EUROPEAN CENTRAL BANK

EUROSYSTEM

WORKING PAPER SERIES

NO 1728 / SEPTEMBER 2014

FINANCIAL STRESS AND ECONOMIC DYNAMICS

THE TRANSMISSION OF CRISES

Kirstin Hubrich and Robert J. Tetlow

In 2014 all ECB publications feature a motif taken from the €20 banknote.

Acknowledgements
We thank, without implication, Geert Bekaert, Francesco Bianchi, Jim Bullard, Marcelle Chauvet, Larry Christiano, Tim Cogley, Marco del Negro, Frank Diebold, Marty Eichenbaum, Jordi Gali, Simon Gilchrist, Anil Kashyap, Alejandro Justiniano, Nobu Kiyotaki, Sharon Kozicki, Thomas Laubach, Jim Nason, Chris Otrok, Hashem Pesaran, Giorgio Primiceri, Frank Schorfheide, Chris Sims, Frank Smets, Andrea Tambalotti, Harald Uhlig and Egon Zakrajšek for helpful comments. We are grateful to Dan Wagonner and Tao Zha for discussions and advice. We benefitted from the remarks of conference participants at the CEF 2012, SNDE 2011, NASM 2011, NBER Summer Institute 2011, EEA 2011, Dynare 2011, Central Bank Macro Modelling 2011, EC(2) 2012, a workshop in Rotterdam, 2012, SED 2013, as well as at seminars at the European Central Bank, Federal Reserve Board, FRB-Atlanta, FRB-New York, International Monetary Fund, UC-Irvine, UC-Riverside, University of Southern California, Maastricht University and the University of Pennsylvania. Peter Chen, Trevor Davis, Jens Kruk, Vesela Ivanova and Cristina Manea provided capable research assistance. Tetlow thanks the IMF for hosting him for his leave year during which part of this paper was written. The views expressed in this paper are solely the responsibility of the authors and should not be attributed to the European Central Bank or the Federal Reserve Board.

Kirstin Hubrich
European Central Bank; e-mail: kirstin.hubrich@ecb.europa.eu

Robert J. Tetlow
Federal Reserve Board; e-mail: rtetlow@frb.gov

Address	Kaiserstrasse 29, 60311 Frankfurt am Main, Germany
Postal address	Postfach 16 03 19, 60066 Frankfurt am Main, Germany
Telephone	+49 69 1344 0
Internet	http://www.ecb.europa.eu

ISSN	1725-2806 (online)
ISBN	978-92-899-1136-8
EU Catalogue No	QB-AR-14-102-EN-N (online)

Abstract

A financial stress index for the United States is introduced—an index that was used in real time by the staff of the Federal Reserve Board to monitor the financial crisis of 2008-9—and the interaction with real activity, inflation and monetary policy is demonstrated using a richly parameterized Markov-switching VAR model, estimated using Bayesian methods. A "stress event" is defined as a period where the latent Markov states for both shock variances and model coefficients are adverse. Results show that allowing for time variation is economically and statistically important, with solid (quasi) real-time properties. Stress events line up well with financial events in history. A shift to a stress event is highly detrimental to the outlook for the real economy, and conventional monetary policy is relatively weak during such periods.

Keywords: Nonlinearity, Markov switching, Financial crises, Monetary policy.

JEL Codes: E44, C11, C32

Non-technical Summary

The financial crisis and ensuing Great Recession in the United States revealed some important shortcomings in macroeconomic models in particular in explaining the amplification and feedback effects between the financial sector and the macroeconomy. In normal times, even finding a role for financial factors in affecting the real economy, once monetary and other factors have been accounted for, has been a challenge. In this paper, we argue that one reason why macroeconomically important linkages between the financial sector and the macroeconomy have been elusive is because the importance of financial factors has tended to be *episodic* in nature.

Our contention is that the U.S. economy has been sporadically, if not frequently, affected by what we call “stress events.” These events — manifestations of the amplification and propagation of financial shocks through the financial sector and the macroeconomy — lead us to examine the issue in a nonlinear, multivariate framework. In particular, we build on the work of Sims, Waggoner and Zha (2008) by employing a richly parameterized Markov switching vector autoregression (MS-VAR) model, an excercise made feasible by exploiting recent developments in Bayesian econometrics. Model assessment is done by a mixture of goodness-of-fit criteria and economic criteria - mostly marginal data densities, the properties of the model and the model’s ability to explain key events in recent financial and economic history.

We also introduce an index of financial stress—one that was used by the staff of the Federal Reserve Board to monitor and model financial developments in real time during the crisis—and assess its efficacy.

Our primary focus is on whether the economy behaves differently during periods of high stress, as the story sketched above suggests. Does the economy propagate shocks—transmit crises—differently during such periods? In other words, is there empirical evidence for nonlinearities in the linkages between the financial sector and macroeconomic dynamics? To answer these questions we investigate whether the coefficients and stochastic shocks of the VAR shift over time, and whether these shifts coincide with established events in U.S. economic and financial history. Allowing for *variance switching* is important not only to isolate the contribution of "bad luck" in explaining the volatility and performance of the economy during stressful times, but also to avoid biasing the outcomes toward the erroneous finding of *coefficient switching*.

We find substantial evidence of nonlinearities or non-Gaussian shock processes: The linkage between financial stress and the macroeconomy is not well described by a time-invariant Gaussian VAR benchmark model. It follows from this that inference drawn from time-invariant linear models may be misleading for some questions of interest. Second, variance switching alone is not sufficient to characterize departures from the benchmark model; unlike the business cycle characterization of Sims and Zha (2006), or the depiction of the drivers of the most recent recession described by Stock and Watson (2012), both of which explain the phenomena under study as arising from unusual sequences of shocks, we find that coefficient switching—and hence, nonlinear dynamics—is an important part of the mechanism linking financial stress and macroeconomic outcomes. Third, we find that the Fed staff appears to have been well served by its reliance on the financial stress index we study here since it appears to be a useful tool that can aid in capturing periods of financial stress in quasi-real time. More generally, the fact that switches in state appear to be reliably inferrable in quasi real time suggests that financial stability policies could be implemented in a timely fashion. Fourth, turning to the quantitative outcomes, we find that an important precursor to adverse economic events is a switch to what we call a *stress event*: a period in which the shock variance is at a relatively high-stress level, and the coefficient state is also at a high-stress level. It is often the case that stress events occur when shock volatility begins to rise and is followed by the change in coefficient state. We also find that stress is of quantitatively negligible importance in "normal" times, but of critical importance when the economy is in a high-stress coefficient state. Moreover, our results

suggest that conventional monetary policy —both systematic and in terms of policy shocks–is not particularly effective in times of high financial stress; a much more powerful tool is to induce a switch from a high-stress state back to "normal times," although how this could be achieved is outside the scope of the paper.

Taken together, we argue that these results have meaningful implications for the construction of dynamic stochastic general equilibrium (DSGE) models. In particular, while linearized DSGE models may be useful for thinking about normal business fluctuations, to the extent that one is interested in the sort of dynamics that underscored the 2008-9 financial crisis the usefulness of linearized DSGE models is limited. Rather, Markov Switching-DSGE models or fully articulated nonlinear models that are solved with global methods are more appropriate for that task. On the empirical side, it also follows that inference regarding the relationship between financial stress and the macroeconomy that is gleaned from a constant-parameter model may be inappropriate.

Lastly, we uncovered an interpretation emphasizing risky spreads as a key component of financial stress on the one hand, and durable goods as a real variable on the other, with implications for structural modelling.

1 Introduction

Financial factors have long been recognized as being important for understanding macroeconomic dynamics; see, e.g., Bernanke and Blinder (1988) and Kashyap et al. (1993). Yet the inclusion of financial frictions within dynamic stochastic general equilibrium (DSGE) models has been a notably recent phenomenon. One reason why modeling financial frictions was neglected is that it is empirically challenging. As the survey articles by Kashyap and Stein (1994) and Hubbard (1998) make clear, it has been remarkably difficult to uncover significant effects of financial frictions in macroeconomic time-series data. Indeed, with the noteworthy exceptions of Carlstrom and Fuerst (1997) and Bernanke et al. (1999), DSGE models with financial frictions have arisen after the experience of the recent financial crisis and subsequent recession.

In this paper, we argue that a reason why statistically significant and macroeconomically important linkages have been elusive is because the importance of financial factors tends to be *episodic* in nature. In "normal times," firms make investment decisions on the basis of whether a project's expected rate of return exceeds the user cost of capital, and then having made that decision, seek the financing. In such times, the financing decision is, in some sense, subordinate to the real-side decisions of the firm; credit "doesn't matter." In other times, when the financial system is not operating normally, financial frictions become important as lending terms and standards tighten, making the interest rate a much less reliable metric of the cost of funds, broadly defined. During such times, which we will call *stress events*; credit can seem like it is the only thing that matters.

Our contention that there are stress events that are episodic in nature, together with the associated interdependency of the financial sector and the macroeconomy, leads us to examine the issue in a nonlinear, multivariate framework. In particular, we build on the work of Sims, Waggoner and Zha (SWZ 2008) by employing a richly parameterized Markov switching vector autoregression (MS-VAR) model, estimated with Bayesian methods. Our primary focus is on whether the economy behaves differently during periods where the latent Markov state is one of high stress, as the story sketched above suggests. Does the economy propagate shocks differently—transmit crises—during such periods? Thus we investigate whether the VAR coefficients shift over time, and whether these shifts coincide with established events in U.S. economic and financial history.

Mindful of the possibility that financial stress could arise from outsized shocks, we also explicitly allow for switching in the variances of shocks—or *variance switching*, for short. Besides being an issue in its own right, allowing for variance switching is important to avoid biasing results toward the erroneous finding of *coefficient switching*. As in the literature on the origins of the Great Moderation, variance switching and coefficient switching are rivals in explaining the data.

A second contribution of this research, is the public introduction and assessment of a financial stress index, one that covers a broad range of financial market phenomena, that was formulated and used by the Federal Reserve Board staff during the crisis—on the fly, as it were—to analyze financial conditions and their macroeconomic consequences.

Ours is not the first paper in this area, broadly defined. Since the onset of the crisis, a second generation of DSGE models with financial frictions have sprung up, including Curdia and Woodford (2009), Jermann and Quadrini (2012) and Gilchrist et al. (2014). These and other papers have added insight to thinking about financial frictions as a source of shock amplification, but in most instances, their depiction of model economies allows for a single time-invariant steady state; no role for instability, volatility dynamics or important nonlinear effects is considered. There are also Markov switching DSGE models, including Liu et al. (2011) and F. Bianchi (2013). However these papers focus on business cycle phenomena, rather than financial stress.

The noteworthy empirical models in the area have included Lown and Morgan (2006), who examine the interaction of real variables and the responses to the Fed's Senior Loan Officers' Opinion Survey in a quarterly time-invariant VAR. An MS-VAR is arguably preferable to model the abrupt, discrete changes in economic dynamics of the recent crisis. Among the very few Markov switching models that pay attention to financial stress that we are aware of is Davig and Hakkio (2010) who, like us, employ an index of financial stress; however, their model is much simpler than ours and omits any consideration of monetary policy or price determination.

To presage the results, taking as a benchmark the standard, time-invariant Gaussian VAR model, we find substantial evidence of non-Gaussian shock processes and nonlinearities; the linkage between financial stress and the macroeconomy is not well described by the benchmark model. Second, variance switching alone is not sufficient to model the departures from the benchmark model; unlike the business cycle characterization of SZ (2006), or the depiction of the drivers of the most recent recession described by Stock and Watson (2012), both of which explain the phenomena under study as arising from unusual sequences of shocks, we find that coefficient switching—and hence, nonlinear dynamics—is important. Third, we find that the financial stress index we use (and that the Federal Reserve Board's staff used during the crisis) is indeed a useful tool that can aid in capturing periods of financial stress in real time. Fourth, our results suggest that conventional monetary policy is not particularly effective in times of high financial stress; a much more powerful mechanism is to induce a switch from a high-stress state back to "normal times."

While linearized DSGE models may be useful for thinking about the role of financial factors in business cycle fluctuations, we argue that if one is interested in the type of dynamics that underscored the 2008-9 financial crisis, linearized DSGE models will be of limited applicability. Rather, MS-DSGE models, such as F. Bianchi (2013), or fully articulated nonlinear models, solved with global methods, are better equipped for the job. Examples of the latter include Brunnermeier and Sannikov (2014), Mendoza (2010), He and Krishnamurthy (2012), J. Bianchi (2011) and Boissay et al. (2013). [1] On the empirical side, it also follows that inference regarding the relationship between financial stress and the macroeconomy that is gleaned from a constant-parameter model may be inappropriate.

The remainder of the paper proceeds as follows. In section 2, we discuss the history of financial stress in the United States. We also introduce our data and link historical events to the data. The third section discusses our modeling framework and econometric strategy while the fourth presents our results. The fifth section explores the economic interpretation of our results, in part through an analysis of robustness, while the sixth demonstrates the macroeconomic properties of the base case model. A seventh and final section sums up and concludes. An online appendix provides details on data and computation as well as more results.

2 Measuring financial stress

We begin with a bit of recent financial history for the U.S. before turning to a discussion of the Financial Stress Index.

2.1 Some history

To casual observers, financial stress might seem like a recent phenomenon, but it has been more prevalent than one might think. Students of banking history know that there were banking

[1] Taking Brunnermeier and Sannikov (2014) as an example, models of this class can allow for instabilities and periodic episodes of volatility, driven in part by ocassionally binding financial constraints. Such models emphasize the highly non-linear amplification effects caused by leverage and feedback effects from asset prices. Risk is sometimes endogenous in such models so that financial innovations can lead to better sharing of exogenous risk, but higher endogenous systemic risk as agents optimally respond to the safer environment they find themselves in. Externalities can lead to socially inappropriate levels of leverage, excess volatility and higher correlations of asset prices.

crises in the U.S. in 1837, 1857, 1873, 1907 and 1933. It is only recently that crises have become rare. Nevertheless, the rarity of full-blown crises does not mean that there has not been episodes of financial stress. Table 1 lays out some events over the last twenty years that have buffeted financial markets.

Table 1
Selected Financial Events Affecting the US Economy, 1986-2011_

	Event description	Date(s)
a	Savings & loan (S&L) crisis and its aftermath	1986-1992
b	Iraqi invasion of Kuwait	August 2, 1990
c	Mexican peso crisis	Dec. 1994-1995
d	Asia crisis	July 1997-1999
e	Collapse of Long-Term Capital Management (LTCM)	May-Sept. 1998
f	Russian debt default	Aug. 1998
g	Technology bubble bursts (NASDAQ descent)	Mar '00-Apr '01
h	Enron scandal and bankruptcy	Oct.-Nov. 2001
i	Argentine financial crisis	Dec. 2001-2002
j	Bear Stearns halts redemptions from two of its funds	July 17, 2007
k	Run on the repo market starts, according to Gorton (2010)	Aug. 9, 2007
l	Fed announces Term Auction Facility (TAF)	Dec. 12, 2007
m	TSLF and PDCF initiated; Bear Stearns sold	March 2008
n	AIG announces imminent bankrupty, gets bailed out	Sept. 16, 2008
o	Lehmann Brothers declares bankruptcy	Sept. 14, 2008
p	Congress passes Troubled Asset Relief Program (TARP)	Oct. 3, 2008
q	Term Asset-backed Securities Facility (TALF) announced	Nov. 25, 2008
r	Treasury department announces stress tests	Feb. 10, 2009
s	US bank stress test results released	May 7, 2009
t	Greece admits deficit-to-GDP ratio of 12 percent	Oct 18, 2009
u	First Eurozone-IMF rescue plan completed	May 2, 2010
v	Eurpoean FSB cleared to purchase sovereign bonds	July 2011
w	ECB offers massive loans to distressed banks	Dec. 21, 2012

As the table notes, there were financial incidents long before troubles at hedge funds owned by Bear Stearns showed up in the spring of 2007.

