

News Shocks and the Slope of the Term Structure of Interest Rates*

André Kurmann

Université du Québec à Montréal

Christopher Otrok

University of Virginia

First Version: May 29, 2009

This Version: January 18, 2010

Abstract

We provide a new structural interpretation of the relationship between the slope of the term structure of interest rates and macroeconomic fundamentals. We first adopt an agnostic identification approach that allows us to identify the shocks that explain most of the movements in the slope. We find that two shocks are sufficient to explain virtually all movements in the slope. Impulse response functions for the first shock, which explains 70-90 percent of the movements in the slope, lead us to interpret this main shock as a news shock about future productivity. We confirm this interpretation by formally identifying such a news shock as in Barsky and Sims (2009) and Sims (2009). We then assess to what extent a New Keynesian DSGE model is capable of generating the observed slope responses to a news shock. We find that augmenting DSGE models with a term structure provides valuable information to discipline the description of monetary policy and the model's response to news shocks in general.

*An earlier version of this paper was circulated as 'The Sources of Movements in the Term Structure of Interest Rates'. We thank seminar participants at Cardiff University, the Federal Reserve Board, the Federal Reserve Banks of Atlanta, Philadelphia, San Francisco and St Louis, Koc University, Wharton and conference participants at the 2009 Joint Statistical Meetings. We thank Eric Sims for providing his productivity data, and Wen Yao and Chen Han for excellent research assistance. Kurmann gratefully acknowledges financial support from the SSHRC and the hospitality of The Wharton School where part of this project was completed.
Contact information: kurmann.andre@gmail.com and otrok@virginia.edu.

1 Introduction

The slope of the term structure – commonly defined as the spread between the yield on a long-term treasury bond and a short-term bill rate – has drawn the attention of many separate literatures. In forecasting, it is well established that the slope provides valuable predictive content for future economic activity (e.g. Estrella and Hardouvelis, 1991).¹ In finance, there is a large literature trying to explain both the average size and the time-variation of the slope with either latent factor no-arbitrage models (e.g. Duffie and Kan, 1996) or consumption-based asset pricing models (e.g. Piazzesi and Schneider, 2006).² In macroeconomics, the slope of the term structure plays a central role for the transmission of monetary policy (e.g. Clarida, Gali and Gertler 1999). In recent years, a rapidly growing number of papers has attempted to bridge the gap between these different literatures. While these papers have uncovered strong linkages between term structure and macroeconomic dynamics (e.g. Ang and Piazzesi, 2003; Diebold, Rudebusch and Aruoba, 2006), important questions remain unanswered. In particular, what are the fundamental sources of movements in the term structure slope? Do these fundamentals look like typical macroeconomic shocks? Can modern dynamic stochastic general equilibrium (DSGE) macroeconomic models replicate the observed slope responses to these shocks?

In this paper we provide answers to these questions. We first use a macro-finance VAR to show that over 70% of movements in the slope are due to news shocks about future innovations to total factor productivity (TFP). A key driver of this result is the *endogenous* response of monetary policy. After a positive news shock, the Federal Funds rate, and with it the short-end of the term structure, drops. Since the long-end of the term structure doesn't react to the news shock, the slope increases and only gradually returns to its initial value. The structure we identify provides a unified explanation for a number of stylized facts of the term structure: (i) variations in the slope are primarily due to fluctuations in the short-end of the term structure; (ii) steep yield curves (i.e. large slopes) generally predict future economic growth; and (iii) *systematic* monetary policy plays an important role for the linkage between macroeconomic and term structure dynamics. We then assess to what extent a modern medium-scale DSGE model is capable of generating the observed

¹See Ang, Piazzesi and Wei (2006) for a recent application and an extensive review of the literature.

²Other important latent factor no-arbitrage contributions include Knez, Litterman and Scheinkman (1994) and Dai and Singleton (2000). Recent consumption-based contributions are Wachter (2006) and Bansal and Shaliastovich (2007).

slope responses to a news shock. We find that the term structure provides valuable information to discipline the description of monetary policy. At the same time, the model falls well short of matching simultaneously both macroeconomic and term structure responses to TFP news shocks. In our view, this failure represents an important challenge for modern macroeconomic models.

Our interest in uncovering the structural shocks that drive the term structure slope is motivated by the recent macro-finance literature that has moved in the direction of linking simple atheoretical term structure factor models with macroeconomic data. This literature has provided some interesting linkages between the term structure and macroeconomic data. For example, Ang and Piazzesi (2003) find that movements in the slope are related to macro factors such as inflation. Diebold, Rudebusch and Aruoba (2005) find that there is close connection between movements in the slope and changes in the Federal Funds rate. However, as Rudebusch and Wu (2008) argue, as useful as this literature is, it is still in its 'infancy'. In particular, they argue that the current state of the macro-finance literature is an 'intermediate step' towards the goal of building a deep theoretical model linking the term structure with the macroeconomy. In this paper we take that next step by uncovering the fundamental shocks and structure that propagates these shocks.

We adopt a novel approach for the macro-finance literature in the search for a structural explanation for slope movements. Instead of postulating a particular type of shock and then analyzing its effects, our strategy consists of first uncovering (in a statistical sense) the main innovations of movements in the slope of the term structure and then trying to provide an economic interpretation of these shocks. As in existing papers, we start by combining term structure variables with prominent macroeconomic aggregates in a VAR. We then apply a methodology developed by Uhlig (2003) to extract the exogenous shocks that explain as much as possible of the Forecast Error Variance (FEV) of a target variable in the VAR, which in our case is the slope. That is, we first look for a quantitatively important shock, and then interpret it. We do so by analyzing the impulse responses of the different variables in the VAR and contrasting them with the theoretical implications of different types of macroeconomic shocks.

Nothing in our approach requires that a small number of shocks accounts for a large fraction of slope variations or that these shocks have an appealing interpretation. Yet, when applying our empirical strategy to the 1959-2005 period, we find that one single shock can account for 70% to 90% of all unpredictable fluctuations in the term structure slope over a 10-year horizon. Furthermore, we find that this slope shock closely resembles a news shock about future innovations in TFP as

proposed in Beaudry and Portier (2006), Jaimovich and Rebelo (2009) or more recently Barsky and Sims (2009) and Sims (2009). Specifically, TFP and consumption barely move on impact of the shock but gradually increase to a new permanent level thereafter. At the same time, inflation and the Federal Funds rate drop sharply and remain below their initial level for more than 2 years. The gradual but permanent long-run reaction of TFP and consumption together with the inverse reaction of *both* inflation and the Federal Funds rate rules out alternative interpretations of the slope shock such as exogenous monetary policy shocks, demand shocks, marginal rate of substitution shocks or contemporaneous TFP shocks.

To investigate the TFP news shock interpretation more formally, we follow Barsky and Sims (2009) and Sims (2009) and identify a TFP news shock directly as the innovation that accounts for most of the FEV of TFP over a 10-year horizon but is orthogonal to contemporaneous TFP movements. Even though this identification procedure is completely different from our slope shock identification, we find that the extracted TFP news shock is highly correlated with the slope shock and generates almost identical impulse responses of the slope and macroeconomic aggregates. These results remain unchanged for a battery of robustness checks. We conclude that the main driver of fluctuations in the slope of the term structure is *news about future innovations to TFP*.

To shed more light on the transmission mechanism from TFP news shocks to the term structure, we decompose slope movements into a part due to the Expectations Hypothesis and a part due to variations in term premia. We find that term premia increase significantly on impact of a positive TFP news shock. This is consistent with the general statistical rejection of the Expectations Hypothesis in the finance literature (e.g. Campbell and Shiller 1991, Cochrane and Piazzesi 2005). At the same time, the Expectations Hypothesis remains empirically relevant: the negative response of the Federal Funds rate and thus the short-end of the term structure to the TFP news shock is much stronger than the reaction of the discounted sum of future expected short rates (i.e. the long rate under the Expectations Hypothesis). The resulting difference accounts for more than half of the total increase in the slope. Hence, the systematic response of monetary policy is an important channel through which TFP news shocks transmit to movements in the slope.

In the final section of the paper we evaluate the extent to which a medium-scale DSGE model (e.g. Smets and Wouters 2007) can account for term structure movements in response to a TFP news shock. As has been shown elsewhere, this class of model is not capable of generating sizable and variable term premia (e.g. Rudebusch and Swanson 2008). This fact motivates us to limit

our analysis to a log-linear environment (where term premia are by definition constant) and ask whether the model can at least account for the expectational part of the term structure response. The estimated DSGE model is relatively successful in matching the response of macroeconomic aggregates to the TFP news shock. The model is also capable of generating a drop in the Federal Funds rate on impact of the shock and a gradual return back to steady state. The Expectations Hypothesis thus implies a positive response of the slope. While these responses are consistent with the data in a qualitative sense, the model falls well short of delivering the magnitude of the term structure response that we observe in the data.

