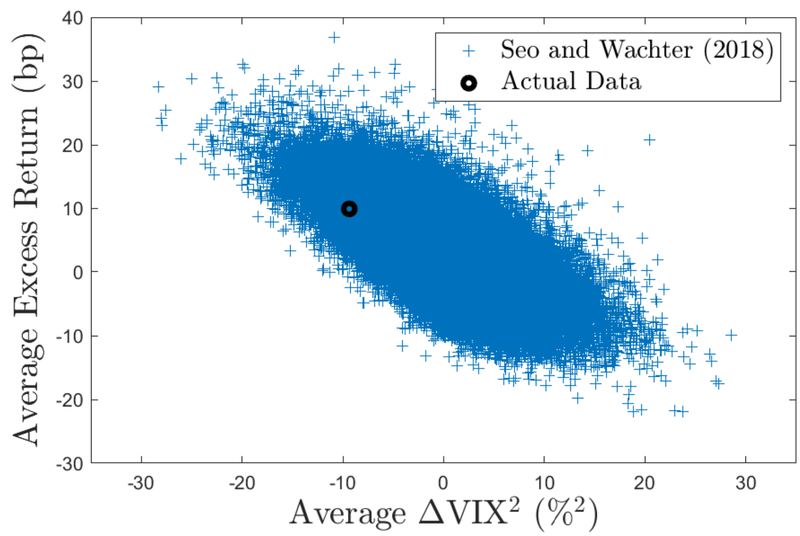

Note: This figure displays the relation between average excess returns and average changes in the squared VIX in a model that does not feature an announcement premium. Specifically, we simulate the model of Seo and Wachter (2018b) and generate one million 29-year paths at the daily frequency. Then, we randomly select 915 days from each no-disaster path and calculate the average excess return and the average change in the squared VIX (each blue cross). The black hollow circle indicates the actual data point from the macro-announcement sample.

Table 3.1: Main Estimation Results

Panel A: Parameter estimates

ρ	μ (bp)	γ (bp)	$\bar{v}$ (%^{2})	ϕ	σ (%)	k_0 (%^{2})	k_v
-0.711	2.953	-1.744	1.043	0.985	0.192	0.301	1.198
[0.011]	[0.920]	[2.204]	[0.163]	[0.004]	[0.006]	[0.007]	[0.016]

Panel B: Excess return vs. risk premium

	Excess return		Risk premium	
	Mean (bp)	Share (%)	Mean (bp)	Share (%)
Announcement days	9.93	52.35	1.21	5.54
Non-announcement days	1.29	47.65	2.95	94.46
All days (1990-2018)	2.38	100.00	2.73	100.00

Notes: This table shows the estimation results for our benchmark model with $\mathbb{M} = \mathbb{M}_{\text{All}}$ based on the sample period of 1990-2018. Panel A reports parameter estimates together with their robust standard errors in square brackets. Panel B compares the resulting average equity premiums with the average excess returns in the data over macro-announcement days, non-announcement days, and all days. The average equity premiums/excess returns are expressed in basis points. Their shares of the full equity premium/excess return, accumulated over the period of 1990-2018, are reported in percentage points.

Table 3.2: Average Variance and Return Innovations

	Variance		Return	
	u_{t+1}	$\sigma\sqrt{v_t}u_{t+1}$ (%^{2})	ϵ_{t+1}	$\sqrt{v_t}\epsilon_{t+1}$ (bp)
Announcement days	-0.15	-0.03	0.09	8.72
Non-announcement days	0.04	0.01	-0.02	-1.66
All days (1990-2018)	0.01	0.00	-0.01	-0.35

Notes: This table reports the averages of variance shocks (u_{t+1}), variance innovations ($\sigma\sqrt{v_t}u_{t+1}$), return shocks (ϵ_{t+1}), and return innovations ($\sqrt{v_t}\epsilon_{t+1}$) over macro-announcement days, non-announcement days, and all days in the sample period of 1990-2018.

