The Macroeconomic Shock with the Highest Price of Risk

Gabor Pinter*

August 10, 2016

Preliminary – Comments Welcome

Abstract

A single structural shock, that demands the highest possible risk premium in a standard macroeconomic VAR, explains more of the cross-section of stock returns than the 3-factor Fama-French model can. This shock is highly (>70%) correlated with technology news shocks studied by the macroeconomics literature. This is striking given that my identification strategy has nothing to do with the strategies used to identify news shocks, and my VAR does not even contain a measure of technology as an observable. My results provide strong independent support for the role of technology news shocks in explaining business cycles as well as the cross-sectional variation in stock returns. In addition, I highlight the empirical link between two largely unconnected literatures on consumption based asset pricing and on macroeconomic news shocks.

JEL Classification*:* C32, G12

Keywords: Stock Returns, VAR, Identification, Technology News Shocks, Consumption Based Asset Pricing

*Bank of England; email: gabor.pinter@bankofengland.co.uk
The author would like to thank, but do not implicate, Saleem Bahaj, Andy Blake, John Cochrane, Wouter den Haan, Clodo Ferreira, Rodrigo Guimaraes, Ralph Koijen, Peter Kondor, Ricardo Reis and Shengxing Zhang for helpful comments. The views expressed in this paper are those of the author, and not necessarily those of the BoE, the MPC or the FPC.

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 stock returns. I argue that part of this challenge has been caused by the fact that macrovariables and innovations in them 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 using a macroeconometric identification strategy which confirms that one single structural shock can explain about as much of the average excess returns of the 25 portfolios of Fama and French (1993) sorted on book-to-market and size (FF25 henceforth) as the 3-factor model proposed by these authors. This shock, which I will refer to as a λ-shock, behaves like a news shock about future total factor productivity (TFP) – an object which has been subject to increasing interest among macroeconomists. In fact, the correlation between the λ-shock I identify and the TFP news shock, studied by Kurmann and Otrok (2013) amongst others, is 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 news shocks, and my model does not even contain a measure of TFP as an observable. My results therefore provide indirect, though very strong, empirical support for the importance of TFP news shocks for understanding asset prices and business cycles in general. In addition, I will highlight the strong empirical link between TFP news shocks and innovations in the consumption process that have a more delayed or even long run effect on consumption growth, as studied by Bansal and Yaron (2004), Parker and Julliard (2005) amongst others.

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

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

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 when pricing the cross-section of stock returns. Mechanically, this 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 value for the associated factor risk premium. I will also show that 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.

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 a TFP news shock, provides compelling evidence on the important role of TFP news shocks in not only driving business cycles but also in explaining the cross-section of stock returns. And this is the main contribution of this paper.

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).[3] To the best of my knowledge, only few of these papers sought

[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.

[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.

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 macroeconomics 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 implications 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 expanding macroeconomics literature on the effects of news shocks on business cycle dynamics. A partial list includes Beaudry and Portier (2006, 2014), Jaimovich and Rebelo (2009), Barsky and Sims (2011), Schmitt-Grohe and Uribe (2012), Blanchard, L'Huillier, and Lorenzoni (2013) Kurmann and Otrok (2013), 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 impact and long-run restrictions as identification schemes) to identify news shocks.

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 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. $A_0 = Q^\star\tilde{A}_0$. The 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 regression of each of the k portfolios' excess return on the implied candidate structural shock $\tilde{\varepsilon}_t$:

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

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 $\tilde{Q}$:

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

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 demand the highest estimated value of λ in 2.4. 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 chose the Euler-angles of the Givens rotation appropriately such that the corresponding second-pass λ is maximised.[5]

[4]This is also the starting point for a range of identification strategies in the macroeconometrics 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); Kurmann and Otrok (2013)), using external instruments as proxies for structural shocks (Mertens and Ravn (2013); Gertler and Karadi (2015)) etc.

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

It is important to note that while assumptions about identification determines risk exposures and prices, 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 bank 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 bank 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 1963Q3-2008Q4 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

[6] I would like to thank Professor Amit Goyal for updating and sharing his dataset on his website.

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; 1963Q3-2008Q4, 182 quarters.

	Low	2	3	4	High
Small	0.35	2.06	2.22	2.83	3.09
2	0.88	1.77	2.49	2.65	2.79
3	1.00	1.90	2.08	2.26	2.92
4	1.31	1.32	1.74	2.22	2.22
Large	0.98	1.27	1.08	1.41	1.59

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 0.35% per quarter for the lower B/M portfolio (extreme growth stocks) to more than 3% per quarter for the highest B/M portfolio (extreme value stocks). In contrast, for the largest stocks (megacaps), average excess returns rise from about 1% to only about 1.6%. 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 λ-shock is a TFP News 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

[7] I would like to thank Professor Ken French for making the data available on his website.

paper, because it confirms that identifiable macroeconomic shocks (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