2.2 A Financial Stress Index

As the financial crisis began to take hold in 2007, as a complement to existing models, and to capture the higher frequency dynamics that no quarterly model could absorb in real time, a Financial Stress Index (FSI) for the United States was constructed. One contribution of this paper will be our assessment of the efficacy of the FSI as a useful real-time tool for the Board's staff during this critical period.[2]

[2] The FSI discussed in this section is based on an index described in Nelson and Perli (2005), modified to allow a longer historical series. Note that our goal is not to construct the best, ex post, measure of financial stress; an index that is data mined to "explain" historical financial events would likely turn out to be fragile.

The index focusses on capital market measures of stress, as opposed to banking measures. There are costs and benefits associated with this focus. As we noted in the introduction, financial stress manifests itself through both price and non-price channels, and in both capital markets and in banking. A common source of data for (something like) stress in banking is the Senior Loan Officer Opinion Survey (SLOOS), however, its quarterly periodicity, time lag to release, and short sample are significant drawbacks. There are measures of banking stress that trade in capital markets, such as the well-known TED spread, but these too have their problems.[3] Finally, there are other indexes of financial stress that mostly use principal components analysis of fairly large numbers of series, including some series we use, as well as banking related data, and the levels of interest rates which we prefer to avoid.[4] They share some similarities to the one we use, but these typically do not go back as far as the FSI.

Table 2 below describes the constituent parts of the FSI. As can be seen, the index includes two variables that measure risky spreads on bonds (#1 and 2), two that capture liquidity premiums on bonds (#6 and 7),[5] three variables that capture market volatility as measured from options prices (#4, 5 and 9) in bond and equity markets, a variable measuring the slope of the term structure at the short end (#3), and finally a measure of the equity premium (#8). Data availability limits the start date to 1988:12; the last observation we use is 2011:12, leaving 277 observations.

[3] The TED spread is the difference between interbank lending rates and the rate on short-term US Treasury securities. In normal times, these should be very close substitutes, but when counterparty risk is an issue, the spread between the two can widen. The definition of the TED spread has changed over time. The LIBOR-OIS spread, which is arguably better than the TED spread for some purposes, only goes back to 2001. Both of these indexes measure only a subset of the phenomena captured by the FSI.

[4] The St. Louis Fed's STLFSI is the first principal component of a variety of variables, some of which that are also in the FSI, plus the levels of some interest rates. It starts in 1993. For details, see Kliesen and Smith (2010). The Cleveland Fed's CFSI uses daily data from credit, foreign exchange, equity and interbank markets and dates back to 1994. See also Oet et al. (2011). The Kansas City Fed's index (KCFSI) is constructed using principal components of 11 monthly financial market variables. See Hakkio and Keeton (2009) for details. The Chicago Fed produces an index (NFCI) that is a dynamic factor of an unbalanced panel of mixed frequency indicators of financial activity. See Brave and Butters (2012) for details.

[5] The on-the-run premium is the difference in yield between just-issued Treasury bonds and the identical bond from the previous auction, corrected for the difference in term to maturity.

Table 2
Components of the Federal Reserve Board staff's Financial Stress Index*

#	Description	Source	Stddev
1.	AA rate-Treasury spread, const. maturity	Merrill & Bloomberg	66.3
2.	BBB rate-Treasury spread, const. maturity	Merrill & Bloomberg	96.2
3.	Federal funds rate less 2-yr Treasury yield	FRB & Bloomberg	0.70
4.	10-year Treasury bond implied volatility	Bloomberg	1.40
5.	Private long-term bond implied volatility	Bloomberg	2.30
6.	10-year Treasury on-the-run premium	Bloomberg	9.43
7.	2-year Treasury on-the-run premium	Bloomberg	3.60
8.	S&P 500 earnings/price less 10-year Treasury	I/B/E/S & FRB	2.01
9.	S&P 100 implied volatility (VIX)	Bloomberg	8.53

* The FSI is a simple demeaned sum of the nine components shown, weighted as a function of the inverse of their sample standard deviations.

The components of the FSI capture different aspects of risk and uncertainty in capital markets. Risk premiums, for example, reflect default risk whereas liquidity premia capture unwillingness to trade. The two concepts are likely to be associated but are not the same. In general, the components are correlated, of course, and sometimes quite strongly, but not so much that one would argue that a series is redundant. We explore modifications of, and alternatives to, the FSI in Section 5 and in the online appendix.

Figure 1 shows the FSI at a monthly frequency. The first thing to notice about the index itself is that it does not look like a stationary process with Gaussian disturbances; rather, the index appears to have lengthy periods of low readings with modest fluctuations, together with shorter episodes of high levels and volatility. This impression is reinforced by our overlay of some of the key dates in U.S. financial history discussed in the previous subsection. Clearly, the periods of what the unaided eye sees as high stress are associated with well-known events in financial history, with the period beginning with the forced merger of Bear Stearns standing out as one of particularly high stress. That said, it is worth noting that not every recession—the NBER datings of which are marked in gray in the figure—is associated with financial stress, and not every period of high levels of the FSI is associated with a recession. And finally it is not the case that every headline generating event manifests itself in high stress: the Peso crisis in 1994-95 generated much discussion, and a great deal of activity at the U.S. Treasury, and yet resulted in scarcely any movement in the FSI. The level of FSI is not a sufficient statistic for assessing economic outcomes; as we show below, the interaction of stress with the rest of the economy is key to understanding the role of stress.

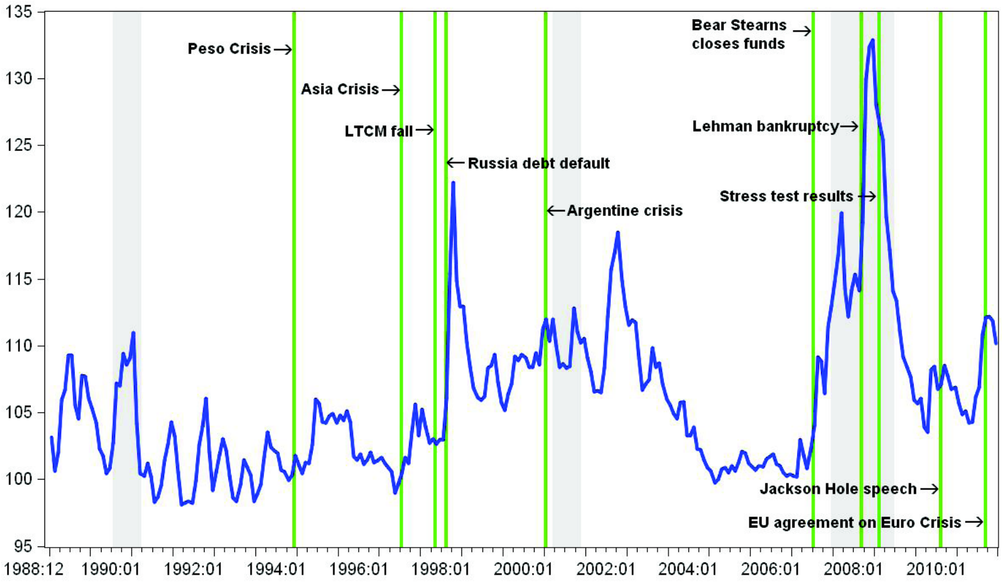

Figure 1: The Federal Reserve Board staff's Financial Conditions Index (FCI), 1988:12-2011:12

3 Model Specification, Estimation and Evaluation

The Markov-switching framework we employ is ideal for our purposes for several reasons. First, and most obviously, it provides a formal framework to investigate the presence of nonlinearities. Moreover, it does so by allowing discrete shifts, which is more appropriate than, say, a time-varying-parameter framework since drifting parameters will be unable to pick up the flight-to-safety phenomena that often occurs in financial markets. Second, it can distinguish between variance switching and coefficient switching. Regime switching in coefficients would suggest either that agents change their behavior during episodes of high financial stress, or that the environment they face is materially different; taken at face value, variance switching suggests that financial crises are a matter of happenstance. And third, the MS-VAR framework allows us to investigate feedback and amplification effects between the real and the financial sector.

The combination of high dimensionality of the model, combined with the relatively short sample of data with which we must work, presents a challenge from an econometric point of view. We address these challenges by employing state-of-the-art Bayesian econometric tools for MS-VAR models, as developed by SWZ (2008). In this section, we lay out the basic model and discuss our methodology.

3.1 The model

We consider (possibly) nonlinear vector stochastic processes of the following form:

$$y_t' A_0(s_t^c) = \sum_{l=1}^{p} y_{t-l}' A_l(s_t^c) + z_t' C(s_t^c) + \varepsilon_t' \Xi^{-1}(s_t^v), \tag{1}$$

where y is an $n \times 1$ vector of endogenous variables; s^m, $m = \{v, c\}$ are unobservable (latent) state variables, one each for *v*ariances, v, and intercepts and *c*oefficients, c; p is the VAR's lag length; z is a matrix of exogenous variables which we are going to take as 1_n—that is, a column vector of constants. A_0 is an $n \times n$ matrix of parameters describing contemporaneous relationships between the elements of y, $C(k)$ is an $1 \times n$ vector of parameters of the exogenous variables and $A_l(k)$ is a $n \times n$ matrix of parameters of the endogenous variables. The values of s_t^m are elements of $\{1, 2, ...h^m\}$ and evolve according to a first-order Markov process:

$$\Pr(s_t^m = i | s_{t-1}^m = k) = p_{ik}^m, \quad i, k = 1, 2, ...h^m. \tag{2}$$

Letting $A_+' = [A_1(k)', A_2(k)', \ldots A_p(k)', C(k)']$ and $x_t' = [y_{t-1}', \ldots y_{t-p}', z_t']$, the model can then be written as

$$y_t' A_0(s_t^c) = x_t' A_+(s_t^c) + \varepsilon_t' \Xi^{-1}(s_t^v), \quad t = 1, 2...T \tag{3}$$

where T is the sample size. Let us designate $Y^t = \{y_0, y_1, ...y_t\}$ as the vector y stacked in the time dimension. We assume that the structural disturbances are normal, conditional on the state $p(\varepsilon_t | Y^{t-1}, s_t^m, A_0, A_+) \sim N(0_{n\times 1}, I_n)$. The reduced-form system is then:

$$y_t' = x_t' B(s_t^c) + u_t'(s_t^v, s_t^c), \qquad t = 1, 2, ...T \tag{4}$$

with

$$B(s_t^c) = A_+(s_t^c) A_0^{-1}(s_t^c) \tag{5}$$

$$u_t'(s_t^v, s_t^c) = A_0'^{-1}(s_t^c) \varepsilon_t' \Xi^{-1}(s_t^v) \tag{6}$$

$$E(u_t(s_t) u_t(s_t)') = (A_0(s_t^c) \Xi^2(s_t^v) A_0'(s_t^c))^{-1}. \tag{7}$$

As can be seen in equations (5) through (7), the reduced form contains structural parameters and shocks that make distinguishing regime switching impossible, whereas it is possible in the structural form, equations (3). More important for our application, notice that switching in the coefficients, s^c, imparts switching in the reduced-form residuals, equations (7), as does switching in the structural variance-covariance matrix, through s^v. To see the significance of this, consider a model in which only coefficient switching is permitted, so that s^v drops out of equations (6) and (7). There is still time variation in reduced-form shocks and coefficients, (5)-(7), but that variation is inextricably tied by a single Markov process. Now consider switching in structural shock variances only, so that s^c drops out of (5)-(7). In this instance, the reduced-form coefficients, (5), are fixed, but the shocks can vary in an unstructured way.

At one level of abstraction, fitting a Markov switching model is an exercise in giving interpretation and meaning to what, in the context of a single-regime model, would be considered outliers; allowing arbitrary non-normalities in shock processes is a highly flexible way of doing this, whereas coefficient switching is less so. It follows that empirical evidence of coefficient switching is likely to be harder to obtain than for variance switching.[6] It should be clear from equations (4) to (7) that for a given dataset, the more s^v accounts for variability in the data, the smaller the role of s^c to explain the variability in the data, and vice versa. Thus it will be important to ensure that variance switching is not wrongly attributed to coefficient switching; it also follows that a finding of coefficient switching in a model that also allows for variance switching will be a noteworthy outcome.

In December 2008 the Federal Reserve lowered the federal funds rate to the zero lower bound (ZLB) where it stayed for the remainder of our sample. Our model handles the ZLB bound in two ways. First, and most straightforwardly, the ZLB can be thought of as simply another regime which the model can pick out, if warranted. Specifically, once the ZLB is obtained, the perception, if applicable, that the funds rate can fall no further would be captured by switching in coefficients that would rule out shocks from equations other than the federal funds rate equation resulting in negative values of the funds rate, plus switching in shock variances such that negative shocks to the funds rate do not obtain.[7] Second, there could be a change in the relationship between the federal funds rate and the stock of money either directly because of the ZLB, or because of nonstandard monetary policy measures that stand in for conventional monetary policy. Indeed, this is one reason why money growth is included in our model. Thus, the model can, in principle, pick out new states to capture the ZLB.

3.2 Estimation and evaluation

To estimate the model, we employ a blockwise optimization algorithm to find the posterior mode. As described in SWZ08, this methodology improves over, for example, the MCEM

[6] The importance of this issue is demonstrated by the debate between Cogley and Sargent (CS 2002) and SZ (2006) on the origins of the Great Moderation. CS (2002) argued that "good policy" as captured by drifting in the parameters of their VAR explained the Great Moderation; SZ (2006) showed that the omission of time variation in shock variances could bias results toward shifts in coefficients: "good luck" was responsible. CS (2005) revisited the issue allowing for stochastic volatility, and found "substantial variation" in all contributors, including coefficients. They also showed that tests of the time-invariance of coefficients of VARs in the presence of stochastic volatility have low power.

[7] Sveriges Riksbank, the central bank of Sweden, established that the nominal policy rate can be less than zero when it reduced the deposit rate to -0.25 percent in July 2009.

method proposed by Chib (1996), particularly for large-dimensional systems. In a first step, parameters are divided into blocks and the resulting initial guesses for the parameters are used in a hill-climbing quasi-Newton optimization routine. To be sure that the estimated posterior mode is a robust maximum, we perturb each maximum point with both large and small steps to generate new starting points from which we recommence the optimization process. The posterior modes described in the paper are the peak values obtained from this process.

Two sets of priors are applicable for the model, one for the VAR parameters, the other for the state transition matrix. Following SWZ (2008) we use a standard Minnesota prior for the VAR parameters. However the priors we employ for the VAR parameters are weaker than the ones suggested by SZ (2006) for monthly data. For the state transition matrix, the Dirichlet prior is used. The key prior here is the prior probability of remaining in the same state in the next period as in the current period. A prior that is reasonable for the problem under study, is one that does not promote, *a priori*, a finding of more switching in one part of the model over switching in another—in this context, switching in shock variances versus switching in coefficients. The online appendix provides some more remarks on priors.

