The Macroeconomic Shock with the Highest Price of Risk

Gabor Pinter*

December 2, 2016

First version: August 10, 2016

Comments Welcome

Abstract

This paper proposes a new orthogonalisation method that uses the cross-section of returns to construct a macroeconomic shock. This λ-shock demands the highest possible risk price per unit of exposure, or equivalently, best approximates the stochastic discount factor with VAR residuals. When applying the method to the HML-SMB-industry portfolios, a robust feature of the λ-shock is the delayed effect on output and the sharp impact on the term spread and interest rates. The estimated λ-shock closely resembles (>70% correlation) monetary policy and technology news shocks studied by macroeconomists. In contrast, the λ-shock implied by momentum portfolios is markedly different.

JEL Classification*:* C32, G12

Keywords: Stochastic Discount Factor, Vector Autoregression, Shocks, Technology News, Monetary Policy, Cross-section of Returns, Momentum

*Bank of England; email: gabor.pinter@bankofengland.co.uk. I would like to thank Saleem Bahaj, Andy Blake, Martijn Boons, Svetlana Bryzgalova, 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, Adam Szeidl, Andrea Tamoni, Harald Uhlig, Stjin Van Nieuwerburgh, Garry Young and Shengxing Zhang for helpful comments. I would also like to thank Kenneth French and Amit Goyal for making their data available on their websites. The views expressed in this paper are those of the author, and not necessarily those of the Bank of England or its committees.

1 Introduction

> "We would like to understand the real, macroeconomic, aggregate, nondiversifiable risk that is proxied by the returns of the HML and SMB portfolios."
> (pp. 442 Cochrane (2005))

The literature is yet to find a compelling macroeconomic explanation behind the cross-sectional variation of asset returns. Part of this challenge is caused by the fact that innovations in macroeconomic variables are reduced-form objects: they are combinations of orthogonal structural shocks that may offset each other over the business cycle and demand possibly very different levels of risk premia. Using reduced-form variables such as output growth 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 the following econometric strategy in a vector autoregression (VAR) model: instead of starting with economic assumptions and testing their asset pricing implications, I start by using a given asset portfolio to construct an orthogonal shock that has the highest risk premium in absolute value when pricing the given portfolio. Equivalently, this shock is the best possible approximation of the stochastic discount factor (SDF) that can be recovered from the space of residuals of a given VAR model when pricing the given portfolio. Only then I check the macroeconomic characteristics of the resulting shock by inspecting the associated impulse response functions, forecast error variance decomposition and the estimated time-series of the shock.

The method is general and could be applied to any VAR and any test assets. When applying it to a simple five-variable macroeconomic VAR and to the 25 portfolios of Fama and French (1993) augmented with the 30 industry portfolios as prescribed by Lewellen, Nagel, and Shanken (2010) (FF55 henceforth), I find that the obtained shock, which I refer to as a λ-shock, closely resembles well-known structural shocks, studied by the macroeconomic literature. The shock triggers a delayed reaction in aggregate quantities such as GDP and consumption, and has a sharp impact on the short-term interest rate and the term spread. These features make the λ-shock similar to monetary policy shocks and also to what macroeconomists refer to as news shocks about future total factor productivity (TFP). In fact, the correlation between the λ-shock series and the *TFP news shock* series, estimated by Kurmann and Otrok (2013), and the *monetary policy shock* series, proposed by Romer and Romer (2004), are more than 70%. This is striking given that my orthogonalisation strategy, as explained further below, has nothing to do with the strategies used to identify monetary policy or TFP news shocks, as my VAR model does not even contain a measure of TFP as an observable. When applying the method to other test portfolios such as size-operating profitability or size-investment, I find that the economic properties of the λ-shock are largely unchanged.

In contrast, the obtained macroeconomic shock is markedly different when using the same VAR model but applying the orthogonalisation method to *momentum* portfolios. In this case, a positive λ-shock induces an immediate positive jump in output, which coincides with a sharp increase in the short-term interest rate. The shock can be interpreted as a strong, aggregate demand-type shock that is strongly counteracted by endogenous monetary policy reactions. Overall, these applications of my method suggest that a risk-based explanation of momentum could be based on structural macroeconomic forces that are very different from the structural shocks that are related to the "macroeconomic, aggregate, nondiversifiable risk proxied by the returns of the HML and SMB portfolios".

The Orthogonalisation Strategy The starting point of my analysis is a standard VAR including a small set of macroeconomic variables. The finance literature often used Cholesky decomposition to obtain triangularised innovations in the spirit of the Intertemporal CAPM (Merton, 1973)[1]. 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 point by exploring the entire space of possible orthogonalisations, given the estimated time-series of reduced-form residuals, with the aim to find the best approximation of the SDF from linear combinations of these residuals. Mechanically, the λ-shock is constructed as the one that, if used as a factor in the two-pass procedure of Fama and MacBeth (1973) applied to the given test portfolios, would generate the highest estimated factor risk premium in absolute value.

My approach does not make 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 (FEV) of a target variable in the VAR over a pre-specified horizon. Compared to these approaches, my method can be thought of as much more agnostic. Hence, there is no direct reason to believe that the obtained structural λ-shock should capture 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.

In addition, approximating the SDF with VAR residuals may have a possible advantage over standard no-arbitrage methods of estimating the SDF. The VAR framework and its rich machinery allows one to explore the link between the SDF and macroeconomic dynamics in more detail, making full use of the traditional macroeconometric toolkit. Impulse response function (IRF) analysis can be used to estimate how the λ-shock propagates through the economy in comparison with structural shocks traditionally identified in the macroeconomic literature. FEV decomposition can be used to estimate the con-

[1] See Campbell (1996), Petkova (2006) and Boons (2016) amongst others.

tribution to business cycle dynamics of shocks that do not demand risk compensation, according to the given test portfolios, compared to shocks that do (λ-shock). These are just two of the examples of how the proposed framework can potentially provide a better understanding of the links between asset prices and business cycles.

Related Literature My paper is related to the finance literature on 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), Boons and Tamoni (2015), He, Kelly, and Manela (2016). In addition, consumption based asset pricing (CCAPM) models also had success in explaining the cross-section of returns by introducing conditioning variables ((Jagannathan and Wang, 1996; Lettau and Ludvigson, 2001; Lustig and Nieuwerburgh, 2005; Santos and Veronesi, 2006; Yogo, 2006)) or 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; Bryzgalova and Julliard 2015).

A number of recent papers explored factors that are less reduced-form and are more tied to macroeconomic primitives. Modern macroeconomic models interpret business cycles as the outcome of 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); Campbell, Pflueger, and Viceira (2015); Greenwald, Lettau, and Ludvigson (2015); Ludvigson, Ma, and Ng (2015); Kliem and Uhlig (2016)). In this spirit, more recent explanations of the cross-sectional variation of returns involve macroeconomic surprises related to monetary policy (Weber 2015; Ozdagli and Velikov 2016) and production technology (Papanikolaou 2011; Kogan and Papanikolaou 2014; Garlappi and Song 2016) among others. My paper builds on these developments, and the results from applying my orthogonalisation strategy to the FF55 portfolios are consistent with the empirical findings of these two literatures.

Further, the method I propose builds heavily on the structural VAR literature (Sims 1980; Stock and Watson 2001). More specifically, the implementation of my orthogonalisation theme draws on the more recent identification themes that use sign restrictions to identify structural shocks (Uhlig 2005; Rubio-Ramirez, Waggoner, and Zha 2010; Fry and Pagan 2011). As mentioned, finance papers using VARs (Campbell, 1996; Petkova, 2006; Boons, 2016) typically applied Cholesky decomposition to the estimated reduced-form variance covariance matrix. While the obtained innovations had success in explaining the cross-section of returns, it has been difficult to assign macroeconomic interpretations to these innovations. Moreover, the idea of using observed asset prices to select a structural shock draws on the long-standing literature of no-arbitrage estimation of the SDF (Hansen and Singleton 1982; Ait-Sahalia and Lo 2000; Rosenberg and Engle 2002; Chernov 2003; Ross 2015; Ghosh, Julliard, and Taylor 2016). Building on these papers, my method to explore the entire space of possible orthogonalisations in a VAR and to find the structural shock based on approximating the SDF given a set of portfolios is, to the best of my knowledge, novel in the literature.