Table 3.3: Different Macro-Announcement Samples: Model Parameters

Panel A: Parameter estimates

(1) FOMC		(2) Employment		(3) Price index	
μ (bp)	γ (bp)	μ (bp)	γ (bp)	μ (bp)	γ (bp)
2.814	-2.252	2.842	-2.095	2.845	-2.156
[0.902]	[3.761]	[0.893]	[3.996]	[0.912]	[3.175]

Panel B: Excess return vs. risk premium

	Excess return		Risk premium	
	Mean (bp)	Share (%)	Mean (bp)	Share (%)
(1) FOMC				
Announcement days	25.13	33.58	0.56	0.65
Non-announcement days	1.63	66.42	2.81	99.35
All days (1990-2018)	2.38	100.00	2.74	100.00
(2) Employment				
Announcement days	6.38	12.80	0.75	1.30
Non-announcement days	2.18	87.20	2.84	98.70
All days (1990-2018)	2.38	100.00	2.74	100.00
(3) Price index				
Announcement days	4.29	8.60	0.69	1.20
Non-announcement days	2.28	91.40	2.84	98.80
All days (1990-2018)	2.38	100.00	2.74	100.00

Notes: This table shows the estimation results for our benchmark model with $\mathbb{M} = \mathbb{M}_{\text{FOMC}}$ (FOMC decisions), $\mathbb{M}_{\text{Empl}}$ (employment figures), and $\mathbb{M}_{\text{Price}}$ (price indices) based on the sample period of 1990-2018. Panel A reports the parameter estimates for μ and γ together with their robust standard errors in square brackets. Panel B compares the resulting average equity premiums with the average excess returns in the data over macro-announcement days, non-announcement days, and all days. The average equity premiums/excess returns are expressed in basis points. Their shares of the full equity premium/excess return, accumulated over the period of 1990-2018, are reported in percentage points.

Table 3.4: The Role of Asymmetric Volatility

	(1) $\rho = 0$		(2) $\rho = -0.25$	
	Mean (bp)	Share (%)	Mean (bp)	Share (%)
Announcement days	7.75	40.82	5.48	28.89
Non-announcement days	1.61	59.18	1.93	71.11
All days (1990-2018)	2.38	100.00	2.38	100.00
	(3) $\rho = -0.50$		(4) $\rho = -0.75$	
	Mean (bp)	Share (%)	Mean (bp)	Share (%)
Announcement days	3.05	16.06	0.43	2.27
Non-announcement days	2.28	83.94	2.66	97.73
All days (1990-2018)	2.38	100.00	2.38	100.00

Notes: This table calculates the average equity premiums over macro-announcement days, non-announcement days, and all days by fixing the value of ρ at 0, -0.25, -0.50, and -0.75 in the estimation. The sample period is 1990-2018. The average equity premiums are expressed in basis points. Their shares of the full equity premium, accumulated over the period of 1990-2018, are reported in percentage points.

Table 3.5: Different Risk Premium Specifications

Panel A: Parameter estimates

(1) $\mu_t = \mu + \eta v_t + \gamma I_t$			(2) $\mu_t = \eta v_t + \gamma I_t$	
μ (bp)	γ (bp)	η	γ (bp)	η
2.941	-1.743	0.028	-0.815	2.386
[1.251]	[2.205]	[1.296]	[2.189]	[0.945]

Panel B: Risk premium

	(1) $\mu_t = \mu + \eta v_t + \gamma I_t$		(2) $\mu_t = \eta v_t + \gamma I_t$	
	Mean (bp)	Share (%)	Mean (bp)	Share (%)
Announcement days	1.23	5.60	2.08	9.59
Non-announcement days	2.97	94.40	2.81	90.41
All days (1990-2018)	2.75	100.00	2.72	100.00

Notes: This table shows the estimation results when we adopt the following two alternative equity premium specifications: (1) $\mu_t = \mu + \eta v_t + \gamma I_t$ and (2) $\mu_t = \eta v_t + \gamma I_t$. Panel A reports the estimates for the parameters that are associated with the conditional equity premium μ_t, together with their robust standard errors in square brackets. Panel B calculates the average equity premiums over macro-announcement days, non-announcement days, and all days under the two specifications. The average equity premiums are expressed in basis points. Their shares of the full equity premium, accumulated over the period of 1990-2018, are reported in percentage points.

Table 3.6: Expected Rise and Fall in Variance Around Announcements

(1) All				(2) FOMC			
μ (bp)	γ (bp)	$\bar{v}$ (%^{2})	ξ (%^{2})	μ (bp)	γ (bp)	$\bar{v}$ (%^{2})	ξ (%^{2})
2.417	2.666	1.041	0.013	2.583	5.093	1.042	0.023
[0.971]	[2.960]	[0.108]	[0.005]	[0.909]	[5.903]	[0.134]	[0.010]
(3) Employment				(4) Price Index			
μ (bp)	γ (bp)	$\bar{v}$ (%^{2})	ξ (%^{2})	μ (bp)	γ (bp)	$\bar{v}$ (%^{2})	ξ (%^{2})
2.632	2.371	1.042	0.014	2.795	-1.094	1.042	0.003
[0.901]	[5.446]	[0.136]	[0.008]	[0.919]	[4.139]	[0.156]	[0.008]

Notes: This table shows the estimation results when we adopt a different specification of v_t in equation (3.4.1), which accommodates an expected rise and a subsequent fall in the conditional variance around macro-announcements. The estimation is conducted with four definitions of the macro-announcement sample: $\mathbb{M} = \mathbb{M}_{\text{All}}$ (All announcements), $\mathbb{M}_{\text{FOMC}}$ (FOMC decisions), $\mathbb{M}_{\text{Empl}}$ (employment figures), and $\mathbb{M}_{\text{Price}}$ (price indices). The sample period is 1990-2018.