To evaluate models in terms of goodness of fit, we compare the marginal data densities (MDDs) of candidate specifications, consistent with accepted practice. A number of alternative methods have been promoted for computing MDDs, beginning with the standard modified harmonic mean (MHM) calculation of Gelfand and Dey (1994). However, it has been established that the MHM computation is not likely to work well with models whose posterior distributions may be far from Gaussian as is the case with many Markov switching models. At least three alternatives have been proposed that use weighting functions to approximate the unknown posterior distribution, including the bridge method of Meng and Wong (1996), a method suggested by Ulrich Müller of Princeton University in an unpublished paper and detailed in Liu, Waggoner and Zha (2011, Section V.1), and a method by Waggoner and Zha (2012, Appendix B). We found in experiments using artificial data that the method of SWZ (2008) was the most reliable for our purposes.[8]

4 Macro-financial Linkages and Financial Stress

We focus on five-variable MS-VARs identified using the well-known Choleski decomposition. In particular, let $y_t = \begin{bmatrix} C & P & R & M & S \end{bmatrix}'$ where C is the monthly growth in personal con-

[8] The W-Z method is designed to reduce the sensitivity of the MDD computations to the construction of the weighting matrix by taking into account the overlap between the weighting function and the posterior distribution.

sumption expenditures (PCE); P is CPI inflation, excluding food and energy prices (hereinafter, core inflation); R is the nominal federal funds rate; M is growth in the nominal M2 monetary aggregate; and S represents the financial stress index. All variables are monthly (or monthly averages of daily rates, where applicable), seasonally adjusted, and expressed at annual rates. The data run from 1988:12 to 2011:12.[9]

We are interested primarily in three questions: first, whether there are periods of high financial stress, and if those periods are marked by different dynamics than more normal times; second, if there is evidence of regime switching, whether it is confined to variance switching, as SZ (2006) find in a different context, or whether differences in economic behavior, as captured by coefficient switching, better explain the data; and third, whether any regime switching is confined to specific equations—such as the stress equation alone, or the monetary policy response to stress—as opposed to applying to all equations.

With regard to model selection, Bayesian econometrics lends itself to model assessment on the basis of comparing the marginal data density (marginal likelihoods) of alternative models. While we carry out comparisons of this nature, we use broader criteria for model selection, placing some weight on the plausibility of the model, as captured by the state probabilities and the economic interpretation of their timing and duration in the light of past events.

4.1 Financial stress regimes: Is it just the shocks or do agents change behavior?

At this point, it is useful to introduce a bit of notation in order to facilitate the presentation of results. We designate $\#v, \# = 1, 2, 3$ to indicate the number of independent Markov states governing *variance switching*, and $\#c$ to indicate the number of states governing *coefficient switching* (that is, slope and intercept parameters). Also, when shifts in structural parameters are constrained to a particular equation(s), the restriction is indicated by prefixing the letter of the variable, $l = \{\}, C, P, R, M, S$, with $\{\}$ representing a null entry. So, for example, an MS-VAR with two Markov states in the variances and two in coefficients with the latter restricted to the financial stress variable would be designated as $2vS2c$.

Our results are summarized in Table 3. Let us focus, for the moment, on panel (a) which shows outcomes for "general models", in which switching is entertained in all equations but could be in either variance switching alone or in variances and coefficients. The first line of the panel shows the MDDs. The second line reports the difference in MDD for the applicable model

[9] The limiting factor in taking the data back further in history is the financial stress index. No meaningful extension of the index further back in time is possible without unduly narrowing the composition of the FSI.

from that of the best fitting model in the same table. The third line is essentially a reference item that shows the rankings of models by posterior mode thereby allowing the reader to see whether the method we employ for computing MDDs materially affects the ranking: they do not.

There are a number of interesting observations that can be taken from panel (a). First, a model with constant coefficients and constant shock variances, the $1v1c$ model, shown in column [1]–is not favored by the data: extensions of the model to add a second state in variances—column [2]—or in coefficients—column [4]—improve the fit, and substantially so. It follows from this that the transmission of stress in the US economy is properly thought of as a nonlinear phenomenon, or a non-Gaussian one, or both. Second, while Stock and Watson (2012) using a dynamic factor model, argue that the Great Recession arose from an unusual sequence of shocks, we can say with some assurance that allowing for coefficient switching is beneficial, a result that we show below to be robust.[10] The comparison of the $2v1c$ model in column [2] with that of the $2v2c$ model in column [5] provides an example: the improvement in fit from adding switching in coefficients is of the order of 60 in terms of log MDDs, which is very large; by comparison, adding a third Markov state for variances, as in column [3], improves the fit only in relatively small ways. Thus, the transmission of crises is not merely a non-Gaussian phenomena, but a non-linear one as well. Third, of the models shown in panel (a), the best model, on goodness-of-fit criteria, is the $3v2c$ model, shown in column [6].[11]

Table 3
MS-VAR estimation results

(a) general models

	[1]	[2]	[3]	[4]	[5]	[6]
model –>	$1v1c$	$2v1c$	$3v1c$	$1v2c$	$2v2c$	$3v2c$
MDD	-2569.7	-2438.4	-2425.0	-2464.7	-2366.9	-2349.1
- diff. from best	*-220.6*	*-89.3*	*-75.9*	*-115.6*	*-17.8*	*0*
posterior density	-2286.9	2213.8	-2113.6	-2169.2	-2076.2	-2047.6

(b) restricted models

	[7]	[8]	[9]	[10]	[11]	[12]
model –>	$3vS2c$	$3vSC2c$	$3vSCP2c$	$3vSRM2c$	$3vRM2c$	$3vRMC2c$
MDD	-2438.1	-2397.1	-2370.4	-2408.4	-2438.1	-2383.5
- diff. from best	*-89.0*	*-48.0*	*-21.3*	*-59.3*	*-89.0*	*-34.4*
posterior density	-2115.8	-2102.5	-2055.1	-2098.5	-2067.2	-2078.0

Notes: marginal data densities (MDDs) and posterior modes are in logarithms.

[10] The Stock and Watson (2012) approach has the advantage of taking into account a wider range of information than we use, but does not formally account for nonlinearities as in our model.

[11] Based solely on MDD computations, an even more elaborate model, the $3v3c$ specification, is better still, albeit only slightly. However, the $3v3c$ model's economic dynamics are difficult to interpret. And, unlike the models shown in the table, the ranking of models based on the posterior densities does not accord with the rankings by MDDs for the $3v3c$ specification.

4.2 Whence switching: is it just in stress or everywhere?

We now turn our attention to models with coefficient switching restricted to certain equations, and compare their goodness of fit to the $3v2c$ base case. Financial crises could be associated with different financial sector behavior, but with macro and policy responses unchanged; or it could be that changes in financial sector behavior induce changes in monetary policy, but the real side of the economy responds normally; or something else.

An assortment of restricted models was entertained, the most relevant of which are summarized in panel (b) of Table 3. Our primary focus is on restrictions of coefficient switching to the financial stress equation, either alone, S, or in combination with the real economy, SC; or in combination with monetary policy, SRM . From the perspective of the monetary authority, a shift to a period of high financial stress is an exogenous event that puts the authority in a quandary: does it stick to its policy rule because consistency in policy is important, or does it switch to a policy that is germane to the conditions of the day? If the former is the case, switching will be observed in the S equation but not in the policy equations; otherwise both sets of equations will exhibit switching. Finally, we looked at cases of switching in monetary policy either alone, RM, or in policy and the real economy, RMC.

Panel (b) shows that the data strongly favor switching in all equations, over the restricted specifications. This means that the dynamics of monetary policy have differed over recent monetary history, and these changes have coincided with changes in financial stress and other variables. Indeed, although this causality cannot be formally tested, it seems reasonable to assume that changes in the behavior of financial stress induced concomitant changes in the operation of monetary policy. At the same time, the limits to what monetary policy can do are indicated by the fact that shifts in monetary policy induced by shifts in financial stress were insufficient to leave the behavior of the real economy and inflation unchanged.

Omitted from formal presentation here are results for models that restrict shock variance switching to subsets of equations. We consistently found that models embodying such restrictions were inferior, in terms of goodness of fit, to unrestricted alternatives. This finding supports the argument, advanced in section 3.1, that it might be the flexibility of (unrestricted) variance switching that allows it to "push out" coefficient switching as a source of time variation in the data. That we find that coefficient switching is helpful in explaining the data even in the presence of unrestricted variance switching is thus all the more noteworthy.[12]

[12] Also of interest is the fact that models that restrict variance switching to the monetary variables are not favored by the data, as was the case for switching in the coefficients for those equations. This suggests that the Fed's nonstandard policy measures—large-scale asset purchase programs, interest on required reserves, and maturity extension and reinvestment policies—do an adequate job of standing in for conventional policy. Or it

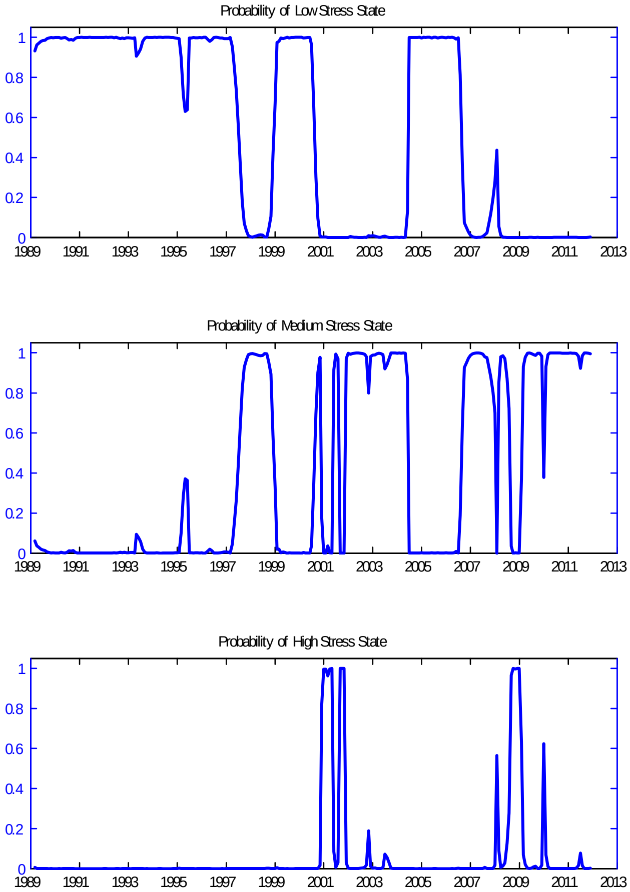

Figure 2: Probabilities of shock variance states, smoothed estimates, $3v2c$ model specification

4.3 The economic history of stress: state probabilities

Figure 2 below shows the (smoothed, or two-sided) estimated state probabilities for shock variances for the preferred $3v2c$ specification, which we treat as our base case. As can be seen, what we will call the *high-stress variance state*, shown in the bottom panel, is not a common one, although there are periods other than the crisis of 2008-9 that are identified. The first cluster of high-stress variance states begins in December 2000 when the tech-stock boom was cresting and ends in September 2001; the second has a spike in February 2008, when Northern Rock was nationalized by the British government, and another in September 2008, the month that Lehman Brothers declared bankruptcy.

Of greater interest is the probability of being in a *high-stress coefficient state*, because to be in such a state suggests fundamental differences in economic behavior—differences in the transmission of crises—as opposed to just enhanced volatility. As shown in Figure 3, there have been perhaps six periods of high stress in coefficients. The first is a cluster in the early

could simply mean that the period of the ELB is too short to be picked out of the data.

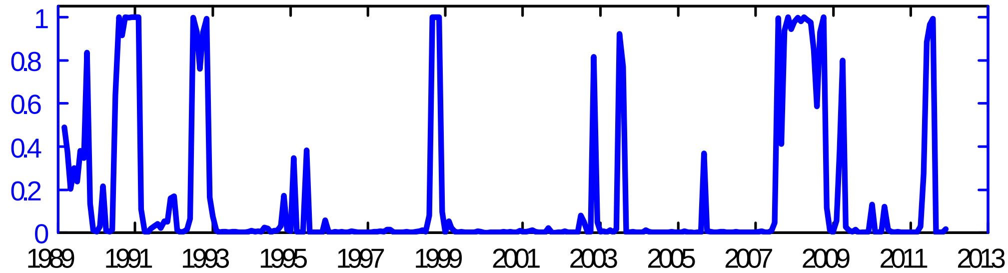

Figure 3: Probability of high-stress coefficient state, smoothed estimates, $3v2c$ model specification

part of the sample beginning in July 1990 with German reunification and ending in February 1991, at the end of the Persian Gulf war. The second begins in July 1992 and lasts until November 1992, around the time when Britain and Italy were forced by speculative attacks off of the European Exchange Rate Mechanism. The third period, in 1998, corresponds with the Russian debt default and the collapse of Long-Term Capital Management. The fourth period, two short-lived spikes in November 2002 and July 2003, matches up well with the aftermath of the Argentine debt default, or perhaps the bankruptcy of Worldcom, while the fifth, which begins in August 2007 and ends in April 2009, is the period of the 2008-9 financial crisis and associated recession. Of note is that former date, August 2007, matches exactly the beginning of the run on the repo market described by Gorton (2010), while the latter date corresponds with the leaking of the results of U.S. bank stress tests. Finally, there is a short-lived spike beginning in June 2011 which lines up with a variety of developments in the European sovereign debt crisis. Overall there are 4 periods in which a medium- or high-stress variance state prevailed *and then* the economy transitioned into the high-stress coefficient state: September 1998, July 2003, August 2007 and June 2011, all dates of prominence in U.S. financial history. There are no periods during which a high-stress coefficient state preceded a jump in the shock variance state to medium or high stress from a lower level.[13]

Taking Figures 2 and 3 together helps us understand the Great Recession. From Figure 2 we see that the period from 2004 to 2006 was a lengthy one of the low-stress variance state (the upper panel of the figure); Figure 3 shows that this was also a period in which the coefficient state was low stress as well. Figure 1 tells us that this was also the period in which the FSI itself was at a very low level—and showed little variation over time. In addition, the level of interest rates was very low and stable. It is commonly alleged that financial firms

[13] A comparison of figures 1 and 3 reveal that it is *not* the case that one need only observe a high level of the FSI to conclude that one is in a high-stress coefficient state, or vice versa. It is the joint behavior of the system that determines the Markov state.

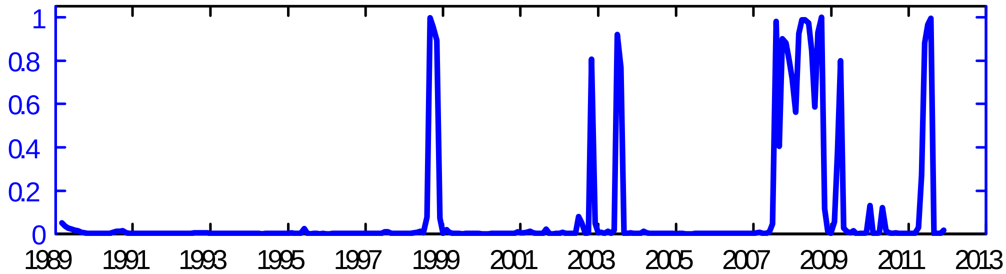

Figure 4: "Stress events" in recent U.S. economic history, defined as high-stress coefficient states coincident with high- or medium-stress shock variance states

began "chasing yield" in response to this state of affairs, increasing leverage in order to magnify returns; see, e.g., Geanakoplos (2010). Back on Figure 2, the economy then transitions in late 2006—about the time that prices of existing homes at the national level crested—to the medium-stress variance state (the middle panel). The crisis begins in earnest when the economy transitions in August 2007 to the high-stress coefficient state and finally reaches full bore in September 2008 when the variance state also jumps to high stress (the lower panel of Figure 2). All this leads to a proposed definition of a *stress event*: when the shock variance state is *either* medium or high, *and* the coefficient state is high. As can be seen in Figure 4 below, this definition eliminates the periods of high-stress coefficients in the early 1990s at which time there was apparently insufficient turbulence to create severe difficulties for the real economy (although there was a mild recession). Also omitted from this status is the September 11, 2001 attacks and the associated extraordinary provision of liquidity by the Federal Reserve that followed those attacks.