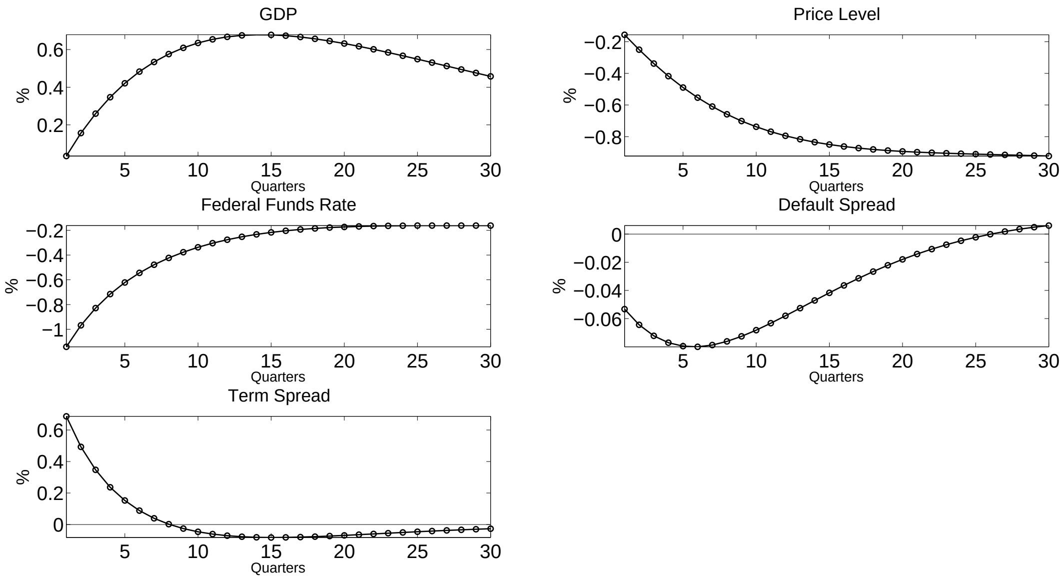

Notes: The vertical axes are %-deviations from steady-state, and the horizontal axes are in quarters.

Figure 1 displays the IRFs of the five variables to a 1 percent innovations in the λ-shock. 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, implying an endogenous monetary policy reaction that is strongly expansionary in response to the shock. Interestingly, the λ-shock has virtually no effect on GDP on impact, but the effect increases substantially with the horizon. As shown by Figure 6 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 empirical framework can possibly shed further light on the drivers of the slow consumption adjustment shocks that are the main source of macroeconomic risks. Is it demand, supply or monetary policy-type shocks that generate most of the consumption dynamics relevant to pricing? The VAR model I use includes information on the aggregate price level as well as the instrument of the monetary policy authority – two additional variables that are necessary to isolate demand, supply and monetary policy-type shocks from each other.

Overall, Figure 1 suggests that the λ-shock behaves like a supply-type shock with aggregate production moving into the opposite direction compared to the price level and the short-term interest rate. However, the delayed expansion of output makes the λ-shock clearly distinct from a positive unanticipated technology shock which tends to have an immediate positive impact on output and consumption, as traditionally studied by the Real Business Cycle (RBC) and the subsequent New Keynesian literature. 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) indeed finds that the contribution of unanticipated TFP shocks to the variance of aggregate stock market wealth is close to zero.[8] What kind of a supply-shock is the λ-shock then?

The reader by now may have noticed that these IRF results appear to be very similar to those obtained for identified TFP news shocks in the macroeconomics literature. For example, Figure 4 of Kurmann and Otrok (2013) shows results for an identified TFP news shocks 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 (FEV) of TFP over a given forecast horizon, and they force this shock to be orthogonal to contemporaneous movements in TFP.

To formally show the correspondence between the TFP news shock and the λ-shock that I identify from the cross-section of stock returns, Figure 2 plots the time-series of both shocks together. The correlation coefficient between the TFP news shocks (red dashed line) as identified in Kurmann and Otrok (2013) and the λ-shock (blue solid line) is 0.71. This correlation coefficient is computed based on the overlapping estimation period 1963Q4–2005Q2.

To reiterate, my identification (i) makes no assumption about the λ-shock's contribution to the FEV of any of the variables, (ii) does not impose any zero-type restrictions and (iii) does not even include TFP as an observable in the VAR. Not to mention that the empirical model of Kurmann and Otrok (2013) is slightly different from mine in terms of lag structure, sample period and variables used in the VAR. The fact that I come close to

[8] Greenwald, Lettau, and Ludvigson (2015) identifies 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: Comparison of TFP News Shock (Kurmann and Otrok, 2013) and the λ-shock: Correlation = 0.71