Despite its quantitative failure, the estimated DSGE model offers important insights into how TFP news shocks transmit through the economy to the slope of the term structure. Specifically, our VAR analysis implies that a positive TFP news shock triggers a fall in both inflation and real activity, to which the Fed reacts systematically by lowering the Fed funds rate. Two crucial ingredients are necessary for the DSGE model to generate these responses. First, as in Barsky and Sims (2009), the model requires forward-looking price setting and a high degree of wage rigidity. Second, monetary policy needs to react strongly to both inflation and output growth. This description of monetary policy is consistent with the Taylor rule as long as monetary policy reacts primarily to output growth instead of the output gap. In sum, augmenting DSGE models with news shocks and term structure variables provides valuable information to discipline the description of monetary policy and the structure of DSGE models in general.

The literature on the linkages of macroeconomics and the term structure is too large to survey in this paper. However, there are some closely related papers that are important to discuss in the context of our results. Evans and Marshall (2007) consider a series of macroeconomic shocks in a macro-finance VAR and assess their impact on the yield data. While they find important effects of their shocks on the level of the term structure, the shocks they identify do not provide a quantitative explanation for the majority of slope movements. This result in part motivates our approach of first finding the shocks that are quantitatively important for the slope, and then interpreting them. Evans and Marshall (2007) have already eliminated the 'usual suspects' from contention.

A natural potential source of movements in the term structure is shocks to monetary policy. Mechanically, if there is a monetary policy shock (and no reaction of long rates) then the slope will move. In an innovative study Piazzesi (2005) shows how to use high frequency data to trace the affect of monetary shocks onto yield data. Her work provides important insights into the nature of

how monetary shocks work their way into the yield curve. Yet, despite the rather large literature on the term structure there are few claims to providing structural shocks that quantitatively explain *most* of the movements in the slope. Indeed, the motivation of our empirical approach of extracting the one shock that is quantitatively the most important is the lack of consensus in the literature on the main sources of slope fluctuations. However, it is important to note that we do not rule out the possibility that other shocks play a role in slope movements. Our news shock, while a dominant driver of the slope, still leaves up to 40 percent of the variation unexplained. It would be interesting in future work to combine these shocks simultaneously in a DSGE model and do a fuller shock accounting exercise.

There are a few related papers using DSGE models to understand the importance of various macroeconomic shocks on the term structure. Both Rudebusch and Wu (2008) and Bekaert, Cho, Moreno (2010) combine small New-Keynesian models with no-arbitrage term structure models to investigate the role of various shocks on yields. Bekaert, Cho and Moreno (2010) conclude that monetary policy shocks explain a large portion of the movements in the slope. In contrast, De Graeve, Emiris and Wouters (2009) use a larger Smets-Wouters style DSGE model with many shocks (but no news shocks). They find a much smaller role for monetary shocks in explaining movements in the slope. Instead they find that demand shocks, defined as innovations to the intertemporal consumption Euler equation, explain up to 50% of movements in the slope. We interpret our results relative to these papers as follows. Rudebusch and Wu (2008) and Bekaert, Cho and Moreno (2010) have a small model with a few shocks. If there are missing shocks in the model, the Taylor rule setup can attribute movements in the short rate (which mostly drive the slope) to monetary shocks since these shocks are the residuals in the Taylor equation. We add an important structural shock into the model, lowering the residual claims of the monetary shock. De Graeve, Emiris and Wouters (2009), find a smaller role for monetary shocks, which is not inconsistent with our results (our news shock only accounts for at most 60 percent of slope movements, leaving at least 40 percent for other shocks.) We interpret their Euler equation shock as an measurement error to be explained, not as a deep structural shock. Our news shock is one with a clear economic motivation and interpretation, and we argue that this is the 'deep' structural explanation for slope movements.

The remainder of the paper proceeds as follows. Section 2 explains our empirical approach. Section 3 provides information about the data and VAR specification. Section 4 presents our empirical results. Section 5 examines the dynamics of term premia. Section 6 presents the term structure

DSGE model and estimates the model conditional on TFP news shocks. Section 7 concludes.

2 Identifying Structural Shocks: Two VAR Approaches

In this section we present two approaches to VAR identification. The first approach, proposed by Uhlig (2003), is purely statistical and extracts the largest 1 or 2 (or 3 or 4) shocks that explain the maximal amount of the forecast error variance (FEV) in a target variable, which in our case is the slope of the term structure. We then analyze what this shock does to the impulse response functions (IRFs) of the variables contained in the VAR in the hope of providing an economic interpretation of the shock.

The second identification approach is motivated by a key result from the first identification: news about future TFP play an important role in explaining movements in the slope. To assess this interpretation formally, we follow Barsky and Sims (2009) and Sims (2009) who extend the FEV maximization approach of Uhlig (2003) using TFP as the target variable and placing the extra restriction that the identified shock is orthogonal to contemporaneous TFP.

2.1 Review of VAR basics

We begin by discussing the general issue of identifying shocks in a VAR framework. This issue is well-known but we present these results for completeness and because the notation is useful for understanding the ensuing identification strategy. Consider a reduced-form VAR of the form

$$Y_t = B_1 Y_{t-1} + B_2 Y_{t-2} + \dots + B_p Y_{t-p} + u_t, \quad (1)$$

where Y_t is a $m \times 1$ vector of variables observed at time t ; and u_t is a $m \times 1$ vector of one-step-ahead prediction errors with variance-covariance matrix $E[u_t u_t'] = \Sigma$. Constant terms are dropped to save on notation. The objective is to impose restrictions on equation (1) to identify structural shocks; i.e. innovations that are mutually orthogonal to each other. Identifying all m shocks in the VAR requires a minimum of $m(m-1)/2$ restrictions. However, it is well-known that one can instead place restrictions to identify fewer than m shocks.

In order to more clearly see the identification issue it is useful to rewrite equation (1). Under the assumption that Y_t is covariance-stationary, we can invert this VAR to express it as a moving

average process

$$Y_t = [B(L)]^{-1}u_t = C(L)u_t, \quad (2)$$

where $B(L) \equiv I - B_1L - \dots - B_pL^p$, and $C(L) \equiv I + C_1L + C_2L^2 + \dots$

This moving average representation is of course the impulse response function for the VAR. Identification of the structural shocks amounts to decomposing the vector of prediction errors u_t into m mutually orthogonal innovations v_t with normalized variance-covariance matrix $E[v_tv_t'] = I$. In other words, in identifying VAR shocks we are trying to find a mapping A between the reduced-form and structural shocks (i.e. $u_t = Av_t$).

In this mapping the i -th column of the $m \times m$ matrix A describes the contemporaneous effect of the i -th innovation in the structural shock vector v_t on the different variables in Y_t . By definition, A needs to satisfy $\Sigma = E[Av_tv_t'A'] = AA'$. This restriction, however, is not sufficient to identify A because for any matrix A , there exists some other matrix $\tilde{A}$ that satisfies the restriction that the covariance matrix be respected. This alternative matrix provides a different map from u_t into $\tilde{v}_t$; i.e. $u_t = \tilde{A}\tilde{v}_t$.³ Thus, the set of statistically valid 'structural' identifications of the VAR is quite large. To choose which identification restriction to use, one then typically uses some sort of economic theory. One prominent example is to use the Cholesky decomposition to restrict A to be lower triangular. Economically, this amounts to ordering the variables in the VAR in terms of the timing with which variables can respond to various structural shocks.

2.2 Extracting the Most Important Shocks

An alternative to the traditional approach of placing economic restrictions to identify a shock and then checking to see if the shock is important is to move in the reverse direction, as proposed by Uhlig (2003). This approach to identification is purely statistical and consists of finding the innovation(s) that explain(s) as much as possible of the FEV of some variable in Y_t over a chosen horizon $\underline{k}$ to $\bar{k}$. One then tries to provide an economic interpretation of the shock (conditional on it being important) by studying the full set of IRFs. As part of this procedure one learns how many shocks are needed to explain a given variable. That is, do we need a DSGE model with many shocks, or is a more parsimonious model able to explain a given time series?

³To see this, consider an orthogonal matrix Q with $QQ' = I$ and define $A = \tilde{A}Q$ and $Qv_t = \tilde{v}_t$. Then, $\Sigma = E[\tilde{A}Qv_tv_t'Q'\tilde{A}'] = E[\tilde{A}\tilde{v}_t\tilde{v}_t'\tilde{A}'] = \tilde{A}\tilde{A}'$ because $E[\tilde{v}_t\tilde{v}_t'] = QE[v_tv_t']Q' = QQ' = I$.

More formally, Uhlig's procedure searches for the n largest shocks to explain the FEV of one variable in the VAR. Thus we need to find the $m \times n$ submatrix A_1 for the n most important innovations in v_t such that $A = [A_1 \ A_2]$ with $AA' = \Sigma$ for some $m \times (m - n)$ submatrix A_2 . Given an initial decomposition $\tilde{A}$, this amounts to computing $A_1 = \tilde{A}Q_1$ where Q_1 is the $m \times n$ partition Q_1 of $Q = [Q_1 \ Q_2]$ that satisfies our statistical criteria.