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

2 The Econometric Framework

2.1 The Geometry of the λ-shock

Before presenting the VAR model, it is instructive to first summarise the intuition behind finding the λ-shock. To do so and to highlight the geometrical nature of the ideas, I try to map some of the relevant mathematical background into a simplified 3-dimensional graph shown in Figure 1. There is an underlying probability space, and L_2 denotes the collection of all random variables with finite variances defined on that space. L_2 is a Hilbert-space with the associated norm $\|p\| = \left(E\left(p^2\right)\right)^{1/2}$ for $p \in L_2$. Let P denote the space of portfolio excess returns (zero-price payoffs) that is assumed to be a closed linear subspace of L^2.[2] P is represented by the red plane in Figure 1. An admissible stochastic discount factor is a random variable m in L^2 such that the inner product of the excess return and m satisfies $0 = E\left(pm\right)$ for all $p \in P$. The set of all admissible SDFs denoted by M is represented by the black line which goes through the origin and perpendicular to the red plane in Figure 1.[3]

Let S denote the set of reduced-form innovations from a VAR (the blue solid arrows in the Figure) and denote D the space spanned by these innovations. D is assumed to be a closed subspace of L_2, and it is represented by the blue plane in the Figure. The Gram-Schmidt orthogonalisation procedure allows the reduced-form innovations that span D to be transformed into a set of orthonormal vectors that also span D. The blue dashed arrows in Figure 1 represent two possible elements of the infinite sequence of orthogonalisations. The set of the all admissible orthogonalisations is denoted by O and is represented by the blue circle with unit radius in the Figure.

The space of VAR innovations is unlikely to contain an SDF because of model misspecification or measurement error associated with observing SDFs (Roll (1977)). Loosely

[2]See Hansen and Jagannathan (1991, 1997) for a detailed discussion.

[3]As is well known, all SDFs can be represented as the sum of the minimum norm SDFs (the intersection of the black line and the red plane in Figure 1) and of a random variable that is orthogonal to the space P of excess returns (Hansen and Richard (1987); Cochrane (2005)).

Figure 1: A Simplified Geometry of Finding the λ-shock

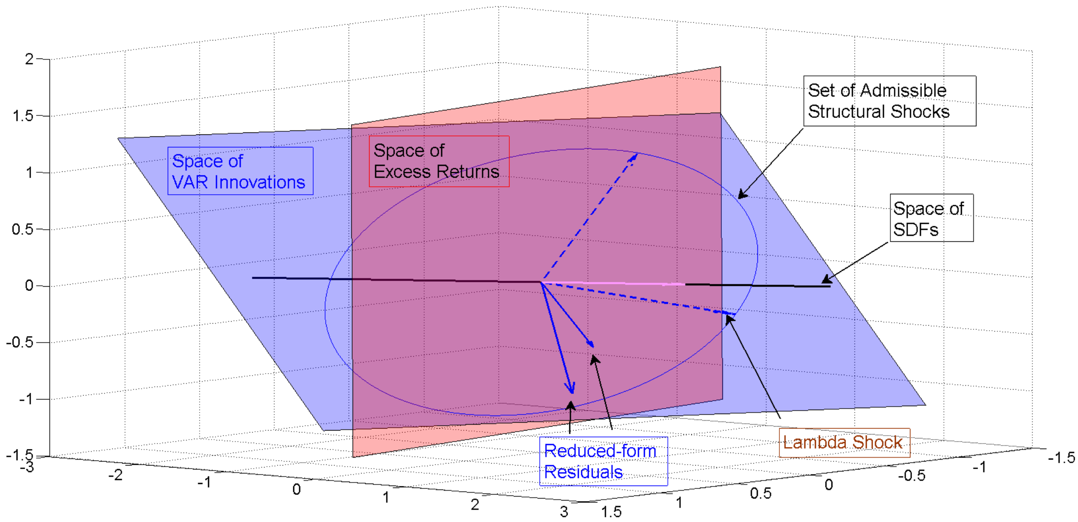

speaking, the tilted nature of the blue plane prevents all elements of O to be orthogonal to the space of excess returns, i.e. $M \cap O = \emptyset$. Nevertheless, one can find an element in O that is closest to M in the spirit of Hansen and Jagannathan (1997) by applying the classical Projection Theorem.[4] This has important implications for linear models of the SDF that use structural innovations from VAR models as pricing factors: there is *one* particular orthogonalisation of the reduced-form VAR residuals that delivers a structural shock, which is closer to the SDF than all the other structural shocks in the VAR. This is the blue arrow labelled as the λ-shock in Figure 1, whose projection onto the space of SDFs is the magenta line. Given that this shock is the best possible approximation of the SDF, it summarises all the relevant information contained in all the reduced-form residuals of the VAR model. The next proposition for the two-dimensional case highlights that it is in fact easy to find the rotation which generates the λ-shock.

Proposition 1 *Given the linear combination:* $m = af_1 + bf_2$*, where* $a, b \in \mathbb{R}$*,* $m, f_1, f_2 \in \mathbb{R}^2$*,* $\|f_1\| = \|f_2\| = 1$ *and* $\langle f_1, f_2 \rangle = 0$*, there exists a rotation* $r_\theta = \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix}$ *with* $0 < \theta < 2\pi$ *such that* $m = a^\star f_1^\star + b^\star f_2^\star$*, where* $a^\star \neq 0$*,* $b^\star = 0$ *and* $f_i^\star = r_\theta f_i$ *for* $i = 1, 2$*.*

Proof. It suffices to find an angle $\theta^\star$ and associated rotation $r_{\theta^\star}$ such that m will be a scaled multiple of any one of the rotated vectors denoted by $f_1^\star$. If $\theta^\star$ exists then $b^\star = 0$ because $f_1 \perp f_2$ and $r_{\theta^\star}$ is an orthonormal transformation. The angle $\theta^\star = \arctan\left(\frac{b}{a}\right)$

[4]That is, assuming that O is a complete linear subspace of H, there exists a unique vector $m_0 \in O$, corresponding to any vector $x \in M$, such that $\|x - m_0\| \leq \|x - m\|$ for all $m \in O$. See pp. 50-51 of Luenberger (1969) for a classic treatment and pp. 608-609 of Hansen and Richard (1987) for a conditional version of the theorem.

satisfies $f_1^\star = r_{\theta^\star} f_1$ so that $m = a^\star f_1^\star + b^\star f_2^\star$ with the associated scalars $a^\star = \frac{\|m\|}{\|f_1^\star\|}$ and $b^\star = 0$. ■

Extending proposition 1 to higher dimensions is straightforward. While the proposition may seem a trivial piece of linear algebra, it has important implications for using orthonormalised shocks from VAR models as pricing factors in linear pricing models. It is a well known theorem that beta pricing models are equivalent to linear models for the SDF (pp. 106-107 Cochrane (2005)). Denoting the SDF, the pricing factor, the excess returns and the first- and second-stage regression coefficients from a linear pricing model by m, f, R^e, β and λ, respectively, I re-state the version of the theorem when the test assets are all excess returns:

Theorem 2 (Cochrane 2005) *Given the model*

$$\begin{aligned} m &= 1 + [f - \mathbb{E}(f)]' b \\ 0 &= \mathbb{E}(mR^e), \end{aligned} \tag{2.1}$$

one can find λ such that

$$\mathbb{E}(R^e) = \beta'\lambda, \tag{2.2}$$

where β are the multiple regression coefficients of excess returns R^e on the factors. Conversely, given λ in 2.2, we can find b such that 2.1 holds.

It is shown by Cochrane (2005) that λ and b are related $\lambda = -var(f)\, b$. This result simplifies greatly when working with pricing factors (such as orthonormalised VAR residuals) that have zero mean and unit variance. In this case, $\lambda = -b$ and $\mathbb{E}(f) = 0$. The following example highlights that finding the orthonormalised shock in a VAR model that demands the highest price of risk when pricing a given portfolio of assets is equivalent to finding the shock that summarises all the information (contained in the residuals of the given VAR model) relevant to pricing the given portfolio.[5]

Example 3 (A Two-variable VAR Model) *Suppose the pricing factors are arbitrarily rotated orthonormalised residuals from a two-variable VAR model. Given a linear pricing model 2.2, the model for the SDF (m) is written as a linear combination of the two innovation series:*

$$m = a + \lambda_1 f_1 + \lambda_2 f_2, \tag{2.3}$$

where a is a constant and λ_1 and λ_2 are the estimated prices of risk associated with f_1 and f_2. Because $f_1 \perp f_2$ and $var(f_1) = var(f_2) = 1$, the variance of the SDF is simply

[5] Another way of saying this is that the cross-sectional R^2-measure associated with a pricing model that includes all the reduced-form residuals from the VAR is the same as the R^2-measure associated with the one-factor model which uses the appropriately orthonormalised shock. This will be confirmed during the empirical application of the method (Panel A and B of Tables 3–8.)

the sum of the squared values of the estimates of prices of risk associated with each one of the two VAR shock series:

$$var\,(m) = \lambda_1^2 + \lambda_2^2. \tag{2.4}$$

Rotation does not affect the overall information content in the VAR, that is, the volatility of the implied SDF is determined by the specification of the VAR and not by rotating the variance-covariance matrix of the residuals. The main implication of proposition 1 is that the information contained in the VAR residuals can be summarised by merely one structural shock after applying an appropriate rotation to the variance-covariance matrix. To put it simply, there exists a rotation $r_\theta = \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix}$ *such that using* $f_i^\star = r_\theta f_i$ *for* $i = 1, 2$ *as pricing factors, one of the estimated prices of risk would be non-zero* $\lambda_1^\star \neq 0$ *and in fact it would be* $\lambda_1^\star = \sqrt{var\,(m)}$*, as the other one corresponding to the rotated factor* $f_2^\star$ *would be zero* $\lambda_2^\star = 0$*. This implies that the best approximation of the SDF is found,* $f_1^\star = m$*, and* r_θ *can be used to perform structural analysis in the VAR.*

The next subsection will show how to implement the method in practice.