Table 3.7: Excess Returns vs. Changes in the Squared VIX

Panel A: Univariate regression			
	const.	ΔVIX^2	R^2
Announcement days	2.904	-0.751	0.529
	[2.551]	[0.023]	
Non-announcement days	2.292	-0.715	0.522
	[0.939]	[0.009]	
All days (1990-2018)	2.411	-0.719	0.523
	[0.881]	[0.008]	

Panel B: Fixed effects regression					
	const.	ΔVIX^2	I	$\Delta\text{VIX}^2 \times I$	R^2
All days (1990-2018)	2.292	-0.715	0.612	-0.036	0.523
	[0.942]	[0.009]	[2.671]	[0.025]	

Notes: This table examines the empirical relation between excess returns and changes in the squared VIX. Panel A presents the results of regressing the excess return (i.e. rx_{t+1}) on the contemporaneous change in the squared VIX (i.e. ΔVIX^2_{t+1}) in the macro-announcement sample, in the non-announcement sample, and in the full sample:

$$rx_{t+1} = \text{const.} + \beta\Delta\text{VIX}^2_{t+1} + e_{t+1}.$$

Panel B runs the following fixed effects regression:

$$rx_{t+1} = \text{const.} + \beta_1\Delta\text{VIX}^2_{t+1} + \beta_2 I_t + \beta_3\Delta\text{VIX}^2_{t+1} \times I_t + e_{t+1},$$

where I_t is a dummy variable that takes a value of 1 if time $t+1$ is a pre-scheduled macro-announcement day. The standard errors are reported in square brackets. The sample period is from 1990 to 2018.

Table 3.8: MCMC Estimation: 1990-2018

Panel A: Parameter estimates

	ρ	μ (bp)	γ (bp)	$\bar{v}$ ($\%^2$)	ϕ	σ (%)	k_0 ($\%^2$)	k_v
Mean	-0.685	2.480	-0.968	1.173	0.987	0.198	0.316	1.248
5%	-0.695	1.166	-3.805	0.927	0.983	0.195	0.307	1.214
50%	-0.685	2.481	-0.969	1.154	0.987	0.198	0.316	1.252
95%	-0.675	3.795	1.864	1.480	0.990	0.202	0.323	1.268

Panel B: Excess return vs. risk premium

	Excess return		Risk premium	
	Mean (bp)	Share (%)	Mean (bp)	Share (%)
Announcement days	9.93	52.35	1.51	8.03
Non-announcement days	1.29	47.65	2.48	91.97
All days (1990-2018)	2.38	100.00	2.36	100.00

Notes: This table shows the estimation results for our model via MCMC based on the sample period of 1990-2018. Panel A reports the posterior mean as well as 5%, 50%, and 95% percentile values for each model parameter. Panel B compares the resulting average equity premiums with the average excess returns in the data over macro-announcement days, non-announcement days, and all days. The average equity premiums/excess returns are expressed in basis points. Their shares of the full equity premium/excess return, accumulated over the period of 1990-2018, are reported in percentage points.

Table 3.9: MCMC Estimation: 1961-2018

Panel A: Parameter estimates								
	ρ	μ (bp)	γ (bp)	$\bar{v}$ (%^{2})	ϕ	σ (%)	k_0 (%^{2})	k_v
Mean	-0.662	2.035	1.537	0.878	0.978	0.186	0.308	1.319
5%	-0.672	1.147	-0.609	0.789	0.975	0.183	0.300	1.303
50%	-0.662	2.034	1.544	0.875	0.978	0.186	0.308	1.322
95%	-0.652	2.933	3.667	0.976	0.980	0.189	0.315	1.331

Panel B: Excess return vs. risk premium

	Excess return		Risk premium	
	Mean (bp)	Share (%)	Mean (bp)	Share (%)
Announcement days	8.64	58.05	3.57	20.83
Non-announcement days	0.94	41.95	2.03	79.17
All days (1961-2018)	1.94	100.00	2.24	100.00

Notes: This table shows the estimation results for our model via MCMC based on the sample period of 1961-2018. Panel A reports the posterior mean as well as 5%, 50%, and 95% percentile values for each model parameter. Panel B compares the resulting average equity premiums with the average excess returns in the data over macro-announcement days, non-announcement days, and all days. The average equity premiums/excess returns are expressed in basis points. Their shares of the full equity premium/excess return, accumulated over the period of 1961-2018, are reported in percentage points.

