BANK OF ENGLAND

Staff Working Paper No. 616
The macroeconomic shock with the highest price of risk

Gabor Pinter

September 2016

BANK OF ENGLAND

Staff Working Paper No. 616
The macroeconomic shock with the highest price of risk

Gabor Pinter[(1)]

Abstract

I propose a new method of constructing a macroeconomic shock based on its ability to explain the cross-section of asset returns. The only identifying assumption is that this λ-shock demands the highest risk price per unit of exposure, or equivalently, minimises the associated sum of squared pricing errors, when pricing a given asset portfolio. When applying the method to the stock portfolios studied by Fama-French, a robust economic feature of the λ-shock is the delayed effect on aggregate quantities such as output and consumption and a sharp impact on the short-term interest rate and the term spread. The estimated λ-shock bears strong resemblance both with monetary policy shocks and with technology news shocks studied by the macroeconomic literature.

Key words: Stock returns, VAR, identification, shocks, technology news, monetary policy.

JEL classification: C32, G12.

(1) Bank of England. Email: gabor.pinter@bankofengland.co.uk

The views expressed in this paper are those of the author, and not necessarily those of the Bank of England or its committees. I am grateful to Saleem Bahaj, Andy Blake, Gino Cenedese, John Cochrane, Wouter den Haan, Clodo Ferreira, Rodrigo Guimaraes, Campbell Harvey, Ravi Jagannathan, Ralph Koijen, Peter Kondor, Sydney Ludvigson, Stefan Nagel, Morten Ravn, Ricardo Reis, Andrea Tamoni, Harald Uhlig, Garry Young and Shengxing Zhang for helpful comments.

Information on the Bank's working paper series can be found at
www.bankofengland.co.uk/research/Pages/workingpapers/default.aspx

Publications Team, Bank of England, Threadneedle Street, London, EC2R 8AH
Telephone +44 (0)20 7601 4030 Fax +44 (0)20 7601 3298 email publications@bankofengland.co.uk

ISSN 1749-9135 (on-line)

1 Introduction

> "We would like to understand the real, macroeconomic, aggregate, nondiversifiable risk that is proxied by the returns of the HML [High-minus-Low] and SMB [Small-minus-Big] portfolios." (pp. 442 Cochrane (2005))

The literature is yet to find a compelling macroeconomic explanation behind the cross-sectional variation of stock returns. I argue that part of this challenge has been caused by the fact that innovations in macrovariables are reduced-form objects: they are linear combinations of orthogonal structural shocks that can offset each other over the business cycle and demand possibly very different levels of risk premia. Using reduced-form variables such as unexpected changes in output or inflation, as often done in the empirical asset pricing literature, can therefore pose an insurmountable challenge to estimate risk exposures and risk prices associated with structural macroeconomic forces.

My paper aims to solve this problem by proposing a macroeconometric identification strategy in a simple vector autoregression (VAR) model: instead of starting with macroeconomic assumptions and testing their asset pricing implications, I start by using the returns of a given asset portfolio to construct an orthogonal shock that has the highest risk premium in absolute value, or equivalently the best cross-sectional fit, when pricing the given portfolio. Only then I check the macroeconomic characteristics of the resulting shock by inspecting the associated impulse response functions and the estimated time-series of the shock. While performing this strategy, I condition on the information set contained in the VAR, therefore my strategy is not data mining even though it does have a reverse engineering nature. When applying the method to the 25 portfolios of Fama and French (1993) sorted on book-to-market and size (FF25 henceforth), I find that the obtained shock closely resembles well-known structural shocks, traditionally studied by the macroeconomic literature, while explaining as much of the average excess returns of the FF25 portfolios as the 3-factor model can.

This shock, which I will refer to as a λ-shock, triggers a delayed reaction in aggregate quantities such as GDP and consumption which is consistent with recent empirical papers on consumption based asset pricing since Bansal and Yaron (2004) and Parker and Julliard (2005). Moreover, the shock has an immediate impact on the short-term interest rate and the term spread. These features make the λ-shock bear a strong similarity with what the macroeconomic literature refers to as a news shock about future total factor productivity (TFP), meanwhile the statistical properties of the shock are also similar to monetary policy shocks. In fact, the correlation between the λ-shock I identify and the *TFP news shock* series, estimated by Kurmann and Otrok (2013), and the *monetary policy shock* series, constructed by Romer and Romer (2004), are more than 70%. This is quite striking given that my identification strategy, as explained further below, has nothing to do with the strategies used to identify monetary policy shocks or TFP news shocks, as my model does not even contain a measure of TFP as an observable.

The starting point of my analysis is a standard vector autoregression (VAR) model of a small set of macroeconomic variables. The finance literature often applied Cholesky decomposition to the estimated reduced-form variance covariance matrix of similar VAR models to obtain triangularised innovations in the spirit of the Intertemporal CAPM (Merton, 1973).[1] While the obtained innovations had success in explaining the returns on the FF25 portfolios, it has been difficult to assign macroeconomic interpretations to these innovations. Moreover, triangularisation is merely one of the infinite number of identification strategies to transform the reduced-form variance-covariance matrix to a structural form. I build on this last point by exploring the entire space of possible orthogonalisations, given the estimated time-series of reduced-form residuals.

I make only one assumption in my identification scheme: I propose to directly look for a single structural shock that demands the highest possible level of risk premium in absolute value when pricing the cross-section of stock returns. Intuitively, the λ-shock captures the type of aggregate, non-diversfiable risks that are most 'feared' by market investors, as they expect the highest risk compensation from holding assets whose returns are exposed to this shock. Mechanically, the λ-shock is identified as the one that, if used as a factor in the two-pass procedure of Fama and MacBeth (1973) applied to the FF25 portfolios, would generate the highest estimated factor risk premium in absolute value. I will also show that (i) this identification strategy is equivalent to searching for the structural shock whose corresponding 1-factor model has the lowest possible sum of squared pricing errors, and (ii) the identification of the λ-shock is robust to changing the test portfolios by augmenting the FF25 with the 30 Industry portfolios, thereby addressing the critique of Lewellen, Nagel, and Shanken (2010).

Nothing in my approach makes any of the assumptions that macroeconometricians tend to make when identifying structural shocks, e.g. restrictions regarding the short/long-run effects of the shock, or regarding the shock's contribution to the forecast error variance of a target variable in the VAR over a pre-specified horizon.[2] Compared to these approaches, my method can be thought of as much more agnostic. My only identifying assumption is to construct a macroeconomic shock which demands the highest risk price per unit of exposure according to the FF25 portfolios. Hence, there is absolutely no a priori reason to believe that the obtained structural λ-shock captures any of the economic forces studied by the structural VAR literature. The fact that it does, by closely resembling the statistical features of well-known macroeconomic shocks, could provide strong evidence on the relevance of those shocks in not only driving business cycles but also in explaining the cross-section of stock returns.

[1] See Campbell (1996); Petkova (2006); Maio and Santa-Clara (2012); Boons (2016) amongst others.

[2] The latter type of restriction has been increasingly popular (since its development by Uhlig (2004)), particularly in the context of the identification of news shocks.

My paper relates to two strands of literature. First, it builds on the finance literature that aimed at finding macroeconomic factors that drive the cross-sectional variation of risk premia. A partial list includes Chen, Roll, and Ross (1986), Ferson and Harvey (1991), Campbell (1996), Cochrane (1996), Vassalou (2003), Brennan, Wang, and Xia (2004), Petkova (2006), Liu and Zhang (2008), Maio and Santa-Clara (2012), Koijen, Lustig, and van Nieuwerburgh (2012), Kan, Robotti, and Shanken (2013), Boons and Tamoni (2015), He, Kelly, and Manela (2016).[3] To the best of my knowledge, only few of these papers sought orthogonalised, economically meaningful shocks rather than reduced-form innovations or macroeconomic variables themselves to price the cross-section of returns. This is somewhat in contrast with the empirical macroeconomic literature which has showed that business cycle fluctuations are caused by the simultaneous realisations of various structural disturbances with potentially very different quantities and prices of risk (Smets and Wouters (2007); Justiniano, Primiceri, and Tambalotti (2010); Rudebusch and Swanson (2012); Borovicka and Hansen (2014); Kliem and Uhlig (2016)). An implication of my results is that identification is key to understanding the macroeconomic forces behind cross-sectional variation in stock returns.

Second, my paper relates to the vast literature on macroeconomic shocks on business cycle dynamics. A partial list regarding news shocks includes Beaudry and Portier (2006, 2014), Jaimovich and Rebelo (2009), Barsky and Sims (2011), Schmitt-Grohe and Uribe (2012), Kurmann and Otrok (2013), Barsky, Basu, and Lee (2014), Christiano, Motto, and Rostagno (2014), Malkhozov and Tamoni (2015). While news shocks have been found to be important in explaining business cycles, yet, to the best of my knowledge, their role in explaining the cross-section of stock returns has been unexplored. This is particularly interesting given that actually some of the early work such as Beaudry and Portier (2006) used information in aggregate stock price movements (together with observed TFP measures and certain short-run and long-run restrictions as identification schemes) to identify news shocks. A partial list regarding monetary policy shocks includes Sims (1980), Thorbecke (1997), Christiano, Eichenbaum, and Evans (1999), Romer and Romer (2004), Coibion (2012) and Gertler and Karadi (2015) amongst many others.

The remainder of the paper is as follows: Section 2 explains my empirical approach, Section 3 presents the empirical results and Section 4 concludes.

[3]In addition, consumption based asset pricing (CCAPM) models also had success in explaining the returns on the Fama-French portfolios either by (i) introducing conditioning variables (Jagannathan and Wang (1996); Lettau and Ludvigson (2001); Lustig and Nieuwerburgh (2005); Santos and Veronesi (2006); Yogo (2006); Jagannathan and Wang (2007)), (ii) focusing on the long-run component of consumption risk (Bansal and Yaron (2004); Parker and Julliard (2005); Hansen, Heaton, and Li (2008); Constantinides and Ghosh (2011)) or (iii) focusing on non-durable consumption risk (Piazzesi, Schneider, and Tuzel (2007)). The empirical performance of these various approaches were subsequently criticised by Lewellen, Nagel, and Shanken (2010), Beeler and Campbell (2012) amongst others. Nevertheless, some macroeconomists may argue that the CCAPM is still silent about what underlying structural shocks drive movements in consumption that are priced in the cross-section.