2.2 The VAR Model

To implement the ideas above, I follow the macro-finance literature (Campbell and Shiller (1988); Campbell, Giglio, and Polk (2013)) by using a linear reduced-form VAR model to describe the evolution of the macroeconomic state:

$$y_t = c + \sum_{j=1}^{p} B_j y_{t-j} + u_t, \tag{2.5}$$

where y_t is an $n \times 1$ vector of observed endogenous variables, c is an $n \times 1$ vector of constants, p denotes the number of lags, B_j 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. $\hat{\Sigma} = \left(Q\tilde{A}_0\right)' Q\tilde{A}_0$.[6]

[6] 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)) etc.

The key step in finding the λ-shock is the following: 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 given test 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)):

$$\mathbb{E}\left(R_t^e\right) = \beta\left(\varepsilon_t^\star\right) \times \lambda\left(\varepsilon_t^\star\right), \tag{2.6}$$

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 find. 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 $\tilde{Q}$, the first stage is an OLS estimation of the time-series regression of each of the k portfolios' excess returns on the implied candidate structural shock $\tilde{\varepsilon}_t$:

$$R_{it}^e = a_i + \tilde{\varepsilon}_t \beta_i + \epsilon_{it}, \tag{2.7}$$

where β_i represents the ith element in β. Given 2.7, the second stage is a cross-section regression of average portfolio returns on the estimated betas associated with the candidate matrix $\tilde{Q}$:

$$\bar{R}_i^e = \tilde{\beta}_i \times \lambda + \alpha_i, \tag{2.8}$$

where $\bar{R}_i^e = \frac{1}{T}\sum_{t=1}^{T} R_{it}^e$, 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 $\varepsilon_t^\star$ which will generate the highest estimated λ in absolute value in 2.8. Finding $\varepsilon_t^\star$ 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.[7]

It is worth reiterating 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. As the previous subsection showed, it is straightforward to find a single shock that is the best candidate for the SDF. And it will summarise all the relevant

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

information, contained in the reduced-form VAR innovations, that is relevant to pricing the given cross-section of portfolios.

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 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). 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. 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 baseline VAR model includes two lags as suggested by the Schwarz Information criterion.

The sample period for the estimation is 1963Q3-2008Q3 and the data are at quarterly frequency. The start of the estimation period is chosen by 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 FF55 portfolios, 25 of them (FF25 henceforth) are formed from independent sorts of stocks into five size groups and five B/M groups as described in Fama and French (1993). The other 30 portfolios are four-digit SIC code level industry portfolios. 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). Augmenting the FF25 with the 30 industry portfolios follows prescription 1 (pp. 182) of Lewellen, Nagel, and Shanken

(2010), thereby relaxing the tight factor structure of Size- B/M. Moreover, I apply the orthogonalisation method to additional test portfolios such as the 25 portfolios sorted on size and operating profitability, the 25 portfolios sorted on size and investment, and the 10 momentum portfolios sorted on the cumulative returns of stocks from 12 months before to one month before the formation date using a one-month gap before the holding period. All the portfolios are value-weighted and are taken from Ken French's data library. As an alternative, I also use the 10 momentum portfolios as constructed in Daniel and Moskowitz (2016).[8]

3.2 The Economic Characteristics of the λ-shock

Using the OLS estimates of the VAR, I compute impulse response functions (IRF) after performing the orthogonalisation strategy described in Section 2. This is to understand the macroeconomic impact of the λ-shock which is by construction the structural shock that best approximates the SDF given the 5-variable VAR(2) model and the FF55 portfolios. Section B of the Appendix describes a Bayesian treatment of the computation of the IRFs in order to explore the role of parameter uncertainty in the VAR model.

Figure 2: Impulse Responses to a λ-shock

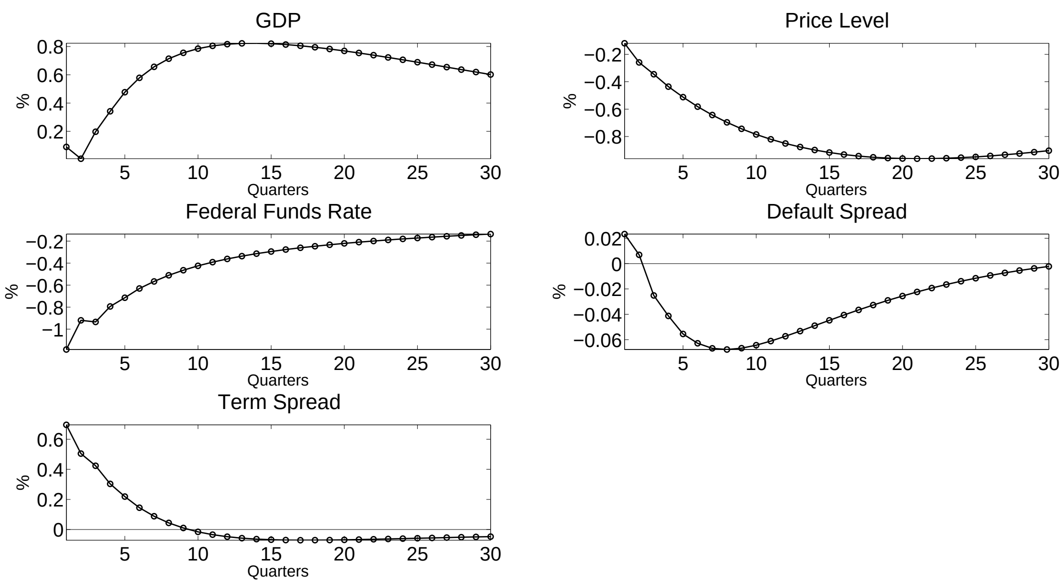

Notes: The vertical axes are %-deviations from steady-state, and the horizontal axes are in quarters.

Figure 2 displays the IRFs of the five variables to a one standard deviation structural innovation. The term spread jumps by about 70 basis points on impact and there is a

[8] As explained in Daniel and Moskowitz (2016) the biggest difference is that the portfolio breakpoints for the portfolios constructed by Ken French are set so that each of the portfolios has an equal number of NYSE firms. In contrast, Daniel and Moskowitz (2016) set their breakpoints so that there are an equal number of firms in each portfolio. This mainly affects the low-momentum returns.

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 from the third quarter onwards and reaching a peak impact of about 0.8% approximately 12-15 quarters after the shock hits. As shown by Figure 7 in the Appendix, the shape of these IRFs is similar when the lag length is changed or when output is replaced by consumption in the VAR. In addition, Section D of the Appendix illustrates the equivalence between two algorithms to construct the λ-shock: maximising the price of risk and maximising the cross-sectional R^2.

To assess the contribution of the λ-shock to business cycles, in comparison with other structural shocks that have zero covariance with the implied SDF, I compute FEV decomposition over different horizons and for different lag structures of the VAR. Table 1 presents the results for VAR(2), VAR(3) and VAR(4) models using the FF55 portfolios as test assets. The results suggest that the λ-shock explains less than 10% of output fluctuations over the one-year horizon, but the shock explain around 40-75% of fluctuations over longer (4-9 years) horizon. While these number are substantial, there is some unexplained fraction of output fluctuations that is driven by structural shocks exposures to which do not demand risk compensation according to the given test assets. Moreover,

Table 1: The Contribution of the λ-shock to Business Cycles:
Results from Forecast Error Variance Decomposition

	Output	CPI	FFR	Def. Spread	Term Spread
			VAR(2)		
4Q	8.5	18.3	72.8	5.0	61.2
8Q	40.5	24.7	70.5	20.0	59.2
16Q	57.9	28.6	69.8	27.6	58.6
24Q	72.3	33.3	69.2	31.9	58.3
36Q	75.3	34.1	69.2	31.8	58.3
			VAR(3)		
4Q	5.0	15.9	65.5	5.0	57.4
8Q	25.1	16.5	55.9	19.2	50.2
16Q	40.5	16.3	51.9	25.7	49.1
24Q	51.2	15.6	48.5	26.3	48.8
36Q	54.3	14.9	48.0	26.1	48.2
			VAR(4)		
4Q	3.8	16.6	71.2	5.5	67.8
8Q	26.4	16.3	63.6	24.1	62.9
16Q	43.8	14.4	58.6	33.4	62.3
24Q	51.6	10.3	52.6	31.5	61.4
36Q	53.4	8.5	51.5	31.4	60.4

Notes: The table shows the % fraction of the total forecast error variance that is explained by the λ-shock over different forecast horizons. The FF55 portfolios are used as test portfolios for each VAR model.