Bibliography

Abel, Andrew, 1990, Asset prices under habit formation and catching up with the Joneses, *American Economic Review Papers and Proceedings* 80, 38–42.

Adelino, Manuel, Antoinette Schoar, and Felipe Severino, 2016, Loan originations and defaults in the mortgage crisis: The role of the middle class, *The Review of Financial Studies* 29, 1635–1670.

Adelino, Manuel, Antoinette Schoar, and Felipe Severino, 2018, Dynamics of housing debt in the recent boom and great recession, *NBER Macroeconomics Annual* 32, 265–311.

Ai, Hengjie, 2010, Information quality and long-run risk: asset pricing implications, *The Journal of Finance* 65, 1333–1367.

Ai, Hengjie, and Ravi Bansal, 2018, Risk preferences and the macroeconomic announcement premium, *Econometrica* 86, 1383–1430.

Albanesi, Stefania, Giacomo De Giorgi, and Jaromir Nosal, 2017, Credit growth and the financial crisis: A new narrative, Working paper, National Bureau of Economic Research.

Andersen, Torben G, Tim Bollerslev, Francis X Diebold, and Clara Vega, 2007, Real-time price discovery in global stock, bond and foreign exchange markets, *Journal of international Economics* 73, 251–277.

Avdis, Efstathios, and Jessica A Wachter, 2017, Maximum likelihood estimation of the equity premium, *Journal of Financial Economics* 125, 589–609.

Backus, David, Mikhail Chernov, and Stanley Zin, 2014, Sources of entropy in representative agent models, *The Journal of Finance* 69, 51–99.

Balduzzi, Pierluigi, and Fabio Moneta, 2017, Economic risk premia in the fixed-income markets: The intraday evidence, *Journal of Financial and Quantitative Analysis* 52, 1927–1950.

Bandi, Federico M., and Roberto Reno, 2012, Time-varying leverage effects, *Journal of Econometrics* 169, 94–113.

Bansal, Ravi, Dana Kiku, and Amir Yaron, 2012, An Empirical Evaluation of the Long-Run Risks Model for Asset Prices, *Critical Finance Review* 1, 183–221.

Bansal, Ravi, and Amir Yaron, 2004, Risks for the long-run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.

Baron, Matthew, and Wei Xiong, 2017, Credit expansion and neglected crash risk, *The Quarterly Journal of Economics* 132, 713–764.

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

Barro, Robert J., and Tao Jin, 2011, On the size distribution of macroeconomic disasters, *Econometrica* 79, 1567–1589.

Barro, Robert J., and José F. Ursúa, 2008, Macroeconomic crises since 1870, *Brookings Papers on Economic Activity* no. 1, 255–350.

Barro, Robert J., and José F. Ursúa, 2017, Stock-market crashes and depressions, *Research in Economics* 71, 384–398.

Bekaert, Geert, and Guojun Wu, 2000, Asymmetric volatility and risk in equity markets, *Review of Financial Studies* 13, 1–42.

Benzoni, Luca, Pierre Collin-Dufresne, and Robert S. Goldstein, 2011, Explaining asset pricing puzzles associated with the 1987 market crash, *Journal of Financial Economics* 101, 552 – 573.

Beraja, Martin, Andreas Fuster, Erik Hurst, and Joseph Vavra, 2019, Regional heterogeneity and the refinancing channel of monetary policy, *The Quarterly Journal of Economics* 134, 109–183.

Berkman, Henk, Ben Jacobsen, and John B. Lee, 2011, Time-varying rare disaster risk and stock returns, *Journal of Financial Economics* 101, 313–332.

Bernanke, Ben S, Mark Gertler, and Simon Gilchrist, 1999, *The financial accelerator in a quantitative business cycle framework*Elsevier vol. 1 pp. 1341–1393.

Bernanke, Ben S, and Kenneth N Kuttner, 2005, What explains the stock market's reaction to Federal Reserve policy?, *The Journal of finance* 60, 1221–1257.

Black, Fischer, 1976, Studies of stock price volatility changes, *Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economic Statistics Section* pp. 177–181.

Boguth, Oliver, Vincent Grégoire, and Charles Martineau, 2019, Shaping expectations and coordinating attention: The unintended consequences of FOMC press conferences, *Journal of Financial and Quantitative Analysis* 54, 2327–2353.