4.4 Real-time properties

As we noted, the FSI was constructed and used by the Fed staff in real time during the financial crisis. Figure 5 looks at the real-time efficacy of the index, showing with the lighter, cyan-colored lines, the real-time estimates of the state probabilities for the high-stress coefficient state; that is, the probability measured at each point in time based on information up to the current period.[14] Two noteworthy conclusions may be drawn from this figure. First, the switches in coefficients indicated in ex post data, the black line, were revealed in the real-time estimates, the colored lines; that is, false negatives are negligible. Second, while there are hints of false positives—for example in 1996 and 2002—at no time did the real-time data adamantly

[14] These are *quasi*-real-time estimates. There is no complete set of real-time data that would allow a full real-time assessment. That said, the FSI and the core CPI are not subject to revision. The money and real PCE data are subject to revision however.

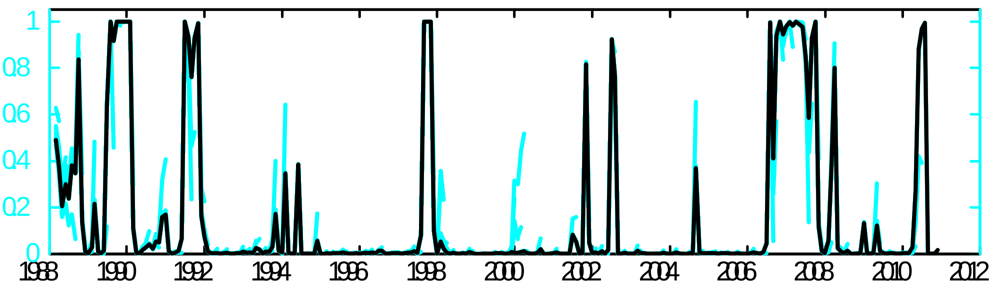

Figure 5: Probability of high-stress coefficient state, quasi-real-time estimates (lighter tone), and ex post (black)

call for a switch that was rescinded, ex post.[15] All in all, we would argue that the model does remarkably well in real time.

5 Some interpretation of results and their robustness

Our objective in this section is three-fold: first, to provide some interpretation of the structural mechanisms that are likely behind our results, mostly through an investigation of the real macro variable we use; second, to report on experiments that show the value-added of the FSI; and third, to discuss experiments that demonstrate the robustness of our results to alternative measures of stress. The economic properties of the base case model are discussed in Section 6. In order to conserve on space, the discussion here will be brief; most of the results are relegated to the online appendix.

5.1 The real variable: investment, durable goods and labor market variables

A common narrative in discussions of the Great Recession is the connection between financial stress, credit availability and expenditures, particularly expenditures on goods for which credit is seen as a strong complement, such as consumer durables, housing and business fixed investment. Sometimes the story is told in terms of the amplification and propagation of shocks because of costly state verification and associated leverage constraints, as in Bernanke et al. (1999). In other frameworks, it is collateral constraints that matter. For example, Chaney et al. (2012), and Liu et al. (2013), describe empirically and model structurally, respectively, how the value of real estate played a role in the decline of business fixed investment during the Great Recession.

One way for us to cast light on this mechanism is to study how well our model works for real variables other than PCE, in particular, for classes of durable goods. For this and a number

[15] Charts of the real-time performance of the variance states are broadly similar.

of other exercises described in this section, it is necessary to specify the basis for comparison with the results of our base case. For us, the fact that the estimated dates of coefficient switching coincide with known events in U.S. economic and financial history is compelling and so we rely on comparisons of smoothed coefficient state probabilities of our alternatives, compared with the base case we showed in Figure 3. (The online appendix also compares conditional forecasts of some alternative specifications.) We begin by splitting growth in PCE into durables and nondurable goods and services plus footware. Figure 6 shows the smoothed high-stress coefficient state probabilities for these two series, alongside our aggregate PCE base case. Note that the figure uses offset vertical scales so that the precise dates of climbs and falls of probabilities can be distinguished.As the figure shows, PCE durables picks up many

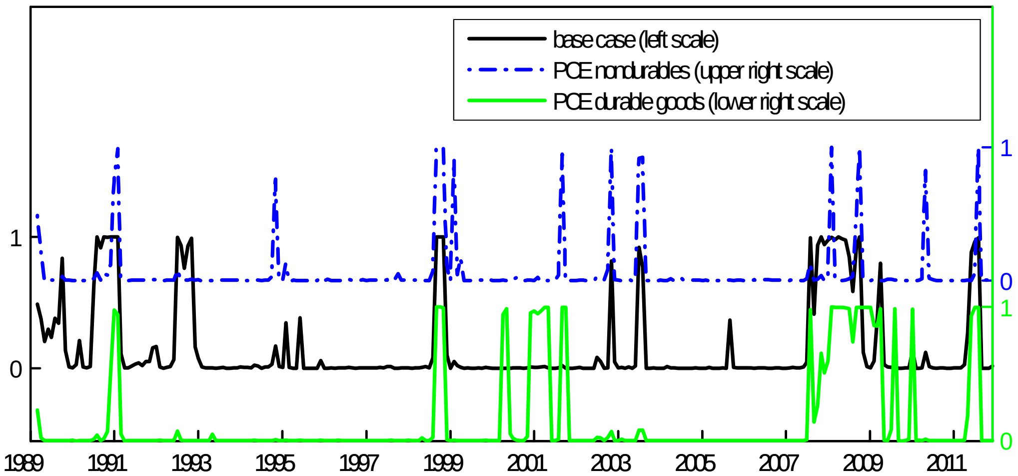

Figure 6: Probability of high-stress coefficient state, with real variable PCE nondurables, services and footware (dot-dashed blue line, upper-right scale), PCE durable goods (lighter green solid line, lower-right scale) and the base case total PCE (black solid line, left scale).

of the same key switching dates as does the base case, while nondurables performs less well. The online appendix demonstrates that this is not unique: broadly similar results obtain for a monthly interpolated version of investment in equipment and intangibles. Taken together, these results suggest to us that the role of nonprice rationing of credit, defined to include collateral constraints, and its role in durable goods expenditures, is central in propagating financial crises.

Another conventional story of the Great Recession is that extraordinary dynamics in labor markets were in play, either through mismatch in employment, perhaps connected with the sharp decline in construction and finance industries, or more generally in the outsized drop in employment, relative to output, and an associated rise in precautionary saving. To examine this proposition, we reestimated our model using each of the unemployment rate, growth in payroll employment and unemployment insurance claims as our real variable.[16] Results for all three cases, by our metric, were universally inferior to the base case.[17]

5.2 The stress index: composition and construction

In an attempt to identify which aspects of the FSI are critical to our results and thereby cast some light on the stories conventionally offered to explain the transmission of crises, we engaged in two broad classes of investigation regarding the weighting of the components of the FSI and in its construction. On the logic that financial stress is only important when it is systemic, one might argue that instead of weighting the components in the *ad hoc* way that the FSI does, a method that chooses weights to explain the maximum amount of variation of the nine components collectively would be efficacious. To test this, we reconstructed the index using the first principal component of the constituent parts of the FSI and substituted this measure in our full system, obtaining results that were very similar to our base case. All told, this result tells us, first, that the FSI's ability to capture the phenomena of interest is not an artifact of the construction of the index, and second, that the *ad hoc* weighting of the FSI turns out to have been a good one.

For our second class of experiments, we considered financial stress indexes that excluded from the index one of five blocks of components of the FSI, relative to our base case results. These five classes are risky bond rate spreads (rows 1 and 2 of Table 2), the term spread (row 3), implied bond rate volatilities (lines 4 and 5), on-the-run premiums (line 6 and 7), and equity market factors (lines 8 and 9). We found that the results we obtained for our base case were largely unchanged from exclusion of the term spread, implied bond-rate volatilities, and on-the-run premiums, and that there were some modest differences from excluding the equity premium. The more interesting differences, for a variety of reasons, were obtained from exclusion of the risky spreads.

Of the variables that comprise the FSI, risky spreads exhibit the highest correlation with the aggregate FSI; thus, it would not be surprising if these variables turned out to be critical for our findings. There are, moreover, results showing that default premiums on bonds are predictors of financial distress, with a nascent literature on the proper measurement of these premiums; see, e.g., Gilchrist and Zakrajšek (2012). To address this issue, we estimated two models focussing

[16] The Great Recession was marked by a more substantial decline in labor markets than in GDP. It seems plausible—if beyond the scope of this paper—that relatively low frequency movements in labor market conditions exacerbated the duration of the switch to the high-stress coefficient state in 2008.

[17] A concise summary of results for labor market variables is available in the online appendix.

on an index of risky spreads (hereinafter, *sprd*), created from the first two components of the FSI shown in Table 2, one in which *sprd* was *excluded* from the FSI, and the other where *sprd* *substitutes* for the FSI. The results, which appear in the online appendix, show that omitting *sprd* from the index results is a fairly substantial deterioration in estimated probabilities of high-stress coefficient state, missing some key events in history, underestimating the duration of the 2008-9 event and posting a false positive at the turn of the century. Thus, the inclusion of *sprd* is necessary for our results. The model with *spdm* alone suggests substantially fewer period of high-stress coefficients and, more importantly, misses the onset of the 2008-9 financial crisis by several months. We conclude that risky spreads alone are insufficient to pick up the information contained within the FSI.

We also tested whether the FSI is even necessary to obtain results similar to our base case, or whether the FSI alone is sufficient. In summary, in the case of systems that included macroeconomic variables, but omitted stress, we found that we tended to pick up switching somewhere during the 2008-9 episode, but not much of anything else. In the complementary case of stress alone, we found a substantial deterioration of model performance: macroeconomic variables are important for our results. These cases are covered in the online appendix.

6 The transmission of financial stress

To illustrate some properties of the model and provide some historical perspective, we carry out two classes of simulations on the model. The first are counterfactual simulations, some of which are designed to illustrate the unique features of our model in a compact and intuitive fashion, others are set around the 2008-9 financial crisis. The second class of simulations are conditional forecasts initiated from the end of the sample period. These exercises provide very much the same information as do impulse responses except more compactly, and in a more intuitive and historically appealing context.

Markov switching aside, the unique aspect of our model is the financial stress index. To illustrate how financial stress affects the economy, we carry out two counterfactual simulations involving alternative paths for stress (S in the figures), one carried out during a period when the latent state is one of low stress, the other from more strained conditions.

Figure 7 shows the effects of an autonomous increase in stress during a low-stress period in July 1989. The noteworthy aspects are two-fold: first, the monetary response is slight, with the federal funds rate (R) falling only marginally, relative to the data. The implications for real activity, measured here by growth in personal consumption expenditures (ΔC) in the upper-

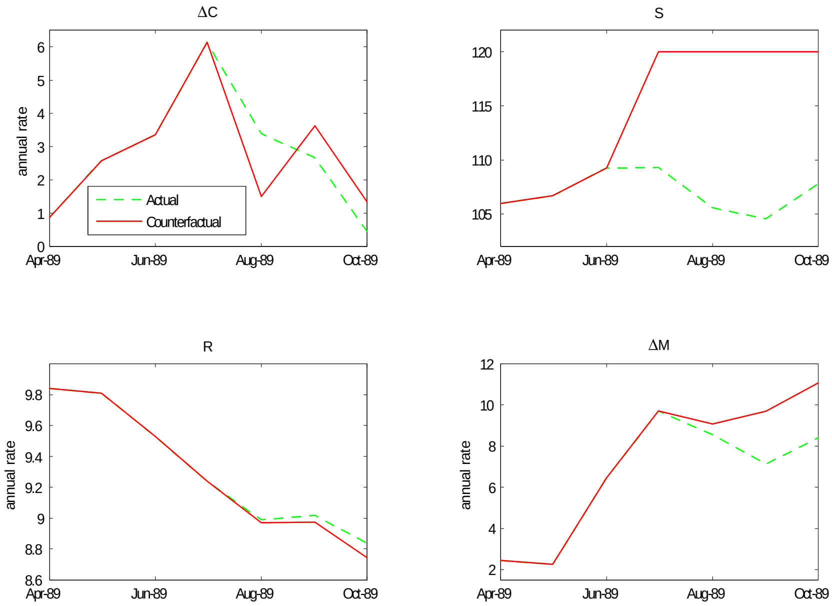

Figure 7: Counterfactual experiment where financial stress, (S), rises to 120 in July 1989, a normal-times period, base-case $3v2c$ specification

left panel are relatively small and short lived. Thus, this exercise ratifies our assertion that financial stress has been underappreciated through much of economic history as an important factor in the transmission of business cycles because in normal times—that is, through the bulk of history—stress has not been a major driver of events.

Figure 8 carries out a broadly similar exercise, this time from August 1998, during the Russian debt default and associated collapse of LTCM. In the data, S climbed rapidly and substantially with the onset of the crisis; our counterfactual imagines that stress had instead remained low. Unlike in Figure 7, in this instance there is a substantial monetary policy response, offsetting the expansionary implications of the lower level of stress. The implications for real activity end up being quite modest. What this says is that monetary policy, *when it has the capacity to do so*, is well disposed to respond to increases in stress, holding constant the stress regime, *when those increases are moderate and temporary,* as was the case in 1998. Arguably, actions by the Federal Reserve to elicit an orderly reorganization of LTCM ensured that this stress event was brief, and monetary policy defined in terms of setting the federal funds rate was in a position to ease. The contrast with the 2008-9 financial crisis is fairly stark. The shock in the latter instance was larger, as shown in Figure 3, the stress event lasted longer, and conventional monetary policy was limited in its ability to respond.

Let us now turn to the recent financial crisis and consider counterfactual changes in regime.

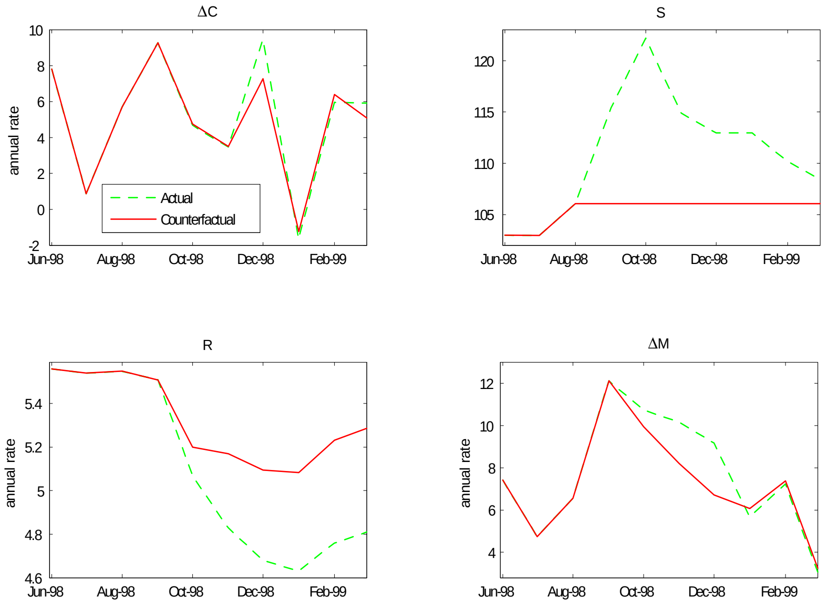

Figure 8: Counterfactual experiment where financial stress, (S), is held at its August 1998 level, high-stress coefficient state, base case $3v2c$ specification

Model estimates show, and Figure 9 confirms, that a stress event began in the second half of 2007. The economy had already switched to the medium-stress variance state late in 2006—by itself not a big deal but sometimes a precursor to worse things—followed by a persistent switch to high-stress coefficients in October 2007; then, in September 2008, the state switched to high-stress variances together with the already existing high-stress coefficients. In Figure 9 we pose the question, what would have happened, according to the model, if the state had remained in the low-stress coefficient state?