To find Q_1 , we let $\tilde{A}\tilde{A}' = \Sigma$ be the Cholesky decomposition of the reduced for VAR covariance matrix.⁴ We then define the impulse responses $\tilde{R}(L)$ associated with the innovations $\tilde{v}_t$ identified by this decomposition as

$$\tilde{R}(L) = C(L)\tilde{A} = \tilde{R}_0 + \tilde{R}_1L + \dots,$$

with $\tilde{R}_0 = \tilde{A}$. The impulse responses associated with the targeted innovations v_t are thus given by

$$R(L) = C(L)\tilde{A}Q = \tilde{R}(L)Q.$$

The k -step ahead forecast error of Y_{t+k} is then given by

$$u_{t+k}(k) = Y_{t+k} - E_{t-1}[Y_{t+k}] = \sum_{l=0}^k \tilde{R}_l Q v_{t+k-l},$$

and its variance-covariance matrix is given by

$$\begin{aligned} \Sigma(k) &= \sum_{l=0}^k \left(\tilde{R}_l [q_1 \ q_2 \dots q_m] \right) \left(\tilde{R}_l [q_1 \ q_2 \dots q_m] \right)' \\ &= \sum_{l=0}^k \left[(\tilde{R}_l q_1)(\tilde{R}_l q_1)' + (\tilde{R}_l q_2)(\tilde{R}_l q_2)' + \dots + (\tilde{R}_l q_m)(\tilde{R}_l q_m)' \right] \\ &= \sum_{j=0}^m \sum_{l=0}^k (\tilde{R}_l q_j)(\tilde{R}_l q_j)', \end{aligned}$$

where q_i , $i = 1 \dots m$ are the $m \times 1$ column vector partitions of Q . The term $\sum_{l=0}^k (\tilde{R}_l q_j)(\tilde{R}_l q_j)'$ thus describes the contribution of the j -th orthogonal shock to the variance-covariance matrix $\Sigma(k)$ of the k -step ahead forecast error $u_{t+k}(k)$. This division into m parts is possible because the v_t are *iid* innovations and the columns q_i , $i = 1 \dots m$ are orthogonal.

⁴Any other triangular factorization would do as well, but the Cholesky is particularly easy to implement.

Our objective is to find the innovation(s) that explain(s) as much as possible of the sum of the k -step ahead forecast error variance of the i -th variable in Y over some horizon $\underline{k} \leq k \leq \bar{k}$

$$\sigma_i^2(\underline{k}, \bar{k}) = \sum_{k=\underline{k}}^{\bar{k}} \Sigma(k)_{ii}.$$

Formally, to identify the most important innovation, we want to find the orthogonal vector q_1 with length 1 (i.e. $q_1' q_1 = 1$) that maximizes⁵

$$\begin{aligned} \sigma_i^2(\underline{k}, \bar{k}; q_1) &= \sum_{k=\underline{k}}^{\bar{k}} \sum_{l=0}^k \left[(\tilde{R}_l q_1) (\tilde{R}_l q_1)' \right]_{ii} \\ &= \sum_{k=\underline{k}}^{\bar{k}} \sum_{l=0}^k \text{trace} \left[E_{(ii)} (\tilde{R}_l q_1) (\tilde{R}_l q_1)' \right] \\ &= q_1' S q_1, \end{aligned} \tag{3}$$

with

$$S \equiv \sum_{k=\underline{k}}^{\bar{k}} \sum_{l=0}^k \tilde{R}_l' E_{(ii)} \tilde{R}_l,$$

where $E_{(ii)}$ is a matrix with zeros everywhere except for the i, i -th position; and where the definition of S takes advantage of the fact that $\text{trace} \left[E_{(ii)} (\tilde{R}_l q_1) (\tilde{R}_l q_1)' \right] = \text{trace}[(\tilde{R}_l q_1)' E_{(ii)} (\tilde{R}_l q_1)]$. This maximization problem can thus be written as a Lagrangian

$$L = q_1' S q_1 - \lambda(q_1' q_1 - 1) \tag{4}$$

with first-order condition

$$S q_1 = \lambda q_1.$$

Inspection of this solution reveals that this is simply the definition of an eigenvalue decomposition, with q_1 being the eigenvector of S that corresponds to the eigenvalue λ . Furthermore, since $q_1' q_1 = 1$, we can rewrite the first-order condition as $\lambda = q_1' S q_1 = \sigma_i^2(\underline{k}, \bar{k}; q_1)$. The partition q_1 that maximizes the variance is therefore the eigenvector associated with the largest eigenvalue λ ; i.e. q_1 is the first principal component of S . Likewise, q_2 is the second principal component and so forth for all the n

⁵For notational convenience, we order this vector first in Q .

components of Q_1 that we want to extract. The submatrix A_1 that we seek is then

$$A_1 = \tilde{A}Q_1.$$

2.3 Identifying News Shocks

Following Barsky and Sims (2009) and Sims (2009) (hereafter BSS), the shock that we seek to identify is news about future TFP. In their procedure, TFP is placed in a VAR with a selection of other macroeconomic variables. The assumption underlying the identification procedure is that TFP is an exogenous process. Therefore shocks to other variables in the system, such as monetary policy shocks, will not impact TFP at any horizon. Furthermore, it is assumed that TFP is driven by two shocks. One is an unforecastable shock to current TFP, while the second is a shock that represents news about future TFP.

The BSS identification approach extends Uhlig’s (2003) approach with additional restrictions. To implement the procedure we choose TFP as the variable in the VAR for which we would like extract the shocks to maximize the amount of the FEV explained. Further, the number of shocks is restricted to two (i.e. $n = 2$ in section 2.2). We then impose the additional restriction on the Lagrangian in (4) that the news shock must have zero impact on TFP contemporaneously.⁶ In other words, a news shock is defined as the innovation that explains most of future movements in TFP but nothing of current TFP. The other shock is then necessarily a contemporaneous TFP shock and can be identified by a Cholesky decomposition with TFP ordered first in the VAR.⁷

3 Data and VAR Estimation Procedure

The VAR we estimate combines term structure and macroeconomic variables. For the term structure data we use two time series. The first is the Federal Funds rate. The second is the term spread which is measured as the difference between the 60-month Fama-Bliss unsmoothed zero-coupon yield from the CRSP government bonds files and the Federal Funds rate. We choose the 60-month

⁶In addition, one could impose that two shocks, the news shock and innovation to current productivity, account for all of the FEV of TFP. Following BSS, we do not impose this restriction explicitly. Yet, as it turns out, two shocks account for virtually all of the FEV of TFP for all horizons.

⁷At any point in time TFP moves for three possible reasons. First, there may be current innovations to productivity. Second, past news shocks are realized as subsequent movements in productivity. Third, past productivity innovations propagate forward through the lag structure in the VAR.

yield as our long rate because it is available back to 1959:2, whereas longer-term yields such as the 120-month yield become available only in the early 1970s. We use the Federal Funds rate as the short term rate in order to be consistent with the macroeconomic model that we examine in Section 6. The DSGE model does not differentiate between the monetary policy rate and the short-end of the Treasury yield curve (e.g. a 3-month bill rate).⁸ To check for robustness, we ran our simulations with alternative measures of the slope and the short rate and found all of the main results to be unchanged.⁹

For the macroeconomic data we use two datasets. The first is a small set of macroeconomic variables consisting of TFP, consumption and inflation. Our measure of TFP is a quarterly version of the series constructed by Basu, Fernald and Kimball (2006). This series exploits first-order conditions from a firm optimization problem to correct for unobserved factor utilization and is thus preferable to a simple Solow residual measure of TFP.¹⁰ Our measure of consumption is the log of real chain-weighted personal consumption expenditures. For inflation, we use the growth rate of the GDP deflator.

The second dataset is a larger dataset that adds four variables to our smaller dataset. These variables are the log of real chain-weighted GDP, the log of real chain-weighted gross private domestic investment and the log of the S&P 500 composite index deflated by the consumer price index.

All of the macroeconomic series are obtained from the FRED II database of the St. Louis Fed and are available in quarterly frequency. The term structure and stock market data are available in daily and monthly frequency. We convert them to quarterly frequency by computing arithmetic averages over the appropriate time intervals. The sample period is 1959:2-2005:2 (with the start date limited by the availability of 60-month yield). Both the baseline VAR and the extended VAR are estimated in levels with 4 lags of each variable. To improve precision, we impose a Minnesota

⁸This approximation seems reasonable since in practice, the Federal Funds rate and short-end bill rates move very closely together. More precisely, the correlation coefficient of the Federal Funds rate and the 3-month bill rate over the 1959:2-2005:2 period is 0.984. The Federal Funds rate is slightly more volatile and has a higher mean than the 3-month bill rate. For our VAR and DSGE exercises, these differences are not important.

⁹There are two important alternative measures of the slope. First, we replaced the 60-month yield with the 120-month zero-coupon yield as computed by Gurkaynak, Sack and Wright (2007) and the Federal Funds rate by the 3-month bill rate. Second, we used a Nelson-Siegel style slope factor as computed in Diebold and Li (2006).