Bollerslev, Tim, Julia Litvinova, and George Tauchen, 2006, Leverage and volatility feedback effects in high-frequency data, *Journal of Financial Econometrics* 4, 353–384.

Bollerslev, Tim, Natalia Sizova, and George Tauchen, 2012, Volatility in equilibrium: Asymmetries and dynamic dependencies, *Review of Finance* 16, 31–80.

Bollerslev, Tim, George Tauchen, and Hao Zhou, 2009, Expected stock returns and variance risk premia, *Review of Financial Studies* 22, 4463–4492.

Bollerslev, Tim, and Hao Zhou, 2006, Volatility puzzles: a simple framework for gauging return-volatility regressions, *Journal of Econometrics* 131, 123–150.

Brandt, Michael W., and Qiang Kang, 2004, On the relationship between the conditional mean and volatility of stock returns: A latent VAR approach, *Journal of Financial Economics* 72, 217–257.

Brandt, Michael W, Qi Zeng, and Lu Zhang, 2004, Equilibrium stock return dynamics under alternative rules of learning about hidden states, *Journal of Economic Dynamics and Control* 28, 1925–1954.

Broadie, Mark, Mikhail Chernov, and Michael Johannes, 2007, Model specification and risk premia: Evidence from futures options, *Journal of Finance* 62, 1453–1490.

Broadie, Mark, Mikhail Chernov, and Michael Johannes, 2009, Understanding index option returns, *The Review of Financial Studies* 22, 4493–4529.

Brusa, Francesca, Pavel Savor, and Mungo Wilson, 2020, One central bank to rule them all, *Review of Finance* 24, 263–304.

Campbell, John Y., and Ludger Hentschel, 1992, No news is good news: An asymmetric model of changing volatility in stock returns, *Journal of Financial Economics* 31, 281–318.

Chabi-Yo, Fousseni, and Johnathan Loudis, 2020, The conditional expected market return, *Journal of Financial Economics*.

Chen, Hui, Michael Michaux, and Nikolai Roussanov, 2020, Houses as ATMs: mortgage refinancing and macroeconomic uncertainty, *The Journal of Finance* 75, 323–375.

Cheung, Yin Wong, and Lilian K. Ng, 1992, Stock price dynamics and firm size: An empirical investigation, *Journal of Finance* 47, 1985–1997.

Christie, Andrew A., 1982, The stochastic behavior of common stock variances: Value, leverage and interest rate effects, *Journal of Financial Economics* 10, 407–432.

Christoffersen, Peter, Steve Heston, and Kris Jacobs, 2006, Option valuation with conditional skewness, *Journal of Econometrics* 131, 253–284.

Christoffersen, Peter, and Kris Jacobs, 2004, Which GARCH model for option valuation?, *Management science* 50, 1204–1221.

Cieslak, Anna, 2018, Short-rate expectations and unexpected returns in treasury bonds, *The Review of Financial Studies* 31, 3265–3306.

Cieslak, Anna, Adair Morse, and Annette Vissing-Jorgensen, 2018, Stock returns over the FOMC cycle, *The Journal of Finance*.

Cochrane, John H., 2017, Macro-finance, *Review of Finance* 21, 945–985.

Collin-Dufresne, Pierre, Michael Johannes, and Lars A Lochstoer, 2016, Parameter learning in general equilibrium: The asset pricing implications, *American Economic Review* 106, 664–98.

Constantinides, George M., 2002, Rational Asset Prices, *The Journal of Finance* 57, 1567–1591.

Constantinides, George M., 2008, Discussion of "Macroeconomic crises since 1870", *Brookings Papers on Economic Activity* no. 1, 336–350.

Cox, John C., Jonathan C. Ingersoll, and Stephen A. Ross, 1985, A theory of the term structure of interest rates, *Econometrica* 53, 385–408.

David, Alexander, and Pietro Veronesi, 2014, Investors' and central bank's uncertainty embedded in index options, *The Review of Financial Studies* 27, 1661–1716.

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

Di Maggio, Marco, and Amir Kermani, 2017, Credit-induced boom and bust, *The Review of Financial Studies* 30, 3711–3758.

Donaldson, R. Glen, Mark J. Kamstra, and Lisa A. Kramer, 2010, Estimating the equity premium, *Journal of Financial and Quantitative Analysis* 45, 813–846.

Drechsler, Itamar, and Amir Yaron, 2011, What's vol got to do with it, *Review of Financial Studies* 24, 1–45.

Duffee, Gregory R., 1995, Stock returns and volatility a firm-level analysis, *Journal of Financial Economics* 37, 399–420.