We allow all the shocks borne by the economy to remain in play; the only thing that is counterfactual here is the set of coefficients through which those shocks play out. The figure shows that financial stress itself (S), would have been much lower than otherwise; this, in turn, would have obviated the need for very easy monetary policy, so that the federal funds rate (R) ends up about 2-1/2 percentage points higher than in history by mid-2008, and money growth would have been lower.[18] Tighter monetary policy notwithstanding, real growth would have been notably stronger than the historical experience. Clearly, the implications for the economy of a persistent, adverse switch in Markov states—that is, a stress event—are substantial.

Figure 10 considers a different counterfactual carried out over the same period beginning in

[18] Inflation, not shown here, would have been higher in this scenario. We omit that panel of this and other charts, to keep the figure compact.

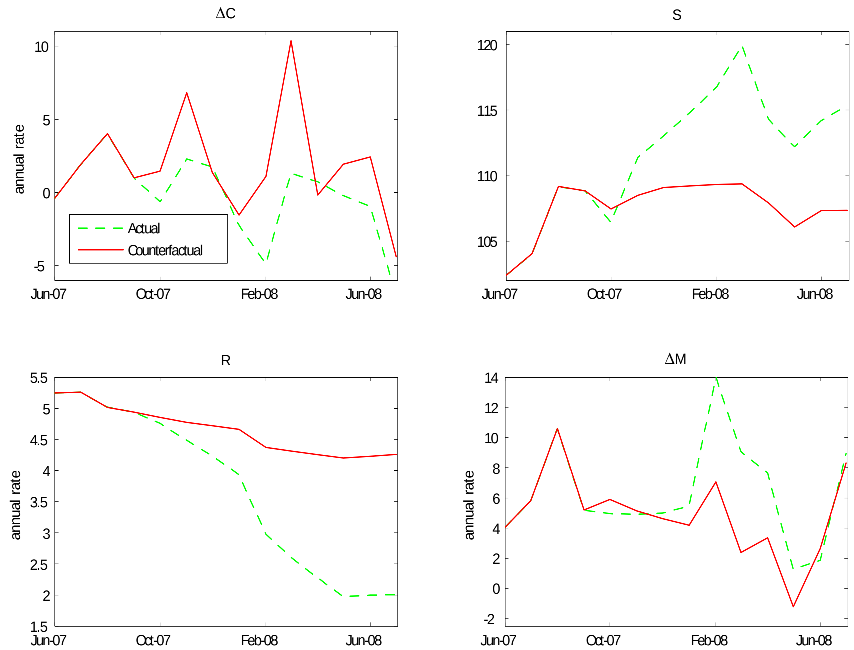

Figure 9: Counterfactual experiment where the latent state returns ot the low-stress coefficient state in October 2007, base case $3v2c$ specification

October 2007. We suppose that the Federal Reserve could have foreseen the grave conditions that were to come and thus immediately reduced the federal funds rate to the de facto zero lower bound of 0.12 percent.

As can be seen from the bottom-left panel, this is a large intervention, which induces a very large increase in money growth, the bottom-right panel. The effect on real activity is relatively small, however. The upper-right panel gives an indication of why this is so: financial stress *rises* measurably and persistently with the policy intervention. Evidently, in high-stress situations, agents regard conventional policy actions that would normally be beneficial as confirmation of incipient financial difficulties. The resulting higher levels of stress choke off the salutary effects of easy monetary policy. This observation may help explain why the recessions caused by financial crises tend to be long lasting; see, e.g., Reinhart and Rogoff (2009). We emphasize that this result is germane to stress events: in normal times, a surprise reduction in the federal funds rate reduces financial stress, as one might expect. We conclude that conventional monetary policy actions, in the absence of actions to alleviate the fundamental causes of the stress event, or actions to arrest increases in financial stress, will only be modestly helpful for economic performance. At one level, this should not be surprising: it is received wisdom in economics that would-be policy cures should be tailored to the ultimate causes of the problem as opposed

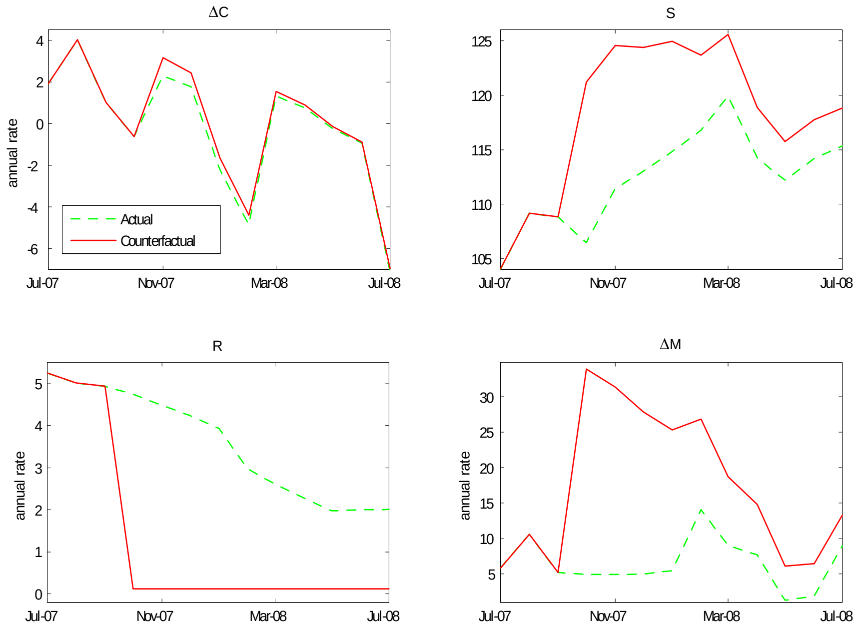

Figure 10: Counterfactual experiment where the federal funds rate, (R), falls to 0.12 percent in October 2007, base case $3v2c$ specification

to the symptoms that those causes engender.

Finally, we turn to our second class of experiments, conditional forecasts that illustrates the importance of latent state for economic outcomes. These conditional forecasts are carried out, as were the counterfactuals described above, using our base case model; however Section 5 (and the online appendix) show that similar results obtain when durables are used as the real variable. Figure 11 shows two forecast paths beginning immediately at the end of our sample in 2011:12, one (the red solid line) conditional on a high stress regimes in both coefficients and variances, the other (the blue dashed line) on a low stress in both coefficients and variances. All else is held constant, and unlike in the counterfactuals, there are no shocks in the simulation period.

As can be seen, PCE growth is much weaker in the high-stress world and this low growth is accompanied by elevated levels of financial stress, particularly in comparison with the low-stress world. Of significance is that the high-stress state is associated with *higher* price inflation than in the low-stress state, a finding that is consistent with an interpretation of a stress event as a negative supply shock that reduces real output and puts upward pressure on prices, all else equal, an interpretation that is in line with that of Jermann and Quadrini (2012) and de Fiore and Tristani (2013). All else is not equal here: monetary policy, as measured by the federal funds rate (or the growth rate of M2, not shown) is easier in the high-stress world than

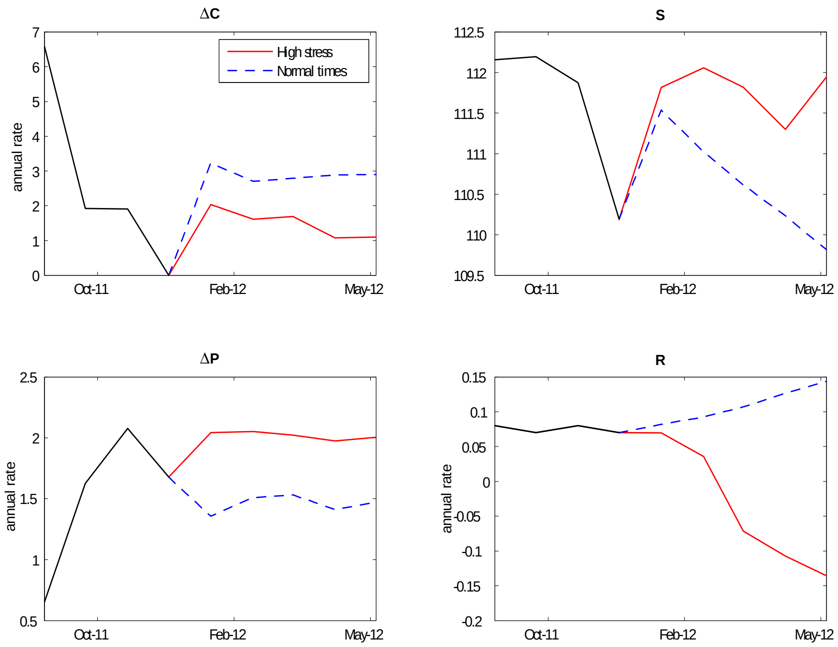

Figure 11: Model forecast, conditional on the state, from 2011:12, base case $3v2c$ specification. High-stress coefficient state (red solid lines) versus low-stress coefficient state (blue dashed lines)

otherwise; but with the interpretation of reduced potential output, this easy monetary policy is seen as something of a palliative that reduces the pain only modestly.

7 Conclusions

This paper has considered the implications of financial stress for the macroeconomy using a richly specified Markov-switching vector autoregression model, estimated with state-of-the-art Bayesian methods, and exploiting a unique series for financial stress constructed and used in real time by the staff of the Federal Reserve Board.

Our analysis showed substantial evidence that a single-regime model of the macroeconomy and financial stress is inadequate to capture the dynamics of the economy. We demonstrated that there have been periodic shifts not just in the stochastic shocks that have buffeted the economy, but also in the dynamic propagation of shocks, with all equations of the model showing evidence of switching. It follows that inference regarding the conduct of monetary policy that is gleaned from a constant-parameter Gaussian model may be inappropriate for periods when the policy is conditioned on movements in financial stress.

Quantitatively, we find that output reacts differently to financial shocks in times of financial stress than in normal times: Stress is of negligible importance in "normal" times, but of critical importance when the economy is in the high-stress coefficient state. We also found that an important precursor to particularly adverse economic events is a switch to what we call a stress event: a period in which the latent state for shock variances is relatively high and the latent Markov state for coefficients is also at a high-stress level. And we showed that the Federal Reserve Board staff's use of the financial stress index described in this paper appears to have been an efficacious choice. Our results also suggest that conventional monetary policy is not particularly effective in times of high financial stress.

Lastly, in digging deeper into our results, we uncovered an interpretation emphasizing risky spreads as a key component of financial stress on the one hand, and durable goods as a real variable, on the other. This suggests to us that structural models aimed at explaining the phenomena studied in this paper would be well advised to assign a prominent role to perceptions of default risk, their role in eliciting occasionally binding constraints on lending, and contagion across markets and over time.

References

[1] Bernanke, B.S., Blinder A., 1988. Credit, money and aggregate demand.American Economic Review 78, 435-439.

[2] Bernanke, B.S., Gertler, M., Gilchrist, S., 1999. The financial accelerator in a quantitative business cycle framework. In: Taylor, J.B., Woodford, M. (Eds.) Handbook of Macroeconomics. Elsevier, Amsterdam, pp. 1341-1393.

[3] Bianchi, F., 2013. Regime switches, agents' beliefs and post-World War II U.S. macroeconomic dynamics. Review of Economic Studies 80, 463-490.

[4] Bianchi, J., 2011. Overborrowing and systemic externalities in the business cycle. American Economic Review 101, 3400-3426.

[5] Boissay, F., Collard, F., Smets, F., 2013. Booms and Systemic Banking Crises. European Central Bank Working Paper 1514.

[6] Brave, S., Butters, R. 2012. Diagnosing the financial system: Financial conditions and financial stress. International Journal of Central Banking 8, 191-239.

[7] Brunnermeier, M., Sannikov, Y., 2014. A macroeconomic model with a financial sector. American Economic Review 104, 379-421.]

[8] Carlstrom, C.T., Fuerst, T.S., 1997. Agency costs, net worth and business fluctuations: A computable general equilibrium approach. American Economic Review 87, 893-910.

[9] Chaney, T., Sraer, D., Thesmar, D., 2012. The collateral channel: How real estate shocks affect corporate investment. American Economic Review 102, 2381-2409.

[10] Chib, S., 1996. Calculating posterior distributions and model estimates in Markov mixture models. Journal of Econometrics 75, 79-97.

[11] Cogley, T., Sargent, T.J., 2002. Evolving post-world war II U.S. inflation dynamics. In: Bernanke, B.S., Rogoff, K. (Eds.) NBER Macroeconomics Annual 2001. MIT Press, Cambridge, pp. 331-73.

[12] Cogley, T., Sargent, T.J., 2005. Drifts and volatilities: Monetary policies and outcomes in the post WWII US. Review of Economic Dynamics 8, 262-302

[13] Curdia, V., Woodford, M., 2009. Credit frictions and optimal monetary policy. Columbia University Department of Economics Discussion Papers 0809-02.

[14] Davig, T., Hakkio, C., 2010. What is the effect of financial stress on economic activity? Federal Reserve Bank of Kansas City Economic Review 95, 35-62.

[15] de Fiore, F., Tristani, O., 2013. Optimal monetary policy in a model of the credit channel. Economic Journal 103, 906-931

[16] Geanakoplos, J., 2010. The leverage cycle. In: Acemoglu, D., Rogoff, K., Woodford, M. (Eds.), NBER Macroeconomics Annual 2010. University of Chicago Press, Chicago, pp. 1-65.

[17] Gelfand, A., Dey, D., 1994. Bayesian model choice: Asymptotics and exact calculations. Journal of the Royal Statistical Society (Series B) 56, 501-514.

[18] Gilchrist, S., Zakrajšek, E., 2012. Credit spreads and business cycle fluctuations. American Economic Review 102, 1692-1720.

[19] Gilchrist, S., Sim, J.W., Zakrajšek, E. 2014. Uncertainty, financial frictions, and investment dynamics. NBER Working Paper 20038.

[20] Gorton, G., 2010. Questions and answers about the financial crisis. Manuscript prepared for the U.S. Financial Crisis Inquiry Commission.

[21] Hakkio, C.S., Keeton, W.R., 2009. Financial stress: What is it, how can it be measured, and why does it matter? Federal Reserve Bank of Kansas City Review 2, 1-46.

[22] He, Z., Krishnamurthy, A., 2012., A Model of Capital and Crises. Review of Economic Studies 79, 735-777.

[23] Hubbard, R.G., 1998. Capital-market imperfections and investment. Journal of Economic Literature 36, 193-225.

[24] Jermann, U., Quadrini, V., 2012. Macroeconomic effects of financial shocks. American Economic Review 102, 238-271.

[25] Kashyap, A., Stein, J., 1994. Monetary policy and bank lending. In: Mankiw, N.G. (Ed.) Monetary Policy. University of Chicago Press, Chicago, pp. 221-256.

[26] Kashyap, A., Stein, J., Wilcox, D., 1993. Monetary policy and credit conditions: Evidence from the composition of external finance. American Economic Review 83, 78-98.

[27] Kliesen, K.L., Smith, D.C., 2010. Measuring financial market stress. Federal Reserve Bank of St. Louis National Economic Trends (January)

[28] Liu, Z., Waggoner, D.F., Zha, T., 2011. Source of macroeconomic fluctuations: A regime-switching DSGE approach. Quantitative Economics 2, 251-301.

[29] Liu, Z., Wang, P., Zha, T., 2013. Land-price dynamics and macroeconomic fluctuations. Econometrica 81, 1147-1184.