Table 1 also shows that the λ-shock drives around 50-70% of interest rate and term spread fluctuations and around 10-30% of fluctuations in the aggregate price level and the default spread. The explained variation in the FEV in the interest rate and the term spread seems to decrease over the forecast horizon, whereas it increases for output and the default spread.

To place these results in the literature, 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 2 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.[9] However, a news-type technology shock that typically triggers a delayed reaction in aggregate quantities may be perfectly consistent with Figure 2.[10] 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 2 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.

[9]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 orthogonalisation theme adopted in this paper.

[10]A partial list of the rapidly increasing macroeconomic literature on 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), Malkhozov and Tamoni (2015).

Figure 3: Comparing the λ-shock to the TFP News Shock Series of Kurmann and Otrok (2013) (Correlation: 78%) and to the Monetary Policy Shock Series of Romer and Romer (2004) (Correlation: 73%).

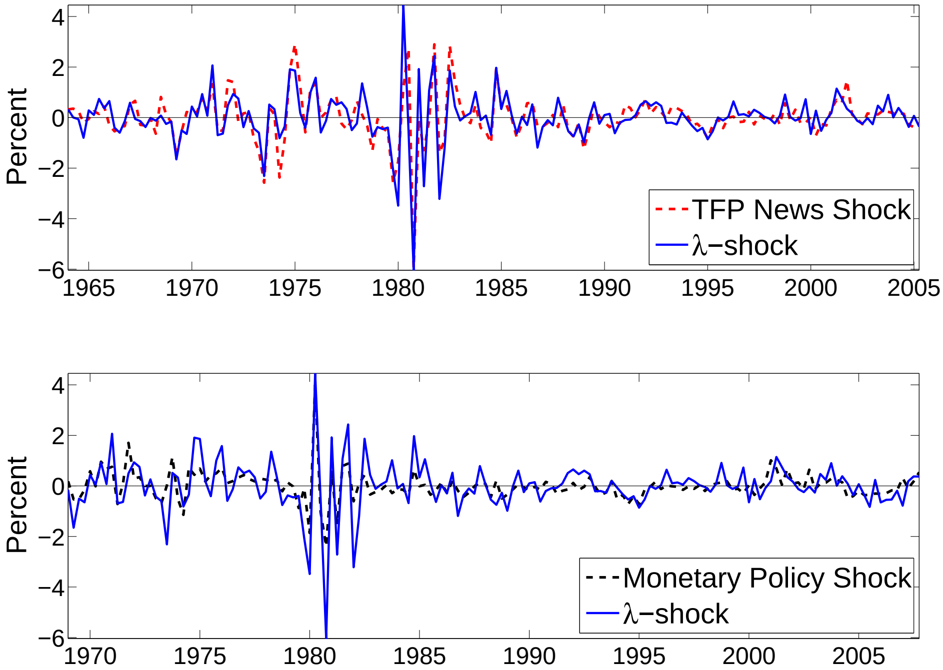

Notes: The TFP news shock series are the ones plotted in Figure 5 on pp. 2625 of Kurmann and Otrok (2013) who apply the method of Uhlig (2004) to identify a TFP news shock over the period 1959Q2-2005Q2. The monetary policy shock series are originally developed by Romer and Romer (2004) and updated by Tenreyro and Thwaites (2016) to the period 1969Q1-2007Q4.

An alternative interpretation of Figure 2 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 and some well-known structural shocks studied by macroeconomists, Figure 3 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 identified 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 0.78. Based on the overlapping estimation period 1969Q1–2007Q4, the correlation coefficient between the monetary policy shock series (black dashed line) as identified in Romer and Romer (2004) (and updated by

Tenreyro and Thwaites (2016)) and the λ-shock series is 0.73.

Table 2: Robustness of the λ-shock to Changing the VAR Model

Correlation Coefficients						
λ-shock Baseline	λ-shock VAR(3)	λ-shock VAR – C	λ-shock VAR – IP	λ-shock VAR – CPI	λ-shock VAR – Spr	External Shocks
1.00						
0.93	1.00					
0.95	0.91	1.00				
0.91	0.84	0.91	1.00			
0.98	0.92	0.95	0.90	1.00		
0.92	0.79	0.81	0.84	0.88	1.00	
0.78	0.69	0.80	0.77	0.78	0.67	TFP News
0.73	0.76	0.79	0.74	0.74	0.64	Monetary Policy

Notes: The table reports the correlation coefficients among λ-shocks from the baseline (Column 1), the baseline VAR with 3 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). In all cases, the FF55 portfolios were used as test assets.

To check whether these results are robust to changing the specification of the VAR model, Table 2 shows the cross-correlations among the various λ-shock series the TFP news 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. Overall, I find that changing the specification of the VAR does not have a material impact on the results. Though it may be worth noting that replacing GDP with aggregate consumption increases the correlation of the λ-shock with TFP news and monetary policy shocks up to around 0.80.

To reiterate, my orthogonalisation 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[11], (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.

[11] 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.

3.3 The λ-shock Implied by Other Test Portfolios

To check the robustness of the findings above, I explore how the behaviour of the λ-shock changes when the same VAR model and the orthogonalisation method are applied to other test assets. A natural choice is the 25 portfolios double sorted on size-profitability and size-investment. These portfolios feature prominently in the most recent empirical asset pricing studies (Fama and French, 2015, 2016), because the standard 3-factor model of Fama and French (1993) turned out to miss much of the variation in average returns related to profitability and investment. In addition, I also compute the IRFs for the λ-shock implied by the benchmark FF25 portfolios, sorted on size-B/M, that have been the most studied test assets to date.

Figure 4: Impulse Responses to a λ-shock, Implied by *FF25, 25 Profitability-Size and 25 Investment-Size* Portfolios

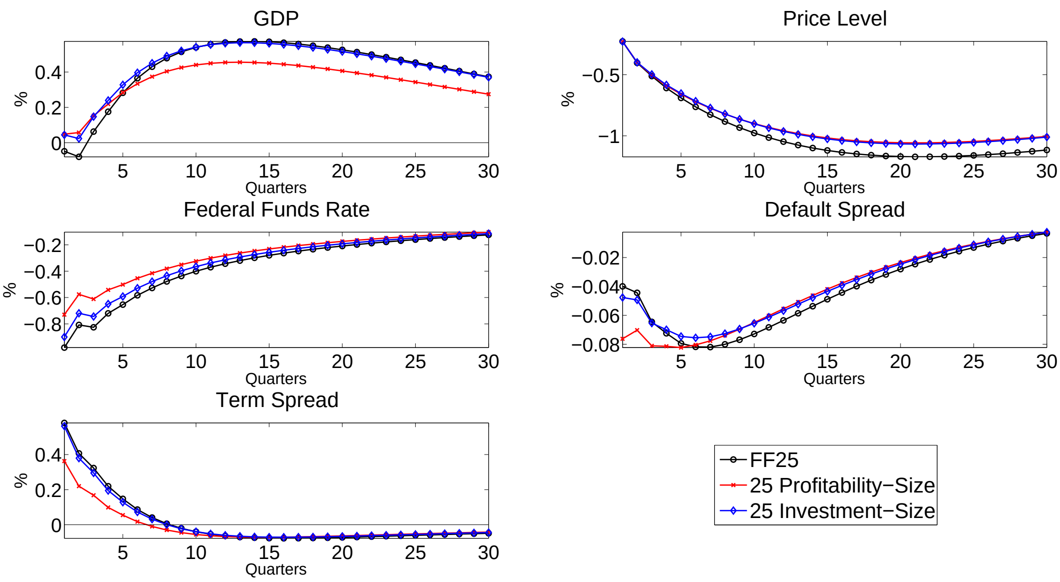

Notes: The vertical axes are %-deviations from steady-state, and the horizontal axes are in quarters.

Figure 4 shows the IRFs for all three sets of test portfolios. The results suggest that the economic behaviour of the λ-shock implied by these portfolios is very similar to the λ-shock implied by the baseline FF55 as shown in Figure 2. The only quantitative difference is that the baseline results imply a larger peak effect on output and a more delayed effect on the default spread compared to Figure 4. Overall, these results suggest that the average returns of the equity portfolios studied so far capture approximately the same source of macroeconomic risk. As shown in the next subsection, these results change markedly when using momentum portfolios to construct the λ-shock.