Duffie, Darrell, Jun Pan, and Kenneth Singleton, 2000, Transform analysis and asset pricing for affine jump-diffusions, *Econometrica* 68, 1343–1376.

Eggertsson, Gauti B, and Paul Krugman, 2012, Debt, deleveraging, and the liquidity trap: A Fisher-Minsky-Koo approach, *The Quarterly Journal of Economics* 127, 1469–1513.

Epstein, Larry, and Stan Zin, 1989, Substitution, risk aversion and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica* 57, 937–969.

Eraker, Bjørn, 2004, Do stock prices and volatility jump? Reconciling evidence from spot and option prices, *Journal of Finance* 59, 1367–1404.

Eraker, Bjørn, Michael Johannes, and Nicholas Polson, 2003, The impact of jumps in volatility and returns, *Journal of Finance* 58, 1269–1300.

Ernst, Rory, Thomas Gilbert, and Christopher M Hrdlicka, 2021, More than 100% of the equity premium: How much is really earned on macroeconomic announcement days?, *Available at SSRN 3469703*.

Fama, Eugene F., and Kenneth R. French, 2002, The Equity Premium, *The Journal of Finance* 57, pp. 637–659.

Farhi, Emmanuel, and Xavier Gabaix, 2016, Rare disasters and exchange rates, *The Quarterly Journal of Economics* 131, 1–52.

Foote, Christopher L, Kristopher S Gerardi, and Paul S Willen, 2012, Why did so many people make so many ex post bad decisions? The causes of the foreclosure crisis, Working paper, National Bureau of Economic Research.

Foote, Christopher L, Lara Loewenstein, and Paul S Willen, 2016, Cross-sectional patterns of mortgage debt during the housing boom: evidence and implications, Working paper, National Bureau of Economic Research.

French, Kenneth R., G. William Schwert, and Robert F. Stambaugh, 1987, Expected stock returns and volatility, *Journal of Financial Economics* 19, 3–29.

Gabaix, Xavier, 2012, An exactly solved framework for ten puzzles in macro-finance, *Quarterly Journal of Economics* 127, 645–700.

Gebhardt, William R, Charles MC Lee, and Bhaskaran Swaminathan, 2001, Toward an implied cost of capital, *Journal of accounting research* 39, 135–176.

Ghent, Andra C, and Michael T Owyang, 2010, Is housing the business cycle? Evidence from US cities, *Journal of Urban Economics* 67, 336–351.

Ghysels, Eric, Pedro Santa-Clara, and Rossen Valkanov, 2005, There is a risk-return trade-off after all, *Journal of Financial Economics* 76, 509–548.

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

Gillman, Max, Michal Kejak, and Michal Pakoš, 2014, Learning about rare disasters: implications for consumption and asset prices, *Review of Finance* 19, 1053–1104.

Glaeser, Edward L, and Charles G Nathanson, 2017, An extrapolative model of house price dynamics, *Journal of Financial Economics* 126, 147–170.

Glosten, L.R., R. Jagannathan, and D. Runkle, 1993, On the relation between the expected value and the volatility of the nominal excess return on stocks, *Journal of Finance* 48, 1779–1801.

Gomes, João F, Marco Grotteria, and Jessica A Wachter, 2019, Cyclical dispersion in expected defaults, *The Review of Financial Studies* 32, 1275–1308.

Gorbachev, Olga, 2011, Did household consumption become more volatile?, *American Economic Review* 101, 2248–70.

Gourio, François, 2012, Disaster risk and business cycles, *American Economic Review* 102, 2734–2766.

Greenwood, Robin, and Samuel G Hanson, 2013, Issuer quality and corporate bond returns, *The Review of Financial Studies* 26, 1483–1525.

Guerrieri, Veronica, and Guido Lorenzoni, 2017, Credit crises, precautionary savings, and the liquidity trap, *The Quarterly Journal of Economics* 132, 1427–1467.

Guerrieri, Veronica, and Harald Uhlig, 2016, *Housing and credit markets: booms and busts*Elsevier vol. 2 pp. 1427–1496.

Guo, Hui, and Robert F. Whitelaw, 2006, Uncovering the risk-return relation in the stock market, *The Journal of Finance* 61, 1433–1463.

Harvey, Campbell R., 2001, The specification of conditional expectations, *Journal of Empirical Finance* 8, 573–637.

Hasler, Michael, and Roberto Marfe, 2016, Disaster recovery and the term structure of dividend strips, *Journal of Financial Economics* 122, 116–134.

Heston, Steven L., 1993, A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options, *Review of Financial Studies* 6, 327–343.