[30] Lown, C., Morgan, D.P., 2006. The credit cycle and the business cycle: New findings using the Survey of Loan Officer Opinion Survey. Journal of Money, Credit and Banking 38, 1575-1597.

[31] Mendoza, E., 2010. Sudden stops, financial crisis, and leverage. American Economic Review 100, 1941-66.

[32] Meng, X.L., Wong, W.H., 1996. Simulating ratios of normalising constants via a simple identity: A theoretical exploration. Statistica Sinica 6, 831–860

[33] Nelson, W., Perli, R. 2005. Selected indicators of financial stability. IFC Bulletin 23, 92-106.

[34] Oet, M.V., Eiben, R., Bianco, T., Gramlich, D., Ong, S.J., 2011. The financial stress index: Identification of systemic risk conditions. Federal Reserve Bank of Cleveland Working Paper 2011-30.

[35] Reinhart, C.M., Rogoff, K.S., 2009. This time is different: Eight centuries of financial folly. Princeton University Press, Princeton.

[36] Sims, C.,Waggoner, D.F., Zha, T., 2008. Methods for inference in large multi-equation Markov-switching models. Journal of Econometrics 146, 255-274

[37] Sims, C., Zha, T., 1998. Bayesian methods for dynamic multivariate models. International Economic Review 39, 949-968.

[38] Sims, C., Zha, T., 2006. Were there regimes shifts in U.S. monetary policy? American Economic Review 96, 54-81.

[39] Stock, J., Watson, M. 2012. Disentangling the channels of the 2007-2009 recession. Brookings Papers on Economic Activity 1, 81-130.

[40] Waggoner, D.F., Zha, T., 2012. Confronting misspecification in macroeconomics. Journal of Econometrics 171, 167-184.

Appendix A

This appendix is devoted to the base case estimates summarized in the main text. It contains information on model priors, the data used and aspects of computation. A second appendix follows that discusses alternative results; that is, the robustness of our estimates and related issues.

A.1 Priors

There are two sets of priors of relevance to our model, one on the reduced-form parameters of the VAR conditional on a state, s, and the other on the transition matrix. The priors on the reduced-form VAR are simply the standard Minnesota prior of Litterman (1986) on the lag decay dampening the influence of long lags. In other words, this prior shrinks the model towards a random walk. Furthermore, it seems reasonable that the importance of a variance decreases with lag length; and that priors on exogenous and deterministic variables, z, be relatively uninformative. Let the relative tightness on the prior on the own lags, non-own lags, and exogenous or deterministic variables be μ_1 through μ_3 respectively. The prior variances of the parameters are then specified as:

$$Var(x_i) = \begin{cases} \mu_1/p & \text{for own lags} \\ \mu_2\sigma_i^2/p\sigma_j^2 & \text{for lags } i \neq j \\ \mu_3\sigma_i^2 & \text{variables } z. \end{cases}$$

The priors that apply to switching are a little less straightforward. Even without restrictions of some sort, $A_0(s_t)$ and $A_+(s_t)$ could, in principle, be estimated straightforwardly, using the method of Chib (1996) for example, but as n or h grows, the curse of dimensionality quickly sets in. The problem is particularly acute in situations where one (or more) of the unobserved states lasts for only a short proportion of the number of total observations, as may be the case for us. The matrix A_+can be rewritten as

$$A_+(s_t) = D(s_t) + \hat{S}\ A_0(s_t) \quad where \quad \hat{S} = [\ I_n \quad 0_{(m-n)\times n}\ ] \tag{A.1}$$

which means that a mean-zero prior can be placed on D which centers the prior on the usual reduced-form random-walk model that forms the baseline prior for most Bayesian VAR models; see e.g. Sims and Zha (1998) for details. The relationship defining B in the main text, namely equation (5): $B(s_t^c) = A_+(s_t^c)A_0^{-1}(s_t^c)$,means that a prior on D tightens or loosens the prior on a random walk for B.

The fact that the latent state, s,is discrete and that the transition probabilities of states must sum to unity lends itself toward the priors of the Dirichlet form. Dirichlet priors also have the advantageous property of being conjugate. Letting α_{ij} be a hyperparameter indexing the expected duration of regime i before switching to regime $k \neq i$, the prior on P can be written:

$$p(P) = \prod_{k \in H} \left[\frac{\Gamma(\sum_{i \in H} \alpha_{ik})}{\Pi_{i \in H}\ \Gamma(\alpha_{ik})} \right] \times \prod_{i \in H} p_{ik})^{\alpha_{ik}-1} \tag{A.2}$$

where $\Gamma(.)$ is the gamma distribution. The Dirichlet prior enables a flexible framework for a variety of time variation including, for example, once-and-for-all shifts and, by letting h become arbitrarily large, diffusion processes. Our application will not consider absorbing states and will keep the number of states small. We will, however, allow for switching in shock variances originating from a separate process from the one controlling shifts in parameters.

For our baseline specification, we use priors that are well-suited for a monthly model. In particular, we specify μ_k $k = 1, 2, ...6 = \{0.57, 0.13, 0.1, 1.2, 10, 10\}$ and Dirichlet priors of 5.6

for the two coefficient states and 11.33 for the three shock variance states. With the values of μ_k we begin with what Sims and Zha (2006) suggest for monthly data, except μ_1 where we use a lower number, and μ_2 which is slightly higher. The value for $\mu_1$reflects that we are interested less in shrinkage toward the random walk and more for allowing persistence. The Dirichlet priors we use are looser than what would be usually used for monthly data. They imply an 85 percent prior probability, for both shock variances and coefficients that the economy will, in the next period, continue in the same state as it is in the current period. This strikes us a fairly low probability, consistent with the notion that shifts are associated with jumps in asset prices.[19]

A.2 Robustness of priors selection

In broad terms, our preferred model is resilient to moderate changes in model priors. For example, if we alter the priors governing VAR coefficients that we used following SZ (2006) with alternatives, such as those that SZ (2006) recommend for a quarterly model, we get, once again, several periods of high-stress coefficients and many periods of switching in variances. Altering the Dirichlet prior such that higher persistence of regimes is somewhat favored returns what looks like the same results as we showed for our preferred model.

A.3 Data transformations

As noted in the main text, we use levels of the federal funds rate and the stress index and growth rates of real personal consumption expenditures (PCE), money and prices. Unit roots tests on the stationarity of these growth rates tend to be mixed, with many tests unable reject the null hypothesis of a unit root. The sole exception is money growth where the bulk of the tests reject the unit root. Similar criteria were used for data transformations of the alternative real variables that are summarized in Appendix B below.

A.4 More on the data

In the main text, we noted without proof that the risky spreads were the components of the FSI that bore the highest correlation with the index itself, and more generally that the components of the FSI are correlated, sometimes strongly so. Table A.1 shows the correlation matrix. The final row of the table shows the correlation of the components with the index as a whole. Indeed, the risky spreads, AA and BBB, stand out as being highly correlated with the FSI as a whole, followed by the VIX and then several of the liquidity premiums.

[19] There are a number of methods outlined in the literature for computing MDDs when the posterior distribution is likely to be far from Gaussian. The alternatives are all based on constructing weighting distributions as initial approximations from which the posterior distribution can be computed. The method of Waggoner and Zha (2011) that we used is designed to reduce the sensitivity of MDD calculations to the construction of the weighting matrix by measuring and taking into account the overlap between the weighting function and the posterior distribution.

Table A.1
Correlation coefficients on components of Financial Stress Index*

	risky spreads		term slope	implied volatilities		on-the-run premiums		equity prem.	
	AA	BBB	$ff-2yr$	$Tbond$	$pbond$	10 liq	2 liq	$equity$	VIX
AA $spread$	1								
BBB $spread$	0.94	1							
$ff\ -2yr$ $slope$	0.27	0.15	1						
$Tbond$ $volatility$	0.53	0.61	-0.20	1					
$pbond$ $volatilty$	0.67	0.73	-0.12	0.86	1				
$10-yr$ $liquidity$	0.69	0.75	-0.04	0.56	0.57	1			
$2-y$ $liquidity$	0.22	0.21	0.25	0.06	0.04	0.28	1		
$equity$ $premium$	0.55	0.47	0.14	0.24	0.52	0.09	-0.30	1	
VIX	0.76	0.77	0.25	0.55	0.64	0.67	0.32	0.20	1
FSI	**0.92**	**0.93**	**0.28**	**0.69**	**0.81**	**0.75**	**0.33**	**0.49**	**0.85**

* Variables appear in the same order as in Table 2 of the main text.

A.5 On the estimated state probabilities

To provide further justification for our selection of the $3v2c$ specification as the preferred one, consider Table A.2 which shows the estimated transition probabilities taken from the posterior mode of the distribution for selected model specifications. By comparing the first and third lines of the table, we see that the introduction of a second state in coefficients to what would otherwise be the $3v1c$ model changes the probabilities of the variance states quite dramatically. This finding illustrates the fact that switching in shock variances and switching in coefficients are rivals in explaining the data; as SZ (2006) have emphasized, failing to account adequately for one will bias estimates of the other. The fact that the $2v2c$ model and the $3v2c$ model are economically similar is demonstrated by the fact that the state probabilities that the two models have in common does not change markedly with the introduction of the third state in variances. In both specifications, it is the case that the high-stress coefficient state is short-lived in duration, on average. The severity of the 2008-9 episode is therefore marked by two unusual phenomena by historical standards: the fact that the high-stress coefficient state lasted as long as it did, and the fact that it was also associated with a period of high-stress shock variances. Figure 4 in the main text showed our estimates of stress events defined in this way. That figure revealed that the early sample periods of high-stress coefficients were not terribly consequential in macroeconomic terms because they were not associated with shock-variance regimes that were conducive to widespread contagion.[20]

[20] Campbell *et al.* (2013) show that default spreads—which are a part of the FSI—have regime-switching like properties for asset returns in that modest levels of volatility are good for stockholders, because they are the residual claimants on firm assets, but once volatility gets large, the effect switches sign, because the viability of the firm comes into question. This characterization of conditional dynamics is very much in the spirit of the findings in this paper.

Table A.2
Estimated transition matrix
(posterior mode)

model	variances			coefficients	
	q^v_{hh}	q^v_{mm}	q^v_{ll}	q^c_{hh}	q^c_{ll}
$3v1c$	0.80	0.89	0.89	-	-
$2v2c$	-	0.92	0.95	0.76	0.95
$3v2c$	0.83	0.93	0.97	0.73	0.95

A.6 Computation

In our MCMC computations, we use 100,000 proposal draws and 500,000 posterior draws, net, retaining every tenth posterior draw in order to minimize correlation across draws. A Markov-switching Bayesian VAR can have a very non-Gaussian likelihood surface, with multiple peaks and ridge lines. To ensure that our solutions are robust, we explored the parameter space by doing random global perturbations first with relatively larger perturbations, and then, once the neighborhood of the posterior mode is found, with smaller perturbations. When those perturbations direct the algorithm to a different region, the process is continued until convergence is achieved. This can be thought of as randomizing over the initial conditions from which the block-wise computation of the posterior mode is done. Computation of a specification's posterior mode and the marginal data density took a minimum of 6 hours in clock time and can take as long as 8 days, depending on the specifics of the run. Adding lags, imposing restrictions on switching on variances and restricting switching in equation coefficients is costly in terms of computing times.

Appendix B

This appendix contains more information on estimates of the high-stress coefficient state with alternative indices or constructions of financial stress. It also covers results using selected alternatives to aggregate real growth in personal consumption expenditures (PCE) that was used in the base case. In some instances, we merely repeat the material in the main text but add a chart that is referred to but not included; in other instances, new material is added. In what follows, we compare our alternative results to the base case from the main text, defined as the $3v2c$ specification of the model using growth in real PCE as the real activity variable and the FSI as the measure of financial stress.

B.1 Aggregation of the FSI

The main text of this article noted that the construction of the FSI, with its averaging of the nine components of the FSI, weighted as a function of the inverse of sample standard deviations, is not critical to our results. Figure B.1 below demonstrates this point. In this figure, like several that follow, we show the (smoothed) probability of the high-stress coefficient state—in our view, the most consequential part of our analysis—for the base case $3v2c$ specification, in black. We compare this against, in this case, the state probabilities estimated from the same index constructed as the first principal component on the nine constituent pieces of the aggregate index, the lighter green line. The two lines are vertically offset, and double scaled, for ease of comparison of the dates at which probabilities climb or descend. As can be seen, the estimated switching dates of the 1st PC and the base case are very similar.

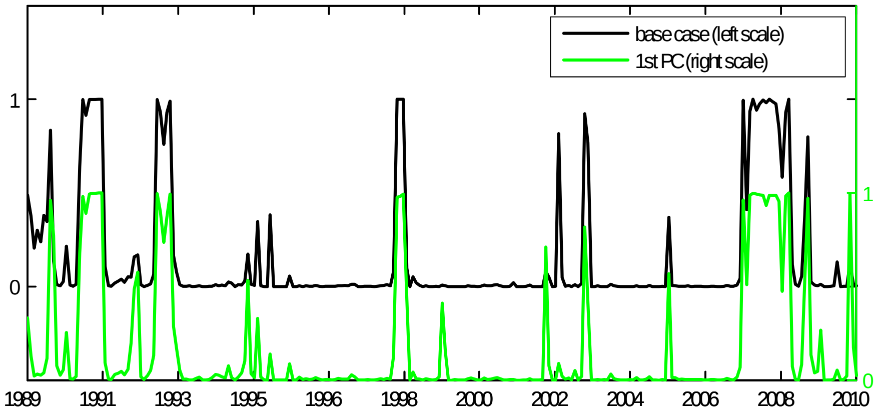

Figure B.1 : Probability of high-stress coefficient state, 1st principal component of nine FSI consituent units (green, right scale) versus base case FSI construction (black, left scale)

B.2 Risky spreads

A combination of the spread of the AA rate and the BBB rate over the 10-year Treasury note rate. It makes no difference how these rates are combined. Figure B.2 below shows the data.

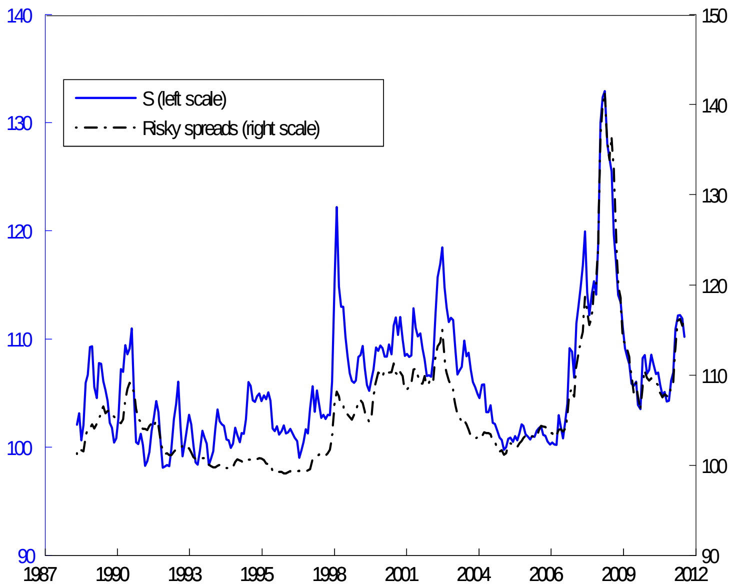

Figure B.2 : Index of risky spreads (*sprd*), black dot-dashed line, right scale, and the FSI (S), blue solid line, left scale, 1988:12-2011:12

The performance of the FSI excluding the *sprd*, measured as always in terms of the high-state switching probabilities, compared with the base case, is shown in Figure B.3, while the case where *sprd* substitutes for the financial stress index, is shown in Figure B.4. As can be seen, the omission of risky spreads harms the performance of the model in some ways, but it still picks up some critical episodes in financial and economic history, albeit tentatively in the case of the 2008-9 period. The replacement of the FSI by *sprd*, on the other hand, leads to a substantial deterioration in performance. Evidently, risky spreads are an important part of the story of financial stress and its transmission, but not a dominant part.