¹⁰Basu, Fernald and Kimball (2006) also make use of industry level data to correct for differences in returns to scale. Since this industry level data is available only on an annual basis, our quarterly TFP measure does not include this returns to scale correction. See Sims (2009) for details.

prior on the estimation and compute error bands by drawing from the posterior.¹¹

4 What Moves the Slope of the Term Structure?

In this section we answer the main question of the paper. We do so by first extracting the shocks that explain most of the movements in the slope of the term structure, our 'target' variable in the VAR. Second, we look for different possible interpretations of this shock. In particular, we pursue the hypothesis that this shock captures news about future innovations to TFP. Third, we show that our results are robust to a variety of alternative assumptions.

4.1 Slope Shocks

As described in Section 2, we extract the shocks that maximizes the fraction of the FEV of the slope that is explained by those shocks. We set the forecast horizon to $0 \leq k \leq 40$ quarters, weighing the importance of each of the forecasts equally. This choice is motivated by the fact that we want to capture short- and medium-run movements in the term structure slope while providing at the same time reliable estimates at the long end of the forecasting horizon. We limit our analysis to two shocks ($n = 2$) because we find that two shocks explain virtually all the movements in the slope. The following results refer to the small VAR described above.

Figure 1 displays the fraction of the FEV of the different variables explained by the first shock. The solid line corresponds to the median estimate, while the dotted lines denote the 16%-84% error bands. As the top left panel shows, this first shock explains more than 85% of all slope movements over the entire 0 to 40 quarter forecast horizon. The second shock (not shown) accounts for virtually all of the remaining fraction of the FEV of the slope. This result is robust across many different VAR specifications. For example, in the extended VAR that we examine at the end of this section, one shock explains about 75% of all slope movements and the second shock accounts for almost all of the remaining 25%. In other words, two shocks are sufficient to understand all movements in the slope and to an approximation, the first shock is by far the most relevant. We thus focus on the

¹¹We performed a battery of robustness checks with other macroeconomic variables including data that allowed us to estimate the VAR on monthly frequency. We discuss the responses of some of the added variables in the next section but note that none of the main conclusions is affected by the different changes in VAR specification. Also, we dropped the Minnesota prior and estimated the VAR with OLS instead, computing the error bands by bootstrapping from the estimated VAR. Details are available from the authors upon request.

properties of this first shock only.¹²

The other panels in Figure 1 show that the slope shock also explains about 50% of the Federal Funds rate over the entire horizon, suggesting that slope movements are largely driven by variations in the short end of the term structure. For the macroeconomic variables the slope shock explains very little of variations in TFP, consumption and inflation at short horizons. As the forecast horizon increases, however, the slope shock gradually accounts for a larger fraction of the movements in these variables. In particular, the shock explains more than 40% of the consumption variation at a 20 quarter horizon and about 30% of TFP variations 40 quarters ahead (with this latter fraction increasing towards 50% for forecast horizons beyond 40 quarters). This confirms earlier findings by Ang and Piazzesi (2003), Diebold, Rudebusch and Aruoba (2006) and Evans and Marshall (2007) that there are important linkages between slope movements and macroeconomic fluctuations. Our analysis adds the qualification, however, that these linkages are mostly present for medium- and longer-term macroeconomic fluctuations whereas high-frequency variations in macroeconomic variables are almost completely orthogonal to slope innovations.

The second step in our approach is to provide an economic interpretation of the slope shock. We do this by examining the IRFs of the different variables to an innovation in the slope shock. Figure 2 displays the results. The term spread increases on impact of the shock, while the long end of the term structure remains roughly constant on impact before becoming slightly negative.¹³ The strong reaction of the spread is driven largely by the drop in the Federal Funds rate. Interestingly, the slope shock has no significant impact on either TFP or consumption on impact, but within 2 quarters of the shock, both of these variables start to increase significantly to what appears to be a permanently higher level. Finally, inflation drops significantly on impact of the slope shock and remains below its initial rate for more than 2 years. This drop in inflation is, however, smaller than the drop in the Federal Funds rate, implying that the *real* Federal Funds rate also turns negative.

How do we interpret this shock? The apparent permanent response of TFP and consumption

¹²As explained in section 2, identifying the two most important shocks amounts to finding the innovations associated with the two largest eigenvalues of the matrix S defined in (3); i.e. the first shock corresponds to the eigenvector with the largest eigenvalue and the second shock corresponds to the eigenvector with the second largest eigenvalue. As Uhlig (2003) explains, however, there are two other pairs (or rotations) of eigenvectors that explain together an equal fraction of the total FEV of the slope. We compared all of our results with these two other rotations and found that for each rotation, there exists one shock that explains over 50% of slope movements. This shock has very similar properties than the first shock we consider in the text.

¹³The long bond rate is not in our estimated VAR. Its IRF is constructed as the sum of the term spread and the Federal Funds rate.

suggests that the slope shock captures technological innovations leading to an increase in productive capacity in the future. Such a supply-side interpretation also rationalizes why, despite the loosening of monetary policy, inflation falls and remains persistently lower for more than two years. More specifically, the macroeconomic dynamics in Figure 2 look very much like the responses to a news shock about future TFP as identified in Barsky and Sims (2009) and Sims (2009). These two papers report that TFP news shocks lead to a delayed but permanent increase in TFP and consumption and a sharp drop in both inflation and short-term interest rates. Both papers also find that TFP news shocks explain almost none of high-frequency variations in TFP and consumption but account for 40% or more of the two variables at horizons of 20 quarters or more.

Before examining this TFP news interpretation in more detail, it is important to consider whether other prominent macroeconomic shocks are consistent with these IRFs. Monetary shocks are often considered in both macroeconomic studies as well as term structure studies (e.g. Piazzesi 2005). The monetary shock interpretation appears clearly inconsistent with the IRFs in Figure 2. If the drop in the Federal Funds rate was related to an *exogenous* monetary policy intervention, then inflation should increase rather than decrease and there should be no permanent effect on either consumption or TFP (e.g. Christiano, Eichenbaum and Evans, 2005). Our technology news hypothesis, by contrast, implies that monetary policy reacts *endogenously* to the drop in inflation and is thus only indirectly the main driver of the slope.¹⁴ Taken together, the IRFs to a slope shock are inconsistent with a monetary shock interpretation.

A second type of shock considered in the macroeconomics literature are demand shocks, either in the form of exogenous changes in government deficits (Evans and Marshall, 2007; Dai and Phillippon, 2008) or exogenous changes to the effective interest rate that applies to savings and investment decisions (De Graeve, Emiris and Wouters 2008). Similar to exogenous monetary policy shocks, such demand shocks should not have a permanent positive effect on either consumption or TFP. Likewise, we know of no theory of demand shocks that produces a prolonged decline in both inflation and the Federal Funds rate in response to a positive demand shock.

A third type of shock from the macro-labor literature is a shock to the marginal rate of substitution (MRS) between consumption and leisure. Evans and Marshall (2007) study the impact of this shock on the term structure and find that this shock has a statistically insignificant affect

¹⁴This result is consistent with Evans and Marshall (2007) who also conclude that the systematic reaction of monetary policy is an important channel through which macroeconomic shocks affect the term structure.

on the slope and inflation while increasing both real activity and the Federal Funds rate. These predictions are inconsistent with the IRFs in Figure 2. We conclude that MRS shocks cannot be an interpretation of our slope shock.¹⁵

A fourth type of macroeconomic shocks is a contemporaneous innovation to TFP as traditionally assumed in the business cycle literature. We identify a contemporaneous TFP shock by ordering TFP first in our VAR and extracting the first column of a Cholesky decomposition. Figure 3 display IRFs to this contemporaneous TFP shock. Notably, TFP rises dramatically on impact and so does consumption. Both of these responses are persistent but ultimately transitory. Furthermore, the shock has no significant effect on the term spread and only a delayed but negligible effect on the Federal Funds rate. All of these IRFs are inconsistent with our slope shock, thus suggesting that it is indeed *news about future productivity innovations* that are a main driver of the slope of the term structure.

4.2 Slope Shocks are News Shocks

We now refine the TFP news interpretation of the slope shock by formally identifying a news shock. News shocks about future productivity have recently been resuscitated as a potential source of business cycle fluctuations in recent work by Beaudry and Portier (2006), Jaimovich and Rebelo (2009) and Schmidt-Grohe and Uribe (2008), among others. A news shock is information about the future level of TFP. Beaudry and Portier (2006) model the process for TFP as:

$$a_t = v_t + D_t, \tag{5}$$

where a_t is the log of TFP; and v_t and D_t are two independent exogenous components. The component v_t captures potentially persistent but transitory surprise movements in TFP. The component D_t is non-stationary and assumed to follow a distributed lag process in *past* innovations; i.e. $D_t = d(L)\eta_t$ with $d(0) = 0$ and $d(1) = 1$. Innovations η_t are interpreted as news shocks about future productivity because they do not affect TFP contemporaneously, but only with a delay of one or more periods.¹⁶

¹⁵Evans and Marshall (2007) find that MRS shocks are primarily important for variations in the level of the term structure but have no significant impact on the slope.