Heston, Steven L., and Saikat Nandi, 2000, A closed-form GARCH option valuation model, *The review of financial studies* 13, 585–625.

Hu, Grace Xing, Jun Pan, Jiang Wang, and Haoxiang Zhu, 2020, Premium for heightened uncertainty: explaining pre-announcement market teturns, Working paper.

Iacoviello, Matteo, 2005, House prices, borrowing constraints, and monetary policy in the business cycle, *American economic review* 95, 739–764.

Ibbotson, Roger G, and Peng Chen, 2003, Long-run stock returns: Participating in the real economy, *Financial Analysts Journal* 59, 88–98.

Jacquier, Eric, Nicholas G Polson, and Peter E Rossi, 2004, Bayesian analysis of stochastic volatility models with fat-tails and correlated errors, *Journal of Econometrics* 122, 185–212.

Jermann, Urban, and Vincenzo Quadrini, 2012, Macroeconomic effects of financial shocks, *American Economic Review* 102, 238–71.

Johannes, Michael, Lars A Lochstoer, and Yiqun Mou, 2016, Learning about consumption dynamics, *The Journal of Finance* 71, 551–600.

Johannes, Michael, and Nicholas Polson, 2010, MCMC methods for continuous-time financial econometrics, in *Handbook of Financial Econometrics: Applications* (Elsevier,).

Jordà, Òscar, Moritz Schularick, and Alan M Taylor, 2019, The total risk premium puzzle, Working paper, National Bureau of Economic Research.

Ju, Nengjiu, and Jianjun Miao, 2012, Ambiguity, learning, and asset returns, *Econometrica* 80, 559–591.

Jurek, Jakub W., 2014, Crash-neutral currency carry trades, *Journal of Financial Economics* 113, 325–347.

Justiniano, Alejandro, Giorgio E Primiceri, and Andrea Tambalotti, 2019, Credit supply and the housing boom, *Journal of Political Economy* 127, 1317–1350.

Kalnina, Ilze, and Dacheng Xiu, 2017, Nonparametric estimation of the leverage effect: A trade-off between robustness and efficiency, *Journal of the American Statistical Association* 112, 384–396.

Kaplan, Greg, Kurt Mitman, and Giovanni L Violante, 2017, The housing boom and bust: Model meets evidence, Working paper, National Bureau of Economic Research.

Kiyotaki, Nobuhiro, and John Moore, 1997, Credit cycles, *Journal of political economy* 105, 211–248.

Leamer, Edward E, 2007, Housing is the business cycle, Working paper, National Bureau of Economic Research.

Lettau, Martin, and Sydney C. Ludvigson, 2001, Consumption, aggregate wealth and expected stock returns, *Journal of Finance* 56, 815–849.

Lettau, Martin, Sydney C. Ludvigson, and Jessica A. Wachter, 2008, The declining equity premium: What role does macroeconomic risk play?, *Review of Financial Studies* 21, 1653–1687.

Lucca, David O, and Emanuel Moench, 2015, The pre-FOMC announcement drift, *The Journal of Finance* 70, 329–371.

Ludvigson, Sydney C., and Serena Ng, 2007, The empirical risk-return relation: A factor analysis approach, *Journal of Financial Economics* 83, 171–222.

Lundblad, Christian, 2007, The risk return tradeoff in the long run: 1836-2003, *Journal of Financial Economics* 85, 123–150.

Manela, Asaf, and Alan Moreira, 2017, News implied volatility and disaster concerns, *Journal of Financial Economics* 123, 137–162.

Martin, Ian, 2017, What is the expected return on the market?, *Quarterly Journal of Economics* 132, 367–433.

Mehra, Rajnish, and Edward Prescott, 1985, The equity premium puzzle, *Journal of Monetary Economics* 15, 145–161.

Mendoza, Enrique G, 2010, Sudden stops, financial crises, and leverage, *American Economic Review* 100, 1941–66.

Merton, Robert C., 1973, An intertemporal capital asset pricing model, *Econometrica* 41, 867–887.

Merton, Robert C., 1974, On the pricing of corporate debt: The risk structure of interest rates, *Journal of Finance* 29, 449–470.

Mian, Atif, Kamalesh Rao, and Amir Sufi, 2013, Household balance sheets, consumption, and the economic slump, *The Quarterly Journal of Economics* 128, 1687–1726.

Mian, Atif, and Amir Sufi, 2009, The consequences of mortgage credit expansion: Evidence from the US mortgage default crisis, *The Quarterly Journal of Economics* 124, 1449–1496.

Mian, Atif, and Amir Sufi, 2011, House prices, home equity-based borrowing, and the US household leverage crisis, *American Economic Review* 101, 2132–56.