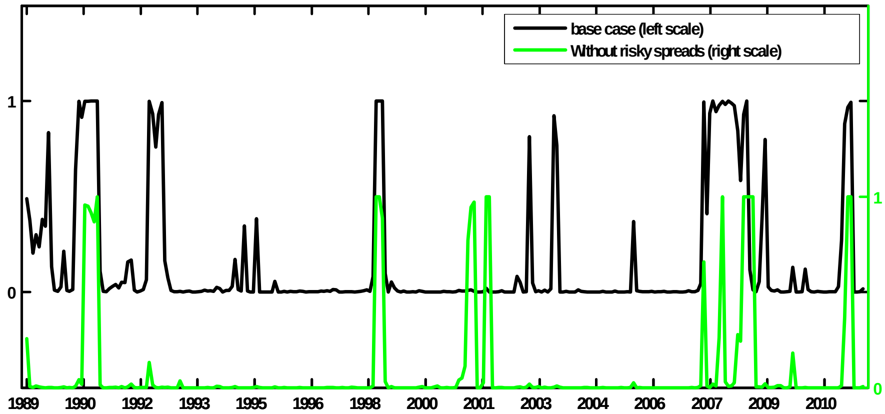

Figure B.3 : Probability of high-stress coefficient state, FSI excluding *sprd* (green, right scale), versus base case (black, left scale)

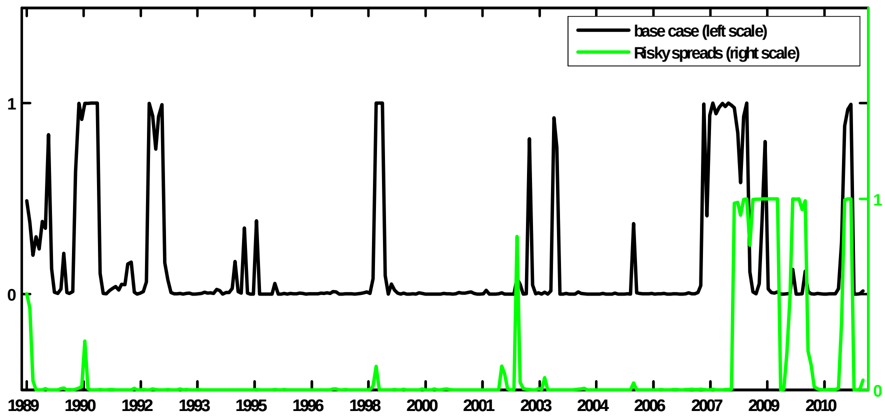

Figure B.4 : Probability of high-stress coefficient state, risky spread (*sprd*) as stress measure (green, right scale) versus base case (black, left scale).

B.3 The equity premium

The stock market is conventionally thought of as a bellwether for all manner of financial and economic activity. It seem relevant, therefore, to consider whether stock market pricing that is out of line with risk-free bond rates is a critical variable for measuring financial stress. Figure B.5 below shows the probability of the high-stress coefficient state when the equity premium is excluded from the measure of the FSI. The figure shows that the equity premium is not particularly important.

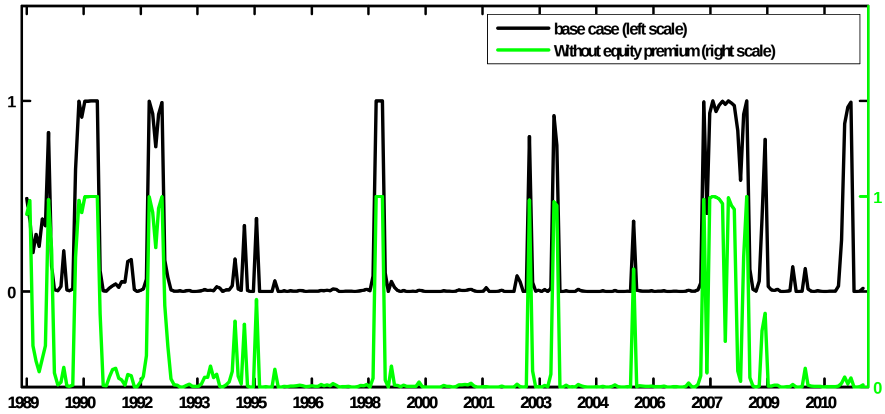

Figure B.5 : High-stress coefficient state probability, FSI excluding *eqprem* (in green, right scale), versus base case (in black, left scale)

B.4 Analysis of the contribution of the FSI

To examine the contribution of the FSI to the results for the system as whole—whether it is the only thing that matters or whether it matters at all—we conduct reestimations of two classes of experiments. In one class, we remove variables from the system. Ideally, we would reduce the system to the FSI alone, however for technical reasons it is difficult to do this. As a very close substitute, we reduce the system to the FSI and the variable that we have concluded is the least consequential to the dynamics, namely price inflation. We show the estimated high-stress coefficient state probabilities for two cases with this specification. The first is for a model that allows switching in coefficients only–not in shock variances—that is, a $1v2c$ specification. The main text established that in the full system, the data are well described by regime switching. Sims and Zha (2006) note that not allowing for switching in shock variances can lead to the erroneous conclusion of switching in coefficients; in other words, it can bias results in favor of coefficient switching.

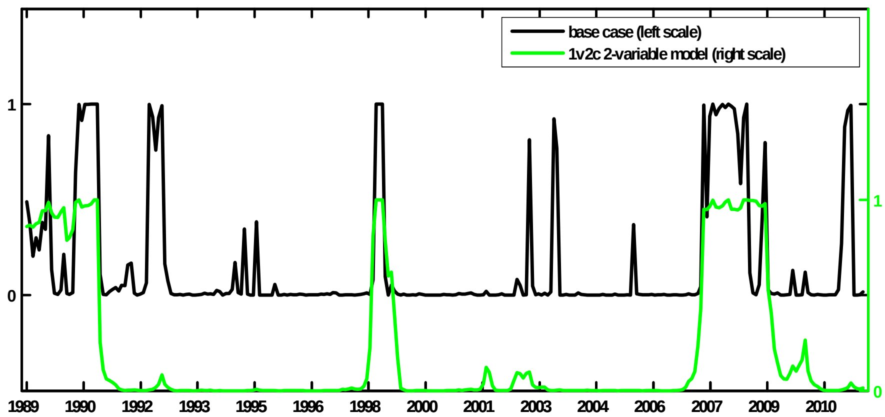

Figure B.6 : Probability of high-stress coefficient state, $1v2c$ 2-variable model with prices (green, right scale) versus base case (black, left scale)

Figure B.6 shows that even after accepting this possible bias as a design feature of the experiment, the reduced dimension model misses many high-stress coefficient states that the base case model picks up. In particular, it misses the key 2010 high-stress episode during the European public debt crisis. Figure B.7 considers the same two-variable model but allowing for the same three states for shock variances and two for coefficients that we use for the base case model. As can be seen, when switching in shock variances is permitted in this way episodes of high-stress coefficients are almost obliterated.

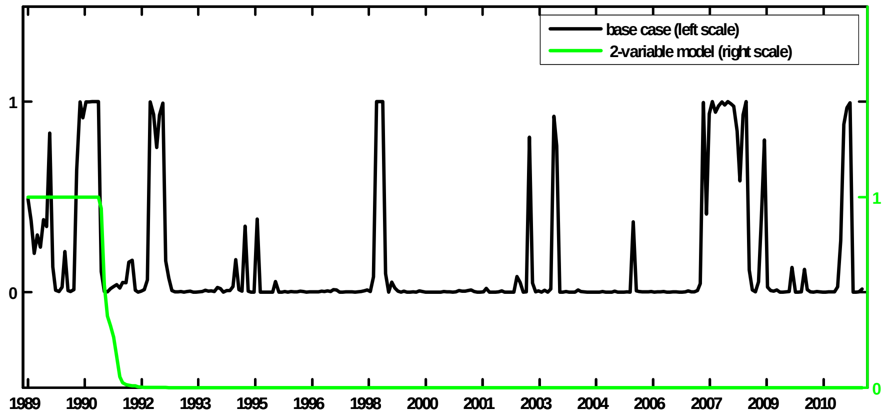

Figure B.7 : Probability of high-stress coefficient state, $3v2c$ 2-variable model with prices (green, right scale) versus base case (black, left scale)

The second class of experiments exploring the role of the FSI to system dynamics is the complement to the case described immediately above: it preserves all the macro variables in the base-case MS-VAR except for the FSI which is removed. The results for this exercise are shown in Figure B.8. In this case, we see that the system without the FSI does pick the 2008-9 financial crisis, but that is about all it picks up. To us this merely suggests that the crisis was severe enough that the omitted variable is picked up by other variables in this circumstances. That this version of the model fails to pick up on other episodes of known importance but less gravity makes this model unsatisfactory, in our view.

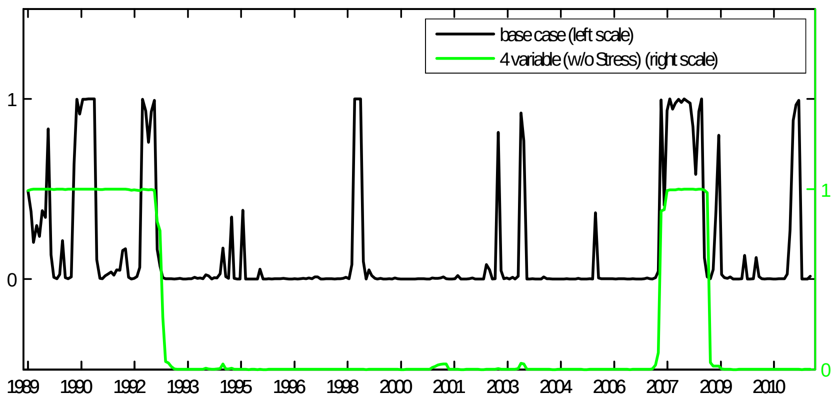

Figure B.8 : High-stress coefficient state probability, 4-variable model, without S (in green, right scale) versus base case (in black, left scale)

B.5 Investigating the real variable in the system

How financial stress affects the real economy might depend on which real variable one considers. To the extent that results differ depending on the real variable could reveal information on what channel is at work in the propagation mechanism. For example, if the effects of financial stress were stronger, in some sense, for industrial production than for base case using aggregate PCE, one might conjecture that this is because the role of working capital in facilitating production is important, as opposed to, say, something to do with consumer credit as a source of funds or as a means by which consumers can substitute intertemporally. Or if business fixed investment was more empirically persuasive as a measure of real activity in the model it might suggest something to do with the availability of credit to firms, or the costs and terms of credit, as an important channel.

In this subsection, we investigate alternative measures of our real variable focussing on two commonly articulated mechanisms by which financial shocks are sometimes thought to be transmitted. The first story builds around the observed volatility of expenditures on (or production of) durable goods to ask whether the complementarity of credit and durables is a major source of the propagation and magnification of shocks. The second is broader, and concerns what could be a lower frequency (and thus possibly less switch-like) mechanism, namely the transmission of shocks via labor markets. In this story, it is less financial disruption that is at work, and more either mismatch in labor market, as could be the case in the 2008-9 recession in the United States given the concentration of the shock in the construction and financial industries, or induced changes in savings behavior in the form of household "balance sheet restructuring" that somehow manifests in extended periods of unemployment.[21] To investigate, however imperfectly, these stories, we re-estimated the model substituting various measures of durable goods on the one hand, and labor market variables on the other, as our real variable. As before, to assess these alternative specifications, we compare the high-coefficient state probabilities—analogous to Figure 3 in the main text—with those of our base case; however, given the subject matter of this investigation, as a compact demonstration of model properties, in a few cases we also show conditional forecasts of the model—analogous to Figure 11. In the construction of conditional forecasts in this section, as in the main text, we simply fix the latent Markov states as appropriate, and simulate out of sample, beginning in January of 2012, without shocks, holding all else constant.

B.5.1 Durables goods

We looked at durables in three different aspects of the macroeconomy. First, we examined durables in consumption, by splitting our base-case real variable, real PCE, into PCE on durable goods, and PCE on services and nondurable goods including footware.[22] Second, we examined durables in factor inputs, by using growth in investment in equipment and intangibles (hereinafter, simply investment).[23] And third we explored durables in production, in the form of growth in total industrial production. We consider these three cases in order.

Figure B.9 reprises part of the information contained in Figure 6 in the main text, showing that PCE durables produces many of the same coefficient switches as the base. In particular, it captures the 2008-9 crisis—albeit haltingly given that it retraces its climb for a time in 2008—and the 2011 euro area sovereign debt crisis. At the same time, the model with PCE durables misses some earlier episodes and misinterprets a period at the turn of the century as a period of high stress. Figure B.10 shows forecasts of the model conditional on the coefficient state—high stress or low. As can be seen, the forecasts are entirely conventional and quite similar to Figure 11 in the main text. Forecasts conditional on the high-stress coefficient latent Markov

[21] The latter story is a bit tricky in that an autonomous increase in private savings need not cause the labor market to fail to clear. Completing the circle on that story requires something more, such as workers not recognizing that the market clearing wage has declined and thus electing to tolerate longer spells of unemployment than otherwise instead of bidding down the real wage.

[22] Created by chain-weighting the nominal series with the appropriate price indexes by the authors. Details are available on request.

[23] In 2013, the BEA substituted equipment & intangibles for equipment & software as that part of business fixed investment that excludes investment in nonresidential structures.

state render higher levels of stress itself, and markedly lower (negative) growth in PCE durables expenditures ($\Delta PCEdur$), despite substantially easier monetary policy (R). The only outcome that is materially different from the base case is the inflation response (ΔP) which, in the base case showed uniformly higher inflation under the high-stress coefficient latent Markov state than under the low-stress state, whereas the pattern is less straightforward here. We noted in the main text that the inflation response shown in Figure 11 supported the interpretation that a switch to the high-stress state is akin to a negative productivity shock; the effect is more equivocal in the case of PCE durables.[24]

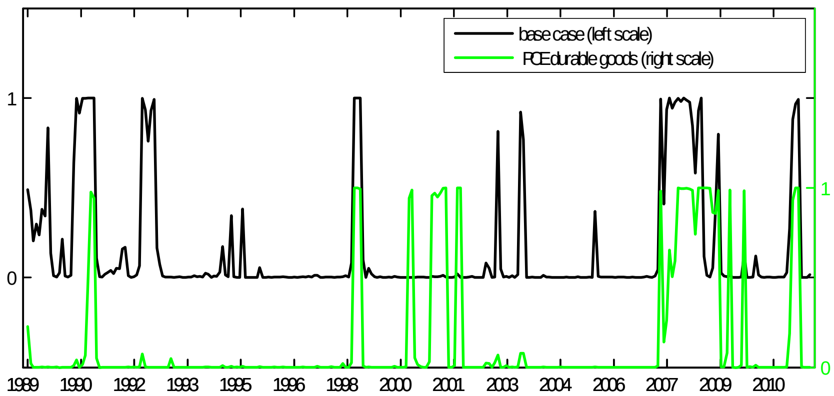

Figure B.9 : Probability of high-stress coefficient state, growth in expenditures on PCE durable goods ($PCEdur$) (green, right scale) versus base case (black, left scale)

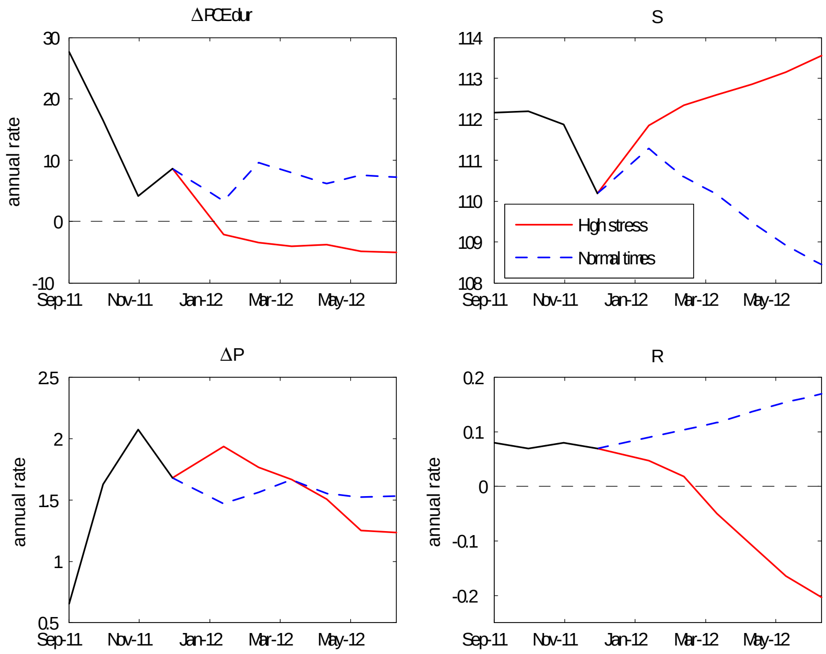

Figure B.10 : Model forecast, conditional on the state, consumer expenditures on durable goods ($pcedur$), from 2011:12. High-stress coefficient state (red solid lines) versus low-stress coefficient state (blue dashed lines)

[24] As in Figure 11 in the main text, for this figure and others we omit the response of ΔM in order to keep the chart compact.