¹⁶TFP news shocks resemble recent characterizations of technological adoption by Rotemberg (2003) or Comin and Gertler (2006). While neither of these papers imposes the zero restriction on impact of the shock, they both

Rather than following the empirical approach of Beaudry and Portier (2006) who identify news shock with a mix of short- and long-run restrictions on stock prices and TFP, we adopt the more recent identification approach proposed by BSS. As described in the previous section, the BSS approach is similar in spirit to our statistical extraction of the slope shock and consists of identifying the shock that explains most of TFP variations over a given forecast horizon but is orthogonal to contemporaneous innovations in TFP. As such, the BSS identification satisfies the definition of news about future productivity (i.e. $d(0) = 0$) in Beaudry and Portier (2006) and also allows for news to have a permanent effect on TFP (i.e. $d(1) = 1$ is possible but not required).

Figure 4 displays the fraction of the FEV of the variables in the VAR explained by the TFP news shock. As we found for the slope shock, the TFP news shock explains almost none of the movements in macroeconomic variables on impact but up to 50% of consumption variations after 20 quarters and about 40% of TFP variations after 40 quarters. The shock also explains over 60% of term spread movements over all horizons and between 60% and 80% of Federal Funds rate movements. In other words, the TFP news shock seems to be a major driver of term structure movements.¹⁷

Figure 5 reports the IRFs of the different variables to the TFP news shock (solid blue lines) and reproduces the IRFs to the slope shock from Figure 2 for comparison (dashed red lines). The similarity in results is striking. In particular, the slope jumps up significantly on impact and then returns back to its pre-shock value after 10 to 15 quarters; TFP increases gradually from zero (by construction for the news shock identification) to a permanently higher level (even though no constraint on long-run effects is imposed); consumption increases slightly (but insignificantly) on impact and then gradually increases to a permanently higher level; and both inflation and the Federal Funds rate drop markedly on impact and remain below their initial value for more than 15 quarters. As with the slope shock, the drop in the Federal Funds rate is larger than the drop in inflation, implying a decline in the real Federal Funds rate.

To further illustrate the correspondence between the TFP news shock and our slope shock, we

argue that it takes on average several years for new technologies to be adopted even though these innovations are known to exist and be commercially valuable. See Rotemberg (2003) for an extensive discussion of evidence about the slow diffusion of technological innovations.

¹⁷Figure 4 also shows that the TFP news shock explains almost nothing of high-frequency fluctuations in inflation and only about 20% of inflation fluctuations at horizons beyond 5 quarters. Barsky and Sims (2009), by contrast, report that TFP news shocks explain more than 60% of high-frequency variations in inflation and between 40-55% at horizons of 4 quarters and higher. This difference is due to the fact that they compute inflation from the CPI deflator while we use the GDP deflator. We prefer the latter because it represents a broader measure of aggregate prices, does not suffer from substitution bias, and is less affected by large temporary swings in food and energy prices.

extract the time series of each of the two shocks and plot them together. As Figure 6 shows, the slope shock is slightly more volatile than the TFP news shock but overall, the two shocks move closely together. In fact, the correlation of the two is 0.87. This close correspondence is surprising because the identification criteria behind the two shocks are very different from each other. The slope shock is extracted by maximizing the FEV of the *slope* while the TFP news shock is extracted by maximizing the FEV of *TFP* subject to the additional constraint that the shock is orthogonal to contemporaneous TFP movements. Hence, there would be no a priori reason to believe that the two innovations capture the same economic shock.

Finally, to assess the empirical relevance of the TFP news shock, Figure 7 plots the historical time series of the different variables in the VAR against the simulated times series conditional on the TFP news shock (i.e. assuming that TFP news shocks are the only stochastic innovation). As the two top panels show, TFP news shocks explain very little of the high-frequency fluctuations in TFP and consumption, which is consistent with our conclusions from the FEV decompositions. TFP news shocks also miss most of the high-frequency variations in inflation but capture quite a lot of the medium-frequency movements in inflation, especially during the 1970s and early 1980s. TFP news shocks do a surprisingly good job in accounting for fluctuations in the Federal Funds rate and the slope. In particular, TFP news shocks account for almost all of the large swings in the slope during the 1970s and also rationalize the increase in term spreads during the early 1990s and the early 2000s. This close fit is striking and leads to two important lessons. First, a large part of Federal Funds rate fluctuations are driven by news about future supply-side innovations. Second, through the endogenous response of the Federal Funds rate, TFP news shocks are a main driver of the slope of the term structure.

4.3 Robustness

In this section we show that our empirical results are robust to a number of possible concerns. The first potential issue with our results concerns mismeasurement of technological progress. In particular, advances in technology may not come through increases in TFP but rather through technological progress that is embodied in new capital. Hence, if capital services are not appropriately measured, our identification may mistake embodied (i.e. capital-specific) technological progress for TFP improvements. This concern is motivated by recent empirical evidence from Fischer (2005)

who reports that embodied technological shocks are a main driver of business cycle fluctuations. To address this issue, we add Fischer’s (2005) relative price deflator series for investment and equipment goods to our VAR and rerun both the slope shock identification and the TFP news shock identification.¹⁸ In response to the slope shock, both relative price deflators increase slightly on impact and then decrease significantly after about 10 quarters to a permanently lower level. In response to the TFP news shock, by contrast, neither of the relative price deflators reacts significantly. All of the other results remain unaffected. This suggests, on the one hand, that TFP news shocks are not erroneously capturing capital-specific embodied technological progress. On the other hand, the slope shock seems to pick up not only news about future TFP increases, but also news about future embodied technological progress. This could be one of the reasons why the extracted slope shock is slightly more volatile than the TFP news shock.

A second issue is the extent to which our results are robust to alternative VAR specifications. We estimated many different VAR specifications and found our results to be generally robust. For space reasons, we report here only one of these alternative specifications, which extends the baseline VAR with output, investment and the S&P 500 composite index. We choose this particular extension because it allows us draw comparisons with the recent empirical literature on news shocks and because it provides a useful benchmark for the DSGE model that we introduce in Section 6. Figure 8 reports the IRFs to the TFP news shock for this extended VAR.¹⁹ As in the smaller VAR, the TFP news shock has a gradual but permanent effect on real variables. Consumption now increases somewhat on impact of the shock. Output declines slightly on impact, but the change is not very significant. Investment, by contrast, contracts significantly over the first two periods. The real stock market index increases on impact and remains significantly higher for about four years before slowly returning back to its initial value. Finally, both inflation and the Federal Funds rate drop markedly on impact and remain persistently below their initial value for 15 to 20 quarters. Since the long rate barely moves, the term spread increases on impact of the shock and then gradually returns to its average value. Overall, these results look very similar to the results obtained above

¹⁸The relative price series we use are updated by DiCecio (2008).

¹⁹In the interest of conciseness, we do not plot the fraction of FEVs explained by TFP news shocks for the different variables of this extended VAR. Interestingly, the TFP news shock accounts for an even larger fraction of TFP and consumption movements at the 20-40 quarter horizon. Similarly, the shock explains almost nothing of output and investment fluctuations on impact but about 50% of both variables after 20 quarters and more. For the term structure, in turn, the shock explains between 40% and 50% of movements in the slope and the Federal Funds rate over the entire horizon. This is somewhat less than in the baseline VAR but still very sizable. Finally, the TFP news shock explains roughly 20% of inflation and stock market movements over the entire horizon.

with the baseline VAR.

The small inverse reaction of consumption and output on impact of the TFP news shock matches closely the findings in Sims (2009).²⁰ At the same time, these results contradict Beaudry and Portier (2006) who find that consumption and real activity (measured by either hours or investment) *both* display large positive reactions almost immediately after a TFP news shock. Furthermore, Beaudry and Portier's (2006) TFP news shocks account for a large part of the high-frequency fluctuations in real aggregates whereas this is not the case in our analysis. As Sims (2009) shows, this difference in results is due to the different identification approach employed by Beaudry and Portier (2006). They identify a TFP news shock either as the VAR innovation that may have a permanent long-run effect on TFP, or the innovation that is orthogonal to current TFP but may affect stock prices contemporaneously.

There are several advantages to the BSS identification approach over the ones employed by Beaudry and Portier (2006). First, long-run restrictions have been shown to have very poor finite sample properties (e.g. Faust and Leeper, 1997). According to Monte-Carlo simulations by Francis, Owyang and Roush (2007), FEV-based criteria such as the one employed by BSS perform significantly better at identifying technology shocks. Second, the BSS approach imposes that news shocks account for the maximum variation in TFP over the entire short- to medium-run horizon. The BSS approach is thus more inclusive than a long-run restriction and directly addresses the problem that shocks identified by long-run restrictions commonly account for only a modest fraction of TFP fluctuations even at forecast horizons of 10 or more years. Third, stock prices react to many different shocks and are thus a relatively uninformative measure of future technology innovations. For all these reasons, we prefer the BSS approach to identify TFP news shocks.