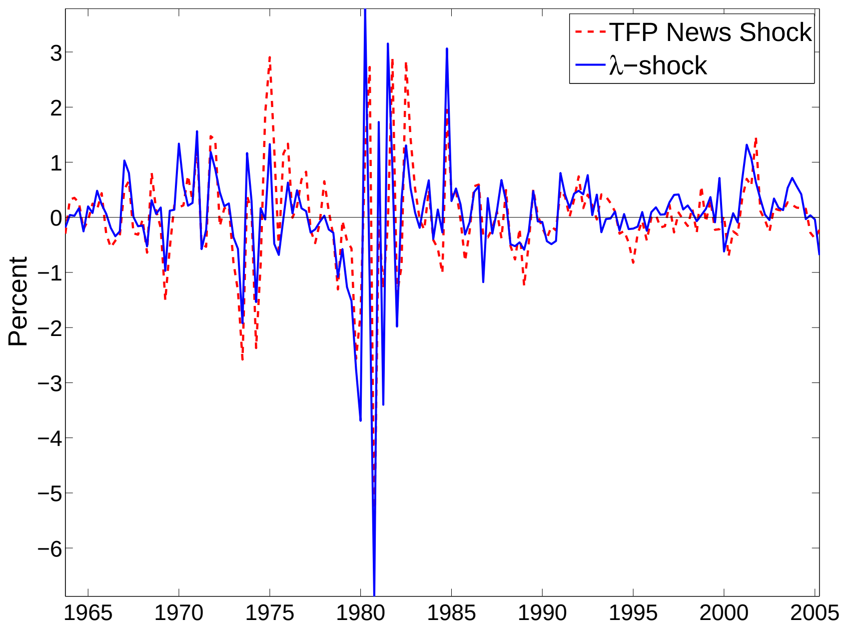

Notes: The TFP news shock series is the one plotted in Figure 5 on pp. 2625 of Kurmann and Otrok (2013) who apply the method of Uhlig (2004) to identify a TFP news shock over the period 1959Q2-2005Q2.

reconstructing the object the TFP news literature has studied (by applying a completely different and relatively more agnostic methodology, based on the magnitude of the risk price associated with the shock) provides strong empirical support for the relevance of TFP news shocks in driving business cycles and 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 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 (2.44) have more than 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)

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	1.34	1.80	1.84	2.44	0.39	1.20	1.66	1.95	2.27
2	0.69	1.44	1.96	2.13	1.98	0.53	1.34	2.07	2.05	2.09
3	1.14	1.60	1.65	2.13	1.81	0.98	1.65	1.78	2.33	2.19
4	1.00	1.07	1.80	2.10	2.16	0.96	1.06	2.04	2.34	2.45
Large	0.83	0.57	0.31	1.20	1.42	1.02	0.75	0.35	1.77	2.29

Notes: The table reports loadings on the identified λ-shock computed in time-series regressions for the FF25 portfolios sorted by size and book-to-market. The sample period is 1963Q3-2008Q4. 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.

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 1.1–1.2%-points additional excess returns per quarter. The estimated constant is not statistically different from zero, implying negligible 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 very similar to those obtained in the literature (e.g. Petkova (2006)).

A notable difference between the two models is the size of the estimate of the common pricing error parameter α_i: the estimated constant is an order of magnitude smaller for the 1-factor model (0.23) compared to the 3-factor model (3.26). This explains why the adjusted-R^2 measure somewhat drops (from 0.76 to 0.71) for the 3-factor model when the constant is not included in the second-pass regression, whereas the adjusted-R^2 measure is virtually unchanged for the 1-factor model irrespective of whether the constant is included or not.

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.

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			
0.23	1.10			0.78
(0.81) [1.22]	(0.24) [0.36]			
	1.24			0.77
	(0.44) [0.70]			
Panel B: The Fama-French 3-factor Model				
Constant	MKT	HML	SMB	
3.26	-2.09	1.40	0.54	0.76
(0.99) [1.05]	(1.18) [1.24]	(0.43) [0.44]	(0.43) [0.44]	
	1.09	1.48	0.59	0.71
	(0.64) [0.65]	(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.

Overall, these results lead to a remarkable conclusion: a single macroeconomic shock, that resembles the behaviour of a TFP news shock, seems to perform as 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 driving the cross-section of returns 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 are the type of macroeconomic risk that explains 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 8 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 miserably when used separately to price the cross-section. The least bad performance ($R^2 = 0.44$) is displayed by the 1-factor model that uses innovations in the Federal Funds rate as the pricing factor.[9] However, this result in itself

[9]This is in line with the reason why most empirical asset pricing paper, written in the spirit of the ICAPM, used innovations in the short term interest rate as one of the pricing factors (Campbell (1996);