2 VAR and Identification

The starting point of my empirical approach follows the macro-finance literature (Campbell and Shiller (1988); Campbell and Vuolteenaho (2004); Campbell, Giglio, and Polk (2013)) by using a first-order reduced-form VAR to model the evolution of the macroeconomic state:

$$y_t = c + By_{t-1} + u_t, \qquad t = 1, \ldots, T, \tag{2.1}$$

where y_t is an $n \times 1$ vector of observed endogenous variables, c is an $n \times 1$ vector of constants, B is an $n \times n$ matrix of coefficients and u_t is a $T \times n$ matrix of reduced-form residuals with a variance-covariance matrix Σ. Given that the estimated $\hat{\Sigma}$ is positive definite, there exists a non-unique decomposition $A_0 A_0' = \hat{\Sigma}$ such that the relationship between the reduced-form and structural errors can be written as $u_t = A_0 \varepsilon_t$, where ε_t is a $T \times n$ matrix of structural errors and A_0 is an $n \times n$ structural impact matrix to be determined. To find A_0, I first apply Cholesky decomposition to the estimated reduced-form variance-covariance matrix $\hat{\Sigma} = \tilde{A}_0' \tilde{A}_0$. It is known that one can take any orthonormal matrix Q to obtain a new structural impact matrix $A_0 = Q\tilde{A}_0$, thereby obtaining a new set of structural shocks, which is still consistent with the reduced-form variance covariance matrix, i.e. $\hat{\Sigma} = \left(Q\tilde{A}_0\right)' Q\tilde{A}_0$.[4]

The main assumption of my identification is as follows. I select the matrix $Q^\star$ from the space of all Q matrices such that the implied ε_t matrix of structural shocks contains one $T \times 1$ vector of shocks $\varepsilon_t^\star$ with the following property: if it were to be used as a factor to price the cross-section of FF25 portfolios, it would command the largest possible risk premium from the set of all possible structural shocks, consistent with $\hat{\Sigma}$, i.e. $A_0 = Q^\star \tilde{A}_0$. To put it formally, denote the $T \times k$ matrix of portfolio excess returns, R_t^e and write the beta representation as (Chapter 9 of Cochrane (2005)):

$$\mathbb{E}\left(R_t^e\right) = \beta\left(\varepsilon_t^\star\right) \times \lambda\left(\varepsilon_t^\star\right), \tag{2.2}$$

where $\beta\left(\varepsilon_t^\star\right)$ is a $k \times 1$ vector of factor betas, and $\lambda\left(\varepsilon_t^\star\right)$ is the associated factor risk premium. The notation aims to emphasise that both the factor betas and the risk premium are naturally functions of the underlying structural λ-shock, $\varepsilon_t^\star$, that I aim to identify. I proceed by searching through the entire space of $n \times n$ orthonormal matrices and estimate the associated candidate λs using the two-stage procedure of Fama and MacBeth (1973). Given a candidate matrix $\tilde{Q}$, the first stage is an OLS estimation of the time-series regres-

[4] This is also the starting point for a range of identification strategies in the macroeconometric literature, e.g. sign restrictions (Uhlig (2005); Rubio-Ramirez, Waggoner, and Zha (2010), see Fry and Pagan (2011) for a survey), identification of news shocks (Barsky and Sims (2011); Pinter, Theodoridis, and Yates (2013); Kurmann and Otrok (2013)), using external instruments as proxies for structural shocks (Mertens and Ravn (2013); Gertler and Karadi (2015)) etc.

sion of each of the k portfolios' excess return on the implied candidate structural shock $\tilde{\varepsilon}_t$:

$$R^e_{it} = a_i + \tilde{\varepsilon}_t \beta_i + \epsilon_{it}, \tag{2.3}$$

where β_i represents the ith element in β. Given 2.3, the second stage is a cross-section regression of average portfolio returns on the estimated betas associated with the candidate matrix $\tilde{Q}$:

$$\bar{R}^e_i = \tilde{\beta}_i \times \lambda + \alpha_i, \tag{2.4}$$

where $\bar{R}^e_i = \frac{1}{T}\sum_{t=1}^{T} R^e_{it}$, and $\tilde{\beta}_i$ is the OLS estimate obtained in the first stage and α_i is a pricing error. To sum up, I will select matrix $Q^\star$ from all $\tilde{Q}$ candidate matrices to generate the time-series $\varepsilon^\star_t$ which will generate the highest estimated λ in absolute value in 2.4. Finding $\varepsilon^\star_t$ is done via the following optimisation routine: I span the space of n-dimensional orthonormal matrices that are rotations with an n-dimensional Givens rotation. I then choose the Euler-angles of the Givens rotation appropriately such that the corresponding second-pass λ is maximised.[5]

It is important to note that while assumptions about identification determines risk exposures and prices of risk, it does not at all affect the overall cross-sectional (R^2-type) fit of the transformed residuals, if all the structural shocks were to be used for pricing the cross-section of returns. After all, the structural shocks are merely different linear combinations of the reduced-form residuals, thereby containing exactly the same information set. However, I will use only the one structural shock (that my proposed strategy identifies based on the *magnitude* of the associated risk premium) when subsequently pricing the cross-section. It is perhaps not obvious how well the identified structural shock should fit the cross-section of returns. Nevertheless, the next Section will show that all the information, contained in the VAR innovations, that is relevant to pricing the cross-section will in fact be captured by the λ-shock that I identify.

3 The Empirical Results

3.1 Data

To operationalise the VAR model described in Section 2, one needs to specify the variables to be included in the state vector. I opt for a parsimonious model with the following five, completely standard state variables: output, aggregate price level, the policy interest rate, the default spread and the term spread. Data on the following four series are from the Federal Reserve Bank of St. Louis (FRED): output is measured as quarterly seasonally adjusted real GDP (FRED code: GDPC1), price level is measured as the personal consumption expenditures (chain-type) price index (FRED code: PCEPI), the

[5] See Fry and Pagan (2011) for further details on Givens rotations in the context of sign restrictions.

policy interest rate is the Federal Funds Rate (code: FEDFUNDS) and the default spread is the difference between the AAA (FRED code: AAA) and BAA (FRED code: BAA) corporate bond yields. The term spread is defined as the difference between the long term yield on government bonds and the T-bill as used in Goyal and Welch (2008).[6] These five variables have long been recognised as good candidates for state variables within the ICAPM framework (Petkova, 2006), and they frequently appear as key variables in macroeconomic forecasting models as well (Stock and Watson, 2003; Ng and Wright, 2013).

When estimating the VAR, I deliberately avoid using financial variables such as aggregate excess returns or various valuation ratios, that are known to increase the overall fit of cross-sectional asset pricing models. The specification of the state vector is motivated by the desire to stay as close as possible to macroeconomic explanations of the cross-section of stock returns, in the spirit of Chen, Roll, and Ross (1986) and subsequent papers.

The sample period for the estimation is 1963Q3-2008Q3 and the data are at quarterly frequency. The start of the estimation period is selected based on the fact that it is used in the majority of empirical asset pricing studies of the cross-section. The end of the estimation period is chosen to exclude the Great Recession period when the Federal Funds Rate hit the zero-lower bound.

As for the FF25 portfolios, they are formed from independent sorts of stocks into five size groups and five B/M groups as described in Fama and French (1992, 1993).[7] The returns are the accumulated monthly returns in excess of the one-month U.S. Treasury bill rate. As studied extensively by the empirical asset pricing literature, average returns typically fall from small stocks to big stocks (size effect), and they rise from portfolios with low to large book-to-market ratios (value effect).

Table 1: Average quarterly percent excess returns for portfolios formed on Size and Book-to-Market; 1963Q3-2008Q3, 181 quarters.

		Book-to-Market				
		Low	2	3	4	High
	Small	0.53	2.23	2.37	2.97	3.31
	2	1.04	1.94	2.63	2.81	2.97
Size	3	1.16	2.05	2.22	2.43	3.02
	4	1.45	1.47	1.91	2.40	2.41
	Large	1.11	1.37	1.21	1.59	1.73

As is well-documented (most recently by Fama and French (2015)), the value effect is stronger among smaller firms. For example, for the microcap portfolios presented in the first row of Table 1, average excess return rises from 0.5% per quarter for the lower B/M portfolio (extreme growth stocks) to more than 3.3% per quarter for the highest

[6] I would like to thank Professor Amit Goyal for updating and sharing his dataset on his website.
[7] I would like to thank Professor Ken French for making the data available on his website.

B/M portfolio (extreme value stocks). In contrast, for the largest stocks (megacaps), average excess returns rise from about 1.1% to only about 1.7%. As shown below, these patterns can be explained by cross-sectional variation in exposures to a macroeconomic shock about future technology that this paper uncovers.

3.2 The Economic Characteristics of the λ-shock

Using the OLS estimates of the VAR, I compute impulse response functions (IRFs) after performing the identification strategy described in Section 2. This is to understand the macroeconomic impact of the λ-shock which is by construction the structural shock demanding the highest possible price of aggregate risk (conditional on the VAR being an accurate representation of the evolution of the macroeconomy) when pricing the cross-section of stock returns. One's reaction to this exercise is that it is somewhat tautologous to study the impact of a structural shock on a set of macrovariables, if the shock itself was constructed from linear combinations of the innovations in the same set of macrovariables with the aim to explain the cross-section of stock returns. However, it is worth reiterating that there is no direct reason to believe that the constructed shock should possess any well-known economic characteristics. The fact that it does is the main finding of this paper, because it confirms that structural shocks with economically meaningful characteristics (and not necessarily the reduced-form innovations) are in fact the relevant factors for cross-sectional asset pricing.

Figure 1: Impulse Responses to a λ-shock

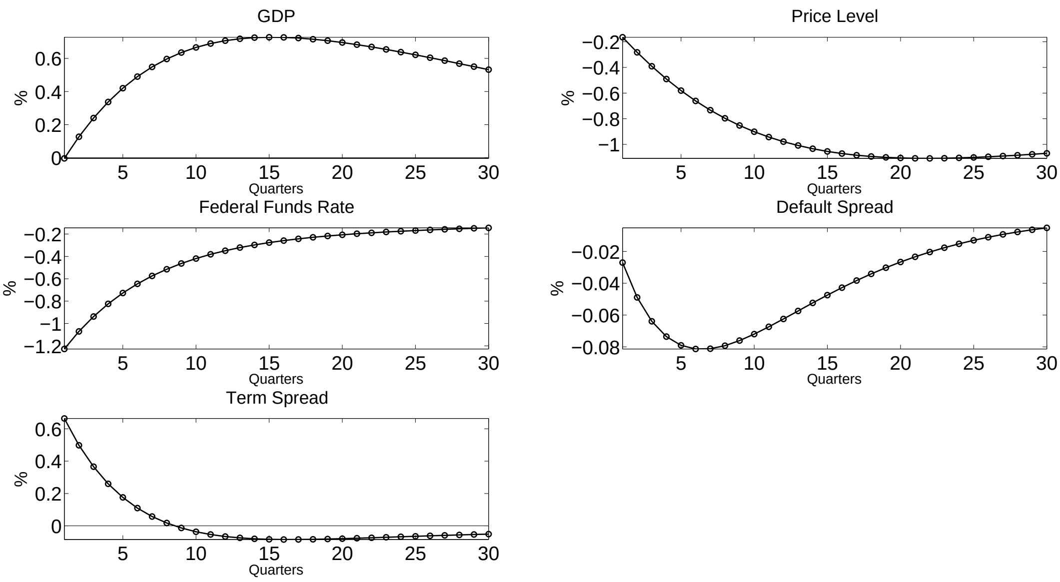

Notes: The vertical axes are %-deviations from steady-state, and the horizontal axes are in quarters.

To understand the causal effect of a λ-shock, Figure 1 displays the IRFs of the five variables to a 1 percent structural innovation. The term spread jumps by about 70 basis points on impact of the shock and there is a very sharp and persistent drop in the Federal Funds Rate. The initial drop in the price level is lower than the drop in the Federal Funds Rate, suggesting a sharp drop in the real interest rate. Interestingly, the λ-shock has virtually no effect on GDP on impact, but the effect increases substantially with the horizon and reaching a peak impact of about 0.7% approximately 12-15 quarters after the shock hits. As shown by Figure 8 in the Appendix, the shape of these impulse response functions is similar when the lag length is increased or when output is replaced by consumption in the VAR.

The delayed response of aggregate quantities in response to innovations that are relevant to asset pricing is a phenomenon that has been documented by the consumption based macro-finance (Parker and Julliard, 2005) and long run risk literatures (Bansal and Yaron, 2004). More recently, Bryzgalova and Julliard (2015) have shown that "slow consumption adjustment shocks" account for about a quarter of the time series variation of aggregate consumption growth, and its innovations explain most of the time series variation of stock returns. My results are consistent with their findings. In addition, my multivariate time-series framework is somewhat richer than their reduced-form consumption growth model, so it can possibly shed further light on the macroeconomic drivers of the slow consumption adjustment shocks that are the main source of aggregate risk.