5 Term Premia versus the Expectations Hypothesis

Our VAR framework allows us to decompose the reaction of the long rate into variations due to term premia and expectations about future short rates (i.e. the Expectations Hypothesis). We can

²⁰Sims (2009) also reports that hours worked decline for the first few quarters after the TFP news shock. We find the same result if we include hours worked as an additional variable in the VAR.

decompose the yield on a T -period yield R_t^T (in our case the 60-month yield) as

$$R_t^T = \frac{1}{T} \sum_{i=0}^{T-1} E_t R_{t+i} + tp_t, \quad (6)$$

where the $E_t R_{t+i}$ are time t expectations of future short rates; and tp_t denotes term premia. The reaction of the long-rate with respect to TFP news shocks may be relatively small either because the Expectations Hypothesis part $1/T \sum_{i=0}^{T-1} E_t R_{t+i}$ and the term premia part do not respond strongly or because variations in the two almost cancel each other out. This, in turn, determines the importance of term premia fluctuations for the reaction of our slope measure.

The technical difficulty with this decomposition is that term premia are inherently unobservable. Here, we follow Campbell and Shiller (1987, 1991) and use our larger VAR to compute expectations of short rates conditional on TFP news shocks. The term premia response to the TFP news shock is then simply the difference between the actual long rate response to the shock and the response as implied by the Expectations Hypothesis computed from the VAR. Figure 9 shows the resulting decomposition both for the long rate (top panels) and the spread (bottom panels). As the top panels show, term premia react positively to a TFP news shock before turning slightly negative after about 10 quarters. This initial jump in term premia occurs because the long rate under the Expectations Hypothesis displays a more marked drop on impact than the actual long rate. Or put differently, the reaction of the long rate to the TFP news shock is relatively small because term premia variations neutralize almost all of the initial drop in the long rate as implied by expectations of future short rates. As the bottom panels show, this implies that almost half of the initial jump in the slope is due to the increase in term premia. The response of the slope under the Expectations Hypothesis (i.e. as implied by short-rate fluctuations) remains, however, significant and returns only gradually back to its initial value. Hence, the endogenous reaction of the Federal Funds rate remains a quantitatively important direct channel through which TFP news shocks affect the slope.

The large and significant reaction of term premia is consistent with the general statistical rejection of the Expectations Hypothesis in the finance literature (e.g. Fama and Bliss, 1987; Campbell and Shiller, 1991; Cochrane and Piazzesi, 2005). At the same time, the Expectations Hypothesis by itself can account for more than half of the slope response to a news shock and thus remains empirically relevant, which is consistent with the basic message of Campbell and Shiller (1987) and more recently King and Kurmann (2002).

6 A DSGE-News Model of the Term Structure

Our final exercise is to evaluate how well a medium-scale New Keynesian DSGE model along the lines of Smets and Wouters (2007) can account for term structure movements in response to news shocks. It is important to note that while the type of DSGE models we analyze provides a good fit of macroeconomic quantities to many types of shocks, these models fail to generate sizable and variable term premia (e.g. Rudebusch and Swanson, 2008).²¹ Given the importance of time-varying term premia for term structure dynamics, this clearly limits our exercise. At the same time, there are several reasons why our exercise remains interesting. First, as documented above, the Expectations Hypothesis remains quantitatively important and accounts for more than half of the response of the slope to a news shock. Second, the model is a useful framework to provide economic intuition to interpret the VAR results. Third, the results offer guidance for future research on macro-finance models of the term structure.

6.1 Model

The model is very similar to the one presented in Smets and Wouters (2007) and contains several real and nominal frictions. Specifically, the model features sticky nominal price and wage setting that allow for indexation to lagged inflation, external habit persistence in consumption, investment adjustment costs, variable capital utilization and fixed costs in production. The main structural difference of our model to the one presented in Smets and Wouters (2007) is that we specify TFP as an exogenous process with a stochastic trend, driven by both a contemporaneous shock and a news shock; i.e.

$$\mu_t = \gamma\mu_{t-1} + \varepsilon_{1,t} + \varepsilon_{2,t-j}, \quad (7)$$

where $\mu_t = a_t - a_{t-1}$ is the growth rate of TFP; $\varepsilon_{1,t}$ is the contemporaneous shock; and $\varepsilon_{2,t-j}$ is the news shock. This news shock impacts actual TFP in period t but is known j periods in advance. Both the contemporaneous shock and the news shock are *i.i.d.* processes with mean zero and variance $\sigma_{\varepsilon_1}^2$ and $\sigma_{\varepsilon_2}^2$. In our VAR, TFP begins to react after the first period after the news shock so we set $j = 1$.

We also use a different monetary policy rule from Smets and Wouters (2007). We describe

²¹Earlier work documenting the bond premium puzzle in more basic DSGE models are Donaldson, Johnson and Mehra (1990) and Den Haan (1995).

monetary policy by an interest rate rule that allows for a separate response of the short term nominal rate R_t to both the output gap $y_{gap,t}$ and output growth Δy_t

$$R_t = \rho R_{t-1} + (1 - \rho)[\theta_\pi E_t \pi_{t+1} + \theta_{ygap} y_{gap,t} + \theta_{\Delta y} \Delta y_t], \quad (8)$$

where the output gap is defined as the difference between actual output and potential output if there were no nominal price and wage rigidities.

We also need to append a long term bond pricing equation to the model. Since this class of models does not generate sizable variations in term premia (see Rudebusch and Swanson, 2008), we confine our analysis to a loglinear environment where term premia are by definition zero. In order to obtain time-varying term premia, we would have to analyze at least a third-order approximation of the model. While this is technically possible it is known that even these approximations do not deliver sufficiently volatile term premia. In our model we impose that the Expectations Hypothesis holds exactly and thus, the long-bond yield is simply equal to the expectational part in (6).

The full set of log-linearized model equations is available by request. The Rational Expectations equilibrium of the resulting system is computed using the numerical methods of King and Watson (1998).

6.2 Estimation

We partition the parameters of our model into two groups. The first group consists of parameters that we calibrate to match long-run moments of the data. The second group is estimated to match the impulse responses to a news shock from our empirical VAR. All values reported are for a quarterly frequency.

The first three values imply a labor share of approximately 0.7 as reported by Gollin (2002); an average annualized real interest rate of 3 percent; and an annual depreciation rate of 10 percent. The unit elasticity of labor supply is a compromise between values suggested in the microeconomic and macroeconomic literatures. The elasticity of substitution across goods implies an average markup for final goods producers of 11% following Basu and Fernald (1997). The elasticity of substitution across labor is set as in Smets and Wouters (2007). The growth rate of TFP μ is set to match the average growth rate of real GDP (1.81% annually). Finally, the fixed cost parameter in production (not reported here) is set so that economy-wide net profits are zero as suggested by Basu and

Fernald (1994) or Rotemberg and Woodford (1995).

The second group of parameters is estimated by minimizing a weighted distance between the model-implied IRFs to a news shock and the empirical counterparts from the larger VAR. Specifically, denote by $\hat{\Psi}$ a vector of empirical IRFs to a news shock over obtained from a VAR. Likewise, denote by $\Psi(\zeta)$ the same vector of IRFs implied by the model, where ζ contains all the structural parameters of the model. The estimator of some parameter subset $\zeta^* \subseteq \zeta$ is the solution to

$$\hat{\zeta}^* = \arg \min_{\zeta^*} \left[\hat{\Psi} - \Psi(\zeta) \right]' \Omega^{-1} \left[\hat{\Psi} - \Psi(\zeta) \right],$$

where Ω is a diagonal matrix with the sample variances of $\hat{\Psi}$ along the diagonal. This limited-information approach is very similar to the one implemented by Christiano, Eichenbaum and Evans (2005) and Altig, Christiano, Eichenbaum and Lindé (2004). Here, we adapt it for our purposes by first estimating the parameters governing the response of TFP to an exogenous news shock and then, in a second step, by estimating the remaining structural model parameters such as to match the IRFs of other variables in the VAR. We adopt this two step approach because we want to evaluate the ability of our model to generate realistic term structure and macroeconomic dynamics to a news shock given the evolution of observed TFP.

We include the IRFs of TFP and all five macroeconomic variables in the objective function. In addition, since our model implies constant term premia, we replace the IRFs of the observed spread and long rate with their counterparts under the Expectations Hypothesis. This leaves us with a total of eight empirical IRFs. For each of these IRFs, we include the 40 quarter horizon in the estimation criteria.

6.3 Results

Our main objective is to see if a DSGE model can match our VAR IRFs. Figure 10 plots the IRFs implied by the model and compares them to the IRFs from the VAR (with the grey-shaded areas demarking the 16%-84% error bands of the VAR responses). As the plot for TFP shows, the stochastic growth process in (7) is capable of matching almost perfectly the gradual increase of TFP after the news shock process.

The estimated model does well at matching the responses of macroeconomic variables. The responses of output and investment closely match the VAR responses. The model also matches

quite well the observed sharp drop of inflation on impact of the shock and the gradual return of inflation to steady state thereafter. On a less positive note, while the model matches the initial increase in consumption, it does not generate the subsequent increase in consumption to the new balanced growth level.