Figure 3: Fitted Expected Returns vs. Average Realised Returns for 1963Q3–2008Q4

(a) 1-factor Model with λ-shock (b) Fama-French 3-factor model

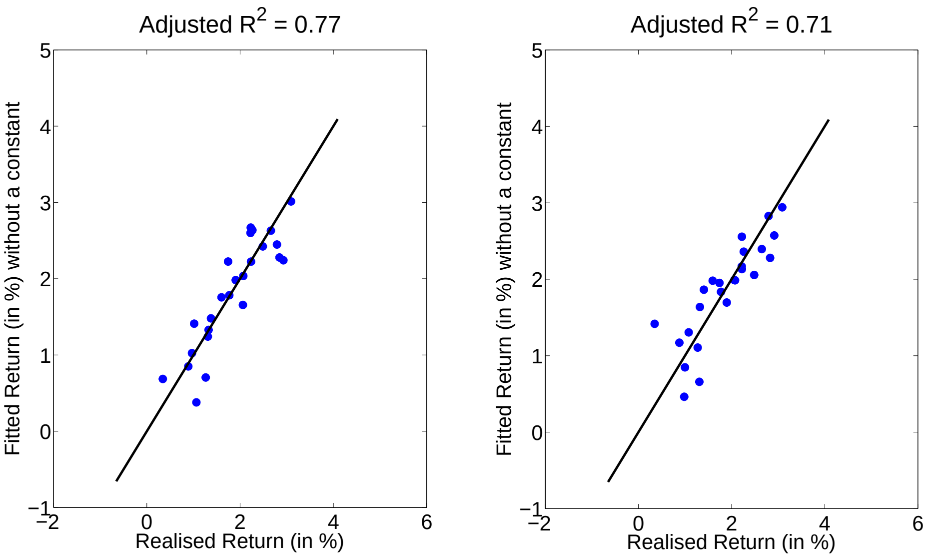

Notes: The R^2 values for the 1-factor model (left panel) and the 3-factor Fama-French model (right panel) are calculated assuming no constant in the second-pass regressions.

could even be misleading as it may suggest that *surprise* monetary policy actions could be the type of macroeconomic risk that is priced in the cross-section. While this statement may be true, interest rate innovations in my empirical framework reflect entirely *endogenous* monetary policy reactions to the λ-shock – an object that is fundamentally supply-related. Therefore all the pricing power of interest rate innovations in my VAR model is attributed to the fact that monetary policy reacts endogenously and aggressively on impact to λ-shocks.[10]

3.4 Extensions and Robustness

3.4.1 The Equivalence between Maximising the Price of Risk and Maximising the Cross-sectional Fit[11]

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

Petkova (2006); Maio and Santa-Clara (2012); Boons (2016) etc.).

[10]Such a strong endogenous response of monetary policy to TFP news type shocks is also documented in Kurmann and Otrok (2013).

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

also be reflected in the overall fit of the model, as measured by the R^2 statistic:

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

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, $\hat{\beta}\left(\varepsilon_t^\star\right) \times \hat{\lambda}\left(\varepsilon_t^\star\right)$ is the model's predicted mean returns and the estimated pricing errors are the residuals, $\hat{\alpha} = \bar{R}^e - \hat{\beta}\left(\varepsilon_t^\star\right) \times \hat{\lambda}\left(\varepsilon_t^\star\right)$. 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 $\varepsilon_t^\star$ 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.

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

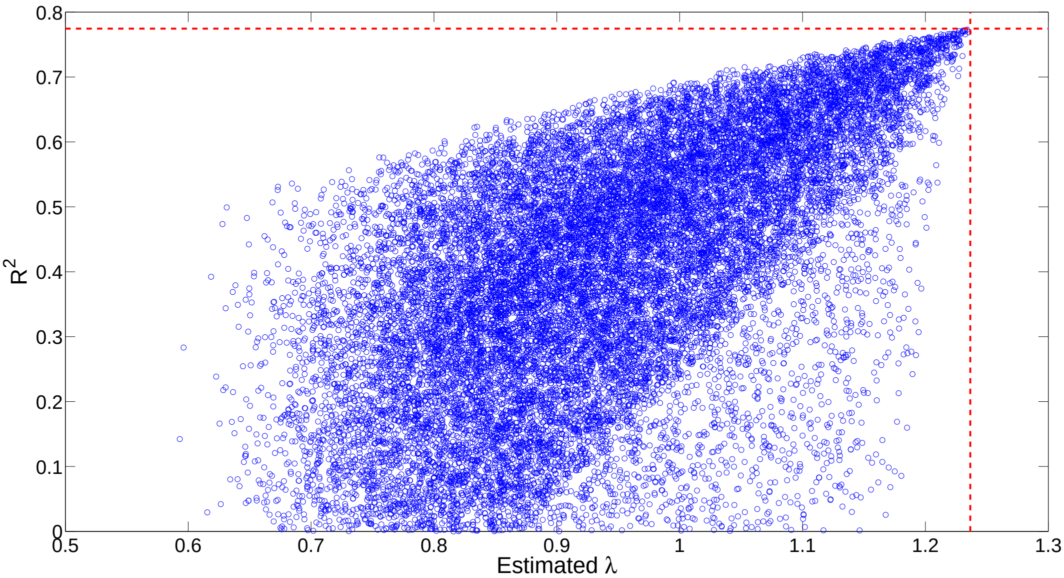

Notes: The scatter plot (based on 20,000 random $\tilde{Q}$ matrices) shows the relationship between the price of risk demanded by $\tilde{\varepsilon}_t$ associated with a given candidate draw $\tilde{Q}$ and the cross-sectional R^2 implied by the corresponding 1-factor model. For presentation purposes, I exclude those rotations that imply negative R^2 (about 38% 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 (1.24) from the five-variable VAR model 2.1, and the horizontal red dashed line is the upper bound (0.7744) 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.