One possible interpretation of Figure 1 is that the λ-shock behaves like a supply-type shock with aggregate production moving in the opposite direction compared to the price level and the short-term interest rate. However, the delayed expansion of output would make the λ-shock clearly distinct from a positive unanticipated technology shock which would have an immediate positive impact on output and consumption, as traditionally studied by the Real Business Cycle (RBC) and the subsequent New Keynesian literature.[8] However, a news-type technology shock that typically triggers a delayed reaction in aggregate quantities may be perfectly consistent with Figure 1. Indeed, Figure 4 of Kurmann and Otrok (2013) shows results for an identified TFP news shock with very similar IRFs to mine. The striking similarity between my Figure 1 and their findings occurs in spite of the fact that they identify a TFP news shock, following Barsky and Sims (2011), by searching for a shock that accounts for most of the forecast error variance of TFP over a given forecast horizon, and they force this shock to be orthogonal to contemporaneous movements in TFP.

[8] Though technology shocks had some theoretical success in explaining aggregate excess returns in an RBC framework (Jermann, 1998), the most recent empirical evidence by Greenwald, Lettau, and Ludvigson (2015) finds that the contribution of unanticipated TFP shocks to the variance of aggregate stock market wealth is close to zero. These authors identify three mutually orthogonal observable economic disturbances that are associated with over 85% of fluctuations in real quarterly stock market wealth. They find that the third triangularised shock from a cointegrated three-variable VAR (including consumption, labor income, and asset wealth) is the main driver of the variance of aggregate stock market wealth. Their identifying assumption implies zero contemporaneous impact on consumption – an assumption that is consistent with the IRF results implied by the more agnostic identification theme adopted in this paper.

Figure 2: Comparing the λ-shock to the TFP News Shock Series of Kurmann and Otrok (2013) (Correlation: 72%) and to the Monetary Policy Shock Series of Romer and Romer (2004) (Correlation: 73%).

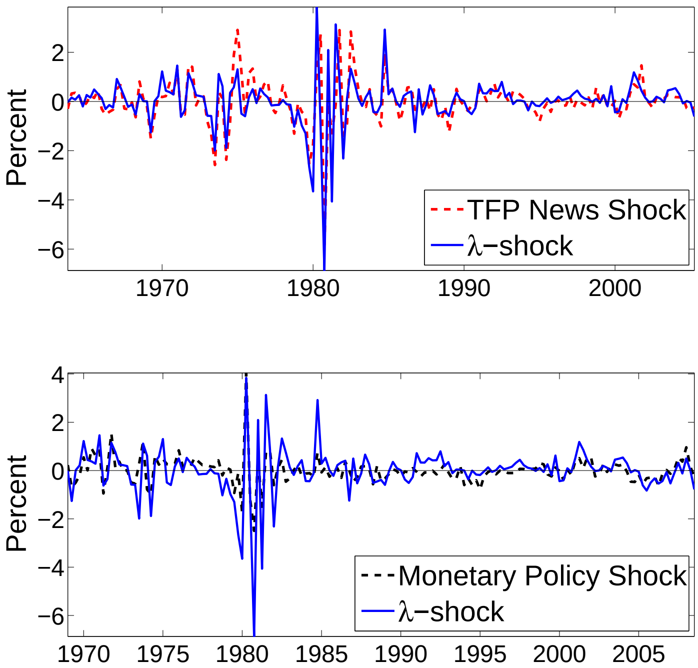

Notes: The TFP news shock series are the ones plotted in Figure 5 on pp. 2625 of Kurmann and Otrok (2013) who apply the method of Uhlig (2004) to identify a TFP news shock over the period 1959Q2-2005Q2. The monetary policy shock series are originally developed by Romer and Romer (2004) and updated by Olivier Coibion to the period 1969Q1-2008Q4.

An alternative interpretation of Figure 1 is that a positive λ-shock behaves like an expansionary monetary policy shock to the extent that it generates an immediate jump in the short-term interest rate and the term spread and a delayed but persistently expansionary reaction in output. Though CPI goes the 'wrong' way, but it is somewhat consistent with the 'price puzzle' (Sims, 1992) associated with early methods of Cholesky orthogonalisation to identify monetary policy shocks as in Christiano, Eichenbaum, and Evans (1999) and others.

To formally show the similarity between the λ-shock that I identify from the cross-section of stock returns and some well-known structural shocks studied by macroeconomists, Figure 2 plots the time-series of the λ-shock against the TFP news shocks identified by Kurmann and Otrok (2013) (upper panel) and against the monetary policy shocks constructed by Romer and Romer (2004) (lower panel). Based on the overlapping estimation period 1963Q4–2005Q2, the correlation coefficient between the TFP news shock series (red dashed line) as identified in Kurmann and Otrok (2013) and the λ-shock series (blue solid line) is 0.72. Based on the overlapping estimation period 1969Q1–2008Q3, the correlation coefficient between the monetary policy shock series (black dashed line) as identified in Romer and Romer (2004) (and updated by Olivier

Coibion) and the λ-shock series is 0.73. As will be shown further below, these high correlations are not a statistical artifact, but a robust feature of the data.

To reiterate, my identification strategy is unrelated to those frequently used in the macroeconomic literature as it (i) makes no assumption about the λ-shock's contribution to the forecast error variance of any of the variables, (ii) does not rely on any narrative measures such as FOMC records, (iii) does not impose any zero-type or sign restrictions and (iv) does not even include TFP as an observable in the VAR. Not to mention the additional differences of my empirical model in terms of lag structure, sample period and variables used in the VAR. The fact that I come close to reconstructing the object the TFP news literature and the monetary policy literature have studied (by applying a completely different and relatively more agnostic methodology) could provide strong empirical support for the relevance of these shocks in driving business cycles as well as asset price dynamics.

3.3 Pricing the Cross-section of Stock Returns

To explore the asset-pricing characteristics of the λ-shock, this subsection examines the performance of the corresponding 1-factor model in explaining average returns on portfolios formed to produce large spread in Size and B/M. During this exercise, I will treat the identified λ-shock as a known factor when estimating the two-pass regression model 2.3–2.4. To estimate the risk premium associated with the λ-shock, I apply the GMM procedure described in Cochrane (2005) and implemented by Burnside (2011).

Table 2: The First-pass Regression: the 1-factor Model with the λ-shock

1-factor model: $R_{it} = a_i + \varepsilon_t^{\star}\beta_{\lambda,i} + \epsilon_{it}$										
	β_λ					t-value				
	Low	2	3	4	High	Low	2	3	4	High
Small	0.55	1.23	1.68	1.70	2.12	0.41	1.17	1.64	1.91	2.26
2	0.73	1.39	1.91	1.99	1.75	0.59	1.36	2.11	2.06	2.06
3	1.17	1.54	1.63	2.01	1.81	1.05	1.67	1.85	2.40	2.25
4	1.03	1.09	1.72	2.01	2.00	1.03	1.11	2.13	2.45	2.55
Large	0.79	0.61	0.22	1.00	1.32	1.03	0.81	0.27	1.83	2.46

Notes: The table reports loadings on the identified λ-shock computed in time-series regressions for the FF25 portfolios sorted by size (in the rows) and book-to-market (in the columns). The sample period is 1963Q3-2008Q3. The t-statistics are computed based on the VARHAC procedure, following den Haan and Levin (2000); Burnside (2011), in order to take into account possible serial correlation in the errors.

Table 2 reports the estimates of the factor loadings computed in the first-pass time-series regressions 2.3. All portfolios have positive loadings on the λ-shock. Furthermore, the overall pattern is that small and value stocks have much larger exposure to surprise news about future technology than large growth stocks. For example, the point estimates suggest that extreme values stocks (2.12) have about four times larger exposures to the λ-shock than extreme growth stocks (0.55).

Table 3 displays the results for the second-pass of the Fama and MacBeth (1973) method. I also estimate factor prices for the 3-factor model proposed by Fama and French (1993) which has become the benchmark in the empirical asset pricing literature. The second-pass regressions are estimated both with and without a constant. To assess and compare the models' fit, I compute cross-sectional R^2-measures that adjust for degrees of freedom.

The results reported in Table 3 show that, in terms of model fit as well as statistical significance of the estimated risk prices, the 1-factor model including the identified λ-shock performs at least as well as the 3-factor model. The risk premium estimates suggest that a unit exposure to a λ-shock demands around 1.2–1.4%-points additional excess returns per quarter. The estimated constant is not statistically different from zero, implying small pricing errors for the 1-factor model. As for the risk price estimates associated with 3-factor model of Fama and French (1993), they are similar to those obtained in the literature (e.g. Petkova (2006)).

Table 3: The Second-pass Regressions: 1-factor Model vs. Fama-French 3-factor Model

Factor prices (λ)				Adj. R^2
Panel A: 1-factor Model with the λ-shock				
Constant	λ-shock			
0.32	1.21			0.761
(0.80) [1.26]	(0.264) [0.40]			
	1.41			0.747
	(0.46) [0.79]			
Panel B: The Fama-French 3-factor Model				
Constant	MKT	HML	SMB	
2.92	-1.62	1.44	0.57	0.766
(1.12) [1.18]	(1.30) [1.35]	(0.43) [0.44]	(0.43) [0.44]	
	1.24	1.44	0.64	0.734
	(0.63) [0.63]	(0.43) [0.44]	(0.43) [0.44]	

Notes: The table reports the cross-sectional regressions using the excess returns on the FF25 portfolios. The coefficients are expressed as percentage per quarter. Panel A presents results for the 1-factor model where the identified λ-shock is used as the sole pricing factor. Panel B presents results for the Fama-French 3-factor model. MKT is the market factor, HML is the book-to-market factor and SMB is the size factor. OLS standard errors are in parentheses, whereas standard errors, computed with the method of Shanken (1992) to adjust for errors-in-variables, are in brackets.

To provide a visual illustration on the remarkable pricing performance of the 1-factor model, even in comparison with the Fama-French 3-factor model, Figure 3 displays the fitted expected return of each FF25 portfolio (without using a constant in the cross-section regression) against its realised average return. The realised average returns are the time-series averages of the portfolio returns. If the model priced the cross-section of returns perfectly, then the points would lie on the 45-degree line through the origin. As shown by Figure 3, the 1-factor model does better than the Fama-French 3-factor model.