The responses of the term structure data (the short rate, the spread and the long rate under the Expectations Hypothesis) have the correct sign and are reasonably persistent. At the same time, the magnitude of these term structure dynamics fall short of the empirical responses implied by the VAR. In particular, the initial drop in the Federal Funds rate is only about half as large as in the data. As a result, the spread as implied by the Expectations Hypothesis increases less than what we estimated from the VAR. Likewise, the drop in the long rate as implied by the Expectations Hypothesis is insufficient. Finally, as the bottom two panels of Figure 10 show, the model remains even further away from matching the dynamics of the actual spread and long rate. While this last result should not come as a surprise given that term premia are constant by definition in the model (and these IRFs were not part of our estimation criteria), it nevertheless illustrates the extent to which the model fails to generate the type of term structure dynamics that we observe in the data.

It is possible to find a set of parameters that allows the model to match the responses of the term structure data. However, this is only possible if we choose parameters to match only the term structure IRFs. The problem is that for such a parameter estimate the model's performance for the macroeconomic variables completely falls apart. In other words, modern New Keynesian DSGE models as proposed by Smets and Wouters (2007) fail to match simultaneously the macroeconomic response and the term structure response to TFP news shocks.²²

The estimated parameters that generate the IRFs in Figure 10 are reported in Table 2. We begin with the parameters defining price and wage stickiness. The estimates of $\kappa_p = 0.145$ and $\omega_p = 0$ indicate that the data favors a purely forward-looking New Keynesian Phillips Curve (NKPC) with a relatively small degree of price rigidity (i.e. with standard Dixit-Stiglitz goods differentiation, a NKPC slope of $\kappa_p = 0.145$ implies an average price duration of about 3 quarters).²³ By contrast,

²²This finding contradicts the conclusion by De Graeve, Emiris and Wouters (2008) who argue that a very similar log-linear macroeconomic model is able to fit the unconditional dynamics of the term structure. The reason for this difference in results is that they estimate their model using a large number of relatively unrestricted business cycle shocks, which provides the estimation with a lot of flexibility to match not only the macroeconomic but also the term structure data. Our exercise is more restrictive because it assesses the model's response conditional on one type of shock, the news shock. This news shock is not present in the De Graeve, Emiris and Wouters model.

²³Since the price indexation parameter ω_p is estimated to be at its lower bound, it would not be meaningful to report a standard error. We thus fix the parameter when computing standard errors for the other parameters. We

the data requires a large degree of nominal wage rigidity with an estimated frequency of wage reoptimization of only $1 - \xi_w = 0.15$ per quarter and complete indexation of non-reoptimized wages to past inflation (i.e. $\omega_w = 1$). These estimates are close to the full-information estimates by Smets and Wouters (2007). The main force behind these estimates is the sharp drop of inflation on impact of the news shock, which the model can generate only if inflation is a mainly forward-looking process (i.e. ω_p is small). In that case, inflation is driven predominantly by current and future expected marginal cost. Marginal cost, in turn, depends positively on wages and negatively on TFP. After a news shock, the negative income effect on labor supply from consumption smoothing puts upward pressure on wages and thus on marginal cost. In subsequent periods, as the expected increase in TFP realizes, marginal cost falls. The drop in inflation on impact and the gradual response thereafter occurs only if there is a lot of wage rigidity (i.e. ξ_w and ω_w large) so that the initial increase in marginal cost is relatively modest and its evolution thereafter relatively smooth.

The IRFs to the news shock are also informative about the parameters defining variable capital utilization and investment adjustment cost. The variable capital utilization parameter σ_u is estimated at its lower bound, which means that capital utilization is proportional to the rental rate of capital.²⁴ As Dotsey and King (2006) show, variable capital utilization reduces the sensitivity of marginal cost. Hence, the smaller the cost of utilization, the less pressure production exerts on marginal cost. This helps the model reconcile the large expansion of production with the persistent drop in inflation in the wake of the news shock. The adjustment cost parameter S'' , in turn, is also estimated at its lower bound; i.e. adjustment costs are zero. This result is driven primarily by the drop of investment on impact of the news shock. If investment adjustment costs were large, then there would be a strong incentive to smooth investment, which in turn would put upward pressure on production and inflation. The zero investment adjustment cost estimate raises an important challenge for the model since it is estimated to be an important ingredient of macroeconomic dynamics conditional on more standard business cycle shocks (e.g. Christiano, Eichenbaum and Evans, 2005).

Our final set of parameter estimates are for the monetary policy rule. The estimates $\rho = 0.68$ and $\theta_\pi = 2.63$ indicate that the Fed smoothens its policy rate considerably and reacts aggressively to inflation expectations. The latter estimate may appear relatively high but is consistent with the full-information estimate reported in Smets and Wouters (2007). Both parameter estimates are

adopt the same approach for all other parameters that are estimated at their respective lower or upper bound.

²⁴For $\sigma_u = 0$, depreciation increases linearly with utilization.

crucial to generate the sharp drop in the Federal Funds rate on impact of the news shock and its persistent evolution thereafter. The estimates $\theta_{ygap} = 0.06$ and $\theta_{\Delta y} = 1.00$, in turn, imply that the Fed does not really respond to the output gap but reacts strongly to output growth. This is consistent with Orphanides (2005) who argues that U.S. monetary policy is better described by a rule that responds to observables such as output growth rather than some theoretical output gap measure. The focus on output growth rather than the output gap turns out to be a second crucial ingredient for the model to generate a fall in the Federal Funds rate. In response to a news shock, the output gap in the model increases whereas output growth falls.²⁵ Hence, if monetary policy responded strongly to the output gap, this would reduce (or even reverse) the already insufficient drop in the Federal Funds rate. A strong response to output growth, by contrast, reinforces the accommodative stance of the Fed thus bringing the model closer to the observed term structure dynamics.²⁶ Simulations with more basic policy rules that feature only the output gap show that the fit of the model under this specification falls apart almost completely, with the Federal Funds rate and the slope hardly responding to the news shock. This illustrates that augmenting DSGE models with term structure variables provides valuable information to discipline the description of monetary policy and the model’s performance with respect to news shocks.

7 Conclusion

In this paper we provide a new structural interpretation of the relationship between the slope of the term structure and macroeconomic fundamentals. Our results show that there exists one single shock that can account for a large part of slope movements at all horizons. We interpret the slope shock as a news shock about future innovations to TFP. We assess this interpretation formally using the identification of Barsky and Sims (2009) and Sims (2009) and find a striking correspondence. In response to a positive news shock real activity does not initially respond. At longer horizons real activity increases gradually towards a higher permanent level. Inflation falls sharply on impact of

²⁵As described above, the output gap is defined as the difference between actual output and potential output in the absence of nominal price and wage rigidities. In response to a news shock, prices drop abruptly, which means that firms’ average markups decrease. Hence, actual output drops less than potential output (for which markups are constant by definition) and the output gap *jumps up*.

²⁶Barsky and Sims (2009) argue in favor of a similar monetary policy rule that does not respond to the output gap. However, their argument is somewhat different, based on their empirical result that real short-term rates in response to a TFP news shock are positive. As we pointed out above, however, real short-term rates are negative after a TFP news shock if inflation is measured by the more inclusive GDP deflator rather than the CPI deflator.

the shock and returns only slowly to its initial level. Monetary policy reacts to the low inflation by lowering the Federal Funds rate. As a result, the short-end of the term structure falls, which in turn leads to a sharp increase in the slope. *Endogenous* monetary policy thus provides an important channel through which TFP news shock transmit to movements in the slope.

Our news interpretation provides a structural interpretation for why the yield curve is such a reliable predictor of future output growth. These movements in the slope are asset markets efficiently responding to news about the future level of productivity. Future productivity is of course a main determinant of future output. This result also provides a structural interpretation for the result in the **existing** macro-finance literature on the strong linkages **between the yield curve**, inflation, and real output.

Our results provide an important benchmark to evaluate theories of the term structure and, more generally, DSGE models. We show that a medium-scale DSGE model along the lines of Smets and Wouters (2007) falls well short of matching the term structure response to a TFP news shock that we see in the data. While we can find parameters to match either the term structure data or the macroeconomic data we cannot do both simultaneously. This failure of generating realistic term structure dynamics is problematic for two reasons. First, asset prices (of longer-term securities in particular) are an important determinant of consumption and investment decisions. If a DSGE model cannot simultaneously match both macroeconomic and asset price dynamics, then this suggests a serious empirical shortcoming of theory. Second, medium-scale DSGE models are increasingly used for monetary policy analysis. If these models fail to generate reasonable term structure dynamics, then it seems difficult to trust them for the evaluation of how monetary policy transmits into the economy. A fruitful path for future research is search for a mechanism to augment DSGE models with that can generate large endogenous variations in term premia and to estimate these models in a full-information context with both term-structure data and news shocks.