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 0.7744. 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 drawn from the multivariate Normal distribution. The vertical red dashed line denotes the maximum achievable price of risk (1.24) associated with the λ-shock given the VAR specification. The horizontal red dashed line denotes the upper bound on the unadjusted R^2 (0.7744), 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.[12] I therefore interpret 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

[12] These results are available upon request.

deviation) of βs and λ, Figure 7 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 (0.0059) 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 Robustness to Changing the VAR and Portfolio Specifications

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 prevailing consensus that these portfolios are a good representation of the aggregate risk present in the cross-section (Chapter 20 of Cochrane (2005)). Nevertheless, I will now show that the identification of the λ-shock is not specific to the FF25 portfolios, and the time-series behaviour of the λ-shock is robust to changing/augmenting the asset pricing portfolios for identification. To do so, I expand the set of test portfolios beyond FF25 by adding the 30 industry portfolios of Fama-French (FF30 henceforth), thereby following prescription 1 (pp. 182) of Lewellen, Nagel, and Shanken (2010). 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 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 when pricing the 55 Size–B/M and Industry portfolios (FF25+FF30).

Panel A of Table 4 shows the cross-correlations for the baseline λ-shock (column 1), the modified λ-shock using the augmented portfolio set (column 2) and the TFP news shock series. The correlation between the baseline and modified λ-shock series is 0.87, suggesting that the same shock that prices the FF25 portfolios prices the larger set of 55 portfolios as well. As a consequence, the modified λ-shock continues to have high (0.67) correlation with TFP new shocks.

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-shocks and the TFP news shocks from different model specifications. 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 increase the correlation with the TFP news shock series. Conversely, increasing the lag length of the VAR, using

Table 4: Robustness of the Identification of the λ-shock

(a) Selection of the Base Portfolios

Correlation Coefficients		
λ-shock Baseline	λ-shock FF25+FF30	TFP News Shock
1.00		
0.87	1.00	
0.71	0.67	1.00

(b) 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	TFP News Shock
1.00						
0.88	1.00					
0.94	0.86	1.00				
0.79	0.80	0.76	1.00			
0.99	0.86	0.93	0.77	1.00		
0.95	0.83	0.89	0.81	0.95	1.00	
0.71	0.63	0.71	0.60	0.72	0.69	1.00

Notes: The table in Panel A reports the correlation coefficients among the baseline λ-shock (using the FF25 portfolios), the modified λ-shock (FF25+FF30) and the TFP news shock as identified in Kurmann and Otrok (2013). The table in Panel B 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 using the real monthly Industrial Production Index (FRED code: INDPRO) averaged over each quarter instead of GDP (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 TFP news shock (Column 7). The values are computed based on the overlapping period 1963Q4–2005Q2 (Panel A) and 1964Q1–2005Q2 (Panel B).

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.

3.4.3 The Link between TFP News and Consumption Based Models

In this subsection, I highlight how TFP news shocks relate empirically to the large macro-finance literature that used consumption based models to explain the cross-section of stock returns and asset prices in general. A key tenet of the CCAPM theory is that equilibrium prices of risk should be determined by the covariance of asset returns with current and future marginal utilities of consumption of the representative household. However, it has long been pointed out (Mehra and Prescott (1985)) that the empirical covariance of asset returns with current consumption growth is not sufficiently large to generate plausible risk premia. As a result, more recent consumption based models (Bansal and Yaron (2004); Parker and Julliard (2005); Hansen, Heaton, and Li (2008); Bryzgalova and Julliard (2015) amongst others) focused on innovations that generate movements in consumption that can happen with a delay and can even occur in the very long run.

The impulse response functions of consumption and output displayed in Figure 1 and 6 and presented in the macroeconomics literature (Barsky and Sims, 2011; Kurmann and Otrok, 2013) suggest that TFP news shocks may be related to the type of innovations that CCAPM models have studied. To explore this issue more rigorously, I use a simple model to illustrate the empirical relationship between CCAPM models and the models of TFP news. I estimate the following reduced-form process for aggregate consumption growth, used by Bansal and Yaron (2004) amongst many others:

$$\begin{aligned} \Delta c_t &= \mu + x_{t-1} + \sigma_c \varepsilon_{c,t} \\ x_t &= \rho x_{t-1} + \sigma_x \varepsilon_{x,t}, \end{aligned} \tag{3.2}$$

where Δc_t is aggregate consumption growth, μ is a constant, x_t is a persistent predictable component, $\varepsilon_{c,t}$ is a contemporaneous innovations in consumption growth and $\varepsilon_{x,t}$ is an innovation in the time-varying component of expected consumption growth. The recent macro-finance literature has explored the role of $\varepsilon_{x,t}$ in explaining a wide range of asset pricing phenomena. I will therefore inspect the statistical relationship between the estimated time-series of $\varepsilon_{x,t}$ and the TFP news shock series.

Table 5: Maximum Likelihood Estimate of the Consumption Process 3.2

	Estimate	SE
μ	0.0057	0.0010
ρ	0.7746	0.0940
σ_c	0.0053	0.0005
σ_x	0.0030	0.0007

Notes: the table shows the maximum likelihood estimates of the parameters in the state-space model 3.2, using the quarterly real consumption growth data from Greenwald, Lettau, and Ludvigson (2015), spanning the period 1963Q3 – 2008Q4. The values in the second and third column are the point estimates and standard errors, respectively.