Overall, these results suggest that a single macroeconomic shock seems to perform as

Figure 3: Fitted Expected Returns vs. Average Realised Returns for 1963Q3–2008Q3

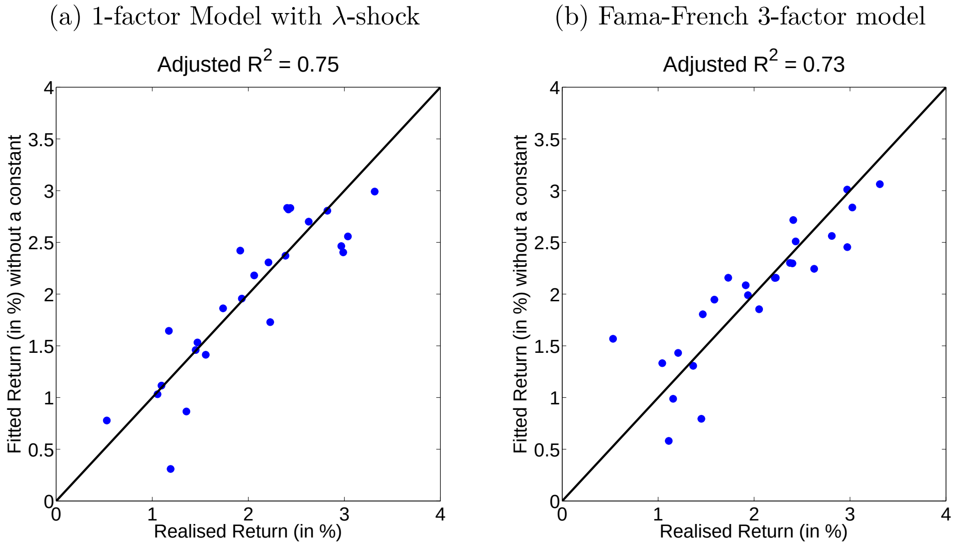

Notes: The R^2 values for the 1-factor model (left panel) and the 3-factor Fama-French model (right panel) are calculated assuming no constant in the second-pass regressions.

well in explaining the FF25 portfolios as the 3-factor model of Fama and French (1993) (which itself was constructed and sorted on the same basis as the FF25 portfolios). To the best of my knowledge, my findings related to the importance of a single structural shock in explaining the FF25 portfolios are novel. This is consistent with Campbell (1996) who argued that "innovations in variables that have been shown to forecast stock returns and labour income should be used in cross-sectional asset pricing studies" (pp. 312), and my results also provide an answer to Cochrane (2005) who calls for understanding "the real, macroeconomic, aggregate, nondiversifiable risk that is proxied by the returns of the HML and SMB portfolios" (pp. 442). In this sense, my findings suggest that innovations related to the technology news process or monetary policy surprises are the types of macroeconomic risks that explain most of the cross-sectional variation in the HML and SMB portfolios.

To highlight the role of structural identification in obtaining these results, I check the cross-sectional fit of the five 1-factor models that include each one of the five reduced-form VAR innovations, separately, as pricing factors. Figure 9 in the Appendix displays the scatter plots of the average versus fitted excess returns of the FF25 portfolios, corresponding to each one of the five 1-factor models. The results confirm that the individual reduced-form innovations fail markedly when used separately to price the cross-section.

3.4 Extensions and Robustness

3.4.1 The Equivalence between Maximising the Price of Risk and Maximising the Cross-sectional Fit[9]

A natural criticism of my identification is its somewhat arbitrary nature, as it searches for a shock based on the magnitude of the associated price of risk. After all, finding the structural shock that best explains the cross-section of stock returns should be equivalent to finding the shock that minimises the sum of squares of the pricing errors. This would also be reflected in the overall fit of the model, as measured by the R^2 statistic:

$$R^2 = 1 - \frac{\left[\bar{R}^e - \hat{\beta}\left(\varepsilon^{\star}\right) \times \hat{\lambda}\left(\varepsilon^{\star}\right)\right]' \left[\bar{R}^e - \hat{\beta}\left(\varepsilon^{\star}\right) \times \hat{\lambda}\left(\varepsilon^{\star}\right)\right]}{\left[\bar{R}^e - \ddot{R}^e\right]' \left[\bar{R}^e - \ddot{R}^e\right]}, \tag{3.1}$$

where $\ddot{R}^e = \frac{1}{k}\sum_{i=1}^{k} \bar{R}_i^e$ is the cross-sectional average of the mean returns in the data, $\hat{\beta}\left(\varepsilon^{\star}\right) \times \hat{\lambda}\left(\varepsilon^{\star}\right)$ is the model's predicted mean returns and the estimated pricing errors are the residuals, $\hat{\alpha} = \bar{R}^e - \hat{\beta}\left(\varepsilon^{\star}\right) \times \hat{\lambda}\left(\varepsilon^{\star}\right)$. The relationship, between the A_0 matrix that maps the reduced-form innovations to the structural shocks and the R^2 measure implied by the 1-factor model that uses the time-series $\varepsilon^{\star}$ as the pricing factor, seems complicated. It is therefore difficult to write down in closed-form the theoretical relationship between the identification strategy that maximises λ and the strategy that minimises the sum of squared pricing errors, or equivalently, maximises the R^2.

I therefore perform a simulation exercise to show the equivalence between the two identification strategies. The first step is to recognise that the unadjusted R^2 (as computed in 3.1) associated with any one of the possible structural shocks obtained from the reduced-form VAR model 2.1 is bounded from above by the unadjusted R^2 of a five-factor model that would use all five reduced-from or structural-form shocks. It is important to note that this bound is determined by the specification of the reduced-form VAR model and does not depend on identifying assumptions. After all, identification merely 'rotates' the information set, and does not augment it. In the case of the baseline model without a constant in the second-pass regression, I obtain an unadjusted R^2-measure of 0.747. Any one of the five structural shocks associated with a candidate draw $\tilde{Q}$ cannot contain more information than that contained in the five reduced-form innovations.

The next step is to explore the space of admissible $\tilde{Q}$ matrices and uncover the relationship between λ and the R^2 implied by the corresponding 1-factor model. The scatter plot in Figure 4 displays this relationship based on 20,000 random admissible matrices, all of which are consistent with the reduced-form variance covariance matrix. To obtain these random draws, I apply Householder transformations to five-dimensional matrices

[9] I would like to thank Professors John Cochrane and Shengxing Zhang for their comments that inspired me to write this subsection.

Figure 4: The Identification of the λ-shock: a Simulation Exercise to Illustrate the Relationship between the Price of Risk and Cross-sectional R^2

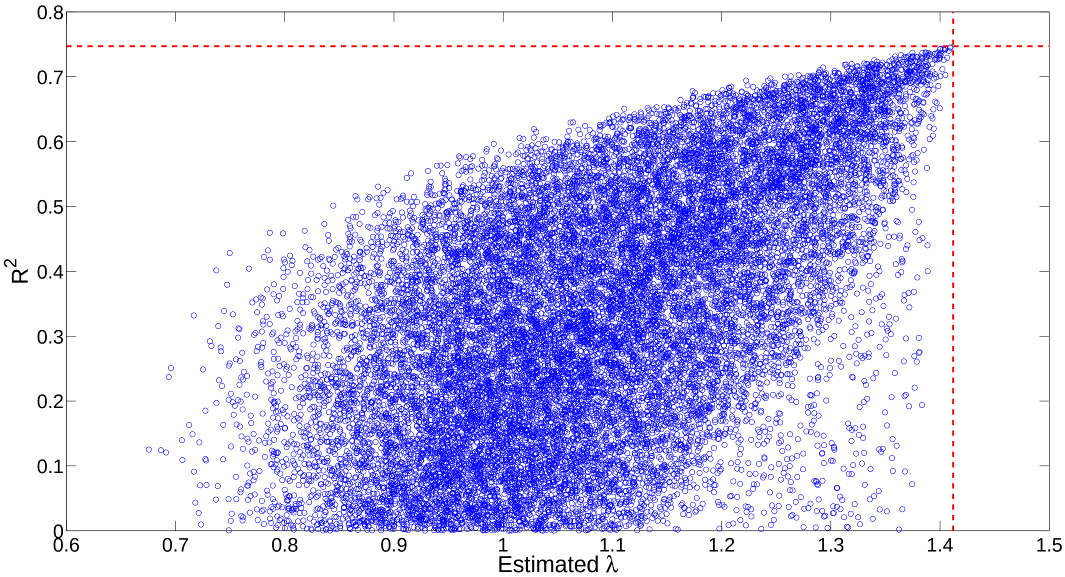

Notes: The scatter plot (based on 20,000 random $\tilde{Q}$ matrices) shows the relationship between the price of risk demanded by $\tilde{\varepsilon}_t$ associated with a given candidate draw $\tilde{Q}$ and the cross-sectional R^2 implied by the corresponding 1-factor model. For presentation purposes, I exclude those rotations that imply negative R^2 (about 48% of all admissible matrices), as it does not cause any loss of generality in the relationship. The vertical red dashed line is the maximum achievable price of risk (1.41) from the five-variable VAR model 2.1, and the horizontal red dashed line is the upper bound (0.747) on the unadjusted R^2-measure associated with any 1-factor model extracted from the VAR model 2.1. To obtain these random draws, I apply Householder transformations to 20,000 five-dimensional matrices drawn from the multivariate Normal distribution.

drawn from the multivariate Normal distribution. The vertical red dashed line denotes the maximum achievable price of risk (1.41) associated with the λ-shock given the VAR specification. The horizontal red dashed line denotes the upper bound on the unadjusted R^2 (0.747), which puts a cap on how well the identified structural shock can explain the cross-section.

There are at least two messages conveyed by Figure 4. First, it suggests that if an admissible model generates a structural shock with a high price of risk, then the corresponding 1-factor model tends to have a high R^2. This observation is based on the darker, densely populated range of the scatter, which most admissible models fall into. Of course, a critic may point out that there are a few admissible models that indeed perform very poorly in pricing the cross section in spite of the fact that they command a high price of risk (bottom-right part of the scatter), and there are also shocks that fare well in asset pricing in spite of the relatively low price of risk they demand (left part of the scatter). Nevertheless, the second message seems very clear: as the random 1-factor models get closer and closer to the upper bound in terms of the implied R^2 values, the associated prices of risk converge to the maximum price of risk that is numerically achievable. Increasing the number of random draws does not change Figure 4.[10] I therefore interpret

[10]These results are available upon request.

this as a numerical proof of the equivalence between maximising the price of risk and maximising the cross-sectional fit – two possible identification strategies to uncover the λ-shock.

This equivalence may however not be surprising: inspecting the definition 3.1 of the cross-sectional fit makes it clear that identification does not affect average excess returns or the maximum achievable R^2. Therefore, conditional on the reduced-form VAR specification, the best possible identification that delivers a structural shock with the highest cross-sectional R^2 must either generate large dispersion in factor loadings in the first stage, or it must generate a high price of risk in the second stage, or both. To explore the relationship between the cross-sectional dispersion (measured by the standard deviation) of βs and λ, Figure 10 in the Appendix shows the results from a simulation exercise similar to the one above. Interestingly, the contour of the dispersion of βs seems quadratic in λ: for a given highest price of risk, there is an infinite number of β dispersions with a bounded range. The exception is numerically at the unique standard deviation value (0.0054) where the associated λ has the highest achievable price of risk – the point where the red dashed line touches the contour. This also suggests that the identification problem can also be possibly reformulated as an optimisation problem in the dimension of β dispersions.

3.4.2 Addressing the Lewellen, Nagel, and Shanken (2010) Critique[11]

In addition, the reader may criticise my identification strategy as it is based on maximising price of risk according to a particular set of stock portfolios: the FF25 portfolios. The choice of using the FF25 as the basis for my baseline identification is motivated by the long prevailing consensus that these portfolios are a good representation of the aggregate risk present in the cross-section (Chapter 20 of Cochrane (2005)). However, recent research such as Lewellen, Nagel, and Shanken (2010) has pointed at the strong factor structure of the FF25 portfolios which makes it relatively easy to find factors that generate high cross-sectional R^2s. The first three principal components (approximately the three Fama-French factors) explain most of the time-series and cross-sectional variation in returns. In fact they show that any proposed factors that are weakly correlated with the SMB and HML factors, but not with the small, idiosyncratic three-factor residuals of FF25 portfolios, are likely to generate high cross-sectional R^2 values.

To address their critique, I follow prescription 1 (pp. 182) of Lewellen, Nagel, and Shanken (2010) by expanding the set of test portfolios beyond FF25 and adding to them the 30 industry portfolios of Fama-French (FF30 henceforth). They argue that such an expansion of the portfolio set serves the purpose of relaxing the tight factor structure of Size-B/M so that it would be much 'harder' for artificial factors to explain expected

[11] I would like to thank Professor Stefan Nagel for helpful comments.