3.4 Momentum

Since it was first documented (Asness, 1994; Jegadeesh and Titman, 1993), momentum returns have been challenging to explain with pricing factors that worked well in pricing the traditional Fama-French portfolios. As a result, many linear factors models since Carhart (1997) included a momentum factor explicitly in their pricing models in order to explain momentum. Even the most recent generation of pricing models such as the five-factor model of Fama and French (2015, 2016) fail badly as descriptions of average returns on momentum portfolios without including a momentum factor in their model. This is particularly puzzling given that momentum is a pervasive phenomenon that appears in many diverse markets and asset classes (Asness, Moskowitz, and Pedersen, 2013).

To explore the potentially different structural macroeconomic risks underlying momentum portfolios, I apply the same VAR and orthogonalisation technique to the 10 momentum (Prior 2-12) portfolios constructed by Ken French and to the 10 momentum portfolios used in Daniel and Moskowitz (2016). Figure 5 shows the impact of a one standard deviation λ-shock implied by the momentum portfolios compared with the λ-shock implied by the FF55 portfolios. The results suggest that the λ-shock implied by momentum has a markedly different dynamic effect on the economy. Output jumps on impact by about 0.4% which is strongly counteracted by endogenous monetary policy raising the short-term interest rate by 80-90 basis points on impact thereby putting downward pressure on the term-spread. The sudden economic expansion and the strongly countercyclical monetary policy response result in a small net effect on CPI and the default spread.

Figure 5: Impulse Responses to a λ-shock, Implied by *Momentum* Portfolios

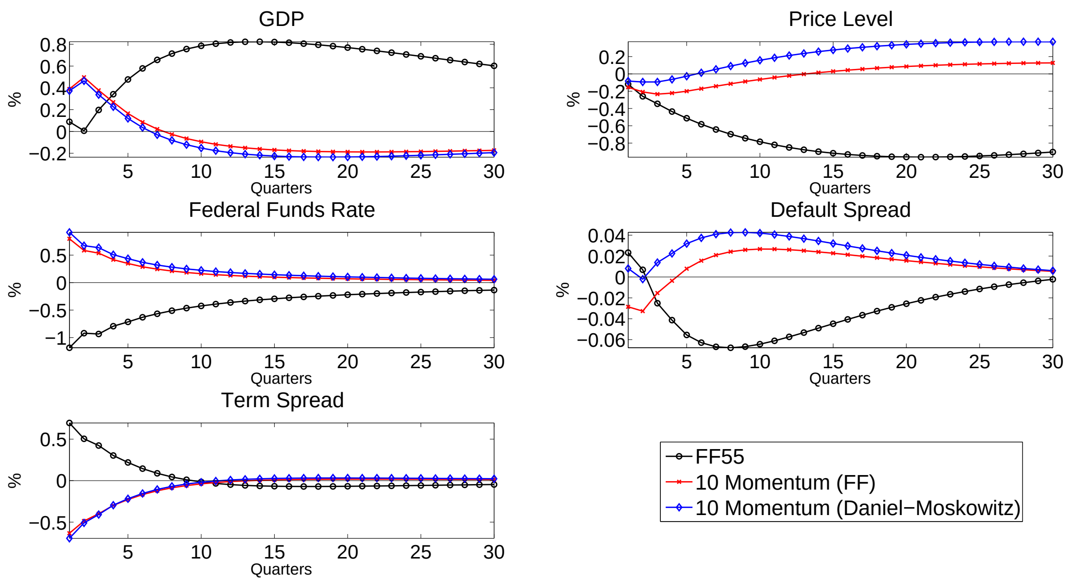

Notes: The vertical axes are %-deviations from steady-state, and the horizontal axes are in quarters.

Overall, these results suggest that the λ-shock implied by momentum resembles an aggregate demand-type shock that induces a much more rapid but more transient reaction in economic activity. This is in sharp contrast with the λ-shock implied by the traditional FF55 portfolios. These findings are consistent with the rational explanation of momentum suggesting that winners are more exposed to economic growth risks than the losers (Johnson, 2002). For example, Liu and Zhang (2008) shows that winners have temporarily higher loadings than losers on the growth rate of industrial production. While the growth rate of industrial production is a reduced-form object and can be driven by various structural shocks, Figure 5 confirms that the λ-shock implied by momentum indeed induces a large contemporaneous impact on the growth rate of output.

Moreover, it is well known that momentum strategies and more traditional trading strategies such as value are negatively correlated.[12] In the present framework this would suggest that the time-series of the macroeconomic shock exposures to which these two strategies have delivered large risk premia historically must be quite distinct. Indeed, the correlation between the time-series of the λ-shock implied by the FF55 portfolios and of the λ-shock implied by the 10 momentum portfolios of Ken French and Daniel and Moskowitz (2016) is -0.53 and -0.40, respectively.

3.5 Pricing the Cross-section of Stock Returns

It is worth emphasising that the focus of this paper is not the asset pricing performance of the λ-shock. Conditional on the VAR model 2.5 being an accurate representation of the economy, the λ-shock itself is the SDF by construction, so there is no longer any sampling or model uncertainty surrounding the implied linear pricing model. Put it differently, the pricing performance of the given λ-shock can easily be improved by changing the specification of the VAR (e.g. including additional variables such as valuation ratios[13]) but not by changing the orthogonalisation assumption. Naturally, the simple five-variable macroeconomic VAR model is not an accurate representation of the economy, but it is an entirely standard and parsimonious way of summarising macroeconomic dynamics which also mitigates potential problems related to fishing bias.

Nevertheless, for the interested reader, I summarise in this subsection the asset-pricing performance of the λ-shock implied by each test portfolios studied above. As argued above, this only is a test as to whether the variables included in the VAR contain information relevant to pricing the given portfolios. Tables 3–8 of the Appendix present the results from the two-pass regression technique of Fama and MacBeth (1973). During this exercise, I treat the uncovered λ-shock as a known factor when estimating the two-pass regression model 2.7–2.8. To estimate the risk premium associated with the λ-shock, I

[12] For example, Table I of Asness, Moskowitz, and Pedersen (2013) shows that value and momentum strategies in the US have had an average correlation about -0.60 over the period 1972/1-2011/7.

[13] These results are available upon request.

apply the GMM procedure described in Cochrane (2005) and implemented by Burnside (2011).

Overall, the pricing performance of the VAR (or equivalently, the λ-shock) is comparable with the 3-factor model of Fama and French (1993).[14] Moreover, as explained in Section 2, finding the λ-shock implies that the other four structural shocks have zero covariance with the implied SDF, and therefore the associated estimated prices of risk are numerically zero, as shown in panel B of Tables 3–8. Relatedly, the R^2 associated with the one-factor model using the λ-shock is identical to the R^2 for the model using any set of five orthogonalised shocks or in fact the model which uses the five reduced-form VAR residuals.

Moreover, the results are also consistent with Lewellen, Nagel, and Shanken (2010) who pointed out the strong factor structure of the FF25 portfolios which makes it relatively easy to find factors that generate high cross-sectional R^2s. Hence, they prescribed to augment the FF25 with the 30 industry portfolios of Fama-French in order to relax the tight factor structure of the FF25. Indeed, the cross-sectional R^2 drops drastically from 0.84 to 0.13 for the 1-factor model without a common constant, and it drops from 0.76 to 0.13 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 FF55 portfolios than for pricing the FF25 portfolios. This may of course lead to a critique of the (lack of) relevant information content of the VAR for pricing the FF55 portfolios, which may call for enriching the information set by adding valuation ratios to the VAR. Nevertheless, changing the VAR may be unnecessary because this poor pricing performance is unlikely to undermine the results of this application of my orthogonalisation strategy: the macroeconomic shock that captures all relevant information for pricing the cross section (irrespective of whether the information content is relatively small or large) bears virtually the same economic characteristics as the λ-shock using the FF25 portfolios. The IRFs are similar for the λ-shock using the FF25 and the FF55 (Figures 2 and 4), and the time-series of the shocks implied by the two portfolios have a high (0.81) correlation coefficient.

3.6 The λ-shock and the Fundamentals

An application of my proposed orthogonalisation strategy to the stock portfolios of FF55 led to the result that the estimated λ-shock bears a close empirical relationship both with TFP news shocks and with monetary policy shocks. This ambiguity of the result might seem an awkward outcome: after all, how can the resulting λ-shock have such a high correlation with two, seemingly distinct structural disturbances? To convince the

[14] Applying the 3-factor model to the FF25 portfolios (Table 4) yields similar results to those obtained in the literature (e.g. Petkova (2006)).

reader that this is not a fault of my orthogonalisation 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 6 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 6: 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.