The state-space model 3.2 is estimated with Maximum Likelihood using the quarterly personal consumption expenditure series of Greenwald, Lettau, and Ludvigson (2015) for the 1963Q3 – 2008Q4 period. The estimation results are presented in Table 5. I use the Kalman-smoother to extract the estimated series $\varepsilon_{x,t}$, which I then annualise, following the practice of Bansal and Yaron (2004), in order to avoid problems related to higher-frequency measurement. Figure 5 plots the obtained annual consumption growth innovation series against the annualised series of TFP news shock. The results clearly confirm that there is a positive relationship between consumption innovations studied by the recent macro-finance literature and TFP news shocks studied by the macroeconomics literature.

One macroeconomic interpretation of these findings is that the structural primitives that underlie the reduced-form innovations studied by the CCAPM models are of fundamentally supply-type in nature. This is consistent with most recent general equilibrium

Figure 5: The Relationship between TFP News Shocks and Reduced-form Innovations in Expected Consumption Growth

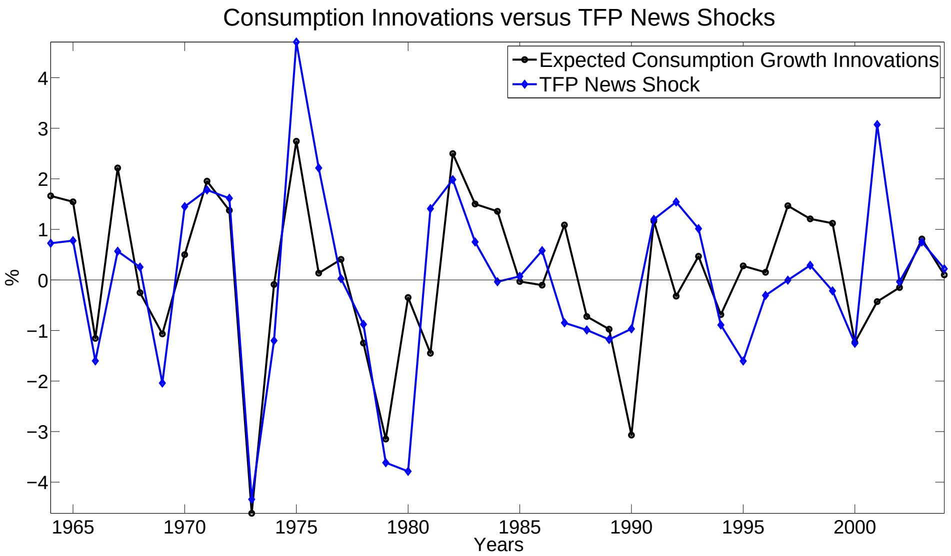

Notes: the figure shows the estimated time-path of the annualised estimated innovations in expected consumption growth ($\varepsilon_{x,t}$ in model 3.2) and the annualised TFP news shock. The correlation coefficient is 0.67.

models that analyse the link between expected consumption growth, risk premia and macroeconomic shocks. For example, the theoretical model studied in Borovicka and Hansen (2014) finds that supply side factors such as the neutral technology channel is a main driver of the consumption dynamics that are relevant to pricing. In addition, Malkhozov and Tamoni (2015) estimates a DSGE model with Bayesian methods and finds that it is the news component in the technology process which explains a large fraction of the dynamics in stock-market valuations and business cycles.

4 Conclusion

This paper has taken on the challenge to link the origins of cross-sectional variation in stock returns to macroeconomic primitives, by seeking to identify an orthogonal structural shock (in the spirit of the ICAPM) that can perform well in pricing the 25 portfolios of Fama and French (1993). The main finding is easy to summarise: innovations that are related to news shocks about future technology are the main source of the aggregate, nondiversifiable risk that is proxied by the returns of the FF25 portfolios. This 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 macroeconomics literature that focuses on unanticipated shocks as sources of business cycles, but consistent with a smaller but rapidly expanding literature (Schmitt-Grohe and Uribe, 2012; Beaudry and Portier, 2014) that emphasised the role of anticipated, news-type shocks in driving aggregate fluctuations.

Second, there is a close empirical relationship between technology news shocks and innovations in the consumption growth process that have a delayed effect on consumption and that are relevant to asset pricing, as shown by recent CCAPM models. Interestingly, I found no evidence for the role of demand-type shocks in explaining the cross-section of stock returns, or driving the type of reduced-form innovations in the consumption process that CCAPM models have focused on.

Third, 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 well-known structural shock. In fact, this paper could uncover only one such a pricing factor: the λ-shock which behaves like a TFP news shock.

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

References

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.

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.

Blanchard, O. J., J.-P. L'Huillier, and G. Lorenzoni (2013): "News, Noise, and Fluctuations: An Empirical Exploration," *The American Economic Review*, 103(7), 3045–3070.

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., 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.

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.

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.

Mehra, R., and E. C. Prescott (1985): "The equity premium: A puzzle," *Journal of Monetary Economics*, 15(2), 145–161.

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.

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.

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.