Figure 5: Addressing the Lewellen, Nagel, and Shanken (2010) Critique: The Identification of the λ-shock is Robust to the Inclusion of the 30 Industry Portfolios

(a) Impulse Responses to a λ-shock: FF25 versus FF25+FF30 Portfolios

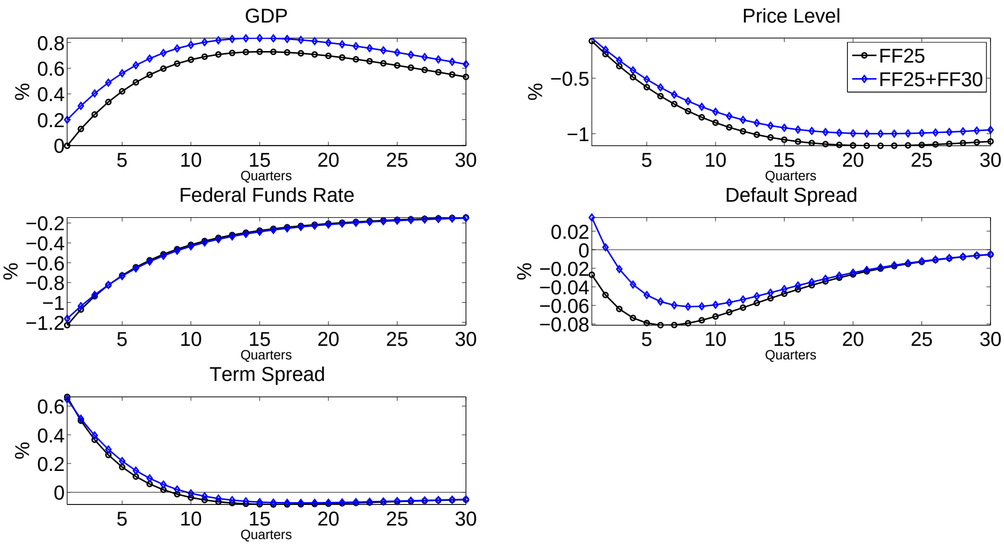

(b) Comparing the Time-series of λ-shocks: FF25 versus FF25+FF30 Portfolios

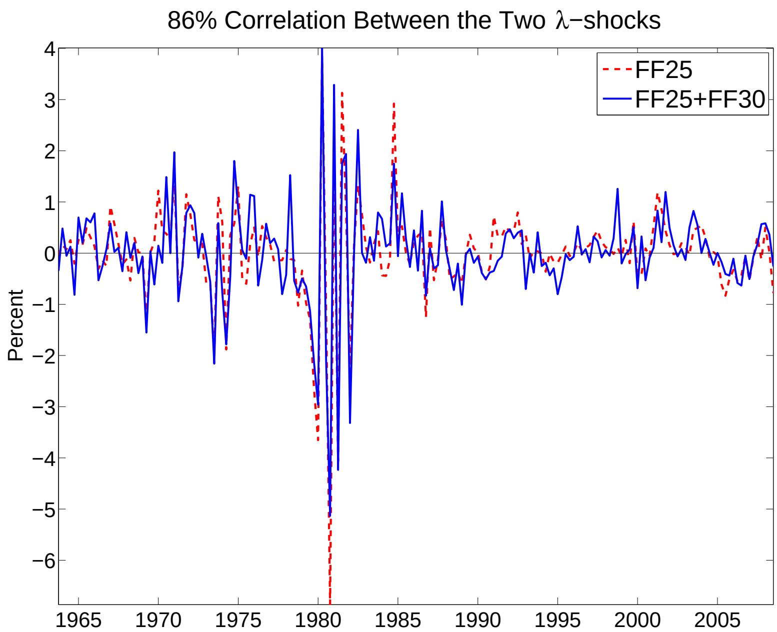

Notes: In panel a, the vertical axes are %-deviations from steady-state, and the horizontal axes are in quarters. Panel b shows the estimated time-series of the λ-shock using the FF25 (red dashed line) and the FF25+FF30 (blue solid line) as testing portfolios for identification.

returns on the resulting 55 portfolios. I therefore estimate a modified λ-shock series by searching for the structural shock that demands the highest possible risk premium in absolute value when pricing the 55 Size–B/M and Industry portfolios (FF25+FF30). The pricing performance of the corresponding 1-factor model is summarised by Table 5 of the Appendix. Indeed, the cross-sectional adjusted R^2 drops drastically: it falls from 0.75 to 0.12 for the 1-factor model without a common constant, and it drops from 0.73 to 0.1 for the 3-factor model of Fama-French without a common constant. This can be interpreted as the relevant information content of the VAR being much smaller for pricing the FF25+FF30 portfolios than for pricing the FF25 portfolios.

However, conditional on this smaller relevant information set, the identification recovers virtually the same shock as the one obtained by using the FF25 only. Panel *a* of Figure 5 plots the impulse response functions using the FF25 (black circled lines) and the FF25+FF30 (blue crossed lines). With the exception of the slight level shift in the output response and the delayed effect on default spreads, the results seem very similar across the two models. Consequently, the time-series of the estimated shocks continue to look alike as shown by panel *b* of Figure 5. Also, the correlation coefficient with respect to the TFP news shock series of Kurmann and Otrok (2013) drops only slightly from 0.72 to 0.67 when using the FF25+FF30 portfolios instead of the FF25. The correlation coefficient with respect to the monetary shock series of Romer and Romer (2004) increases from 0.73 to 0.74 when using the augmented portfolio set.

The overall interpretation of these results can be as follows: including the 30 Industry portfolios may lead to a critique of the (lack of) relevant information content of the VAR for pricing the FF25+FF30 portfolios, but does not give rise to a critique of my identification: the macroeconomic shock, that captures all relevant information for pricing the cross section (irrespective of whether the information content is relatively small or large) continues to bear the same economic characteristics as the λ-shock using the baseline FF25 portfolios.

3.4.3 Changing the VAR Specification and Results from a Bayesian VAR

In addition, I check whether the results are robust to changing the specification of the VAR model. Panel B of Table 4 shows the cross-correlations among the different λ-shock series with the TFP news shock series and monetary policy shock series. I explore increasing the lag length of the VAR and experiment with alternative measures of GDP, the aggregate price level and the term spread. The results suggest that replacing real GDP with the real consumption or using CPI instead of the PCE price index can yield a similar correlation with the TFP news shock series. Conversely, increasing the lag length of the VAR, using alternative measures of the term spread or using the quarterly industrial production index instead of GDP can reduce this correlation. Overall, I find that changing the specification of the VAR does not have a material impact on the results.

Table 4: Robustness of the Identification of the λ-shock to Changing the VAR Model

Correlation Coefficients						
λ-shock Baseline	λ-shock VAR(2)	λ-shock VAR – C	λ-shock VAR – IP	λ-shock VAR – CPI	λ-shock VAR – Spr	External Shocks
1.00						
0.86	1.00					
0.86	0.77	1.00				
0.81	0.80	0.73	1.00			
0.99	0.84	0.81	0.78	1.00		
0.98	0.85	0.85	0.78	0.98	1.00	
0.72	0.64	0.70	0.61	0.71	0.71	TFP News
0.73	0.61	0.70	0.52	0.71	0.71	Monetary Policy

Notes: The table reports the correlation coefficients among λ-shocks from the baseline (Column 1), the baseline VAR with 2 lags (Column 2), the VAR using the consumption measure from Greenwald, Lettau, and Ludvigson (2015) instead of GDP (Column 3), the VAR after replacing GDP with the real monthly Industrial Production Index (FRED code: INDPRO) lead by a month and averaged over each quarter (Column 4), the VAR using CPI (FRED code: CPIAUCSL) as an alternative measure of the aggregate price index (Column 5), the VAR using the difference between the 10-year Treasury constant maturity rate (FRED code: GS10) and the Federal Funds rate as an alternative measure of the term spread (Column 6), and the external shocks (Column 7). The values are computed based on the overlapping period 1963Q4–2005Q2 with Kurmann and Otrok (2013), except the last row which is using data for 1969Q1–2008Q3, dictated by the availability of the monetary policy shock series of Romer and Romer (2004).

Moreover, I explore the role of parameter uncertainty in the VAR model 2.1 by re-estimating the model with Bayesian methods. I use Minnesota-type normal inverted Wishart priors that I impose using the dummy observation approach of Sims and Zha (1998), as implemented in Banbura, Giannone, and Reichlin (2010). To approximate the posterior marginal distribution of the VAR parameters, I set up the Gibbs-sampler whereby I use the well-known analytical formulae for the conditional distributions of the dynamic parameters and the variance covariance matrix of the VAR. To construct a probability distribution for the impulse response functions of the λ-shock, I proceed as follows: (i) I burn the first N_1 draws from the conditional distributions to avoid potential problems of initial values, (ii) draw a $B - \Sigma$ pair of VAR parameters from the conditional distributions, (iii) apply the identification method to these draws and save the resulting IRFs, and (iv) and repeat the Gibbs-iteration and the identification for another N_2 times. The posterior distribution of IRFs is then constructed based on the N_2 draws.

Figure 6 shows the posterior distribution of IRFs of the λ-shock. A one standard deviation expansionary λ-shock continues to have a delayed effect on output. The $16^{th} - 84^{th}$ probability bands suggest that at a 15-quarter horizon output rises around 0.5% above steady-state with 84% probability but does not rise more than 0.9% above steady-state with the same probability. At the same horizon, the price level more than 0.8% below steady-state with 84% probability but does not fall more than 1.3% below steady-state with the same probability. The smallest degree of uncertainty, reflected in the tight probability bands, is around the impact on the term spread. As shown by Figure 11 of

Figure 6: Impulse Responses to a λ-shock: Results from a Bayesian VAR(1)

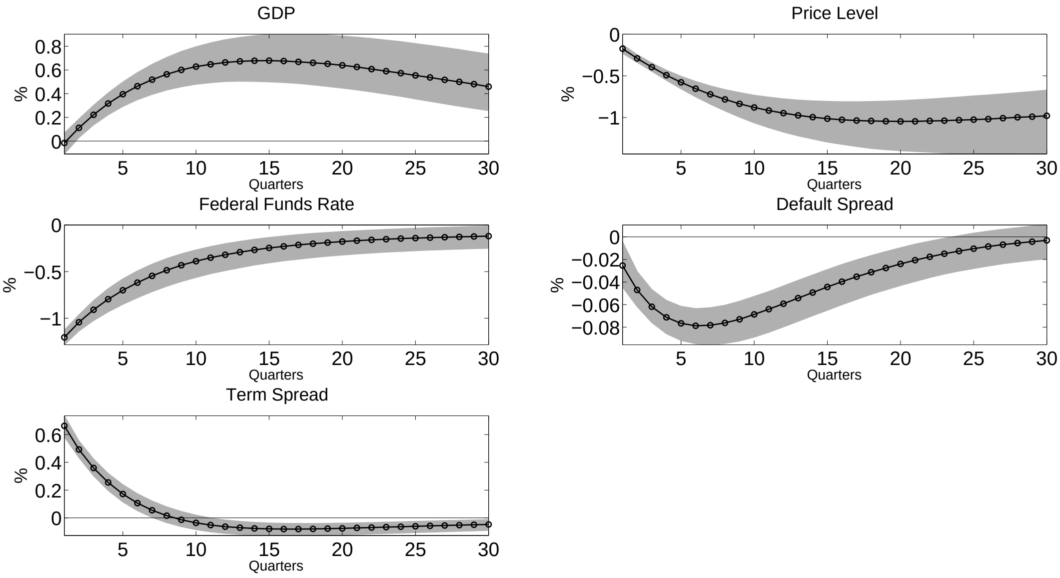

Notes: The sample period is 1963Q3 - 2008Q3. The Minnesota-type normal inverted Wishart priors are implemented following Banbura, Giannone, and Reichlin (2010). The figure shows the pointwise median and 16th-84th percentiles of $N_2 = 5000$ draws (after burning the first $N_1 = 5000$ draws) from the posterior distribution of the impulse responses. The vertical axes are %-deviations from steady-state, and the horizontal axes are in quarters.

the Appendix, these results are quantitatively similar when increasing the lag length of the BVAR from one to two lags.