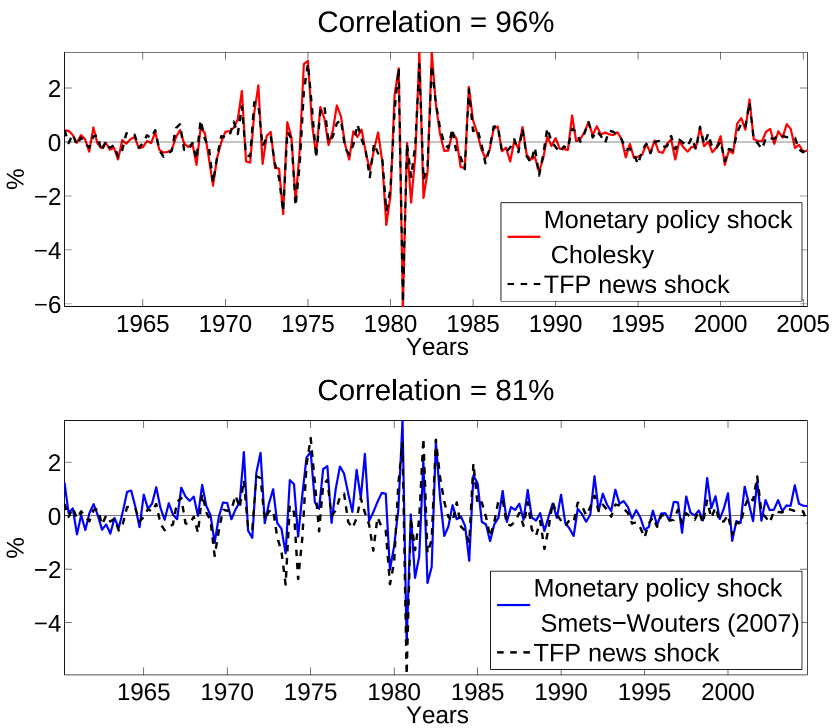

Notes: The TFP news shock series (black dashed line) are the ones plotted in Figure 5 on pp. 2625 of Kurmann and Otrok (2013) who apply the method of Uhlig (2004) to identify a TFP news shock over the period 1959Q2-2005Q2. The monetary policy shock series in the upper panel (red solid line) are identified 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 6 (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 orthogonalisation 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 regularity has not been documented in the literature yet, and it could be subject to further research (Pinter (2016)).

4 Conclusion

This paper proposed a new orthogonalisation theme in a VAR framework based on the ability of the obtained shock to explain the cross section of asset returns. The orthogonalisation theme is motivated by the long-standing challenge to link the origins of the cross-sectional variation in stock returns to macroeconomic primitives. When applying the method to the FF55 portfolios, the obtained shock is found to exhibit meaningful economic characteristics, closely resembling well-known structural shocks studied by the macroeconomic literature. These results have some direct implications for business cycle and asset price dynamics. The structural shock that is responsible for the aggregate risk captured by the FF55 portfolios is related to aggregate shocks that tend to generate a delayed response in aggregate quantities. In contrast, when applying the method to momentum portfolios, the implied λ-shock seems more related to unanticipated shocks that tend to generate immediate jumps in aggregate quantities. These results are consistent with the recent macroeconomic literature (Schmitt-Grohe and Uribe, 2012) that emphasise the role of both anticipated and unanticipated shocks as sources of business cycles.

More generally, the method I propose is not restricted to equity portfolios and could easily be used to study the macroeconomic forces behind aggregate risks underlying portfolios in other asset classes and markets. This could potentially help bridge some of the gap between the macroeconomic and the financial market anomalies literatures (Harvey, Liu, and Zhu, 2015; Bryzgalova, 2015; Fama and French, 2016; Novy-Marx and Velikov, 2016). Equally, the simple linear VAR framework could easily be extended to incorporate time-varying parameters, regime-switching, stochastic volatility and other forms of non-linearities. These extensions could be interesting avenues for future research.

References

Ait-Sahalia, Y., and A. W. Lo (2000): "Nonparametric risk management and implied risk aversion," *Journal of Econometrics*, 94(1-2), 9–51.

Asness, C. (1994): "The power of past stock returns to explain future stock returns," Working paper, University of Chicago.

Asness, C. S., T. J. Moskowitz, and L. H. Pedersen (2013): "Value and Momentum Everywhere," *Journal of Finance*, 68(3), 929–985.

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., 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.

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.

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. (2015): "Spurious Factors in Linear Asset Pricing Models," Job market paper, LSE.

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., C. Pflueger, and L. M. Viceira (2015): "Monetary Policy Drivers of Bond and Equity Risks," Harvard Business School Working Papers 14-031, Harvard Business School.

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.

Carhart, M. M. (1997): "On Persistence in Mutual Fund Performance," *Journal of Finance*, 52(1), 57–82.

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

Chernov, M. (2003): "Empirical reverse engineering of the pricing kernel," *Journal of Econometrics*, 116(1-2), 329–364.

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.

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.

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.

Daniel, K., and T. J. Moskowitz (2016): "Momentum crashes," *Journal of Financial Economics*, 122(2), 221 – 247.

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 (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.

——— (2016): "Dissecting Anomalies with a Five-Factor Model," *Review of Financial Studies*, 29(1), 69–103.

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.

GARLAPPI, L., AND Z. SONG (2016): "Can Investment Shocks Explain the Cross Section of Equity Returns?," *Management Science*, Forthcoming.

GHOSH, A., C. JULLIARD, AND A. P. TAYLOR (2016): "What Is the Consumption-CAPM Missing? An Information-Theoretic Framework for the Analysis of Asset Pricing Models," *Review of Financial Studies*, Forthcoming.

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.

HANSEN, L. P., AND R. JAGANNATHAN (1991): "Implications of Security Market Data for Models of Dynamic Economies," *Journal of Political Economy*, 99(2), 225–62.

——— (1997): " Assessing Specification Errors in Stochastic Discount Factor Models," *Journal of Finance*, 52(2), 557–90.

HANSEN, L. P., AND S. F. RICHARD (1987): "The Role of Conditioning Information in Deducing Testable," *Econometrica*, 55(3), 587–613.

HANSEN, L. P., AND K. J. SINGLETON (1982): "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models," *Econometrica*, 50(5), 1269–86.

HARVEY, C. R., Y. LIU, AND H. ZHU (2015): ". . . and the Cross-Section of Expected Returns," *Review of Financial Studies*, Forthcoming.

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 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.

Jegadeesh, N., and S. Titman (1993): "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency," *Journal of Finance*, 48(1), 65–91.

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

Johnson, T. C. (2002): "Rational Momentum Effects," *The Journal of Finance*, 57(2), pp. 585–608.

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

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

Kogan, L., and D. Papanikolaou (2014): "Growth Opportunities, Technology Shocks, and Asset Prices," *Journal of Finance*, 69(2), 675–718.

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.

Ludvigson, S. C., S. Ma, and S. Ng (2015): "Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?," Working Paper 21803, National Bureau of Economic Research.

Luenberger, D. G. (1969): *Optimization by Vector Space Methods*. John Wiley and Sons.

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.

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

Novy-Marx, R., and M. Velikov (2016): "A Taxonomy of Anomalies and Their Trading Costs," *Review of Financial Studies*, 29(1), 104–147.

Ozdagli, A., and M. Velikov (2016): "Show Me the Money: The Monetary Policy Risk Premium," mimeo, Federal Reserve Bank of Richmond.

Papanikolaou, D. (2011): "Investment Shocks and Asset Prices," *Journal of Political Economy*, 119(4), 639 – 685.

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.

Pinter, G. (2016): "Comment on "News Shocks and the Slope of the Term Structure of Interest Rates" (Kurmann-Otrok, 2013 AER)," mimeo.

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

Roll, R. (1977): "A critique of the asset pricing theory's tests Part I: On past and potential testability of the theory," *Journal of Financial Economics*, 4(2), 129–176.

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

Rosenberg, J. V., and R. F. Engle (2002): "Empirical pricing kernels," *Journal of Financial Economics*, 64(3), 341–372.

Ross, S. (2015): "The Recovery Theorem," *The Journal of Finance*, 70(2), 615–648.

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.

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 (2001): "Vector Autoregressions," *Journal of Economic Perspectives*, 15(4), 101–115.

Tenreyro, S., and G. Thwaites (2016): "Pushing on a String: US Monetary Policy Is Less Powerful in Recessions," *American Economic Journal: Macroeconomics*, 8(4), 43–74.

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.

Weber, M. (2015): "Nominal Rigidities and Asset Pricing," mimeo, University of Chicago.