References

- [1] Altig, D., Christiano, L.J., Eichenbaum, M., Lindé, J. (2004). "Firm-Specific Capital, Nominal Rigidities and the Business Cycle." NBER working paper 11034.
- [2] Ang, Andrew and Monica Piazzesi, (2003), "A No-Arbitrage Vector Autoregression of Term Structure Dynamics with Macroeconomic and Latent Variables", *Journal of Monetary Economics* 50(4), 745-787.
- [3] Ang, A., M. Piazzesi and M. Wei (2006). "What does the Yield Curve tell us about GDP growth?" *Journal of Econometrics* 131, 359-403.
- [4] Bansal, Ravi and Ivan Shaliastovich (2007). "Risk and Return in Bond, Currency, and Equity Markets," Working paper.
- [5] Barsky, Robert and Sims, Eric R. (2009). "News Shocks" Working paper.
- [6] Basu, S. and J. G. Fernald (1997). "Returns to Scale in U.S. Production: Estimates and Implications." *Journal of Political Economy* 105(2), 249-283.
- [7] Basu, Susanto, John Fernald, and Miles Kimball (2006). "Are Technology Improvements Contractionary?" *American Economic Review* 96, 1418-1448.
- [8] Beaudry, Paul and Frank Portier (2006). "News, Stock Prices, and Economic Fluctuations." *American Economic Review* 96, 1293-1307.
- [9] Campbell, J.Y., Shiller, R.J. (1987). "Cointegration and tests of present-value models." *Journal of Political Economy* 95, 1062-1088.
- [10] Campbell, J.Y., Shiller, R.J. (1991). "Yield spreads and interest rate movements: a bird's eye view." *Review of Economic Studies* 58, 495-514.
- [11] Christiano, L. J., Eichenbaum, M., Evans, C. L. (2005). "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113, 1-45.
- [12] Clarida, R., J. Gali and M. Gertler (1999). "The Science of Monetary Policy: A New Keynesian Perspective." *Journal of Economic Literature* 37, 1661-1707
- [13] Cochrane, John and Monika Piazzesi (2005). "Bond Risk Premia," *American Economic Review* 95, 138-60.
- [14] Comin, D. and M. Gertler (2006). "Medium Term Business Cycles." *American Economic Review*.
- [15] Dai, Qiang and Kenneth Singleton (2000). "Specification Analysis of Affine Term Structure Models." *Journal of Finance*, 1943-1978.
- [16] Dai, Q. and T. Philippon (2006). "Fiscal Policy and the Term Structure of Interest Rates." Working paper.

- [17] De Graeve F., M. Emiris and R. Wouters (2009). "A Structural Decomposition of the US Yield Curve." *Journal of Monetary Economics* 56, 545-559.
- [18] Den Haan, W. (1995). "The Term Structure of Interest Rates in Real and Monetary Economies." *Journal of Economic Dynamics and Control* 19, 909-40.
- [19] DiCecio, R. (2008). "Sticky Wages and Sectoral Labor Comovement." *Journal of Economic Dynamics and Control* 33, 538-53.
- [20] Diebold, F.X., Rudebusch, G.D. and Aruoba, B. (2006). "The Macroeconomy and the Yield Curve: A Dynamic Latent Factor Approach." *Journal of Econometrics* 131, 309-338.
- [21] Diebold, F.X. and Li, C. (2006). "Forecasting the Term Structure of Government Bond Yields." *Journal of Econometrics* 130, 337-364.
- [22] Donaldson, J., T. Johnson and R. Mehra (1990). "On the Term Structure of Interest Rates." *Journal of Economic Dynamics and Control* 14, 571-96.
- [23] Dotsey M., King, R.G., 2002. Price, Production and Persistence. NBER working paper 8407.
- [24] Duffie, Darrell and Rui Kan (1996). "A Yield-Factor Model of Interest Rates" *Mathematical Finance* 6, 379-406.
- [25] Evans, Charles L. and Marshall, David A. (2007). "Economic determinants of the nominal treasury yield curve." *Journal of Monetary Economics* 54(7), 1986-2003.
- [26] Estrella, A., Hardouvelis, G.A. (1991). "The term structure as predictor of real economic activity." *Journal of Finance* 46, 555-576.
- [27] Fama, E.F., Bliss, R.R. (1987). "The information in long-maturity forward rates." *American Economic Review* 77, 680-692.
- [28] Faust, Jon and Eric Leeper (1997). "When Do Long Run Identifying Restrictions Give Reliable Results?" *Journal of Business and Economic Statistics* 15, 345-353.
- [29] Fisher, J. D. M. (2006). "The Dynamic Effects of Neutral and Investment-Specific Technology Shocks." *Journal of Political Economy* 114, 413-451.
- [30] Francis, Neville, Michael Owyang, and Jennifer Roush (2007). "A Flexible Finite Horizon Identification of Technology Shocks." Working paper, Federal Reserve Bank of St. Louis.
- [31] Gollin, D. (2002). "Getting Income Shares Right." *Journal of Political Economy* 110(2), 458-474.
- [32] Gürkaynak, Refet, Brian Sack, and Jonathan Wright (2007). "The U.S. Treasury Yield Curve: 1961 to the Present," *Journal of Monetary Economics* 54, 2291-2304.
- [33] Jaimovich, Nir and Sergio Rebelo (2009). "Can News about the Future Drive the Business Cycle?" *American Economic Review* 99(4), 1097-1118.
- [34] King, R. G., Watson, M.W. (1998). "The solution of singular linear difference systems under rational expectations." *International Economic Review* 39, 1015-26.

- [35] King, R. G. and A. Kurmann (2002). "Expectations and the Term Structure of Interest Rates: Evidence and Implications." Federal Reserve Bank of Richmond *Economic Quarterly*.
- [36] Knez, Peter, Robert Litterman and Jose A. Scheinkman (1994). "Explorations Into Factors Explaining Money Market Returns." *Journal of Finance* 49(5), 1861-1882.
- [37] Orphanides, A. (2003). "Historical monetary policy analysis and the Taylor rule." *Journal of Monetary Economics* 50, 983-1022.
- [38] Piazzesi, Monika, and Martin Schneider (2006). "Equilibrium Yield Curves." *NBER Macro Annual*, 389-442.
- [39] Piazzesi, Monika (2005). "Bond yields and the Federal Reserve", *Journal of Political Economy* Volume 113, Issue 2, pp. 311-344.
- [40] Rotemberg, J. J., Woodford, M. (1995). "Dynamic General Equilibrium Models with Imperfectly Competitive Product Markets." In: Cooley, T.F. (ed). *Frontiers of Business Cycle Research*. Princeton: Princeton University Press.
- [41] Rotemberg, Julio (2003). "Stochastic Technical Progress, Smooth Trends, and Nearly Distinct Business Cycles." *American Economic Review* 93, 1543-1559.
- [42] Rudebusch, G. D. and E. T. Swanson (2008). "Examining the Bond Premium Puzzle with a DSGE Model." *Journal of Monetary Economics* 55, 111-26.
- [43] Sims, Eric R. (2009). "Expectations Driven Business Cycles: An Empirical Evaluation." Working paper.
- [44] Schimdt-Grohe and Martin Uribe (2008). "Whats News in Business Cycles" Working paper.
- [45] Smets, F. and R. Wouters (2007). "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review* 97(3), 586-606.
- [46] Uhlig, Harald (2003). "What Drives GNP?" Working paper.
- [47] Wachter, Jessica (2006). "A Consumption-Based Model of the Term Structure of Interest Rates." *Journal of Financial Economics* 79, 365-99.

Table 1: Calibrated parameters

Parameter	Description	Calibration
α	Elasticity of production to labor	0.75
β	Discount factor	0.997
δ	Depreciation rate	0.025
η	Frisch elasticity of labor supply	1
θ_p	Elasticity of substitution across goods	10
θ_w	Elasticity of substitution across labor	3
μ	Steady state growth rate of TFP	1.0045

Table 2: Estimated Parameters

Parameter	Description	Estimate
γ	Persistence of TFP growth	0.837 (0.037)
σ_{ε_2}	Standard deviation of news shock	0.061 (0.018)
κ_p	Marginal cost slope of NKPC	0.145 (0.000)
ω_p	Degree of price indexation	0 (n.a.)
ξ_w	Frequency of wage adjustment	0.844 (0.001)
ω_w	Degree of wage indexation	1 (n.a.)
b	Habit persistence	0.279 (0.000)
σ_u	Capital utilization parameter	0 (n.a.)
S''	Investment adjustment cost	0 (n.a.)
ρ	Persistence of monetary policy	0.682 (0.000)
θ_π	Inflation response	2.628 (0.007)
θ_{ygap}	Output gap response	0.063 (0.000)
$\theta_{\Delta y}$	Output growth response	1.000 (0.004)

Figure 1: Variance Decomposition of First Slope Shock




Figure 2: Impulse Response to First Slope Shock




Figure 3: Impulse Response to Contemporaneous TFP Shock




Figure 4: Variance Decomposition to TFP News Shock




Figure 5: Impulse Response to TFP News Shock




Figure 6: Comparison of First Slope Shock and News Shock




Figure 7: Historical Simulations with TFP News Shock




Figure 8: Impulse Response to TFP News Shock




Figure 9: Impulse Response of Term Premium and Expectations Hypothesis to TFP News Shock




Figure 10: Impulse Response to TFP News Shock in DSGE Model and VAR