Figure B.11 shows the high-stress coefficient state probabilities for PCE nondurables. Results here are much worse than for PCE durables, missing most of the financial crisis, among other problems. Figure B.12 shows the conditional forecast for this model. Despite very different switching dates in history, the out-of-sample forecasts for PCE nondurables are very similar to the base case. To summarize, consumer durables appear to be more important than nondurables for picking up switching behavior, but for a full characterization it appears that it is helpful to have both.

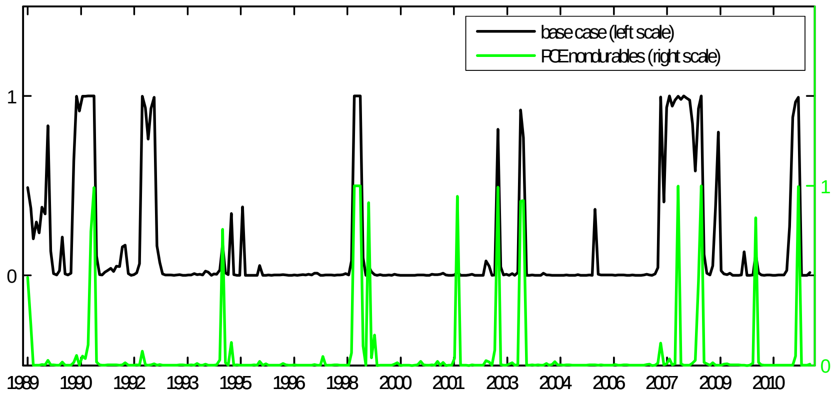

Figure B.11 : Probability of high-stress coefficient state, growth in expenditures on PCE nondurable goods including footware plus services ($PCEnondur$) (green, right scale) versus base case (black, left scale)

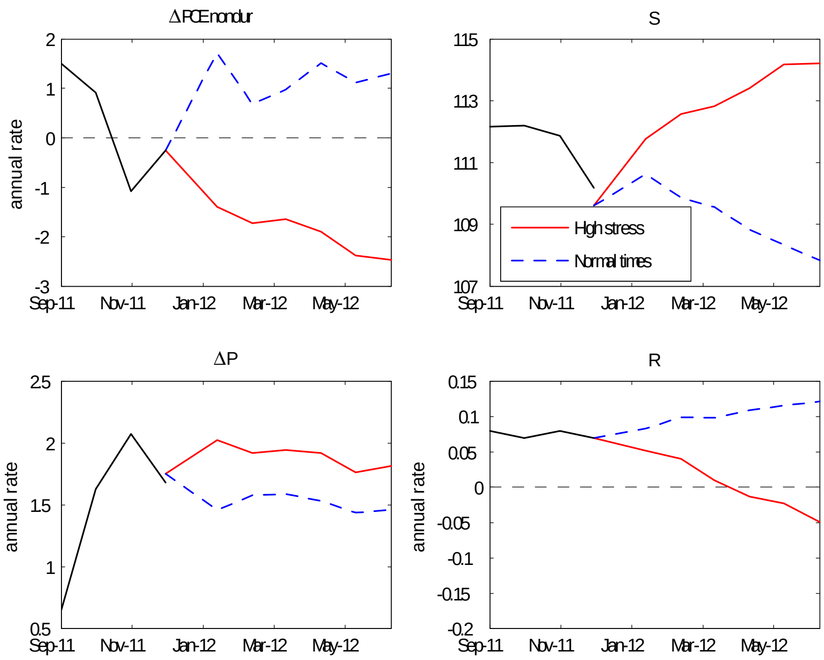

Figure B.12 : Model forecast, conditional on the state, consumer expenditures on services and nondurable goods including footware ($pcenondur$), from 2011:12. High-stress coefficient state (red solid lines) versus low-stress coefficient state (blue dashed lines)

Next we turn to business investment.[25] Figure B.13 shows results for investment that are

[25] The data for investment, or more specifically investment in equipment and intangibles, were taken from a

quite similar to those for PCE durables. Investment picks up the Great Recession, at least in part, the European sovereign debt crisis in 2010-11, as well as the 1998 Russian debt default and the 2002 Argentine sovereign debt crisis; like PCE durables it attributes a high-stress coefficient state to the bursting of the high-tech bubble in 2000-01, that the base case model does not see.

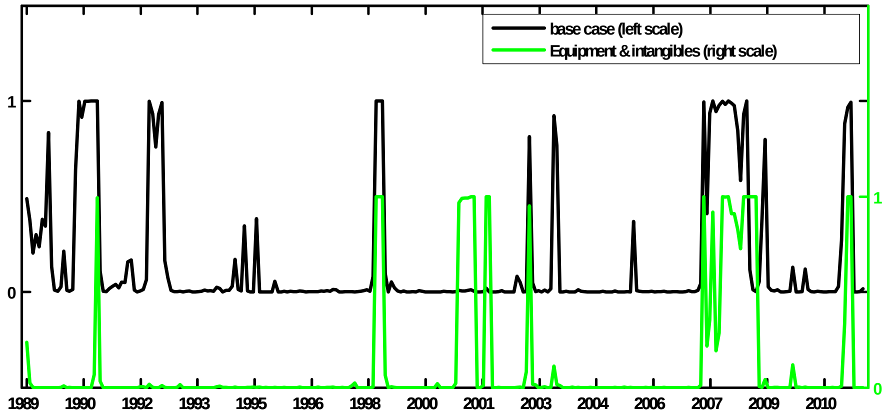

Figure B.13 : Probability of high-stress coefficient state, growth in investment in equipment & intangibles (*eandi*), (green, right scale) versus base case (black, left scale)

In the case of industrial production, Figure B.14 shows that IP does a reasonable job picking out many periods of high-financial stress particularly in the early part of the sample, but misses by a wide margin the onset of the financial crisis. The conditional forecast is much like the others shown in this section, except that it fails to produce much of a monetary policy response (R), a reflection, perhaps, of the small share of the overall economy represented by industrial production as well as its nonrepresentativeness. We conclude that the nonlinearities captured by our base case model are not well represented by industrial production and therefore that the mechanisms that are germane to that sector, such as the availability of working capital do not seem to play an outsized role.

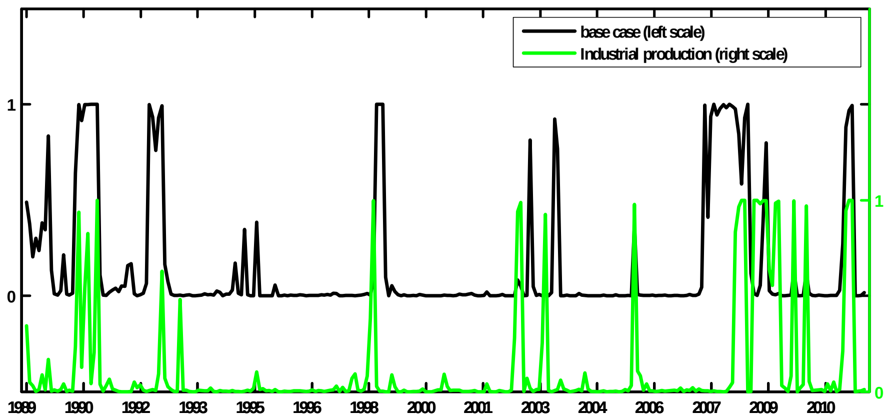

Figure B.14 : Probability of high-stress coefficient state, growth in industrial production (*ip*), (green, right scale) versus base case (black, left scale)

Taken together, these results for consumer durables, investment and industrial production suggest that a good part of the effects of financial stress likely operate through credit conditions or credit availability—nonprice terms, more generally—and their effects on expenditures

recent vintage of the National Income and Produce Accounts and interpolated to monthly frequency using the Chow-Lin (1971)procedure using data for new orders of non-defense capital goods for identification. Details are available from the authors.

on durable goods. For example, it is sometimes argued that disproportionate effects of shocks on durable goods occurs because of irreversibility of investment, irrespective of switching phenomena. However, while irreversibility could be expected to produce large movements in macroeconomic aggregates in response to negative shocks, there would, however, be no expectation that such shocks would produce Markov switching as an empirical phenomenon, unlike the case where credit availability is impinged. One mechanism that is consistent with our results through which credit conditions might operate is through collateral constraints where declines in the value of pledgeable assets would affect firms' ability to finance new capital investment, or households' ability to purchase big ticket durable goods. In this regard, our results are consistent with the empirical observations of Chaney *et al.* (2012) and the theoretical constructs of Liu *et al.* (2013), among other contributions.

B.5.2 Labor markets

Finally, we also re-estimated the model using three labor market variables as the real variable in our system, the unemployment rate, initial claims for unemployment insurance on state programs, and growth in private nonfarm payroll employment.[26] The results were uniformly inferior to the expenditure-based real variables discussed above. Figure B.15 presents coefficient state probabilities for payroll employment, arguably the best of the three models. As can be seen, results for payroll employment are broadly similar to those for investment. At the same time, the conditional forecasts, shown in Figure B.16, are conventional.

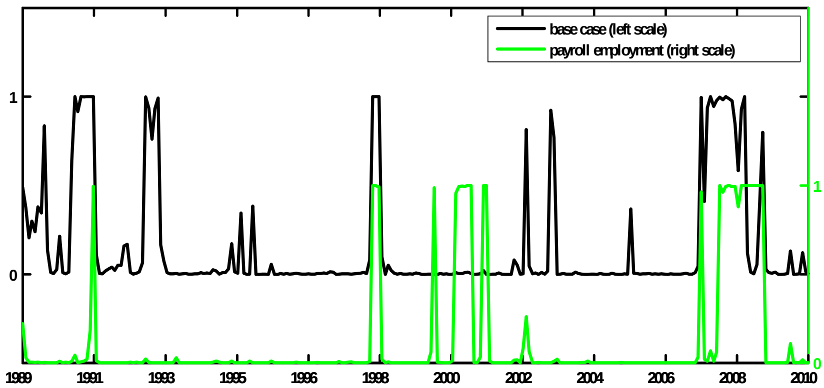

Figure B.15 : Probability of high-stress coefficient state, growth in payroll employment (*pemp*), (green, right scale) versus base case (black, left scale)

At one level, the results for labor market variables are not surprising; there are fewer of the high-frequency discretely-shifting mechanisms in play for labor markets than there are for durable goods. Nevertheless, the fact that labor market variables perform relatively poorly compared to, say, PCE durables in our model does not mean that labor markets are immaterial to the transmission of crises. Rather, it suggests to us (but does not prove) that labor markets are not likely to be an independent source of the switching phenomena studied here, at least at the monthly frequency we use. Taking the 2008-9 Great Recession as our example, it is certainly possible—indeed plausible—that the fragility of household financial conditions—a condition that built up over many years in the U.S. economy owing in part to "jobless recoveries" from the 1991 and 2000 recessions and the large share of residential real estate on both sides of

[26] The unemployment rate and nonfarm payroll employment are both monthly. Initial claims on state unemployment insurance programs is highly scrutinized by financial market participants because they provide a reasonably good real-time indication of the state of the labor market and are available weekly. We collapsed the weekly series to monthly by averaging.

the typical household balance sheet—set the stage for severity of the recession by obligating a period of household deleveraging following the collapse of the housing market. It would certainly be a worthwhile endeavour to construct the datasets that are long enough and detailed enough to investigate the role of these lower frequency, cumulative processes.

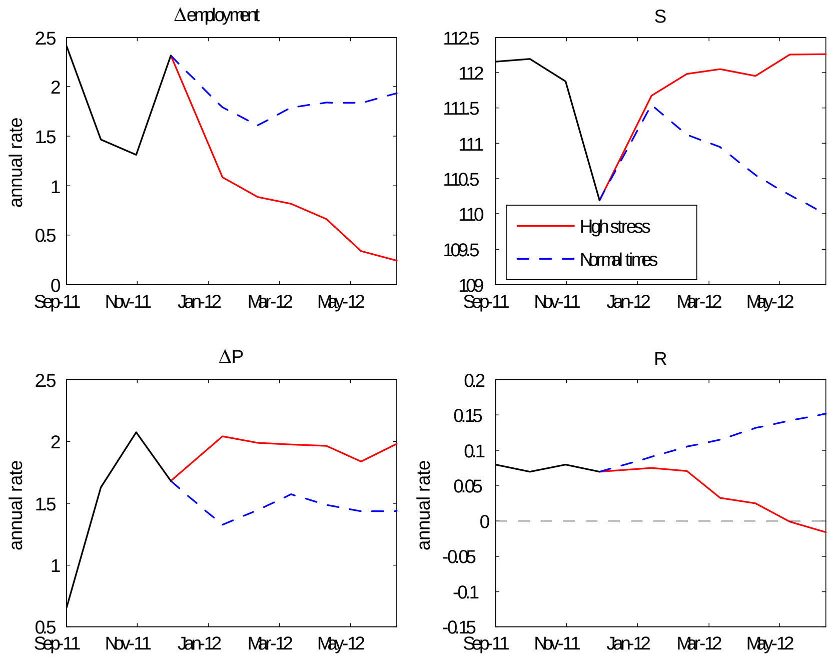

Figure B.16 : Model forecast, conditional on the state, payroll employment (*pemp*), from 2011:12. High-stress coefficient state (red solid lines) versus low-stress coefficient state (blue dashed lines)

References

[1] Campbell, J.Y.; Giglio, S., Polk, C. and Turley, R. 2013. An intertemporal CAPM with stochastic volatility. unpublished manuscript, Department of Economics, Harvard University.

[2] Chow, G.C., Lin, A. 1971. Best linear unbiased interpolation, distribution, and extrapolation of time series by related series. Review of Economics and Statistics 53, 372-375.