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

Appendix

A Including Consumption in the VAR

Figure 7: Impulse Responses to a λ-shock: Output and Consumption

(a) VAR with Output

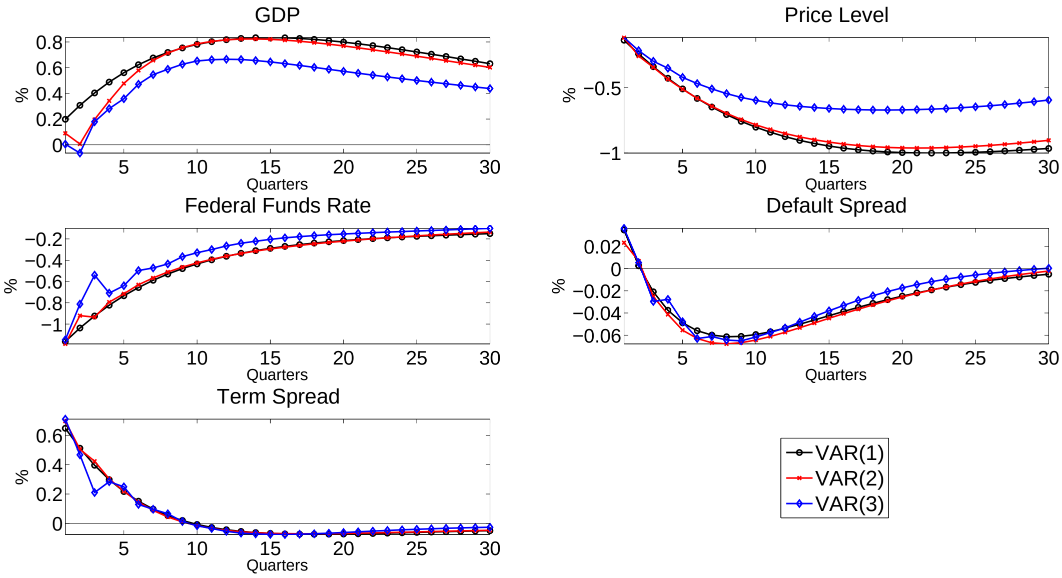

(b) VAR with Consumption

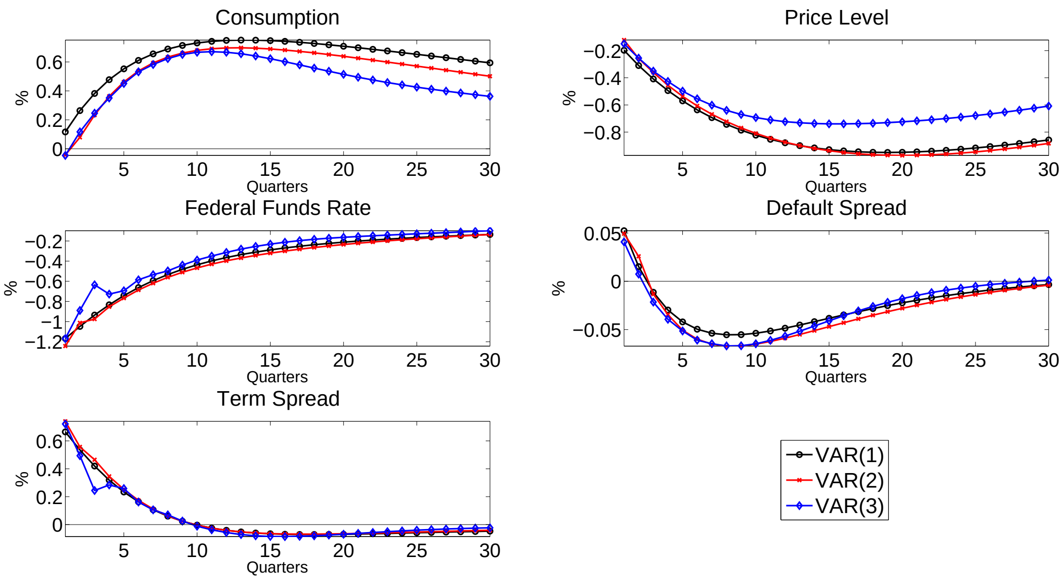

Notes: The vertical axes are %-deviations from steady-state, and the horizontal axes are in quarters. In the upper panel, the baseline IRFs (Figure 2) are shown against the IRFs implied by VAR(1) and VAR(3) models. In the lower panel IRFs for VAR(1), VAR(2) and VAR(3) models are plotted where GDP is replaced by the consumption measure used in Greenwald, Lettau, and Ludvigson (2015). In all cases, the FF55 portfolios were used as test assets.

B Results from a Bayesian VAR

To explore the role of parameter uncertainty in the VAR model 2.5, I re-estimate 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 orthogonalisation method to these draws and save the resulting IRFs, and (iv) and repeat the Gibbs-iteration and the orthogonalisation for another N_2 times. The posterior distribution of IRFs is then constructed based on the N_2 draws.

Figure 8: Impulse Responses to a λ-shock: Results from a Bayesian VAR(2)

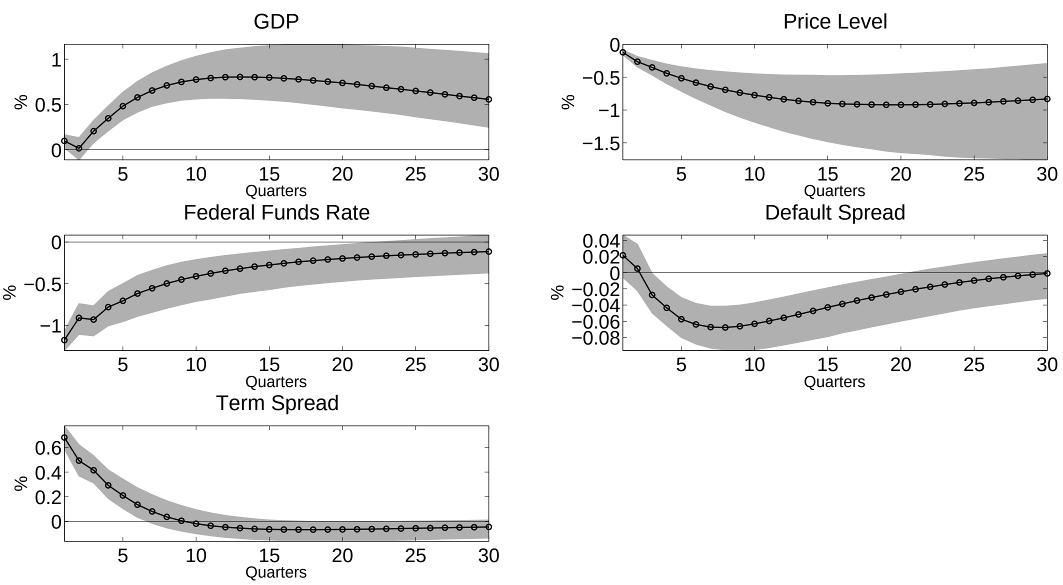

Notes: The sample period is 1963Q3 - 2008Q3. The Minnesota-type normal inverted Wishart priors are implemented following Banbura, Giannone, and Reichlin (2010). The figure shows the pointwise median and 5th-95th percentiles of $N_2 = 5000$ draws (after burning the first $N_1 = 5000$ draws) from the posterior distribution of the impulse responses. The vertical axes are %-deviations from steady-state, and the horizontal axes are in quarters.

Figure 8 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 $5^{th} - 95^{th}$ probability bands suggest that at a 15-quarter horizon output rises around 0.5% above steady-state with 95% probability but does not rise more than 1.2% above steady-state with the same probability. The rest of the IRFs can be interpreted in a similar fashion.

C Pricing Performance

Table 3: Results from the Two-pass Regressions, FF55 Portfolios

Factor prices (λ)						R^2
Panel A: 1-factor Model with the λ-shock						
Constant	λ-shock					
0.69	0.51					0.21
(0.54) [0.70]	(0.24) [0.26]					
	0.81					0.13
	(0.31) [0.38]					
Panel B: 5-factor Model with the λ- and the Other VAR Shocks						
	λ-shock	Shock 2	Shock 3	Shock 4	Shock 5	
	0.81	0.00	0.00	0.00	0.00	0.13
	(0.22) [0.24]					
Panel C: The Fama-French 3-factor Model						
Constant	MKT	HML	SMB			
3.20	-1.78	0.91	0.48			0.37
(0.95) [0.99]	(1.11) [1.16]	(0.45) [0.45]	(0.44) [0.44]			
	1.43	0.92	0.35			0.13
	(0.63) [0.63]	(0.45) [0.46]	(0.44) [0.44]			

Notes: This table reports the cross-sectional regressions using the excess returns on the FF55 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 the results for five-factor model using all structural shocks from the VAR. Panel C 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 VARHAC procedure (following den Haan and Levin (2000); Burnside (2011), in order to take into account possible serial correlation in the errors) are in brackets.