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 6: Impulse Responses to a λ-shock: Output and Consumption

(a) VAR with Output

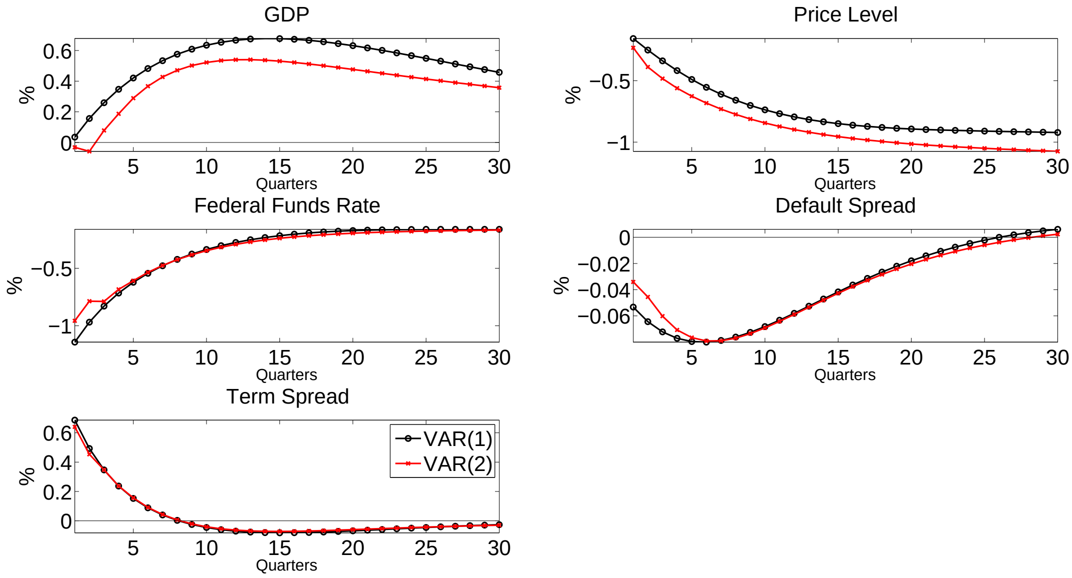

(b) VAR with Consumption

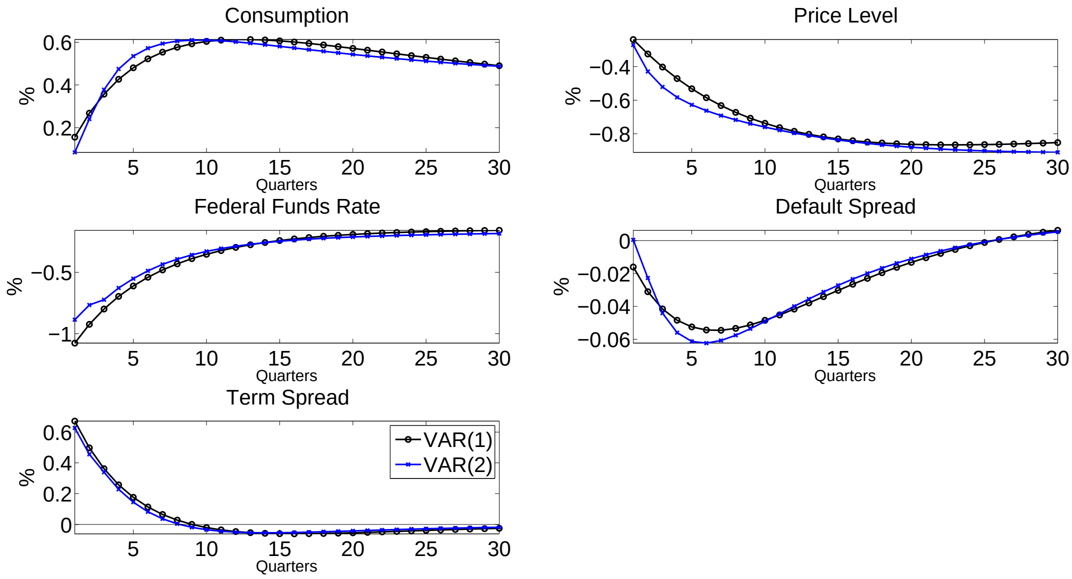

Notes: The vertical axes are %-deviations from steady-state, and the horizontal axes are in quarters.

Figure 7: The Identification of the λ-shock: the Relationship between Price of Risk and Dispersion of Exposures