Mian, Atif, and Amir Sufi, 2014, What explains the 2007–2009 drop in employment?, *Econometrica* 82, 2197–2223.

Mueller, Philippe, Alireza Tahbaz-Salehi, and Andrea Vedolin, 2017, Exchange rates and monetary policy uncertainty, *The Journal of Finance* 72, 1213–1252.

Muir, Tyler, 2017, Financial crises and risk premia, *The Quarterly Journal of Economics* 132, 765–809.

Nakamura, Emi, and Jón Steinsson, 2018, High-frequency identification of monetary non-neutrality: the information effect, *The Quarterly Journal of Economics* 133, 1283–1330.

Nakamura, Emi, Jón Steinsson, Robert Barro, and José Ursúa, 2013, Crises and Recoveries in an Empirical Model of Consumption Disasters, *American Economic Journal: Macroeconomics* 5, 35–74.

Pastor, Lubos, Meenakshi Sinha, and Bhaskaran Swaminathan, 2008, Estimating the intertemporal risk-return tradeoff using the implied cost of capital, *The Journal of Finance* 63, 2859–2897.

Pastor, Lubos, and Pietro Veronesi, 2009a, Learning in Financial Markets, *Annual Review of Financial Economics* 1, 361–381.

Pastor, Lubos, and Pietro Veronesi, 2009b, Technological Revolutions and Stock Prices, *The American Economic Review* 99, pp. 1451–1483.

Piazzesi, Monika, and Martin Schneider, 2016, *Housing and macroeconomics*Elsevier vol. 2 pp. 1547–1640.

Rietz, Thomas A., 1988, The equity risk premium: A solution, *Journal of Monetary Economics* 22, 117–131.

Savor, Pavel, and Mungo Wilson, 2013, How much do investors care about macroeconomic risk? Evidence from scheduled economic announcements, *Journal of Financial and Quantitative Analysis* 48, 343–375.

Savor, Pavel, and Mungo Wilson, 2014, Asset pricing: A tale of two days, *Journal of Financial Economics* 113, 171–201.

Schwert, William G., 1989, Why does stock market volatility change over time?, *Journal of Finance* 44, 1115–1153.

Seo, Sang Byung, and Jessica A Wachter, 2018a, Do rare events explain CDX tranche spreads?, *The Journal of Finance* 73, 2343–2383.

Seo, Sang Byung, and Jessica A. Wachter, 2018b, Do rare events explain CDX tranche spreads?, *Journal of Finance* 73, 2343–2383.

Seo, Sang Byung, and Jessica A Wachter, 2018c, Option prices in a model with stochastic disaster risk, *Forthcoming, Management Science*.

Seo, Sang Byung, and Jessica A. Wachter, 2019, Option prices in a model with stochastic disaster risk, *Management Science* 65, 3449–3469.

Shiller, Robert J, 2014, Speculative asset prices, *American Economic Review* 104, 1486–1517.

Timmermann, Allan, 1996, Excess volatility and predictability of stock prices in autoregressive dividend models with learning, *The Review of Economic Studies* 63, 523–557.

Veronesi, Pietro, 1999, Stock market overreactions to bad news in good times: a rational expectations equilibrium model, *Review of Financial Studies* 12, 975–1007.

Veronesi, Pietro, 2000, How does information quality affect stock returns?, *The Journal of Finance* 55, 807–837.

Veronesi, Pietro, 2004, The Peso problem hypothesis and stock market returns, *Journal of Economic Dynamics and Control* 28, 707–725.

Wachter, Jessica A., 2013, Can time-varying risk of rare disasters explain aggregate stock market volatility?, *Journal of Finance* 68, 987–1035.

Wachter, Jessica A, and Yicheng Zhu, 2019, Learning with rare disasters, *Available at SSRN*.

Wachter, Jessica A, and Yicheng Zhu, 2020, A model of two days: discrete news and asset prices, Working paper.

Wang, Christina D., and Per A. Mykland, 2014, The estimation of leverage effect with high-frequency data, *Journal of the American Statistical Association* 109, 197–215.

Weil, Philippe, 1989, The equity premium puzzle and the risk-free rate puzzle, *Journal of Monetary Economics* 24, 402–421.

Weil, Philippe, 1990, Nonexpected utility in macroeconomics, *Quarterly Journal of Economics* 105, 29–42.

Welch, Ivo, 2016, The (time-varying) importance of disaster risk, *Financial Analysts Journal* 72, 14–30.

Zhou, Hao, 2018, Variance risk premia, asset predictability puzzles, and macroeconomic uncertainty, Forthcoming, *Annual Review of Financial Economics*.