Table 4: Results from the Two-pass Regressions, FF25 Portfolios

Factor prices (λ)						R^2
Panel A: 1-factor Model with the λ-shock						
Constant	λ-shock					
0.18	1.32					0.85
(0.81) [1.49]	(0.28) [0.46]					
	1.44					0.84
	(0.47) [0.83]					
Panel B: 5-factor Model with the λ- and the Other VAR Shocks						
	λ-shock	Shock 2	Shock 3	Shock 4	Shock 5	
	1.44	0.00	0.00	0.00	0.00	0.84
	(0.26) [0.38]					
Panel C: The Fama-French 3-factor Model						
Constant	MKT	HML	SMB			
2.92	-1.62	1.44	0.57			0.80
(1.12) [1.12]	(1.30) [1.27]	(0.43) [0.44]	(0.43) [0.43]			
	1.24	1.44	0.64			0.76
	(0.63) [0.64]	(0.43) [0.44]	(0.44) [0.43]			

Notes: See notes under Table 3.

Table 5: Results from the Two-pass Regressions, 25 Profitability-Size Portfolios

Factor prices (λ)						R^2
Panel A: 1-factor Model with the λ-shock						
Constant	λ-shock					
0.04	1.53					0.67
(0.71) [1.21]	(0.50) [0.82]					
	1.56					0.67
	(0.59) [1.18]					
Panel B: 5-factor Model with the λ- and the Other VAR Shocks						
	λ-shock	Shock 2	Shock 3	Shock 4	Shock 5	
	1.44	0.00	0.00	0.00	0.00	0.84
	(0.48) [0.77]					
Panel C: The Fama-French 3-factor Model						
Constant	MKT	HML	SMB			
2.40	-1.06	1.36	0.58			0.64
(1.01) [0.97]	(1.20) [1.17]	(0.74) [0.77]	(0.45) [0.44]			
	1.21	2.11	0.57			0.59
	(0.63) [0.63]	(0.78) [0.96]	(0.45) [0.44]			

Notes: See notes under Table 3.

Table 6: Results from the Two-pass Regressions, 25 Investment-Size Portfolios

Factor prices (λ)						R^2
Panel A: 1-factor Model with the λ-shock						
Constant	λ-shock					
0.24	1.47					0.66
(0.73) [1.29]	(0.40) [0.60]					
	1.65					0.65
	(0.57) [1.15]					
Panel B: 5-factor Model with the λ- and the Other VAR Shocks						
	λ-shock	Shock 2	Shock 3	Shock 4	Shock 5	
	1.65	0.00	0.00	0.00	0.00	0.65
	(0.29) [0.53]					
Panel C: The Fama-French 3-factor Model						
Constant	MKT	HML	SMB			
3.11	-1.62	1.49	0.43			0.76
(1.31) [1.23]	(1.45) [1.38]	(0.56) [0.55]	(0.44) [0.43]			
	1.38	2.14	0.44			0.71
	(0.63) [0.63]	(0.51) [0.50]	(0.44) [0.43]			

Notes: See notes under Table 3.

Table 7: Results from the Two-pass Regressions, 10 Momentum Portfolios of Ken French

Factor prices (λ)						R^2
Panel A: 1-factor Model with the λ-shock						
Constant	λ-shock					
-0.01	2.56					0.97
(0.78) [2.29]	(0.53) [1.50]					
	2.54					0.97
	(0.69) [2.10]					
Panel B: 5-factor Model with the λ- and the Other VAR Shocks						
	λ-shock	Shock 2	Shock 3	Shock 4	Shock 5	
	2.54	0.00	0.00	0.00	0.00	0.97
	(0.42) [1.14]					
Panel C: The Fama-French 3-factor Model						
Constant	MKT	HML	SMB			
13.52	-12.29	-1.58	1.76			0.84
(2.70) [6.82]	(2.95) [7.29]	(1.15) [3.17]	(1.41) [3.57]			
	2.20	-4.39	-4.16			0.76
	(0.64) [0.71]	(1.07) [2.28]	(0.94) [1.80]			

Notes: See notes under Table 3.

Table 8: Results from the Two-pass Regressions, 10 Momentum Portfolios of Daniel and Moskowitz (2016)

Factor prices (λ)						R^2
Panel A: 1-factor Model with the λ-shock						
Constant	λ-shock					
0.00	2.51					0.95
(0.82) [2.27]	(0.46) [1.19]					
	2.51					0.95
	(0.41) [1.29]					
Panel B: 5-factor Model with the λ- and the Other VAR Shocks						
	λ-shock	Shock 2	Shock 3	Shock 4	Shock 5	
	2.51	0.00	0.00	0.00	0.00	0.95
	(0.41) [1.10]					
Panel C: The Fama-French 3-factor Model						
Constant	MKT	HML	SMB			
5.58	-3.53	-3.87	-3.18			0.93
(1.97) [3.92]	(2.21) [4.40]	(1.15) [2.24]	(1.19) [2.34]			
	2.64	-5.38	-4.16			0.92
	(0.66) [0.85]	(1.14) [2.89]	(1.06) [2.48]			

Notes: See notes under Table 3.

D The Equivalence between Maximising the Price of Risk and Maximising the Cross-sectional Fit

This section provides a numerical illustration that finding the shock that has the highest price of risk is equivalent to maximising the R^2-measure from the associated 1-factor model.[15] I use the baseline VAR(2) model and the FF25 portfolios and explore the space of admissible $\tilde{Q}$ matrices to uncover the relationship between λ and the R^2 implied by the corresponding 1-factor model. The scatter plot in Figure 9 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 drawn from the multivariate Normal distribution. The vertical red dashed line denotes the maximum achievable price of risk (1.44) associated with the λ-shock given the VAR specification. The horizontal red dashed line denotes the upper bound on the unadjusted R^2 (0.84 as in Table 4), which puts a cap on how well the orthogonalised structural shock can explain the cross-section.

Figure 9 makes it clear 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, 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, as the random 1-factor models get closer and closer to the upper bound in terms of the implied R^2 values, the associated price of risk converges to the maximum price of risk that is numerically achievable. Increasing the number of random draws does not change Figure 9.[16] Building on the theoretical discussion of Subsection 2.1, this can also be interpreted as a numerical proof of the equivalence between maximising the price of risk and maximising the cross-sectional fit – two different algorithms to uncover the λ-shock.

[15]The R^2 statistic has been calculated as:

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

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, $\hat{\alpha} = \bar{R}^e - \hat{\beta}\left(\varepsilon^\star\right) \times \hat{\lambda}\left(\varepsilon^\star\right)$.

[16]These results are available upon request.

Figure 9: A Simulation Exercise: Illustrating the Relationship between the Price of Risk and Cross-sectional R^2

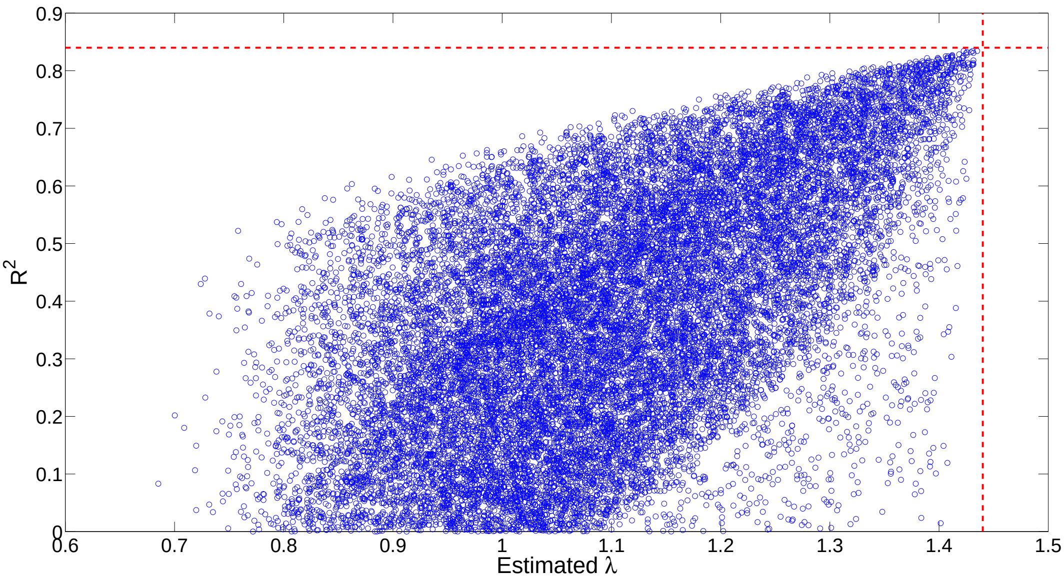

Notes: The scatter plot (based on 20,000 random $\tilde{Q}$ matrices) shows the relationship between the price of risk demanded by $\tilde{\varepsilon}_t$ associated with a given candidate draw $\tilde{Q}$ and the cross-sectional R^2 implied by the corresponding 1-factor model. For presentation purposes, 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 (1.44) from the five-variable VAR model 2.5, and the horizontal red dashed line is the upper bound (0.84) on the unadjusted R^2-measure associated with any 1-factor model extracted from the VAR model 2.5. To obtain these random draws, I apply Householder transformations to 20,000 five-dimensional matrices drawn from the multivariate Normal distribution.