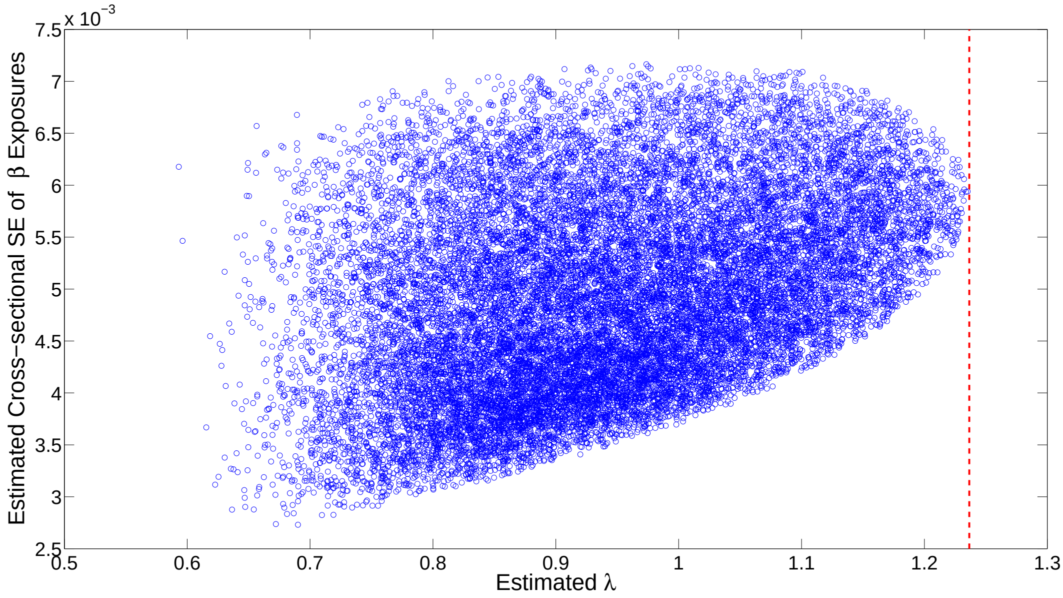

Notes: The scatter plot (based on 20,000 random $\tilde{Q}$ matrices) shows the relationship between the standard deviation of β-exposures to $\tilde{\varepsilon}_t$ associated with a given candidate draw $\tilde{Q}$ and the cross-sectional R^2 implied by the corresponding 1-factor model. For presentation purposes, I exclude those rotations that imply negative R^2 (about 38% 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 (1.24) 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.

Figure 8: The Role of Structural Identification: the Inability of the Individual VAR Innovations to Explain the Cross-section of FF25 Portfolios

(a) GDP Innovation Model

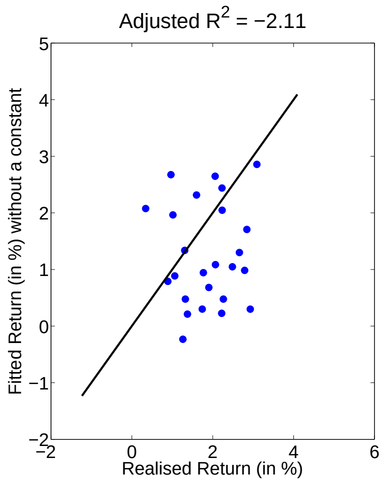

(b) Inflation Innovation Model

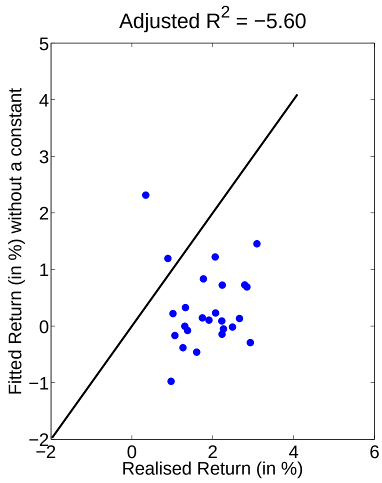

(c) Interest Rate Innovation Model

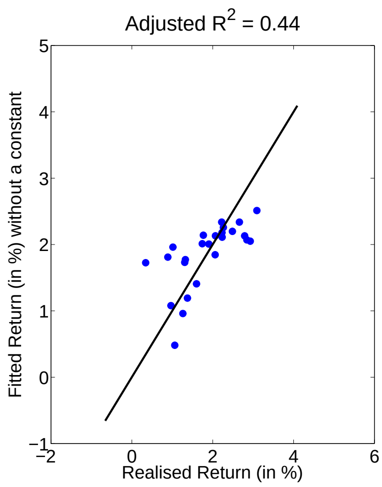

(d) Default Spread Innovation Model

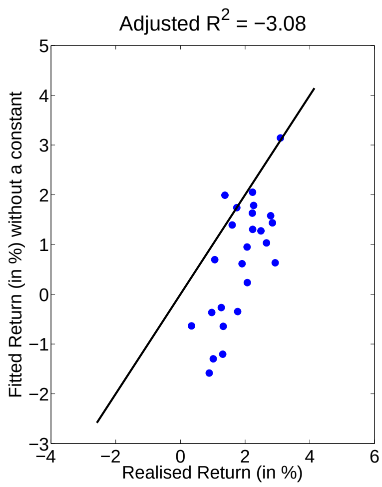

(e) Term Spread Innovation Model

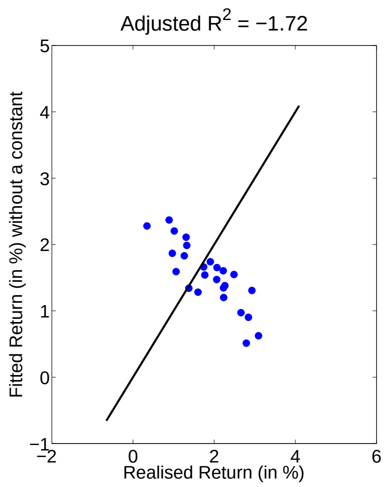

Notes: The R^2 values for the five 1-factor models are calculated assuming no constant in the second-pass regressions.