3.4.4 The λ-shock and the Fundamentals

An application of my proposed identification strategy to the stock portfolios of FF25 led to the result that the estimated λ-shock bears a close empirical relationship both with TFP news shocks and with monetary policy shocks. The reader may justifiably feel uncomfortable with the ambiguous nature of such a conclusion, and may blame the somewhat reverse engineering and overly agnostic nature of my identification strategy for it. After all, how can the resulting λ-shock have such a high correlation with two, seemingly distinct structural disturbances? To convince the reader of the usefulness of my identification strategy, I propose one possible and simple explanation for such an ambiguity: TFP news shocks and monetary policy shocks are in fact highly correlated in the data.

To provide some suggestive evidence for this argument, I use the VAR model of Kurmann and Otrok (2013) to identify a monetary policy shock using Cholesky orthogonalisation as done by Sims (1980), Christiano, Eichenbaum, and Evans (1999) and many others in the monetary policy literature. In this case, I deliberately use exactly the same VAR specification as used by Kurmann and Otrok (2013) when they identified a TFP news shock so that I can learn about differences and similarities across the two identification themes without changing the information set. The upper panel of Figure 7 plots the estimated time-series of the TFP news shocks (black dashed line) against the monetary policy shock series identified with Cholesky orthogonalisation (red solid line). The correlation between the two series is strikingly high (0.96), raising questions about the orthogonality of these shocks with respect to one another.

Figure 7: Comparing TFP News Shocks against Monetary Policy Shocks: Results from Kurmann and Otrok (2013)'s VAR and from Smets and Wouters (2007)'s DSGE Model.

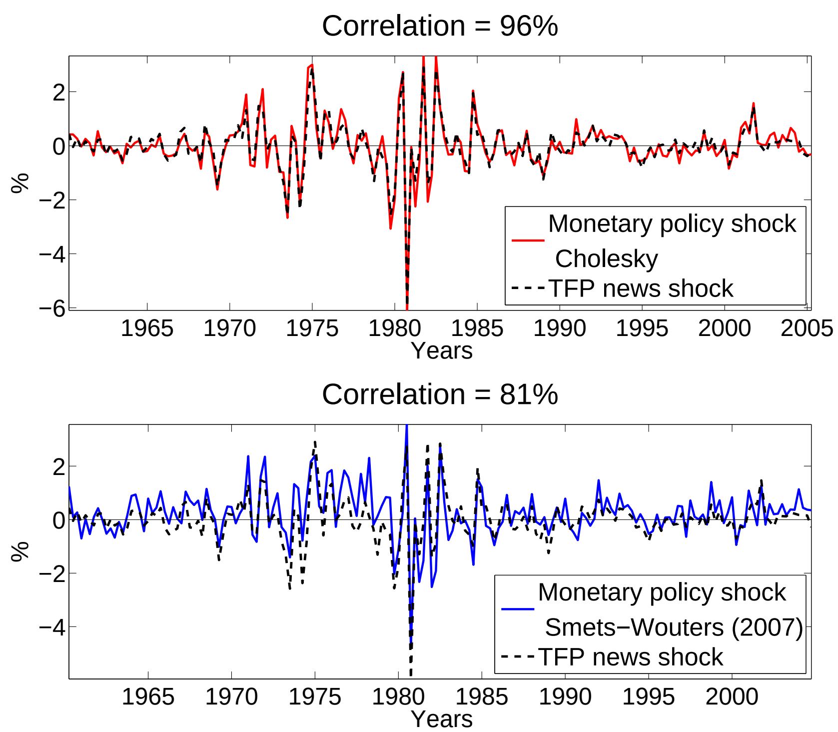

Notes: The TFP news shock series (black dashed line) are the ones plotted in Figure 5 on pp. 2625 of Kurmann and Otrok (2013) who apply the method of Uhlig (2004) to identify a TFP news shock over the period 1959Q2-2005Q2. The monetary policy shock series in the upper panel (red solid line) are identified with Cholesky identification as in Christiano, Eichenbaum, and Evans (1999), using the same variables and lag length as Kurmann and Otrok (2013). The monetary policy shock series in the lower panel (blue solid line) are the estimated time-series of innovations in the Taylor-rule in the DSGE model of Smets and Wouters (2007).

Of course, the identification of monetary policy shocks with Cholesky orthogonalisation is only one of the many possible identification strategies. Therefore, I provide additional evidence from the structural model of Smets and Wouters (2007) which is a dynamic stochastic general equilibrium (DSGE) model estimated with Bayesian methods. Monetary policy shocks in this framework are the estimated innovations in a Taylor-type monetary policy rule. The estimated time-series of these structural innovations from the DSGE model are plotted in the lower panel of Figure 7 (blue solid line) against the TFP news shocks (black dashed line) of Kurmann and Otrok (2013). The correlation between these two series is still remarkably high (0.81).

I interpret these findings as confirmation that the somewhat ambiguous characterisation of the obtained λ-shock is not an outcome of the potential weakness of my identification theme, but it is a result of the high empirical correlation between the two, well-known structural disturbances that the λ-shock resembles. To the best of my knowledge, this empirical regulairty has not been documented in the literature yet, and it could be subject to further research.

4 Conclusion

This paper proposed a new identification theme based on the ability of the obtained orthogonalised shock to explain the cross section of asset returns. The identification theme is motivated by the long-standing challenge to link the origins of cross-sectional variation in stock returns to macroeconomic primitives. When applying the method to the FF25 stock portfolios of Fama and French (1993) or the augmented FF25+FF30 portfolios, the obtained structural shock exhibits meaningful economic characteristics and bears close resemblance with well-known structural shocks studied by the macroeconomic literature. My results may have the following implications.

First, the structural shock that is responsible for most of the aggregate risk captured by the cross-section of stock returns is not related to the unanticipated shocks that tend to generate immediate jumps in aggregate quantities: the IRF analysis made it clear that aggregate output responds with a considerable delay to the λ-shock. This is in sharp contrast with the majority of the macroeconomic literature that focuses on unanticipated shocks as sources of business cycles.

Second, many have regarded the ICAPM as a "fishing license" (Fama, 1991) for empirical multifactor models aiming to explain the cross-sectional variation in stock returns. My results show that macroeconometric identification is key to finding the type of structural innovations in state variables, that macroeconomic general equilibrium models have studied, and that are the major hedging concerns to investors. In this sense, the "fishing license" is largely restricted once we force ourselves to use a pricing factor which also behaves like a structural shock with meaninful economic characteristics. In fact, this paper could uncover only one such a pricing factor: the λ-shock.

Finally, the identification strategy I propose is not restricted to stock returns and could easily be generalised to understanding the macroeconomic forces behind aggregate risks underlying portfolios in other asset classes and markets, e.g. bond portfolios, international currency portfolios, assets sorted on liquidity characteristics. This could be an interesting avenue for future research.

References

Banbura, M., D. Giannone, and L. Reichlin (2010): "Large Bayesian vector auto regressions," *Journal of Applied Econometrics*, 25(1), 71–92.

Bansal, R., and A. Yaron (2004): “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles,” *Journal of Finance*, 59(4), 1481–1509.

Barsky, R. B., S. Basu, and K. Lee (2014): *Whither News Shocks?* pp. 225–264. University of Chicago Press.

Barsky, R. B., and E. R. Sims (2011): “News shocks and business cycles,” *Journal of Monetary Economics*, 58(3), 273–289.

Beaudry, P., and F. Portier (2006): “Stock Prices, News, and Economic Fluctuations,” *American Economic Review*, 96(4), 1293–1307.

——— (2014): “News-Driven Business Cycles: Insights and Challenges,” *Journal of Economic Literature*, 52(4), 993–1074.

Beeler, J., and J. Y. Campbell (2012): “The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment,” *Critical Finance Review*, 1(1), 141–182.

Boons, M. (2016): “State variables, macroeconomic activity, and the cross section of individual stocks,” *Journal of Financial Economics*, 119(3), 489–511.

Boons, M., and A. Tamoni (2015): “Horizon-Specific Macroeconomic Risks and the Cross-Section of Expected Returns,” mimeo, LSE.

Borovicka, J., and L. P. Hansen (2014): “Examining macroeconomic models through the lens of asset pricing,” *Journal of Econometrics*, 183(1), 67 – 90, Internally Consistent Modeling, Aggregation, Inference and Policy.

Brennan, M. J., A. W. Wang, and Y. Xia (2004): “Estimation and Test of a Simple Model of Intertemporal Capital Asset Pricing,” *Journal of Finance*, 59(4), 1743–1776.

Bryzgalova, S., and C. Julliard (2015): “The Consumption Risk of Bonds and Stocks,” mimeo, LSE.

Burnside, C. (2011): “The Cross Section of Foreign Currency Risk Premia and Consumption Growth Risk: Comment,” *American Economic Review*, 101(7), 3456–76.

Campbell, J. Y. (1996): “Understanding Risk and Return,” *Journal of Political Economy*, 104(2), 298–345.

Campbell, J. Y., S. Giglio, and C. Polk (2013): “Hard Times,” *Review of Asset Pricing Studies*, 3(1), 95–132.

Campbell, J. Y., and R. J. Shiller (1988): “The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors,” *Review of Financial Studies*, 1(3), 195–228.

Campbell, J. Y., and T. Vuolteenaho (2004): “Bad Beta, Good Beta,” *American Economic Review*, 94(5), 1249–1275.

Chen, N.-F., R. Roll, and S. A. Ross (1986): "Economic Forces and the Stock Market," *The Journal of Business*, 59(3), 383–403.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (1999): "Monetary policy shocks: What have we learned and to what end?," in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford, vol. 1, chap. 2, pp. 65–148. Elsevier.

Christiano, L. J., R. Motto, and M. Rostagno (2014): "Risk Shocks," *American Economic Review*, 104(1), 27–65.

Cochrane, J. H. (1996): "A Cross-Sectional Test of an Investment-Based Asset Pricing Model," *Journal of Political Economy*, 104(3), 572–621.

——— (2005): *Asset Pricing*. Princeton University Press.

Coibion, O. (2012): "Are the Effects of Monetary Policy Shocks Big or Small?," *American Economic Journal: Macroeconomics*, 4(2), 1–32.

Constantinides, G. M., and A. Ghosh (2011): "Asset Pricing Tests with Long-run Risks in Consumption Growth," *Review of Asset Pricing Studies*, 1(1), 96–136.

den Haan, W. J., and A. T. Levin (2000): "Robust Covariance Matrix Estimation with Data-Dependent VAR Prewhitening Order," NBER Technical Working Papers 0255, National Bureau of Economic Research, Inc.

Fama, E., and K. French (1992): "The Cross-Section of Expected Stock Returns," *Journal of Finance*, 47(2), 427–65.

——— (1993): "Common risk factors in the returns on stocks and bonds," *Journal of Financial Economics*, 33(1), 3–56.

——— (2015): "A five-factor asset pricing model," *Journal of Financial Economics*, 116(1), 1 – 22.

Fama, E. F. (1991): "Efficient Capital Markets: II," *The Journal of Finance*, 46(5), 1575–1617.

Fama, E. F., and J. D. MacBeth (1973): "Risk, Return, and Equilibrium: Empirical Tests," *Journal of Political Economy*, 81(3), 607–36.

Ferson, W. E., and C. R. Harvey (1991): "The Variation of Economic Risk Premiums," *Journal of Political Economy*, 99(2), 385–415.

Fry, R., and A. Pagan (2011): "Sign Restrictions in Structural Vector Autoregressions: A Critical Review," *Journal of Economic Literature*, 49(4), 938–60.

Gertler, M., and P. Karadi (2015): "Monetary Policy Surprises, Credit Costs, and Economic Activity," *American Economic Journal: Macroeconomics*, 7(1), 44–76.

Goyal, A., and I. Welch (2008): "A Comprehensive Look at The Empirical Performance of Equity Premium Prediction," *Review of Financial Studies*, 21(4), 1455–1508.

Greenwald, D. L., M. Lettau, and S. Ludvigson (2015): "Origins of Stock Market Fluctuations," CEPR Discussion Papers 10336, C.E.P.R. Discussion Papers.

Hansen, L. P., J. C. Heaton, and N. Li (2008): "Consumption Strikes Back? Measuring Long-Run Risk," *Journal of Political Economy*, 116(2), 260–302.

He, Z., B. Kelly, and A. Manela (2016): "Intermediary Asset Pricing: New Evidence from Many Asset Classes," Working Paper 21920, National Bureau of Economic Research.

Jagannathan, R., and Y. Wang (2007): "Lazy Investors, Discretionary Consumption, and the Cross-Section of Stock Returns," *The Journal of Finance*, 62(4), 1623–1661.

Jagannathan, R., and Z. Wang (1996): "The Conditional CAPM and the Cross-Section of Expected Returns," *Journal of Finance*, 51(1), 3–53.

Jaimovich, N., and S. Rebelo (2009): "Can News about the Future Drive the Business Cycle?," *American Economic Review*, 99(4), 1097–1118.

Jermann, U. J. (1998): "Asset pricing in production economies," *Journal of Monetary Economics*, 41(2), 257–275.

Justiniano, A., G. Primiceri, and A. Tambalotti (2010): "Investment shocks and business cycles," *Journal of Monetary Economics*, 57(2), 132–145.

Kan, R., C. Robotti, and J. Shanken (2013): "Pricing Model Performance and the Two-Pass Cross-Sectional Regression Methodology," *Journal of Finance*, 68(6), 2617–2649.

Kliem, M., and H. Uhlig (2016): "Bayesian estimation of a dynamic stochastic general equilibrium model with asset prices," *Quantitative Economics*, 7(1), 257–287.

Koijen, R., H. Lustig, and S. van Nieuwerburgh (2012): "The Cross-Section and Time-Series of Stock and Bond Returns," CEPR Discussion Papers 9024, C.E.P.R. Discussion Papers.

Kurmann, A., and C. Otrok (2013): "News Shocks and the Slope of the Term Structure of Interest Rates," *American Economic Review*, 103(6), 2612–32.

Lettau, M., and S. Ludvigson (2001): "Resurrecting the (C)CAPM: A Cross-Sectional Test When Risk Premia Are Time-Varying," *Journal of Political Economy*, 109(6), 1238–1287.

Lewellen, J., S. Nagel, and J. Shanken (2010): "A skeptical appraisal of asset pricing tests," *Journal of Financial Economics*, 96(2), 175 – 194.

Liu, L. X., and L. Zhang (2008): "Momentum Profits, Factor Pricing, and Macroeconomic Risk," *Review of Financial Studies*, 21(6), 2417–2448.

Lustig, H. N., and S. G. V. Nieuwerburgh (2005): "Housing Collateral, Consumption Insurance, and Risk Premia: An Empirical Perspective," *Journal of Finance*, 60(3), 1167–1219.

Maio, P., and P. Santa-Clara (2012): "Multifactor models and their consistency with the ICAPM," *Journal of Financial Economics*, 106(3), 586–613.

Malkhozov, A., and A. Tamoni (2015): "News Shocks and Asset Prices," mimeo, LSE.

Mertens, K., and M. O. Ravn (2013): "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States," *American Economic Review*, 103(4), 1212–47.

Merton, R. C. (1973): "An Intertemporal Capital Asset Pricing Model," *Econometrica*, 41(5), 867–87.

Ng, S., and J. H. Wright (2013): "Facts and Challenges from the Great Recession for Forecasting and Macroeconomic Modeling," *Journal of Economic Literature*, 51(4), 1120–54.

Parker, J. A., and C. Julliard (2005): "Consumption Risk and the Cross Section of Expected Returns," *Journal of Political Economy*, 113(1), 185–222.

Petkova, R. (2006): "Do the Fama-French Factors Proxy for Innovations in Predictive Variables?," *The Journal of Finance*, 61(2), 581–612.

Piazzesi, M., M. Schneider, and S. Tuzel (2007): "Housing, consumption and asset pricing," *Journal of Financial Economics*, 83(3), 531–569.

Pinter, G., K. Theodoridis, and T. Yates (2013): "Risk news shocks and the business cycle," Bank of England working papers 483, Bank of England.

Romer, C. D., and D. H. Romer (2004): "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review*, 94(4), 1055–1084.

Rubio-Ramirez, J. F., D. F. Waggoner, and T. Zha (2010): "Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference," *Review of Economic Studies*, 77(2), 665–696.

Rudebusch, G. D., and E. T. Swanson (2012): "The Bond Premium in a DSGE Model with Long-Run Real and Nominal Risks," *American Economic Journal: Macroeconomics*, 4(1), 105–43.

Santos, T., and P. Veronesi (2006): "Labor Income and Predictable Stock Returns," *Review of Financial Studies*, 19(1), 1–44.

Schmitt-Grohe, S., and M. Uribe (2012): "What's News in Business Cycles," *Econometrica*, 80(6), 2733–2764.

SHANKEN, J. (1992): "On the Estimation of Beta-Pricing Models," *Review of Financial Studies*, 5(1), 1–33.

SIMS, C. A. (1980): "Macroeconomics and Reality," *Econometrica*, 48(1), 1–48.

——— (1992): "Interpreting the macroeconomic time series facts : The effects of monetary policy," *European Economic Review*, 36(5), 975–1000.

SIMS, C. A., AND T. ZHA (1998): "Bayesian Methods for Dynamic Multivariate Models," *International Economic Review*, 39(4), 949–68.

SMETS, F., AND R. WOUTERS (2007): "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 97(3), 586–606.

STOCK, J. H., AND M. W. WATSON (2003): "Forecasting Output and Inflation: The Role of Asset Prices," *Journal of Economic Literature*, 41(3), 788–829.

THORBECKE, W. (1997): "On Stock Market Returns and Monetary Policy," *Journal of Finance*, 52(2), 635–54.

UHLIG, H. (2004): "What moves GNP?," Econometric Society 2004 North American Winter Meetings 636, Econometric Society.

——— (2005): "What are the effects of monetary policy on output? Results from an agnostic identification procedure," *Journal of Monetary Economics*, 52(2), 381–419.

VASSALOU, M. (2003): "News related to future GDP growth as a risk factor in equity returns," *Journal of Financial Economics*, 68(1), 47–73.

YOGO, M. (2006): "A Consumption-Based Explanation of Expected Stock Returns," *Journal of Finance*, 61(2), 539–580.

A Additional Figures

Figure 8: Impulse Responses to a λ-shock: Output and Consumption

(a) VAR with Output

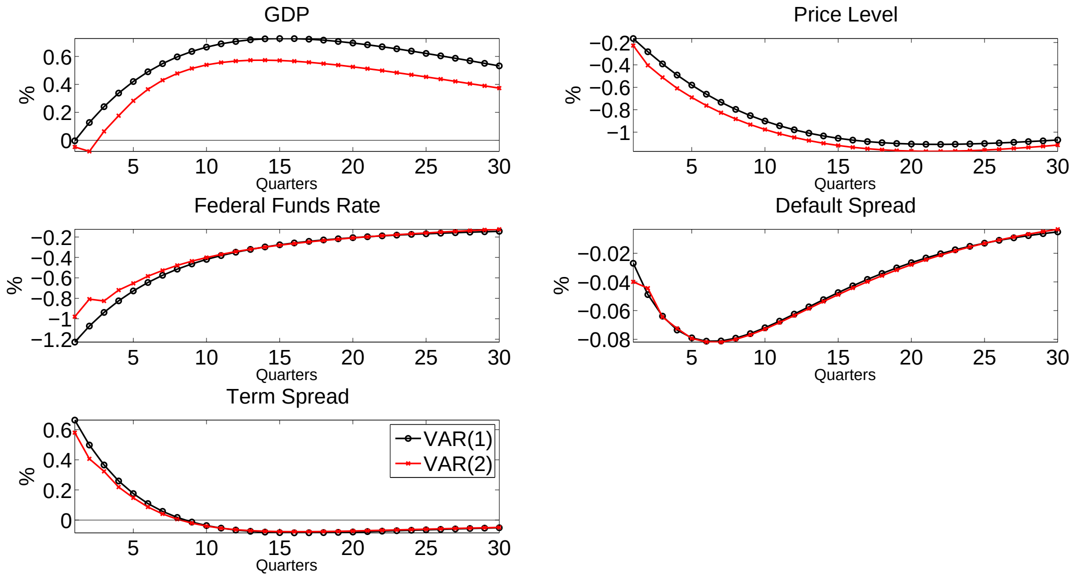

(b) VAR with Consumption

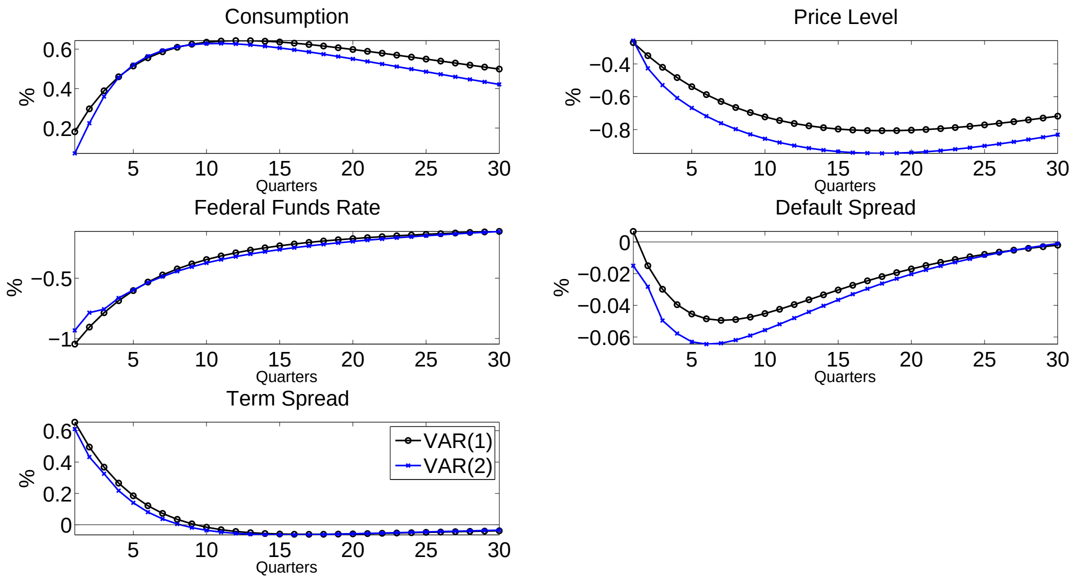

Notes: The vertical axes are %-deviations from steady-state, and the horizontal axes are in quarters.

Figure 9: The Role of Structural Identification: the Inability of the Individual VAR Innovations to Explain the Cross-section of FF25 Portfolios

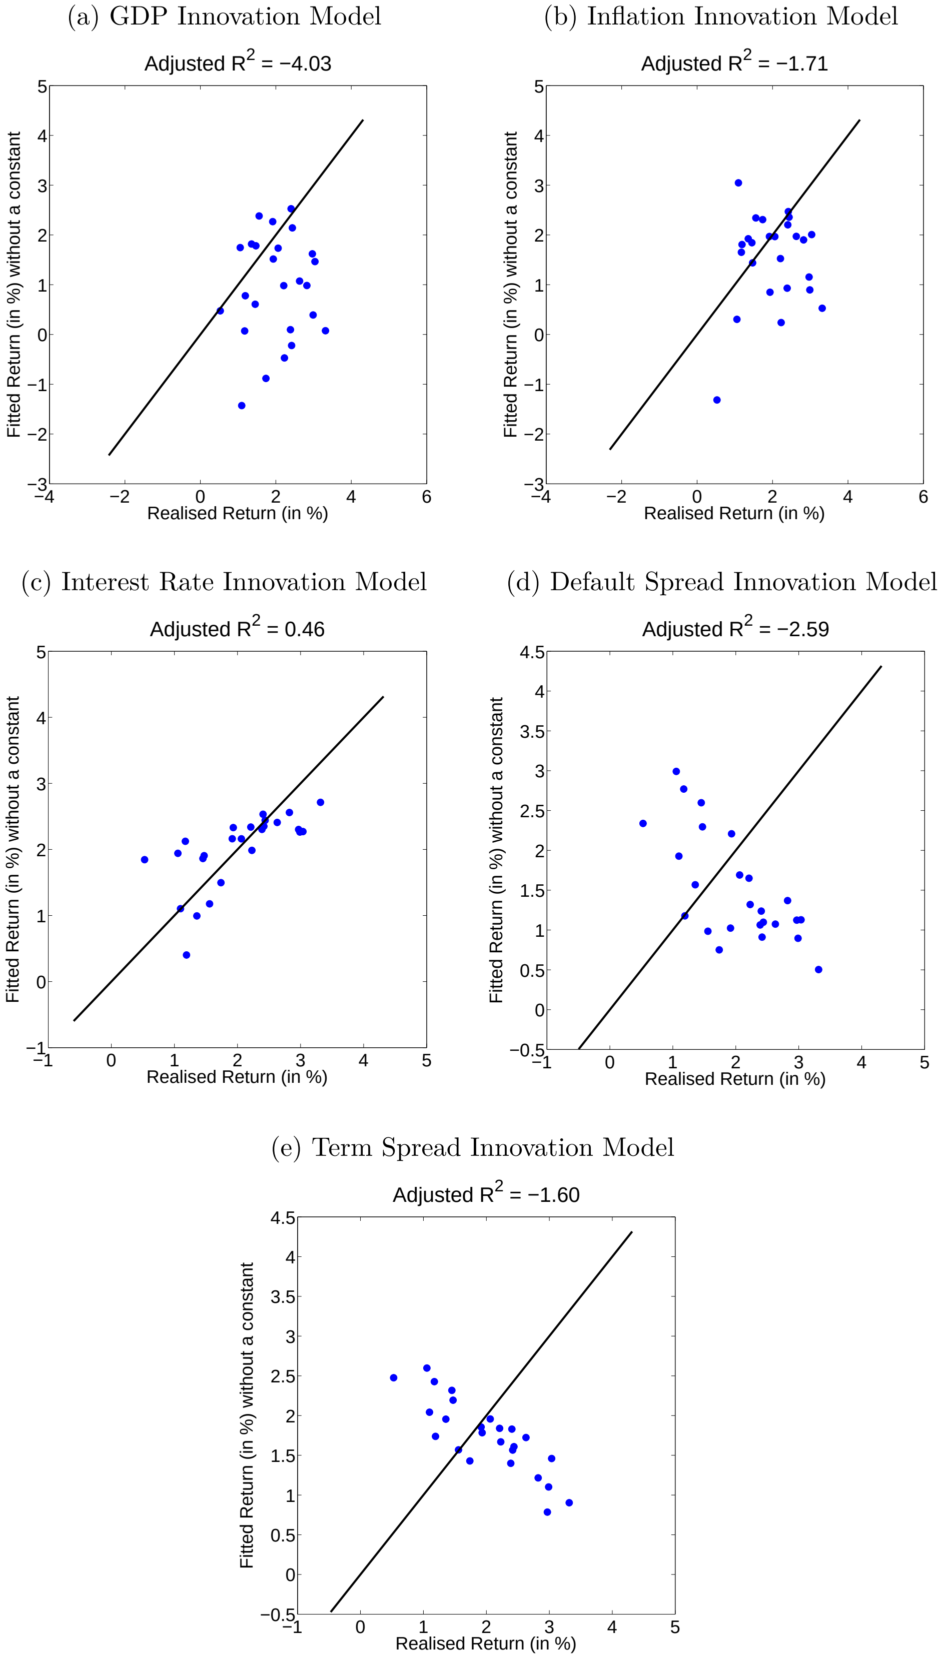

Notes: The R^2 values for the five 1-factor models are calculated assuming no constant in the second-pass regressions.

Figure 10: The Identification of the λ-shock: the Relationship between Price of Risk and Dispersion of Exposures

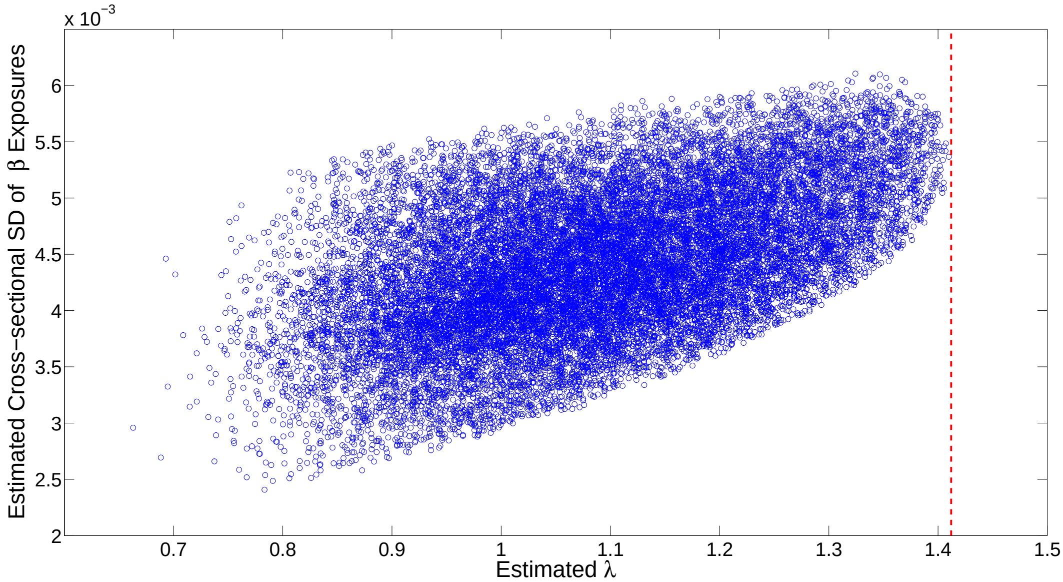

Notes: The scatter plot (based on 20,000 random $\tilde{Q}$ matrices) shows the relationship between the standard deviation of β-exposures to $\tilde{\varepsilon}_t$ associated with a given candidate draw $\tilde{Q}$ and the cross-sectional R^2 implied by the corresponding 1-factor model. For presentation purposes, I exclude those rotations that imply negative R^2 (about 48% of all admissible matrices), as it does not cause any loss of generality in the relationship. The vertical red dashed line is the maximum achievable price of risk (1.41) from the five-variable VAR model 2.1. To obtain these random draws, I apply Householder transformations to 20,000 five-dimensional matrices drawn from the multivariate Normal distribution.

Table 5: The Second-pass Regressions for Pricing the FF25+FF30 Portfolios: 1-factor Model vs. Fama-French 3-factor Model

Factor prices (λ)				Adj. R^2
Panel A: 1-factor Model with the λ-shock				
Constant	λ-shock			
0.72	0.60			0.179
(0.57) [0.66]	(0.23) [0.26]			
	0.97			0.115
	(0.36) [0.50]			
Panel B: The Fama-French 3-factor Model				
Constant	MKT	HML	SMB	
3.20	-1.78	0.91	0.48	0.336
(0.95) [0.99]	(1.11) [1.14]	(0.45) [0.45]	(0.44) [0.44]	
	1.43	0.92	0.35	0.098
	(0.63) [0.63]	(0.45) [0.45]	(0.44) [0.44]	

Notes: The table reports the cross-sectional regressions using the excess returns on the FF25+FF30 portfolios. The coefficients are expressed as percentage per quarter. Panel A presents results for the 1-factor model where the identified λ-shock is used as the sole pricing factor. Panel B presents results for the Fama-French 3-factor model. MKT is the market factor, HML is the book-to-market factor and SMB is the size factor. OLS standard errors are in parentheses, whereas standard errors, computed with the method of Shanken (1992) to adjust for errors-in-variables, are in brackets.

Figure 11: Impulse Responses to a λ-shock: Results from a Bayesian VAR(2)

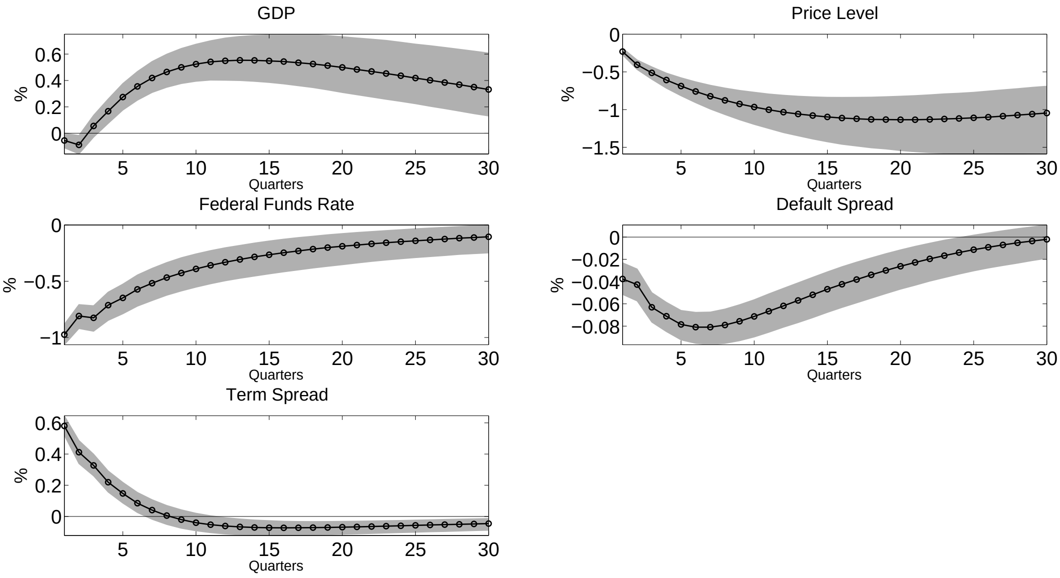

Notes: The sample period is 1963Q3 - 2008Q3. The Minnesota-type normal inverted Wishart priors are implemented following Banbura, Giannone, and Reichlin (2010). The figure shows the pointwise median and 16th-84th percentiles of $N_2 = 5000$ draws (after burning the first $N_1 = 5000$ draws) from the posterior distribution of the impulse responses. The vertical axes are %-deviations from steady-state, and the horizontal axes are in quarters.